Smoothness within ruggedness: the role of neutrality in adaptation

Martijn Huynen[*,†], **Peter Stadler**[‡,†] **and Walter Fontana**[‡,†,§]

[*] Los Alamos National Laboratory
Theoretical Biology and Biophysics, MS-K710
and Center for Nonlinear Studies
Los Alamos, NM 87545 USA

[‡] Institut für Theoretische Chemie
Universität Wien
Währingerstraße 17
A-1090 Wien, Austria

[†] Santa Fe Institute
1399 Hyde Park Road
Santa Fe, NM 87501 USA

[§]To whom reprint requests should be addressed

Abstract

RNA secondary structure folding algorithms predict the existence of connected networks of RNA sequences with identical structure. On such networks evolving populations split into subpopulations which diffuse independently in sequence space. This demands a distinction between two mutation thresholds: one at which genotypic information is lost and one at which phenotypic information is lost. In between diffusion enables the search of vast areas in genotype space, while still preserving the dominant phenotype. By this dynamics the success of phenotypic adaptation becomes much less sensitive to the initial conditions in genotype space.

Introduction

To explain the high fixation rate of nucleotide substitutions in a population, Motoo Kimura argued that the vast majority of genetic change at the level of a population must be neutral, rather than adaptive [1]. Sewall Wright's reaction to Kimura's point was politely neutral [2, page 474]: "Changes in wholly nonfunctional parts of the molecule would be the most frequent ones but would be unimportant, unless they occasionally give a basis for later changes which improve function in the species in question which would then become established by selection". Today, in view of the data generated by comparative sequence analysis, the surprise is no longer over the existence of neutrality, but over how little conservation there is at the sequence level [3, 4, 5, 6]. This makes Wright's point even more pertinent. How are we to imagine the relation between neutral evolution and adaptation? An answer to this question requires a model of the relationship between genotype and phenotype. Such a model is available for RNA secondary structure. The latter can be computed from the sequence by means of procedures based on thermodynamic data which have become standard in the past 15 years [7, 8]. Secondary structure covers the major share of the free energy of tertiary structure formation, and is frequently used to interpret RNA function and evolutionary data. As such the case is a qualitatively important one.

Robust Properties of RNA folding

The mapping from sequences to secondary structures is many-to-one for two reasons: first, there are many more sequences than secondary structures, second, some structures are realized much more frequently than others [9]. Call two sequences "connected", if they differ by one or at most two point mutations. A "neutral network", then, is a set of sequences with identical structure, such that each sequence is connected to at least one other sequence. The crucial point for our discussion comes from a recent study of the standard secondary structure prediction algorithm [9] which showed that such networks exist and that for frequent structures these networks percolate through sequence space. For example, starting at a sequence which folds into a tRNA structure it is possible to traverse sequence space along a connected path, thus changing every nucleotide position without ever changing the structure. Moreover, due to the high-dimensionality of sequence space, networks of frequent structures penetrate each other such that each frequent structure is almost always realized within a small distance of any random sequence. These features seem to be intrinsic to RNA folding, since they are insensitive to whether the folding algorithm is thermodynamic, kinetic or maximum matching [10] or whether one considers one minimum free energy structure or the entire Boltzmann ensemble [11].

A Simple Model for Test Tube Evolution

To assess the consequences of these properties for molecular evolution, we study a model in which the replication rate (fitness) of an RNA sequence depends on its secondary structure. Our folding procedure [12] is a speed-tuned implementation of the Zuker-Stiegler algorithm [8]. The model consists of a population of RNA sequences of fixed length ν which replicate and mutate in a stirred flow reactor. RNA populations manageable in the computer or in the laboratory are tiny compared to the size of the sequence space (4^ν), and a correct simulation must, therefore, resort to stochastic chemical reaction kinetics [13, 14]. A selection pressure is induced by a dilution flow which adjusts over time to keep the total RNA population fluctuating around a constant capacity N [15, 13]. This setup mimics Spiegelman's serial transfer technique [16], where sequences with a replication rate above (below) the average increase (decrease) in concentration.

