

Frustration in biomolecules

Diego U. Ferreiro¹, Elizabeth A. Komives^{2*} and Peter G. Wolynes^{3*}

¹Protein Physiology Lab, Dep de Química Biológica, Facultad de Ciencias Exactas y Naturales, UBA-CONICET-IQUIBICEN, Buenos Aires, Argentina

²Department of Chemistry and Biochemistry, University of California San Diego, La Jolla, CA 92093, USA

³Department of Physics, Department of Chemistry, and Center for Theoretical Biological Physics, Rice University, Houston, TX 77005, USA

Abstract. Biomolecules are the prime information processing elements of living matter. Most of these inanimate systems are polymers that compute their own structures and dynamics using as input seemingly random character strings of their sequence, following which they coalesce and perform integrated cellular functions. In large computational systems with finite interaction-codes, the appearance of conflicting goals is inevitable. Simple conflicting forces can lead to quite complex structures and behaviors, leading to the concept of *frustration* in condensed matter. We present here some basic ideas about frustration in biomolecules and how the frustration concept leads to a better appreciation of many aspects of the architecture of biomolecules, and especially how biomolecular structure connects to function by means of localized frustration. These ideas are simultaneously both seductively simple and perilously subtle to grasp completely. The energy landscape theory of protein folding provides a framework for quantifying frustration in large systems and has been implemented at many levels of description. We first review the notion of frustration from the areas of abstract logic and its uses in simple condensed matter systems. We discuss then how the frustration concept applies specifically to heteropolymers, testing folding landscape theory in computer simulations of protein models and in experimentally accessible systems. Studying the aspects of frustration averaged over many proteins provides ways to infer energy functions useful for reliable structure prediction. We discuss how frustration affects folding mechanisms. We review here how the biological functions of proteins are related to subtle local physical frustration effects and how frustration influences the appearance of metastable states, the nature of binding processes, catalysis and allosteric transitions. In this review, we also emphasize that frustration, far from being always a bad thing, is an essential feature of biomolecules that allows dynamics to be harnessed for function. In this way, we hope to illustrate how *Frustration* is a fundamental concept in molecular biology.

1. Perspectives on frustration 286

- 1.1. Frustration, logic and magnets 289
- 1.2. Frustration in heteropolymers 293

2. Frustration and the energy landscape of real proteins 302

- 2.1. The energy landscape of real proteins 302
- 2.2. A license to learn via bioinformatics: structure prediction and energy functions 303

* Authors for correspondence: Elizabeth A. Komives, Department of Chemistry and Biochemistry, U.C. San Diego, La Jolla, CA 92093-0378, USA; Tel.: (858) 534-3058, Fax: (858) 534-7390; Peter G. Wolynes, Department of Physics, Department of Chemistry, and Center for Theoretical Biological Physics, Rice University, Houston, TX 77005, USA; Tel.: (713) 348-4101, Fax: (713) 348-4109.

2.3. Localizing and visualizing frustration in natural proteins 308

2.4. Patterns of frustration in globular proteins 312

3. Frustration in the evolution of natural proteins and in protein design 314

3.1. Frustration and the complexity of Nature's amino acid code 314

3.2. Frustration in artificially designed proteins 317

4. Folding kinetics and frustration 320

4.1. Kinetics on idealized funnel landscapes 320

4.2. Chemical frustration 321

4.3. Symmetry effects 323

4.4. Frustration in repeat proteins 325

4.5. Anomalous ϕ values 328

4.6. Non-native intermediates 330

4.7. Friction and roughness effects on rates 331

5. Functional consequences of frustration 333

5.1. Binding sites are often frustrated 333

5.2. Frustration in allostery and local conformational changes 334

5.3. Frustration in metastable and multistable proteins 338

5.4. Frustration, dynamics, and catalysis 341

5.5. Frustration and the initiation of aggregation 345

5.6. Frustration in nucleic acids and protein complexes 348

6. Frustration is a fundamental aspect of biochemistry 352**7. Acknowledgements 354****8. References 354****1. Perspectives on frustration**

Life is based on molecular information processing. Thousands of nucleic acids and proteins are organized in cells, cooperate and reproduce themselves as a group and respond to their environment. Most of the biomolecules individually carry out the complex task of decoding their own one-dimensional sequences in order to find three-dimensional structures (Anfinsen, 1973) furthermore yielding four-dimensional dynamical patterns that permit each molecule to carry out its 'functions' (Frauenfelder *et al.* 1991). This decoding task, the first step of which is called folding, is a challenging information processing step that only recently has been successfully carried out for some small proteins by computers (Piana *et al.* 2012). This task, even with computer aid, severely challenges the most clever human gamers (Cooper *et al.* 2010). Contemplating why this task is so hard both for computers and for people but is apparently not hard for the molecules themselves leads us to consider the concept of *Frustration* in biomolecules, the subject of this review.

In human psychology frustration is an emotion 'the experience of non-fulfillment of some wish or needs' (Anonymous, 2013). The emotion of frustration according to psychologists, often arises because a person has conflicting goals that interfere with one another (see Fig. 1). It is this type of internal frustration that can plague information processing rather generally even outside the human context. This sort of frustration is important in formal logic and in computer science and is closely related to the notion of frustration as it is used to describe physical




Fig. 1. Frustration is a general feeling and a deep concept. This cartoon presents part of the problem: a group of people that cannot be seated optimally in a single row. The upper panel shows a social graph that takes into account the interactions among people. The lower panel shows the irritation of the character that tries to calculate the arrangement for him and his seven friends, finding out that there will always be unsatisfied agents. Reproduced with permission from: <http://xkcd.com/173>.

systems, such as disordered magnets (Néel, 1970; Nordblad, 2013; Schiffer & Ramirez, 1996; Vannimenus & Toulouse, 1977; Wannier, 1950a) or biomolecules (Bryngelson & Wolynes, 1987). For these inanimate systems, the 'goals' are set by evolution and are reflected in the systems' own laws of development and motion given the genetic information. These laws are usually encoded in an energy function or 'Hamiltonian', which in turn determines the kinetics and thermodynamics at the macroscopic level.

Frustration slows the performance of computer algorithms and indeed can make some tasks impossible for a computer to carry out in a timely fashion. For physical systems, frustration leads to complex dynamics that often spans many orders of magnitude in time scale. Some disordered magnets, called spin glasses, can never seem to settle into any equilibrium state, no matter how long one waits. This behavior resembles the way some liquids do not crystallize into their lowest energy structure when cooled below their freezing point but instead become trapped in a glassy state. Proteins have some of the characteristics of glasses at low temperatures (Bryngelson & Wolynes, 1987; Frauenfelder *et al.* 1991) but at physiological temperatures the functionally relevant motions are generally fast enough to reliably take place in a cell (Frauenfelder, 1987). Glassy dynamics of a condensed matter system or of a biomolecule resembles a computer

code still in the process of crunching the problem – generally when a system is glassy the actual system's state depends on its history and kinetics. While kinetic control of pathways is paramount at the cellular and network level in biology, the physics of individual biomolecules is largely controlled thermodynamically and is independent of the history of the macromolecular system. We shall see the reason for this is that Nature has evolved biomolecules where the frustration of protein dynamics is minimal. Nevertheless some frustration is present in evolved biomolecules and is crucial for many aspects of their function.

The purpose of this paper is to explain the basic ideas about frustration in biomolecules and how these ideas let us understand many aspects of the architecture of biomolecules and how that architecture connects with their motions. These ideas are simultaneously both seductively simple and perilously subtle to understand completely. To do so without a mathematical framework is impossible. Such a framework is provided by the energy landscape theory of protein folding that indeed gives criteria for quantifying frustration globally in large systems (Bryngelson & Wolynes, 1987; Goldstein *et al.* 1992b; Wolynes, 2005). Formally such criteria would be rigorous for infinite-sized proteins, so that they can only be applied using approximations to finite systems. The global criterion for quantifying frustration can be implemented at several levels of mathematical sophistication and has led to concrete and useful protein structure prediction algorithms (Davtyan *et al.* 2012). A complete mathematical description is quite complicated and this review will only sketch some of the basics (Bryngelson *et al.* 1995; Plotkin & Onuchic, 2002). But this summary will allow us to concentrate on reviewing how a local notion of frustration can also be quantified (Ferreiro *et al.* 2007) and also how localizing frustration can be used to look at protein structures and motion both in a survey sense and to understand the specific examples. The local notion of frustration is not rigorous and the approximate form we use could be severely criticized by statistical mechanics purists. Nevertheless we have found in several investigations, that a very simple estimate of local frustration can give surprising insights into how proteins fold and misfold, how proteins move once they are folded and how they may have evolved to carry out their functions.

The structure of this review is as follows. First, we develop the notion of frustration seeing how it enters the area of abstract logic and how it is used for the simplest condensed matter systems which are magnets. We then discuss how the frustration concept applies specifically to heteropolymers and discuss the role of frustration in the energy landscape theory of protein folding. We further review tests of folding landscape theory in the stylized computer simulation world provided by lattice models of proteins. We then briefly review how understanding frustration globally for real proteins provides a 'license to do bioinformatics', enabling us to extract physical information about molecular forces from structural databases. Studying frustration thus provides practical tools for protein structure prediction and leads to an appreciation of the complexity of the protein-folding code. Some illustrations of how frustration occurs in evolved natural proteins and in artificially designed proteins are then discussed. Following this we review how residual frustration in some biomolecules affects their folding kinetics focusing on some anecdotal examples. *En passant* we discuss broader uses of the frustration concept in folding, describing topological frustration, chemical frustration as it occurs in metalloproteins and the role of symmetry, which surprisingly can frustrate the folding process. The larger part of the review then goes beyond the study of folding to seeing how the notion of frustration help understanding biomolecular function by examining the role of frustration in metastable protein states, binding processes, enzymatic catalysis and allosteric transitions. We also discuss frustration in protein aggregation and in protein–DNA assemblies.

1.1. Frustration, logic and magnets

To get a feeling for frustration it is fun to start with some old, seemingly simple, logic puzzles. Medieval philosophers developed abstract logic using such puzzles anticipating many modern ideas in metamathematics and computational complexity. They are known as ‘sophismata’ (Heytesbury, 1494; Kretzmann, 1982). Consider these two sentences:

1. All men are donkeys or men and donkeys are donkeys.
2. The sentence ‘Socrates says something false’ in the case where Socrates says nothing other than ‘Socrates says something false’.

Are these statements true or false? How do you feel when you try to figure this out? Frustrated? Figuring these puzzles out in the original Latin in which they were written is even more challenging than it is in English translation. Modern students of logic can do better than most students did centuries ago by systematically using the idea that one must first disambiguate the sentences (with appropriate punctuation in English, such as parentheses and commas, not used in classical Latin) and then judge the truth or falsity of the elemental parts e.g., we are all pretty sure that the proposition ‘All men are donkeys’ is false, etc. After generating a ‘truth table’ of possibilities a ‘simple’ search of all the entries tells us the answer. Using more cumbersome argumentation, Albert of Saxony concluded that two different renditions of the truth or falsity of sentence 1 are possible. Sentence two is one of the forms of the ‘liar paradox’ with which even 20th century philosophers such as Whitehead and Russell continued to grapple. In modern computer science, much longer strings of logical statements are needed to answer practical problems such as how to detect objects in a scene that is digitally encoded (Field, 1999). Are there methods that allow a computer to decide the truth or falsity of very complex, long statements? Can a computer find consistent solutions of any set of abstract assertions? Such questions are known as constraint satisfaction problems (Mézard & Mora, 2009; Monasson *et al.* 1999). In general, a huge number of possibilities must be tested to see if the logical constraints in the sophism can be satisfied in some way and the computations can become quite lengthy.

That solving a particular problem must lead to frustration is easiest to see when the choices for the subparts of the problem are few – the binary choice of whether an individual proposition is true or false being the simplest. Binary choice is the ‘atom’ of logic and computer science. The consequences of conflicts for problem solving are most profound, however, when many such multiple choices must be made to specify the solution. This is because the search space through the truth table becomes exponentially large with the number of choices needed to specify the answer. We quickly see the liar paradox has no solution but for people, the task of deciding whether the longer strings of logical statements in typical computer programs actually lead to solutions or whether they do not is much, much harder. As a practical matter we usually just run the programs and check whether they stop or whether they yield a unique solution after trying out all possible inputs.

In physics, frustration entered the lexicon in the study of magnets (Anderson, 1950; Vannimenus & Toulouse, 1977). Like the logical puzzles we have just discussed, the elementary units in magnets, the spins, have a limited number of choices – classically they can ‘choose’ to be up or they can ‘choose’ to be down. Actually in some magnets the quantum mechanical aspect of the spin can give the spins the opportunity not to choose at all but to remain in the limbo of a quantum superposition of up and down, like a resonance hybrid in organic chemistry, but this complication corresponding to non-Aristotelian logic can often be ignored. The analog of the logical-grammatical constructions for the spins in a magnet are the energies of interactions




Fig. 2. Frustration entered the physics lexicon in the study of magnets. The arrows represent spins that can be in any two states: *up* or *down*. Favorable antiferromagnetic interactions between spins are represented by squiggly lines. The potentially frustrating problem is: what is the lowest energy arrangement of the up and down spins that can be reached dynamically? On the left, the particles are arranged on a rectangular lattice. How would you set the spins such that all local interactions are satisfied? What happens if the particles are arranged on a triangular lattice, as shown on the right? There is no way to arrange the spins so that every component interaction can be minimized. The triangle is ‘frustrated’.

between the spins. The simplest energy describing the interactions between two spins, one on the i th atom, S_i , and on another, usually neighboring atom j th, S_j , has the form $E_{ij} = -J_{ij}S_iS_j$ where the two choices for S_i and S_j are ± 1 , i.e. up or down. If J_{ij} is positive, the low-energy state aligns the two spins in the same direction; if J_{ij} is negative the spins prefer to point in opposite directions.

If we have a large magnet the energy can be taken as the sum of the individual (pair) interactions $E = -\sum_{i \neq j} J_{ij}S_iS_j$. The potentially frustrating problem to solve then is, at low temperature, what is the lowest energy arrangement of the up and down spins that can be reached dynamically? (See Fig. 2). This problem is ‘solved’ by the system when it is cooled and random thermal motions allow it to try out various configurations or arrangements of the spins, according to their energy.

If all the J_{ij} are positive, it is easy even for people to see the answer: the lowest energy state has all the spins pointing in the same direction – they are either all up or all down. What if all (or even some) of the J_{ij} ’s are negative? Now things seem to be tricky, but sometimes people can figure out the answer pretty quickly for certain situations. Consider first what turns out to be an easy problem: suppose the atoms bearing the spins are situated on a simple square lattice and all the J_{ij} are negative. In this case, an interlaced checkerboard of up and down spins has the lowest energy and minimizes simultaneously every local interaction E_{ij} . The possible conflict between different local interactions can be resolved (see Fig. 2). The resulting pattern is called an anti-ferromagnet. This solution, while it can be quickly checked by anyone, was not really that easy to see *a priori*, and indeed (along with other insights) figuring out this pattern netted its discoverer, Louis Néel, a Nobel laureate in Physics (Néel, 1970). Although it takes at least a few moments for humans to see the answer for the square lattice antiferromagnet, the system itself has a much easier time figuring out what to do. When an antiferromagnet is cooled below a certain temperature it orders spontaneously and fairly quickly.

But antiferromagnets on other lattices can be more complex than the uniform square lattice antiferromagnet case. Imagine the spins reside on atoms situated on a triangular lattice with all antiferromagnetic interactions (Fig. 2). Already a single triangle of the lattice shows the problem. There is no way to arrange the spins, so that every component interaction can simultaneously be minimized; always at least one interaction must be still in a high-energy state. We would say the individual triangle is ‘frustrated’.

What happens for a big lattice of frustrated triangles? In the strictly two-dimensional system, one finds a huge degeneracy of possibilities. There is no unique ground state. This would lead to




Fig. 3. Frustration leads to barriers in the energy landscape. The arrows represent spins that can be in any two states on a rectangular lattice. Favorable ferromagnetic interactions are represented by straight lines and antiferromagnetic interactions by squiggly lines. In this case most of the bonds are ferromagnetic so at low temperature most of the spins will be parallel. Two choices of spin assignments are shown. These have similar energies, are relatively stable, but they are globally dissimilar. At left most spins are down, at right most spins are up. In both cases, the same number of interactions remains locally unsatisfied (red dots). There are two ways of arranging the ‘misaligned’ spins. To rearrange from one to the other a very complicated set of spin changes must be made. This rearrangement is entropically disfavored at high temperature and energetically disfavored at low temperature. Many large-scale moves must be made for the system to find out which arrangement or set of arrangements are most stable. Although there is a ground state it can be hard to find because of the very slow dynamics that emerges.

an apparent contradiction to the Third Law of Thermodynamics, which requires entropy to vanish at absolute zero, implying a unique ground state. Wannier showed that $T=0$ was a critical point for the triangular antiferromagnet and thus there would be large length scale fluctuations even at low temperature (Wannier, 1950a). In real three-dimensional systems, weaker interactions between planes of spins restore the Third Law and lead to ordering to a unique lowest energy state, but the energies required to excite the system from the ground state now depend only on the weak interactions and so are much smaller than you would otherwise expect. The kind of ordering found in such three-dimensional systems is also very complex and generally hard to predict with precision (Collins & Petrenko, 1997).

You might think the triangular antiferromagnet case is special because of the exact symmetry of the frustration and indeed it is. Without symmetry, frustration and degeneracy still go together but manifest themselves in a very different way: very slow dynamics emerges that gives the system difficulty reaching equilibrium at low temperatures. Although there is a ground state it can be hard to find. Such a system is called a ‘spin glass’. The models that are simplest to describe place spins on a simple lattice and at random assign some bonds to be ferromagnetic and others to be antiferromagnetic (Fig. 3). A particular assignment of bonds is analogous to a sophism having a particular combination of logical connectors (AND, NOT, OR. . .). Although there are numerous approaches for finding the low-energy states of such a magnet, in the general case all known algorithms take a very long time. This time scales up exponentially with the size of the magnet. The reason for this slowness is there are many choices of spin assignments that have similar energies that are relatively stable but that are globally quite dissimilar. It is difficult for the system to carry out such rearrangements between degenerate states. How does frustration lead to barriers? In fact, the locality of moves makes this hard to understand but you can get some intuition from a weakly frustrated spin system with only a few frustrated interactions like the example shown in Fig. 3. Since most of the bonds in Fig. 3a are ferromagnetic, at low temperature most of the spins will be parallel, let us say oriented up. However, some

interactions must be unsatisfied locally. In the example shown there are two ways of arranging these ‘misaligned’ spins. To rearrange from one to the other a very complicated set of spin changes must be made. This rearrangement is entropically disfavored at high temperature and energetically disfavored at low temperature. Many large-scale moves must be made for the system to find out which arrangement or set of arrangements are most stable.

When frustration is widespread the system resembles a glass – the exact arrangement of spins found at any one time depends on the history of the system, just as a liquid cooled into a glass has properties that depend on its annealing schedule. This is the feature of frustration most noted in the condensed matter physics of magnets. The global effect of widespread frustration is that it leads to such completely new phases of matter. The illustrative example, however, focuses attention on a point of special relevance to biochemistry: even a little frustration only locally may lead to new degeneracies that could have dynamical consequences.

Nearly degenerate states will interconvert slowly not only because the frustrated parts of the system have to be moved about, but possibly also because other parts of the system that are not frustrated have to move in order to allow the rearrangement to occur. This entails breaking up strong interactions and is energetically costly and thus difficult to be achieved by random thermal motions.

Although completely random magnets and weakly, locally frustrated magnets are not very familiar real-world systems the binary choices of spins in magnets have made magnetism a convenient playground to understand degeneracies in many other non-magnetic systems. While spin glasses have huge numbers of near-degenerate states, ferromagnets or the simple antiferromagnets have only a few low-energy patterns. We can imagine magnets that exhibit intermediate levels of degeneracy. The most conceptually important of such mildly complex landscapes are models that encode a modest number of spin patterns. Such magnets would reliably order to a specific input pattern if only a fraction of the spins were initially organized to resemble the template. Such magnetic models thus represent a very simple model of how one could construct an associative memory using very simple components. Little (1974) and later Hopfield (1982) began the study of such associative memory models. The energy function of such a magnet has the form, $E = - \sum_{ij} \sum_{\mu=1}^M \xi_i^\mu \xi_j^\mu S_i S_j$, where there are M patterns ξ_i^μ of spins S_i and the sum goes over all pairs. If there were only one such pattern, ξ_i^μ , the magnet would not be at all frustrated but is really just a disguised ferromagnet. Its ground state corresponds to each spin S_i taking on the value ξ_i . If only a few patterns are encoded you can see that patterns close to any one of them remain as fairly low-energy states (see Fig. 4) and the existence and encoding of the other patterns just adds some random energy contributions that probably do not completely destabilize the state, in any event, if there are not too many interferences. If the magnet starts in a configuration resembling one of the individual patterns at low temperature, as the spins reconfigure they will rearrange to resemble the complete encoded ‘memorized’ pattern. Such an easily reachable pattern of spins would resemble the native structure of a protein that can be reached quickly and is fairly robust to environmental changes (Fig. 4). Like a human brain which can retrieve memories after losing many cells, if many bonds of the lattice (but not too many!) are modified the pattern still remains retrievable. The ‘memory’ pattern is stored globally not just locally. This is quite different from a typical digital computer today where memories are put into specific addresses, locally, and recall can be far from robust.

Associative memory magnets can be studied very thoroughly using modern statistical mechanics. A somewhat surprising property of associative memory models is that they have a finite ‘capacity’ to store information: the maximum number of storable patterns scales only linearly

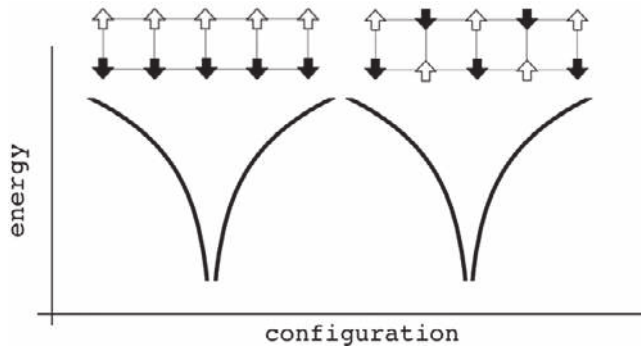


Fig. 4. Patterns can be encoded within magnets. The magnet is composed of different patterns of binary spins arranged on a rectangular lattice. Configurations are sorted by similarity to each other in a natural reaction coordinate measuring their overlap. The energy of each configuration is given as a function of the local interactions. In this case only two patterns are encoded, allowing two globally different ground states. Configurations close to any one of them are fairly low in energy. If there are not too many interference between the patterns, the existence of one pattern only contributes a small random energy that doesn't destabilize the other energy minimum much. If the magnet starts in a configuration resembling one of the individual patterns, at low temperature the spins will rearrange and recollect the complete encoded pattern. These patterns remain retrievable even if many bonds of the lattice (but not too many!) become modified. The memory patterns are thus stored globally, not just locally. There is a finite capacity in the maximum number of storable patterns, as too many of them make the system behave like a random spin glass and robust recollection of the associated input memories becomes impossible.

with system size. If there are too many patterns encoded, robust recollection becomes impossible. When there are too many patterns, the magnet instead behaves like a random spin glass and becomes 'confused' in trying to recall the encoded patterns. The situation is like that recounted in the joke about the old zoology professor who, when meeting a new faculty member at the Fall Faculty Party, says, 'No, no. Please don't tell me your name. Every time I learn a new name I forget a fish!'

Frustration enters into protein biochemistry in ways very analogous to what happens for magnets and for these neural network models of associative memory. Frustration leads to degeneracy. Finite degeneracy allows interesting and controlled emergent behaviors. Too much frustration and degeneracy lead to qualitatively different dynamics in which searching through the degeneracies cannot be carried out in a meaningfully short time. A little frustration can be good, a lot can be disastrous.

1.2. Frustration in heteropolymers

Ferromagnetic interactions or antiferromagnetic; spins up or down, on a fixed lattice – the choices are stark and it is easy to see the origin of frustration in trying to find the lowest energy state for magnets. Similarly we may ask, what is the root of the difficulty in predicting the three-dimensional folds of proteins? How do we specify the choices of structures and how do we judge their quality? The choice space is a bit harder to visualize for proteins than it is for magnets and there are several ways to represent choice space, none perfect. Pairs of peptide bonds are restricted in their relative orientations and side-chain orientations are restricted too, so local dihedral angles seem like a sensible choice. The global fold is however a very non-local function of the

individual backbone dihedral conformational selections – wild swings of portions of the chain are induced even by changing a single backbone dihedral by itself, so interactions (contacts) between the residues generally become completely reorganized by local moves.

Many such simple individual local reconfigurations of the backbone would be impeded by steric clashes and of course the chain connectivity must be preserved at all costs. The allowed configurations of a chain with excluded volume, possibly entangled, therefore require a much more complex encoding than a simple set of binary spins on a lattice. Connectivity and excluded volume by themselves already make computer searching through the allowed configuration space non-trivial (just think about the myriad of patterns in books on knitting or crocheting, or about modern day problems in mathematical knot theory). However, our understanding and study of dense polymer melts suggest the entanglement effect should actually be small for chains of the length of typical protein domains. Knots first begin to appear in random homopolymer chains typically only after hundreds of persistence lengths. Indeed knots appear rather rarely in the natural protein database, perhaps occurring at the 1% level (Sułkowska *et al.* 2012). They are more important for DNA folding.

Frustration due to topology alone is thus only a small part of the story of the difficulty of finding the global free energy minimum for proteins (but as we shall see ‘topological frustration’ does play a role in affecting *in vitro* protein-folding kinetics). The problems of moving the chain around that are encountered in making a single dihedral angle change can often be overcome by a local combination of several such moves, if the chain is not too compact, a so-called ‘bicycle pedal move’. Thus an alternative encoding of the protein configuration in terms of the contact map that specifies which residues are near each other in three-dimensional space (and that therefore contribute to the energy) can be more useful than the backbone encoding. Such an encoding emphasizes the idea that non-local van der Waals or solvent-mediated interactions determine the energy (or scoring function) for proteins. While more complex than the interactions in magnets, these type of interactions still are local in space. There are 20 genetically encoded amino acids, so there are at least 200 possible interaction types compared to the \pm choice for magnets. Still amino acid side chains can be classified into groups whose members share similar chemical properties – big, small, acidic, basic, hydrophobic, hydrophilic, etc. Since Kauzmann, the hydrophobicity of the side chains has been known to be the dominant effect in determining protein architecture (Kauzmann, 1959). The other physicochemical aspects of the amino acids, as we shall see, are also relevant but this caricature of the interaction as being simply hydrophobic *versus* hydrophilic already makes it clear how frustration might arise in heteropolymers. Frustration arises from the combination of the constraints of fixed chain connectivity for a molecule with a defined sequence along with the heterogeneous energies for contact interactions that specify the greater attraction that hydrophobic residues have for each other than they have for hydrophilic residues. A problem of frustrating choice thus can arise when for example a single hydrophilic residue is surrounded in sequence by hydrophobes. The hydrophobes can lower their free energy by becoming buried, but then their nearby hydrophile must become buried too. If this lone hydrophile does not then find another hydrophilic mate in the interior (to form a salt bridge, for example), its burial will be costly in free-energy terms. The energy of the hydrophile could be alternatively chosen to be minimized by allowing the hydrophile to rise to the surface, but then if the backbones are too rigid, this choice will also bring the hydrophobes nearby in sequence up to the surface as well. We see this reconfiguration will be costly for the surrounding hydrophobes. Not everyone can be satisfied at the same time by an arrangement. In general, then low-energy states of a random heteropolymer with no specific sequential pattern of hydrophobes and hydrophiles will be the




Fig. 5. Specific folding becomes possible as sequence complexity increases. Two-dimensional lattice polymers and its associated energy landscapes are shown. The leftmost is a homopolymer. Many configurational states with nearly similar energy can be formed, for example by shifting the ‘strands’. The energy landscape is very rugged with local minima and many barriers. If the temperature is higher than a typical barrier, the polymer will be fluid; at lower temperature the system becomes trapped in many configurationally disparate states and the dynamics gets glassy. At the center a two-letter heteropolymer. The shifting of strands is more difficult as the register can be partially determined by the sequence composition. The energy landscape is still rugged but an overall shape appears. The barriers between any two low-energy configurations are high as energetically favorable contacts must be broken up in order to move. At right most a polymer with three types of monomers. The energy of the configurations can be better specified by the sequence and the energy landscape becomes funnel shaped. Adventitious trapping in high-energy states is reduced and folding to the ground state is robust to both sequence and environmental perturbations. Redrawn with the permission from Wolynes (1997a).

result of compromises. In addition, excluded volume and chain connectivity make it likely that the other different compromised structures that have low energy will actually be quite different in arrangement or topology. To avoid this frustration special patterns of hydrophobicity are needed like those seen in the helical wheel representation of helices (Schiffer & Edmundson, 1967).

Energetic frustration of the type just discussed is easiest to see and to quantify in the world of ‘lattice polymers’. Lattice polymers have been much studied by theorists because the simplicity in visualization and ease of computation of such models brings intellectual benefits that overwhelm our desire for realism. A concrete example from the lattice polymer world illustrates the power of simplicity, at least in pedagogy. In Fig. 5, some low-energy configurations of a polymer chain are laid out on a two-dimensional checkerboard lattice (Wolynes, 1997a). The left panel shows configurations of a homopolymer. If every pair interaction has the same value, each of these structures is nearly degenerate in energy. Only surface costs differentiate between the structures energetically. (of course in a very small system there is always a big surface to volume ratio!) Artificial homopolymers in bulk indeed choose among these degenerate states in a kinetically controlled way, making the materials science of solid polymers a challenging topic. We may say, with modest exaggeration that homopolymers slip and slide on a nearly featureless energy surface like that shown schematically below the homopolymer structure. This situation for the homopolymer is a little like what we found for the triangular antiferromagnet where we have an extremely degenerate landscape even near absolute zero temperature. The middle panel of Fig. 5 shows configurations of a binarily encoded heteropolymer. The energies of the various structures now vary a lot as they reconfigure, even their bulk energy contributions are pinned by the heterogeneity of contact energy. The landscape is rugged and two topologically very different configurations can compete at being possible ground states. The barriers between any two low-energy configurations are high since now many energetically favorable contacts must be broken up in

order to move from one low-energy state to the other low-energy state. The frustration of finding the lowest ground state comes from the conflict of many reasonably similar, but nevertheless highly compromised, choices. The energy landscape of a random heteropolymer is much like that for the magnetic spin glass (Bryngelson & Wolynes, 1987; Bryngelson *et al.* 1995).

Even when there are more than two interaction types, frustration, degeneracy, and slow dynamics remain the norm for complex heteropolymers. But having more choices in the interactions allows molecular evolution to pre-select landscapes to fold better. A canny ‘choice’ (read successful but rare evolutionary outcome!) makes it possible for a single structure to clearly emerge as ground state winner and no longer have significant competitors. An example of such pre-selection is the hydrophobic patterning of α -helices and β -strands that we mentioned earlier. In the lattice world such a pre-selected example is shown in the third panel (Fig. 5). The energy landscape for the ‘evolved’ sequence plotted below this structure now has a larger overall energy scale than for the random heteropolymer but the landscape is locally a bit smoother and at a low-enough temperature only a single structure will dominate. This is much like the disguised ferromagnetic order found at low temperatures in the spin models of associative memory.

