

A First-Principles Model of Early Evolution: Emergence of Gene Families, Species, and Preferred Protein Folds

Konstantin B. Zeldovich¹, Peiqiu Chen^{1,2}, Boris E. Shakhnovich³, Eugene I. Shakhnovich^{1*}

1 Department of Chemistry and Chemical Biology, Harvard University, Cambridge, Massachusetts, United States of America, **2** Department of Physics, Harvard University, Cambridge, Massachusetts, United States of America, **3** Molecular and Cellular Biology, Harvard University, Cambridge, Massachusetts, United States of America

In this work we develop a microscopic physical model of early evolution where phenotype—organism life expectancy—is directly related to genotype—the stability of its proteins in their native conformations—which can be determined exactly in the model. Simulating the model on a computer, we consistently observe the “Big Bang” scenario whereby exponential population growth ensues as soon as favorable sequence–structure combinations (precursors of stable proteins) are discovered. Upon that, random diversity of the structural space abruptly collapses into a small set of preferred proteins. We observe that protein folds remain stable and abundant in the population at timescales much greater than mutation or organism lifetime, and the distribution of the lifetimes of dominant folds in a population approximately follows a power law. The separation of evolutionary timescales between discovery of new folds and generation of new sequences gives rise to emergence of protein families and superfamilies whose sizes are power-law distributed, closely matching the same distributions for real proteins. On the population level we observe emergence of species—subpopulations that carry similar genomes. Further, we present a simple theory that relates stability of evolving proteins to the sizes of emerging genomes. Together, these results provide a microscopic first-principles picture of how first-gene families developed in the course of early evolution.

Citation: Zeldovich KB, Chen P, Shakhnovich BE, Shakhnovich EI (2007) A first-principles model of early evolution: Emergence of gene families, species, and preferred protein folds. *PLoS Comput Biol* 3(7): e139. doi:10.1371/journal.pcbi.0030139

Introduction

Molecular biology has collected a wealth of quantitative data on protein sequences and structures revealing complex patterns of the protein universe, such as markedly uneven usage of protein folds and near-scale-free character of protein similarity networks [1–5]. On a much higher level of biological hierarchy, ecology, evolution theory, and population genetics established a framework for studying speciation, population dynamics, and other large-scale biological phenomena [6–8]. While it is widely accepted that gene families and the protein universe emerged during the course of molecular evolution through selection [9–11], there is a substantial gap in our conceptual and mechanistic understanding of how molecular evolution occurred or what the determinants of selection are. Indeed, evolution, as we understand it, proceeds at the level of organisms and populations but not at the level of genomes. Evolutionary selection at the molecular level occurs due to a relation between genotype and phenotype, although a detailed understanding of this relation and its consequences for molecular evolution remains elusive.

A number of phenomenological models (e.g., Eigen's quasispecies model) were developed where fitness of an organism was related to the sequence of its genome [12–14]. A standard definition of fitness in phenomenological models is the growth rate of a population that is higher for the more fit species. However, the relationship between genotype and phenotype in quasispecies (QS) and similar population genetics (PG) models is purely phenomenological. For example, in single-fitness peak models, one specific genotype

is postulated to be most fit, while deviations from it confer selective disadvantage. Despite providing several important insights, these types of approaches lack a fundamental microscopic connection between fitness and quantities of proteins that are both easily justifiable on biological grounds and measurable (e.g., structure/stability, function, or regulation). Therefore, such models cannot accurately describe molecular evolution of proteins.

On the other hand, a number of models were proposed that focus on emergence and evolution of sequences of model proteins and RNA under direct pressure on their molecular properties such as stability [11,15–18], folding kinetics [19,20], and mutational robustness [21]. Schuster and Stadler [22] first studied the evolution of biological macromolecules, RNA, in the context of population dynamics. Later, in a series of papers, Taverna and Goldstein [10] used an Eigen model of reaction flow to grow populations of proteins modeled as 2-D 25 mers. These authors showed that when the requirement to

Editor: Gustavo Caetano-Anolles, University of Illinois, United States of America

Received: April 4, 2007; **Accepted:** June 4, 2007; **Published:** July 13, 2007

A previous version of this article appeared as an Early Online Release on June 4, 2007 (doi:10.1371/journal.pcbi.0030139.eor).

Copyright: © 2007 Zeldovich et al. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Abbreviations: DPS, dominant protein structure; LUCA, last universal common ancestor; PDUG, protein domain universe graph; PG, population genetics; QS, quasi-species

* To whom correspondence should be addressed. E-mail: eugene@belok.harvard.edu

Author Summary

Here, we address the question of how Darwinian evolution of organisms determines molecular evolution of their proteins and genomes. We developed a microscopic *ab initio* model of early biological evolution where the fitness (essentially lifetime) of an organism is explicitly related to the evolving sequences of its proteins. The main assumption of the model is that the death rate of an organism is determined by the stability of the least stable of their proteins. A lattice model is used to calculate stability of all proteins in a genome from their amino acid sequence. The simulation of the model starts from 100 identical organisms, each carrying the same random gene, and proceeds via random mutations, gene duplication, organism births via replication, and organism deaths. We find that exponential population growth is possible only after the discovery of a very small number of specific advantageous protein structures. The number of genes in the evolving organisms depends on the mutation rate, demonstrating the intricate relationship between the genome sizes and protein stability requirements. Further, the model explains the observed power-law distributions of protein family and superfamily sizes, as well as the scale-free character of protein structural similarity graphs. Together, these results and their analysis suggest a plausible comprehensive scenario of emergence of the protein universe in early biological evolution.

exceed certain stability thresholds is imposed, the resulting distribution of structures in the evolved population appears highly skewed toward more designable structures [23] and more robust (i.e., less susceptible to mutations) proteins [24].

One of the most surprising features of the protein universe is an uneven and broad distribution of proteins over folds, families, and superfamilies. While this fact had been noted by many researchers long ago [1,4,25,26], the quantitative descriptions of these distributions began to emerge only recently. Huynen and van Nimwegen found that sizes of paralogous gene families follow a power-law distribution [3]. Gerstein and coworkers [5] observed a power-law distribution of frequencies of several other properties of gene families as defined in the Structural Classification of Proteins (SCOP) database [27]. Dokholyan et al. [2] studied a network of structural similarities between protein domains (called the protein domain universe graph, or the PDUG) and found that distribution of connectivities within the PDUG follows a power law (within a limited range of connectivity variance), making it a finite size counterpart of a scale-free network. This is in striking variance with an expectation from random distribution of folds, which would result in an (approximately) Gaussian distribution of connectivities of the PDUG.

The ubiquitous nature of power-law dependencies of many characteristics of gene families and of the protein universe may suggest their possible common origin from the fundamental evolutionary dynamics and/or physics of proteins. Huynen and van Nimwegen [3], Gerstein and coworkers [5], and Koonin and coworkers [9,28,29] proposed dynamical models (the version proposed in [28] is called the birth, death, and innovation model [BDIM]) on the basis of gene duplication as a main mechanism of creation of novel types. Such models, while providing power-law distribution of family sizes in some asymptotic cases, are sometimes based on assumptions that call into question their generality. In particular, as pointed out by Koonin and coworkers, for gene

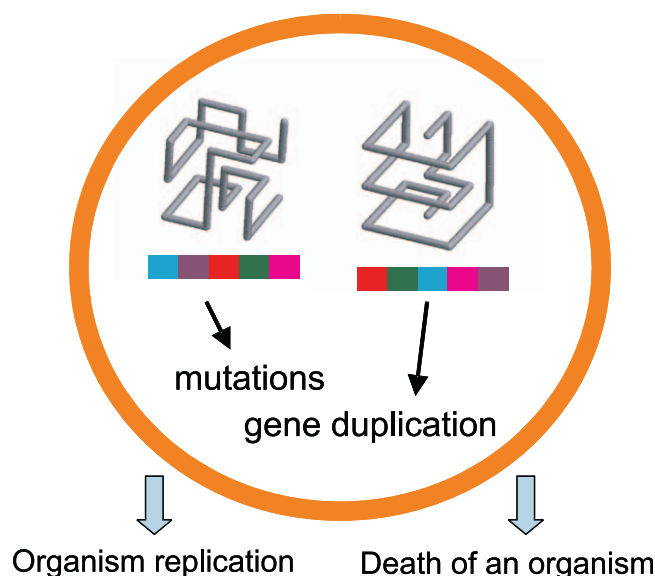


Figure 1. Schematic Representation of the Genome and Population Dynamics in the Model

Individual genes undergo mutations and duplications. Organisms as a whole can replicate, passing their genomes to the progeny, or die, effectively discarding the genome.

doi:10.1371/journal.pcbi.0030139.g001

duplication dynamic models to provide nontrivial power-law distributions of paralogous family sizes, one has to assume that the probability of gene duplication per gene depends in a certain regular way on the size of an already existing gene family. Further, even under this assumption, the power-law distribution in the birth, death, and innovation (BDIM) model arises only asymptotically in a steady state of evolutionary dynamics [29]. In contrast, the duplication and divergence phenomenological model of Dokholyan et al. [2] did not use such dramatic assumptions. However, this model is limited to an explanation of the scale-free nature of the PDUG, and it does not provide any insight as to the nature of the power-law distribution of gene family sizes. In protein sequence space, a similar approach has been employed by Qian et al. [5]. However, models like the ones proposed in [2,3,5,28] and other works are purely phenomenological in nature, whereby proteins are presented as abstract nodes and where sequence–structure relationships are not considered.

Here, we present a microscopic physics-based model of early biological evolution (Figure 1) with a realistic generic population dynamics scenario where fitness (i.e., life expectancy) of an organism is related to a simple necessary requirement of functionality of its proteins—their ability to be in native conformations. Since the latter can be estimated exactly in our model from sequences of evolving genomes, the proposed model provides a rigorous, microscopic connection between molecular evolution and population dynamics. We demonstrate that the model indeed bridges multiple evolutionary timescales, thus providing an insight into how selection of a best-fit phenotype results in molecular selection of proteins and formation of stable long-lasting protein folds and superfamilies. Furthermore, the coupling of molecular and organismal/population scales results in the emergence of species—subpopulations of evolved organisms whose genomes are similar within their groups and dissimilar

between groups. The resulting protein universe features power-law distribution for gene family and superfamily sizes closely matching real ones. The proposed model can be viewed as a first step toward a microscopic first-principles description of emergence and evolution of the protein universe.

Results

Population Dynamics, Fold Discovery, and Punctuated Evolution

Our evolution dynamics runs start from an initial population of 100 organisms, each having the same one primordial gene in their genomes. Initial gene sequence is random. Runs proceed according to evolutionary dynamics rules as described in the Methods section (see also Figure 1). The life expectancy of an organism is directly related to the stability of its proteins as explained in the Methods section; briefly, the death rate d is inversely related to protein stability P_{nat} ,

$$d = d_0(1 - \min_i P_{\text{nat}}^{(i)}). \quad (1)$$

This equation expresses a postulate that all genes of early organisms were essential at the given time; no a priori assumptions about the number of these genes are made.