When a sequence undergoes a replication, each base is copied with fidelity $1 - p$. The overall replication rate of an individual sequence is defined to be a function of the distance between its secondary structure and a predefined target structure. Here the target structure is the tRNA$^\text{Phe}$ clover leaf, but the structure of any randomly chosen sequence would do as well. This corresponds to the artificial *in vitro* selection of a structure with some desired function or affinity to a target [16, 17, 18, 19, 20, 21, 22, 23]. A similar situation, though not with RNA but proteins, occurs in the affinity maturation of the immune response [24]. In both artificial and natural selection there are two sources of neutrality: one is the sequence (genotype) to structure (phenotype) mapping, the other is the structure to replication rate (fitness) mapping. It is the former source which is central to this discussion. Notice, thus, that in the present model the second source of neutrality only arises for sequences whose structures differ from the target.

Genotypic versus Phenotypic Error Threshold

We first consider the case where the initial population consists of, say, $N = 1000$ copies of the tRNA$^\text{Phe}$ sequence ($\nu = 76$). Since that sequence folds into the target structure, no evolutionary optimization will take place. In this case the sequences which are most relevant to the population dynamics are those with structure-distance zero, i.e. those which fold exactly into the target structure. At a single-base error rate of $p = 0.001$ the population first broadens in sequence space as mutants accumulate, then starts drifting away and the initial sequence is lost. Any two populations which are separated in time by about 500 generations do not have a single identical sequence in common. The picture in structure space looks entirely different: the distribution of structures present at any time is firmly dominated by the original tRNA phenotype. In previous studies of evolutionary

dynamics which did not involve a genotype-phenotype model, the loss of sequence information at a critical error rate was identified with the loss of the dominant phenotype [25]. In the presence of percolating neutrality, however, all sequence information is lost at any error rate different from zero. Yet there is another threshold which we term the "phenotypic error threshold", p_c, beyond which the dominant phenotype is lost as well. That is when evolutionary adaptation breaks down. A linear extrapolation of the relation between the fraction of sequences that have the tRNA phenotype and the mutation rate (correlation coefficient 0.98) gave an estimate of $p_c = 0.0031$. This is consistent with findings that for $p_c = 0.003$ the tRNA phenotype remained present in the population for the length of the simulation (1460 generations), whereas for $p_c = 0.0035$ it disappeared shortly after the initiation of the simulation (within 200 generations).

Diffusion in Sequence Space

Figure 1(a) shows the mean squared displacement of the center of the RNA population in sequence space. It is linear for short times. This indicates that the aggregate population undergoes "Brownian Motion": the effects of mutations on the center of the population cancel out on average, and the overall displacement is generated by the variance. For a "flat landscape", where every sequence has the same replication rate a, the diffusion constant D can be approximated as follows. A simple model of our flow reactor consists of repeated events in which the replication of a randomly chosen sequence is followed by the removal of a randomly chosen sequence [26], thereby keeping the total population exactly constant. This is a first order Markov process with reflecting barriers for the number, $n_{i\alpha}$, of nucleotides of type α at position i. The diffusion coefficient for a population of N sequences of length ν replicating at rate a with mutation frequency p per nucleotide is $D_0 = (a\nu/N) \sum_\alpha \langle \delta n_{i\alpha}^2 \rangle$, where δ is the change in the state variable upon one replication and removal event. The calculation yields $D_0 = \frac{a\nu}{N} \left\{ \frac{2}{3}p(4F - 1) + 2(1 - F) \right\}$, where $F = (1/N^2) \sum_\alpha n_{i\alpha}^2$ is the probability that two randomly chosen sequences have the same nucleotide type at position i. The stationary F can be obtained without knowing the full solution of the stochastic process. A calculation along the lines in Ewens [27, page 81] yields for this process $F = (3 + Np)/(3 + 4Np)$ and, therefore, $D_0 = (6a\nu p[1 + 1/N])/(3 + 4Np)$, or

$$D_0 \approx \frac{6a\nu p}{3 + 4Np}.$$

For small mutation rates the diffusion coefficient, D, on the structure-dependent landscape can be approximated by $D = D_0\lambda$, where λ denotes the average fraction of neutral mutants for the dominant structure.