Different degrees of frustration at the global level lead to different kinetic and thermodynamic behaviors for the heteropolymer. These different behaviors show up in direct simulations of heteropolymers in the lattice world (Onuchic *et al.* 1997). In Fig. 6, we show the results of Onuchic *et al.* for the folding times of two minimally frustrated lattice polymers (one with a two-letter code, the other a three-letter code) and one for a rather frustrated but typical two-letter code polymer. These are small systems and folding times of random heteropolymers scale exponentially with size so the contrasts in folding time would be bigger for longer sequences, but you can see the systems with less frustration are able to achieve faster folding under ideal conditions. The contrast is made even clearer when we simultaneously examine the lattice polymers’ thermodynamics: even at the temperature of its fastest folding the typical 27-mer is not thermodynamically stable in the folded state but instead takes up a myriad of unfolded configurations. The minimally frustrated sequences however still fold fast even at temperatures where their folded state wins out thermodynamically over the panoply of other possibilities. When there is a single ground state (as is true when there is little frustration) the global minimum of the lattice polymer can be accessed fairly quickly at low temperature, perhaps with a modest nucleation barrier largely entropy controlled still remaining that can slow folding. In this case, a single exponential rate characterizes the folding process. Folding occurs as a Poisson-distributed random event in what would appear to be a single step in a classic mechanistic investigation. For the typical sequence as shown in Fig. 6, there are several near-degenerate low-energy states. This degeneracy leads to kinetic trapping and the time course of folding therefore exhibits several distinct exponential phases. These phases signal the distributed escape from traps. Non-exponentiality is always a sign of a multiplicity of states and can also occur for minimally frustrated polymers. But for fast, minimally frustrated folders the non-exponentiality coming from intermediate states is not the result of frustration but reflects high energy states involved in direct diffusion to the ground state (down-hill folding) which typically occur on the microsecond time scale. For slower, but minimally frustrated folders with nucleation, longer-lived intermediate ensembles involve species where only part of the chain (a set of ‘foldons’) has become organized and the remainder remains disorganized (Panchenko *et al.* 1996). When there is partial frustration these intermediates that represent entropically stabilized traps can be further stabilized by allowing the formation of non-native but nevertheless stabilizing contacts. Not only does frustration influence the net stability, but also it changes the character of the thermodynamic transition. From the thermodynamic

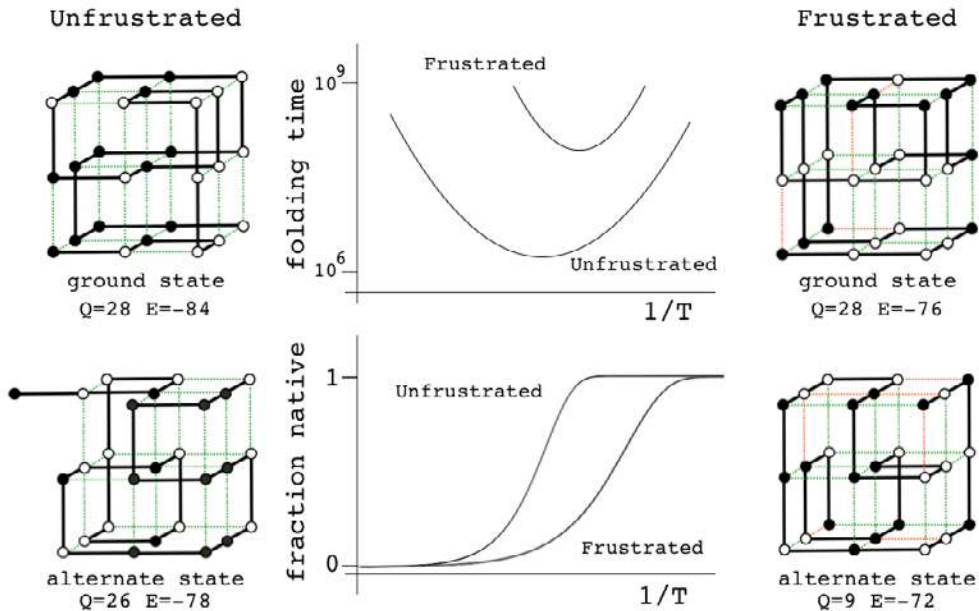


Fig. 6. Frustration in the lattice polymer world. Different degrees of frustration lead to distinct kinetic and thermodynamic behaviors. Here the 27-mer polymers fold into $3 \times 3 \times 3$ cubic lattices. If two monomers are adjacent in space there is an attractive interaction between them. This interaction is strong (-3 energy units, green) if the monomers are of the same type and weak (-1 energy units, red) if they are of different types. The most compact configurations are cube-like and have 28 contacts. The structures of the ground state of two sequences are shown at the top. At left an unfrustrated sequence where all interactions are optimal is in the lowest energy state. At right a sequence for which it is impossible to optimize all interactions in any structure, making it necessary to compromise and have weak contacts in the ground energy state. Q measures the similarity to the ground state as the number of native contacts, E is the pairwise added energy. For the unfrustrated sequences, most of the configurations with energy just above the ground state are very similar to the ground-state configuration. For frustrated sequences however there are configurations with energy just above the ground state that are very different from the ground-state configuration. When the system gets trapped in one of these low-energy states, it takes a long time to completely reconfigure before it can try to fold again, as schematized. The middle panel sketches the results from various folding simulations in the lattice world (Onuchic *et al.* 1997). At top the kinetics measured as the number of Monte Carlo steps required to reach the native structure. Below the thermodynamics of the same systems are shown. Even at its fastest folding temperature the frustrated sequence is not thermodynamically stable in the folded state but takes up other configurations. The unfrustrated sequences fold fast even at temperatures where the folded state wins out thermodynamically over the panoply of other possibilities. When there is a single ground state (as is true when there is little frustration) the global minimum of the lattice polymer can be accessed fairly quickly at low temperature. For the typical frustrated sequence there are several near-degenerate low-energy states, so the folding becomes thermodynamically reliable only at very low temperature. This degeneracy leads to kinetic trapping and the folding time course becomes non-exponential.

viewpoint when there is little frustration, the folding process appears very cooperative – a single very rapid transition occurs sharply with changing the temperature or solvent conditions as seen for the two or three letter minimally frustrated chains (Wolynes, 1997a). With growing frustration, the formation of populated intermediates corresponding to the near-degenerate ground states with non-native contacts causes the transition to be spread out over a range of temperatures so the change appears less cooperative. This is manifested in significant shifting baselines in the observable ensemble averages sometimes seen in equilibrium denaturation experiments.




Fig. 7. The global effects of frustration can be quantified with the energy landscape theory. A frustrated heteropolymer will have the energy distribution drawn from a random energy landscape, a Gaussian with some mean \bar{E} and a standard deviation ΔE . The energy of the lowest energy states E_g can be estimated given the size of the configuration space and the variance of the energy distribution. A large protein whose energy is just E_g will have quite a few kinetic traps of nearly the same stability. Therefore, if a protein is to fold robustly, its ground-state energy must be substantially below E_g . If the energy of the completely folded state E_f is substantially below this estimate, we can predict that frustration effects will be minimal for this sequence (Bryngelson & Wolynes, 1987). The condition that E_f is below E_g is known as the *gap* condition for folding. Comparing E_f and E_g and requiring them to be well separated (δE) provides one way of quantitatively stating the Principle of Minimal Frustration.

Unfrustrated heteropolymers, in other words, exhibit first-order transitions, resembling the freezing of a liquid into a crystal while highly frustrated heteropolymers, in contrast, exhibit thermodynamically continuous transitions that resemble the transition from a viscous liquid to a glass.

Energy landscape theory allows us to quantify the global effects of frustration. In a completely random energy landscape, a simple argument (made first in a more abstract setting by Derrida, 1981) allows us to estimate the energy of the lowest energy states by combining knowledge of the size of the configuration space to be searched with the knowledge of the statistical variation of energy between substantially different configurations (the variance of the energy). Let us call this typical ground-state energy E_g . If the energy of the completely folded state, E_f (for the particular sequence being studied) is substantially below the E_g estimate, we can predict that frustration effects will be minimal for this sequence. On the otherhand if $E_f \approx E_g$ there will likely be at least a few competitors for the ground state and folding will be slow, accounting for the time needed to escape from traps becomes necessary and the ensemble will change more continuously with thermodynamic conditions so folding will not seem very cooperative in equilibrium experiments. Comparing E_f and E_g , requiring them to be well separated (Fig. 7), then provides one way of quantitatively stating the principle of minimal frustration (Bryngelson & Wolynes, 1987). This principle generalizes an earlier idea due to Gō, the ‘consistency principle’ which suggested that secondary structure preference needed to be consistent with tertiary packing forces (Go, 1983). The minimal frustration principle requires also consistency between the different tertiary forces themselves. In addition, as we shall see, the quantitative formulation of the minimal frustration principle using landscape theory leads to practically useful algorithms.

It is worth reprising how one makes the random energy model estimate of the globally frustrated ground-state energy here, because later we will see how this estimate can be used to give a meaning to the notion of having specific locations for frustration in proteins. The random energy landscape argument is simplest for a large system. For a large frustrated heteropolymer we expect there to be, for any given configuration, both a large number of stabilizing interactions and

a large number of destabilizing interactions. We can thus invoke the central limit theorem of probability to conclude that the energy of a given configuration will be chosen from a Gaussian distribution having some mean \bar{E} and a standard deviation ΔE . Thus if we choose a configuration at random, the energy of that chosen specific state is a random variable having a probability distribution

$$P(E) = \frac{1}{\sqrt{2\pi\Delta E^2}} \exp - \frac{[E - \bar{E}]^2}{2\Delta E^2}. \quad (1)$$

The mean \bar{E} and standard deviation ΔE may depend on the properties that characterize globally the sampled ensemble of structures such as their degree of collapse and average secondary structure content. Both these statistical quantities can be estimated from an incomplete sampling of the states so long as it is a fair sampling. This makes the statistical characterization of the landscape feasible even while the exact enumeration of all possibilities would be impractical or even impossible owing to the exponentially large configurational space. We can also understand how these statistical properties vary in different parts of the energy landscape. For example, if contact pair interactions are the dominant energy contribution, the extended configurations (as measured by radius of gyration) will have very few contributing terms, while compact configurations will have many specific pairs in contact to be summed over. In this case, the mean energy \bar{E} depends on the expected number of pairs in that configuration, which in turn depends on the degree of collapse as does ΔE^2 .

$$\bar{E} = \bar{\epsilon}\rho N, \quad (2)$$

$$\Delta E^2 = \Delta\epsilon^2\rho N. \quad (3)$$

If compact configurations are generally stable, $\bar{\epsilon}$ is negative and mostly depends on the average hydrophobicity. In fact, compactness can be guessed quite well from knowledge of sequence composition alone. Intrinsically disordered proteins often are known not to be compact merely on the basis of their composition, which is generally insufficiently hydrophobic. Like $\bar{\epsilon}$, $\Delta\epsilon^2$ also depends on composition because it depends on how many energetically distinct pairs of interactions there are. Therefore also $\Delta\epsilon^2$ is a function of protein composition. $\Delta\epsilon^2 = 0$, for a homopolymer for example. Generally $\Delta\epsilon^2$ also grows with the number of different kinds of amino acids used to encode the protein.

The Gaussian distribution for the energies of states at first looks like it would allow the energy of a state to be arbitrarily low, but because there is a finite, albeit exponentially large, number of protein conformations, an infinitely low energy cannot be realized: there must always be some particular lowest energy state to satisfy the Third Law of Thermodynamics. What is the lowest energy we would expect then? The actual number of states that have energy E is a product of the total number of states regardless of energy (called Ω) times the probability of a specific energy E , $P(E)$. Thus,

$$W(E) = \Omega \frac{1}{\sqrt{2\pi\Delta E^2}} \exp - \frac{[E - \bar{E}]^2}{2\Delta E^2}, \quad (4)$$

Ω depends exponentially on the chain entropy $\Omega = e^{S_0/k_B}$. Since entropy is an extensive quantity, $S_0 = sN$, Ω is the exponential in protein size. This exponential scaling is the essence of the Levinthal calculation of the difficulty of folding on a flat energy landscape (Levinthal, 1968). Again we can account for how the number of chain configurations varies with several measures that stratify the energy landscape of the heteropolymers such as its overall compactness, average

amount of secondary structure, etc. A lot of interesting polymer physics is involved in such calculations (Koretke *et al.* 1996; Luthey-Schulten *et al.* 1995; Plotkin *et al.* 1996). Since both factors in W vary exponentially in protein size, for large proteins, we see that depending on the energy E , $W(E)$ is either exponentially large in N , suggesting that every finite part of the chain still has some finite number of energetically equivalent configurations or else W is exponentially small in N – there would not be expected to be a single viable configuration for any segment of the chain with this energy. The expected ground-state energy E_g then satisfies the equation $W(E_g) \cong 1$, or explicitly:

$$E_g = \bar{\epsilon}\rho N - \sqrt{\rho} \cdot \frac{\Delta\epsilon}{\sqrt{2}} \sqrt{S_0} N. \quad (5)$$

From this formula we see the expected ground state of random heteropolymers becomes deeper in energy, the higher the specific entropy of the chain and the larger the variety that is allowed in the choice of the interactions.

At the value of the energy E_g given by Eq. (5), the number of states is not, in fact, actually precisely one. This is because we have only kept the dominant terms in the exponentials. The number of traps at precisely E_g is known to scale polynomially in polymer length N , i.e. $N(E_g) \approx N^\alpha$. A large protein whose energy is just E_g thus still will have quite a few kinetic traps of nearly the same stability. Therefore, in order for a protein to fold robustly, its ground-state energy must be substantially below E_g on a per residue basis, otherwise trapping will become a problem. Without this gap in energies the actual ground state will be degenerate but perhaps only in short (intrinsically disordered?) segments. The condition that $E_f - \bar{E}$ is below $E_g - \bar{E}$ is sometimes known as the *gap* condition for folding.

The ratio $(E_f - E_g) / \sqrt{\Delta E}$ thus indicates how confident we can be, based on the statistics of the random energy landscape that there are no competing traps for the ground state. Since $E_g - \bar{E}$ is proportional to $\sqrt{\Delta E}$, this ratio is also related to the Z -score used in statistics, $Z = \frac{(E_f - \bar{E})}{\sqrt{\Delta E}}$. Quantitatively we can say the effects of frustration are minimal if the Z -score is very large. See Fig. 7 for a picture of the distribution of energy levels for a minimally frustrated protein. The number of standard deviations below the mean that are needed in order to get a gap depends on the specific entropy S_0 and chain length. S_0 can be estimated for compact polymers with secondary structure (Koretke *et al.* 1996; Luthey-Schulten *et al.* 1995). Note that if there are further *a priori* constraints lowering S_0 , one can get by with a lower Z -score.

The minimal frustration principle can also be stated in terms of the characteristic temperatures at which states with the energies E_f and E_g would be accessed at equilibrium. If there is a large stability gap the folded state will be accessed in a highly cooperative first-order transition at T_f . At this ‘folding temperature’ the entropy of the disordered chain is balanced by the deep ground-state energy. Clearly then T_f is roughly $(E_f - \bar{E}) / S_0$. If the heteropolymer is frustrated, on the other hand, the lowest-energy-trapped states can only be accessed at a temperature T_g , which is roughly $\Delta E / \sqrt{S_0}$. In terms of these characteristic temperatures we can say a heteropolymer is minimally frustrated if $T_f / T_g > 1$. We see that T_f / T_g also depends monotonically on the Z -score. One advantage of this way of stating the criterion for minimal frustration is that this criterion does not explicitly involve chain length. Also this criterion in terms of temperatures turns out to be fairly robust to there being correlations in the energy landscape, which are ignored in the rough argument we have just given (Plotkin *et al.* 1996). The correlations between energy levels predicted by the polymer theory certainly exist as they allow states to be grouped together into basins of




Fig. 8. Quantifying frustration in lattice world. Different degrees of frustration lead to distinct kinetic and thermodynamic behaviors. Folding simulations in enumerable systems have confirmed the criteria for achieving non-degenerate ground states and robust rapid folding. The configurational space of a two-letter polymer folding on a $3 \times 3 \times 3$ lattice can be exhaustively analyzed, so that the ground-state energy E_f can always be found, the energy distribution of all other compact structures can be computed explicitly and the relevant parameters of the energy landscape calculated. Each dot in the figure corresponds to a different sequence. The ratio of the energy gap $\delta E = E_f - E_g$ to the roughness $\bar{E}/\Delta E^2$ defines the relevant temperatures T_f and T_g . The average folding time at the constant temperature for each sequence is shown. At the same time gap/rough correlates with the cooperativity of the thermodynamic folding transition as measured by its width. For a lattice world heteropolymer to behave in a protein-like way, it must conform the minimal frustration principle requirement of having a sufficiently large T_f/T_g ratio as measured by the Z -score statistics. Seeing that this principle works in the controlled world of lattice proteins with known force laws gives us confidence in applying the minimal frustration principle to proteins in the real world, where the force laws are still uncertain. Redrawn with the permission from Mélin *et al.* (1999).

attraction which can interconvert without significant barriers. The number of such basins is much smaller than the number of individual configurations. The criterion in terms of T_f/T_g resembles very much the criteria that have long been used in metallurgy to distinguish those systems that easily crystallize from those alloys that form glasses (Chaudhari & Turnbull, 1978; Greet & Turnbull, 1967; Stevenson & Wolynes, 2010). The success of the idea in the macroscopic context of material science gives us more confidence in its use at the nanoscale of protein folding.

The criteria for achieving non-degenerate ground states and also rapid, reliable folding have also been checked many times by carrying out simulations in the lattice polymer world (Sali *et al.* 1994). The results of a very comprehensive study by Mélin *et al.* (1999) are shown in Fig. 8. The sequences in this study use a two-letter code with a reasonably well-chosen interaction threshold, so that ground-state structures are collapsed 27-mer cubes. Competing, perfectly collapsed structures can be completely enumerated explicitly (there are 103 346 of them) (Shakhnovich & Gutin, 1990), so that the ground state can always be found (and checked to see that there is no competing non-compact structures) and the average \bar{E} and ΔE^2 can be computed explicitly. The ratio $\Delta = (E_f - E_g)/\sqrt{\Delta E}$ can then be computed explicitly. We can see in Fig. 8 that the folding time at constant temperature correlates well with Δ and therefore T_f/T_g . The contrast in rates would be even larger at temperatures where the low stability sequences with low Δ would be thermodynamically stable. The dramatic slowing of folding of low T_f/T_g with cooling has been explicitly confirmed then in the lattice world. At the same

time, Δ correlates with the cooperativity of the thermodynamic folding transition as measured by its width δ . Clearly for a lattice model heteropolymer to behave like proteins do in the laboratory it must obey the minimal frustration principle requirement of having a sufficiently large T_f/T_g ratio or high Z-score. Seeing that this principle works in the controlled world of lattice proteins with known force laws gives us confidence in applying the minimal frustration principle to describe how proteins behave in the real world, where the force laws are still uncertain.

2. Frustration and the energy landscape of real proteins

2.1. The energy landscape of real proteins

The principle of minimal frustration has some very direct consequences for understanding folding mechanisms and kinetics in the laboratory. The most important consequence is that thermodynamic stability and folding kinetics should be correlated for real proteins of similar structure in much the same fashion as these thermodynamic and rate quantities are correlated for lattice models. This correlation between rates and equilibria while not completely perfect in the real world is an especial surprise since our introductory chemistry classes have drilled into our heads the idea that kinetics and thermodynamics are supposed to be completely distinct aspects of chemical reactions. Indeed in small molecule chemistry the relations between changes in rates and in equilibria are often unpredictable, but this is not true for the folding reactions of proteins because the overall folding reaction is itself a concatenation of many elementary steps each step being mostly near to local equilibrium. The smooth funnel-like landscape of completely minimally frustrated proteins thus implies a rate–stability correlation because the landscapes of minimally frustrated proteins should have few energetically generated traps. If there were many significant kinetic traps on the landscape the changes in the stabilities of the trap states when the protein is mutated would change the rate unpredictably and would thus interfere with finding a simple correlation of the overall rates with the stability of the ground-state configuration alone. In a minimally frustrated protein, the final native structure of the molecule itself is generally directly correlated with the energy landscape throughout the thermally accessed configuration space because the native-like interactions are stronger than random and thus are most common along folding paths. This structure–energy correlation is what we mean by there being a ‘funneled’ energy landscape for a minimally frustrated protein. We should remember, although, that the degree of rate/equilibrium correlation will not be uniform across the protein sequence. This non-uniformity of correlation arises even on a perfectly funneled landscape because some parts of the protein make few contacts or if they make many contacts a significant fraction of these contacts are distant in sequence so that forming them becomes entropically costly and thus unfavorable in the transition state for folding. At the same time other parts of the protein may have the bulk of their contacts rather close in sequence and thus these parts could easily form with lower entropy loss and thus will tend to form early. Studying the structural features of the rate/equilibrium correlation is called the ϕ -value analysis (Fersht *et al.* 1992; Matthews *et al.* 1987). ϕ -Value analysis reveals what are the most critical parts of the protein for the rate-limiting steps of the folding process.

The success of the ϕ -value analysis confirms the basic tenet of the landscape theory, the minimal frustration principle. This along with the controlled behavior seen in lattice model simulations gives us confidence (at a global level) that we can use the minimal frustration principle to think about natural proteins. These successes of energy landscape theory are partial, however. They do not tell us precisely what the physical interactions actually are that have allowed such an apparently low level of frustration to have been achieved through evolution. Also we would like

to know how perfect the funneling of the landscape of real proteins is. The level of agreement between purely structure based landscape models and experiment is so good that we cannot tell directly from the existing kinetic experiments how much energetic residual frustration if any actually remains in most real proteins. The data show the overall frustration level is small for most proteins, but how small precisely? Basically, the glassy effects on the thermodynamics and rates of most folding reactions are insufficient to provide an accurate value of the glass transition temperature T_g . A quantitative scale for the energetic ruggedness cannot be easily found using the measured folding kinetics in the laboratory alone. Nevertheless some inferences on ruggedness, quantified by T_g and perhaps more important, the energy of the typical competing near-ground states characterized by E_g , can be made by using indirect experiments on structure and dynamics of molten globules in concert with the random energy model and polymer physics (Onuchic *et al.* 1995). The key quantity that we need to know is the configurational entropy to be associated with molten globules. In other words, how numerous are collapsed structures that already have their secondary structure partially formed? Such ‘molten globule’ structures should make up the bulk of the decoy structures with which typical final native structures have had to compete energetically. Using the theory of helix-coil transition in collapsed states (Luthey-Schulten *et al.* 1995), the combination of collapse and helix formation has been shown to reduce the configurational entropy to about $0.6 k_B$ per residue from a nominal value of $2-3 k_B$ that would be available for an amino acid chain in an expanded random coil. Thus using the random energy model we can conclude that the ground-state energy of a configuration that might compete with the native structure is approximately $E_g = \bar{E} - \Delta E \sqrt{0.6}$. If the energy of the folded state E_f is lower than this value, the native structure will have few competitors. If, on the other hand, E_f is comparable to E_g , traps should appear. To quantify the possibility of there being a multiplicity of states, i.e. to measure frustration quantitatively, the ruggedness energy scale itself thus needs to be determined. This scale can be computed once a reliable energy function is given. We will use this approach throughout this review but it is worth noting that as an alternative we can try to estimate the roughness of the energy landscape from experiment by looking at fluctuation rates in molten globules, as was done rather early on. Motions in molten globule proteins can be measured via NMR. Such dynamical studies have been carried out for surprisingly few systems (Bruylants & Redfield, 2009; Dyson & Wright, 2002; Pervushin *et al.* 2007). Using some measured millisecond fluctuation rates in molten globules in this way Onuchic *et al.* arrived at an estimate for the T_f to T_g ratio for typical helical proteins. They obtained a value of T_f/T_g of about 1.6. Thus for a typical protein, glassy effects would only show up strongly at temperatures well below the freezing point of water, explaining why pure funnel models work so well in practice. This estimate of T_f/T_g must be considered quite crude but it is nevertheless higher than what one sees in the simplest lattice models using a two letter hydrophobic–hydrophilic code. Such lattice simulations therefore are plagued by trapping effects that are actually less common for real proteins. The low level of frustration suggested by this analysis for real proteins can be achieved for lattice proteins however by using a maximally distinct three-letter residue folding code. The pedagogically popular two flavor model (hydrophobic/hydrophilic) of lattice proteins is more frustrated, giving a lower T_f/T_g ratio but the three-letter code may be qualitatively fine for thinking about real proteins.

2.2. A license to learn via bioinformatics: structure prediction and energy functions

Given the difficulty that direct experiments have in quantifying frustration, a purely theoretical approach to estimating ruggedness is desired. A sufficiently reliable energy function, of course,

would allow us to assess computationally how funnel-like a protein landscape is and would allow us to quantify the degree of frustration that real proteins have by determining the free energies of a sampled variety of conformations. One is tempted to develop a computational approach to the ruggedness then by employing the all-atom energy functions that are routinely used to simulate the dynamics of folded proteins (Karplus & McCammon, 1983). Such a direct assault however faces several difficulties today. The first difficulty is that simulating completely atomistic models in their aqueous environment still takes a lot of computing time. Only in the last few years have the design of special purpose machines and the development of algorithms made it possible to carry out fully atomistic simulations in solvent covering the millisecond time scale motions. This is the time scale for accessing truly distinct conformational minima as inferred from NMR (Zerbetto *et al.* 2013). Mere sampling of the structures of molten globules is also insufficient: it is necessary to quantify the solvent averaged free energies (not merely energies!) of the individually sampled configurations – since we know hydrophobic effects arising from solvent averaging are important and these have a large solvent entropy component. Finally even when sampling and solvent averaging can be carried out, this direct approach to quantify ruggedness relies on the accuracy of the input all-atom force field. Alas, the accuracy of atomistic force fields while well documented to be good for configurations near to the folded state is, again for reasons of time scale, incompletely tested very far away from the native configurations (until recently!) (Piana *et al.* 2011, 2012). There are indeed many hints that many of the all-atom molecular force fields used today fail to properly describe the stabilities of incompletely folded states. One hint that there is a problem comes from a meta-analysis of folding simulations made by different research groups: there has been a distinct lack of agreement between the specific results obtained from different simulations of the early stages of folding carried out on the same proteins (Wolynes, 2004). There is often a tendency to ascribe these differences between results from different groups to the sampling protocols each different simulator uses but this is probably not the main source of error. Instead what we are seeing in most of these disagreements is that the landscapes encoded by the popular all-atom force fields (like the more obviously stylized two-letter code lattice models) are generally all more frustrated than the actual protein energy landscapes really are. Thus each landscape's particular set of traps is quite sensitive to model details much as would be expected for a random heteropolymer. The conclusion we may draw from these observations is that many current atomistic models are still too frustrated to be completely realistic to describe folding in detail. This conclusion is buttressed more directly by a recent study by the Shaw group of the thermodynamics of folding for a variety of all-atom force field models. They have shown that some tuning is needed in order to get reliable folding to the correct native structure across the spectrum of small proteins and that the better tuned models were more cooperative in their thermodynamic behavior than are earlier, less well-tuned, models (used however by many other groups) (Piana *et al.* 2012). Lower cooperativity correlates as we have seen with there being greater frustration and more intermediates. Even the most successful energy function studied by the Shaw group still gives less cooperative folding transitions than are found via laboratory calorimetry (Henry *et al.* 2013). Thus we can infer there probably remains still some room for improving all atom force field models. We also conclude that the direction of improvement will involve further reducing the frustration of the landscape. So today's atomistic simulations probably show larger frustration effects on kinetic pathways than really occur in nature. Recently Best *et al.* have shown that the Shaw simulations of folding do conform to a minimally frustrated landscape (Best *et al.* 2013).

Many of the difficulties of quantifying frustration using all-atom modeling would disappear if a coarse-grained but transferable energy model were available that was also sufficiently accurate. There is a tendency to think coarse-grained models intrinsically must be less accurate than all-atom models because they are clearly less complete in having fewer degrees of freedom explicitly treated. But, when it comes to the issue of having a realistic level of frustration of the landscapes this reasoning may be flawed. As we have seen the funneled nature of the landscape must emerge from evolution, but evolution itself has had only coarse-grained abilities to reduce frustration: mutations can change only a single amino acid at a time. In contrast, mutations are not able to change single atoms in proteins one at a time! If fluoroalanine would be the best fitting residue in a protein structure, natural evolution is out of luck if it has to find a way to try to put it in – there is no atomic level mutation facility in the cell. (Well, we might imagine post-translational modifications like phosphorylation are attempts at doing atomic level mutation.) This evolutionary meta-argument suggests therefore that a sufficiently complicated but still coarse-grained energy function should be able to capture at least the bulk of the evolved low frustration level.

One of the best ways to find a good transferable coarse-grained energy function is to use the minimal frustration principle itself. The correlation between structure and energy landscape envisioned in the minimal frustration principle gives a strategy for learning energetics from protein structural databases. The minimal frustration principle provides a ‘license to do structural bioinformatics’ since it implies there is a non-random part of the sequence–structure correlation. Beyond merely providing such a license, the minimal frustration principle, in mathematical form, motivates a number of specific algorithms for optimally learning about energetics from known structures (Goldstein *et al.* 1992b). These machine learning algorithms try to find energy functions that make the landscape of structurally well-characterized proteins as funnel-like as possible (Eastwood *et al.* 2001). These algorithms start by surveying the T_f/T_g ratios using various levels of statistical mechanical approximation to estimate the ratio. This must be done for many examples in the database of proteins with known structures but using a transferable energy function having a common set of adjustable parameters that encode the interactions between the residues which are assumed universal. Then the typical T_f/T_g ratios for the various different force fields can be compared. The more funneled potentials should have higher T_f/T_g . Thus the relative quality of two different coarse-grained force field models can be accessed, provided the T_f/T_g ratio can be estimated computationally with sufficient rapidity. The computer can furthermore carry out an optimization of the parameters used in the coarse-grained potential energy function to find the force fields that are most likely to fold properly. In this way, optimal protein-folding codes have been found (Goldstein *et al.* 1992b). In the lattice world, this learning strategy has been explicitly tested in a clever way. By first simulating lattice folding with a known energy function one is able first to construct an artificial database of lattice structures that can be folded with the assumed energy function. These structures can then be used as input to the learning optimization algorithm (Yue *et al.* 1995). This test shows the learning procedure does reproduce to rather good accuracy the original energy function that was used to evolve (design) the structures. This kind of study has established that the mathematics of the learning algorithm is sound.

Even before the lattice model tests of the learning procedures based on the minimal frustration principle were made, the idea was used in a practical way to find transferable potentials for natural proteins by Goldstein *et al.* (1992a, b). The strategy was first applied to associative memory Hamiltonians (Friedrichs & Wolynes, 1989). These are energy functions whose form resembles the associative memory magnet models that we discussed in Section II. The associative

memory function consists of a sequence-independent backbone interaction term along with a sequence-dependent ‘memory’ term that encodes particular structure–sequence correlations. $E = - \sum_{\text{memories}, \mu} \gamma(S_i S_j; S_i^\mu S_j^\mu) \Theta(r_i - r_j^\mu)$ The transferable coefficients γ depend on the sequence of the target protein to be folded S_i and the sequences of memory protein S_j , which are aligned to the target sequence.

Some of the alignments used in the memory set may be correct, but they may represent assignments in the ‘twilight zone’ where sequence identity by itself is so low that we cannot be sure if a supposed homolog is actually evolutionarily related since completely structurally unrelated proteins can also be similarly well-aligned based on their sequences alone. Other memories used in the associative memory model are simply misalignments to proteins that are incongruent in structure to the target but which are not recognizable *a priori* as being wrong simply based on sequence. The pairwise distances in the example proteins or memories ($r_i - r_j^\mu$) encode tertiary contacts, of course but they also encode more distant spatial relationships such as the distances involved in extended helices, strands or turns. The minimal frustration principle then is used first to find the optimal γ ’s for a training set of proteins. The resulting γ ’s which are supposed to be transferable are then used to construct Hamiltonians for proteins not in the training set. Free-energy profile analysis reveals that the resulting transferable energy landscapes were funneled to near-native structures. Folding simulations with these associative memory Hamiltonians perform well, showing that an algorithm based on learning the γ coefficients could ultimately recognize correct alignments even when the alignment by itself was clearly in the twilight zone.

The optimization learning strategy based on the minimum frustration principle has also been applied to off-lattice Hamiltonians with simple contact Hamiltonians (Koretke *et al.* 1996). The interaction parameters learned in this way turn out to be quite adequate to recognize correct structures from candidate homologs in the twilight zone by direct sequence–structure alignment (Goldstein *et al.* 1992a). Alignment recognition is much easier than molecular dynamics based recognition because the search space for alignment (hundreds of base structures) is smaller. In energy landscape terms, the smaller entropy of possible structures means that a bigger gap results between correct alignments (the ‘folded state’) and the best decoy alignment than would result from searching through all possible polymeric configurations. Still developing alignment Hamiltonians pinpointed the need to estimate better the T_f/T_g ratios in a self-consistent manner; the most troublesome decoys are obviously those that are produced as minima by the same Hamiltonian. This notion of self-consistent optimization is illustrated in Fig. 9 (Koretke *et al.* 1996). This makes the learning process used in the search algorithm iterative and non-linear but the learning process is still practicable since the development of parameters needs only to be done once. Such optimized but simple contact based Hamiltonians also work surprisingly well in off-lattice Monte Carlo structure searches. The server developed by Takada’s group (Jin *et al.* 2003), predicts structures quite well for small proteins. It is based on such a landscape optimized contact energy function but uses molecular dynamics and fragment assembly as search algorithms.

