IOPscience iopscience.iop.org

Home Search Collections Journals About Contact us My IOPscience

Constraint theory and hierarchical protein dynamics

This article has been downloaded from IOPscience. Please scroll down to see the full text article.

2004 J. Phys.: Condens. Matter 16 S5065

(http://iopscience.iop.org/0953-8984/16/44/004)

View the table of contents for this issue, or go to the journal homepage for more

Download details:
IP Address: 129.252.86.83
The article was downloaded on 27/05/2010 at 18:25

Please note that terms and conditions apply.

INSTITUTE OF PHYSICS PUBLISHING JOURNAL OF PHYSICS: CONDENSED MATTER

J. Phys.: Condens. Matter **16** (2004) S5065–S5072 PII: S0953-8984(04)84994-4

Constraint theory and hierarchical protein dynamics

J C Phillips

Department of Physics and Astronomy, Rutgers University, Piscataway, NJ 08854-8019, USA

Received 14 August 2004
Published 22 October 2004
Online at stacks.iop.org/JPhysCM/16/S5065
doi:10.1088/0953-8984/16/44/004

Abstract

The complexity and functionality of proteins requires that they occupy an exponentially small fraction of configuration space (perhaps 10^{-300}). How did evolution manage to create such unlikely objects? Thorpe has solved the static half of this problem (known in protein chemistry as Levinthal's paradox) by observing that for stress-free chain segments the complexity of optimally constrained elastic networks scales not with $\exp N$ (where $N \sim$ 100–1000 is the number of amino acids in a protein), but only with N. Newman's results for diffusion in N-dimensional spaces provide suggestive insights into the dynamical half of the problem. He showed that the distribution of residence (or pausing) time between sign reversals changes qualitatively at $N \sim 40$. The overall sign of a protein can be defined in terms of a product of curvature and hydrophobic(philic) character over all amino acid residues. This construction agrees with the sizes of the smallest known proteins and prions, and it suggests a *universal clock* for protein molecular dynamics simulations.

1. Introduction

Levinthal raised two closely related questions in his discussion of protein dynamics [1] that are now widely referred to as 'Levinthal's paradox'. The first, discussed explicitly, arises from the number of different conformations accessible to an $N^* = 150$-residue protein, which is roughly of order 10^{300}. This number is astronomically large compared to the number (10^8) sampled by folding a natural protein: what mechanisms guide the folding process which effectively select the right path with an accuracy of order 10^{292}? The second question is similar but even more demanding: how can proteins evolve if the number of possible combinations of amino acids is similarly large? Many solutions to Levinthal's first question have been proposed [2–10], but most of them (which follow his suggestion, which relies on 'guiding' by local interactions) do not appear to apply to its close second.

Levinthal's combinatorial puzzle is not merely an idle scholastic exercise: the standard Metropolis method (from 1952) for exploring protein dynamics, molecular dynamics simulations (MDS), is a 'brute force' approach based solely on Newton's equations (1687), which *does* attempt to explore a very large number of different conformations [10]. If such

 © 2004 IOP Publishing Ltd Printed in the UK

exploration is really necessary, then it appears that even with the fastest and largest supercomputers, alone or in tandem, research will necessarily be restricted to quite small (fast-folding) proteins. We are thus led to ask, can we not take advantage of post-Newtonian methods in theoretical physics to generate minimal approaches that will still be realistic enough to describe actual protein functionality (for instance, the transition states characterized by mutagenesis experiments [11])?

The first question that must be answered is whether such minimal approaches are even possible. Levinthal's paradox refers to the fact that protein dynamics belongs to a large class of problems that mathematicians describe as exponentially complex ('NP (not polynomial) complete' in their terms), of which the most famous example is the travelling salesman, who visits cities once each, with minimal distance travelled. Enormous progress has been made in solving such problems, without introducing further specific knowledge, but the present state of the art is not encouraging. Primes can be calculated exactly at the level of Avogadro's number, which is very impressive, but still much too small. Variational (close, but not exact) solutions to the travelling salesman problem are known at the level of 25 000 cities, which sounds encouraging, until one remembers that proteins consist of 20 different amino acids (20 neighbourhoods in each city, and cities overlap!). There is one subject that is encouraging; unexpectedly, that is games, but if one looks on protein dynamics as nature's ultimate game, then that example is quite interesting.