We found that out of 50 simulation runs starting with different starting sequences, 27 runs successfully resulted in a steady exponential growth of the population, whereas in 23 runs the population has quickly gone extinct. A typical behavior of the population growth and protein structure dynamics in a successful evolution run is shown in Figure 2. After a period of “hesitation” lasting for about 100 time steps, a steady exponential growth of the population sets in (Figure 2B). In Figure 2C, we present the mean native state probability P_{nat} of all proteins present in the population at a given time. Owing to mutations and selection, P_{nat} steadily increases with time and dramatically exceeds the mean P_{nat} for random sequences, $P_{\text{nat}}^{\text{rand}} = 0.23$. In contrast to earlier models [10], the selection pressure is applied to whole organisms rather than to individual protein molecules. The genotype–phenotype feedback, which we model by Equation 1 (see also Methods), transfers the pressure from organisms to individual proteins to gene sequences. Figure 2B and 2C show that our selection mechanism works and results in the discovery of stable proteins due to evolutionary pressure.

Using our model, we can follow each structure in the population. In Figure 2A, color hue encodes the number of genes in the population corresponding to each of the 103,346 lattice structures (ordinate) as a function of time (abscissa). Structures marked in green are the most abundant in population at a given time, while black background corresponds to structures not found in any of the evolving organisms. The most important feature of this plot is the appearance of specific structures that correspond to highly abundant proteins comprising a significant fraction of the gene repertoire of the population. In what follows we will call them dominant protein structures (DPSs). Such proteins visually appear as bright lines on Figure 2A. What is the genesis of DPSs, and how is their appearance related to population growth or decay?

To answer this question, let us track the development of the population of structures in time by comparing the structure

repertoire, the population size, and P_{nat} plots. At $t = 0$, the proteome consists of a single sequence–structure combination (a single line on the structural repertoire plot), which corresponds to all individuals in the initial population having that structure in the genome. Over time, random mutations diverge sequences in each organism such that the dominance of a single structure is lost. This can be seen as a smeared line on the structural repertoire plot, as shown in Figure 2A, $t < 100$. However, at a certain point, very favorable sequence–structure combinations are discovered. They represent DPSs whose incorporation into the genome leads to an abrupt increase of P_{nat} and explosive exponential growth of the population through increase in fitness. Shortly after the discovery of that DPS, the diversity of the structural space abruptly collapsed, as most of the organisms converge toward the newly discovered DPS. Such a dramatic event, discovery of a limited number of dominant proteins and ensuing exponential growth of the population, can be called the “Biological Big Bang,” following a loose analogy with astrophysics. As seen on Figure 2A, the emerged dominant folds are very persistent in time. Nevertheless, fold discovery can occur at later stages of evolution. For example, in this particular simulation, at $t \approx 1,300$, new folds were discovered (white arrow in Figure 2A), they become new DPSs, and the initial DPSs are completely replaced by the new folds by $t \approx 1,600$. This switchover, accompanied by a marked increase of P_{nat} , is a clear manifestation of punctuated discoveries of new folds, coupled with selection at the organismal level.

Even though the number of organisms increases exponentially, the number of genes in each genome increases very slowly (and stabilizes after the discovery of DPSs) (Figure S1, red curve). Indeed, large genomes are not very advantageous in our model, as mutations occur in all of the genes, whereas the death rate is controlled by the gene with the lowest P_{nat} . Thus, it is only this gene that bears the brunt of selective pressure. Therefore, the rest of the genome accumulates mutations and is more prone to deleterious mutations. Unless all of the genes are very carefully selected (or formation of pseudogenes is allowed), a larger number of genes means that there is a substantial probability that a point mutation will result in a sequence–structure combination with a very low P_{nat} , immediately killing the organism. The observed slow increase of the size of the genome reflects the subtle balance between the selection pressure and gene duplication and is analyzed in more detail below. Remarkably, the average number of genes in the surviving organisms decreases with increasing mutation rates (Figure S7). Indeed, if every protein is essential, then the probability of organism death due to a deleterious mutation is lower in organisms with shorter genomes. This result allows direct experimental verification and clearly sets the current model apart from the previous sequence evolution simulations [10], which focused on the properties of individual proteins.

Figure S2 shows the structural repertoire and population size of an unsuccessful simulation run, where the population quickly became extinct. This simulation did not result in a discovery of a stable fold, and the structural space was evenly filled until the extinction of the population. We found (unpublished data) that the choice of starting sequence does not have any significance in determining whether a particular simulation run will result in exponential growth or extinction. Furthermore, in the case of most unsuccessful evolution

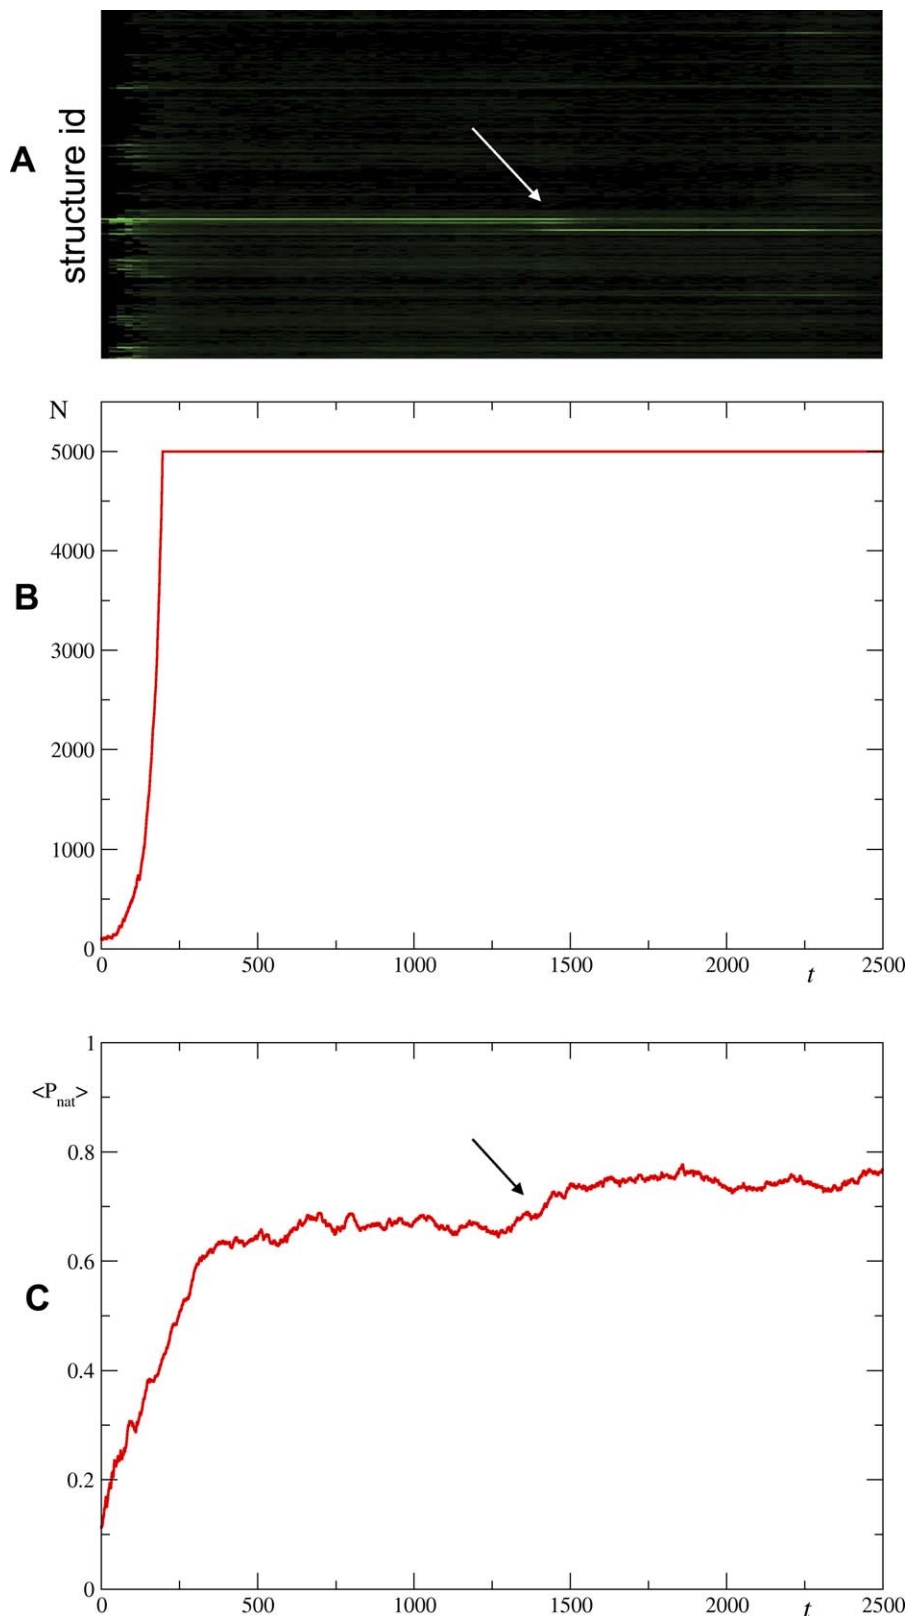


Figure 2. Time Progression of an Evolution Run

(A) Structural repertoire of an exponentially growing population as a function of time (abscissa) is shown. Ordinate represents the sequential number of the structure out of the 103,346 possible, and abundance of a structure at a given time is encoded by color. Bright green, abundant structures; black, rare or nonexistent structures. Arrows point to the discoveries of DPSs (bright lines in the structure repertoire).

(B) Population as a function of time is presented. Exponential growth sets in as soon as stable DPSs have been found.