Seeking greater reliability and accuracy in structure prediction has shown that things are not so simple however. When studying larger protein complexes Papoian, Ulander and Wolynes discovered that the simplest contact-based energy functions while being adequate for some interfaces, were inadequate for describing the binding landscape of many interfaces between proteins (Papoian *et al.* 2003a). They found it was important to allow more distant water-mediated interactions between solvent separated residue pairs to contribute to the energetics (Fig. 10). The same

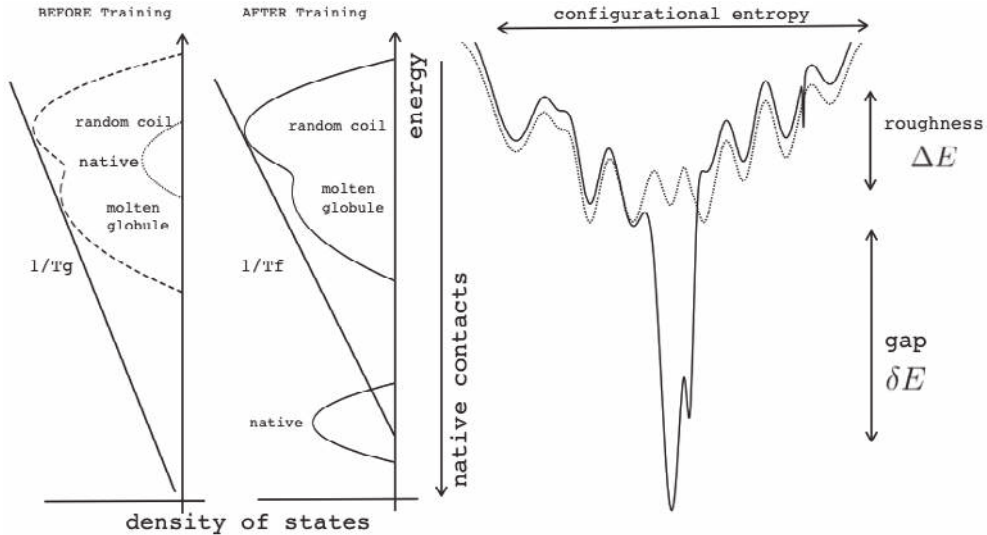


Fig. 9. Reliable energy functions can be learned applying the energy landscape theory. The correlation between structure and energy envisioned in the minimal frustration principle gives a strategy for learning energetics from protein structural databases. Machine learning algorithms find energy functions that make the landscape of natural proteins as funnel-like as possible. The more funneled potentials should have higher T_f/T_g . The parameter of a coarse-grained potential energy function is optimized self-consistently to find the ones most likely to fold properly. In this way, optimal protein-folding codes are found (Goldstein *et al.* 1992b). On the right, a representation of a realistic protein-folding landscape of a globular domain is shown. There are a large number of configurations at the top of the funnel that are nearly random coils, with few non-local contacts. As contacts are made, energy on the average are decreased. Non-specific contacts can be sufficiently favorable, a collapsed but fluid set of configurations becomes thermodynamically relevant and indeed may be a separate phase. At the left a sketch of the histograms of the energies of the configurations before and after training the energy function. The folded configurations are well separated from these disordered configurations by a stability gap from the molten globule.

parameter-learning strategy used for contact models can still be applied to this style of energy function. Structure prediction based on the resulting water-mediated energy function while not perfect, turns out to be reasonably robust and quite good for moderate size proteins (Fig. 11). This coarse-grained energy function used for *de-novo* structure prediction still appears to be more frustrated than the real folding energy landscape in so far as misfolded structures are still sometimes found via molecular dynamics, but the extra frustration largely comes from the ambiguities of the local-in-sequence structure information, which is used as input (secondary structure). Such information is encoded via associative memory terms. When a correct native secondary structure alone is used as the only memory input, very few intermediates are found in folding simulations and for most proteins folding is nearly as cooperative as folding processes in the laboratory. In this review, we will use this rather sophisticated but still coarse-grained energy function to quantify frustration. It is worth remembering therefore that the coarse-grained model we use to illustrate frustration effects in natural proteins still probably errs on the side of exhibiting greater frustration than is present in reality. This being the case, one should take most seriously then indications of very high levels of frustration from this model but probably one should not ascribe too much significance when there are only modest variations or changes in frustration level as measured by this energy function.




Fig. 10. Proteins *live* in water. Amino acid chains interact strongly with their environment that modulates the intrachain forces and conditions the global protein dynamics (Frauenfelder *et al.* 2009). Accurate and reliable force fields must include the solvent, at least implicitly. A sketch of a globular domain is shown on the top, where the protein is partitioned in layers, conceptualizing different interaction modes. An implementation of these ideas was realized with a transferable associative memory Hamiltonian energy function (Papoian *et al.* 2004). The different layers were calculated and projected on a real protein structure (lysozyme). The backbone chain is shown as a continuous trace; the lines correspond to interactions between amino acids. At left only the direct interactions are drawn (brown), next to only the water-mediated interactions (blue). At right an overlay of the molecular surface of the buried (brown) or exposed (transparent blue) residues is shown.

2.3. Localizing and visualizing frustration in natural proteins

Like the elementary landscape theory, the traditional phenomenological picture of folding kinetics using single rate coefficients provides only a global description. By global we mean that the protein as a whole is treated as one single entity having no localized substructure. This global phenomenological description works well for small, well-folded proteins but, even phenomenologically, the global picture begins to break down for larger proteins which are usually made up of identifiable domains that can fold independently as kinetic units. In multidomain proteins, rather often, the individual domains fold in separate steps rather than all at once. The same multistep folding behavior is seen in long repeat proteins (Werbeck *et al.* 2008). These observations impel us to the recognition that the forces in biomolecular energy landscapes act locally even though they have emergent global consequences for the kinetics. How deeply does the locally correlated nature of the dynamics go? The ϕ -value analysis of transition-state ensembles of even small proteins suggests that the more highly folded residues tend to be close to each other in the native state. A simple model of local folding on a minimally frustrated landscape, called the capillarity model (Ferreiro & Wolynes, 2008; Galzitskaya & Finkelstein, 1999; Wolynes, 1997b) likens




Fig. 11. Protein structure predictions. Energy landscape theory provides a mathematical framework to learn energy functions. The associative memory, water-mediated, structure and energy model (AWSEM) is a coarse-grained protein force field optimized with the energy landscape theory (Davtyan *et al.* 2012). AWSEM contains physically motivated terms, such as hydrogen bonding, as well as a bioinformatically based local structure-biasing term, which efficiently takes into account many-body effects modulated by the local sequence composition. With appropriate memories taken from local or global alignments, AWSEM can be used to perform *de novo* protein structure prediction. The figure shows structural alignments of the best ranking predictions for sequences of different sizes in which no *a priori* homology information is used. The predicted structures are shown in blue, and the actual experimental structure in green. The regions where the structures are highly similar as judged quantitatively via structural similarity of the contact maps are colored red for the predicted and orange for the actual structure (Sippl *et al.* 2012). Regions that are not accurate are shown in blue or green. While not perfect, the predictions are reasonably robust and quite good for moderate size proteins. The same model and energy function can also be used to predict the native interfaces of protein complexes (Zheng *et al.* 2012), thus it amounts to a flexible docking algorithm.

folding to crystallization of a small droplet. This (oversimplified) analysis suggests that a compact region encompassing about $4/27 \approx 15$ percent of a single domain protein must come to a near-native configuration before folding can proceed downhill to the native state in a free-energy sense (Wolynes, 1997b). Typically for small proteins such a compact core contains about 20 residues. Local excursions away from the native state can be monitored by hydrogen–deuterium exchange. These unfolding elements are of a similar size to the folding core (Weinkam *et al.* 2005). The size scale of both folding nuclei and unfolding units is then comparable to the protein segments encoded in the exons of eukaryotic genes (Gilbert, 1987). It is plausible then that regions of this size are ‘foldons’ (Panchenko *et al.* 1996), i.e. minimal folding units. These arguments suggest we can meaningfully discuss the level of frustration of the energy landscape of individual foldons rather than consider only the protein, globally, as a whole. Scanning contiguous regions of sequence using the Z-score criterion for the energy in the native configurations can thus be used to assign levels of frustration locally. By doing this, determining the Z-scores of contiguous segments, Panchenko *et al.* showed that there often is a rough correspondence

between those regions that were minimally frustrated based on their internal interactions and thus would be capable of selecting a unique structure on their own and the regions encoded by genetic exons (Panchenko *et al.* 1996).

It is hard to refrain from going still further down in length scale in characterizing the energy landscape in order to ask whether some particular physicochemical interaction between two specific residues in a protein is actually what gives rise to frustration, just as we did for the magnets, where individual plaquettes could be said to be frustrated or not depending on whether as isolated units these clusters of spins would have simple ground states. One must bear in mind though that such an analysis cannot be completely rigorous in proteins any more than it was for magnets: the global structure is determined by the cooperation of many elementary subunits; the clusters are not isolated completely. At the same time we also know that individual interactions in a protein cannot be tuned by mutation – instead several interactions change at once when a mutation is made at the amino acid level, all the interactions involving a given residue change when that residue is mutated. Still this degree of unavoidable spatial coarse graining suggests frustration can be defined at a pretty fine level. Recognizing the mutational limit on the length scale where minimal frustration can be expected suggests a reasonable, but by no means unique strategy for finding and locating the origins of frustration in particular proteins: we can computationally examine the free-energy changes that would be made upon introducing mutations and compare the energy of altered choices with what occurs in the original protein. Such a procedure actually generates a series of measures of localized frustration (Ferreiro *et al.* 2007). The approximate localization of frustration via this approach has proved to be helpful in visualizing regions of high frustration. The localized frustration analysis, while not being rigorous makes vivid the widespread, delocalized nature of the minimally frustrated interactions that do, in the end, allow proteins to behave globally as minimally frustrated systems with smooth funnels.

When trying to assess frustration locally we are ineluctably breaking up the full energy function of the protein into parts. As we have said, this step has some arbitrariness in it. Nevertheless the breakup of the energy function seems straightforward conceptually when the force field is written in the first place explicitly as a sum over pairs. One can then imagine assigning contacts as being frustrated or not just as we did for the few letter lattice models of heteropolymers that we discussed earlier. The best coarse-grained energy functions are not pairwise additive, however. For non-pairwise additive force fields local frustration indices can be obtained using the *in silico* mutational strategy postulated above. A local frustration index should quantify how much a residue or residue pair contributes to the energy in a given structure compared to what it would contribute in a typical decoy or molten globule configuration. Making numerous mutations and changing local environments let us assess the mean and variance of the energies of molten globule configurations (decoys) relative to the native. By normalizing the difference between the native and the average decoy by the contribution of the same residue (or pair of residues) to the variance of the energies of the decoys we thereby can get an idea whether that contribution to the energy is typical of what would be expected in a minimally frustrated protein or more like that in a random unevolved heteropolymer. In the minimally frustrated protein, it is necessary that globally E_{Total} is less than E_g . Requiring each part of the interaction energy to satisfy a similar constraint, $E_{ij}^N > E_{ij,G}^{\text{random}}$ is of course overly stringent since only the total energy needs to be low, so we would expect (and find) that only a fraction of individual sites or of the interacting pairs would be minimally frustrated by themselves when we do not consider their cooperation with their neighbors.

At the level of a single residue, a *frustration index* can be assigned to each residue via such a set of mutations as

$$F_i = \frac{E_i^{T,N} - \langle E_i^{T,U} \rangle}{\sqrt{1/N \sum_{k=1}^n (E_{i'}^{T,U} - \langle E_{i'}^{T,U} \rangle)^2}}. \quad (6)$$

Here $E^{T,N}$ is the total energy of the protein in the native configuration, taken as $E^{T,N} = \sum_{k \neq i}^N (E_{\text{contact}}^{i;k} + E_{\text{water}}^{i;k}) + E_{\text{burial}}^i$, according to the tertiary interaction terms of the associative memory Hamiltonian with water mediated interactions (AMW) energy function (Papoian *et al.* 2003a). This energy considers all the interactions that residue i makes with residues k , either in a direct contact, $E_{\text{contact}}^{i;k}$ or in a water-mediated interaction, $E_{\text{water}}^{i;k}$ and via a single-body burial energy term, E_{burial}^i . The average energy of the decoys $\langle E_i^{T,U} \rangle$ is computed by mutating residue i to every other possible residue. As the 20 genetically coded amino acids are not all equally probable, the decoy energy is calculated with weights according to the amino acid composition of the chain. These mutations are evaluated from the sequence-specific contact and burial terms from the AMW force field with parameters $\lambda_i, r_{i,k}, \rho_i$ that correspond to the amino acid identity, interaction distance, and density, respectively (Papoian *et al.* 2003a). Similar recipes could be used for other coarse-grained energy functions.

In the case of pairs of residues, we ask: how favorable is the actual native pair relative to other possible interactions? To compute the frustration index for interacting pairs of amino acids ij simultaneous mutations on residues i and j are made. We have proposed two related but complementary ways for localizing frustration at the pairwise contact level. These ways differ in how the set of decoys is constructed. In one choice, the decoy set is made by randomizing only the identities of the interacting amino acids ij , keeping all other interaction parameters at their native value. This scheme effectively evaluates every possible mutation of the amino acid pair that forms a particular contact in a robustly fixed structure. We call the resulting index the ‘mutational frustration’:

$$F_{ij}^m = \frac{E_{ij}^{T,N} - \langle E_{ij'}^{T,U} \rangle}{\sqrt{1/N \sum_{k=1}^n (E_{ij'}^{T,U} - \langle E_{ij'}^{T,U} \rangle)^2}}. \quad (7)$$

The decoy energy distribution is calculated by randomly selecting amino acid identities from the protein composition and fixing the density ρ_i and the pairwise distances $r_{i,j}$ to the native conformation. It is worth noting that the energy change upon pair mutation not only comes directly from the particular contact probed but also changes through interactions of each residue with other residues not in the pair, as those contributions may also vary upon mutation. One advantage of the mutational frustration index is that, in principle, this local measure of frustration also could be experimentally determined in the laboratory by combinatorial protein engineering.

A second way of quantifying pairwise local frustration imagines that the residues are not only changed in identity, but also can be displaced in location: how favorable is the native interaction between two residues in the native structure relative to other interactions these residues could form in globally different distinct compact structures? The energy variance thus reflects contributions from the energies of molten globule conformations of the same polypeptide chain. For this index, specially suitable for examining alternative tertiary structures, the decoy set involves

randomizing not just the identities but also the distances $r_{i,j}$ and densities ρ_i of the interacting amino acids. Given a protein sequence and structure this strategy compares the energy of each native contact with the distribution of a randomly interacting set.

$$F_{ij}^c = \frac{E_{ij}^N - \langle E_{i',j'}^U \rangle}{\sqrt{1/N \sum_{k=1}^n (E_{i',j'}^U - \langle E_{i',j'}^U \rangle)^2}}. \quad (8)$$

When $E_{ij}^N = \langle E_{i',j'}^U \rangle$, then the native energy would not be discriminated from the typical energy of a random interaction in the molten globule and $F_{ij}^c \approx 0$. This scheme effectively evaluates the native pair with respect to a set of structural decoys that might be encountered in the folding process. We call the frustration index computed in this way the ‘configurational frustration’.

2.4. Patterns of frustration in globular proteins

In visualizing frustration in proteins, it is useful to simplify the information of the frustration analysis into a discrete form like that used for magnets or lattice models where an interaction can be said to be frustrated or not. Given the continuous nature of the frustration index, physically motivated but ultimately non-unique choices of cutoffs must be made. It is particularly interesting to highlight those interactions that are by themselves minimally frustrated – that is those whose energy is deeper than we would expect in the ground state of a random heteropolymer. Using the entropy estimate for molten states of $0.6 k_B$ alluded to earlier; minimal frustration corresponds to having an F_i index larger than 0.78 (Ferreiro *et al.* 2007). Many of the remaining contacts are neutral, that is they lie near the center of the distribution of possible energies in compact decoy states. While these are important interactions they are so numerous that they would obscure images of the deeper stabilizing contacts. On average if they are consistent they can reinforce the funneling of the landscape but weakly. On the other hand, some contacts appear highly frustrated, that is they are not only somewhat destabilizing but are also, in fact, unusually destabilizing in the native structure. We define highly frustrated contacts to be those having high energies 1σ above the mean. Such interactions would seem to contribute little to the specificity of the folding to the native structure. So why are they there?

To get a feeling for the patterns of frustration, in Fig. 12 we show images of the frustration pattern in the prototypical protein lysozyme. One can see that the protein is networked by a connected set of contacts all of which are minimally frustrated (green) while there are two patches of highly frustrated contacts (red) along with other neutral interactions that are not shown. The frustrated contacts are near to the active site of this protein, which functions as a glycoside hydrolase enzyme.

The bulk of the rest of this review will focus on how frustration patterns of folded proteins can be correlated with folding dynamics and functional aspects of protein dynamics for many specific systems. Before studying particular cases however it is worth making some general observations about the prevalence of frustration when surveying the proteomic database synoptically. First a surprising number of direct contact interactions in native protein structures are individually minimally frustrated. On average about 48% of the direct interactions are individually minimally frustrated and another 46% are neutral. It is easy to see then why proteins as a whole behave as if they are minimally frustrated globally. This also makes clear why the notion used in structure-based folding simulations (that nearly all native contacts are stabilizing) is a good first approximation to an evolved protein’s energy landscape. Only about 6% of the direct contacts turn out to be




Fig. 12. Proteins can be locally frustrated. A sketch of a globular domain is shown on top, with the protein partitioned in regions conceptualizing local frustration patterns. An implementation of this idea can be quantified with a local Frustration Index, calculated on real protein structures with the AMW force field (Papoian *et al.* 2003a). Below the schematic, the results obtained with the prototypical protein lysozyme. The backbone is shown as a continuous trace. The protein is networked by a connected set of minimally frustrated contacts (green), while there are two patches of highly frustrated contacts (red). Neutral interactions are now drawn. Direct interactions are shown with continuous lines, and water-mediated interactions with dotted lines. The surfaces of the amino acids are overlaid, according to their frustration index calculated at the single-residue level. Local frustration patterns were calculated with the frustratometer.tk server (Jenik *et al.* 2012).

highly frustrated. The indirect water-mediated interactions are seemingly less important in guiding the protein to its native state; only 26% of these are minimally frustrated by themselves while a much larger 61% are neutral, leaving 13% highly frustrated. The contribution of water-mediated interactions to global landscape funneling is far from negligible, however. Minimally frustrated residues are generally more buried than highly frustrated residues and are more likely to be found in α -helices or β -sheets rather than in turns or bends. It is possible that some turns actually relieve their frustration in more complex ways not visible in a frustratogram. Side-chain-backbone contacts which are commonly made in ‘turn signals’ may relieve frustration in a way not captured by the coarse-grained energy function (Ferreiro *et al.* 2007).

The pair distribution functions of minimally frustrated and highly frustrated residues are different from each other and are distinct from the pair distribution of all pairs irrespective of the frustration level. As can be seen in Fig. 13 the minimally frustrated residues are more strongly correlated at large distance than are residues in general. This long-range correlation allows the minimally frustrated contacts to provide the main rigidifying elements of the structure. In contrast, highly frustrated contacts are strongly clustered and localized, not spread out: whole localities exhibit a high level of frustration: the question of why these frustrated regions exist takes up much of the rest of this review.

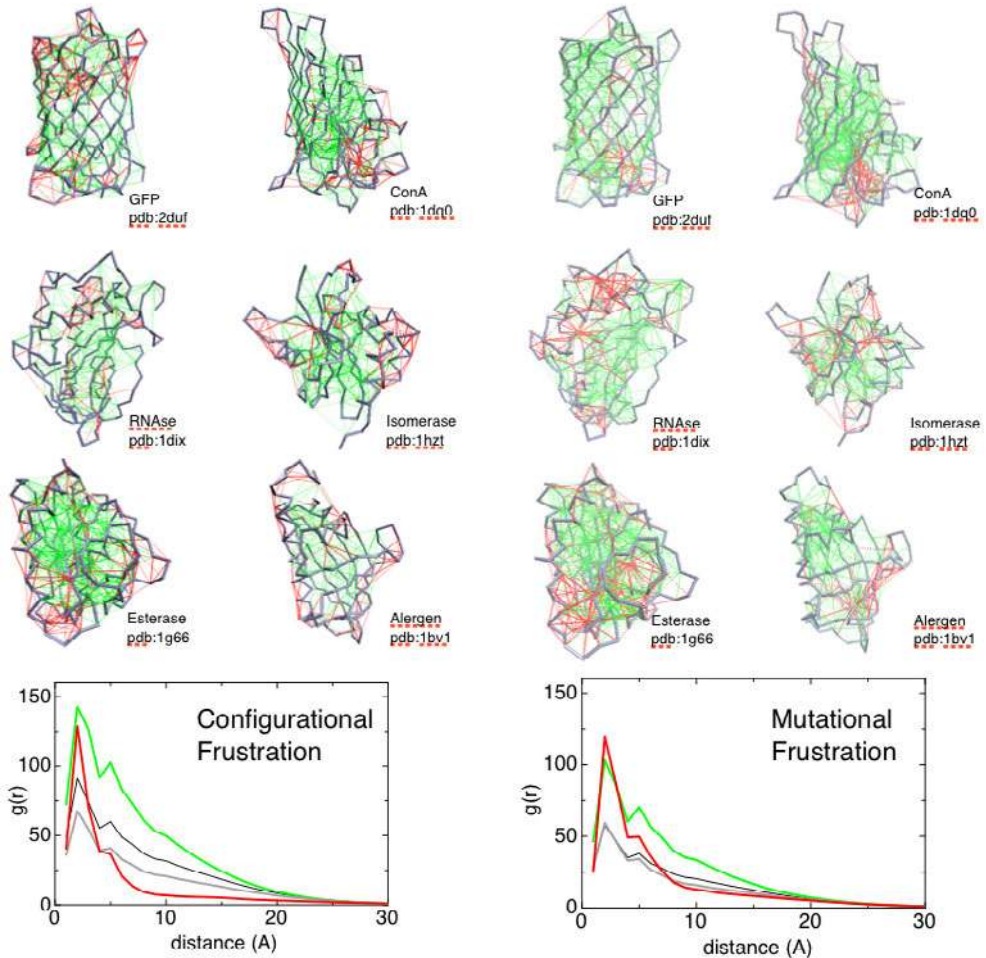


Fig. 13. Local frustration patterns in globular proteins. Examples of *frustratograms* of natural proteins are shown. The structural coordinates from protein models were taken the protein data bank (pdb) (Berman *et al.* 2000). The chain backbone is shown as a continuous trace and interactions between residues are represented by thin lines. Minimally frustrated contacts are green, highly frustrated contacts are red, and water-mediated interactions are shown with dotted lines. Two complementary ways of localized frustration are shown for each protein. At right the patterns are calculated with the *mutational frustration* definition: how favorable is a native interaction relative to the interactions other residues would form in those locations? At left the patterns are calculated with the *configurational frustration* definition: how favorable is the native interaction relative to the interactions these residues would form in other compact structures? A statistical survey was performed with a non-redundant sample of the pdb (Ferreiro *et al.* 2007). The pair distribution functions between the centers of mass of the contacts are shown below for the minimally frustrated (green), neutral (gray), or highly frustrated (red) groups. Neutral interactions follow the distribution of all contacts (black); highly frustrated interactions cluster at short distances; and minimally frustrated interactions form a tight network that percolates the structure.

3. Frustration in the evolution of natural proteins and in protein design

3.1. Frustration and the complexity of Nature's amino acid code

Theoreticians aptly explain many aspects of the fundamentals of protein folding by studying simplified models. Sometimes seen as surreal 'spherical cow' models, these simplifications yield

testable and communicable hypotheses such as those uncovered in the lattice worlds discussed above or that one finds using structure-based models (Clementi, 2008). These are thought of as exaggerations because such extreme simplifications are difficult to construct in the laboratory, although some amusing literally macroscopic models of proteins have been built (Reches *et al.* 2009). In working with real proteins though, one obvious simplification route would be to somehow reduce the apparent complexity of the amino acid sequences. Twenty amino acid types seems like a lot, at least to first year graduate students. It is still an intriguing fact that natural protein sequences appear to be random (Weiss *et al.* 2000), while most random amino acid sequences do not appear to be proteins (Keefe & Szostak, 2001). By most statistical tests natural protein sequences could encode huge amounts of information, so it has never been obvious which signals are strictly necessary and sufficient for folding and which are simply relics of protein natural history (Zuckermandl & Pauling, 1965).

Guided by the 'sequence codes structure' hypothesis, researchers have attempted to simplify the coding by replacing the apparently redundant amino acids, so that proteins can be made via a simpler palette of possibilities. Information redundancy could be experimentally traced by investigating the structure and function of simplified proteins with some of the apparent redundancy removed (Clarke, 1995). Simplified sequences have been made by simultaneously mutating multiple residues to a single amino acid type. When BPTI protein was mutated by replacing 29 of the total 58 amino acid residues with alanine, it was found that the sequences could still fold (Kuroda & Kim, 2000). The effects of a simplifying substitution depend on the position of the mutation in the sequence. Nevertheless, when properly located, simplifying a site only marginally changes overall folding stability of the domain. Most mutants were found to be thermodynamically destabilized and cumulative mutations appear to have additive effects, suggesting that the sites that were chosen to target for mutation do not strongly contribute to the cooperative folding core of this domain (Kato *et al.* 2007). When additional substitutions were tried, the protein became too unstable and usually a molten globule state (i.e. an ensemble of conformations) was populated. This is analogous to the situation described for the lattice model proteins where the overall destabilization of the fold relative to the disordered ensemble occurs when the variety of residue types is reduced. While the native stability often decreases because some favorable native contacts are not formed, the possibility of mis-matches of the homopolymer sections also increases. These effects can be seen as manifestations of frustration in the landscape. One key ingredient is the placing of the simplified sites where alanines were already shown to be tolerated. The point mutation A16 V in an exposed loop of BPTI severely destabilizes the fold but causes no significant conformational change in the mean native structure when folding does occur. In contrast an eight alanine substitution in the protease-binding site does not severely affect stability, but does stabilize other conformations (Cierpicki & Otlewski, 2002). When the effect of the substitution cannot be attributed to structural deviations, some kind of strain is usually invoked by protein engineers. Strain, tension, conflict, and unsatisfaction are all terms that reflect frustration effects on protein folding. Tinkering with protein stability by multiple alanine substitutions has also revealed strain effects in the classical models of T4 lysozyme (Liu *et al.* 2000), and ARC repressor (Brown *et al.* 1999). These studies suggest that the tolerance to mutation is not evenly distributed, but instead is rather polarized in these structures, both in the native stability and contributions to the rate-limiting steps (Gassner *et al.* 1999). A series of multi alanine substitutions carried out in single contiguous stretches show that the lysozyme fold is robust to most changes in helices, turns and loop regions but that some substitutions can lead to strong effects, revealing that local signals such as helix capping act to reduce the frustration in

the natural protein (Zhang *et al.* 2002). When the T_f/T_g ratio is large enough, the protein is robust to substitutions. As the energy gap is lowered, the conflicting effects show up, sometimes catastrophically.

Going beyond making substitutions of a specific and precise character, combinatorial libraries have been exploited to search for foldable sequences using reduced sets of amino acids. These libraries provide a rich source of diversity for the discovery of novel proteins. At the same time these studies begin to answer questions on the minimal set requirements, while of course the practical size of the screening libraries is a concern (Bradley *et al.* 2006). Is the existing genetically encoded set of 20 amino acids the minimal set needed to fold all proteins? As natural protein sequences usually derive from common ancestors and the existing fold classes are believed to be limited (Brenner *et al.* 1997), robust identification of fold class has been used as a constraint to test for whether reductions of coding information is possible. It has been shown that coding alphabets reduced up to 12 letters can still detect 90% of the fold classes. Further reduction rapidly degrades the signal in homology detection (Murphy *et al.* 2000). The success of any reduction approach critically relies on how the residue grouping is done, thus different scoring matrices for alignments are obtained when the criteria changes (May, 1999). It may well be possible that the optimal grouping of the amino acids into fewer types depends on the fold class of the structures, and that the observed natural set of amino acids is the actual minimal one that would satisfy all the functional constraints. There is still no absolute standard against which to compare the groupings of amino acids, and chemical intuition has been the most common inspiration for reduced alphabets.

Energy landscape theory shows that increasing in the alphabet allows for the possibility of decreasing the ruggedness, while still maintaining the stability gap. Hence, bigger alphabets allow for more robust encoding of various specific folds. Once a single overall topology is specified, successful design using few-letter codes becomes easier. When the overall frustration of the starting sequence is low, competing structures are automatically destabilized by the landscape topography. Some of the complexity of protein sequences can then be reduced, as shown in several experimental cases. Riddle *et al.* selected a functional SH3 domain sequence that contained 95% of its amino acids reduced to a five-letter alphabet of Isoleucine-Lysine-Glutamic acid-Alanine (IKEA) and G. Physical measurements on this protein suggest that the simplified protein adopts native-like conformations (Riddle *et al.* 1997) and conserves the overall folding dynamics (Yi *et al.* 2003). Notably, this group also reported a valuable negative result: they were unable to find functional SH3 proteins when the alphabet was further reduced to the three-letter code IKE (Riddle *et al.* 1997). While success has been reported by other groups in reducing α -helical proteins to three-letter codes (Beasley *et al.* 1997) (see below), this apparently is not possible for the β -sheet-rich SH3 topology. Frustration in tertiary interactions can be at play. Figure 14 shows the local frustration patterns of models of this SH3 domain for various simplified sequences. It is clear that just as for the lattice model protein (Fig. 5), the density of minimally frustrated interactions diminishes as the complexity of the alphabet is reduced (Fig. 14). Also coding simplification leads to the appearance of more highly frustrated interactions. Frustration between alternate states limits the ability to reduce coding complexity. A specific structure cannot be encoded in a homopolymer. An extremely simple alphabet nevertheless allows for encoding the simplest, often highly symmetrical folds. Increasing the variety of interactions makes minimally frustrated structures easier to find even when the topology is intricate. Some of the apparent complexity of protein sequences thus arises presumably not just as a confounding result of historical accidents within a highly degenerate folding code, but may actually be needed in order to encode a suitably




Fig. 14. Is it possible to re-encode a protein structure with a reduced alphabet? Once a single overall topology is specified, successful design using oligo-letter codes becomes feasible. When the overall frustration of the starting sequence is low, competing structures are automatically destabilized by the landscape topography. The local frustration patterns of SH3 proteins with simplified sequences are shown. A natural protein is shown at the far right (pdb: 1RLQ), and next to it simplified sequences threaded on same backbone. The five-letter code is the sequence that Riddle *et al.* found by selecting functional SH3 domains (Riddle *et al.* 1997). This sequence contains 95% of its amino acids reduced to a five-letter alphabet of IKEA and G. The number of highly frustrated interactions increases from 3 to 7%, maintaining the minimally frustrated core of interactions. Further reduction of this same sequence did not produce folded SH3 proteins. A specific structure cannot be encoded in a homopolymer. Energy landscape theory shows that increasing in the alphabet allows for the possibility of decreasing the ruggedness while still maintaining the stability gap. Bigger alphabets allow for robust encoding of various specific folds.

wide range of structures with sufficient discrimination in their energy landscapes (Wolynes, 1997a). Within a rich alphabet much more varied structures become ‘designable’, as many sequences can satisfy the constraints and minimize frustration.

3.2. Frustration in artificially designed proteins

Protein design is a promising endeavor that still can frustrate protein engineers. The moving parts of macroscopic machines can be unambiguously cataloged and combined. This is harder to do for proteins where independent parts may be distinguishable on the larger scales such as domains and multimolecular assemblies, but even at the large-scales cooperative effects often surprise us. At the lower levels of engineering single domains, combinatorial libraries of *de novo* amino acid sequences provide an overly rich source of diversity. Even for a relatively small protein, it is impossible to sample all possible sequences, as for a chain of 100 residues composed of the 20 amino acids, there are about 1.2×10^{130} possible polymers. This is a big number. A collection containing one molecule of each sequence would fill a volume larger than Avogadro’s number of universes (Beasley & Hecht, 1997). As randomly generated sequences rarely fold, the effective search space of protein design must be reduced. Even strong simplifications yield libraries still too vast enough to make exhaustive search impractical. The alphabet reduction approaches described in the previous section therefore become a necessary alternative. A ‘binary code’ for protein design was a pioneering reduction to construct libraries where folded proteins were likely to be found (Kamtekar *et al.* 1993). In these libraries, encoding up to five polar and six non-polar amino acids, families of four helix bundles were found to be well ordered, and then could be successfully modified to introduce cofactor-binding, catalytic activity, and other functional properties producing some of the first protein-based nanomaterials (Smith & Hecht, 2011).

Four helix bundles based on only three amino acids have been produced, exemplifying that symmetrical structures are more designable than asymmetric ones (Wolynes, 1996).

The principle of minimal frustration provides a framework for practical engineering questions (Onuchic *et al.* 1997). The essence of landscape engineering is that, in its simplest form, the theory requires quantifying only a few parameters: E_s , the stability gap between the ground (native, folded) state of the protein, E_d , the mean of the excited (misfolded) states and δE , the roughness of the energy landscape. Natural foldable proteins maximize T_f/T_g , suggesting a physically meaningful criterion in the design of energy functions for protein folding *in silico*. The essential feature of a potential function is that the energy of a sequence in the native structure measured from the mean of the misfolded states is much larger than the width of that distribution. In such a case, the folded conformation will automatically have energy lower than that of all alternative conformations.

