Continuity in Evolution: On the Nature of Transitions

Walter Fontana and Peter Schuster

Institut für Theoretische Chemie
Universität Wien
Währingerstrasse 17
A-1090 Wien, Austria

Santa Fe Institute
1399 Hyde Park Road
Santa Fe, NM 87501 USA

International Institute for Applied Systems Analysis
(IIASA)
A-2361 Laxenburg, Austria

Science, **280**, *1451–1455 (1998)*

additional material at http://www.santafe.edu/~walter/RNA/punct.html

Abstract

To distinguish continuous from discontinuous evolutionary change, a relation of nearness between phenotypes is needed. Such a relation is based on the probability of one phenotype being accessible from another through changes in the genotype. This is exemplified by calculating the shape neighborhood of a tRNA secondary structure, and provides a characterization of discontinuous shape transformations in RNA. The simulation of replicating and mutating RNA populations under selection shows that sudden adaptive progress coincides mostly, but not always, with discontinuous shape transformations. The nature of these transformations illuminates the key role of neutral genetic drift in their realization.

A much debated issue in evolutionary biology concerns the extent to which the history of life has proceeded gradually or has been punctuated by discontinuous transitions at the level of phenotypes (*1*). Our goal is to make the notion of a discontinuous transition more precise, and to understand how it arises in a model of evolutionary adaptation.

We focus on the narrow domain of RNA secondary structure, which is currently the simplest computationally tractable, yet realistic phenotype (*2*). This choice enables the definition and exploration of concepts that may prove useful in a wider context. RNA secondary structures represent a coarse level of analysis compared to the three-dimensional structure at atomic resolution. Yet, secondary structures are empirically well-defined and obtain their biophysical and biochemical significance from being a scaffold for the tertiary structure. For the sake of brevity we shall refer to secondary structures as "shapes". RNA combines in a single molecule both genotype (replicatable sequence) and phenotype (selectable shape), making it ideally suited for in vitro evolution experiments (*3, 4*).

To generate evolutionary histories we use a stochastic continuous time model of an RNA population replicating and mutating in a capacity constrained flow reactor under selection (*5, 6*). In the laboratory a goal might be to find an RNA aptamer binding specifically to a molecule (*4*). While in the experiment the evolutionary end product is unknown, we think of its shape as being specified implicitly by the imposed selection criterion. Since our intent is to study evolutionary histories rather than end products, we define a target shape in advance, and assume the replication rate of a sequence to be a function of the similarity between its shape and the target. An actual situation may involve more than one best shape, but this does not affect our conclusions.

Figure 1A shows an instance representing in its qualitative features all simulations we performed. Starting with identical sequences folding into a random shape, the simulation was stopped when the population became dominated by the target, here a canonical tRNA shape. The black curve traces the average distance to the target (inversely related to fitness) in the population against time. Aside from a short initial phase, the entire history is dominated by steps, that is, flat periods of no apparent adaptive progress, interrupted by sudden approaches towards the target structure (*7*). However, the dominant shapes in the population change not only at these marked events, but undergo several fitness-neutral transformations during the periods of no apparent progress. While discontinuities in the fitness trace are evident, it is entirely unclear when and on the basis of what the series of successive phenotypes itself can be called continuous or discontinuous.

A set of entities is organized into a (topological) space by assigning to each entity a system of neighborhoods. In the present case there are two kinds of entities: sequences and shapes, which are related by a thermodynamic folding procedure. The set of possible sequences (of fixed length) is naturally organized into a space, because point mutations induce a canonical neighborhood. The neighborhood of a sequence consists of all its one-error mutants. The problem is how to organize the set of possible shapes into a space. The issue arises, because, in contrast to sequences, there are no physical processes which directly (and inheritably) modify shapes. Rather, transformations of a shape are a complicated consequence of changes in its underlying sequence. To properly frame continuity in the spirit of topology, we must understand how one shape can be considered to be "near" some other. We may then call a temporal succession of sequence/shape pairs (an evolutionary path) continuous, if successive sequences are neighbors in sequence space and their corresponding shapes are neighbors in shape space. A topology is weaker than a metric, since the relation of nearness does not quantify distance (or similarity). We note this to emphasize that continuity does not hinge on the similarity of successive shapes in time. We next define and explore an appropriate relation of nearness for RNA shapes, and then return to the discussion of Figure 1.