(C) Shown is mean native state probability P_{nat} , an equivalent of mean population fitness as a function of time.

doi:10.1371/journal.pcbi.0030139.g002

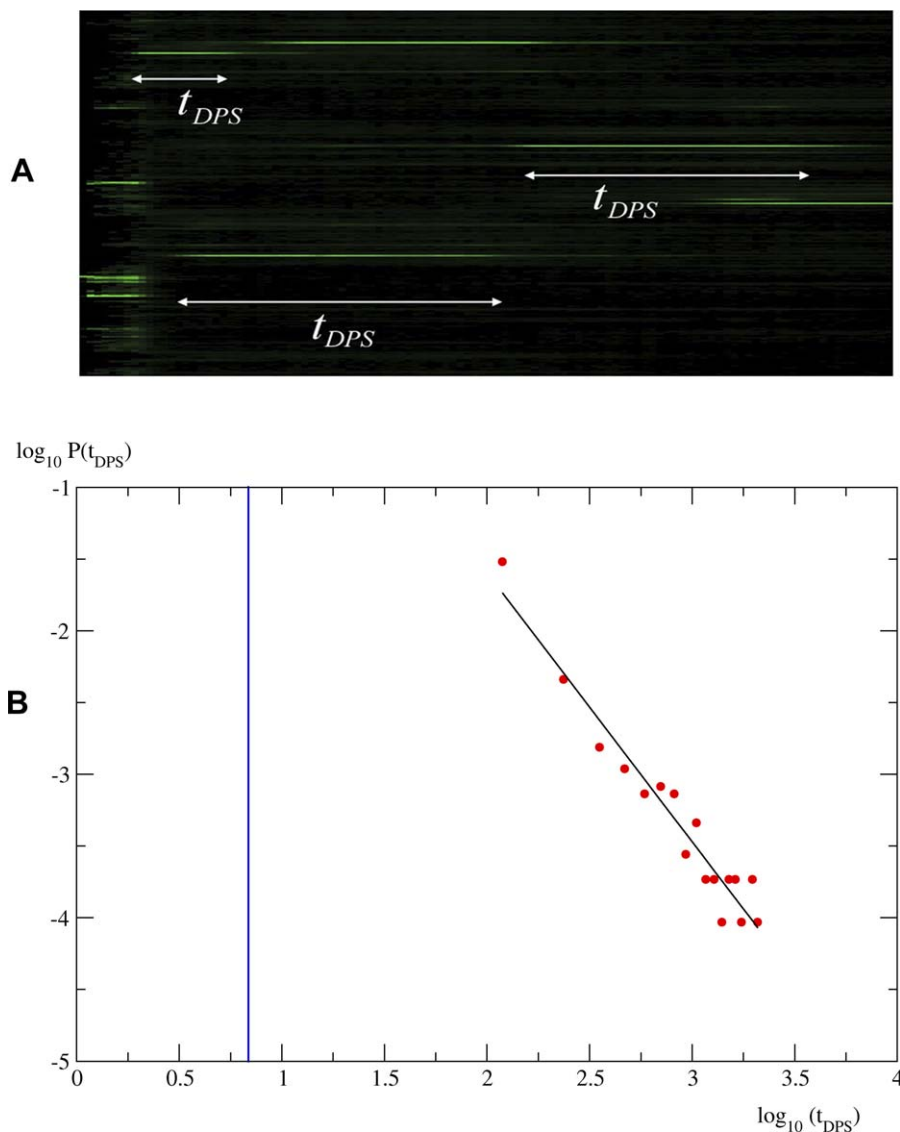


Figure 3. Distribution of Life Times of DPSs

(A) Lifetimes are defined as a span between the emergence of a DPS when it takes over at least 20% of the gene population (bright line) until its extinction as a DPS when it no longer dominates the population.

(B) The lifetime distribution of DPSs approximately follows a power law with exponent 1.87. DPS folds persist over many generations and eventually give rise to protein superfamilies. Blue line, mean lifetime of an organism.

doi:10.1371/journal.pcbi.0030139.g003

runs, the genome size rapidly increases with time (Figure S1, blue curve), decreasing the average evolutionary pressure per gene and making the discovery of DPSs less likely.

Based on these observations, we conjecture that biological evolution, exponential population growth, and existence of stable genomes are possible only after the discovery of a narrow set of specific protein structures.

Emergence of Families and Superfamilies

To quantify the persistence of the DPS during evolution, we calculated the distribution of DPS lifetimes, defined as the timespan during which a structure comprises more than 20% of the genes present in the most populated structure, i.e., the time between emergence of a DPS and its extinction in the population (see Figure 3A). We consider only DPSs that already completed their “lifecycle,” i.e., the DPSs that

emerged and went extinct over the time of an evolutionary simulation. It is clear from Figure 3B that the lifetime of a DPS is much greater than that of an organism or the average time between successive mutations. Moreover, the distribution of DPS lifetimes clearly follows a power-law-like distribution. The long nonexponential tail of the distribution demonstrates that some protein folds are extremely resistant to mutations and may persist over thousands of generations. Over such a long time, diverse protein (super)families are formed around the DPS folds. This is illustrated on Figure 4A, which shows the distribution of sizes of evolved families and superfamilies of proteins. To avoid confusion, we note that families and superfamilies here are defined not necessarily as sets of paralogous sequences but in the same way as they are defined in SCOP (Structural Classification of Proteins) [27]: protein families are defined as sets of all (not necessarily

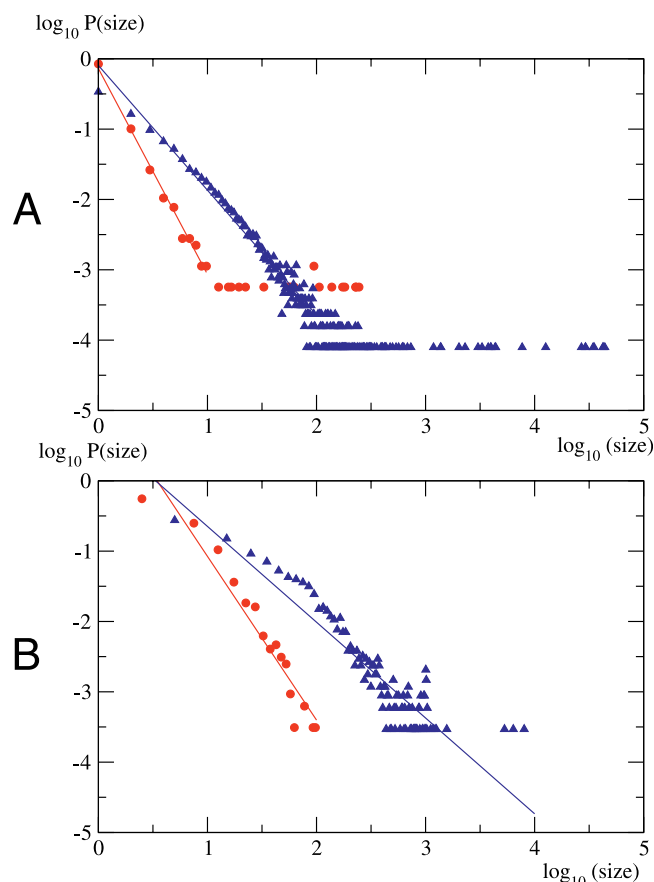


Figure 4. Distributions of Protein Family and Superfamily Sizes in Model Evolution and in Reality

Distribution of family and superfamily sizes (A) model evolution. Blue triangles represent the number of sequences folding into the same structure (gene family); the blue solid line approximates a power law with exponent -1.77 . Red circles, distribution of the number of nonhomologous (Hamming distance $>56\%$) sequences folding into the same structure (superfamilies); red solid line, a power law with exponent -2.92 .

(B) Orthologous gene family and superfamily sizes in real proteins are shown. Red circles, the number of different functions performed by each domain as defined by InterPro (Bin size = 2 and Pearson $R = 0.97$ of fit with slope = -2.2); blue triangles, the number of nonredundant sequences folding into each domain. (Bin size = 10 and Pearson $R = 0.92$ of fit with slope = -1.5)

doi:10.1371/journal.pcbi.0030139.g004

belonging to the same organism) homologous sequences that fold into a given domain structure, and superfamilies are defined as all monophyletic sets of sequences whose homology may not be detectable by sequence comparison methods, but which nevertheless fold into structurally similar domains. The statistics of protein families is dominated by orthologous genes in contrast to paralogous families studied in [5,9]. As shown in Figure 4A, both family and superfamily size distributions of evolved proteins follow almost perfect power laws with power-law exponent being greater for superfamilies (-2.92) than that for families (-1.77).

To compare this result with real proteins, we plotted the distribution of family and superfamily sizes of real proteins. As a measure of family sizes, here we estimated the number of homologous sequences that fold into a given domain (see Methods), and as a proxy for superfamily size we estimated the number of functions performed by each domain. Clearly

the distributions in Figure 4B follow power-law statistics, and as in the model, the exponent for the superfamily distribution (-2.2) is greater than that for families (-1.6). Quantitatively, the slopes of the model and real distributions are similar.

Genome Sizes of Model Organisms

As mentioned above, the genomes of model organisms from exponentially growing populations are rather short (about three genes) in contrast to the extinct populations, where uncontrolled gene duplication is observed (Figure S1). To better understand this phenomenon, one should consider the distribution of protein stabilities P_{nat} before and after a round of mutations. Suppose each genome has N genes, and the fitness (inversely related to the probability of death of an organism) of the genome is defined by $f = \min\{P_{\text{nat}}^{(1)}, \dots, P_{\text{nat}}^{(N)}\}$. At each time step, each gene in the genome has an equal probability of mutation. For simplicity, we assume that the distribution of stability P_{nat} of a lattice protein after a point mutation (i) does not depend on the stability before the mutation and (ii) is uniformly distributed between 0 and 1. Such a crude approximation works surprisingly well for lattice proteins (see Figure S3) and allows for an analytic calculation of the average genome fitness f' after a round of point mutations (see Methods). The average fitness after a point mutation depends on the number of genes N and original fitness f , and since larger genomes accumulate more mutations, they are more prone to a decrease in fitness after the mutation. In particular, if the fitness of an original genome was f , then after one round of point mutations the average new fitness f' will be no less than f only if the genome is sufficiently short, namely if

$$N < \frac{3f + \sqrt{17f^2 - 16f + 8}}{2f}. \quad (2)$$

In other words, on average, organisms with more than N genes will decrease their fitness after a point mutation and will be eventually washed out from the population. Thus, Equation 2 establishes an upper boundary on the number of genes per organism at a given level of stability $f = \min\{P_{\text{nat}}^{(i)}\}$ in the weakest link model of evolution with lattice proteins. In Figure 5, we plotted the predicted boundary from Equation 2 and the results of 50 simulation runs, where we show the scatter between the average number of genes per organism N in a population and average stability P_{nat} of proteins in a population at every time step during each of the runs. As predicted by Equation 2, only organisms with sufficiently short genomes survive at a given level of protein stability; the higher the stability, the lower the maximum possible number of genes per organism. Indeed, in a genome consisting of very stable proteins, most of the mutations are deleterious and confer a lethal phenotype in our evolutionary model. In this particular model, no more than three genes can be present in a genome at very high values of P_{nat} . It should be noted that this consideration applies only to the equilibrium size of the genome at large evolutionary times and does not describe the entire course of its evolution in time (Figure S1).