Inspired by these principles, the Takada laboratory developed an automated procedure to design sequences to fold to a given simple structure (Jin *et al.* 2003). Folding simulations of a natural (target) fold were performed and a set of misfolded structures or ‘decoys’ was selected. Design requires crafting a global funneled landscape by figuring out how to design out a vast number of such misfolded configurations. This can be done with statistical energy landscape theory. Sequences were searched to either simply minimize the energy in the native state or to maximize the energy gap between folded and misfolded forms as landscape theory would endorse. It was found that the sequences optimized taking into account the gap display stable secondary and tertiary interactions. The one whose design was based on simply getting a low-energy native structure did not show a well-dispersed NMR spectrum, suggesting that it did not have a unique fold (Jin *et al.* 2003). Figure 15 shows the local frustration patterns of the designed and the natural proteins. All three actually have a large fraction of minimally frustrated interactions, with the energy-optimized sequence presenting an unusually high number of them. The largest cluster of highly frustrated interactions in the natural protein coincides with the albumin-binding site. This frustrated region (apparently evolved for function) is absent in the designed sequences. To explore the effects of the tertiary energetic frustration in these systems, Truong *et al.* performed folding simulations and structural predictions using AWSEM (Truong *et al.* 2013). They found that small differences in frustration level do impact the quality of the predictions. The fold of the natural sequence is indeed best predicted by simulated annealing, unlike the designed sequences that show lower fidelity in a structure prediction. Unlike the nicely funneled landscape of the natural sequence, landscapes of the designed proteins display complicating features. The main difference between the designed proteins and the natural ones was indeed related to the gap-criterion used for design. The natural sequences turn out to display the largest difference between the native energy and the energies of decoys (Truong *et al.* 2013).

A very nice example showing the possible breadth of protein design was realized by the Baker laboratory in their creation of Top7, a protein with a non-natural fold (Kuhlman *et al.* 2003). The design strategy started with a trace of the topology for which fragments of naturally occurring sequences in other proteins were selected to match the desired secondary structure elements. The sequences were then subjected to further cycles of mutation and selection with the computer scoring the energy of the sequences in the target structure. The selection criteria were based primarily on minimizing the total energy, with some few additional restrictions on the sequences in order to maintain exposed polar residues in an attempt to avoid aggregation. The resulting sequence had no significant similarity to any natural sequence but was shown to fold to a monomeric soluble form whose crystal structure matched the targeted goal (Kuhlman *et al.* 2003).




Fig. 15. Are humans too good protein designers? Although it is still challenging for routine technological applications, protein design has had some success. Top7 is a protein with a novel fold, designed by taking a trace of a target topology and evaluating sequences that minimize the energy in that configuration. The resulting protein is unusually unfrustrated, much less than natural proteins of similar size and secondary structure such as Ribosomal S6. Folding simulations show that S6 folds smoothly in a two-state manner in contrast with Top7 that populates intermediates, even on a perfectly funneled surface (Truong *et al.* 2013). The high stability of the fold is consistent with the original design, but competing low-free-energy structures appear because the sequence was not explicitly designed to avoid them. Energy landscape principles inspired the development of an automated procedure to design sequences by crafting a global funneled landscape to design out a vast number of misfolded configurations (Onuchic *et al.* 1997). Sequences are searched to either simply minimize the energy in the native state or to maximize the energy gap between folded and misfolded forms. An example of a natural albumin-binding domain and two redesigned sequences are shown. The small differences in the frustration level affect the quality of folding (Truong *et al.* 2013). Unlike the nicely funneled landscape of the natural sequence, landscapes of the designed proteins display complicating features, indeed related to the gap-criterion used for the design. Redrawn with the permission from Truong *et al.* (2013).

Study of the folding kinetics of designed proteins showed that Top7 folds fast relative to natural proteins of similar size. Nevertheless Top7 displayed complex multiphasic kinetics unlike its natural counterparts (Scalley-Kim & Baker, 2004). The effect was judged to come from slow interconversion of collapsed folding intermediates that become populated, a sign of frustration in the landscape. Apparently not all protein sequences that can be crystallized have energy landscapes as smooth as those of natural proteins. Perhaps crystal formation itself helps to separate different conformers. The source of frustration was suggested to be in the extreme regularity of the β -strands, which could favor out-of-register pairings. It is also clear that even the basic topology of Top7 makes for its non-cooperative folding (Truong *et al.* 2013; Zhang & Chan, 2010). Local frustration analysis indicates Top7 is unusually unfrustrated in an energetic sense. It has 60% of the interactions classified as minimally frustrated and only 6% highly frustrated. This is much less frustrated than are natural proteins of similar size and secondary structure (Fig. 15). Folding simulations of Top7 and a cognate S6 protein with a non-additive structure-based model show that the topological constraints make the natural counterpart fold smoothly in a two-state manner for a fully unfrustrated landscape, while in contrast the Top7 topology allows there to be populated intermediates even on a perfectly funneled surface (Fig. 15). These kinetic topological traps have a polarized structure, with the C-terminal fragments mostly folded with clusters of minimally frustrated interactions, unlike the transition state of the natural protein that is structurally homogeneous (Fig. 15). The high stability of the fold is consistent with the original design, but the low-free-energy structures appear because the sequence was not explicitly designed to avoid them. Folding simulations with the AWSEM potential also show

that several non-native structures are energetically competitive, having non-native β -swapping as the main source of degeneracy, as earlier suggested. These and other results (Yadahalli & Gosavi, 2014) highlight the idea that the topology of Top7 is in itself a source for conflicts in folding, leading to a kind of ‘topological frustration’.

4. Folding kinetics and frustration

4.1. Kinetics on idealized funnel landscapes

Very often the dominant folding routes as determined via the ϕ -value analysis can be predicted from the native structure alone by assuming a uniform stabilization gain for forming any given contact (Clementi *et al.* 2000). In the folding transition-state ensembles, many contacts usually are formed simultaneously. Therefore the law of large numbers allows the differences in energy between different specific native contacts to become averaged out. This averaging allows a uniform purely structure-based model to work. Even when such an averaged ‘vanilla’ funnel landscape model is inadequate, other models (slightly more nuanced) having non-uniform interaction energies work well. These heterogeneous models (Portman *et al.* 1998, 2001a, b; Shen *et al.* 2005; Shoemaker *et al.* 1999; Zong *et al.* 2006) while nevertheless having completely minimally frustrated landscapes make improved predictions of the variations of rates. At the very least such heterogeneous contact energy models predict why for a particular system there may be an unusual sensitivity of folding route to details. Such sensitivity occurs, for example, when there are two symmetrically related routes to the folded state which compete. The successes of such pure funnel models in predicting experimental results have been reviewed several times before (Oliveberg & Wolynes, 2005; Onuchic & Wolynes, 2004). We also note that the dominance of native contacts in folding routes assumed in idealized funnel models is supported even by recent analyses of state-of-the-art all-atom folding simulations (Best *et al.* 2013; Henry *et al.* 2013; Piana *et al.* 2012).

Folding in idealized funnels does not always progress as smoothly as might be imagined. Real proteins are not geometrically homogeneous systems, so the energy gain for folding some part of the protein may not always completely compensate the chain entropy loss of fixing the residues in the native positions. Thus the free energy changes along the various routes to the native state may differ giving rise to the appearance of distinct families of macroscopic pathways. For the sake of simplicity, contrast the case in which the native interactions are made mostly between residues close in sequence (such as a α -helical bundle) with another set of routes where the native contacts to be formed are distant in sequence space (such as a β -sheet of distant residues). In the helical bundle, fortuitous interactions will be more probable between residues close in sequence, and the entropy loss will be easily compensated with the interaction energy gain. Since the entropy is already reduced, neighboring interactions will be favorably formed, leading to a zipper-like folding mechanism (Zimm & Bragg, 1959). In the second case of the β -sheet, the entropy loss of forming sequence distant interactions is higher, and thus more interactions will be needed to offset the corresponding entropy loss, making it less probable to form a stable enough structure. This argument is at the core of the correlation between folding rate and contact-order which is a rough proxy for loop-entropy (Clementi & Plotkin, 2004; Plaxco *et al.* 1998; Plotkin *et al.* 1997). Even in simple situations, this imperfect cancellation of chain entropy by contact energy can be sufficient to make a sizable free-energy barrier appear, and for the protein to behave as two-state system, populating mostly folded or unfolded configurations and not much in between. So even when energetic frustration is removed, and the contact energies between residues are

taken to be the same for every pair, the mere topology of the native state gets in the way in defining the locations of the barriers and preferred routes in the folding mechanism. This type of conflict between the need for making native interactions and the necessity of preserving chain connectivity is called ‘topological frustration’.

A simple way to investigate the effects of topological frustration is to explore the folding of a protein with structure-based models (Noel *et al.* 2010). These models use native structure as the sole input of the dynamic calculation and assign a favorable energy to every pair of residues that are ‘in contact’. Contact definitions may differ, but usually a distance cutoff is sufficient to determine the most salient characteristics (Noel *et al.* 2012). Molecular dynamics is then performed to explore the phase-space and both thermodynamic and kinetic parameters can be determined because here there is no energetic frustration to slow down the sampling. A natural reaction coordinate here is ‘how native-like the protein is’, which is closely related to the number of native contacts that can be easily counted in the snapshots of the simulation (Cho *et al.* 2006). Although deceptively simple, these schemes based on perfectly funneled landscapes have been proven extremely useful over the past decade. It has been shown that the structures of transition-state ensembles (Ejtehadi *et al.* 2004), the folding rates (Chavez *et al.* 2004), the existence of folding intermediates (Clementi *et al.* 2000), dimerization mechanisms (Levy *et al.* 2004), and domain-swapping events (Yang *et al.* 2004), are often very well predicted with these kind of models.

The study of topological frustration with structure-based models is not only a playground toy to understand the basics of protein folding, but it has also given insights into diverse functional mechanisms. Studies on the inflammatory cytokine interleukin-1 β (IL-1 β) have yielded significant insights about the modulation of the folding landscape by perturbing functional regions. Structure-based models have identified that a functionally relevant β -bulge (located between β -strands 4 and 5) lies in a region that is topologically frustrated. These simulations suggest that contacts with this region direct folding route selection (Gosavi *et al.* 2008). Deletion of the β -bulge maintains high-affinity receptor binding but abrogates signaling activity, effectively converting the agonist into an antagonist molecule. This insight inspired experimental studies of malleability of the folding process that relate this malleability to functional properties, using proteins where the conflicting region was deleted (Capraro *et al.* 2012). Both the structure-based modeling and the detailed folding kinetics show that topological frustration causes premature folding of some parts that ideally would have folded later. Undoing the traps becomes a kinetic bottleneck, a phenomena termed ‘backtracking’ (Gosavi *et al.* 2006). Occupying these traps can be inevitable and in natural environments eventually becomes a functional feature, acting as a signal facilitating export of the slow-unfolding protein, the formation of a functional excited state, etc. Still, topological frustration is not trivial to detect from first principles. We generally still need to run extended simulations to localize the trapping states and characterize the kinetic bottlenecks, all of which are somehow encoded in the native topology. ‘It is nice to know that the computer understands the problem. But I would like to understand it too.’ E. Wigner.

4.2. Chemical frustration

Most of the forces that give three-dimensional shape to one-dimensionally connected protein molecules are individually quite weak, a few $k_B T$ being a typical interaction energy between amino acid side chains. Some protein chains however are partly assembled using stronger forces such as covalent disulfide bonds or the coordination bonding so prevalent in metalloproteins to stabilize the three-dimensional structure. Because these ‘chemical’ interactions are stronger than

most of the interresidue forces giving rise to the folding funnel, they can lead to an especially large degree of frustration and kinetic trapping in the folding mechanism, especially when folding is studied far away from physiological conditions. While these chemical interactions are very strong they are more specific and relatively fewer in number than the other forces making it possible often to enumerate structural possibilities explicitly.

Disulfide-bonded proteins are especially interesting because of their historical place in studying protein folding. Anfinsen's famous experiment (Anfinsen, 1973) did not directly show that the complete three-dimensional fold was spontaneously formed but rather only showed that the cysteines which can pair up in multiple ways when ribonuclease is oxidatively refolded under redox buffered conditions did indeed finally find their way to their proper native pairings – a rather coarse-grained view of the structure! In the language of this review, then, the possible frustration of mispairing disulfides turned out to play a minimal role at least when the disulfide/sulfide exchange is well catalyzed by the glutathione buffer. A historically important next step was the pioneering study of the oxidative folding of bovine pancreatic trypsin inhibitor BPTI carried out by Creighton & Goldenberg (1984). Here chemical trapping via alkylation revealed several populated species, many of which, surprisingly, exhibited non-native pairings: a sign of frustration. Indeed this laboratory observation was very much on the mind of one of us (PGW) when he first began to study frustration effects on folding in the 1980s (Bryngelson & Wolynes, 1987). Of course the free-energy barrier to rearranging the covalent disulfide pairings depends very much on the oxidation environment so it is not a surprise that chemical frustration can play a role when there are barriers to disulfide interchange. Rather soon after landscape ideas were introduced into folding, ironically it was found by Kim and Weissman that the non-native intermediates originally found in Creighton's study were actually not well populated when the disulfide rearrangements were better catalyzed with a different buffering system so that the 'chemical frustration' was actually minimized, leading to a dominantly funneled mechanism for folding via on-pathway intermediates in which primarily native structure was partially formed and where there was very little long-lived non-native structure (Weissman & Kim, 1991, 1992). The disulfide exchange processes, of course, are influenced by tertiary structure formation and premature burial of a disulfide can lead to an intermediate that must 'back-track', undoing and remaking disulfides to complete folding. This scenario seems to apply to hirudin folding as do some topologically frustrated proteins without disulfides (Chang, 2011).

Metalloproteins provide many other examples of chemical frustration. The paradigm for metalloprotein folding, again of great historical importance to the folding field is cytochrome *c*, which has been the subject of many experimental observations (Englander, 2000) and more recently several theoretical studies (Weinkam & Wolynes, 2010a; Weinkam *et al.* 2005, 2008, 2010b). The intense color of the heme made studies of ultrafast folding particularly attractive. While the iron in the heme of cytochrome *c* is permanently ligated to a cysteine, the sixth ligand for the iron, normally a methionine, can transiently be one of many choices as a coordination partner, each of these gives a characteristic optical spectrum for the heme absorption. Under differing solution conditions then different misligated species can be populated and their concentrations monitored. Flooding the system with imidazole or going to low pH, thereby protonating the histidines, prevents such non-native histidine–heme misligation. Then by minimizing this source of chemical frustration essentially by changing solution conditions, completion of folding turns out to speed up by several orders of magnitude (Weinkam *et al.* 2009). In contrast, chemical frustration can be increased by raising the pH. This allows many histidines and lysines to lose their protons and now be able to ligate the iron, stabilizing alternate configurations, eventually leading




Fig. 16. Symmetry provides an efficient way for minimal frustration. In a symmetric structure, the same set of interactions is used multiple times. If a strong interaction is found, it will appear more often in a symmetric molecule than in the asymmetric case, broadening the energy distribution as sketched in the figure. For folding biomolecules the energetic threshold is set by the ground-state energy of random heteropolymers, E_g/N .

to an alkali denatured species. The species populated by alkaline denaturation turn out to have misligated lysines but otherwise are partially structured in a native-like fashion through the remaining minimally frustrated interactions, which of course include a dominant contribution from hydrophobic interactions of the chain with the heme. Models of perfectly funneled landscapes supplemented by just the specific strong chemically frustrated misligation interactions with lysines have been shown to describe the structural properties of the alkali denatured state quite well (Weinkam *et al.* 2005). These models also do a good job of predicting the thermodynamics in terms of bond strengths of the coordinating ligands and the thermal unfolding characteristics of the native proteins (Weinkam *et al.* 2010b).

4.3. Symmetry effects

Symmetry turns out to provide an efficient way for the minimal frustration in biomolecules to have evolved. The basis for this idea is easy to see qualitatively using the notion of the ‘double-or-nothing’ strategy of high stakes gambling. If you are playing a game where there is a minimum threshold that must be met in order to survive so your goal is not just maximizing your expected winning, in a fair game you are very much better off to increase the size of your individual bets and to make proportionately fewer of them. For folding biomolecules the energetic threshold is set by the floor corresponding to the ground-state energy of random heteropolymers, E_g/N must be achieved for survival.

In a symmetric structure, the same set of interactions is used multiple times – this multiplicity is the meaning of symmetry, after all. So if a strong interaction is found in a symmetric molecule it will appear more often than would otherwise be expected. This corresponds to evolution’s making bigger energetic bets, but again of course, fewer independent bets are made for the same size molecule. This is the double-or-nothing strategy in molecular evolution. Because the distribution of energies is thus wider for the symmetric case, low-energy symmetric structures (see Fig. 16) occur much more often than might be naively expected. Symmetric structures are, indeed, prevalent, and often found by chance (symmetric homo-dimers are often found by crystallographers, even when unexpected). In addition, structural symmetries are believed to have evolved from fused gene duplications even when the actual sequences found in present day examples leave




Fig. 17. Symmetric structures tend to be low in energy. ROP dimer folds as a four-helix bundle of paired coiled coils. An alternate non-native pairing of the helices would form a structure that uses similar contact interactions and would be energetically comparable. The energy landscape appears as a dual-funnel pictured on top. The wild-type (wt) protein folds in *anti* topology and the A2L2 mutant in a *syn* topology, as found by crystallography. In these structures, the backbone of one monomer is colored blue and the other orange. The local frustration patterns are displayed as green lines for the minimally frustrated interactions and red lines for the highly frustrated ones. Both forms are energetically competitive. Structure-based simulations predicted that the *syn* topology folds much faster than the native form, i.e. the *anti* form is more topologically frustrated. The prediction that there was a structural change based on preserving basic energy landscape theory was confirmed by single molecule Förster resonance energy transfer (FRET) experiments. At right, the FRET histograms for wild-type protein and the A2L2 mutant obtained in native buffer (top) and in slightly denaturing conditions (bottom). The peaks at efficiencies of 0.45 and 0.75 correspond to the *anti* and *syn* conformations, respectively. Although the wild-type protein remains stable in one form, the A2L2 mutant occupies a mixed ensemble of states at 0.6 M GuHCl. Because of the near symmetry of a macromolecule, mutations can cause a conformational switch to a nearly degenerate, yet distinct, topology or lead to a mixture of both topologies. (Redrawn with the permission from Gambin *et al.* 2009; Levy *et al.* 2005).

no exact repetition or apparent symmetry in the sequence but only in the three-dimensional structure. A rather indirect argument along these lines was made by Wolynes (1996) and was more elegantly formalized by Wales (1998). A more direct version is due to André *et al.* (2008).

The tendency for symmetric structures to be low in energy means they can show up, seemingly out of the blue, in folding kinetics studies as well as in the natural history of proteins. A cause celebre of this type is the folding of the ROP dimer, a four helix bundle made up of two sets of paired helical coiled coils (Fig. 17). The usual native structure is made by an antiparallel packing of one helical pair on the partner pair of helices. The unusual kinetic anomaly noted by Regan and co-workers was that in some cases making alanine mutations in the individual monomers leads to a dramatic speeding up of the folding rate, while at the same time leaving the equilibrium stability unchanged (Munson *et al.* 1997). This is not consistent with the idea of there being a unique folding funnel and was cited as a problem for the landscape theory. Although a variety of alternative hypotheses were put forward (Lum *et al.* 1999), Levy *et al.* conjectured that symmetry was the problem (Levy *et al.* 2005). An alternate non-native parallel pairing in the dimer structure was

predicted by structure-based simulations to fold much faster than would the usual native anti-parallel form, i.e. the parallel form is less topologically frustrated. Both parallel and antiparallel forms use similar contact interactions and are thus energetically comparable and can act as dual funnels. While originally the fast alanine mutants were thought to have the same structure since like the natural protein they bound RNA *in vitro*, in fact, the mutant was not active *in vivo* where the RNA concentration was lower. This prediction that there was a structural change based on basic energy landscape theory was confirmed by recent single molecule FRET experiments on ROP (Gambin *et al.* 2009). By labeling the monomers it was possible to confirm that both parallel and antiparallel packing coexist for the mutant (Fig. 17). The *in vitro* RNA binding, owing to the high RNA concentration, pulls the structural ensemble over to the natural one by the law of mass action.

4.4. Frustration in repeat proteins

Many natural proteins contain tandem repeats of similar amino acid stretches. These have been broadly classified in groups according to the length of the minimal repeating units. Short repeats up to five residues usually fold into fibrillar structures such as collagen or silk, while repeats longer than about 60 residues usually fold as independent globular domains (Kajava, 2012). There is a class of repeat proteins that lie in between these for which the repeating units couple their folding. For these proteins unique ‘domains’ are not obvious to define (Parra *et al.* 2013). Repeat proteins are believed to be ancient folds. Their biological activity is usually attributed to mediating specific protein–protein interactions. Their modular structure allows a versatility of recognition parallel to that of antibodies. Successful design of repeat-protein domains with novel functions based on simple sequence statistics (Binz *et al.* 2003; Tamaskovic *et al.* 2012) suggests that the folding and the functional signals present in their sequences can be partially segregated. Still, when studied in detail the coupling between folding and binding of natural proteins turns out to be intimately related to their biological function (Ferreiro & Komives, 2010). Typical repeat-proteins are made up of tandem arrays of ~20–40 similar amino acid stretches that fold up into elongated architectures of stacked repeating structural motifs. Quasi-one-dimensional, these non-globular folds are stabilized only by interactions within each repeat and between neighboring adjacent repeats, with no obvious contacts between residues more distant in sequence. The internal symmetries of repeat-proteins suggest that the overall folding properties of a complete ‘domain’ (the stability and cooperativity of the array) may be derived from a microscopic description of the energy balance within each folding element and its interactions with its neighbors (Aksel & Barrick, 2009). Because of the delicate energetic balance in each subunit, subtle variations in the interactions in and between the repeats can give the impression of major changes in the folding landscape while in fact it is well funneled (Ferreiro & Komives, 2007). Such variations may ‘decouple’ the folding of the elements, and partially folded species become populated in kinetic studies. Sufficient information about the population of these states can yield rather quantitative models of the energetic distribution along the protein (Ferreiro *et al.* 2008b; Schafer *et al.* 2012). Behind the apparent simplicity of repeating amino acid sequences, an extremely rich behavior emerges.

Energy landscape theory argues that the three-dimensionally connected globular proteins must fold along a funneled landscape. One-dimensionality weakens this necessity (Luthey-Schulten *et al.* 1995). Folding simulations based on perfectly funneled model landscapes nevertheless have revealed how finer details of the energetic contributions contribute to repeat-protein folding

and are able to recapitulate the experimental results (Barrick *et al.* 2008). Mello & Barrick, (2004) have derived a funneled energy landscape from equilibrium unfolding data on the Notch seven-ankyrin-repeat domain and a simple energy function based on the folding of each repeat and the interaction with its neighbors. Werbeck & Itzhaki, (2007) have shown that the energy landscape of a large ankyrin-repeat protein is funneled yet the kinetics can be described as that for two approximately independent six-repeat subdomains. The different stabilities of these subdomains hint at unevenness in the energy landscape that results in a broad ensemble of species of similar free energies that cannot be distinguished from one another. Single-point mutations cause further landscape unevenness (Werbeck *et al.* 2008) and give rise to multistate folding kinetics (Itzhaki & Lowe, 2012). It is important to note that this ‘unevenness’ is of a different sort than the ‘roughness’ discussed in landscape theory, by which is meant the energetic frustration that leads to non-native interactions.

Hagai *et al.* have recently compared the folding kinetics of two repeat-protein systems that, besides having similar contact-order, display very different folding rates in both experiments and simulations. They used energetically unfrustrated models and showed that modulating the energetic strength of the interface between repeats (relative to that of the intra-repeat) can drastically change the stability, folding rate, and cooperativity of the repeating array (Hagai *et al.* 2012). Their study suggested that a conflict appears between forming interfaces to the left or to the right of an already folded internal repeat-unit. They showed that the formation of one interface is attenuated when the other is already formed. They argue this topological frustration will often arise as a consequence of the quasi-monodimensional geometry of repeat proteins. According to their analysis, these systems can be designed to fold faster by decreasing the strength of the interface or parts of it, resulting in less frustration in the formation of the two interfaces, and consequently in a higher probability for forming the intra repeat contacts. This reduces the conflict involved in forming the two interfaces of a given repeat unit and in turn gives rise to a less cooperative behavior (Hagai *et al.* 2012).

Unlike globular domains, solenoid repeat architectures can be severely truncated and nevertheless still retain cooperative folding properties. Javadi & Main (2009) have described the folding kinetics of a series of designed TPR-repeat proteins assembled with identical repeats, from two to ten repeats, and have described how the energy landscape changes with the addition of repeat units. Consistent with theoretical predictions (Ferreiro *et al.* 2008b), they observed that the longer proteins display more complex multistate relaxations. Moreover, although the initial folding event of all these proteins involves a nucleus with similar solvent accessibility, the subsequent folding of the other regions depends directly on the repeat number, with the longest protein populating an apparent off-pathway intermediate (Javadi *et al.* 2009). Other repeat protein topologies show congruent folding characteristics. Long arrays of Ankyrin and Heat repeats display complex denaturant-dependent relaxations, indicative of sequential transition states and parallel pathways (Itzhaki & Lowe, 2012). Kinetic studies on site-directed mutants suggest that the sequential transition states arise from uneven unzipping of the repeat array with local trapping, and parallel pathways result from the inherent symmetry of repeat-protein structures that initiate folding in either or both N-terminal to C-terminal, depending on the local stability of the initially nucleating repeats. Moreover, the addition of terminal stabilizing repeats can shift the transition-state ensemble toward those regions, rerouting the folding of the whole repeat array (Tripp & Barrick, 2008). Local energetics are thus crucial in determining the routes actually followed and the traps populated in these quasi one-dimensional proteins. Experimentally tracing the kinetics of repeat-protein folding further details the richness of mechanisms that can be gained by molding




Fig. 18. The energy landscape of repeat-proteins has internal symmetries. The repeating sequence encodes similar structural elements that interact with nearest neighbors, depicted by many folding elements that merge upon interaction and comprise an overall funneled landscape (Ferreiro & Komives, 2007). High-energy unfolded configurations are on top and the fully folded state at the bottom. In between, intermediate ensembles of partially folded states can become populated at equilibrium (red shadow). Asymmetries in the local energy distribution mold the landscape roughness that defines the kinetic routes. The thick blue lines depict the preferred folding routes of the 12-ankyrin repeat protein D34, for which the intermediate ensemble is polarized toward the C-terminal repeats of the molecule (Werbeck *et al.* 2008). The crystal structure of D34 is pictured at the bottom, colored according to its local frustration pattern. Highly frustrated regions are located at the ends, and in a central region between repeats 5 and 6. These breaks occur where the experiments suggest the intermediate's folding boundary is. The collective influence of local interactions along the one-dimensional scaffold allows sites to thermodynamically modulate each other even at considerable distance. Small perturbations, such as mutation or ligand binding, may change the interactions in or between the modules causing large effects that can be experimentally dissected. Redrawn with the permission from Ferreiro *et al.* (2008b).

the traps. As mentioned, these traps may arise purely topologically or as the effects of local energetic frustration.

Energetic frustration seems to be at work in determining the traps in the folding of the ankyrin repeat protein D34. This protein unfolds via a fast process that yields an intermediate that can also be detected at equilibrium (Werbeck & Itzhaki, 2007). The barrier separating the native from the intermediate states appears 'broad,' consistent with the picture of neighboring repeats unfolding sequentially (Fig. 18). A slower unfolding phase corresponds to unravelling the C-terminal subdomain, where mutations affect both the folding rate and its urea dependence in a manner consistent with unfolding by two parallel routes. How does this change in folding regime appear? Local frustration analysis points out that there are three regions enriched in frustrated interactions in an otherwise strongly cross-linked web of minimally frustrated contacts (Fig. 18). One of the highly frustrated regions collocates with the central region where the experiments suggest the intermediate's folding boundary is found (Ferreiro & Wolynes, 2008). The highly frustrated interactions are apparently responsible for the roughness that gives rise to a long-lived folding intermediate. Modification of this interface so that it is minimally frustrated should destabilize the trap, making the D34 landscape appear much smoother.

Experiments redesigning repeat-proteins so as to conform to minimally frustrated ‘consensus’ sequences show that indeed traps are destabilized and kinetic barriers are considerably reduced (Kramer *et al.* 2010; Street & Barrick, 2009). The overall picture emerging from repeat-protein-folding studies is that the differences in folding rates are related to escape from local traps. These traps appear as manifestations of frustration, either topological, energetic or both. There is a fine balance between these signals that can be at play at the secondary structure and/or the tertiary contact levels. A local effect is felt globally because the near-neighbor interactions are extraordinarily important in stabilizing the repeats, and weak biases can tip the balance to complete folding. Single substitutions can affect local biases (such as helix propensity), and thus exert profound effects on the overall folding of these domains. The larger repeating arrays are more likely to tolerate ‘cracks’ and the folding at the repeats ends may become anticorrelated (Ferreiro *et al.* 2008a). Such sensitivity allows for specific encoding of traps by means of making few sequence or environmental modifications. This may be used in intracellular targeting and sorting. Owing to the symmetry of the repeating array, long arrays can be fine tuned to populate partially folded states, and these ensembles can be coopted in functional mechanisms. Changing the stability of a single repeating element (by post-translational modifications or binding of other macromolecules) may affect the behavior at a distant site, providing a coupling mechanism that can transmit allosteric signals to long distances within a single repeating geometry.

4.5. Anomalous ϕ values

In the laboratory, ϕ -values are determined by measuring the changes in both overall stability and folding speed of a protein and single point mutant (s) of it and comparing them (Fersht *et al.* 1992). The stability change is $\Delta\Delta G_{\text{mut-wt}}^{\text{eq}} = \Delta G_{\text{mut}}^{U-F_{\text{mut}}} - \Delta G_{\text{wt}}^{U-F_{\text{wt}}}$; and the change in the folding rate usually can be ascribed to the effect of the mutation on the activation free energy, $\Delta\Delta G_{\text{mut-wt}}^{\text{TSE}} = -RT \ln(k_{\text{mut}}/k_{\text{wt}})$. This analysis assumes that the transmission coefficient, related to configurational diffusion and internal friction, remains constant upon mutation, which may not always be the case. The ϕ -value is defined as $\Delta\Delta G^{\text{TSE}}/\Delta\Delta G^{\text{eq}}$. For minimally frustrated proteins, realistic pictures of the rate-limiting steps in folding can be constructed by measuring the effect of mutations not just at the folding midpoint but over a range of thermodynamic conditions. In contrast to the situation for a minimally frustrated protein, on a very rugged landscape the ϕ -values become very difficult if not impossible to interpret in terms of the native folded protein structure alone. This is because the kinetics of re-organization is determined by the stability of unknown non-native contacts in the transition-state ensemble. Since these need not be present in the final native structure for highly frustrated systems, their energetic change would be unrelated to stability changes of the protein as a whole. The globally unfrustrated character of the energy landscape of natural proteins indicates that there will be relatively few thermally accessible states that can be associated with such folding ‘mistakes’ at the folding temperature (T_f). When T_f/T_g is large, the structures will have mostly native interactions formed and only a few weak non-native interactions are anticipated. Partially folded forms of the protein may become populated, but highly misfolded alternate structures are not expected to occur in general. This simplifying feature of a minimally frustrated heteropolymers’ landscape is what allows one to infer the existence of a relation between the energetic contributions a residue pair makes in the native state with the energetic contribution to the populated intermediate or transition-state ensembles in the folding routes. The thermodynamic contribution of an interaction will be simply scaled by the fraction of the time that the pair contact is actually formed in the ensemble, and the folding rates will




Fig. 19. Anomalous ϕ -values often involve frustrated sites. In a minimally frustrated landscape, there is a relation between the energetic contributions of residues in the native state with the energetic contribution to the transition-state ensembles along the primary folding routes. These can be experimentally measured by analyzing the change in stability and folding speed upon single-site mutation by the ϕ -value analysis. Consistent with the notion of widespread minimal frustration, more than 1000 reliable measurements (Naganathan & Muñoz, 2010) lie in the range $0 < \phi < 1$ as expected on a perfectly funneled landscape. The figure shows the frustratograms for proteins in which deviations have been confidently measured, with the anomalous residues highlighted in orange. CI2 mutant A16 G ($\phi = 1.12$); FKBP12 mutant F36A ($\phi = -0.1$); fynSH3 mutants F26I ($\phi = -0.4$) A39 V ($\phi = 2.34$) A39F ($\phi = -0.16$); srcSH3 mutant I34A ($\phi = 2.01$).

largely follow native-state stability giving rise to ‘rate equilibrium free-energy relations’ (Leffler, 1953). This regularity is predicted by a funneled landscape and provides the ultimate theoretical basis of the analysis of protein-folding kinetics using ϕ -values (Matouschek *et al.* 1989). Strong deviations from this simple behavior then are likely manifestations of localized frustration in the energy landscape.