5

The diffusion of finite populations in sequence space observed here is a phenomenon which can be related to Kimura's neutral theory [28]. The latter stresses a different aspect: the number of nucleotide substitutions that reach fixation per generation, k, also referred to as the "rate of evolution". The theory yields $k = a\ p\ \nu\ \lambda$, independent of population size, and for small p ($p \leq 0.0005$) we find our model to agree with this very well. Thus, $D = 6\,k/(3 + 4Np)$.

In a laboratory setting, sequencing a diverse population yields a virtual consensus sequence. The squared distance between the consensus sequences at different times divided by the time interval yields an estimate for the diffusion coefficient. If the error rate is known, the average degree of neutrality can be determined. For example, from the $D(p)$-data in Figure 1(b) we obtain a λ of 0.27 ± 0.04 for small p. This agrees with independent calculations based on a random sample of sequences generated by inverse folding [9] of the tRNA structure (0.28 ± 0.06).

The Structure of Sequence Populations

Replication, mutation and selection on an extended neutral network induce a clustered structure of the population in sequence space (Figure 2). Analytical calculations by Derrida and Peliti [29] have predicted this effect for the extreme case where every sequence folds into the same structure ("flat landscape"). Randomly amplified neutral sequences become origins of temporarily localized clouds of related mutants. Clusters are defined with respect to a threshold Hamming distance d using the single-linkage method (nearest neighbor joining) [30]. For the average number of clusters Derrida and Peliti derived a $1/d$ scaling law for small d. On the neutral tRNA network we find a scaling relation of $1/d^{\alpha}$ with $1.7 < \alpha \leq 2$ depending on population size and error rate. The population structure, therefore, reflects qualitatively a flat landscape, but there is less fragmentation on average as a result of the network boundaries.

This population structure has profound consequences for evolutionary adaptation. All sequences compete for concentration shares in the chemostat, while at the same time the population splits into well-defined subpopulations in the abstract high-dimensional sequence space. These subpopulations share the same dominant phenotype, but undergo independent diffusion on the currently prevailing neutral network, thereby exploring vastly different parts of sequence space simultaneously. Connected neutrality endogenously induces a divide and conquer type of parallel local search. Yet selection is non-local in sequence space: once an entry point to a neutral network of a more favorable phenotype has been found, the population is quickly amplified around that point by adaptive selection.

6

Implications for Adaptation

The implications of neutrality for adaptive evolution are shown in Figure 3, where we monitor the evolutionary progression toward our tRNA target, starting from a homogeneous population consisting of a single random sequence. The average structure distance from the target shows plateaus separated by transitions. As opposed to the previous example (in which the target phenotype was present at the outset), every plateau is now characterized by more than one dominant phenotype, because of the degeneracy introduced by the structure-distance function. From [9] we know, however, that almost any structure assumed by a random sequence has its own extended neutral network. A diffusion process of the kind described above for the tRNA case will occur on the neutral network belonging to any such structure. Diffusion enhances the likelihood that a region in sequence space is found where the network of the currently dominant structure comes close to the network of a better structure. In this contact region a selection induced transition will occur like those seen in Figure 3. We believe, therefore, that the scenario of Figure 3 is typical even for the (unlikely) case of a fitness function which would assign a unique value to each structure.