A more realistic distribution of changes of P_{nat} upon mutation and a more detailed consideration of the effect of mutations on the fate of organisms result in realistic

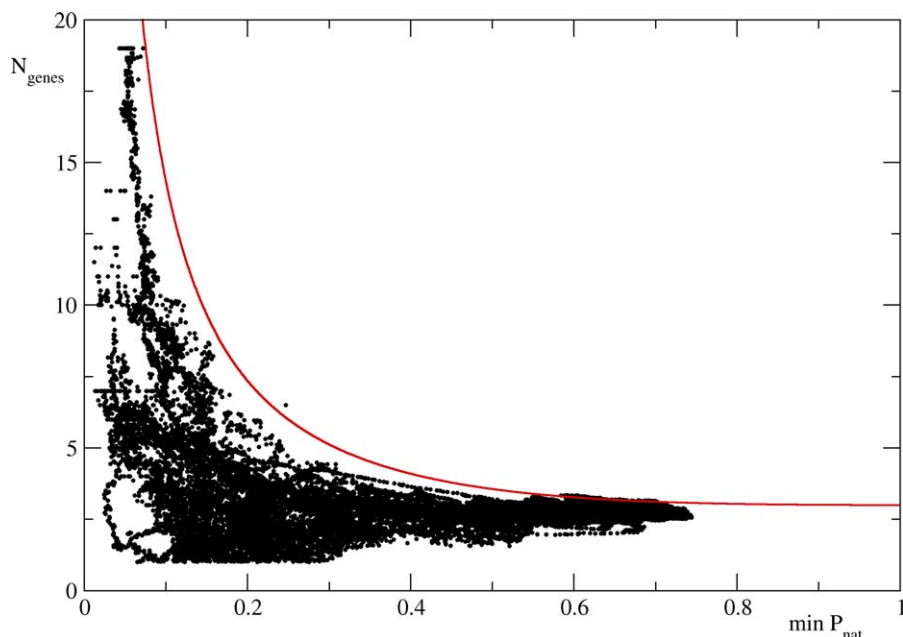


Figure 5. Analytic Prediction for the Maximum Number of Genes in an Organism as Function of the Mean Protein Stability P_{nat} (f in Equation 2) in the Analytical Model and the Results of Simulations

The data from 50 simulation runs, both exponentially growing and extinct, have been combined. Red curve, analytical model; black dots, results of the simulations.

doi:10.1371/journal.pcbi.0030139.g005

estimates of genome sizes for real organisms (K. B. Zeldovich, P. Chen, and E. I. Shakhnovich, unpublished data).

Emergence of Clonal Lines or Species

The exact nature of the model gives us direct access to the genomes of all evolved organisms, and an interesting question is whether all evolved genomes are similar (monoclonal or single-species population), and, if not, can they be clustered into distinct clonal lines or species. It turns out that the number of DPSs in the evolved population is a very good indicator of species formation. In many cases, there is only one DPS in the evolved population. Then, the genomes of all organisms are similar, and the population is monoclonal. A more interesting case is presented in Figure 6A, where two different DPSs corresponding to structures 10,107 (“A”) and 15,550 (“B”) (in our arbitrary numbering) have evolved. Are these structures randomly distributed between organisms, or are there groups of organisms preferentially using structure A, but not B, and vice versa? In the latter case, one could argue that two clonal lines or species have evolved, as each of the groups will have its own and distinct set of protein structures, and correspondingly, sequences. It turns out that in 1,536 organisms in the population, at least one gene encodes for the structure A, in 2,767 organisms, at least one gene encodes for B, but there are no organisms that include both A and B in their genomes. A total of 697 organisms have neither A nor B in their genomes. Therefore, organisms having the DPS of fold A in their genomes are very distinct from the organisms with the fold B. This difference is further illustrated in Figure 6B, where we plotted the histograms of pairwise Hamming distances between the sequences encoding for structures A and B. The black curve represents the distribution of all pairwise Hamming distances between the sequences encoding for structure A; the red curve corre-

sponds to structure B. Both curves are shifted toward lower values of the Hamming distance, illustrating a certain degree of similarity of multiple sequences encoding for the same structure. However, the Hamming distance between the sequences encoding for A and sequences encoding for B (green curve) is much larger and is very close to that of purely random sequences. Therefore, in sequence space, we can identify two groups of sequences that are similar within each group and dissimilar across the groups. Thus, the genomes of our model organisms can be classified according to their membership in the two well-defined groups of sequences, which is our model analogue of genome-based taxonomy. It is interesting to note that since our model is purely divergent and lateral gene transfer is not allowed, the evolving lines (or species) of organisms remain isolated, each evolving around its own DPS.

Structural Similarity Network of Evolved Proteins

In the discussion of structures of evolved proteins, an important global characteristic of the set of evolved proteins is the PDUG [2]. In this graph, nonhomologous proteins are linked by an edge if their structural similarity score exceeds a certain threshold. It is known that in natural proteins [2], the size of the largest cluster (giant component) of the PDUG abruptly shrinks at some value of the threshold, similar to the percolation transition. The degree distribution of the graph, (i.e., the probability $p(k)$ that a protein has k structurally similar neighbors), is a power law at the transition point. The scale-free character of this graph is believed to be a consequence of divergent evolution [2,30,31] as suggested by simple phenomenological “duplication and divergence” models [2]. Therefore, it is important to test whether our model can reproduce the global features of the natural

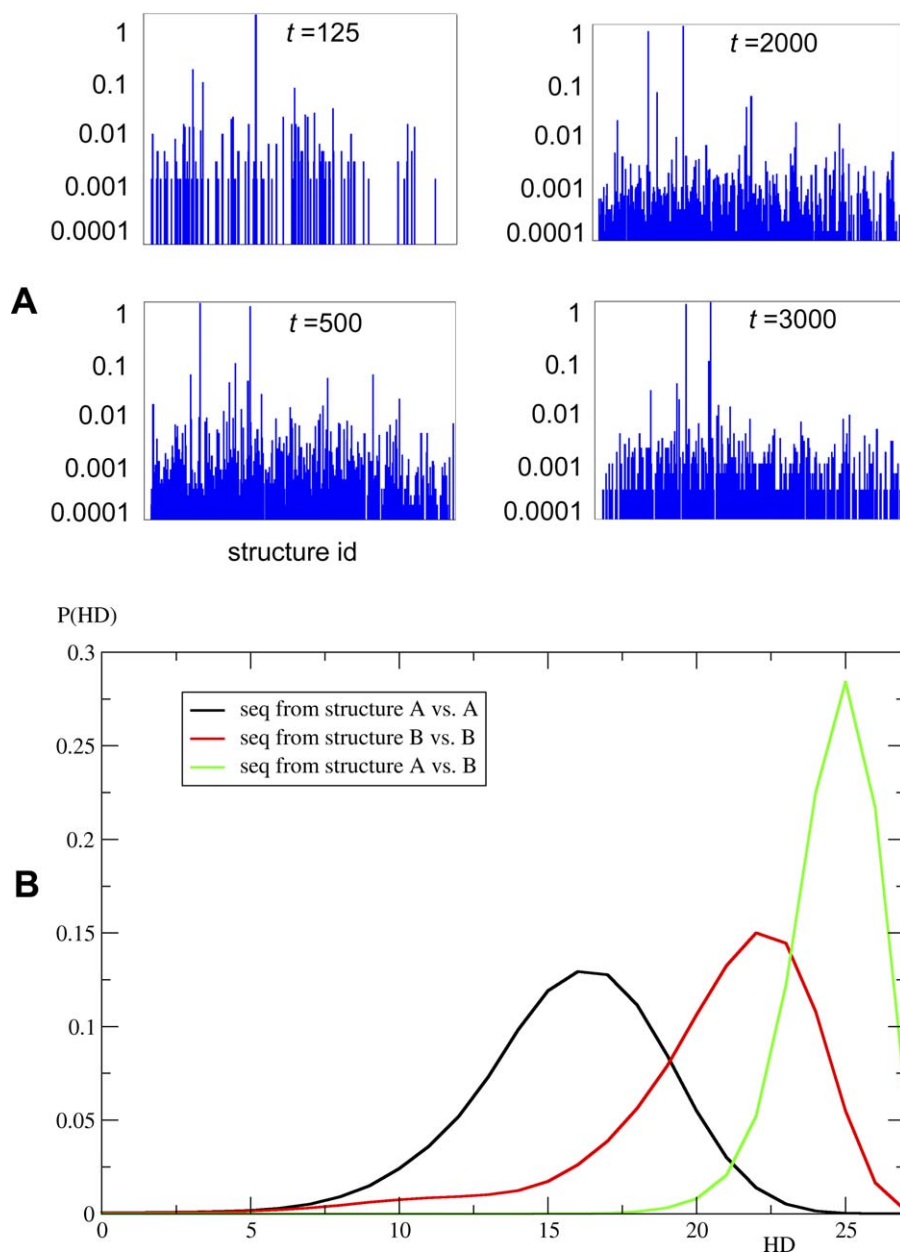


Figure 6. Emergence of Species

(A) Structural repertoire of an evolution run developing two DPSs is shown. The height of the bars represents the number of sequences folding into a given structure; the structure numbers are arbitrary.

(B) Histograms of pairwise Hamming distances between sequences corresponding to the two DPSs (black and red curves) demonstrate sequence similarity within the structure's superfamily. The histogram of Hamming distances between the sequences folding into one DPS and the sequences folding into another DPS (green) shows a lack of sequence similarity. As each organism bears only one of the two DPSs, one can say that this evolution run resulted in the formation of two different strains or species of organisms.

doi:10.1371/journal.pcbi.0030139.g006

protein universe that are manifest in the unusual properties of the PDUG.

Here we plot the PDUG of evolved proteins using Q score, the number of common contacts between a pair of proteins, as a structural similarity measure [31]. The degree distribution of the evolved PDUG at similarity threshold $Q = 17$ (the midtransition in giant component of the evolved graph, see Figure S4) is shown in Figure 7. The degree distribution plot clearly shows that the graph consists of two components, a scale-free-like component at lower k and a small but very

highly connected component at high k . As a control, we computed $p(k)$ for a divergent model without the genotype-phenotype feedback, with the fixed death rate of organisms equal to the death rate in the exponential growth regime of the evolution model. The degree distribution of the PDUG obtained in this control simulation where death rate is constant and independent of the stability of evolving proteins is shown in Figure S5. The control graph is weakly connected, indicating randomness of the discovered structures. The degree distribution of the control graph is well approximated

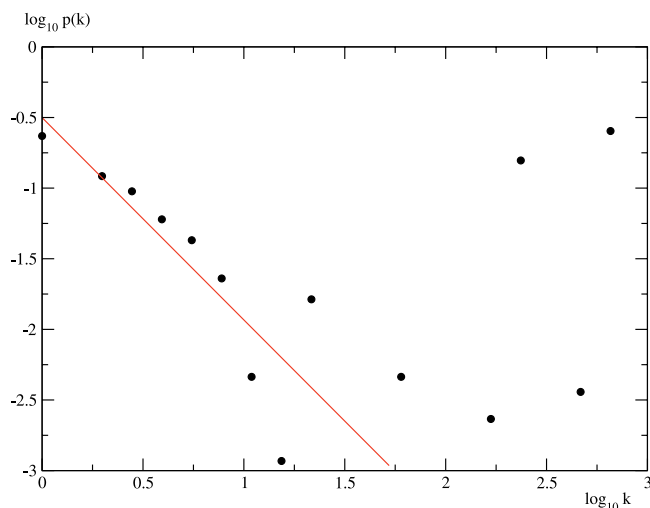


Figure 7. Degree Distribution of Structure Similarity Graph (PDUG) of the Surviving Populations in the Evolution Model

The similarity threshold was set to $Q = 17$, corresponding to the transition point in the largest cluster size (the giant component) of the graph. The slope of the linear approximation is -1.4 for $\log k < 1.75$. doi:10.1371/journal.pcbi.0030139.g007

by a Gaussian distribution, in contrast to the one obtained from evolution simulation or the real PDUG [2].