Many natural proteins with simple kinetic behavior have been subjected to the ϕ -value determination and analysis. Almost 1000 mutations have been independently evaluated and the general behavior is in accordance with the expectations outlined above (Naganathan & Muñoz, 2010). Most of the ϕ -values indeed lie in the range $0 < \phi < 1$, i.e. they are consistent with an ensemble averaged fractional occupancy of local native structure (reviewed elsewhere (Onuchic *et al.* 1997)). In general, the ϕ -value increases linearly with the stability perturbation and, in conditions of isostability, where the unfolded and folded forms are equally populated, most of the ϕ -values cluster about $\phi = 0.36 \pm 0.11$ (Naganathan *et al.* 2010). This average value is anticipated from both mean field and capillarity theories on perfectly funneled landscapes if the activation free-energy barrier originates from an imperfect cancellation of the entropy loss by the energy gain upon folding, as we have discussed previously. ϕ -values are not completely uniform, however. Site-specific variations of ϕ -values can be roughly predicted by knowing the topology of native structure under the assumption of a minimally frustrated energy landscape (Clementi *et al.* 2000; Muñoz & Eaton, 1999). These predictions are pretty good when the conformations belonging to the transition state ensemble (*TSE*) have a large connected web of native contacts established. Sometimes, especially for α -helical proteins, inhomogeneities of the energies do make the vanilla predictions less accurate (Cho *et al.* 2009). If the interactions around the probed site are sparse, mutations can change the location of the TSE dramatically and the perturbations can no longer be treated as linear free-energy perturbations (Sánchez & Kiefhaber, 2003b). Very large deviations from structure-based predictions or ϕ -values outside the expected 0–1 range are signatures that energetic frustration could be at play. Sometimes, the deviations can be directly attributed to highly frustrated interactions that one can identify in a native structure. Figure 19 shows the local frustration pattern of four example proteins where specific mutations perturb the folding rate more

than would be expected by the native stability change ($\phi > 1$, Fig. 19a) or where the mutations have opposite effects on the TSE and the native state ($\phi < 0$ Fig. 19b). In both CI2 and FKBP12, the sites indicated by these anomalous ϕ -values map to regions that are enriched in highly frustrated interactions in the native state (Fig. 19). However, in the cases of the fynSH3 and srcSH3 domains, the mutations for which anomalous ϕ -values were measured do not coincide with frustrated interactions for tertiary contacts (Fig. 19c, d). For these sites the anomaly must be attributed to other sources of frustration. One major component of the stabilization energy that is not visible in the frustratograms of Fig. 19 is the secondary structure energy term. Strongly stabilizing tertiary interactions may locally conflict with the secondary structure propensity of a backbone that will otherwise occupy other conformations if unrestrained. Conflict between secondary structure formation and tertiary interactions has been commonly observed in the SH3-fold class (Klimov & Thirumalai, 2002). Note that in these cases the regions of anomalous ϕ -values are enriched in minimally frustrated tertiary interactions, perhaps compensating for the secondary/tertiary conflict (Fig. 19c, d).

4.6. Non-native intermediates

Immunity Protein 7 (Im7) is an 86-residue-long protein that *in vitro* folds through an on-pathway intermediate state. Kinetic studies suggested that this intermediate contains some non-native interactions of three of the four helices around a specific hydrophobic core (Capaldi *et al.* 2002). This exception to the behavior expected for minimally frustrated heteropolymers hints that a functional constraint may be conflicting with rapid folding. Does frustration change the way the Im7 intermediate is formed? Sutto *et al.* explored this question using simulations that include native bias between residues close in sequence but that allowed for non-native (and possibly frustrated) interactions between residues farther apart (Sutto *et al.* 2007). These simulations predict an intermediate state stabilized by favorable non-native interactions formed by residues that, in the native form, cluster on the surface to form a locally frustrated region (Fig. 20). Further mapping of the structural transitions using protein engineering combined with kinetic folding experiments and restrained molecular dynamics characterized more completely the details of the basins and bottlenecks, and confirm that indeed the energy landscape of Im7 is unusually rugged (Friel *et al.* 2009). Rather than folding by progressively increasing the number of native contacts, the sequence of Im7 contains other signals that constrain efficient folding. Many of the residues identified as forming non-native contacts during the early stages of folding of Im7 lie in regions that have a functional role: they are involved in the recognition and inactivation of colicin toxins forming the interaction interface with them. By coupling protease susceptibility and antibiotic resistance traits, Foit *et al.* developed a novel assay to screen for protein stability *in vivo* (Foit *et al.* 2009). When random mutagenesis and selection for antibiotic resistance was conducted on the Im7 protein, most of the mutations that were found to enhance thermodynamic stability map again to the predicted binding surface as would be expected from the local frustration analysis (Fig. 20). It appears that the functional constraint for specific binding restricts the sequences and gives rise to energetic conflicts with the rest of the polymer. Smooth folding of this small domain is compromised by frustration and in the case of Im7, traps with non-native structure give rise to an intermediate, which is observed both in experiment and simulations (Figueiredo *et al.* 2013; Whittaker *et al.* 2011). The observation of clusters of frustrated interactions in the native state points the way to a general mutational strategy to reduce the ruggedness of folding landscapes. The most effective approach to destabilize the intermediate of Im7 should be to




Fig. 20. Non-native intermediates highlight frustrated sites. Functional constraints for the specific binding constrain the polypeptide sequences and can give rise to energetic conflicts with the rest of the polymer. Im7 is an example in which smooth folding is compromised by function, leading to local frustration. In the top panel is shown the average structure of the native protein where highly frustrated interactions are marked with red lines. Traps with non-native structure give rise to a folding intermediate, which is observed both in experiment and simulations. The intermediate is stabilized by favorable non-native interactions formed by residues that, in the native form, cluster on the surface and form a locally frustrated region. Minimally frustrated contacts present in the native structure (green) are distinguished from those that are non-native (blue). A distinct non-native cluster can be observed involving interactions between helix IV and the helix III region. Native (red) and non-native (orange) frustrated contacts surround the core. Random mutagenesis identified sites where substitutions increase the thermodynamic stability (purple). Many of the residues identified as forming non-native contacts during the early stages of folding of Im7 lie in regions that have a functional role: they are involved in the recognition and inactivation of colicin toxins forming the interaction interface with them. The observation of clusters of frustrated interactions in the native states points the way to a general mutational strategy to reduce the ruggedness of real protein-folding landscapes.

stabilize the native state by alleviating the energetic conflicts in this region (Sutto *et al.* 2007). When redesigned mutants were simulated, the intermediate is no longer populated when residual native-state frustration is reduced. The redesign strategy based on minimizing native-state frustration is more effective than the specific negative design of the intermediate. A mutant specifically redesigned to destabilize the wild-type intermediate did not fold without an intermediate but rather retained an intermediate having a different structural ensemble, but nevertheless involving some non-native contacts (Sutto *et al.* 2007).

4.7. Friction and roughness effects on rates

The Im7 example shows that one of the effects of frustration is to introduce kinetic traps or non-native intermediates into the folding mechanism, a possibility first pointed out by Bryngelson & Wolynes (1987). If the system is small the number of trapping states will be




Fig. 21. Frustration can affect proteins' internal friction. Search in a rough energy landscape is affected by the population of local traps. When there are many traps these can be thought of as a source of friction on the diffusive folding. The structurally similar Spectrin repeats were shown to fold with very different rates (Wensley *et al.* 2010). The frustratograms of a natural protein that include these repeats is shown. R16 and R17 have a central cluster of highly frustrated residues and fold three orders of magnitude slower than does R15, which is more uniformly minimally frustrated. This central cluster of frustrated residues is involved in the main crossing from an initial transition state for folding to a later exit transition state for folding (Borgia *et al.* 2012). The frictional effect can be revealed by the viscosity dependence of the folding reactions, which can be studied by varying the solvent conditions and keeping the relative stabilities unchanged. For R15 the folding rate scales inversely to the solvent viscosity while for the frustrated systems R16 and R17 the folding rates are independent of solvent viscosity (Wensley *et al.* 2010).

small and can be described as arising from the specific intermediates. The structural uniqueness of these intermediates is however in each case an open issue. As we saw in the Im7 example destabilizing one particular misfolded structure, generally just gives rise to an alternate trap, if frustration still remains in the native structure. What is the effect of these ensemble traps on the folding kinetics? In general, searching through the traps slows folding down. Bryngelson & Wolynes (1987) showed how this slowing effect, especially where there are many traps, could be thought of as a source of 'friction' on the diffusive progress of the protein as it is organized by the remaining guiding forces of the funnel. A very nice experimental system that seems to illustrate this idea is the extensive study of the folding of spectrin domains by Clarke and co-workers (Wensley *et al.* 2010). The experiments of Clarke and co-workers have shown that structurally highly homologous members of the spectrin domain family fold with very different rates. The frustratograms of the three natural systems that have been studied are shown in Fig. 21. R16 and R17 that have a central cluster of frustrated residues fold three orders of magnitude slower than does R15, which is more uniformly minimally frustrated. This central cluster of frustrated residues has been shown to be involved in the main step crossing from an initial transition state for folding to a later exit transition state for folding (Borgia *et al.* 2012; Wensley *et al.* 2012). As in other quasi-one-dimensional proteins such as the repeat-proteins, this extended transition region entails what may be described either as a broad malleable transition-state region or a high-energy intermediate (Sánchez & Kiefhaber, 2003a).

One of the effects of having localized frustration in the part of the protein that is being studied at this key kinetic bottleneck is thermodynamic – the transition-state region is destabilized by the frustrated contacts and, by itself, this effect diminishes the folding rate. This thermodynamic effect might be said, with more precision, to be an effect of energetic heterogeneity rather than frustration *per se* since no alternate structure is involved. Yet a residual search or friction effect contributes to the slowing as well. This frictional effect of frustration is revealed by the viscosity dependence of the folding reactions, which can be studied by varying the solvent, while keeping the stability unchanged. For the unfrustrated R15, the rate scales inversely to the solvent viscosity while for the frustrated systems R16 and R17 the folding rates are independent of solvent viscosity (Wensley *et al.* 2010).

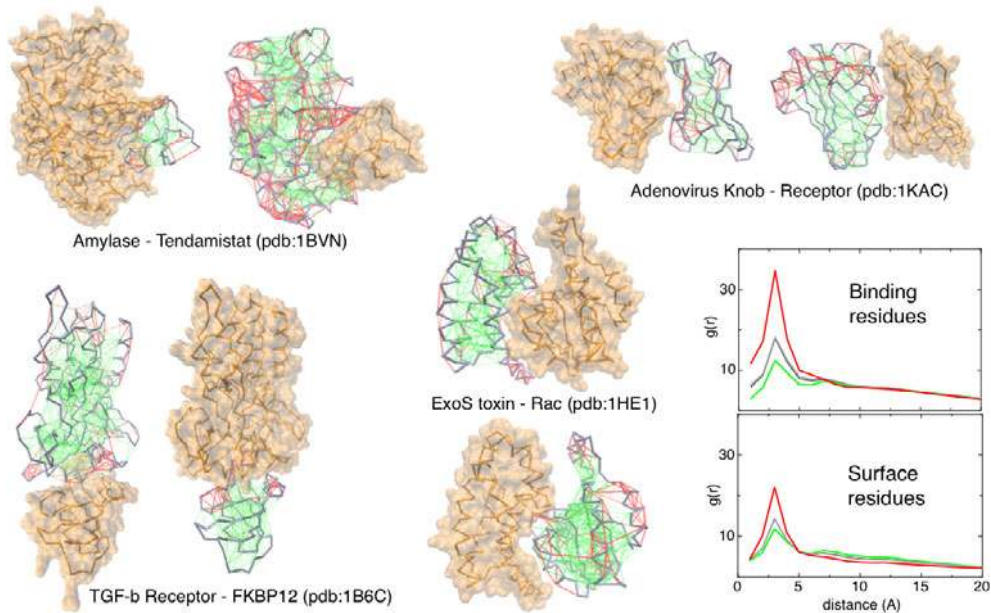


Fig. 22. Frustration guides protein–protein association. The structural coordinates of natural protein complexes were taken from the pdb, and local frustration was calculated on the unbound monomers. The configurational frustration pattern of one ligand is shown together with the surface of the other. Minimally frustrated contacts are green, highly frustrated contacts in red, water-mediated interactions with dotted lines. A statistical survey was performed with a non-redundant sample of 85 complexes (Ferreiro *et al.* 2007). The pair distribution functions between the centers of mass of the contacts and the Ca of either binding or non-binding surface residues is shown. Whereas many of the highly frustrated interactions are close to the binding residues as defined in the co-crystal structures, there are additional patches of frustrated interactions on the protein surfaces, suggesting that there may be other functional reasons to retain these conflicting residues.

Rather general arguments suggest that solvent viscosity effects depend on how much chain motion occurs as the transition state is traversed (Frauenfelder & Wolynes, 1985; Portman *et al.* 2001b). On the smooth R15 landscape, large amounts of chain motion occur but when frustrated traps are present, as in R16 and R17, many individual smaller steps must be made searching through the misfolds. These smaller reconfiguration events displace little solvent. The idea that several misfolded configurations are involved in the frustrated system is supported by atomistic simulations of the R16 system that show explicitly helix sliding events coming from mismatches in the registry of the helices in this frustrated region (Best, 2012).

5. Functional consequences of frustration

5.1. Binding sites are often frustrated

Landscape theory approaches have already proven useful for describing protein–protein interactions (Zheng *et al.* 2012). From the point of view of the localized frustration index, a natural question is whether and how often the residues that can conflict with folding of a single domain are there specifically to allow protein–protein association. Ferreiro *et al.* analyzed the distribution of the frustration index in protein assemblies on a curated set of 85 non-redundant heterodimeric protein complexes (Ferreiro *et al.* 2007). Figure 22 shows the frustration indices of the

contacts and residues in some examples from the database. When the calculations of local frustration are performed in the unbound monomers, indeed one sees that near the binding sites the proteins are enriched in highly frustrated interactions (Fig. 22). The statistics over the whole database showed that the highly frustrated contacts both cluster and tend to be closer to the binding residues than they are to other surface residues. Minimally frustrated interactions are also particularly excluded from the vicinity of binding residues and neutral interactions follow the native topological distribution (Ferreiro *et al.* 2007). Interestingly, when the non-binding surface is examined, some clusters of highly frustrated interactions can still be detected (Fig. 22). Whereas many of the highly frustrated interactions are close to the binding residues as defined in the co-crystal structures, these additional frustrated patches on the protein surface suggest there may also be other functional reasons to retain conflicting residues. These reasons may be binding to other unknown molecular partners or, perhaps, these frustrated regions are sites relevant for allostery (*vide infra*). A corollary to this observation is that if the binding surfaces are too big with respect to the area of the domain, they are expected to enrich the total amount of frustrating interactions, compromising the folding of the unbound domain. It is interesting to note that coupling between folding and binding is more prevalent when the interaction surface is large (Papoian & Wolynes, 2003b).

Changes in frustration upon binding may help the guide-specific association. A survey of the change in local frustration upon binding was performed by calculating the frustration indices for proteins in the complexes and comparing these to the frustration indices of the unbound monomers making up those complexes (Ferreiro *et al.* 2007). When calculated over a non-redundant set of hetero-oligomer complexes, the newly formed interactions at the binding interface have on average a local frustration distribution similar to that for contacts in single domains. This implies that the specific association can as well be accounted by the principle of minimal frustration, as suggested earlier (Papoian & Wolynes, 2003b). Most of the contacts strictly within the monomers do not change their frustration index upon binding, but approximately 25% of them do become less frustrated upon association. A minor fraction about 7%, change in the opposite direction, becoming more frustrated. The decrement in frustration index is nearly constant and mostly comes from the change in the solvent accessibility upon binding, captured by the burial term and the contact type in the AMW energy function. In summary, the frustration index of the regions close to the binding site changes upon association, usually becoming less frustrated in the complex, and the newly formed interfaces have a similar distribution of local frustration to the interiors of unbound domains.

5.2. Frustration in allostery and local conformational changes

If the minimal frustration principle were to be satisfied everywhere, and the temperature sufficiently low, the protein molecule becomes a beautiful sculpture that would tend to move as a rigid body. For real proteins this *rigor mortis* cannot be the whole story, as the basic functions of proteins usually involve responding to changes in their environment. The detection and logical control capabilities of proteins rely on their ability to amplify the effect of a relatively small environmental change so as to produce a large discrete transition in a different chemical realm. Monod emphasized how the apparent gratuity of such 'transductions' provide the essential computational capabilities of protein molecules (Monod, 1973). We can thus expect that the minimal frustration principle will be at least locally violated by natural functional proteins. Local violations

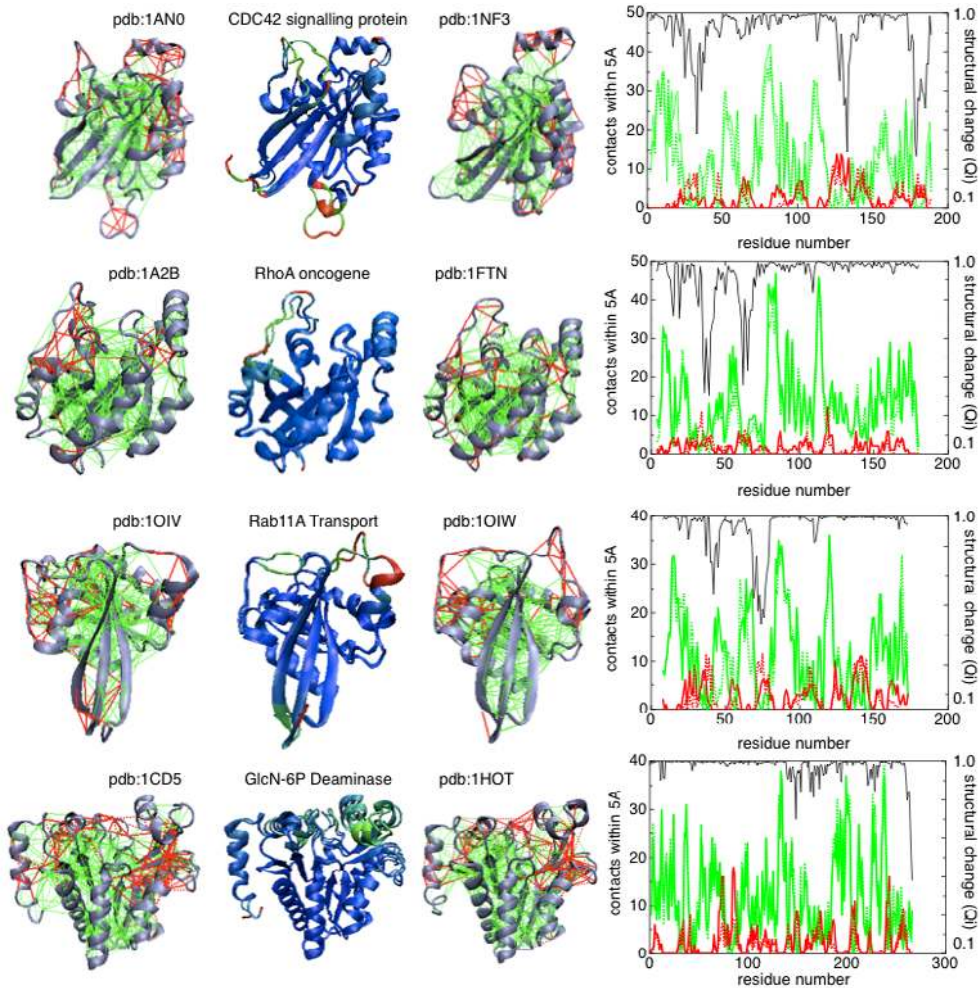


Fig. 23. Local frustration in allosteric domains. Proteins can change their structural forms responding to signals in their environment; some have been crystallized in two different conformations. Examples of structural alignment of experimentally determined conformations are shown above, colored according to the structural deviation (blue low, red high). The frustratograms for the individual conformations are shown at the sides, with the minimally frustrated interactions in green lines, the highly frustrated interactions in red lines. At far right a quantification of the local frustration projected on the linear sequence of the protein, minimally frustrated interactions (green) or highly frustrated interactions (red) in the vicinity of each residue, in either forms (solid or dashed). The structural deviation between both states is quantified by a local Q_i score. Clusters of high local frustration collocate with residues whose local environment shifts between the two structures. Redrawn with the permission from Ferreiro *et al.* (2011).

open possibilities for more complex energy landscapes, such as those needed for allostery and large-scale conformational changes.

A statistical survey of the local frustration patterns of allosteric protein domains indeed does show that the regions that reconfigure are often enriched in patches of highly frustrated interactions (Ferreiro *et al.* 2011). In contrast, the more rigid parts of the proteins, which are locally structurally superimposable in both forms, are connected by a dense web of minimally frustrated interactions (Fig. 23). Regions that are highly frustrated typically are where the action is and are

seen to reconfigure locally between the two different forms. In some cases, frustrated regions display rather extensive reconfigurations of compact regions, whereas in others, the frustrated clusters localize around apparent pivot points joining the more rigid elements. In general, clusters of high local frustration colocalize with residues whose local environment shifts between the two structures (see Fig. 23). It is apparent that the major structural ‘core’ is conserved between the large-scale taxonomic conformational substates, and differences between them tend to localize on the protein surface. The changing frustrated regions often include the binding site of an allosteric effector. While conformational changes enter, keeping about 80% of the pairwise interactions, 10% of the interactions are exclusively found in one conformational substate or the other. The distribution of local frustration for contacts that are common to both structures is similar to the distribution for contacts exclusively found in only one substate. The interactions that are found in both substates do not dramatically change, but instead new contacts are formed in one substate that does not exist in the other. Thus the conformational change is brought about by a balanced exchange of local frustration upon switching substates. In surveying several allosteric domains, there are examples in which only one structure of a pair has a relevant frustrated cluster but where the frustration is relieved in the other conformation (Ferreiro *et al.* 2011). In other cases, both structures display highly frustrated interactions in the same region, and these residues interchange the interactions among each other. There are no examples of structure pairs of monomeric proteins where the locally displaced regions are completely minimally frustrated in both configurations.

The term allostery was in fact first used only to describe the regulation of protein activity through changes in quaternary structure of multisubunit complexes (Changeux, 2013). Landscape theory underscores how preserving symmetry is an easy path for achieving low-energy near-degenerate conformations, an idea taken as axiomatic by Monod, Wyman, and Changeux. Using the same counting argument that makes symmetric structures individually more easily designable, pairs of symmetric structures provide a statistically favored way of achieving a switch-like behavior. The local frustration distribution in the multimeric forms of the allosteric proteins show that, on average, interfacial (quaternary) interactions are less frustrated than the monomeric (tertiary) interactions (Ferreiro *et al.* 2011). The interfaces in multimers appear in fact thus to be depleted of highly frustrated interactions, suggesting that the symmetry is reflected in the local frustration patterns when alternate configurations are used as decoys. The primal example of allostery, hemoglobin is shown in Fig. 24. Here the interface between the subunits is not frustrated in either form; rather a rigid-body rotation allows symmetrically related interactions to occur. These distant structural forms involve a different knob-in-hole packing of the helices of one monomer on the others, and both packings appear to be minimally frustrated (Fig. 24). Multiple funnels to structurally distinct low-free-energy states can be achieved with nearly rigid subunits packing in a number of symmetrically equivalent ways. In this sense, the natural ‘engineering’ of frustration may be not essential for allostery. The classic Wyman–Monod view of symmetry of multimeric assemblies leading to near degeneracy does not depend on frustration at all. Because this usually involves only a few $k_B T$ of free energy, weakly assembled complexes may accidentally exhibit allostery without any strong frustration, as recently argued by Kuriyan & Eisenberg (2007).

The precise mechanism by which localized frustration allows allostery clearly deserves further investigation. Local frustration may simply allow a small discrete set of configurations interconverting by local motions. The allosteric dynamical transition may then involve specifically defined paths of transformation much like a small molecule or a macroscopic machine




Fig. 24. Frustration in prototypic allosteric proteins. Hemoglobin is a tetrameric protein in which near rigid-body rotation allows symmetrical packing of subunits to occur. In this multimeric protein, the interface between the subunits is not frustrated in either form. Highly frustrated interactions do occur internally in each subunit near the heme-binding site. This is a classical example of a Wyman–Monod view of symmetry leading to near degeneracy, for which allostery does not depend on frustration. Below, the case of adenylate kinase, a protein that undergoes a large-scale conformational transition upon binding substrates. Steered molecular dynamics identified locations where ‘hinge’ motions are believed to occur, shown here with blue arrows (Henzler-Wildman *et al.* 2007). An extensive minimally frustrated network of contacts rigidifies the molecule in the closed form, and highly frustrated regions co-locate with the hinges. Motion of adenylate kinase along the low-frequency normal modes contributing to the closure accumulates stress in some regions (Miyashita *et al.* 2003). A high-stress region can ‘crack’ or locally unfold releasing the strain and catalyzing the motion. This region is highly frustrated in both forms. The presence of interactions that conflict with folding an enzyme is a general theme in the realization of effective catalysts. Redrawn with the permission from Ferreiro *et al.* (2011).

with hinges. Alternatively, the local frustration may destabilize a part of the protein in favor of an ensemble of rather high entropy. In other words, a local region can locally unfold or ‘crack’. This possibility has been entertained previously (Miyashita *et al.* 2003) and elegantly explains the observations of the way denaturants can catalyze conformational changes (Zhang *et al.* 1997). Resolving experimentally the issue of hinges *versus* cracks requires an analysis of the effect that local mutations exert on the conformational kinetics (Whitford *et al.* 2008).

Adenylate kinase is celebrated example of large-scale conformational change related to a functional transition. The gross opening and closing of this protein requires at least two reaction coordinates to be taken into account. Steered molecular dynamics calculations have identified

several locations where ‘hinge’ motions are believed to occur (Henzler-Wildman *et al.* 2007). The frustratograms of both substates are shown in Fig. 24. Nearly rigid-body motions of the lid and the core are involved in the closure. An extensive minimally frustrated network of contacts rigidifies the molecule in the closed form, and the places identified as hinges indeed co-locate with highly frustrated regions. Miyashita *et al.* have examined the motion of adenylate kinase along the low-frequency normal modes contributing to closure and concluded that there were regions where high stress accumulates at the transition state for conformational change (Miyashita *et al.* 2003). This high-stress region can however crack or locally unfold releasing the strain and catalyzing the motion. This high-strain region is highly frustrated in both forms. The presence of these interactions that conflict with folding the enzyme may be a general theme in the design of effective catalysts.

Protein phosphorylation is a paramount example of allosteric control of enzymatic activity. The activation of protein kinases usually occurs in a sharp switch-like fashion upon phosphorylation, making the ‘kinome’ a rich source of computational cellular complexity where recurring networks can be constructed. These days particular attention has come to some members of the multitude of kinases, because mutations of kinases have been found to be associated with pathological cell growth. In some cancers the activating oncogene involves substitutions on ABL and EGFR kinases quite often. Oncogenic mutations perturb the minimally frustrated core of these proteins, while the locally frustrated clusters appear related to highly dynamic regions that are involved in allosteric activation (Dixit *et al.* 2011). In particular, the cancer-related mutations were found to both amplify dynamics of the inactive state and relieve local frustration in the active state (Dixit *et al.* 2011). Thus, the redistribution of local frustration may facilitate conformational transitions and at the same time change the relative stability of the substates.

5.3. Frustration in metastable and multistable proteins

The universe of conformational substates of any protein molecule is huge. Therefore the complete energy landscape of proteins is complex, yet many experiments at physiological conditions can be interpreted using a simpler free-energy spectrum of states, which can be considered as ensembles of low-temperature substates (Hegler *et al.* 2008). The low-free-energy part of the spectrum of excitations corresponds with partially folded ensembles that get naturally populated as a consequence of the hierarchical nature of the mostly funneled landscape. Often the properties of these excited states can be predicted by native topology-based models (see above). As we have seen, frustration, the conflict between inconsistent stabilizing interactions evolved for other purposes than optimizing folding, offers another mechanism for forming low-free-energy excitations. If these states have large structural differences, then the interactions that stabilize one form are expected to destabilize the other, making natural proteins metastable in certain conditions. Given that minimal frustration corresponds to folding stability, does high local frustration correspond to metastability?

Hemagglutinin (HA) is a fascinating example of metastability. The process of viral infection starts with the interaction of multiple HA molecules on the infecting virion interacting with multiple sialic acids on the cell to be infected (Luo, 2012). This interaction triggers endocytosis, and the virion is transported to the endosome where the pH is 1.5–2 units lower than outside the cell. At this lower pH, the HA molecule undergoes an irreversible and dramatic conformational change that was first mapped-out by Wiley and co-workers in the early 1990s (Bullough *et al.* 1994). The molecule rearranges so that parts of it move by as much as 100 Å from one structure




Fig. 25. Local frustration in metastable proteins. Hemagglutinin (HA) is a viral protein that undergoes a dramatic conformational change in response to an environmental pH change. The figure shows the local frustration patterns of the HA molecule crystallized in different conditions. The protein is a trimer that protrudes from the membrane of coated viruses, rearranging so that parts of it move by as much as 100 Å, labeled here A, B, C and colored on a monomer to ease visualization. Several regions of high frustration are identified at pH 7 (red lines). The local frustration of these region changes upon rearrangement and is diminished overall in the structure at pH 5. The regions involved in the major conformational change are already destabilized in the pre-fusion state and are more likely to ‘crack’. There is also a region of high local frustration of the fusion peptide (FP, orange) in the pH 7 structure that conflicts with the core of the HA trimer. This peptide interacts with the membrane of the host endosome upon triggering by pH.

to the other. The local frustration patterns of the HA molecule crystallized at pH 7 show several regions of high frustration (Fig. 25). The frustration of these region changes upon rearrangement and is diminished in the structure crystallized at pH 5. Thus, the regions that are involved in the conformational change are already destabilized by local interactions, and are more likely to ‘crack’, as in the cases of allostery describe above. Another interesting feature of these frustratograms is the high local frustration of the ‘fusion peptide’ in the pH 7 structure (Fig. 25). This peptide is poised to interact with the membrane of the endosome upon triggering, and promote the fusion of the virion membrane with that of the infected cell’s endosome. It is evident that when the peptide is occluded in the pre-transition structure, it conflicts with the core of the HA trimer. Presumably this structure can be held in place by the minimally frustrated network of interactions that the other residues of the core contribute.

Serpins are another large group of metastable proteins that undergo a dramatic conformational change when cleaved by target proteases (Huntington *et al.* 2000). The protease is itself strongly inhibited by the cleaved serpin. The cleavage occurs on an exposed amino acid sequence in the serpin molecule (called the reactive-center loop, RCL), and the resulting cleaved loop becomes an additional β -strand in the central β -sheet (colored orange in Fig. 26). The uncleaved serpin is




Fig. 26. Frustration in Serpin strand swapping. Metastable proteins can undergo large conformational changes upon backbone connectivity changes. When a protease cleaves the α 1-antitrypsin molecule in the RCL loop (top left, purple and orange) the whole molecule rearranges and this part of the polypeptide is found as a β -strand inserted in the central β -sheet (top right, orange). The local frustration patterns show that the RCL is ‘strained’ in the pre-cleavage form, where local destabilization promotes conformational change. Accommodating the new β -strand requires sub-optimal interactions among residues in the β -sheet region of the protein, which is also locally frustrated. Some serpins can spontaneously undergo a conformational rearrangement without RCL cleavage. The structure of AT-III strand-swapped dimer is shown at the bottom, with one monomer colored cyan and the other gray. In a classic domain-swap fashion, the polypeptides exchange part of the β -sheet with one another. The strand that gets swapped corresponds to the RCL region that can get cleaved. Regions of high local frustration in the dimer correspond with the regions that undergo rearrangements, the base of the swapped β -strand in addition to the α -helix that is dislodge from the face of the β -sheet for strand insertion.

thought to be more stable than the cleaved form (Im *et al.* 2000). Part of the energy difference comes from the cleavage of a peptide bond in the RCL, but other energies are also implicated. The ‘strain’ of the metastable form of the serpin is evident in the frustatograms. The RCL itself is enriched in highly frustrated interactions. High mutational frustration is observed both before and after cleavage in the region of the structure where the conformational change occurs, again

highlighting the role of local destabilization promoting conformational change. The structural ‘openness’ required to accommodate a newly inserted β -strand apparently requires sub-optimal interactions among residues in that region of the protein. There is also high local frustration in the α -helix that stacks over the face of the central β -sheet (Fig. 26). These regions have unusually rapid hydrogen/deuterium exchange dynamics (Tsutsui *et al.* 2008) and may remain metastable even in the cleaved form. Also it has been suggested that the frustrated helix may be easy to detach from the face of the β -sheet so as to initiate serpin aggregation.