For a shape β to succeed a shape α, β must obviously be accessible from α. Accessibility means that a sequence whose shape is β arises by mutation from a sequence whose shape is α. The issue of accessibility logically precedes any reasoning about the fitness of β, although fitness will strongly influence the fate of the mutant in a population under selection. We shall call a shape β "near" a shape α, if β is very likely to be accessible from α. The issue, then, becomes one of estimating the statistical frequency with which a mutation in α's sequence yields the mutant shape β. It is here that neutrality comes crucially into play (*8*). When a shape α is realized by a large class of sequences, "nearness" of β to α comes to mean that β must arise from α with a high probability when averaged over all sequences folding into α. Only then is the neighborhood of α a robust property of α itself, independent of a particular sequence.

This notion of neighborhood is illustrated by considering a tRNA-like shape of length 76 (*9*) (inset Figure 2A). A sample of the many sequences folding into this shape is obtained by an inverse folding procedure (*10, 11*). For every sequence in the sample we compute all shapes realized by its 228 one-error mutants (the sequence neighborhood). From this data we determine the fraction of sequence neighborhoods in which a mutant shape appeared at least once. The totality of these mutant shapes, irrespective of how often they occurred, is termed the (shape space) boundary of the tRNA.

When rank-ordering the boundary shapes with decreasing frequency, we obtain Figure 2A. The most salient feature is a marked change in the scaling exponent, suggesting a natural cut-off point for the definition of neighborhood. In the present case, the high frequency range comprises some 20 shapes, which we define to be near the tRNA shape (*12*). These shapes constitute the characteristic set of the tRNA, that is, its most specific neighborhood. The topmost 12 shapes are listed in Figure 2B, and exhibit two properties we found to hold for all shapes whose neighborhoods we studied. First, most shapes in the characteristic set of a shape α are highly similar to α, typically differing in a stack size by single base pairs (*13*). Second, some shapes, such as tRNA_8 (the shape ranked 8th in Figure 2B), differ by the loss of an entire stack. The latter finding illustrates that nearness does not imply similarity. More importantly, it illustrates that nearness is not a symmetric relation. In fact, the tRNA shape was not found in the characteristic set of the $tRNA_8$, and it did not even occur in its boundary sample. Not surprisingly, the destruction of a structural element through a single point mutation is easier than its creation. While the high frequency of the event is surprising, it is ultimately a consequence of the average base pair composition of stacks and the markedly different stacking energies of AU and GC base pairs (*12*).

The tRNA boundary has an intriguing property. Intersections with large samples of coarse grained random shapes of the same length support the conjecture that all common coarse grained shapes occur in the boundary of any common shape (*9, 14*). This conjecture was verified in the case of the exhaustively folded binary (GC-only) sequence space of length 25.

We may visualize the neighborhood structure (the topology) on the set of all shapes as a directed graph. Each shape is represented by a node. Directed edges fan out from a node α to the nodes in its characteristic set. We can think of a continuous transformation of shape α into shape β as a connected path in the graph which follows the direction of the edges. Discontinuous transformations are transitions between disconnected components of the graph.

The preceding data enable us to characterize continuous transformations as those structural rearrangements which fine tune a shape architecture in a sequential fashion by lengthening or shortening stacks, or which destroy a stack element and the loop implied by it (Figure 3). Discontinuous transformations are characterized by the two remaining possible structural changes: (i) the creation of a long stack in a single step, and (ii) generalized shifts (Figure 3). For example, one strand of a stacked region slides past the other by a few positions (simple shift). Notice here that structural similarity does not imply nearness. Both types of discontinuous transformations require the synchronous participation of several bases (or base pairs) in a fashion that cannot be sequentialized on thermodynamic grounds (*15*).

A pertinent issue is whether the folding map from sequences to shapes is continuous in our topology, that is, whether the shapes realized in the sequence neighborhood of a particular sequence folding into α are in the neighborhood of α. It turns out that the folding map is almost nowhere continuous. Many of the frequent shapes assumed by the one-error mutants of a sequence folding into α are not members of the characteristic set of α, and those who are don't always occur with high frequency. Each sequence folding into α has, therefore, its own specific set of accessible shapes. Yet, the local peculiarities disappear and a shape-specific neighborhood obtains when averaging over a sufficiently large sample of sequences folding into α.