Chess is played with six different pieces on an 8×8 lattice. This is not so much different from a small protein, and the rules of the game have been empirically designed to challenge, but not overwhelm, our intelligence. Today one can buy for $50 a computer program that operates on a PC and is competitive at the level of an international master. It took about 30 years to develop the program. It does not use only 'brute force' methods; the earlier programs, which did, played at the Class D (outright beginner) level. This success, based on combining empirically derived principles with explicit calculation, is somewhat deceptive, because a game of chess is similar to unfolding a protein from its (known) native state: the game begins with all pieces on the board, which is half filled, and the pieces are gradually removed. There is another game, the Japanese game of Go, played with one kind of piece on a 19×19 lattice, which is equally (or slightly more) challenging, which begins with an empty board that is gradually filled. Go is analogous to protein folding from a denatured state, and no progress has been made here. The best computer programs still play Go at the Class D level.

There are many lattice models of protein dynamics, including recent studies of solvent-mediated inter-residue interactions [12]. Using direct, or brute force methods, there are enormous computational advantages to working on lattices. However, as discussed below, lattice models have yielded disappointing results for the much simpler case of network glasses. Recent theories and experiments have shown that the generic properties of network glasses are much more easily and reliably obtained topologically off-lattice, for reasons that are well understood.

The lessons of these examples are that both brute force and minimal lattice models are inadequate. Moreover, if we are to achieve success, we must find the principles behind protein dynamics, but those principles cannot be stated merely descriptively, but must be convertible into quantitative measures that can be processed by the computer. Here we argue that these principles are indeed hierarchical, as suggested in a resolution of Levinthal's paradox [5]. But the quantitative hierarchy is not that given by alpha helices, beta sheets, etc, in other words, the larger scale building blocks of the protein in real space, or loops of those blocks (such as σ 'contact order' [13]); these features are consequences of, and may partially reflect, hierarchical ordering, not its causes. Instead the relevant hierarchy occurs in an abstract space determined by the marginal 'π' interactions of the protein-bending forces and hydrogen bonds.

This immediately presents us with a problem: how do we identify the hierarchical factors that are involved in such abstract spaces?

One approach to this problem is to study a simpler case first. Proteins can be regarded as the most complex and sophisticated off-lattice networks known: what do we know about the properties of simpler off-lattice networks, such as network glasses? Here the results are most encouraging, because these simpler off-lattice networks satisfy some of the same demanding conditions satisfied by proteins. In particular, proteins fill space very efficiently, much like icosahedra [14], and so do network glasses. The glass-forming regions of network glasses can be mapped out, and it turns out that their density maxima often are located at the centres of these regions. Network glasses at the compositions of their density maxima exhibit striking properties that shed a great deal of light on the nature of protein dynamics.

A hierarchical theory explains these generic 'coincidences', and indeed it has played a major part in their discovery. First, a caveat and a claim. Glass technology, that is, glass practice, predates modern science by millennia. Glass technology has, in fact, long resisted the incursions of modern science (not that modern science has made many attempts to understand glass), in part for proprietary reasons. In this environment a largely empirical and fragmentary ceramics literature developed, apparently quite satisfactorily, until scientists made major breakthroughs (microporous glasses, optical fibers, sol–gel synthesis). The hierarchical theory discussed here represents the first effort to understand generically both glasses and proteins *as networks*; it remains to be seen whether this theory is part of a new wave, or merely an isolated effort that lies outside the mainstream of the narrowly specialized sociology of science as presently practiced.

2. Network glasses

The hierarchical theory of network glasses is axiomatic. It originated with the observation [15] that the glass-forming tendency is maximized in chalcogenide network glass alloys when the number of valence bond constraints per atom (N_c) is equal to the number of degrees of freedom (N_d). (Lagrange (1789) first recognized the convenience of using constraints to describe restricted dynamics, but he considered only macroscopic examples where the number of constraints was small.) The axiomatic condition

$$N_c = N_d \tag{1}$$

very economically describes *several* properties of ideal glass networks. First, the number of degrees of freedom/atom N_d is just the dimensionality d (ordinarily $d = 3$) of the space in which the network is embedded. In other words, already at this earliest stage, one of the most important aspects of glass networks is basic to the theory, namely, their space-filling character. The number N_c of valence bond constraints is obtained, in the manner of Pauling, by studying the crystal structures of corresponding compounds, together with partial radial distribution functions in the glass, as given by diffraction experiments. In chalcogenide glass alloys composed of atoms of similar size (for example, Ge–As–Se), the valence bond constraints that are intact when the glass is formed by quenching from the melt are usually all the two-body, nearest neighbour bond-stretching interactions, and the three-body second nearest neighbour bond-bending interactions. Moreover, all the bonds are single bonds and the atoms are coordinated as one would expect from valence chemistry, with $8 - N$ nearest neighbours (Ge, 4; As, 3, and S or Se, 2).