Certain serpins (for example, plasminogen activator inhibitor-1 and antithrombin III (AT-III)) can spontaneously undergo a conformational rearrangement without RCL cleavage to form an inactive latent protein. Sometimes, ‘beads-on-a-string-type’ structures have been observed in electron micrographs, suggesting that this latent form is an ordered oligomer. In addition, mutations in some serpins promote the formation of highly stable, inactive polymers in the endoplasmic reticulum of liver cells. It was originally thought that the mutation might subvert the conformational flexibility required for the inhibitory activity of serpins, and that polymerization occurred through the insertion of the RCL of one serpin molecule into the β -sheet of another, and so on. However, Yamasaki *et al.* recently solved the structure of an antithrombin domain-swapped dimer, the structure of which revealed how any serpin adopting a similar configuration could form long-chain polymers. The domain swap involves a third of the β -sheet, which forms an equivalent region in another molecule (Fig. 26). In Yamasaki and colleagues’ dimer structure, the two molecules simply exchange part of the β -sheet with one another (a classic domain swap). The strand that swaps is the same strand that inserts during serpin maturation after cleavage (colored orange in Fig. 26). Examination of the frustration in the AT-III monomer both before and after thrombin cleavage reveals that as was seen in the α 1-antitrypsin, the base of the β -strand that swaps is frustrated (Fig. 26). In addition, the helix that is thought to have to dislodge from the face of the β -sheet in order to allow strand insertion and/or swapping is highly frustrated. Comparison of the ATIII monomer before and after cleavage reveals that the observed frustration is present both before and after strand insertion hinting that this part of the protein can readily unfold even after strand insertion. A few highly frustrated contacts are observed in the center of the swapped strand, but it is not obvious why the strand swap is favored. While the ubiquity of domain swapping is a consequence of minimal frustration (Yang *et al.* 2004), a frustrated region of a protein can nevertheless unfold more readily, encouraging partial domain swapping if the protein contains frustrated elements. To form a polymer capable of incorporating more than two monomers, extensive domain swapping among multiple serpin monomers is required, but it is readily apparent how this can occur from the structure (Yamasaki *et al.* 2008). The fact that the overlaying helix remains frustrated in all of the serpin forms suggests that domain swapping may occur even after serpin activation.

5.4. Frustration, dynamics, and catalysis

We have reviewed here how natural sequences are shaped via constraints imposed by the interaction of the protein with its immediate environment. These create conflicts in the folding of a polypeptide that can even compromise its foldability, become multistable or, in the extreme case, appear unfoldable, ‘intrinsically disordered’. On the other hand, these frustrating interactions facilitate excursions out of the thermodynamic minima so that the polypeptide can explore low-lying excited states that can become fundamental for function. For enzymes, an obvious requirement for activity is, in addition to foldability, catalytic power, which in the cellular context must

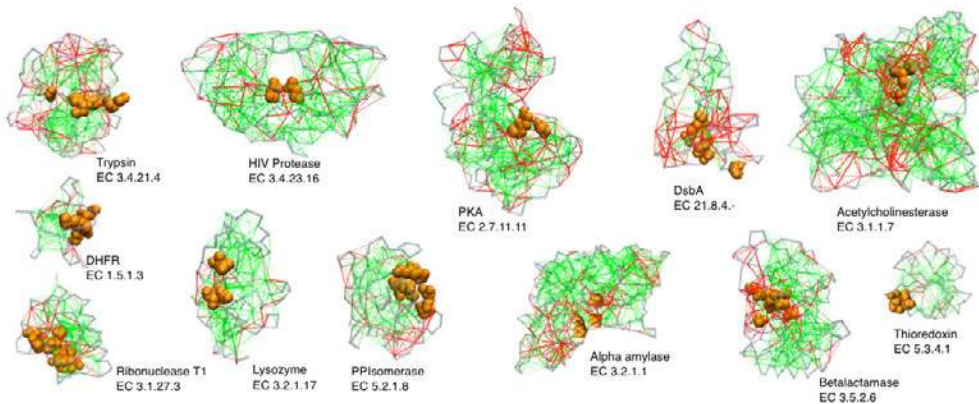


Fig. 27. Local frustration in enzymes. Catalytic sites are stringent, usually bringing together in space residues that would otherwise adopt different interactions. Examples of local frustration patterns in prototypic enzymes are shown. Orange spheres mark the catalytic residues annotated in the Catalytic Site Atlas (Porter *et al.* 2004). The catalytic residues are typically involved in highly frustrated interactions. Additionally, the catalytic sites are often surrounded by a dense minimally frustrated network of interactions. It is apparent that large enzymes also contain other surface patches of highly frustrated interactions that could mark other functional sites, such as the binding of an allosteric effector. These general aspects of frustration in enzymes appear to hold irrespective of enzyme class or catalytic mechanism, as classified by the E.C. numbers.

be jointly met with the requirement for regulation. Usually the regulation comes about through an allosteric site, allowing otherwise disconnected activities to become coupled, gratuitously, in functional space (Monod, 1973).

Catalysis requires detailed chemistry. Reaction rates strongly depend on the precise arrangement of reacting groups at sub-Å resolution. Because of this, catalytic sites are generally stringently conserved, bringing together in space residues that would otherwise adopt different interactions. Thus, in principle, we would expect natural enzymes to have locally frustrated catalytic sites. Some examples of frustration patterns of enzymes are shown in Fig. 27. In most cases the catalytic residues are involved in highly frustrated interactions either between them or with the rest of the molecule. Additionally however these sites are often surrounded by a dense minimally frustrated network of interactions that could be, in principle, related to the requirement of second-shell interactions to be favorable enough to compensate for the frustrating catalytic groups keeping it in place. It is apparent that large enzymes also contain other surface patches of highly frustrated interactions that mark other functional sites, such as the binding of an allosteric effector (Fig. 27). These general aspects of frustration in enzymes appear to hold irrespective of enzyme class or catalytic mechanism. This observation is consistent with the fact that mutations at catalytic sites often make the protein fold more stable (Meiering *et al.* 1992; Shoichet *et al.* 1995). Conversely, mutations that decrease the stability of folds are often associated with increased catalytic power, usually referred to as to stability-function trade-offs (Tokuriki *et al.* 2008). Many studies have noted that the evolution of new enzymatic activities is accompanied by a loss of protein stability. Also mutations that modulate activity are as destabilizing as any other random mutation (Tokuriki & Tawfik, 2009). On the other hand, substitutions in the substrate-binding pockets compromise stability to a larger extent than surface mutations that underline neutral, non-adaptive evolutionary changes. It is suspected that new enzymatic activities arise with the help of compensatory stabilization from otherwise silent mutations (Tokuriki *et al.* 2012).

These compensatory changes can be located in regions of the protein that were neutral to the stability of the original enzymes. The core of the protein would therefore act as a buffer for the accumulation of mutations in other regions. Enzymes are polypeptidic foldable systems for which many functional requirements simultaneously need to be met: bind substrate/s, catalyze, release product/s, bind effectors, switch on/off, and evolve. Each of these activities may require compromise with robust folding, resulting in sites of higher local frustration. In the cases of the enzymes shown in Fig. 27, it is still not possible to deconvolute the contribution of each of these effects to the observed frustration patterns. Many of the activities in the catalytic mechanism overlap: the site where the substrate is bound is necessarily close to the catalytic residues. Thus we cannot distinguish whether the local frustration comes about just for thermodynamic binding, for kinetic catalysis or for both. Experimentally these differing causes could be distinguished by mapping perturbations that differentially affect catalysis and substrate recognition. Perhaps these effects can be computationally distinguished by analyzing the changes in energy distributions upon binding, transition-state formation and conversion. Because of the sensitivity of chemistry to structural details, such a study requires fine-grained energy functions along with extensive sampling of conformational substates.

A common denominator to any catalytic mechanism is the requirement for motion. The details of how enzyme dynamics affect the activity have aroused much controversy (Karplus & McCammon, 1983; Olsson *et al.* 2006). The regulation of enzymatic activity is perhaps less controversial than the catalytic acts themselves. An interesting example is thrombin. Research on the coagulation protease thrombin, has suggested that its activity is allosterically regulated by changes in its dynamics. The core serine protease domain structure typified by trypsin and chymotrypsin is highly conserved in thrombin. It consists of two interacting β -barrel structures. Compared to trypsin and chymotrypsin, however, thrombin contains insertions at several of the loops connecting the β -strands. These loop insertions create a bipartite substrate (fibrinogen) binding site where fibrinogen is bound at the catalytic site as well as at a second site on the backside of the protease domain termed anion-binding exosite-1. This exosite is also the binding site of the allosteric regulator protein, thrombomodulin, and thrombomodulin binding alters the catalytic activity of thrombin toward a different substrate, protein C. Binding studies have demonstrated thermodynamic linkage between these two sites, but no differences are seen in the various bound structures determined by x-ray crystallography. Evidently the large change in enzymatic activity comes about from changes in the dynamics of the enzyme, which arise from a rather malleable ensemble of conformations, a situation that has been termed 'entropic allostery'. Recently, a combination of NMR and MD studies on thrombin has revealed the large extent of changes in motions on binding, particularly in the active site loops, spanning time scales from picoseconds to milliseconds (Fuglestad *et al.* 2013). To capture the longer time scale motions, accelerated molecular dynamics protocols were implemented, from which the average residual dipolar coupling constants for the backbone NH groups of the Boltzmann-reweighted ensembles could be back calculated and compared with NMR observables. By ascertaining which acceleration level produced the ensemble that best matched the NMR data, it was possible to explore excited states that were otherwise inaccessible but in fact occur in solution. Accelerated-MD of thrombin has identified that the changes upon active site occupation occur in groups of residues linked through networks of correlated motions and physical contacts (Fuglestad *et al.* 2013). Upon active site ligation, loop motions are quenched, but new ones appear connecting the active site with distal sites where allosteric regulators bind. The analysis of local frustration revealed that all of these loops are highly frustrated in the mean native structure of thrombin (Fig. 28).




Fig. 28. Slow conformational dynamics and local frustration. Catalytic mechanisms require motion. The dynamics of the protease thrombin change upon ligand binding. In particular, the active site loops display motions over several orders of magnitude – from picoseconds to milliseconds time scales. The ensembles of structures can be visualized with experimentally calibrated molecular dynamics. An ensemble of the representative structures of this protein is shown at top left, with dynamic loops identified with different colors. The local frustration pattern of the lowest energy structure from the ensemble is shown on the right. Minimally frustrated contacts are shown in green, highly frustrated ones in red, thin lines for water-mediated contacts. Below, a quantitative comparison of the order parameters S^2 reflecting nanosecond time scale motions S_{ns}^2 (gray) and longer time scale motions S_{AMD}^2 (red, black), with the local frustration distribution (cyan, blue). The average per residue fraction of highly frustrated contacts in the three lowest energy structures is shown with error bars corresponding to 1 s.d. The results from the inhibitor-bound ensemble are shown in the lower panel, the structural ensemble of the free protein in the top panel. Redrawn with the permission from Fuglestad *et al.* (2013).

We can use the computed structural models of the excited states of thrombin to look for the changes in frustration along the ensembles. Both the inhibitor bound and apo thrombin have regions of high frustration in the surface loops (Fig. 28). The regions of high frustration map to dynamic regions, as seen with NMR dynamic measurements obtained at different time scales. On the one hand, the order parameters derived from conventional MD simulations agree with order parameters derived from NMR relaxations experiments proving motions in the nanosecond time regime. On the other hand, order parameters derived from the NMR-calibrated accelerated MD identify disorder resulting from motions on longer time scales, typically microsecond to milliseconds. The fast dynamics mapped in this system do not correspond well to the regions of high local frustration, while the correspondence with the dynamics on longer time

scales is remarkable (Fig. 28). These results highlight the notion that regions of high local frustration are evolved features of the proteins that allow control of the motions on the slower time scales, which can occur with relative energetic ease. Energetically unfavorable interactions within the surface loops facilitate large amplitude motions, smoothing conversion between the ensembles of states needed for catalysis.

5.5 Frustration and the initiation of aggregation

Most of our biochemical illustrations of the frustration concept have focused on frustration within a single biomolecule or functional association, but what about frustration between biomolecules in the proteome which are not supposed to interact? Biomolecules are never alone during their life in the cell. Proteins are not even always born singly, but in prokaryotes especially are born in a litter – they are synthesized in ‘bursts’ when several ribosomes gather to translate a single messenger RNA. During this burst the local protein concentration can be quite high. After being born and becoming first folded, they will occasionally unfold and have to re-fold, again in the presence of neighbors, some of whom might be unfolded themselves. How do these unfolded proteins then avoid binding to the wrong partners? The principle of minimal frustration helps explain why even within the cell re-folding can be robust. (By the way it is obviously not always the case under all thermodynamic conditions – we can scramble eggs and do not expect them to unscramble when they cool!) The lack of frustration in finding partners is just what would be expected if achieving minimal frustration is a result of evolution. While the native contacts within a protein are quite stable, random contacts between two different proteins will not be as stable, so it will be rare for proteins with distinct sequences to be able to interact strongly with each other. Transient aggregates can unfold and eventually settle in their proper natively folded states. Clearly the dangerous partners that a protein must avoid are other copies of itself: native-like interactions between distinct partners can lead to association. This is called ‘domain swapping’ and if it goes on sequentially indeed an aggregate can form, as we have seen in the Serpin example (Fig. 26).

Domain swapping indeed does occur and seems to play a role in the initiation of some aggregation processes. The possibility of misfolding via domain swapping is a symmetry effect, the possibility of which cannot be easily eliminated by evolution. Simply symmetrizing the interactions of a structure-based pure funnel model for a monomer allows one to predict possible domain-swapped dimers (Yang *et al.* 2004). Different dimers differ in stability primarily through topological effects that determine the fraction of contacts that are prevented from being made in the more crowded dimer situation. Transient domain swapping seems to be a universal phenomenon when a protein is present at high concentrations but transient domain swapping resolves itself, suggesting the monomeric native form still has at least an entropic advantage at the concentrations found in most cells. A problem can occur if the local concentration is high. One of the worst cases of high concentration is in multidomain tandem repeat proteins, where the covalent connection means two identical domains can never get far away from each other. Evolution does seem to take notice. Han *et al.* have shown by surveying multidomain sequences in several genomes that in such multidomain repeat proteins the nearest-neighboring domains, even if structurally similar, are very distinct in sequence, while more sequentially separated domains often can still have pretty high-sequence identity (Han *et al.* 2007).

Domain swapping is likely to be a major reason aggregates formed in disease states usually involve a single protein. Zheng *et al.* computationally studied the misfolding of dimers of domains

in titin (Zheng *et al.* 2013). This system had been experimentally studied in the laboratory by Clarke and provided support for the ideas about domain swapping found in the bioinformatic survey (Wright *et al.* 2005). This simulation study showed that there was another source of misfolding – the ‘self-recognition’ of short sequences in two copies of the same protein. Figure 29 shows both the domain-swapped dimer and a misfolded dimer with the self-recognizing segments paired up. This self-recognition resembles the first steps in the formation of an amyloid. Indeed the self-recognizing fragment is recognized by software trained to recognize amyloidogenic fragments bioinformatically. In Fig. 29, we show histograms of the energies per residue of hexapeptide fragments in the I27 fused dimer system along with the native energy. The amyloid fragment comes perilously close to competing with the native structure locally. But it does not compete globally. A histogram of the global energies of native structures, self-recognizing structures, and domain-swapped structures are shown in Fig. 29. Here the native structure clearly wins. At low temperature folding should be fine but the amyloid-like structures can act as kinetic traps at higher temperature due to their disordered parts giving them entropic stability (Zheng *et al.* 2013). Globally the energy landscape is a funnel but still misfolding can give a kinetic trap, which can act as the initial state of further aggregation. We see that frustration can manifest itself more strongly in partially disordered ensembles even when it may not be always apparent in the final crystal structure. This example shows that looking at pair-level frustration as in typical frustratograms does not reveal all situations where energetic frustration can show up.

Pair-level frustration in the monomer can however encourage aggregation. An example is insulin. The metastability of insulin is revealed in its natural tendency to self-assemble into oligomers, a feature that was first appreciated when the structure was determined in the 1960s (Adams *et al.* 1969; Baker *et al.* 1988). The insulin monomer contains two polypeptide chains, an A chain of 21 amino acids and a B chain of 30 amino acids, which contain helices between residues A1 and A8, A13 and A19, and B9 and B19. The two chains are covalently linked by disulfide bridges between A7 and B7 and A20 and B19. Identical hydrophobic surfaces of two monomers come together forming an antiparallel beta-strand between residues B40 and B46 of each monomer. In the crystal, insulin is a hexamer consisting of three dimers aggregated around two zinc ions forming a soluble, globular protein structure (Baker *et al.* 1988). The hexamer has been seen to adopt one of three conformational states known as T6, T3R3, and R6, depending on the arrangement of the first eight residues of the B chain. In the T state (Fig. 30, PDB code 3P2X; residues B9–B12 are colored orange and residues B16–B19 are colored yellow), residues B9–B12 adopt an elongated conformation, while in the R state (PDB code: 1MPJ; same coloring scheme), they form part of the B-chain helix. Residues B9–B12 do not appear to be making many contacts with the rest of the insulin molecule, as has been suggested by NMR studies on the T-state (Hua *et al.* 2008). The A-chain is highly frustrated in all structures, and one might surmise that it is mainly held-together by the disulfide bonds it makes with the B-chain. Researchers have identified a single residue, a proline at position B28, which is largely responsible for making favorable contacts with residues B20–B23 of the adjacent monomer. Experiments with B-chain C-terminus substitutions and truncations have shown that residues B26–B30 are essential for dimer formation, and single substitutions at residue B28 markedly increase the tendency to form monomeric units upon dilution (Brems *et al.* 1992). We explored the residual local frustration in the B28 Pro – Asp mutant, which crystallized in an R6-like conformation (PDB code: 1ZEH). Plots of the fraction of contacts that are either minimally or maximally frustrated show that the region near the site of mutation is highly frustrated (Fig. 30*b*). Close examination of this region reveals that the Asp makes inter-chain electrostatic contacts that form an extended




Fig. 29. Frustration and the initiation of aggregation. Self-recognition of short sequences in two copies of the same protein sequence can lead to native-like interactions, resembling the first steps in the formation of an amyloid. The folding of tandem copies of TitinI27 or SH3 was simulated using AWSEM (Zheng *et al.* 2013). Besides native-like structures, misfolded structures with a significant amount of self-recognition interactions were found. The different levels of frustration in the tertiary contacts is illustrated in two representative structures for I27 dimer (left), SH3 dimer (right), with the backbone of the sub-domains in blue and yellow. The swapped contacts formed at the interface of SH3–SH3 are minimally frustrated, as expected from the principle of minimal frustration. The self-recognition contacts formed at the domain interface of I27–I27 are also minimally frustrated, indicating that these contacts are stronger than random contacts. Below comparisons of the stability (energy per residue) among various structures and ensembles of thermally sampled structures of I27 (left) and SH3 (right). The stability for the native monomeric structure is shown by a green vertical line, the strongest non-native hexapeptide pairing in magenta. Misfolded configurations are typically less stable than the native structure, indicating that misfolding by inappropriate pairing of strands will be unlikely in folding of the monomers for both I27 and SH3. Blue bars show the distribution of the stability of all of the self-recognition hexapeptide pairs, calculated with the AWSEM-Amylometer based on the energies of a β -sheet formed from the hexapeptide (Zheng *et al.* 2013). If the stability of the strongest self hexapeptides pair is competitive with the native structure, as in the case of I27–I27, the particular self-pair is responsible for the misfolding of the fused protein and can trigger further aggregation. For SH3–SH3, the fused protein folds as well as the monomer, because all self-hexapeptides pairings are weaker than the most stable non-native hexapeptides pairing in the monomer. At low temperature, folding to the native state is robust, but the amyloid-like structures can act as kinetic traps. At higher temperature their disordered parts give them an entropic stability advantage. They can act in the initial states for further aggregation. Frustration can be manifested more strongly in partially disordered ensembles and may not always be apparent in the final crystallographic structure. These examples show that looking at the pair-level frustration as in typical frustratograms may not reveal all situations, whereas energetic frustration can show up. Redrawn with the permission from Zheng *et al.* (2013).




Fig. 30. Frustration in a monomer can encourage aggregation. Insulin crystallizes in either the (a) T-state (PDB code 2P2X), exemplified by the loss helical structure in the N-terminal residues of the B-chain helix (colored orange) as compared to (b) the R-state (PDB code 1ZEH), which is stabilized by the Pro(B28)Asp mutation and by binding of phenol or cresol molecules to pocket formed between the N-terminal residues of the B-chain helix and the the C-terminal residues of the B-chain (purple). Such structural malleability is typically indicative of frustration, and indeed insulin is highly frustrated. On occasion, insulin can cause an amyloid disease and Weiss managed to crystallize the fibrillar form (c) (Weiss, 2013). Interestingly, the molecular interactions observed in the crystal structure showed the participation of the A-chains (cyan in all structures), the most frustrated part of the insulin structure. The fibrillar form, however, appears to remain frustrated.

backbone structure; however, these contacts are frustrated. The increased local frustration is likely the cause of the increased conformational flexibility in the B-chain C-terminus observed in this mutant (Whittingham *et al.* 1998). It is interesting to speculate that the frustration might partly be ameliorated by the binding of small aromatic molecules such as phenol or cresol, which bind in this region of the protein, and may provide stabilization of the R6 conformation.

Insulin can, on occasion cause an amyloid disease. This phenomenon was investigated recently by Weiss and colleagues who managed to crystallize the fibrillar form of insulin (Weiss, 2013). Interestingly, the molecular interactions observed in the crystal structure showed the participation of the A-chains, the most frustrated part of the insulin structure. No obvious alterations in the amount of minimally frustrated contacts were observed in the fibrillar form, however (Fig. 30c).

5.6. Frustration in nucleic acids and protein complexes

We have reviewed here why proteins are such a special kind of polymers: minimally frustrated globule-like structures can be encoded within an amino acid sequence. The richness of fold types is constrained at every level by the variety of the alphabet and the environmental conditions where the correct decoding occurs. Energy landscape theory shows that in the case of polypeptides the present 20-mer alphabet is more than sufficient to allow robust folding of a vast collection of topological families. However, proteins are not the only polymers found in living organisms. How does the concept of frustration apply to other biological polymers?

In terms of dry mass, polysaccharides abound in cells. Some of them are simple linear concatenations of the same monosaccharide units (cellulose), others display branched covalent

structures (glycogen, starch), and even more complex varieties are still being characterized. Admittedly more needs to be done with respect to their understanding. Yet it seems that most polysaccharides do not spontaneously fold to specific structures, but rather form disordered colloids. Although the variety of saccharide units is very high due to local modifications, the chemical ligation properties of the individual units is sufficiently uniform along the chain, so that the interactions between any two parts is as nearly (de)stabilizing as any other. In energy landscape terms, the energetic variance of alternative structures ΔE^2 is small but the gap to the ground state appears to be nearly negligible, and a glass transition is expected to be predominant (Honey is a good example!). Here frustration is the rule that plagues the landscape with local minima. Not surprisingly specific synthesis and handling of even the most simple biological polysaccharides is challenging, and much could be gained if this were not the case (imagine making paper out of plain sugar!).

However, there is a special sort of polysaccharide that does spontaneously fold. Nucleic acids are essentially linear phosphopentoses in which each unit is modified with a nitrogen-rich base, typically one out of four possibilities. In these polymers, the unit–unit interaction modes are restricted, with electrostatics dominating repulsion and hydrophobic base–base interactions favoring association. The most stable structures are periodic packings of bases between two or more chains forming helices. Ever since the proposal of the antiparallel double-helix (Watson *et al.* 1953), it has been evident that this structure is to a certain extent nearly independent of the sequence of the monomers, as long as specific base-pairing is maintained. This makes nucleic acids the prime molecular species where faithful templated polymerization can occur. Spontaneous formation of the double-helical structure is at the core of genetic engineering technologies, yet our present understanding of the folding landscape of DNA is limited. The melting temperature of short chains can be well approximated with simple empirical rules that take into account the unit composition and the saltiness of the solution (Schildkraut *et al.* 1965). This corresponds with a landscape in which the ‘native’ structure can be reached with little competition of alternative states, with a large gap dominating the energetics. However, longer chains display more complicated behaviors, with marked hysteresis, history and concentration-dependent kinetics, all hints of glassiness of the landscape. It should be kept in mind that the differences in stabilization energy of specific base-pairing interactions is low compared to the mere stacking of the bases, and thus the discrimination between correct and incorrect base-pairing is expected to become more difficult as the chain grows longer. In contrast to proteins, nucleic acid folding is dominated by the formation of continuous secondary structure stretches, with small relative contribution of long-range tertiary interactions, which under ordinary conditions are discouraged by electrostatics. The apparent simplicity of the short oligonucleotide folding allows for engineering large assemblies by adjusting local rules (Rothmund, 2006). Indeed statistical mechanical theories of the nucleic acid structure formation suggest that due to its quasi-one-dimensional character the landscape is naturally hierarchical (Fernández & Shakhnovich, 1990).

The most common source of frustration in nucleic acids is the competition between base-pairing continuous stretches of a single chain and between similar chains. The formation of alternative, metastable stem-loop structures was early suggested by Crick as a plausible mechanism for gene regulation (Crick, 1971). In recent years, single molecule spectroscopy has allowed for detailed characterization of the overall time scales in which such transitions occur, providing direct evidence of the existence of many collapsed intermediate structures (Ma *et al.* 2007). The free-energy landscapes drawn out from these and other experiments portray pictures of overall smooth surfaces with discrete metastable traps (Woodside *et al.* 2006). The location of the

traps and energy barriers separating them can be readily manipulated by adjusting the hairpin sequence, such that the population of folding intermediates can be partially controlled introducing mismatches and changing loop length. Frustration between these structures provides the grounds for engineering nucleic acid polymers that can act as switches, with ingenious applications envisioned (Viasnoff *et al.* 2006).

The natural occurrence of functional aspects of frustration in nucleic acids is best exemplified by RNA, not DNA. Small stereochemical restrictions in the dihedral angles in the backbone make the small alphabet of ribonucleotide sequences generally not sufficient to specify a set of secondary and tertiary interactions. Studies on the folding of catalytic RNAs (ribozymes) have highlighted the peculiar propensity of RNA to adopt alternate conformations in solution. This conformational heterogeneity is attributed to misfolded structures or long-lived folding intermediates and it is frequently observed in natural isolates. The high probability of even random RNA sequences to have fairly stable stretches of base-pairs makes the number of alternate structures significant. Still, the RNA-folding complexity is reduced by its more hierarchical nature as compared to proteins. Typically secondary structures form fast and last for minutes to hours at physiological temperatures, whereas tertiary interactions are comparatively weak, interconverting in the time scale of milliseconds to seconds (Solomatin *et al.* 2010). This marked difference in stability and kinetics constrains the accessible tertiary structures. Conversely, isolated secondary structural elements of larger RNAs often fold properly in absence of tertiary interactions. This phenomenon underlies the successful prediction of RNA structure by two-dimensional mixing and matching of fragments, and also gives rise to deep folding intermediates separated by large barriers. These long-lived alternatives can be coopted to perform the specific functions and should be treated as alternative 'native' states. Such functional alternatives have been described in the autocatalytic Tetrahymena Group I intron, sarcin-ricin loop, the ribosome, and *in vitro* selected aptamers (Marek *et al.* 2011). Furthermore, the stability of these frustrated states can depend on the differential binding to other molecules, as exemplified in riboswitches (Serganov & Nudler, 2013). In these examples, binding of a small ligand to a segment of mRNA is sufficient to affect a structural transition in the mRNA that in turn regulates the translation rate, making them effective modulators of protein abundance.

Frustration has an important functional role when proteins meet up with DNA. Some proteins display the remarkable ability to differentially recognize a specific stretch of DNA out of a vast excess of structurally similar sites. Transcription factors locate their cognate site in a genome with rates that appear to break the diffusion speed limit. Delbrück first envisioned a mechanism to explain this puzzle proposing that proteins may bind to any site on a linear chain of DNA and then slide along it, reducing the three-dimensional search problem to one dimension. Since then the existence of such 'facilitated diffusion' mechanisms has been confirmed in a variety of systems (von Hippel & Berg, 1989). This mechanism requires the protein to have a relatively small variance of the sequence-dependent energy landscape of interaction, so that fast sliding can occur. Gerland *et al.* have analyzed this with the same mathematics employed by Bryngelson and Wolynes for folding proteins (Gerland *et al.* 2002). At the same time however a DNA-binding protein must also display a relatively large thermodynamic stability at its target site. The requirement of both fast search and stability of the cognate protein-DNA complex impose conflicting constraints on the landscape: whereas rapid search requires a fairly smooth landscape, the stability of the protein-DNA complex requires low energy at the target site. Slutsky & Mirny (2004) showed that this apparent paradox can be ameliorated if the protein occupies (at least) two distinct conformational states that differ in the variance of the binding

energy distribution. In one state, the protein binds DNA sites with a small variance, thus it cannot strongly discriminate different sequences but allows sliding to occur rapidly. The other state is characterized by a large variance with greater differences in affinity for the various sequences. The differences in free energy and the switching rate between the protein states define an energy gap. In this model, the search time to stably bind a cognate site is optimized when the protein spends most of the time in the state that favors sliding but can undergo fast conformational transitions to the other state. In this scenario, frustration in protein–DNA binding effectively creates a landscape with only a little ruggedness and a few alternate sites with low energies (Slutsky & Mirny, 2004).

Biochemical studies have highlighted the fact that DNA-binding proteins bind non-specifically to any DNA stretch (or any polycationic surface for that matter), with electrostatic interactions predominating the energetics. The surfaces of DNA-binding proteins usually display positively charged patches that complement the negatively charged DNA backbone, providing a rationale for the lack of sequence specificity. On the other hand, the structures of several specific protein–DNA complexes show that these are stabilized by intimate contacts between residues and bases that could discriminate the subtle structural differences between different DNA sequences, mostly present at the bottom of the major groove (Norambuena & Melo, 2010). The differential ‘binding modes’ may thus be related to the formation of distinct interaction regions. A recent survey of the structures of DNA-binding proteins has found that some proteins utilize one interface to specifically recognize a cognate sequence while another region is employed to establish non-specific associations, resembling two distinct active sites. In contrast, sometimes these two patches overlap (Marcovitz & Levy, 2011). The system now has to choose between these alternatives, frustrating the protein–DNA interaction. The free-energy barrier between the binding modes can in principle be related to the overlap of these binding surfaces. MD simulations of the sliding/binding process show that as the local frustration increases, the rate of switching between sliding and binding decreases (Marcovitz & Levy, 2013). On the one hand, high frustration within the protein smoothes the landscape for sliding, but it makes the switching to the specific binding mode slower, which slows down the overall target location rate. On the other hand, if frustration is lowered too much switching becomes rare and the protein often misses the cognate site. The situation encountered is reminiscent to the effect frustration has in folding a polypeptide chain: a little is helpful, lots can be disastrous.

The formation of a specific protein–DNA complex may be coupled to the folding of the protein. While protein flexibility is involved in induced-fit recognition and to a larger extent in intrinsically unstructured DNA-binding proteins, the electrostatic field of the DNA can induce flexibility and a partial unfolding of the protein (Levy *et al.* 2007). Protein disorder facilitates the three-dimensional diffusive search through the ‘fly-casting’ mechanism where the protein binds DNA non-specifically in a partially folded state and completes its folding when it binds the specific site (Shoemaker *et al.* 2000). Indeed, several DNA-binding proteins are known to be partially unstructured in the unbound state, which only fold upon binding to the target (Spolar & Record, 1994). The interplay between fly-casting and electrostatics is observed even for weak electrostatic forces and is expected to vary with the electrostatic screening due to salt and the intrinsic folding barrier, both of which can be modulated. Increasing the strength of the electrostatic forces destabilizes the isolated protein due to frustration between general electrostatic interactions and the specific native interactions. In the presence of DNA, a larger destabilization effect is seen as a result of competition between folding and binding (Levy *et al.* 2007). This latter frustration reflects a ‘tidal force’ from electrostatics tending to pull the molecule apart.

Under strong electrostatic forces, the negatively charged DNA pulls the protein apart, and the protein–DNA association can therefore advance protein folding.

Papillomavirus E2c protein is a remarkable model in which the effects of frustration in protein–DNA association have been experimentally studied. When confronted with a small double-stranded oligonucleotide the E2c protein finds its specific site using two parallel routes (Sánchez *et al.* 2011). In one of them the process is well described as a two-state reaction with a small activation barrier of entropic origin. The effect of single-point mutations on the DNA-contacting surface to the kinetics and thermodynamics indicated that the formation of direct contacts between residues and bases constitutes the bottleneck of complex formation (Ferreiro *et al.* 2008a). Few non-specific ionic interactions mainly stabilize the final consolidated complex. The binding landscape along this route is smooth, consistent with a hydrated protein–DNA interface. In contrast, binding along the other route involves several steps. Attaining an initial, diffuse, transition-state ensemble of structures with some non-native contacts is followed by formation of extensive non-native interactions that drive the system into a kinetic trap (Sánchez *et al.* 2010). These non-native contacts then slowly rearrange into native-like interactions. Dissimilar protein–DNA interfaces are formed in the intermediate and the final complex, as anticipated from the computational studies described above. Frustration between the binding modes creates a degeneracy of interactions that drives the complex into long-lived non-native species (Sánchez *et al.* 2010). It has been speculated that this slowing down of the binding reaction allows for extra functionality, as this metastable species exposes different surfaces that could be linked with other binding activities, as usually seen in eukaryotic protein–nucleic acid complexes.