Equipped with this fitness-independent notion of (dis)continuous shape transformations, we resume the discussion of Figure 1. To obtain an evolutionary path in shape space we record during a simulation all mutation events that produce a new shape. New means here that the shape is not present in the population at the time it is produced, although it may have been present in the past. For each shape ever seen we obtain a series of presence intervals delimited by the shape's entrance and exit times in the population. We define an evolutionary path, $\alpha_n\alpha_{n-1}\alpha_{n-2}\cdots\alpha_{i+1}\alpha_i\cdots\alpha_1\alpha_0$, retrospectively by searching in the history log for the shape α_{n-1} which first gave rise to the target shape α_n, next obtaining the shape α_{n-2} which started that presence interval of α_{n-1} during which α_n was produced, and so on until an initial shape α_0 is reached. This backtrack reconstructs the unique uninterrupted chain of shape innovations which led from an initial shape to the evolutionary end product. Note that this chain is neither defined with regard to fitness nor with regard to the frequency of a shape in the population (*16*). The path is continuous at the ith succession, if the sequences underlying the ith shape innovation differ by a single point mutation (which they typically do at high replication accuracy), and if shape α_{i+1} is near α_i in the sense defined above.

The evolutionary path (green trace) of Figure 1A comprises 43 shapes (not shown). Their presence intervals during the entire history are shown in red, one horizontal level for each shape. The patterns of presence intervals confirm and nicely visualize the nearness relation just developed. When a shape α is succeeded by a shape β that is near α, β is present intermittently in the population well prior to becoming part of the path (Figure 1B). That is, once α is present, β is unavoidable, and a transition to β is continuous. Conversely, at a discontinuous transition, when α is succeeded by a shape β that is not near α, β has almost always its first ever appearance just prior to that transition (Figure 1B). Seen together, the presence intervals of successive shapes on the path form blocks of continuous (within-neighborhood) transitions, separated by discontinuous transitions (neighborhood escapes).

In all computer simulations we observed a few basic patterns of events which combine to form particular histories. When starting with a random shape, there is a short initial phase of a few discontinuous transitions rapidly decreasing the distance to the target. This is understood by noting that many modifications of a random shape increase its similarity to a (random) target, and by recalling that such modifications are accessible in the local neighborhood of any random sequence (discontinuity of the folding map). Both properties effectively establish a funnel in shape space enabling fast relaxation to a level of similarity beyond which adaptation becomes harder. Then the character of evolutionary dynamics changes.

In the second phase, the population level (as monitored by distance to target) is entirely dominated by punctuation events. The point is that these events do mostly, but not always, line up with discontinuous transitions on the evolutionary path. In Figure 1A, events (a) and (b) are rapid (*17*) successions of continuous transitions shortening and elongating stacks by single base pairs. This shows that sudden changes in fitness do not imply discontinuous phenotypic transformations. The reverse isn't true either, as shown by the discontinuous shift event (c) which is silent in terms of fitness. All remaining fitness changes do, however, coincide with discontinuous transitions in shape space. These are the simple shift events (e), (g), (h), (i), the double flip (d), and the flip (f) (*18*). Interestingly, an ancestral shape which has been on the path in the distant past is reoccurring (but not on the path) several discontinuous transitions thereafter (events (j) in Figure 1A), arising by a single point mutation from shapes currently on the path. This is a molecular version of atavism.

Given its nature, a discontinuous transformation can be triggered by a single point mutation only if the rest of the sequence provides the appropriate context. Such sequences are severely constrained and hence rare. When a phenotype is under strong selection, neutral drift is the only means for producing the required genotypic context (*6, 19*). This is why discontinuous transitions are preceded by extended periods of neutral drift in Figure 1A.

The concept of evolutionary continuity cannot be separated from an understanding of the relationship between genotype and phenotype. It is indeed defined by it. A necessary step towards formalizing the concept of punctuated equilibrium is the study of the fitness-independent topological structure of phenotype space induced by the genotype-phenotype map. In a final analysis punctuation may turn out to be a phenomenon intrinsic to an evolving entity, and less dependent on external contingencies than hitherto assumed.

References and notes

1. N. Eldredge, S. J. Gould, in *Models in Paleobiology,* T. J. M. Schopf, Ed., (Freeman, San Francisco, 1972), pp. 82-115; S. J. Gould, N. Eldredge, *Nature* **366**, 223 (1993); S. F. Elena, V. S. Cooper, R. E. Lenski, *Science* **272**, 1802 (1996).

2. Let i, j, k, l denote positions of bases in the linear sequence and (i, j) a base pair. The secondary structure of an RNA sequence is defined as the set P of allowed base pairs (here Watson-Crick pairs plus GU) which minimize free energy, subject to a no-knot condition requiring that if (i, j) and (k, l) are both in P, then $i < k < j$ implies $i < l < j$ (i.e. base pairs don't cross). The secondary structure is computed with our implementation (*10*) of a dynamic programming algorithm, originally due to (*20*), which is widely used in laboratories to assist in the prediction of secondary structures. The procedure is based on empirical energy parameters (*21*).