Therefore, evolutionary selection has a profound effect on the global structure of the evolved protein universe. In the model, the structural similarity graph (PDUG) splits into a scale-free-like part and a highly connected part, corresponding to the DPS, populated by many dissimilar sequences.

Our simulations predict that new folds emerge as offspring of DPSs, and in this picture the DPSs serve as prototypes of the first ancient folds. Following this logic, one should expect that ancient protein folds, being closer to prototypical DPSs, should be highly clustered and more connected than later diverged folds. To test this prediction, we analyzed the subgraph of the PDUG corresponding to the last universal common ancestor (LUCA) domains [32]. There are 915 LUCA domains. We compared the connectivity and clustering coefficient in the PDUG subgraph corresponding to LUCA domains with distributions for the same characteristics for 915 randomly selected domains as a control. The null hypothesis is that a random subset of protein domains has connectivity and clustering coefficients similar to that of the LUCA domains. In Figure S6 we present the histograms of mean connectivity k (average degree of the node) and clustering coefficient C found in 20,000 subsets of $n = 915$ randomly chosen protein domains (out of a total of 3,300 distance matrix alignment [DALI] domains constituting the PDUG, see [2]). For random subsets of 915 domains from the PDUG, $k = 2.91$ and $C = 0.197$, while the average values of the same parameters for the 915 LUCA protein domains are: $k = 4.61$ and $C = 0.267$. The values of k and C for the LUCA domains are statistically much greater than corresponding values for the random subsets (see Figure S6), yielding extremely low p -values ($p < 10^{-10}$). LUCA domains are connected and clustered just as a random subset of the PDUG (assuming Gaussian distributions of mean connectivities and clustering coefficients for random subsets of the PDUG in Figure S6). This proves that LUCA domains are

statistically more connected and clustered than an equivalent set of random protein domains as predicted from our simulations.

In Figure 8, we summarize the divergent evolution scenario as observed in our model. Divergence and selection lead to the infrequent discovery of new protein folds (dashed circles). Within these folds, mutations result in the formation of protein (super)families. The size of protein families steadily increases with time, so older families are generally larger. However, fold formation can occur at any time, branching off any family, so the newly formed families will be necessarily small. At the same time, the structures corresponding to superfamilies are all pairwise similar to each other, and for that reason they are highly clustered in the PDUG. Therefore, at any moment, the snapshot of the evolving protein universe will comprise tightly clustered families of all sizes. This picture of protein families that are “tightly knit” within each fold leads to a prediction of a peculiar property of the PDUG: that each node, (i.e., protein domain), with connectivity k , is primarily connected with nodes of similar connectivity, i.e., members of its own fold family.

To test this prediction, we follow the approach proposed by Maslov and Sneppen [33]. Connectivity correlation in the PDUG is defined as the probability $P(k_1, k_2)$, that two evolved proteins that have k_1 and k_2 structural neighbors are structurally similar to each other, i.e., are themselves connected in the PDUG. To normalize $P(k_1, k_2)$, we created 1,000 realizations of the rewired graph where each node has exactly the same connectivity as in the original graph of evolved structures, but with randomly reshuffled links to other nodes. The rewired graphs allow us to calculate the average value $P_r(k_1, k_2)$ and the standard deviation $\sigma_r(k_1, k_2)$ of the probability that nodes with connectivities k_1 and k_2 are connected in a particular network. In Figure 9A, we present the Z score for connectivity correlations $Z(k_1, k_2) = (P(k_1, k_2) - P_r(k_1, k_2)) / \sigma_r(k_1, k_2)$ for the natural PDUG. It follows from this plot that in PDUG, nodes of similar degree tend to be connected to each other: high values of $Z(k_1, k_2)$ (red) are grouped along the diagonal $k_1 = k_2$. While at low k , this property is simply a consequence of the transitivity of the measure of structural similarity (if structure A is similar to B, and B is similar to C, then C must be similar to A), it is highly nontrivial to observe this property for highly connected nodes. The pattern of connectivity correlations where similarly connected nodes tend to be connected to each other is very different from the one found in protein-protein interaction, communication, and social networks, where low-connected nodes tend to be connected with highly connected hubs, but not to each other [33]. As seen from Figure 9B, our simple evolution model perfectly reproduces this unusual pattern of connectivity correlations. The reason for such unusual property of connectivity correlation is in the punctuated character of fold discovery and evolution both in the model and in real PDUG.

Discussion

Here, we introduced a microscopic physics-based model of early biological evolution, which directly relates evolving protein sequences and structures to the life expectancy of the organism. We used a simple physical model of protein thermodynamics and a simple Malthusian model of the

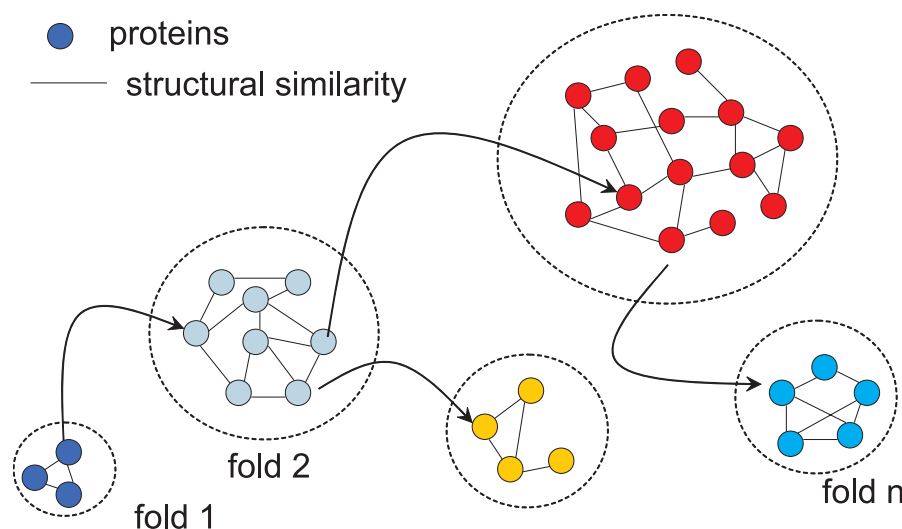


Figure 8. Schematic Representation of the Formation of Protein Folds and Superfamilies by Punctuated Jumps in the Divergent Model

Invention of new folds and their spread in population is a rare event of which the timescale exceeds the lifetime of organisms and the mutation timescale. On a shorter timescale, mutations that do not change protein structure significantly occur and fix in the population, which gives rise to protein families (on the shortest timescales) or superfamilies (on timescales longer than mutational but shorter than fold innovation). Evolutionary time increases from left to right.

doi:10.1371/journal.pcbi.0030139.g008

population dynamics. The main assumption of our minimalistic model is that the necessary condition of survival of a living organism is that its proteins adopt their native conformations. Therefore, the death rate of the organisms decreases when their proteins become more stable against thermal denaturation or unfolding. In other words, we assume that all genes of our model organism are essential. Biological function is not explicitly present in the model, but protein stability is the necessary condition for its evolution. Genes in our model have high mutation rates, conducive to rapid innovation. As such, our model can be directly applicable to (and can be experimentally tested on) the evolution of RNA viruses, which often encode for a handful of proteins, all of which are essential for the virus. The absence of an error correction mechanism results in very high mutation rates and heterogeneous, quasispecies-like populations of RNA viruses [34,35] similar to what is found in this model. Rapid evolution makes RNA viruses an ideal system for experimental studies along the lines of our model, where the simulation algorithm propagates model organisms almost like an infected host cell produces new viral particles. The low number of genes (three to ten proteins per genome depending on simulations conditions, see Figure 5) observed in our model is in part related to the extremely high mutation rates, about six mutations per genome per replication. In modern life, such a high rate is observed only in populations of RNA viruses that lack the error correction mechanism. Remarkably, the genomes of RNA viruses are rather short and normally encode for fewer than ten proteins. More complex DNA-based viruses and all cellular organisms invariably possess much lower mutation rates due to error correction and, correspondingly, longer genomes [36], in qualitative agreement with our model. Thus, our model suggests that protein stability requirements, together with mutation rates, play a crucial role in determining the size of the genomes of surviving organisms.

There is a common belief that the experimentally observed

moderate stability of natural proteins is a result of positive selection for function. However, no experimental proof for this conjecture is available. Rather, a circular argument that natural proteins are not extremely stable is offered to support this claim [37]. On the contrary, a recent study demonstrated that the higher the stability of a protein, the more likely that it confers selective advantage to the protein by making it more evolvable, by enhancing its ability to tolerate more mutations and as a result evolve a new function [38]. A more plausible explanation of moderate stability of natural proteins is that it is a direct result of a tradeoff between stability in the native conformation and entropy in sequence space, which opposes an evolutionary optimization beyond necessary levels [39]. We observe exactly this phenomenon in our model: while organisms with more stable proteins have selective advantage, the opposing factor—enormity of search in sequence/structure space—results in a compromise level of stability that corresponds to stable but not overstabilized proteins (see Figure 2C). By not “overstabilized” we mean here that for the same structure, standard sequence design methods [40,41] can provide sequences with P_{nat} values that are much closer to 1 than observed in evolved model proteins (unpublished data).

Unlike in many previous attempts, our model explicitly describes the interplay of the evolution of individual genes and that of genomes (organisms) as a whole, since the death of an organism leads to a complete loss of its genome. The model gives important insights into the interplay between molecular evolution, protein fold evolution, and population dynamics. In combination with selection pressure, random diffusion in sequence and structure spaces eventually leads to the discovery of specific structures, DPSs, that are resistant to mutations and form very evolvable proteins. This, in turn, immediately leads to the “Big Bang” event whereby discovery of viable proteins is coupled to exponential population growth, as mutations are no longer a big threat to viability. The DPSs persist over many generations and may be