6. Frustration is a fundamental aspect of biochemistry

Some biologists take for granted that proteins just ‘do-their-job’ as if they were robots made by some blind adaptive designer. Physicists are often amazed by the intricacies of how this peculiar kind of chemical scum can organize itself at the temperatures the biosphere experiences. We have explored the foundations of the concept of *Frustration* and how it bridges these fields in the study of biomolecules. The impossibility of simultaneously satisfying constraints leads to degeneracy of states. Finite degeneracy allows for captivating emergent behaviors where apparently unrelated parts of a system respond to subtle changes in local energetics. Too much frustration and degeneracy leads to qualitatively different dynamics where searching through the many states cannot be carried out in a meaningful time and the system freezes out in a glass. On the other hand, systems that lack any frustration at all are well behaved, display sharp phase transitions, get ordered spontaneously in the right conditions and, much like crystals, are beautiful but rather dull. As in many other realms of natural processes, interesting things happen within a balance between orderly and chaotic regimes (Frauenfelder *et al.* 2003; Kauffman, 2013).

The principle of minimal frustration has led to an understanding of the general aspects of protein folding. It has fruitfully led to testable hypotheses, successful interpretations of a diverse range of experiments and to the development of useful computational tools. As with any good theory, the deviations from its expectations mark situations in which discoveries await. Frustration opens up possibilities for degeneracy, recognizing proteins as metastable systems that can compute and modify their structures according to environmental inputs. In turn, the fact that these computations are robust to small perturbations make proteins evolvable, such that the overall behavior can be finely tuned, and trained, over time (Bray, 1995). Present-day protein molecules contain information both about historical chances and physical necessities

imprinted in their structures. It is still challenging for humans to distinguish the relevant information from the inevitable noise that overwhelms our current intellectual capabilities.

Studying the aspects of frustration averaged over many proteins provides ways to infer energy functions for structure prediction and has allowed us to appreciate the apparent complexity of the protein-folding code. Frustration affects laboratory folding kinetics and deviations from perfectly funneled landscapes reflect previously unrecognized functional constraints. We have reviewed here how a large part of the functions of proteins is related to subtle local frustration effects and how frustration influences the appearance of metastable states, the nature of binding processes, catalysis, and allosteric transitions. ‘Doing something’ is the prime characteristic of protein molecules. The teleonomic, apparent purposefulness of biological function often involves several ‘activities’ coming together in a coordinated way (Monod, 1973). At any scale it is challenging to disentangle what are the meaningful parts of living things and what are the relevant interactions between them (Simondon, 2005). The apparent individuality of the polypeptides we can purify from the present-day organisms is restricted by our skewed view and becomes highlighted when a theory based on statistics is available. The study of life beyond the scale of single proteins has to treat basic aspects of the stability and dynamics of finite states, and as such we envision that frustration is a fundamental concept that will help (often frustrated) researchers for many years.

The huge number of proteins in a cell requires that proteins must interact with a relatively high degree of specificity, otherwise most of the cellular matter would be tied up in a mess of inappropriate interactions, and instead of cells we would have scrambled egg cells. This functional specificity has been achieved by the coevolution of sequences that also lead to little scrambling even within each individual protein molecule. Low-free-energy excited state ensembles for a biomolecule can come about in several ways: harmonic motions of large length scale or special harmonic motions utilizing weakly connected links, local folding/unfolding on a funneled landscape, frustrated local interactions giving energetically degenerate structures and sequence symmetry that gives rise to alternate structures. Various combinations of such motions with low-free-energy costs occurring at the same time will also give low-free-energy excited states and kinetically dominant reaction paths. An example of such a combination mode is ‘cracking’ in which the system makes excursions along a low-frequency normal mode (not costly in energy), during this excursion the molecule locally unfolds (not too costly in free energy) and then refolds. If the region that unfolds is frustrated, so much the better: its unfolding will be still faster and cheaper in free-energy terms. Cracking seems to be a likely mechanism for many physiologically relevant biochemical processes.

The problem of self-organization for large nanoscale objects is not stability but the kinetics of getting to a stable structure. Kinetic barriers can be avoided if sufficient entropy remains to fluidize the motion. Up to a point, robust self-assembly in a cell therefore often relies on folding upon binding – a sort of a just-in-time policy for macromolecular economy. Nevertheless, kinetic barriers for folding (even for folding upon binding) tend to scale with the size of the assembled system. The larger barriers occur for frustrated systems but also may grow with size even when the minimal frustration principle is respected, as in the capillarity picture. At some size scale the assembly rates generally become too slow and far-from-equilibrium assembly takes over. This crossover already seems to occur when the largest proteins are folded *in vivo*. Rubisco, a very large protein, requires chaperones in order to fold. It is likely the chaperoned catalysis of proper folding is actually a kind of kinetic proofreading, in which misfolded protein molecules are ripped apart by unfolding machines and given a chance to try again (Gulukota & Wolynes, 1994; Todd *et al.* 1996). The role of the ATP-induced motions of a chaperone can be thought of as

increasing the effective temperature for a limited fraction of the degrees of freedom of the protein.

Cells are not formless bags of molecules. Organization persists even at the cellular scale and changes its nature at this supramolecular or ultra structural level. We are only beginning to probe the atomically weak but colloidal strong forces that act on this larger scale length. It is also clear that in living things, motion on this size scale is not just equilibrium Brownian motion. Instead, chemical energy is flowing through the system: various molecular motors carry parts around in a hurly-burly reminiscent of a large city or ant colony (Wang & Wolynes, 2011). This motorization probably allows the problem of large barriers and high glass transition temperatures to be overcome much as in kinetic proofreading by chaperones where ATP is used. By comparing responses to fluctuations on the colloidal scale (Lau *et al.* 2003), an effective temperature of the cytoskeleton has been measured to be 10000 K. Obviously this high temperature applies to only a minute fraction of the degrees of freedom. Studying the organization of living matter on this longer length scale will require new tools, both experimental and theoretical. We can anticipate, however, that the strong foundation of energy landscape concepts learned in the study of molecular folding and function, will be of great help in finding the required conceptual framework for this far-from-equilibrium regime.

7. Acknowledgements

We thank Weihua Zheng, Nick Schafer, Bobby Kim, Ha Truong, and Rocío Espada for stimulating discussions and for providing help with figures. Supported by Grants R01 GM44557 and P01 GM071862 from the National Institute of General Medical Sciences. Additional support was also provided by the D. R. Bullard-Welch Chair at Rice University, the Agencia Nacional de Promoción Científica y Tecnológica (ANPyCT) and the Consejo Nacional de Investigaciones Científicas y Técnicas de Argentina (CONICET).

8. References

- ADAMS, M. J., BLUNDELL, T. L., DODSON, E. J., DODSON, G. G., VJAYAN, M., BAKER, E. N., HARDING, M. M., HODGKIN, D. C., RIMMER, B. & SHEAT, S. D. (1969). Structure of rhombohedral 2 zinc insulin crystals. *Nature* **224**, 491–495.
- AKSEL, T. & BARRICK, D. (2009). Analysis of repeat-protein folding using nearest-neighbor statistical mechanical models. *Methods in Enzymology* **455**, 95–125. doi: 10.1016/S0076-6879(08)04204-3.
- ANDERSON, P. W. (1950). Antiferromagnetism. Theory of superexchange interaction. *Physical Review* **79**, 350–356. doi: 10.1103/PhysRev.79.350. URL: <http://link.aps.org/doi/10.1103/PhysRev.79.350>.
- ANDRÉ, I., STRAUSS, C. E. M., KAPLAN, D. B., BRADLEY, P. & BAKER, D. (2008). Emergence of symmetry in homooligomeric biological assemblies. *Proceedings of the National Academy of Sciences of the United States of America* **105**, 16148–16152. doi: 10.1073/pnas.0807576105.
- ANFINSEN, C. B. (1973). Principles that govern the folding of protein chains. *Science* **181**, 223–230.
- Anonymous (2013). Frustration. *Wikipedia, the free encyclopedia*. URL: <http://en.wikipedia.org/wiki/Frustration> (visited on 2013).
- BAKER, E. N., BLUNDELL, T. L., CUTFIELD, J. F., CUTFIELD, S. M., DODSON, E. J., DODSON, G. G., HODGKIN, D. M., HUBBARD, R. E., ISAACS, N. W. & REYNOLDS, C. D. (1988). The structure of 2Zn pig insulin crystals at 1.5 Å resolution. *Philosophical Transaction of the Royal Society B, Biological Sciences* **319**, 369–456.
- BARRICK, D., FERREIRO, D. U. & KOMIVES, E. A. (2008). Folding landscapes of ankyrin repeat proteins: experiments meet theory. *Current Opinion in Structural Biology* **18**, 27–34. doi: 10.1016/j.sbi.2007.12.004.
- BEASLEY, J. R. & HECHT, M. (1997). Protein design: the choice of *de novo* sequences. *Journal of Biological Chemistry* **272**, 2031–2034.
- BERMAN, H. M., WESTBROOK, J., FENG, Z., GILLILAND, G., BHAT, T. N., WEISSIG, H., SHINDYALOV, I. N., & BOURNE, P. E. (2000). The protein data bank. *Nucleic Acids Research* **28**, 235–242.

- BEST, R. B. (2012). 'Atomistic molecular simulations of protein folding. *Current Opinion in Structural Biology* **22** (1), 52–61. doi: 10.1016/j.sbi.2011.12.001.
- BEST, R. B., HUMMER, G. & EATON, W. A. (2013). Native contacts determine protein folding mechanisms in atomistic simulations. *Proceedings of the National Academy of Sciences of the United States of America* **110**(44), 17874–17879.
- BINZ, H. K., STUMPP, M. T., FORRER, P., AMSTUTZ, P. & PLÜCKTHUN, A. (2003). Designing repeat proteins: well-expressed, soluble and stable proteins from combinatorial libraries of consensus ankyrin repeat proteins. *Journal of Molecular Biology* **332**, 489–503.
- BORGIA, A., WENSLEY, B. G., SORANNO, A., NETTELS, D., BORGIA, M. B., HOFFMANN, A., PFEIL, S. H., LIPMAN, E. A., CLARKE, J. & SCHULER, B. (2012). Localizing internal friction along the reaction coordinate of protein folding by combining ensemble and single-molecule fluorescence spectroscopy. *Nature Communications* **3**, 1195. doi: 10.1038/ncomms2204.
- BRADLEY, L. H., THUMFORT, P. P. & HECHT, M. H. (2006). *De novo* proteins from binary patterned combinatorial libraries. *Methods in Molecular Biology* **340**, 53–69. doi: 10.1385/1-59745-116-9:53.
- BRAY, D. (1995). Protein molecules as computational elements in living cells. *Nature* **376**, 307–312. doi: 10.1038/376307a0.
- BREMS, D. N., ALTER, L. A., BECKAGE, M. J., CHANCE, R. E., DIMARCHI, R. D., GREEN, L. K., LONG, H. B., PEKAR, A. H., SHIELDS, J. E. & FRANK, B. H. (1992). Altering the association properties of insulin by amino acid replacement. *Protein Engineering* **5**, 527–533.
- BRENNER, S. E., CHOTHIA, C. & HUBBARD, T. J. (1997). Population statistics of protein structures: lessons from structural classifications. *Current Opinion in Structural Biology* **7**, 369–376.
- BROWN, B. M. & SAUER, R. T. (1999). Tolerance of Arc repressor to multiple-alanine substitutions. *Proceedings of the National Academy of Sciences of the United States of America* **96**, 1983–1988.
- BRYUJANTS, G. & REDFIELD, C. (2009). 15N NMR relaxation data reveal significant chemical exchange broadening in the α -domain of human α -lactalbumin. *Biochemistry* **48**, 4031–4039.
- BRYNGELSON, J. D., ONUCHIC, J. N., SOCCI, N. D. & WOLYNES, P. G. (1995). Funnels, pathways, and the energy landscape of protein folding: a synthesis. *Proteins* **21**, 167–195.
- BRYNGELSON, J. D. & WOLYNES, P. G. (1987). Spin glasses and the statistical mechanics of protein folding. *Proceedings of the National Academy of Sciences of the United States of America* **84**, 7524–7528.
- BULLOUGH, P. A., HUGHSON, F. M., SKEHEL, J. J. & WILEY, D. C. (1994). Structure of influenza haemagglutinin at the pH of membrane fusion. *Nature* **371**, 37–43. doi: 10.1038/371037a0.
- CAPALDI, A. P., KLEANTHOUS, C. & RADFORD, S. E. (2002). Im7 folding mechanism: misfolding on a path to the native state. *Nature Structural Biology* **9**, 209–216. doi: 10.1038/nsb757.
- CAPRARO, D. T., ROY, M., ONUCHIC, J. N., GOSAVI, S. & JENNINGS, P. A. (2012). beta-Bulge triggers route-switching on the functional landscape of interleukin-1 beta. *Proceedings of the National Academy of Sciences of the United States of America* **109**, 1490–1493. doi: 10.1073/pnas.1114430109.
- CHANG, J.-Y. (2011). Diverse pathways of oxidative folding of disulfide proteins: underlying causes and folding models. *Biochemistry* **50**, 3414–3431. doi: 10.1021/bi200131j.
- CHANGEUX, J.-P. (2013). 50 years of allosteric interactions: the twists and turns of the models. *Nature Reviews Molecular Cell Biology*. doi: 10.1038/nrm3695.
- CHAUDHARI, P. & TURNBULL, D. (1978). Structure and properties of metallic glasses. *Science* **199**, 11–21.
- CHAVEZ, L. L., ONUCHIC, J. N. & CLEMENTI, C. (2004). Quantifying the roughness on the free energy landscape: entropic bottlenecks and protein folding rates. *Journal of the American Chemical Society* **126**, 8426–8432. doi: 10.1021/ja049510+.
- CHO, S. S., LEVY, Y. & WOLYNES, P. G. (2006). P versus Q: structural reaction coordinates capture protein folding on smooth landscapes. *Proceedings of the National Academy of Sciences of the United States of America* **103**, 586–591. doi: 10.1073/pnas.0509768103.
- CHO, S. S., LEVY, Y. & WOLYNES, P. G. (2009). Quantitative criteria for native energetic heterogeneity influences in the prediction of protein folding kinetics. *Proceedings of the National Academy of Sciences of the United States of America* **106**, 434–439. doi: 10.1073/pnas.
- CIERPICKI, T. & OTLEWSKI, J. (2002). NMR structures of two variants of bovine pancreatic trypsin inhibitor (BPTI) reveal unexpected influence of mutations on protein structure and stability. *Journal of Molecular Biology* **321**, 647–658.
- CLARKE, N. D. (1995). Sequence 'minimization': exploring the sequence landscape with simplified sequences. *Current Opinion in Biotechnology* **6**, 467–472.
- CLEMENTI, C. (2008). Coarse-grained models of protein folding: toy models or predictive tools? *Current Opinion in Structural Biology* **18**, 10–15. doi: 10.1016/j.sbi.2007.10.005.
- CLEMENTI, C., NYMEYER, H. & ONUCHIC, J. N. (2000). Topological and energetic factors: what determines the structural details of the transition state ensemble and 'en-route' intermediates for protein folding? An investigation for small globular proteins. *Journal of Molecular Biology* **298**, 937–953. doi: 10.1006/jmbi.2000.3693.
- CLEMENTI, C. & PLOTKIN, S. S. (2004). The effects of non-native interactions on protein folding rates: theory and simulation. *Protein Science* **13**, 1750–1766. doi: 10.1110/ps.03580104.

- COLLINS, M. F. & PETRENKO, O. A. (1997). Review/synthèse: triangular antiferromagnets. *Canadian Journal of Physics* **75**, 605–655.
- COOPER, S., KHATIB, F., TREUILLE, A., BARBERO, J., LEE, J., BEENEN, M., LEAVER-FAY, A., BAKER, D., POPOVIĆ, Z. & PLAYERS, F. (2010). Predicting protein structures with a multiplayer online game. *Nature* **466**, 756–760. doi: 10.1038/nature09304.
- CREIGHTON, T. E. & GOLDENBERG, D. P. (1984). Kinetic role of a meta-stable native-like two-disulphide species in the folding transition of bovine pancreatic trypsin inhibitor. *Journal of Molecular Biology* **179**, 497–526.
- CRICK, F. (1971). General model for the chromosomes of higher organisms. *Nature* **234**, 25–27.
- DAVYAN, A., SCHAFER, N. P., ZHENG, W., CLEMENTI, C., WOLYNES, P. G. & PAPOIAN, G. A. (2012). AWSEM-MD: protein structure prediction using coarse-grained physical potentials and bioinformatically based local structure biasing. *Journal of Physical Chemistry B* **116**, 8494–8503. doi: 10.1021/jp212541y.
- DERRIDA, B. (1981). Random-energy model: an exactly solvable model of disordered systems. *Physical Review B* **2**, 2613.
- DIXIT, A. & VERKHIVKER, G. M. (2011). ‘The energy landscape analysis of cancer mutations in protein kinases. *PLoS ONE* **6**, e26071. doi: 10.1371/journal.pone.0026071.
- EASTWOOD, M. P., HARDIN, C., LUTHEY-SCHULTEN, Z. & WOLYNES, P. G. (2001). Evaluating protein structure–prediction schemes using energy landscape theory. *IBM Journal of Research and Development* **45**, 475–497.
- EJTEHADI, M. R., AVALL, S. P. & PLOTKIN, S. S. (2004). Three-body interactions improve the prediction of rate and mechanism in protein folding models. *Proceedings of the National Academy of Sciences of the United States of America* **101**, 15088–15093. doi: 10.1073/pnas.0403486101.
- ENGLANDER, S. W. (2000). Protein folding intermediates and pathways studied by hydrogen exchange. *Annual Review of Biophysics and Biomolecular Structure* **29**, 213–238.
- FERNÁNDEZ, A. & SHAKHNOVICH, E. I. (1990). Activation-energy landscape for metastable RNA folding. *Physical Review A* **42**, 3657–3659.
- FERREIRO, D. U., HEGLER, J. A., KOMIVES, E. A. & WOLYNES, P. G. (2007). Localizing frustration in native proteins and protein assemblies. *Proceedings of the National Academy of Sciences of the United States of America* **104**, 19819–19824. doi: 10.1073/pnas.0709915104.
- FERREIRO, D. U., HEGLER, J. A., KOMIVES, E. A. & WOLYNES, P. G. (2011). On the role of frustration in the energy landscapes of allosteric proteins. *Proceedings of the National Academy of Sciences of the United States of America* **108**, 3499–3503. doi: 10.1073/pnas.1018980108.
- FERREIRO, D. U. & KOMIVES, E. A. (2007). The plastic landscape of repeat proteins. *Proceedings of the National Academy of Sciences of the United States of America* **104**, 7735–7736. doi: 10.1073/pnas.0709915104.
- FERREIRO, D. U. & KOMIVES, E. A. (2010). Molecular mechanisms of system control of NF-kappaB signaling by IkappaBalpha. *Biochemistry* **49**, 1560–1567. doi: 10.1021/bi901948j.
- FERREIRO, D. U., SÁNCHEZ, I. E. & DE PRAT GAY, G. (2008a). Transition state for protein–DNA recognition. *Proceedings of the National Academy of Sciences of the United States of America* **105**, 10797–10802. doi: 10.1073/pnas.0802383105.
- FERREIRO, D. U., WALCZAK, A. M., KOMIVES, E. A. & WOLYNES, P. G. (2008b). The energy landscapes of repeat-containing proteins: topology, cooperativity, and the folding funnels of one-dimensional architectures. *PLoS Computational Biology* **4**, e1000070. doi: 10.1371/journal.pcbi.1000070.
- FERREIRO, D. U. & WOLYNES, P. G. (2008). The capillarity picture and the kinetics of one-dimensional protein folding. *Proceedings of the National Academy of Sciences of the United States of America* **105**, 9853–9854. doi: 10.1073/pnas.0805287105.
- FERSHT, A. R., MATOUSCHEK, A. & SERRANO, L. (1992). The folding of an enzyme. I. Theory of protein engineering analysis of stability and pathway of protein folding. *Journal of Molecular Biology* **224**, 771–782.
- FIELD, D. J. (1999). Wavelets, vision and the statistics of natural scenes. *Philosophical transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* **357**, 2527–2542. doi: 10.1098/rsta.1999.0446.
- FIGUEIREDO, A. M., WHITTAKER, S. B.-M., KNOWLING, S. E., RADFORD, S. E. & MOORE, G. R. (2013). Conformational dynamics is more important than helical propensity for the folding of the all alpha-helical protein IM7. *Protein Science* doi: 10.1002/pro.2372.
- FOIT, L., MORGAN, G. J., KERN, M. J., STEIMER, L. R., VON HACHT, A. A., TITCHMARSH, J., WARRINER, S. L., RADFORD, S. E. & BARDWELL, J. C. A. (2009). Optimizing protein stability *in vivo*. *Molecular Cell* **36**, 861–871. doi: 10.1016/j.molcel.2009.11.022.
- FRAUENFELDER, H. (1987). Function and dynamics of myoglobin. *Annals of the New York Academy of Sciences* **504**, 151–167. ISSN: 1749-6632. doi: 10.1111/J.1749-6632.1987.tb48730.x.
- FRAUENFELDER, H., CHEN, G., BERENDZEN, J., FENIMORE, P. W., JANSSON, H., MCMAHON, B. H., STROE, I. R., SWENSON, J. & YOUNG, R. D. (2009). A unified model of protein dynamics. *Proceedings of the National Academy of Sciences of the United States of America* **106**, 5129–5134. doi: 10.1073/pnas.0900336106.
- FRAUENFELDER, H., MCMAHON, B. H. & FENIMORE, P. W. (2003). Myoglobin: the hydrogen atom of biology and a paradigm of complexity. *Proceedings of the National*

- Academy of Sciences of the United States of America* **100**, 8615–8617. doi: 10.1073/pnas.1633688100.
- FRAUENFELDER, H., SLIGAR, S. G. & WOLYNES, P. G. (1991). The energy landscapes and motions of proteins. *Science* **254**, 1598–1603.
- FRAUENFELDER, H. & WOLYNES, P. G. (1985). Rate theories and puzzles of hemeprotein kinetics. *Science* **299**, 337–345.
- FRIEDRICH, M. S. & WOLYNES, P. G. (1989). Toward protein tertiary structure recognition by means of associative memory Hamiltonians. *Science* **246**, 371–373.
- FRIEL, C. T., SMITH, D. A., VENDRUSCOLO, M., GSPONER, J. & RADFORD, S. E. (2009). The mechanism of folding of Im7 reveals competition between functional and kinetic evolutionary constraints. *Nature Structural & Molecular Biology* **16**, 318–324. doi: 10.1038/nsmb.1562.
- FUGLESTAD, B., GASPER, P. M., MCCAMMON, J. A., MARKWICK, P. R. L. & KOMIVES, E. A. (2013). Correlated motions and residual frustration in thrombin. *Journal of Physical Chemistry B* **117**, 12857–12863. doi: 10.1021/jp402107u.
- GALZITSKAYA, O. V. & FINKELSTEIN, A. V. (1999). A theoretical search for folding/unfolding nuclei in three-dimensional protein structures. *Proceedings of the National Academy of Sciences of the United States of America* **96**, 11299–11304.
- GAMBIN, Y., SCHUG, A., LEMKE, E. A., LAVINDER, J. J., FERREON, A. C. M., MAGLIERY, T. J., ONUCHIC, J. N. & DENIZ, A. A. (2009). Direct single-molecule observation of a protein living in two opposed native structures. *Proceedings of the National Academy of Sciences of the United States of America* **106**, 10153–10158. doi: 10.1073/pnas.0904461106.
- GASSNER, N. C., BAASE, W. A., LINDSTROM, J. D., LU, J., DAHLQUIST, F. W., & MATTHEWS, B. W. (1999). Methionine and alanine substitutions show that the formation of wild-type-like structure in the carboxy-terminal domain of T4 lysozyme is a rate-limiting step in folding. *Biochemistry* **38**, 14451–14460.
- GERLAND, U., MOROZ, J. D. & HWA, T. (2002). Physical constraints and functional characteristics of transcription factor-DNA interaction. *Proceedings of the National Academy of Sciences of the United States of America* **99**, 12015–12020. doi: 10.1073/pnas.192693599.
- GILBERT, W. (1987). The exon theory of genes. In *Cold Spring Harbor Symposia on Quantitative Biology*, vol. **52**. New York, NY: Cold Spring Harbor Laboratory Press, pp. 901–905.
- GO, N. (1983). Theoretical studies of protein folding. *Annual Review of Biophysics and Bioengineering* **12**, 183–210. doi: 10.1146/annurev.bb.12.060183.001151.
- GOLDSTEIN, R. A., LUTHEY-SCHULTEN, Z. A. & WOLYNES, P. G. (1992a). Protein tertiary structure recognition using optimized Hamiltonians with local interactions. *Proceedings of the National Academy of Sciences of the United States of America* **89**, 9029–9033.
- GOLDSTEIN, R. A., LUTHEY-SCHULTEN, Z. A. & WOLYNES, P. G. (1992b). Optimal protein-folding codes from spin-glass theory. *Proceedings of the National Academy of Sciences of the United States of America* **89**, 4918–4922.
- GOSAVI, S., CHAVEZ, L. L., JENNINGS, P. A. & ONUCHIC, J. N. (2006). Topological frustration and the folding of interleukin-1 beta. *Journal of Molecular Biology* **357**, 986–996. doi: 10.1016/j.jmb.2005.11.074.
- GOSAVI, S., WHITFORD, P. C., JENNINGS, P. A. & ONUCHIC, J. N. (2008). Extracting function from a beta-trefoil folding motif. *Proceedings of the National Academy of Sciences of the United States of America* **105**, 10384–10389. doi: 10.1073/pnas.0801343105.
- GREET, R. J. & TURNBULL, D. (1967). Glass transition in o-terphenyl. *Journal of Chemical Physics* **46**, 1243.
- GULUKOTA, K. & WOLYNES, P. G. (1994). Statistical mechanics of kinetic proofreading in protein folding *in vivo*. *Proceedings of the National Academy of Sciences of the United States of America* **91**, 9292–9296.
- HAGAI, T., AZIA, A., TRIZAC, E. & LEVY, Y. (2012). Modulation of folding kinetics of repeat proteins: interplay between intra- and interdomain interactions. *Biophysical Journal* **103**, 1555–1565. doi: 10.1016/j.bpj.2012.08.018.
- HAN, J.-H., BATEY, S., NICKSON, A. A., TEICHMANN, S. A. & CLARKE, J. (2007). The folding and evolution of multi-domain proteins. *Nature Reviews Molecular Cell Biology* **8**, 319–330. doi: 10.1038/nrm2144.
- HEGLER, J. A., WEINKAM, P. & WOLYNES, P. G. (2008). The spectrum of biomolecular states and motions. *HFSP Journal* **2**, 307–313.
- HENRY, E. R., BEST, R. B. & EATON, W. A. (2013). Comparing a simple theoretical model for protein folding with all-atom molecular dynamics simulations. *Proceedings of the National Academy of Sciences of the United States of America* **110**, 17880–17885. doi: 10.1073/pnas.1317105110.
- HENZLER-WILDMAN, K. A., THAI, V., LEI, M., OTT, M., WOLF-WATZ, M., FENN, T., POZHARSKI, E., WILSON, M. A., PETSKO, G. A., KARPLUS, M., HÜBNER, C. G. & KERN, D. (2007). Intrinsic motions along an enzymatic reaction trajectory. *Nature* **450**, 838–844. doi: 10.1038/nature06410.
- HEYTESBURY, W. (1494). Regule solvendi sophismata. *Bonetus Locatellus*.
- HOPFIELD, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the United States of America* **79**, 2554–2558.
- HUA, Q.-X., NAKAGAWA, S. H., JIA, W., HUANG, K., PHILLIPS, N. B., HU, S.-Q. & WEISS, M. A. (2008). Design of an active ultrastable single-chain insulin analog: synthesis, structure, and therapeutic implications. *Journal of Biological Chemistry* **283**, 14703–14716. doi: 10.1074/jbc.M800313200.

- HUNTINGTON, J. A., READ, R. J. & CARRELL, R. W. (2000). Structure of a serpin–protease complex shows inhibition by deformation. *Nature* **407**, 923–926. doi: 10.1038/35038119.
- IM, H., AHN, H. Y. & YU, M. H. (2000). Bypassing the kinetic trap of serpin protein folding by loop extension. *Protein Science* **9**, 1497–1502. doi: 10.1110/ps.9.8.1497.
- ITZHAKE, L. S. & LOWE, A. R. (2012). From artificial antibodies to nanosprings: the biophysical properties of repeat proteins. *Advances in Experimental Medicine and Biology* **747**, 153–166. doi: 10.1007/978-1-4614-3229-6_10.
- JANE DYSON, H. & WRIGHT, P. E. (2002). Insights into the structure and dynamics of unfolded proteins from nuclear magnetic resonance. *Advances in Protein Chemistry* **62**, 311–340.
- JAVADI, Y. & MAIN, E. R. G. (2009). Exploring the folding energy landscape of a series of designed consensus tetratricopeptide repeat proteins. *Proceedings of the National Academy of Sciences of the United States of America* **106**, 17383–17388. doi: 10.1073/pnas.0907455106.
- JENIK, M., PARRA, R. G., RADUSKY, L. G., TRUCHANSKI, A., WOLYNES, P. G. & FERREIRO, D. U. (2012). Protein frustratometer: a tool to localize energetic frustration in protein molecules. *Nucleic Acids Research* **40**(Web Server issue), W348–W351. doi: 10.1093/nar/gks447.
- JIN, W., KAMBARA, O., SASAKAWA, H., TAMURA, A. & TAKADA, S. (2003). *De novo* design of foldable proteins with smooth folding funnel: automated negative design and experimental verification. *Structure* **11**, 581–590.
- KAJAVA, A. V. (2012). Tandem repeats in proteins: from sequence to structure. *Journal of Structural Biology* **179**, 279–288. doi: 10.1016/j.jsb.2011.08.009.
- KAMTEKAR, S., SCHIFFER, J. M., XIONG, H., BABIK, J. M. & HECHT, M. H. (1993). Protein design by binary patterning of polar and nonpolar amino acids. *Science* **262**, 1680–1685.
- KARPLUS, M. & MCCAMMON, J. A. (1983). Dynamics of proteins: elements and function. *Annual Review in Biochemistry* **52**, 263–300. doi: 10.1146/annurev.bi.52.070183.001403.
- KATO, A., YAMADA, M., NAKAMURA, S., KIDOKORO, S.-I. & KURODA, Y. (2007). Thermodynamic properties of BPTI variants with highly simplified amino acid sequences. *Journal of Molecular Biology* **372**, 737–746. doi: 10.1016/j.jmb.2007.06.066.
- KAUFFMAN, S. (2013). Beyond reductionism twice: no laws entail biosphere evolution, formal cause laws beyond efficient cause laws. *arXiv:1303.5684*.
- KAUZMANN, W. (1959). Forces responsible for maintaining the native configurations of proteins. *Advances in Protein Chemistry* **14**, 33–35.
- KEEFE, A. D. & SZOSTAK, J. W. (2001). Functional proteins from a random-sequence library. *Nature* **410**, 715–718. doi: 10.1038/35070613.
- KLIMOV, D. K. & THIRUMALAI, D. (2002). Stiffness of the distal loop restricts the structural heterogeneity of the transition state ensemble in SH3 domains. *Journal of Molecular Biology* **317**, 721–737. doi: 10.1006/jmbi.2002.5453.
- KÖRETKE, K. K., LUTHEY-SCHULTEN, Z. & WOLYNES, P. G. (1996). Self-consistently optimized statistical mechanical energy functions for sequence structure alignment. *Protein Science* **5**, 1043–1059.
- KRAMER, M. A., WETZEL, S. K., PLÜCKTHUN, A., MITTL, P. R. E. & GRÜTTER, M. G. (2010). Structural determinants for improved stability of designed ankyrin repeat proteins with a redesigned C-capping module. *Journal of Molecular Biology* **404**, 381–391. doi: 10.1016/j.jmb.2010.09.023.
- KRETMANN, N. (1982). Syncategoremata, exponibilia, sophismata. In *The Cambridge History of Later Medieval Philosophy* (eds. N. KRETMANN, A. KENNY & J. PINBORG), pp. 211–245. Cambridge: Cambridge University Press.
- KUHLMAN, B., DANTAS, G., IRETON, G. C., VARANI, G., STODDARD, B. L. & BAKER, D. (2003). Design of a novel globular protein fold with atomic-level accuracy. *Science* **302**, 1364–1368. doi: 10.1126/Science.1089427.
- KURIYAN, J. & EISENBERG, D. (2007). The origin of protein interactions and allostery in colocalization. *Nature* **450**, 983–990. doi: 10.1038/nature06524.
- KURODA, Y. & KIM, P. S. (2000). Folding of bovine pancreatic trypsin inhibitor (BPTI) variants in which almost half the residues are alanine. *Journal of Molecular Biology* **298**, 493–501. doi: 10.1006/jmbi.2000.3622.
- LAU, A. W. C., HOFFMAN, B. D., DAVIES, A., CROCKER, J. C. & LUBENSKY, T. C. (2003). Microrheology, stress fluctuations, and active behavior of living cells. *Physics Review Letters* **91**, 198101.
- LEFFLER, J. E. (1953). Parameters for the description of transition states. *Science* **117**, 340–341. doi: 10.1126/science.117.3039.340.
- LEVINTHAL, C. (1968). Are there pathways for protein folding. *Journal de Chimie Physique* **65**, 44–45.
- LEVY, Y., CHO, S. S., SHEN, T., ONUCHIC, J. N. & WOLYNES, P. G. (2005). Symmetry and frustration in protein energy landscapes: a near degeneracy resolves the Rop dimer-folding mystery. *Proceedings of the National Academy of Sciences of the United States of America* **102**, 2373–2378.
- LEVY, Y., ONUCHIC, J. N. & WOLYNES, P. G. (2007). Fly-casting in protein–DNA binding: frustration between protein folding and electrostatics facilitates target recognition. *Journal of the American Chemical Society* **129**, 738–739. doi: 10.1021/ja065531n.
- LEVY, Y., WOLYNES, P. G. & ONUCHIC, J. N. (2004). Protein topology determines binding mechanism. *Proceedings of the National Academy of Sciences of the United States of America* **101**, 511–516. doi: 10.1073/pnas.2534828100.