3. S. Spiegelman, *Quart. Rev. Biophys.* **4**, 213 (1971); G. F. Joyce, *Gene* **82**, 83 (1989); A. D. Ellington, J. W. Szostak, *Nature* **346**, 818 (1990); A. A. Beaudry, G. F. Joyce, *Science* **257**, 635 (1992); D. P. Bartel, J. W. Szostak, *Science* **261**, 1411 (1993).

4. C. Tuerk, L. Gold, *Science* **249**, 505 (1990).

5. W. Fontana, P. Schuster, *Biophys. Chem.* **26**, 123 (1987).

6. M. A. Huynen, P. F. Stadler, W. Fontana, *Proc. Natl. Acad. Sci. USA* **93**, 397 (1996).

7. The same phenomenon has been observed in optimization problems of a quite different nature, such as the evolution of particle based computation in cellular automata using genetic algorithms (*22*).

8. The folding of RNA sequences is notoriously redundant (*23*).

9. In RNA folding some shapes are realized much more frequently than others (*11, 24*). They were termed common shapes (*25*), and our statements, based on statistics, can be expected to hold only for these common shapes.

10. I. L. Hofacker et al., *Mh. Chem.* **125**, 167 (1994). A public domain version of the Vienna RNA package is available from http://www.tbi.univie.ac.at.

11. P. Schuster, W. Fontana, P. F. Stadler, I. L. Hofacker, *Proc. Roy. Soc. London* **B 255**, 279 (1994).

12. W. Fontana, P. Schuster, Santa Fe Institute Working Paper #97-11-081 (Santa Fe Institute, Santa Fe, 1997).

13. The shape α itself occurs in every sequence neighborhood of the α-sample (omitted from figure 2). This reflexivity of the nearness relation is the topological way of expressing neutrality.

14. Coarse grained shapes are derived from secondary structures by ignoring the size of stacks and loops, keeping only their relative arrangement. Our tRNA boundary sample (see caption to Figure 2A) contained 5,882 coarse grained shapes. A pool of 11,000 random sequences yielded 1,578 distinct coarse grained shapes, 90.4% of which were found in the tRNA boundary.

15. An example for a discontinuous transition of type (i) is the formation of a multiloop (a loop issuing more than two stacking regions). Generally, the free energy gain upon formation of a stack must offset the free energy loss from the loop caused by it. A stack closing a multiloop must, therefore, come into existence with some minimum length (typically more than 5 bp) in a single step. Likewise, the discontinuity of generalized shifts (type ii) has thermodynamic and structural origins. Shifting a stack by sequentially shifting its base pairs in random order would cause severe sterical conflicts, besides violating the formal no-knot condition. As a consequence, the shifting of a stack requires that all base pairs move synchronously.

16. Most, but not all, phenotypes on the path are highly populated. A path inferred from the fossil record almost certainly misses the low populated ones.

17. Note that the average number of replication events per time unit depends on the average replication rate constant in the population. The plateaus preceding events (a) and (b) have a duration comparable to those preceding events (g) and (h), but during the former about 4,300 replications occur per time unit, while during the latter this number rises to 10,800.

18. Discontinuous transitions may trigger a cascade of continuous events. On a few occasions continuous transtitions hitchhike on discontinuous ones. For example, a major rearrangement, such as a double flip, may involve the simultaneous elongation of a stack formed in the event.

19. M. A. Huynen, *J. Mol. Evol.* **43**, 165 (1996).

20. R. Nussinov, G. Piecznik, J.R. Griggs, D.J. Kleitman, *SIAM J. Appl. Math.* **35**, 68 (1978); M. S. Waterman, *Advances in Mathematics. Supplementary studies* **1**, 167 (1978); M. Zuker, P. Stiegler, *Nucleic Acids Research* **9**, 133 (1981).

21. D. H. Turner, N. Sugimoto, S. Freier, *Annual Review of Biophysics and Biophysical Chemistry* **17**, 167 (1988); A. Walter et al., *Proc. Natl. Acad. Sci. USA* **91**, 9218 (1994).

22. J. P. Crutchfield, M. Mitchell, *Proc. Natl. Acad. Sci. USA* **92**, 10742 (1995).

23. J. Pütz, J. D. Puglisi, C. Florentz, R. Giegé, *Science* **252**, 1696 (1991); M. Sassanfar, J. W. Szostak, *Nature* **364**, 550 (1993); H. Komiyama, G. Miyazaki, J. Tame, K. Nagai, *Nature* **373**, 244 (1995).