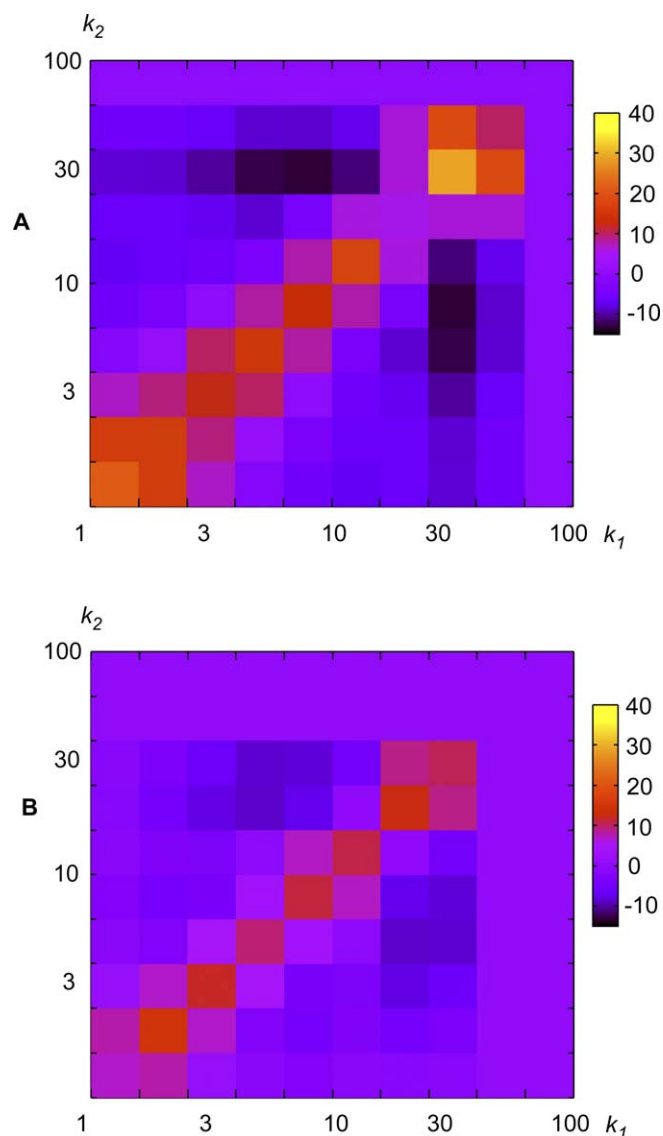


Figure 9. Node Degree Correlations in Evolved and Natural PDUG

(A) Shown is the Z score for the probability $P(k_1, k_2)$ of the two nodes with degree k_1, k_2 being connected to each other in the natural PDUG. Unlike in other networks, nodes of similar degree tend to be connected. (B) The Z score plot of $P(k_1, k_2)$ for the structure similarity graph obtained in the evolution model is remarkably similar to the actual one. doi:10.1371/journal.pcbi.0030139.g009

infrequently replaced or augmented by other even more favorable structures, in a process similar to punctuated evolution. The remarkable separation of timescales between frequent mutations and rare DPS formation allows for the formation of the protein families and superfamilies. Our model suggests that the DPSs may be superseded by more advantageous folds during evolution. A similar domain-loss phenomenon has been discussed in [42] in the context of structure-based prokaryotic phylogenies.

The model and simulations presented here provide a quantitative first-principles description of evolution of the universe of protein families. Despite the simplicity of the structural model of proteins and the phenotype-genotype relation invoked, it is able to quantitatively reproduce the power-law distributions that are observed in the natural

protein universe. Earlier phenomenological models reproduced some aspects of power-law behavior, always at the expense of invoking dramatic assumptions about the dependence of the rates of gene duplication on the sizes of already existing gene families. Here no such assumptions are made, as the model is fully microscopic in nature. Furthermore, our simulations are capable of reproducing not only “power-law”-like behavior but also marked deviation from it. Indeed, as seen on Figure 4B, there is an inflection point in the distribution of family sizes (blue curve) where the apparent slope changes. A similar inflection point was found in a recent clustering analysis of more than 7 million global ocean sampling sequences [43]. Strikingly, the distribution of family sizes of evolved proteins (Figure 4A) features a noticeable inflection as well. It is not clear whether phenomenological duplication-growth models are capable of reproducing such fine details of the family size distribution.

The most intriguing (and relevant) question is the origin of the universally observed power-law distributions in our model. Clearly an explanation proposed in many phenomenological models [3,28] is not applicable here because the rates of all processes, including gene duplication, are constant in the model and do not depend on sizes of already existing gene families. Therefore there are no ad hoc assumptions about the gene birth/death dynamics in the model that could result in power-law distributions. The only plausible reason may be that the underlying dynamics in sequence and structure spaces, coupled with selective pressure, is responsible for the emerging power-law distributions. Indeed, our key finding concerns dynamics of fold discovery and death; that the lifetimes of DPSs are power-law distributed (Figure 3B). The size of a protein family (and superfamily on longer timescales) is proportional to DPS lifetime. Indeed, power-law exponents for family size distribution and DPS lifetimes are very similar. While these observations are suggestive, a more detailed future analysis of our model will make it possible to find a definite answer as to the origin of ubiquitous power-law distributions in sequence and fold statistics.

Several earlier studies modeled the evolution of proteins by applying pressure directly on the proteins, assuming that the probability of replication of a protein in a population of proteins depends on its molecular properties such as stability [10,11,22,39], folding kinetics [19,44], or both [16]. In contrast, in the present model, biological (or as will be argued below “physiological”) constraints are applied to organisms as a whole, not to individual proteins. Evolutionary simulations and simple theory presented here highlight the importance of this distinction; the genome sizes are closely connected with maximum and average stability of evolved proteins. Therefore, biological pressure is “distributed” in the genome, and all genes act in concert in response to it. Furthermore, no DPSs were found in earlier simulations [10], despite the fact that the overall population of evolved lattice proteins was somewhat skewed toward more designable structures. In contrast, our key finding is that evolution of population is strongly coupled with protein evolution, as population growth is contingent upon discovery of a very limited set of protein structures. The difference here may be due to the fact that simpler, 2-D lattice models were used in previous simulations, or due to the differences in how biological

pressure is applied—on whole organisms here and on individual proteins in earlier works [10,45].

The presented model is markedly different from standard models of PG such as Fisher-Wright and QS [46,47]. PG and QS models are phenomenological descriptions of evolution, attributing certain values of fitness for the genomes with predetermined combinations of alleles. These models conveniently sidestep the important question of the molecular origins of the change of fitness upon recombination or mutation. The genotype–phenotype relationship in our model is not phenomenological but physiological: when a gene product loses stability (and by implication functionality), the whole organism is likely to die. This assumption is justified by recent high-throughput experiments that use RNAi to determine the impact of gene knockout on phenotype [48,49]. As in an experiment where knockout of essential genes results in death of an organism, in our model the deterioration of stability of any gene of an organism confers the lethal phenotype. In the present implementation, our model assumes that all genes are essential. However, this assumption can be relaxed (unpublished data), making it possible to study differentially the impact of biological constraints on the evolution of genes [50].

Another critical distinction between our approach and traditional phenomenological models is that in PG and QS approaches a single genotype is assumed to be advantageous [46,51]. While the outcome may be that genomes of the populations are peaked around the most fit one (as in the standard QS model [47]) or that a broader distribution among genotypes may emerge (as in the “survival of the flattest” scenario [52]), it is always an implication of the key assumption that a certain genotype confers the highest fitness. In contrast, the present model makes no a priori assumptions about the fitness advantage of a certain genotype. Strikingly, sets of organisms distributed around dominant genomes and proteomes—species—emerge here as a result of evolution at longer evolutionary times. A key factor determining the emergence of species in this model is that productive evolution occurs only when the structural diversity of proteins collapses into a small set of DPSs.

Our model of natural selection is minimalistic and is limited in its scope. It does not take into account such important biological processes as horizontal gene transfer, gene recombination, sexual reproduction, “death of a gene” (via pseudogenisation), and Darwinian selection due to competition of populations for limited resources. Also, to make the minimum possible number of assumptions, the modern amino acid alphabet is used in the model, although it has been suggested that the amino acid alphabet itself had evolved over time [53–55]. However, we believe that our model is an important step toward the unification of microscopic physics-based models of protein structure and function and the macroscopic (so far, phenomenological) description of the evolutionary pressure. Its extensions are straightforward and may include a more explicit consideration of protein function, protein–protein interactions, and fitness function that rewards functional (and therefore, structural) innovations. Furthermore, since habitat temperature enters the model explicitly, it can be used to study thermal adaptation of organisms as well as adaptation to variable mutation rates. This work is in progress.

Methods

Population dynamics and genotype–phenotype relationships. In our model, an organism is completely described by the set of its genes. The genetic code then defines amino acid sequences, and the exact nature of the lattice protein-folding model makes it possible to find the native structures of the encoded proteins. We assume that for an organism to function properly, it is imperative that its proteins spend a significant part of the time in their native conformations at a given environmental temperature. Let $P_{\text{nat}}^{(i)}$ be the thermodynamic probability that protein i is in its native conformation (see protein model below). As the simplest approximation, we assume that the probability that an organism is alive is proportional to the lowest $P_{\text{nat}}^{(i)}$ across all of its proteins:

$$P_{\text{alive}} \propto \min_i P_{\text{nat}}^{(i)} \quad (3)$$

i.e., the longevity of an organism is determined by the least stable protein in the genome (“weakest link” model).

Our model of population and genome dynamics includes four elementary events: (i) random mutation of a nucleotide in a randomly selected gene with constant rate m per unit time per DNA length; mutations leading to the stop codon are rejected to ensure the constant length of protein sequences; (ii) duplication of a randomly selected gene within an organism’s genome with constant rate u ; (iii) birth of an organism via duplication of an already existing organism with constant rate b (the genome is copied exactly); and (iv) death of an organism with the rate d per unit time (Figure 1). For simplicity, we do not allow for the formation of pseudogenes or any other mechanism of removal of the genes from a genome; in every organism the number of genes increases (or remains constant) with time. However, the average number of genes per organism in the population can either decrease or increase due to enhanced survival of organisms with shorter (longer) genomes.

Condition (Equation 3) translates into the dependence of organism death rate d on the stability of its proteins:

$$d = d_0 (1 - \min_i P_{\text{nat}}^{(i)}), \quad (4)$$

where d_0 is the reference death rate. This relation gives rise to an effective selection pressure on proteins since organisms that have at least one unstable protein live shorter and thus produce less progeny. This simple, direct, and physically plausible relationship between the genotype (thermodynamic properties of the proteins) and the phenotype (life expectancy) is the key novel feature of our model. Another implication of this relationship is in the “collective punishment” effect that genes do not evolve independently; a very unfavorable mutation in a gene will likely lead to a quick death of an organism, so its complete genome will not be able to proliferate. Such cooperativity creates an important selection pressure toward mutation-resistant genes encoding stable and evolvable (see below) proteins. Interestingly, purely physical factors ensure that resistance to mutations, evolvability of a new function, and thermostability are well correlated [24,38], so little or no tradeoff may be needed to satisfy both requirements. To ensure that a sufficient selection pressure is applied, we set $d_0 = b/(1 - P_{\text{nat}}^{(0)})$, where $P_{\text{nat}}^{(0)}$ is the native state probability of a protein encoded by the primordial gene, which is the single gene in all organisms from which evolution runs start. Therefore, the Malthus parameter $b - d$ of population growth is zero for neutral mutations (not changing P_{nat} with respect to the primordial sequence), positive for favorable mutations that increase P_{nat} , and negative for deleterious mutations. In principle, the relationship between growth rate and protein stability can be experimentally verified by analyzing the growth rate of bacteria at elevated temperatures. While the exact biochemical mechanisms leading to slower replication and eventual death are complicated, they all originate in the loss of protein function or enzymatic activity due to thermal denaturation [56]. A sequence evolution model, also using the protein stability P_{nat} as fitness parameter has been recently proposed by Goldstein and coworkers [54].

Simulation algorithm. In our model, each organism is represented by a list of its genes, 81 nucleotide sequences that are translated into amino acid sequences according to the genetic code. There can be up to 100 genes per organism; the gene duplication rate is chosen so that this limit is never reached in a simulation; typically, organisms have fewer than ten genes each at the end of a simulation. Initially, 100 organisms are seeded with one and the same primordial gene; $P_{\text{nat}}^{(0)}$ is the native state probability of the protein encoded by the primordial gene.