Algorithms for folding RNA sequences into secondary structures predict extended connected networks of sequences with identical structure. The actual existence and extension of such networks must be assessed empirically. Such assessment pending, the present work predicts striking consequences for adaptation, based on a fairly realistic model of test-tube evolution. (1) Finite populations diffuse along neutral networks. After a sufficiently long period of time (set by the diffusion coefficient) all sequence information is lost, yet the phenotype is conserved. It is the maintenance of a phenotype, not of a genotype, which defines the mutation threshold beyond which adaptation breaks down. (2) On a single neutral network the population splits into well separated clusters. A population is not a single localized quasispecies in sequence space [25], but rather a collection of different quasispecies. Each undergoes independent diffusion, while all share the same dominant phenotype. (3) Neutral networks of different structures are interwoven. While drifting on a neutral network a population produces a fraction of mutants off the network and thereby explores new phenotypes. A selection induced transition between two structures occurs in regions of sequence space where their networks come close to one another. The independent diffusion of subpopulations increases the likelihood that a population encounters such transition regions. The success of an evolutionary search process for a target-phenotype is, therefore, substantially less dependent on the location in sequence space of the initial population. (4) The model invites to use care when characterizing the ruggedness of "fitness landscapes" by their correlation statistics. In high dimensional spaces, a small correlation length indicating a very rugged landscape with many local optima tells little about their connectedness. When connected, "local optima" are

7

no longer local, and this is reflected in the dynamics of adaptation.

References

[1] Kimura, M. *Nature* **217**, 624–626 (1968).

[2] Provine, W. B. *Sewall Wright and Evolutionary Biology.* The University of Chicago Press, Chicago, (1986).

[3] Pütz, J., Puglisi, J. D., Florentz, C., and Giegé, R. *Science* **252**, 1696–1699 (1991).

[4] Sassanfar, M. and Szostak, J. W. *Nature* **364**, 550–553 (1993).

[5] Schwienhorst, A. *Grundlagen einer evolutiven Biotechnologie.* PhD dissertation, Universität Göttingen, Germany, (1993).

[6] Komiyama, H., Miyazaki, G., Tame, J., and Nagai, K. *Nature* **373**, 244–246 (1995).

[7] Waterman, M. S. and Smith, T. F. *Math. Biosc.* **42**, 257–266 (1978).

[8] Zuker, M. and Stiegler, P. *Nucleic Acids Research* **9**, 133–148 (1981).

[9] Schuster, P., Fontana, W., Stadler, P. F., and Hofacker, I. L. *Proc. Roy. Soc. London* **B 255**, 279–284 (1994).

[10] Bornberg-Bauer, E., Tacker, M., and Schuster, P., unpublished data.

[11] Bonhoeffer, S., McCaskill, J. S., Stadler, P. F., and Schuster, P. *Eur. Biophys. J.* **22**, 13–24 (1993).

[12] Hofacker, I. L., Fontana, W., Stadler, P. F., and Schuster, P. RNA folding package available by anonymous ftp from ftp.itc.univie.ac.at in /pub/RNA.

[13] Fontana, W., Schnabl, W., and Schuster, P. *Phys. Rev. A* **40**, 3301–3321 (1989).

[14] Gillespie, D. T. *J. Chem. Phys.* **81**, 2340 (1977).

[15] Eigen, M. *Naturwissenschaften* **58**, 465–526 (1971).

[16] Mills, D. R., Peterson, R. I., and Spiegelman, S. *Proc.Natl.Acad.Sci.USA* **58**, 217–224 (1967).

[17] Levisohn, R. and Spiegelman, S. *Proc.Natl.Acad.Sci.USA* **63**, 807–811 (1969).

[18] Joyce, G. F. *Gene* **82**, 83–87 (1989).

[19] Robertson, D. L. and Joyce, G. F. *Nature* **344**, 467–468 (1990).

[20] Tuerk, C. and Gold, L. *Science* **249**, 505–510 (1990).

[21] Ellington, A. D. and Szostak, J. W. *Nature* **346**, 818–822 (1990).

[22] Beaudry, A. A. and Joyce, G. F. *Science* **257**, 635–641 (1992).

[23] Bartel, D. P. and Szostak, J. W. *Science* **261**, 1411–1418 (1993).

[24] Berek, C. and Ziegner, M. *Immunology Today* **14**, 400–404 (1993).

[25] Eigen, M., McCaskill, J. S., and Schuster, P. *Advances in Chem. Phys.* **75**, 149–263 (1989).