The calculation of the σ bond-stretching contribution to N_c is very simple: each atom contributes $N/2$, but the calculation of the bond-bending contribution to N_c is more subtle; as one would expect, non-central π forces are much more difficult to handle (which is

why they are nearly always replaced by central forces in many models of glasses, even though it is known from analysis of both hydrocarbon and semiconductor vibrational spectra that valence force fields are much more accurate and economical when bending forces are included explicitly). However, by addressing this difficulty directly constraint theory reaches a remarkable conclusion. One might think that the number of bond angles at a given N-fold coordinated atom would be given simply by $N(N-1)/2$, but for larger N this formula fails, because not all the bond angles are linearly independent. The correct single bond formula for $d = 3$ is $2N - 3$, so that the number of independent bond angles is linear in N, not quadratic: again this result is a consequence of space-filling. (Note that space-filling, which seems to be a complex geometrical question, has been reduced to a very simple (*linear*) algebraic formula.) For atoms with an average of N single bonds/atom, with stretching and bending constraints, (1) becomes

$$N_c = 5N/2 - 3 = 3 \tag{2}$$

or $N = 2.40$.

Are the results for the simplest case, Ge–As–Se, universal? By no means! However, one can use these results as the starting point for constraint theory, and gradually expand the axioms to include more and more complex cases. The insights obtained from constraint theory are very different from those obtained by lattice models, even though both approaches are designed to be consistent with space-filling, because few lattice models recognize the existence of non-central forces, or include dimensionality explicitly, so that space filling takes place generically and is not imposed artificially. Moreover, constraint theory is explicitly hierarchical, as the bond-stretching forces are much stronger than the bond-bending ones. From this many have concluded that bond-bending forces are only marginal and can be ignored in favour of spherical or central-force pair models, but this is a dangerous assumption. It is analogous to assuming that chemical reactivities of aromatic hydrocarbons are determined by σ bonds, but Hueckel (and many others since) showed that it is actually the marginal π interactions that are decisively important. Similarly, in chalcogenide network glasses, the forces that are critical to determining chemical trends in properties of glass transitions are usually the marginal noncentral bond-bending forces.

There are four special cases that have been identified so far that extend the theory of chalcogenide network glasses to oxides. The first concerns bond-bending constraints in the presence of large size differences. Silica (SiO_2) is the most important oxide glass, and there is a large size difference between Si and O. If all the bond-stretching and bond-bending constraints are intact, then one finds that the optimal average coordination number $\langle N \rangle$ for forming a three-dimensional network glass is $\langle N \rangle = 2.40$. Silica is an excellent glass-former, but $\langle N \rangle = 2.67$. In this special case, the oxygen bond-bending constraints are broken, as shown by a wide distribution of Si–O–Si bond angles, and a very narrow distribution of tetrahedral O–Si–O bond angles exhibited in radial distribution functions. However, this case appears to be exceptional: in window glass (0.74 SiO_2, 0.16 Na_2O, 0.10 CaO), where the network has been diluted to a eutectic by adding Na ($N = 1$ Pauling resonating single bond), again $\langle N \rangle = 2.40$. In addition to soda, lime (CaO) is included in window glass to promote chemical stability and increase resistance to etching. The proportions of soda and lime are determined by a second space-filling condition, namely that the average ring size is the same (6) as in the parent material SiO_2. With these two conditions, off-lattice (generic) constraint theory predicts the composition of window glass, one of nature's most remarkable materials, *exactly* ($\sim$1%), and uses *no adjustable parameters*: the composition is determined entirely topologically [16].

The second correction involves one-fold coordinated atoms, while the third involves redundant bond-bending constraints (for example, edge-sharing tetrahedral). The fourth

correction involves interfacial bond-bending constraints. These are often broken, for example, across the technologically very important Si–SiO_2 interface [17]. The latter is uniquely ideal (only one defect/10^4) for inorganic interfaces, on a scale that is known only in proteins created by the magic hand of evolution. Normally one does not expect that structural theories will be accurate on this scale, and certainly lattice models, with their multiple approximations, cannot expect to achieve such accuracies. The accuracy is a direct result of the applicability of the 'limbo' condition $N_c = N_d$.