At each time step of the evolution, each organism can undergo one

of the five events: no event at all or the four events described in the main text (duplication of an organism with probability $b = 0.15$, death with rate d , gene duplication with probability $u = 0.03$, and point mutation of a randomly chosen gene with probability $m = 0.3$ per gene). The organism death rate is calculated according to Equation 4, $d = d_0(1 - \min_i P_{nat}^{(i)})$, with $d_0 = b/(1 - P_{nat}^{(0)})$.

Every 25 time steps, an entire set of genes of all currently living organisms is recorded for analysis. The simulation stops after 3,000 time steps. Whenever the population size N exceeds 5,000, we randomly remove $N - 5,000$ organisms to ensure constant population size, simulating a turbidostat; despite the artificially constrained population size, the growth regime remains exponential.

Protein model. To simulate the thermodynamic behavior of evolving proteins, we use the standard lattice model of proteins, which are compact 27-unit polymers on a $3 \times 3 \times 3$ lattice [57]. The residues interact with each other via the Miyazawa-Jernigan pairwise contact potential [58]. It is possible to calculate the energy of a sequence in each of the 103,346 compact conformations allowed by the $3 \times 3 \times 3$ lattice and the Boltzmann probability of being in the lowest energy native conformation,

$$P_{nat}(T) = \frac{e^{-E_0/T}}{\sum_{i=0}^{103345} e^{-E_i/T}}, \quad (5)$$

where E_0 is the lowest energy among the 103,346 conformations, E_i are the energies of the sequence in the remaining 103,345 conformations, and T is the environmental temperature (in the simulation, we assumed $T = 0.5$ in Miyazawa-Jernigan dimensionless energy units).

Analytic calculation of genome sizes of model organisms. Suppose each genome has N genes, and the fitness of the entire genome is then defined by $f = \min\{P_{nat}^{(1)}, \dots, P_{nat}^{(N)}\}$. Based on the sequence design simulation, we find that it is a reasonably good approximation to assume that in our lattice model the distribution of stability P_{nat} of a lattice protein after a point mutation (i) does not depend on the stability before the mutation and (ii) is uniformly distributed between 0 and 1. Performing a point mutation, we can either mutate the gene with the lowest fitness value, with probability $\frac{1}{N}$ (case A), or select any one of the other more stable genes and mutate it with probability $(1 - \frac{1}{N})$, case B.

In case A, because the mutated gene was the original least-stable gene, there are two possible outcomes after the mutation: (i) if the new gene fitness value is less than f , then this new gene fitness value would be the new minimum among all gene fitnesses in the genome, therefore this new fitness will become the fitness of the new genome. This occurs with a probability f , and since the new fitness follows a uniform distribution in the region $[0, 1]$, the expectation value in this case is $f/2$. So this part's contribution to the expectation value of the new genome fitness is $A_1 = \frac{1}{N}f/2$. (ii) If the new gene stability is greater than f , which happens with a probability of $(1 - f)$, we can calculate the probability distribution for the new genome fitness being x is

$$p(x) = \left(\frac{1-x}{1-f}\right)^{N-1} \text{ for } f < x < 1. \quad (6)$$

The significance of this equation is that when one gene has fitness $f < x < 1$ and is the new minimum, also under the condition that all fitnesses are within the region of $[f, 1]$, the probability for all the other $(N - 1)$ genes has to have fitness greater than x is $\left(\frac{1-x}{1-f}\right)^{N-1}$. The multiplicity of this condition is $M_{N,f} = \frac{N}{1-f}$. This is because we can pick any one of the N genes to be the new least-fit gene, and the fitness value is within the region $[f, 1]$ with uniform probability distribution. The contribution for the new genome fitness in this situation is therefore the total product of the probability of this situation $\frac{1}{N}(1 - f)$, the multiplicity $M_{N,f}$, and the expectation value $\int_f^1 xp(x)dx$,

$$A_2 = \frac{1}{N}(1 - f) * M_{N,f} \int_f^1 xp(x)dx. \quad (7)$$

In case B, we also have two possible situations, situation B_1 states when the mutated gene has a fitness less than f , and situation B_2 states when the mutated gene has a fitness greater than f . In situation B_1 , similar to the derivation in case A, the probability for the new stability to be smaller than f is f , and the expectation value of the new genome fitness in this case is $f/2$. So we have $B_1 = (1 - \frac{1}{N})f/2$.

In situation B_2 , if the stability of the mutated gene is greater than f ,

then the original gene with stability f would remain the least stable in the genome. Therefore, the genome fitness in this situation is still f , so the value of B_2 reads $B_2 = (1 - \frac{1}{N})(1 - f)f$, where $(1 - \frac{1}{N})$ is the probability to choose one of the $(N - 1)$ genes with stability greater than f , $(1 - f)$ is the probability to mutate this gene with fitness greater than f , and f is the expectation of the final fitness under this condition.

Finally, summing up A_1 , A_2 , B_1 , and B_2 , we obtain the expectation value of the genome fitness after one point mutation:

$$f' = \frac{2 - 4f + f^2(2 + 3N - N^2)}{2N(N + 1)} + f \quad (8).$$

Now, if the average genome fitness after a single point mutation must be greater than the original fitness, the condition $f' - f > 0$ must be satisfied. Solving this inequality, we find an upper limit on the number of genes in a genome (Equation 2).

Family and superfamily size estimate for real proteins. We took sequences of all structurally characterized domains from HSSP [59]. We used BLAST [60] with threshold 10^{-10} to identify all sequences with significant homology to each HSSP domain in a nonredundant sequence database NRDB90 [61]. We combined each set of sequences with homology into a single gene family. The number of non-redundant sequences matching the domain is the number considered in that family. We then used cross-indexing between NRDB90 [61], Swiss-Prot [62], and InterPro [63] to define the set of different functions each gene family performs. The number of different functions as defined by InterPro becomes the number of superfamilies folding into the same domain.

Family and superfamily size estimate for model proteins. In the model, the superfamily size is defined as the number of non-homologous sequences with all mutual pairwise Hamming distance of 16 or more (i.e., 40% sequence identity or less) having the same native conformation. The family size is defined as the number of all sequences folding into a given structure, without removing the homologous sequences.

Protein Domain Universe Graph. To construct the PDUG from the simulation data, we considered only the nonhomologous amino acid sequences. The selection is based on the Hamming distance between the sequences, which should exceed 18 (i.e., less than 33% sequence identity).

To calculate the structure similarity in the PDUG, we used the Q score similarity measure. The Q score measure between the two structures i and j is the number of all pairs of monomers (k, m) that are in contact both in structure i and structure j . As there are always 28 contacts in compact 27 mers, Q score varies from 0 for completely dissimilar structures to 28 for two identical structures. The Q score is analogous to the distance matrix alignment (DALI) Z score, used as a structural similarity measure for real proteins.

Definition of LUCA domains. The simplest construction of the LUCA that still yields useful information is the delineation of the very old domains. Any domain shared by the three kingdoms of life can be placed in the LUCA [64]. If any such domain were not placed in the LUCA, multiple independent discovery (or horizontal transfer) events would be required to explain the occurrence of this domain in all kingdoms. The "extra" evolution involved in this case would result in a less parsimonious scenario. Inclusion of other domains is more probabilistic and depends on the exact form and method of parsimony construction used [64]. We thus define the structural content of the LUCA to be all domains that have homologs in at least one archaeal, at least one prokaryotic, and at least one eukaryotic species. This yields approximately one-third of the PDUG members.

Supporting Information

Figure S1. Mean Genome Size (Number of Genes per Organism) in an Exponentially Growing Population Is almost Constant Owing to the Balance between Gene Duplication and Selection Pressure

In a simulation run where population becomes extinct, the genome size grows linearly with time (blue curve). Red curve, exponentially growing population.

Found at doi:10.1371/journal.pcbi.0030139.sg001 (26 KB EPS).

Figure S2. Time Progression of an Evolution Run Where Population Becomes Extinct

(A) Abundance of different structures in the proteomes in an unsuccessful evolution run as a function of time. Red corresponds to abundant structures, and cyan to rare or nonexistent ones.

(B) Size of population as a function of time is shown.
 (C) Mean native state probability P_{nat} as a function of time is shown. DPSs are never found in this run, resulting in extinction of the population due to random diffusion in sequence space.
 Found at doi:10.1371/journal.pcbi.0030139.sg002 (1.0 MB JPG).

Figure S3. Probability Distributions of Lattice Protein Stability P_{nat} after a Point Mutation in Sets of Sequences with a Given Stability P_{nat}^0

The peak at $P_{\text{nat}} = P_{\text{nat}}^0$ corresponds to mutations that do not alter the stability; at high P_{nat}^0 , the long constant-level tail of the distribution makes it possible to approximate the distribution by a uniform one. For each plot, we performed all possible $19 \times 27 = 513$ mutations in 100 different sequences with $P_{\text{nat}}^0 - P_{\text{nat}} < 0.03$.

Found at doi:10.1371/journal.pcbi.0030139.sg003 (83 KB EPS).

Figure S4. Fraction of the Giant Component of the PDUG as a Function of Similarity Cutoff Q for Evolution Simulations

Found at doi:10.1371/journal.pcbi.0030139.sg004 (15 KB EPS).

Figure S5. Degree Distribution of Structure Similarity Graph (PDUG) in a Control Simulation Where Genotype–Phenotype Relationship Does Not Exist

The similarity threshold was set to $Q = 17$ corresponding to the transition point in the largest cluster size (the giant component) of the graph.

Found at doi:10.1371/journal.pcbi.0030139.sg005 (12 KB EPS).

Figure S6. Probability Distribution of the Average Connectivity and

Clustering Coefficient for Random Subsets of 915 Protein Domains from the PDUG and the Value of These Parameters $k = 4.61$ and $C = 0.267$ for the LUCA Domains

The distribution is drawn over 20,000 random selections of 915 subsets out of a total of 3,300 PDUG domains. (A) Average connectivity. (B) Clustering coefficient. Red line, LUCA domains.

Found at doi:10.1371/journal.pcbi.0030139.sg006 (79 KB JPG).

Figure S7. Average Number of Genes per Organism for Two Evolution Runs Proceeding at Different Mutation Rates

Organisms evolving at a higher mutation rate possess shorter genomes. The temperature is $T = 0.8$. Red and black curves, different mutation rates.

Found at doi:10.1371/journal.pcbi.0030139.sg007 (57 KB EPS).

Acknowledgments

The authors acknowledge stimulating discussions with I. N. Berezovsky.

Author contributions. KBZ and EIS conceived and designed the experiments and wrote the paper. KBZ performed the experiments. KBZ, PC, BES, and EIS analyzed the data. PC and BES contributed reagents/materials/analysis tools.

Funding. The authors received financial support from the US National Institutes of Health for this study.

Competing interests. The authors have declared that no competing interests exist.