[26] Moran, P. A. P. *Proc.Camb.Phil.Soc.* **54**, 60–71 (1958).

[27] Ewens, W. J. *Mathematical Population Genetics.* Springer-Verlag, New York, (1979).

[28] Kimura, M. *The Neutral Theory of Molecular Evolution.* Cambridge University Press, Cambridge, (1983).

[29] Derrida, B. and Peliti, L. *Bull.Math,Biol.* **53**, 355–382 (1991).

[30] Sneath, P. H. A. *Journal of General Microbiology* **17**, 184–200 (1957).

[31] Havel, T. F., Kuntz, I. D., and Crippen, G. M. *Bull. Math. Biol.* **45**, 665–720 (1983).
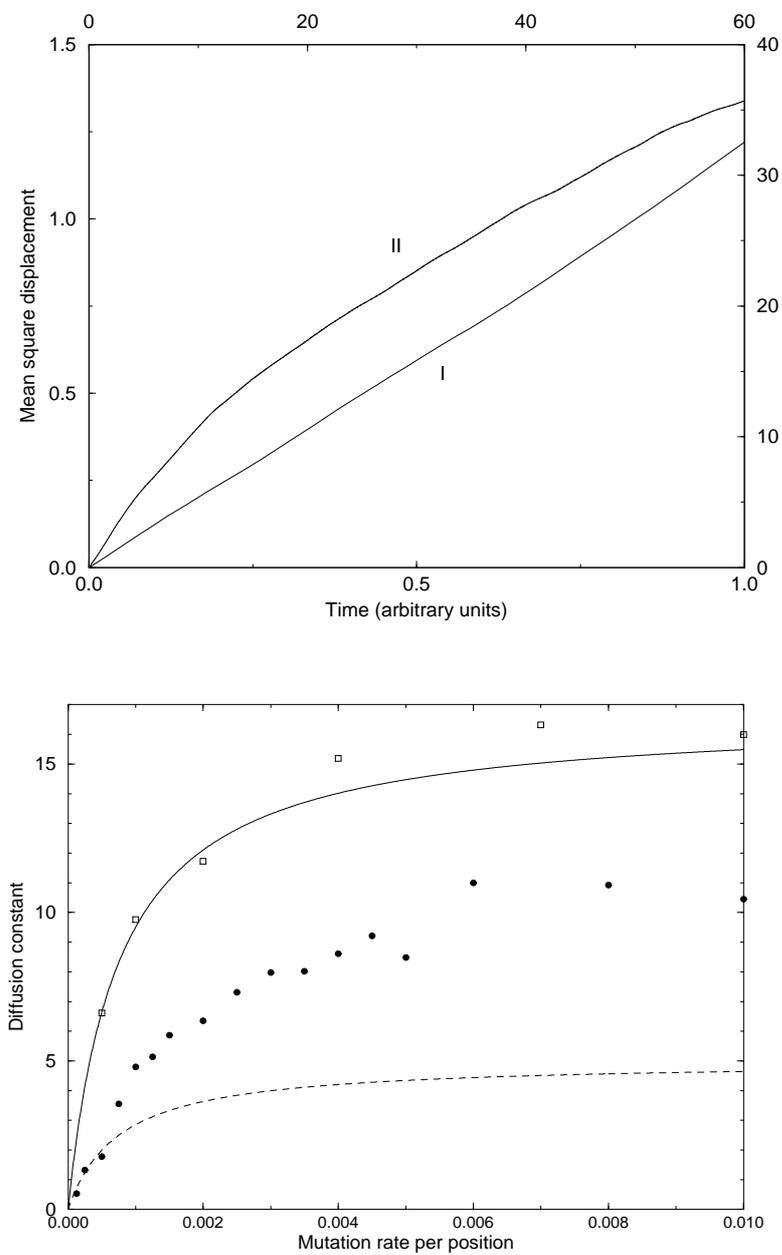
Figures and Legends

Figure 1: Diffusion in sequence space.