What happens when the 'limbo' condition $N_c = N_d$ is not satisfied? This question can be answered at two levels: using mean field theory, and exactly (including all local field corrections). Numerical simulations provided and checked the mean field answer [18]. When $N_c = N_d$ the network is isostatic (marginally rigid). When $N_c < N_d$ the network is undercoordinated, and it becomes soft. Neglecting all weaker forces, the network develops a large number of zero-frequency modes. These modes were called cyclical modes by Hamilton (1830), and today physicists usually associate them with specific symmetries. However, in a large disordered network with three-body bond-bending forces, there may be a large number of cyclical modes with no discernible symmetry, lattice or otherwise. Thorpe [18] has called these modes 'floppy'; in practice, when weaker forces are included, these floppy modes appear in the measured vibrational spectrum as a large peak [19] at lower frequencies whose strength scales with $N_c - N_d$. These soft modes facilitate diffusion of oxygen or water molecules to the compact globular regions of proteins; they have been called breathing modes [20]. When $N_c > N_d$ the network hardens and is overcoordinated. The extra constraints tend to shift vibrational modes from peaks to valleys, but this effect is difficult to observe.

3. Strain energies are nonlocal

Any attempt to develop minimal models of protein folding should discuss the issue of whether interactions are local or nonlocal. It turns out, as some have suspected [9], that nonlocality in general, specifically of strain energies, is one of the keys to resolving Levinthal's paradox. Nonlocality of strain energies is a general property associated with mechanical equilibrium. It is often assumed, especially in contact models, that if only short-range covalent forces are involved, then strain energies must be local. Because of the space-filling nature of protein chains, what actually happens is quite different. Suppose that we have an inhomogeneously constrained polypeptide chain with side groups such that some segments are overconstrained and rigid, while others are underconstrained and floppy. Then the actual motion of the entire chain will *redistribute* the excess constraints of the rigid segments into the constraint space of adjacent floppy segments, a process that is intrinsically nonlocal (it extends far beyond the range of the molecular forces).

The great strength of constraint theory is that it takes qualitative ideas like this, for which there is already considerable evidence (it has been estimated [5] that the average length of rigid protein segments is 25–30 residues), and organizes it in the context of a rigorous and generic theory of off-lattice networks. This theory has evolved over the last 150 years; it was initially developed by Maxwell to discuss the mechanics of scaffolds (buildings, bridges, etc) with only central forces, but the principles are much more general. Maxwell's mean field ideas have been developed by mathematicians to describe local field corrections in the context of what they call graph theory. An algorithm for applying these ideas has been developed by Thorpe [21], in what he calls the 'pebble game' (pebbles are used to count excess constraints, and to redistribute them from overconstrained to underconstrained regions).

We now come to the most important results of the pebble game. Within a hierarchically ordered constraint space, the pebble game eventually redistributes all excess

constraints *uniquely*. Thus one has one's cake and one eats it too: the interactions are nonlocal, and explain the origin of large structural subunits or 'loops', but at the same time there is a well-defined 'funnel' that always leads to the folded state. Moreover, the number of configurational 'pebble' steps N needed to reach the folded state is not exponential in N^*, the number of residues, or even a polynomial in N^* with an exponent of order 30; it is [22] rigorously *linear* in N^*! (Of course, there will always be configurational barriers to these steps that impede the real-time dynamics of protein unfolding, but the number of these barriers scales with N^*, not $\exp N^*$.) These properties will be characteristic not only of protein folding, but of any dynamical process, such as evolutionary candidates for new proteins, formed by hierarchically combining available proteins. At each step joining smaller old proteins to form larger new ones will be simplified from exponentially unlikely in N to linear in N by the fact that the smaller ones are already in mechanical equilibrium (nearly isostatic in their transition states defined below). This is the mechanism that rigorously resolves Levinthal's paradox.

4. Defining transition states

When the isostatic condition $N_c = N_d$ is satisfied, the network undergoes a stiffness transition from floppy to rigid; observation of the stiffness transition in network glasses is discussed elsewhere in this volume by Vempati and Boolchand, who also shows that the single stiffness transition defined in mean field theory usually divides into two transitions because of local field corrections. These two transitions define the mechanical intermediate phase, which is nearly identical to the thermal intermediate phase discussed below. In this phase an isostatic backbone percolates through a floppy network glass matrix: the situation closely resembles a protein immersed in (floppy) water, except that the protein is probably not perfectly isostatic. In any case, the network glass analogue of the transition state of a protein is very well defined [22]: let the number of floppy modes be N_f, and the average number of constraints be N_c. The transition state occurs at the inflexion point of $N_f(N_c)$

$$d^2 N_f / dN_c^2 = 0. \tag{3}$$

The actual transition states of proteins are somewhat less well defined, but must be very similar mechanically and thermally to the intermediate phase of network glasses.