- LITTLE, W. A. (1974). The existence of persistent states in the brain. *Mathematical Biosciences* **19**, 101–120.
- LIU, R., BAASE, W. A. & MATTHEWS, B. W. (2000). The introduction of strain and its effects on the structure and stability of T4 lysozyme. *Journal of Molecular Biology* **295**, 127–145. doi: 10.1006/jmbi.1999.3300.
- LUM, K., CHANDLER, D. & WEEKS, J. D. (1999). Hydrophobicity at small and large length scales. *Journal of Physical Chemistry B* **103**, 4570–4577.
- LUO, M. (2012). Influenza virus entry. *Advances in Experimental Medicine and Biology* **726**, 201–221. doi: 10.1007/978-1-4614-0980-9_9.
- LUTHEY-SCHULTEN, Z., RAMIREZ, B. E. & WOLYNES, P. G. (1995). Helix-coil, liquid crystal, and spin glass transitions of a collapsed heteropolymer. *Journal of Physical Chemistry* **99**, 2177–2185.
- MA, H., WAN, C., WU, A. & ZEWAIL, A. H. (2007). DNA folding and melting observed in real time redefine the energy landscape. *Proceedings of the National Academy of Sciences of the United States of America* **104**, 712–716. doi: 10.1073/pnas.0610028104.
- MARCOVITZ, A. & LEVY, Y. (2011). Frustration in protein–DNA binding influences conformational switching and target search kinetics. *Proceedings of the National Academy of Sciences of the United States of America* **108**, 17957–17962.
- MARCOVITZ, A. & LEVY, Y. (2013). Weak frustration regulates sliding and binding kinetics on rugged protein–DNA landscapes. *Journal of Physical Chemistry B* **112**, 13005–13014. doi: 10.1021/jp402296d.
- MAREK, M. S., JOHNSON-BUCK, A. & WALTER, N. G. (2011). The shape-shifting quasispecies of RNA: one sequence, many functional folds. *Physical Chemistry Chemical Physics* **13**, 11524–11537. doi: 10.1030/c1cp20576e.
- MATOUSCHEK, A., KELLIS, J. T. Jr., SERRANO, L. & FERSHT, A. R. (1989). Mapping the transition state and pathway of protein folding by protein engineering. *Nature* **340**, 122–126. doi: 10.1038/340122a0.
- MATTHEWS, C. R. & HURLE, M. R. (1987). Mutant sequences as probes of protein folding mechanisms. *Bioessays* **6**, 254–257. doi: 10.1002/bies.950060603.
- MAY, A. C. (1999). Towards more meaningful hierarchical classification of amino acid scoring matrices. *Protein Engineering* **12**, 707–712.
- MEIERING, E. M., SERRANO, L. & FERSHT, A. R. (1992). Effect of active site residues in barnase on activity and stability. *Journal of Molecular Biology* **225**, 585–589.
- MÉLIN, R., LI, H., WINGREEN, N. S. & TANG, C. (1999). Designability, thermodynamic stability, and dynamics in protein folding: a lattice model study. *Journal of Chemical Physics* **110**, 1252.
- MELLO, C. C. & BARRICK, D. (2004). An experimentally determined protein folding energy landscape. *Proceedings of the National Academy of Sciences of the United States of America* **101**, 14102–14107. doi: 10.1073/pnas.0403386101.
- MÉZARD, M. & MORA, T. (2009). Constraint satisfaction problems and neural networks: a statistical physics perspective. *Journal of Physiology – Paris* **103**, 107–113.
- MIYASHITA, O., ONUCHIC, J. N. & WOLYNES, P. G. (2003). Nonlinear elasticity, proteinquakes, and the energy landscapes of functional transitions in proteins. *Proceedings of the National Academy of Sciences of the United States of America* **100**, 12570–12575. doi: 10.1073/pnas.2135471100.
- MONASSON, R., ZECCHINA, R., KIRKPATRICK, S., SELMAN, B. & TROYANSKY, L. (1999). Determining computational complexity from characteristic ‘phase transitions’. *Nature* **400**, 133–137.
- MONOD, J. (1973). Le hasard et la nécessité: essai sur la philosophie naturelle de la biologie moderne. Éditions du Seuil.
- MUÑOZ, V. & EATON, W. A. (1999). A simple model for calculating the kinetics of protein folding from three-dimensional structures. *Proceedings of the National Academy of Sciences of the United States of America* **96**, 11311–11316.
- MUNSON, M., ANDERSON, K. S. & REGAN, L. (1997). Speeding up protein folding: mutations that increase the rate at which Rop folds and unfolds by over four orders of magnitude. *Folding and Designing* **2**, 77–87.
- MURPHY, L. R., WALLQVIST, A. & LEVY, R. M. (2000). Simplified amino acid alphabets for protein fold recognition and implications for folding. *Protein Engineering* **13**, 149–152.
- NAGANATHAN, A. N. & MUÑOZ, V. (2010). Insights into protein folding mechanisms from large scale analysis of mutational effects. *Proceedings of the National Academy of Sciences of the United States of America* **107**, 8611–8616. doi: 10.1073/pnas.1000988107.
- NÉEL, L. (1970). Magnetism and the local molecular field. In *Nobel Lectures, vol. Physics 1963–1970* (ed. G. EKSPONG), pp. 318–341. Amsterdam: Elsevier.
- NOEL, J. K., WHITFORD, P. C. & ONUCHIC, J. N. (2012). The shadow map: a general contact definition for capturing the dynamics of biomolecular folding and function. *Journal of Physical Chemistry B* **116**, 8692–8702. doi: 10.1021/jp300852d.
- NOEL, J. K., WHITFORD, P. C., SANBONMATSU, K. Y. & ONUCHIC, J. N. (2010). SMOG@ctbp: simplified deployment of structure-based models in GROMACS. *Nucleic Acids Research* **38**(Web Server issue), W657–W661. doi: 10.1093/nar/gkq498.
- NORAMBUENA, T. & MELO, F. (2010). The protein–DNA Interface database. *BMC Bioinformatics* **11**, 262. doi: 10.1186/1471-2105-11-262.
- NORDBLAD, P. (2013). Competing interaction in magnets: the root of ordered disorder or only frustration? *Physica Scripta* **88**, 058301.
- OLIVEBERG, M. & WOLYNES, P. G. (2005). The experimental survey of protein-folding energy landscapes. *Quarterly Reviews of Biophysics* **38**, 245–288.

- OLSSON, M. H. M., PARSON, W. W. & WARSHIEL, A. (2006). Dynamical contributions to enzyme catalysis: critical tests of a popular hypothesis. *Chemical Reviews* **106**, 1737–1756. doi: 10.1021/cr040427e.
- ONUCHIC, J. N., LUTHEY-SCHULTEN, Z. & WOLYNES, P. G. (1997). Theory of protein folding: the energy landscape perspective. *Annual Review of Physical Chemistry* **48**, 545–600. doi: 10.1146/annrev.physchem.48.1.545.
- ONUCHIC, J. N. & WOLYNES, P. G. (2004). Theory of protein folding. *Currrent Opinion in Structural Biology* **14**, 70–75. doi: 10.1016/j.sbi.2004.01.009.
- ONUCHIC, J. N., WOLYNES, P. G., LUTHEY-SCHULTEN, Z. & SOCCI, N. D. (1995). Toward an outline of the topography of a realistic protein-folding funnel. *Proceedings of the National Academy of Sciences of the United States of America* **92**, 3626–3630.
- PANCHENKO, A. R., LUTHEY-SCHULTEN, Z. & WOLYNES, P. G. (1996). Foldons, protein structural modules, and exons. *Proceedings of the National Academy of Sciences of the United States of America* **93**, 2008–2013.
- PAPOIAN, G. A., ULANDER, J., EASTWOOD, M. P., LUTHEY-SCHULTEN, Z. & WOLYNES, P. G. (2004). Water in protein structure prediction. *Proceedings of the National Academy of Sciences of the United States of America* **101**, 3352–3357. doi: 10.1073/pnas.0307851100.
- PAPOIAN, G. A., ULANDER, J. & WOLYNES, P. G. (2003a). Role of water mediated interactions in protein-protein recognition landscapes. *Journal of the American Chemical Society* **125**, 9170–9178. doi: 10.1021/ja034729u.
- PAPOIAN, G. A. & WOLYNES, P. G. (2003b). The physics and bioinformatics of binding and folding—an energy landscape perspective. *Biopolymers* **68**, 333–349. doi: 10.1002/bip.10286.
- PARRA, R. G., ESPADA, R., SÁNCHEZ, I. E., SIPPL, M. J. & FERREIRO, D. U. (2013). Detecting repetitions and periodicities in proteins by tiling the structural space. *Journal of Physical Chemistry B* **117**, 12887–12897. doi: 10.1021/jp402105j.
- PERVUSHIN, K., VAMVACA, K., VÖGEL, B. & HILVERT, D. (2007). Structure and dynamics of a molten globular enzyme. *Nature Structural and Molecular Biology* **14**, 1202–1206.
- PIANA, S., LINDORFF-LARSEN, K. & SHAW, D. E. (2011). How robust are protein folding simulations with respect to force field parameterization? *Biophysical Journal* **100**, L47–L49. doi: 10.1016/j.bpj.2011.03.051.
- PIANA, S., LINDORFF-LARSEN, K. & SHAW, D. E. (2012). Protein folding kinetics and thermodynamics from atomistic simulation. *Proceedings of the National Academy of Sciences of the United States of America* **109**, 17845–17850. doi: 10.1073/pnas.1201811109.
- PLAXCO, K. W., SIMONS, K. T. & BAKER, D. (1998). Contact order, transition state placement and the refolding rates of single domain proteins. *Journal of Molecular Biology* **277**, 985–994. doi: 10.1006/jmbi.1998.1645.
- PLOTKIN, S. S. & ONUCHIC, J. N. (2002). Understanding protein folding with energy landscape theory. Part I: basic concepts. *Quarterly Reviews of Biophysics* **35**, 111–167.
- PLOTKIN, S. S., WANG, J. & WOLYNES, P. G. (1996). Correlated energy landscape model for finite, random heteropolymers. *Physical Review E* **53**, 6271.
- PLOTKIN, S. S., WANG, J. & WOLYNES, P. G. (1997). Statistical mechanics of a correlated energy landscape model for protein folding funnels. *Journal of Chemical Physics* **106**, 2932–2948. doi: http://dx.doi.org/10.1063/1.473355.
- PORTER, C. T., BARTLETT, G. J. & THORNTON, J. M. (2004). The Catalytic Site Atlas: a resource of catalytic sites and residues identified in enzymes using structural data. *Nucleic Acids Research* **32**(Database issue), D129–D133. doi: 10.1093/nar/gkh028.
- PORTMAN, J. J., TAKADA, S. & WOLYNES, P. G. (1998). Variational theory for site resolved protein folding free energy surfaces. *Physical Review Letters* **81**, 5237.
- PORTMAN, J. J., TAKADA, S. & WOLYNES, P. G. (2001a). Microscopic theory of protein folding rates. I. Fine structure of the free energy profile and folding routes from a variational approach. *Journal of Chemical Physics* **114**, 5237–5240.
- PORTMAN, J. J., TAKADA, S. & WOLYNES, P. G. (2001b). Microscopic theory of protein folding rates. II. Local reaction coordinates and chain dynamics. *Journal of Chemical Physics* **114**, 5082.
- RECHES, M., SNYDER, P. W. & WHITESIDES, G. M. (2009). Folding of electrostatically charged beads-on-a-string as an experimental realization of a theoretical model in polymer science. *Proceedings of the National Academy of Sciences of the United States of America* **106**, 17644–17649. doi: 10.1073/pnas.0905533106.
- RIDDLE, D. S., SANTIAGO, J. V., BRAY-HALL, S. T., DOSHI, N., GRANTCHAROVA, V. P., YI, Q. & BAKER, D. (1997). Functional rapidly folding proteins from simplified amino acid sequences. *Nature Structural Biology* **4**, 805–809.
- ROTHEMUND, P. W. K. (2006). Folding DNA to create nanoscale shapes and patterns. *Nature* **440**, 297–302. doi: 10.1038/nature04586.
- SALI, A., SHAKHNOVICH, E. & KARPLUS, M. (1994). Kinetics of protein folding. A lattice model study of the requirements for folding to the native state. *Journal of Molecular Biology* **235**, 1614–1636.
- SÁNCHEZ, I. E. & KIEFHABER, T. (2003a). Evidence for sequential barriers and obligatory intermediates in apparent two-state protein folding. *Journal of Molecular Biology* **325**, 367–376.
- SÁNCHEZ, I. E. & KIEFHABER, T. (2003b). Non-linear rate-equilibrium free energy relationships and Hammond behavior in protein folding. *Biophysical Chemistry* **100**, 397–407.
- SÁNCHEZ, I. E., FERREIRO, D. U., DELLAROLE, M. & DE PRAT-GAY, G. (2010). Experimental snapshots of a

- protein–DNA binding landscape. *Proceedings of the National Academy of Sciences of the United States of America* **107**, 7751–7756. doi: 10.1073/pnas.0911734107.
- SÁNCHEZ, I. E., FERREIRO, D. U. & DE PRAT GAY, G. (2011). Mutational analysis of kinetic partitioning in protein folding and protein–DNA binding. *Protein Engineering Design and Selection* **24**, 179–184. doi: 10.1093/protein/gzq064.
- SCALLEY-KIM, M. & BAKER, D. (2004). Characterization of the folding energy landscapes of computer generated proteins suggests high folding free energy barriers and cooperativity may be consequences of natural selection. *Journal of Molecular Biology* **338**, 573–583. doi: 10.1016/j.jmb.2004.02.055.
- SCHAFFER, N. P., HOFFMAN, R. M. B., BURGER, A., CRAIG, P. O., KOMIVES, E. A. & WOLYNES, P. G. (2012). Discrete kinetic models from funneled energy landscape simulations. *PLoS ONE* **7**, e50635. doi: 10.1371/journal.pone.0050635.
- SCHIFFER, M. & EDMUNDSON, A. B. (1967). Use of helical wheels to represent the structures of proteins and to identify segments with helical potential. *Biophysical Journal* **7**, 121–135.
- SCHIFFER, P. & RAMIREZ, A. P. (1996). Recent experimental progress in the study of geometrical magnetic frustration. *Comments on Condensed Matter Physics* **18**, 21–50.
- SCHILDKRAUT, C. & LIFSON, S. (1965). Dependence of the melting temperature of DNA on salt concentration. *Biopolymers* **3**, 195–208. doi: 10.1002/bip.360030207.
- SERGANOV, A. & NÜDLER, E. (2013). A decade of riboswitches. *Cell* **152**, 17–24. doi: 10.1016/j.cell.2012.12.024.
- SHAKHNOVICH, E. & GUTIN, A. (1990). Enumeration of all compact conformations of copolymers with random sequence of links. *Journal of Chemical Physics* **93**, 5967.
- SHEN, T., HOFMANN, C. P., OLIVEBERG, M. & WOLYNES, P. G. (2005). Scanning malleable transition state ensembles: comparing theory and experiment for folding protein U1A. *Biochemistry* **44**, 6433–6439.
- SHOEMAKER, B. A., PORTMAN, J. J. & WOLYNES, P. G. (2000). Speeding molecular recognition by using the folding funnel: the fly-casting mechanism. *Proceedings of the National Academy of Sciences of the United States of America* **97**, 8868–8873. doi: 10.1073/pnas.160259697.
- SHOEMAKER, B. A., WANG, J. & WOLYNES, P. G. (1999). Exploring structures in protein folding funnels with free energy functionals: the transition state ensemble. *Journal of Molecular Biology* **287**, 675–694. doi: 10.1006/jmbi.1999.2613.
- SHOICHET, B. K., BAASE, W. A., KUROKI, R. & MATTHEWS, B. W. (1995). A relationship between protein stability and protein function. *Proceedings of the National Academy of Sciences of the United States of America* **92**, 452–456.
- SIMONDON, G. (2005). L'individuation: à la lumière des notions de forme et d'information. Editions Jérôme Millon.
- SIPPL, M. J. & WIEDERSTEIN, M. (2012). Detection of spatial correlations in protein structures and molecular complexes. *Structure* **20**, 718–728. doi: 10.1016/j.str.2012.01.024.
- SLUTSKY, M. & MIRNY, L. A. (2004). Kinetics of protein–DNA interaction: facilitated target location in sequence-dependent potential. *Biophysical Journal* **87**, 4021–4035. doi: 10.1529/biophysj.104.050765.
- SMITH, B. A. & HECHT, M. H. (2011). Novel proteins: from fold to function. *Current Opinion in Chemical Biology* **15**, 421–426.
- SOLOMATIN, S. V., GREENFELD, M., CHU, S. & HERSCHLAG, D. (2010). Multiple native states reveal persistent ruggedness of an RNA folding landscape. *Nature* **463**, 681–684. doi: 10.1038/nature08717.
- SPOLAR, R. S. & RECORD, M. T. Jr (1994). Coupling of local folding to site-specific binding of proteins to DNA. *Science* **263**, 777–784.
- STEVENSON, J. D. & WOLYNES, P. G. (2010). The ultimate fate of supercooled liquids. *Journal of Physical Chemistry A* **115**, 3713–3719.
- STREET, T. O. & BARRICK, D. (2009). Predicting repeat protein folding kinetics from an experimentally determined folding energy landscape. *Protein Science* **18**, 58–68. doi: 10.1002/pro.9.
- SULKOWSKA, J. I., NOEL, J. K. & ONUCHIC, J. N. (2012). Energy landscape of knotted protein folding. *Proceedings of the National Academy of Sciences of the United States of America* **109**, 17783–17788.
- SUTTO, L., LÄTZER, J., HEGLER, J. A., FERREIRO, D. U. & WOLYNES, P. G. (2007). Consequences of localized frustration for the folding mechanism of the IM7 protein. *Proceedings of the National Academy of Sciences of the United States of America* **104**, 19825–19830. doi: 10.1073/pnas.0709922104.
- TAMASKOVIC, R., SIMON, M., STEFAN, N., SCHWILL, M. & PLÜCKTHUN, A. (2012). Designed ankyrin repeat proteins (DARPs) from research to therapy. *Methods in Enzymology* **503**, 101–134. doi: 10.1016/B978-0-12-396962-0.00005-7.
- TODD, M. J., LORIMER, G. H. & THIRUMALAI, D. (1996). Chaperonin-facilitated protein folding: optimization of rate and yield by an iterative annealing mechanism. *Proceedings of the National Academy of Sciences of the United States of America* **93**, 4030–4035.
- TOKURIKI, N., JACKSON, C. J., AFRIAT-JURNOU, L., WYGANOWSKI, K. T., TANG, R. & TAWFIK, D. S. (2012). Diminishing returns and tradeoffs constrain the laboratory optimization of an enzyme. *Nature Communications* **3**, 1257. doi: 10.1038/ncomms2246.
- TOKURIKI, N., STRICHER, F., SERRANO, L. & TAWFIK, D. S. (2008). How protein stability and new functions trade

- off. *PLoS Computational Biology* **4**, e1000002. doi: 10.1371/journal.pcbi.1000002.
- TOKURIKI, N. & TAWFIK, D. S. (2009). Stability effects of mutations and protein evolvability. *Current Opinion in Structural Biology* **19**, 596–604. doi: 10.1016/j.sbi.2009.08.003.
- TRIPP, K. W. & BARRICK, D. (2008). Rerouting the folding pathway of the Notch ankyrin domain by reshaping the energy landscape. *Journal of the American Chemical Society* **130**, 5681–5688. doi: 10.1021/ja0763201.
- TRUONG, H. H., KIM, B. L., SCHAFER, N. P. & WOLYNES, P. G. (2013). Funneling and frustration in the energy landscapes of some designed and simplified proteins. *Journal of Chemical Physics* **139**, 121908.
- TSUTSUI, Y., KURI, B., SENGUPTA, T. & WINTRODE, P. L. (2008). The structural basis of serpin polymerization studied by hydrogen/deuterium exchange and mass spectrometry. *Journal of Biological Chemistry* **283**, 30804–30811. doi: 10.1074/jbc.M804048200.
- VANNIMENUS, J. & TOULOUSE, G. (1977). Theory of the frustration effect. II. Ising spins on a square lattice. *Journal of Physics C: Solid State Physics* **10**, L537. URL: <http://stacks.iop.org/0022-3719/10/i=18/a=008>.
- VIASNOFF, V., MELLER, A. & ISAMBERT, H. (2006). DNA nanomechanical switches under folding kinetics control. *Nano Letters* **6**, 101–104. doi: 10.1021/nl05216c.
- VON HIPPEL, P. H. & BERG, O. G. (1989). Facilitated target location in biological system. *Journal of Biological Chemistry* **264**, 675–678.
- WALES, D. J. (1998). Symmetry, near-symmetry and energetics. *Chemical Physics Letters* **285**, 330–336.
- WANG, S. & WOLYNES, P. G. (2011). On the spontaneous collective motion of active matter. *Proceedings of the National Academy of Sciences of the United States of America* **108**, 15184–15189.
- WANNIER, G. H. (1950a). Antiferromagnetism. The Triangular Ising Net. *Physical Review* **79**, 357–364. doi: 10.1103/PhysRev.79.357. URL: <http://link.aps.org/doi/10.1103/PhysRev.79.357>
- WATSON, J. D. & CRICK, F. H. C. (1953). Molecular structure of nucleic acids. *Nature* **171**, 737–738.
- WEINKAM, P., ROMESBERG, F. E. & WOLYNES, P. G. (2009). Chemical frustration in the protein folding landscape: grand canonical ensemble simulations of cytochrome c. *Biochemistry* **48**, 2394–2402. doi: 10.1021/bi802293m.
- WEINKAM, P. & WOLYNES, P. G. (2010a). The folding landscapes of metalloproteins; in protein folding and metal ions: mechanisms, biology and disease. In *Protein Folding and Metal Ions: Mechanisms, Biology and Disease* (eds. P. WITTUNG-STAFSHED & M. C. GOMES) pp. 247–273. Boca Raton, FL: Taylor and Francis.
- WEINKAM, P., ZIMMERMANN, J., ROMESBERG, F. E. & WOLYNES, P. G. (2010b). The folding energy landscape and free energy excitations of cytochrome c. *Accounts of Chemical Research* **43**, 652–660. doi: 10.1021/ar9002703.
- WEINKAM, P., ZIMMERMANN, J., SAGLE, L. B., MATSUDA, S., DAWSON, P. E., WOLYNES, P. G. & ROMESBERG, F. E. (2008). Characterization of alkaline transitions in ferricytochrome c using carbon–deuterium infrared probes. *Biochemistry* **47**, 13470–13480. doi: 10.1021/bi801223n.
- WEINKAM, P., ZONG, C. & WOLYNES, P. G. (2005). A funneled energy landscape for cytochrome c directly predicts the sequential folding route inferred from hydrogen exchange experiments. *Proceedings of the National Academy of Sciences of the United States of America* **102**, 12401–12406. doi: 10.1073/pnas.0505274102.
- WEISS, M. A. (2013). Diabetes mellitus due to the toxic misfolding of proinsulin variants. *FEBS Letters* **587**, 1942–1950. doi: 10.1016/j.febslet.2013.04.044.
- WEISS, O., JIMENEZ-MONTANO, M. A. & HERZEL, H. (2000). Information content of protein sequences. *Journal of Theoretical Biology* **206**, 379–386.
- WEISSMAN, J. S. & KIM, P. S. (1991). Reexamination of the folding of BPTI: predominance of native intermediates. *Science* **253**, 1386–1393.
- WEISSMAN, J. S. & KIM, P. S. (1992). Kinetic role of non-native species in the folding of bovine pancreatic trypsin inhibitor. *Proceedings of the National Academy of Sciences of the United States of America* **89**, 9900–9904.
- WENSLEY, B. G., BATEY, S., BONE, F. A. C., CHAN, Z. M., TUMELTY, N. R., STEWARD, A., KWA, L. G., BORGIA, A. & CLARKE, J. (2010). Experimental evidence for a frustrated energy landscape in a three-helix-bundle protein family. *Nature* **463**, 685–688. doi: 10.1038/nature08743.
- WENSLEY, B. G., KWA, L. G., SHAMMAS, S. L., ROGERS, J. M., BROWNING, S., YANG, Z. & CLARKE, J. (2012). Separating the effects of internal friction and transition state energy to explain the slow, frustrated folding of spectrin domains. *Proceedings of the National Academy of Sciences of the United States of America* **109**, 17795–17799. doi: 10.1073/pnas.1201793109.
- WERBECK, N. D. & ITZHAKI, L. S. (2007). Probing a moving target with a plastic unfolding intermediate of an ankyrin-repeat protein. *Proceedings of the National Academy of Sciences of the United States of America* **104**, 7863–7868. doi: 10.1073/pnas.0610315104.
- WERBECK, N. D., ROWLING, P. J. E., CHELLAMUTHU, V. R. & ITZHAKI, L. S. (2008). Shifting transition states in the unfolding of a large ankyrin repeat protein. *Proceedings of the National Academy of Sciences of the United States of America* **105**, 9982–9987. doi: 10.1073/pnas.0705300105.
- WHITFORD, P. C., ONUCHIC, J. N. & WOLYNES, P. G. (2008). Energy landscape along an enzymatic reaction trajectory: hinges or cracks? *HFSP Journal* **2**, 61–64. doi: 10.2976/1.2894846.
- WHITTAKER, S. B.-M., CLAYDEN, N. J. & MOORE, G. R. (2011). NMR characterisation of the relationship between frustration and the excited state of IM7.

- Journal of Molecular Biology* **414**, 511–529. doi: 10.1016/j.jmb.2011.09.038.
- WHITTINGHAM, J. L., EDWARDS, D. J., ANTONSON, A. A., CLARKSON, J. M. & DODSON, G. G. (1998). Interactions of phenol and m-cresol in the insulin hexamer, and their effect on the association properties of B28 pro – Asp insulin analogues. *Biochemistry* **37**, 11516–11523. doi: 10.1021/bi980807s.
- WOLYNES, P. G. (1996). Symmetry and the energy landscapes of biomolecules. *Proceedings of the National Academy of Sciences of the United States of America* **93**, 14249–14255.
- WOLYNES, P. G. (1997a). As simple as can be? *Nature Structural Biology* **4**, 871–874.
- WOLYNES, P. G. (1997b). Folding funnels and energy landscapes of larger proteins within the capillarity approximation. *Proceedings of the National Academy of Sciences of the United States of America* **94**, 6170–6175.
- WOLYNES, P. G. (2004). Latest folding game results: protein a barely frustrates computationalists. *Proceedings of the National Academy of Sciences of the United States of America* **101**, 6837–6838. doi: 10.1073/pnas.0402034101.
- WOLYNES, P. G. (2005). Energy landscapes and solved protein-folding problems. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **363**, 453–464.
- WOODSIDE, M. T., ANTHONY, P. C., BEHNKE-PARKS, W. M., LARIZADEH, K., HERSCHLAG, D. & BLOCK, S. M. (2006). Direct measurement of the full, sequence-dependent folding landscape of a nucleic acid. *Science* **314**, 1001–1004. doi: 10.1126/science.1133601.
- WRIGHT, C. F., TEICHMANN, S. A., CLARKE, J. & DOBSON, C. M. (2005). The importance of sequence diversity in the aggregation and evolution of proteins. *Nature* **438**, 878–881. doi: 10.1038/nature04195.
- YADAHALLI, S. & GOSAVI, S. (2014). Designing cooperativity into the designed protein Top7. *Proteins* **82**, 364–374. doi: 10.1002/prot.24393.
- YAMASAKI, M., LI, W., JOHNSON, D. J. D. & HUNTINGTON, J. A. (2008). Crystal structure of a stable dimer reveals the molecular basis of serpin polymerization. *Nature* **455**, 1255–1258. doi: 10.1038/nature07394.
- YANG, S., CHO, S. S., LEVY, Y., CHEUNG, M. S., LEVINE, H., WOLYNES, P. G. & ONUCHIC, J. N. (2004). Domain swapping is a consequence of minimal frustration. *Proceedings of the National Academy of Sciences of the United States of America* **101**, 13786–13791. doi: 10.1073/pnas.0403724101.
- YI, Q., RAJAGOPAL, P., KLEVIT, R. E. & BAKER, D. (2003). Structural and kinetic characterization of the simplified SH3 domain FP1. *Protein Science* **12**, 776–783. doi: 10.1110/ps.0238603.
- YUE, K., FIEBIG, K. M., THOMAS, P. D., CHAN, H. S., SHAKHNOVICH, E. I. & DILL, K. A. (1995). A test of lattice protein folding algorithms. *Proceedings of the National Academy of Sciences of the United States of America* **92**, 325–329.
- ZERBETTO, M., ANDERSON, R., BOUGUET-BONNET, S., RECH, M., ZHANG, L., MEIROVITCH, E., POLIMENO, A. & BUCK, M. (2013). Analysis of 15N-1H NMR relaxation in proteins by a combined experimental and molecular dynamics simulation approach: picosecond-nanosecond dynamics of the Rho GTPase binding domain of plexin-b1 in the dimeric state indicates allosteric pathways. *Journal of Physical Chemistry B* **117**, 174–184. doi: 10.1021/jp310142f.
- ZHANG, H. J., SHENG, X. R., PAN, X. M. & ZHOU, J. M. (1997). Activation of adenylate kinase by denaturants is due to the increasing conformational flexibility at its active sites. *Biochemical and Biophysical Research Communications* **238**, 382–386. doi: 10.1006/bbrc.1997.7301.
- ZHANG, X.-J., BAASE, W. A. & MATTHEWS, B. W. (2002). A helix initiation signal in T4 lysozyme identified by polyalanine mutagenesis. *Biophysical Chemistry* **101–102**, 43–56.
- ZHANG, Z. & CHAN, H. S. (2010). Competition between native topology and nonnative interactions in simple and complex folding kinetics of natural and designed proteins. *Proceedings of the National Academy of Sciences of the United States of America* **107**, 2920–2925.
- ZHENG, W., SCHAFFER, N. P., DAVTYAN, A., PAPOIAN, G. A. & WOLYNES, P. G. (2012). Predictive energy landscapes for protein-protein association. *Proceedings of the National Academy of Sciences of the United States of America* **109**, 19244–19249. doi: 10.1073/pnas.1216215109.
- ZHENG, W., SCHAFFER, N. P. & WOLYNES, P. G. (2013). Frustration in the energy landscapes of multidomain protein misfolding. *Proceedings of the National Academy of Sciences of the United States of America* **110**, 1680–1685. doi: 10.1073/pnas.1222130110.
- ZIMM, B. H. & BRAGG, J. K. (1959). Theory of the phase transition between helix and random coil in polypeptide chains. *Journal of Chemical Physics* **31**, 526.
- ZONG, C., WILSON, C. J., SHEN, T., WOLYNES, P. G. & WITTUNG-STAFSHED, P. (2006). ϕ -Value analysis of apo-azurin folding: comparison between experiment and theory. *Biochemistry* **45**, 6458–6466.
- ZUCKERKANDL, E. & PAULING, L. (1965). Molecules as documents of evolutionary history. *Journal of Theoretical Biology* **8**, 357–366.