The center $\vec{C}(t)$ of a population at time t is a real valued consensus vector specifying the fraction $x_{i\alpha}(t)$ of each nucleotide $\alpha \in \{$A,U,G,C$\}$ at every position i. Figure 1a plots the mean squared displacement of the barycenter, $\overline{\Delta(\tau)^2}$, as calculated

from one simulation by averaging the expression $\Delta(t,\tau)^2 = (C(t+\tau) - C(t))^2 = \sum_i \sum_\alpha (x_{i\alpha}(t+\tau) - x_{i\alpha}(t))^2$ over t for fixed lag τ. $\nu = 76$, $N = 1000$, $p = 0.00025$, target structure is the secondary structure of tRNA$^{\text{Phe}}$. Curve II shows $\overline{\Delta(\tau)^2}$ in the τ-range from 0 to 60 reactor time units (upper abscissa, right ordinate). Curve I zooms into the time interval $[0,1]$ of curve II (lower abscissa, left ordinate). Due to the finiteness and the geometry of sequence space, the displacement must eventually level off to $3/4\,\nu$ (here 57), corresponding to the average distance between two random sequences.

Figure 1b shows the dependency of the diffusion coefficient D on the mutation rate per nucleotide, p. $D = \lim_{\tau \to 0} d\overline{\Delta(\tau)^2}/d\tau$ is calculated as the slope of the mean square displacement. Each data point (filled circles) is calculated from one computer simulation with $N = 1000$ which is run for at least 20 time units (a time unit corresponds to 146 generations) to allow for convergent averaging of the squared displacement. The solid line is the theoretical D_0 for a flat landscape where every sequence is assumed to replicate with the same rate as sequences with tRNA structure in our structure-dependent simulations; the open squares are simulations with our flow reactor on the flat landscape. The dotted line is $D_0\lambda$, where $\lambda = 0.3$ is the estimated fraction of neutral mutants.
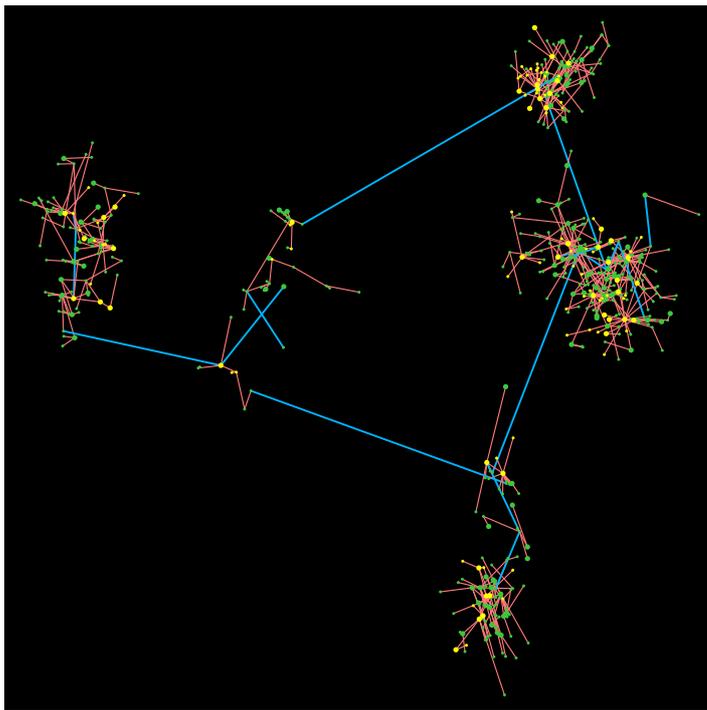

Figure 2: Population structure in sequence space.