5. A universal clock for molecular dynamics simulations

To establish the validity of constraint theory for proteins Thorpe and co-workers have stressed its computational capabilities, for instance its successful prediction of transition state structures. Here we have reviewed at a simpler level the way in which constraint theory resolves Levinthal's paradox, as well as its successes in *defining* transition states, which are ambiguous in conventional molecular dynamics simulations. Whatever may be the limitations of the latter, they are certain to occupy the attention of many workers for some time to come. If one reviews various efforts [23], all focused on a common problem such as the transition state of protein A (a 60 residue fast folder) [11], one is struck by the fact that even for selected 'simple' cases there is no standard quantitative procedure for presenting results. Clearly what is needed is a way of benchmarking protein dynamics that will go beyond merely qualitative descriptions of the relative motions of rigid subunits in some vaguely defined configurational landscape.

A simple way to benchmark protein dynamics should emphasize the role played by hydrophobic(philic) interactions in protein collapse from the denatured state through the transition state to the fully folded native state. It should also stress marginal π interactions explicitly. The centre of mass $\mathbf{r}_n$ of any local peptide unit based on three consecutive C_α centres

$(n-1, n, n+1)$ determines the inside of that unit. One can define a local hydrophobicity π_n using the scalar product of two vectors, $(\mathbf{R}_n - \mathbf{r}_n)$ and $(\mathbf{S}_n - \mathbf{R}_n)$, where $\mathbf{R}_n$ is the location of $C_{\alpha n}$, and $\mathbf{S}_n$ is the centre of mass of the amino acid sidegroup attached to $C_{\alpha n}$. This product should then be weighted by a suitable sidegroup hydrophobicity. (If the latter is close to zero, as for glycine and serine, the corresponding triplet can be omitted from the definition of Π given below.) Thus if all the hydrophobic sidegroups are locally inside, while all the hydrophilic sidegroups are locally outside, each π_n will be positive. In fact, this is not the case even in the native state, where there is still a partial balance of hydrophobic(philic) forces, but it is clear that as the protein collapses from the denatured state, the number of positive π_n will increase at the expense of the number of negative ones. In other words, sign reversals will occur, and counting these as they occur provides a natural topological multiclock that can be used for comparing the results of different simulations. It might be that different triplet curvatures stabilize their appropriate signs (folding) or lose them (unfolding) in different sequences, and that these differences could be used to distinguish between different pathways.

6. Punctuated equilibrium

A very attractive feature of triplet curvatures is that they can be added to most simulation programs with very little effort. The addition can be made in two ways: for individual curvelets, or for the protein as a whole. In the latter case we should measure $\mathbf{R}_n$ and $\mathbf{S}_n$ from the centre of mass of the entire protein, and define an order parameter for the protein as the product

$$\Pi = \Pi_n \pi_n. \tag{4}$$

These ideas seems to be almost trivial, but they could be far from trivial if the persistence times for sign reversals either of individual π_n, or especially for their product exhibit the peculiar crossover properties predicted for diffusion in N-dimensional configuration spaces [24].

If we assume that evolution has designed proteins so that they are optimally functional, then each peptide unit in the protein can be regarded as an independent unit, in other words, a separate dimension in configuration space, with no hidden symmetries that would reduce the dimensionality, as often occurs in non-living molecules and solids. The persistence probabilities (unchanged signs of N variables between -1 and 1) for N-dimensional diffusion are dominated by short-time events for small N, but exhibit a crossover to qualitatively different behaviour for large N, where they are dominated by rare (intermittent) events. This kind of behaviour is called punctuated equilibrium [25, 26], originally conjectured as a mechanism for evolution, for which evidence is still inconclusive at the macroscopic level [27]. The microscopic crossover behaviour identified in numerical simulations for the diffusive model [24] occurs for N between 30 and 50, which is strongly suggestive. The smallest size for living proteins appears to be uncertain (some examples could be only fragments), but values for this threshold appear to fall just in the range between $N \sim 30$ and 50 residues. If one thinks of protein functionality as requiring almost perfect reversibility, then that reversibility (without entanglement) might be best achieved by pausing between conformational events long enough to achieve approximate equilibrium.