References

- Chothia C (1992) Proteins. One thousand families for the molecular biologist. *Nature* 357: 543–544.
- Dokholyan NV, Shakhnovich B, Shakhnovich EI (2002) Expanding protein universe and its origin from the biological Big Bang. *Proc Natl Acad Sci U S A* 99: 14132–14136.
- Huynen MA, van Nimwegen E (1998) The frequency distribution of gene family sizes in complete genomes. *Mol Biol Evol* 15: 583–589.
- Orengo CA, Jones DT, Thornton JM (1994) Protein superfamilies and domain superfolds. *Nature* 372: 631–634.
- Qian J, Luscombe NM, Gerstein M (2001) Protein family and fold occurrence in genomes: Power-law behaviour and evolutionary model. *J Mol Biol* 313: 673–681.
- Hartl DL, Kafatos FC, O'Brien SJ (1995) Genome evolution comes of age. *Curr Opin Genet Dev* 5: 705–708.
- Hartl DL, Taubes CH (1998) Towards a theory of evolutionary adaptation. *Genetica* 102–103: 525–533.
- Lynch M, Koskella B, Schaack S (2006) Mutation pressure and the evolution of organelle genomic architecture. *Science* 311: 1727–1730.
- Koonin EV, Wolf YI, Karev GP (2002) The structure of the protein universe and genome evolution. *Nature* 420: 218–223.
- Taverna DM, Goldstein RA (2000) The distribution of structures in evolving protein populations. *Biopolymers* 53: 1–8.
- Taverna DM, Goldstein RM (2000) The evolution of duplicated genes considering protein stability constraints. *Pac Symp Biocomput*: 69–80.
- Eigen M, Schuster P (1982) Stages of emerging life—Five principles of early organization. *J Mol Evol* 19: 47–61.
- Bloom JD, Silberg JJ, Wilke CO, Drummond DA, Adami C, et al. (2005) Thermodynamic prediction of protein neutrality. *Proc Natl Acad Sci U S A* 102: 606–611.
- Tannenbaum E, Shakhnovich EI (2004) Solution of the quasispecies model for an arbitrary gene network. *Phys Rev E Stat Nonlin Soft Matter Phys* 70: 021903.
- Dokholyan NV, Shakhnovich EI (2001) Understanding hierarchical protein evolution from first principles. *J Mol Biol* 312: 289–307.
- Tiana G, Shakhnovich BE, Dokholyan NV, Shakhnovich EI (2004) Imprint of evolution on protein structures. *Proc Natl Acad Sci U S A* 101: 2846–2851.
- Bastolla U, Porto M, Roman HE, Vendruscolo M (2005) Looking at structure, stability, and evolution of proteins through the principal eigenvector of contact matrices and hydrophobicity profiles. *Gene* 347: 219–230.
- Bastolla U, Farver J, Knapp EW, Vendruscolo M (2001) How to guarantee optimal stability for most representative structures in the Protein Data Bank. *Proteins* 44: 79–96.
- Mirny LA, Abkevich VI, Shakhnovich EI (1998) How evolution makes proteins fold quickly. *Proc Natl Acad Sci U S A* 95: 4976–4981.
- Li L, Mirny LA, Shakhnovich EI (2000) Kinetics, thermodynamics and evolution of non-native interactions in a protein folding nucleus. *Nat Struct Biol* 7: 336–342.
- Bornberg-Bauer E, Chan HS (1999) Modeling evolutionary landscapes: Mutational stability, topology, and superfunnels in sequence space. *Proc Natl Acad Sci U S A* 96: 10689–10694.
- Schuster P, Stadler PF (1994) Landscapes: Complex optimization problems and biopolymer structures. *Comput Chem* 18: 295–324.
- Li H, Helling R, Tang C, Wingreen N (1996) Emergence of preferred structures in a simple model of protein folding. *Science* 273: 666–669.
- Taverna DM, Goldstein RA (2002) Why are proteins so robust to site mutations? *J Mol Biol* 315: 479–484.
- Finkelstein AV, Pitsyn OB (1987) Why do globular proteins fit the limited set of folding patterns? *Prog Biophys Mol Biol* 50: 171–190.
- Thornton JM, Orengo CA, Todd AE, Pearl FM (1999) Protein folds, functions and evolution. *J Mol Biol* 293: 333–342.
- Hubbard TJ, Murzin AG, Brenner SE, Chothia C (1997) SCOP: A structural classification of proteins database. *Nucleic Acids Res* 25: 236–239.
- Karev GP, Wolf YI, Rzhetsky AY, Berezovskaya FS, Koonin EV (2002) Birth and death of protein domains: A simple model of evolution explains power law behavior. *BMC Evol Biol* 2: 18.
- Karev GP, Wolf YI, Berezovskaya FS, Koonin EV (2004) Gene family evolution: An in-depth theoretical and simulation analysis of nonlinear birth–death–innovation models. *BMC Evol Biol* 4: 32.
- Deeds EJ, Shakhnovich EI (2005) The emergence of scaling in sequence-based physical models of protein evolution. *Biophys J* 88: 3905–3911.
- Deeds EJ, Dokholyan NV, Shakhnovich EI (2003) Protein evolution within a structural space. *Biophys J* 85: 2962–2972.
- Shakhnovich BE, Deeds E, Delisi C, Shakhnovich E (2005) Protein structure and evolutionary history determine sequence space topology. *Genome Res* 15: 385–392.
- Maslov S, Sneppen K (2002) Specificity and stability in topology of protein networks. *Science* 296: 910–913.
- Domingo E, Sabo D, Taniguchi T, Weissmann C (1978) Nucleotide sequence heterogeneity of an RNA phage population. *Cell* 13: 735–744.
- Holland J, Spindler K, Horodyski F, Grabau E, Nichol S, et al. (1982) Rapid evolution of RNA genomes. *Science* 215: 1577–1585.
- Anderson JP, Daifuku R, Loeb LA (2004) Viral error catastrophe by mutagenic nucleosides. *Annu Rev Microbiol* 58: 183–205.
- DePristo MA, Weinreich DM, Hartl DL (2005) Missense meanderings in sequence space: A biophysical view of protein evolution. *Nat Rev Genet* 6: 678–687.
- Bloom JD, Labthavikul ST, Otey CR, Arnold FH (2006) Protein stability promotes evolvability. *Proc Natl Acad Sci U S A* 103: 5869–5874.
- Taverna DM, Goldstein RA (2002) Why are proteins marginally stable? *Proteins* 46: 105–109.
- Shakhnovich EI, Gutin AM (1993) Engineering of stable and fast-folding sequences of model proteins. *Proc Natl Acad Sci U S A* 90: 7195–7199.
- Morrissey MP, Shakhnovich EI (1996) Design of proteins with selected thermal properties. *Fold Des* 1: 391–405.
- Deeds EJ, Hennessey H, Shakhnovich EI (2005) Prokaryotic phylogenies inferred from protein structural domains. *Genome Res* 15: 393–402.
- Yooseph S, Sutton G, Rusch DB, Halpern AL, Williamson SJ, et al. (2007)

- The sorcerer II global ocean sampling expedition: Expanding the universe of protein families. *PLoS Biol* 5: e16. doi:10.1371/journal.pbio.0050016
44. Gutin AM, Abkevich VI, Shakhnovich EI (1995) Evolution-like selection of fast-folding model proteins. *Proc Natl Acad Sci U S A* 92: 1282–1286.
 45. Bloom JD, Raval A, Wilke CO (2007) Thermodynamics of neutral protein evolution. *Genetics* 175: 255–266.
 46. Eigen M (1971) Self-organization of matter and the evolution of biological macromolecules. *Naturwissenschaften* 58: 465–523.
 47. Swetina J, Schuster P (1982) Self-replication with errors. A model for polynucleotide replication. *Biophys Chem* 16: 329–345.
 48. Fraser AG, Kamath RS, Zipperlen P, Martinez-Campos M, Sohrmann M, et al. (2000) Functional genomic analysis of *C. elegans* Chromosome I by systematic RNA interference. *Nature* 408: 325–330.
 49. Giaever G, Chu AM, Ni L, Connelly C, Riles L, et al. (2002) Functional profiling of the *Saccharomyces cerevisiae* genome. *Nature* 418: 387–391.
 50. Shakhnovich BE, Koonin EV (2006) Origins and impact of constraints in evolution of gene families. *Genome Res* 16: 1529–1536.
 51. van Nimwegen E, Crutchfield JP, Huynen M (1999) Neutral evolution of mutational robustness. *Proc Natl Acad Sci U S A* 96: 9716–9720.
 52. Wilke CO, Wang JL, Ofria C, Lenski RE, Adami C (2001) Evolution of digital organisms at high mutation rates leads to survival of the flattest. *Nature* 412: 331–333.
 53. Trifonov EN, Kirzhner A, Kirzhner VM, Berezhovsky IN (2001) Distinct stages of protein evolution as suggested by protein sequence analysis. *J Mol Evol* 53: 394–401.
 54. Williams PD, Pollock DD, Blackburne BP, Goldstein RA (2006) Assessing the accuracy of ancestral protein reconstruction methods. *PLoS Comput Biol* 2: e69. doi:10.1371/journal.pcbi.0020069
 55. Jordan IK, Kondrashov FA, Adzhubei IA, Wolf YI, Koonin EV, et al. (2005) A universal trend of amino acid gain and loss in protein evolution. *Nature* 433: 633–638.
 56. Hochachka PW, Somero GN (2002) Biochemical adaptation. Mechanism and process in physiological evolution. New York: Oxford University Press. 466 p.
 57. Shakhnovich E, Gutin A (1990) Enumeration of all compact conformations of copolymers with random sequence of links. *J Chem Phys* 93: 5967–5971.
 58. Miyazawa S, Jernigan RL (1996) Residue–residue potentials with a favorable contact pair term and an unfavorable high packing density term, for simulation and threading. *J Mol Biol* 256: 623–644.
 59. Dietmann S, Park J, Notredame C, Heger A, Lappe M, et al. (2001) A fully automatic evolutionary classification of protein folds: Dali domain dictionary version 3. *Nucleic Acids Res* 29: 55–57.
 60. Altschul SF, Madden TL, Schaffer AA, Zhang J, Zhang Z, et al. (1997) Gapped BLAST and PSI-BLAST: A new generation of protein database search programs. *Nucleic Acids Res* 25: 3389–3402.
 61. Holm L, Sander C (1998) Touring protein fold space with Dali/FSSP. *Nucleic Acids Res* 26: 316–319.
 62. Boeckmann B, Bairoch A, Apweiler R, Blatter MC, Estreicher A, et al. (2003) The SWISS-PROT protein knowledgebase and its supplement TrEMBL in 2003. *Nucleic Acids Res* 31: 365–370.
 63. Apweiler R, Attwood TK, Bairoch A, Bateman A, Birney E, et al. (2000) InterPro—An integrated documentation resource for protein families, domains and functional sites. *Bioinformatics* 16: 1145–1150.
 64. Mirkin BG, Fenner TI, Galperin MY, Koonin EV (2003) Algorithms for computing parsimonious evolutionary scenarios for genome evolution, the last universal common ancestor and dominance of horizontal gene transfer in the evolution of prokaryotes. *BMC Evol Biol* 3: 2.