The support of a population in sequence space is the set of sequences present in at least one copy. The population support can be pictured in two dimensions using

12

some theorems from distance geometry [31]. We compute the metric matrix M with entries $m_{ij} = (d_{0i}^2 + d_{0j}^2 - d_{ij}^2)/2$, where d_{ij} is the Hamming distance between sequences i and j and 0 is the "center of mass" of the support. Sequences are expressed in principal axes coordinates by diagonalizing M. Only the components corresponding to the largest two eigenvalues are kept, yielding a projection onto the plane that captures most of the variation. The dots in Figure 2 represent a static snapshot of $N = 2000$ individuals after 135 time units replicating with $p = 0.002$. Among the 2000 individuals 631 are different and among them 301 fold into different structures. To help correct for the distortions of the projection, the dots are connected by the edges of the minimum spanning tree. Edges connect closest points. Red (blue): Hamming distance less (more) than 6. Dot size large (small): more (less) than 4 copies in the population. Yellow (green): sequences which do (do not) fold into the tRNA target structure.
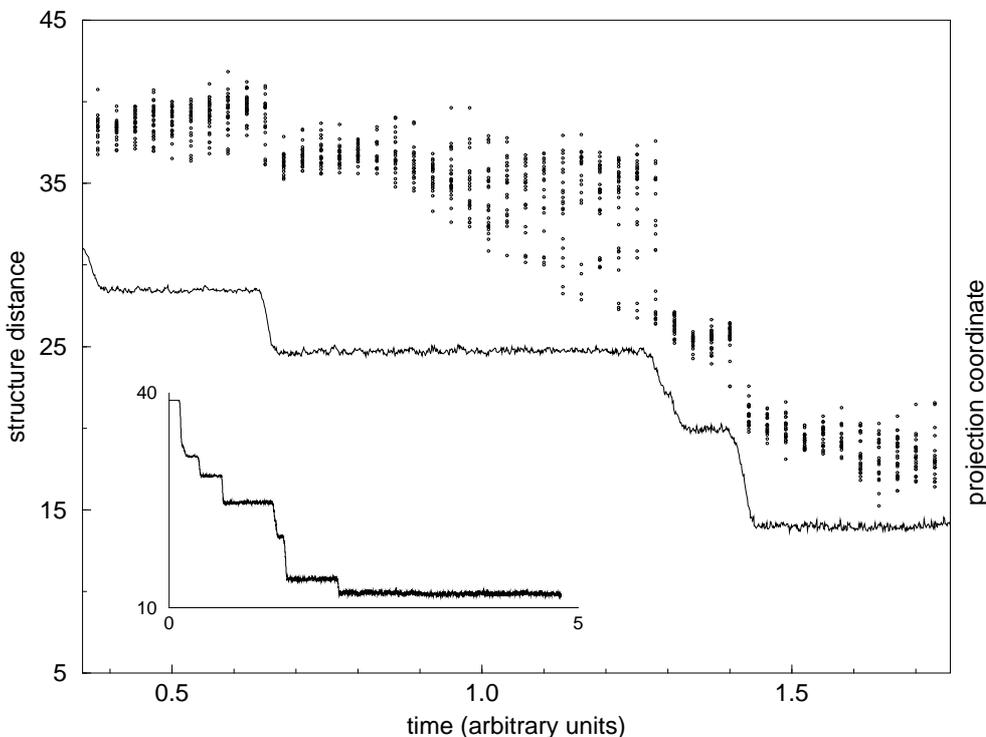

Figure 3: Evolutionary optimization.

A flow reactor with capacity $N = 1000$ is initialized with that many copies of a random sequence of length $\nu = 76$. The mutation rate is $p = 0.001$ and the target secondary structure is the tRNA$^{\mathrm{Phe}}$ cloverleaf, the replication rate function is $A(d) = 1.06^{146-d}$, where d is the tree-edit distance [9] to the target structure. The population average of the distance to the target is plotted against time (full line) for a specific interval of the entire run (shown in the inset). The superimposed

13

series of dots render the evolution of the population structure over time. The dots at one time slice are a one (!) dimensional projection (see caption to figure 2) of the population of sequences present in more than 10 copies at that time. Collecting all time slices yields a unique glimpse of the cluster dynamics.

The same qualitative picture of punctuated equilibria occurs with all parameter settings and random target structures we tried, for both linear and exponential fitness functions $A(d)$.