Persistence probabilities for reversal of $\Pi = \Pi_n \pi_n$ can be easily obtained in any MD simulation, and these can be studied as a function of n (especially for fast folders). The theory predicts many results: first, the persistence crossover itself as a function of N; second, the absence of punctuated equilibrium in artificial protein mimics (in other words, only true proteins that exhibit punctuated equilibrium are alive); and so on. In fact, it appears that this construction may answer (in one way) a famous question [28].

There is an interesting feature of the curvelet construction: it can be used either to define a relaxation time τ_n for each peptide unit from π_n, or an average relaxation time τ for the entire

protein from Π. As the protein folds τ will gradually grow. The individual τ_n will also grow, but not uniformly. Each curvelet will behave differently. The curvelets associated with loops should increase gradually, while the curvelets associated with the formation of rigid units (α helices and β sheets) should exhibit abrupt growth as the rigid units nucleate. The τ_n can be plotted as functions of n, and the resulting plots should correlate well with the predictions of constraint theory. One can conjecture that the Arrhenius activation energy $\tau(t)$ should show a change in slope at the transition state; it is even possible that stretched exponential relaxation would be observed.

References

[1] Levinthal C 1969 How to fold graciously *Mossbauer Spectroscopy in Biological Systems* ed P Debrunner, J Tsibris and E Munck (Urbana: Univ. Illinois) pp 22–4
[2] Vendruscolo M, Zurdo J, MacPhee C E and Dobson C M 2003 *Phil. Trans. R. Soc.* A **361** 1205
[3] Lopez P, Forterre P and Philippe H 1999 *J. Mol. Evol.* **49** 496
[4] Szathmary E 2000 *Phil. Trans. R. Soc.* B **355** 1669
[5] Berezovsky I N and Trifonov E N 2002 *J. Biomol. Struct. Dyn.* **20** 315
[6] Finkelstein A V 2002 *J. Biomol. Struct. Dyn.* **20** 311
[7] Grosberg A 2002 *J. Biomol. Struct. Dyn.* **20** 317
[8] Kloczkowski A and Jernigan R L 2002 *J. Biomol. Struct. Dyn.* **20** 323
[9] Rooman M, Dehouck Y, Kwasigroch J M, Biot C and Gilis D 2002 *J. Biomol. Struct. Dyn.* **20** 327
[10] Fernandez A, Belinky A and de las Mercedes Boland M 2002 *J. Biomol. Struct. Dyn.* **20** 331
[11] Sato S, Religa T L, Daggett V and Fersht A 2002 *Proc. Natl Acad. Sci.* **101** 6952
[12] Leonhard K, Prausnitz J M and Radke C J 2003 *Biophys. Chem.* **106** 81
[13] Baker D 2000 *Nature* **405** 39
[14] Soyer A, Chomilier J, Mornon J-P, Jullien R and Sadoc J F 2000 *Phys. Rev. Lett.* **85** 3532
[15] Phillips J C 1979 *J. Non-Cryst. Solids* **34** 153
[16] Kerner R and Phillips J C 2001 *Solid State Commun.* **117** 47
[17] Lucovsky G and Phillips J C 2004 *Appl. Phys.* A **78** 453
[18] He H and Thorpe M F 1985 *Phys. Rev. Lett.* **54** 2107
[19] Kamitakahara W A, Cappelletti R L, Boolchand P, Gompf F, Neumann D A and Mukta H 1991 *Phys. Rev.* B **44** 94
[20] Fersht A 1999 *Structure and Mechanism in Protein Science* (New York: Freeman) p 47
[21] Jacobs D J, Rader A J, Kuhn L A and Thorpe M F 2001 *Proteins Struct. Funct. Genet.* **44** 150
[22] Rader A J, Hespenheide B M, Kuhn L A and Thorpe M F 2002 *Proc. Natl Acad. Sci. USA* **99** 3540
[23] Wolynes P G 2004 *Proc. Natl Acad. Sci.* **101** 6837
[24] Newman T J and Loinaz W 2001 *Phys. Rev. Lett.* **86** 2712
[25] Gould S J 1977 *Paleobiology* **3** 135
[26] Bak P and Sneppen K 1993 *Phys. Rev. Lett.* **71** 4083
[27] Bokma F 2002 *J. Evol. Biol.* **15** 1048
[28] Schrodinger E 1945 *What is Life?* (London: Cambridge University Press)