24. W. Grüner et al., *Mh. Chem.* **127**, 355 (1996).

25. P. Schuster, *Biophys. Chem.* **66**, 75 (1997).

26. Financial support was provided by the Austrian Fonds zur Förderung der wissenschaftlichen Forschung (Projects P-10578 and P-11065), by IIASA Laxenburg, Austria, by the Commission of the European Union (Contract Study PSS*0884), and by the integrative core research at the Santa Fe Institute.

Figure captions

Figure 1

(A): Simulation of an RNA population evolving towards a tRNA target shape (inset of Figure 2A) in a flow reactor logistically constrained to a capacity of 1,000 sequences on average. The replication accuracy per position is 0.999. The replication rate (=fitness) of a sequence whose shape is α is given by $(0.01 + d(\alpha, \mathrm{tRNA})/l)^{-1}$, where $l = 76$ is the sequence length and d the distance between α and the target. Linear or exponential functions did not affect the character of the dynamics. The initial population consisted of 1,000 identical sequences folding into a random shape. The target was reached after approximately $11 \cdot 10^6$ replications. The black trace shows the average structure distance of the shapes in the population to the target. The chain of shape innovations linking the initial shape to the target (evolutionary path) comprises 43 shapes. To each of these corresponds one horizontal level placed above the black curve. The topmost (bottom) level belongs to the initial (target) shape. For these levels only the time axis has a meaning. At each level a series of red intervals represents the time periods during which the corresponding shape was present in the population. The green step curve indicates the transitions between shapes, and hence the time spent by each shape on the evolutionary path. Each transition was caused by a single point mutation in the underlying sequences. The vertical dotted lines and the labels mark transitions referred to in the text.
(B): Enlargement of the evolutionary path around event (e) of part (A). The transition indicated on the left (α to β) is continuous. This is seen by the fact that β is present (intermittently) in the population well before becoming a link in the evolutionary path (green trace). In other words, β's presence is stochastically correlated with that of α, because it is near α in shape space. The intersecting neighborhood disks (see Figure 3) illustrate schematically that the continuous transition from α to β stays within the neighborhood of α. In contrast, the transition from β to γ is discontinuous, as seen by the fact that γ's presence does not correlate with β's (mutants of sequences folding into β don't typically fold into γ). Here γ has a fitness advantage and almost immediately becomes the next link in the evolutionary path. Note that β remains intermittently present after γ's takeover. This is because β is near γ, despite the fact that γ is not near β. The topological relationship of nearness need not be symmetric.

Figure 2

(A): Rank-ordered frequency distribution of shapes in the tRNA boundary. A sample of 2,199 sequences whose minimum free energy secondary structure is a tRNA cloverleaf (inset) was generated. All their one-error mutants (501,372 sequences) were folded. 28% of the mutants retained the original structure (i.e. were neutral). The remaining 358,525 sequences realized 141,907 distinct shapes. The frequency $f(\alpha)$

is the number of one-error neighborhoods in which α appeared at least once, divided by the number of sequences in the sample. The log-log plot shows the rank of α versus $f(\alpha)$. Rank n means the nth most frequent shape. The dotted line indicates a change in the slope which we take to naturally delimit the high frequency domain (to the left) whose shapes form the characteristic set of the tRNA. (B) lists the 12 highest ranked shapes (left to right, top to bottom) in the characteristic set.

Figure 3

The strings illustrate transformations between RNA secondary structure parts. Solid (dashed) arrows indicate continuous (discontinuous) transformations in our topology. Three groups of transformations are shown. Top: the loss and formation of a base pair adjacent to a stack are both continuous. Middle: the opening of a constrained stack (e.g. closing a multiloop) is continuous, while its creation is discontinous. This reflects the fact that the formation of a long helix between two unpaired random segments upon mutation of a single position is a highly improbable event, whereas the unzipping of a random helix is likely to occur as soon as a mutation blocks one of its base pairs. Bottom: generalized shifts are discontinuous transformations in which one or both strands of a helix shift ending up with or without an overlap. Accordingly, we partition generalized shifts into the four classes shown. The intersecting disks are a schematic representation of continuous and discontinuous transitions between two shapes α and β. The disk with center α (β) stands for the set of shapes that are near α (β). If β is a member of α's disk (neighborhood), a transition from α to β is continuous (solid arrow). A discontinuous transition leaves the neighborhood of α (dashed arrow). Note that even if α and β are highly dissimilar, α might nontheless be transformed continuously into β through intermediate shapes whose neighborhoods have sufficient overlap.

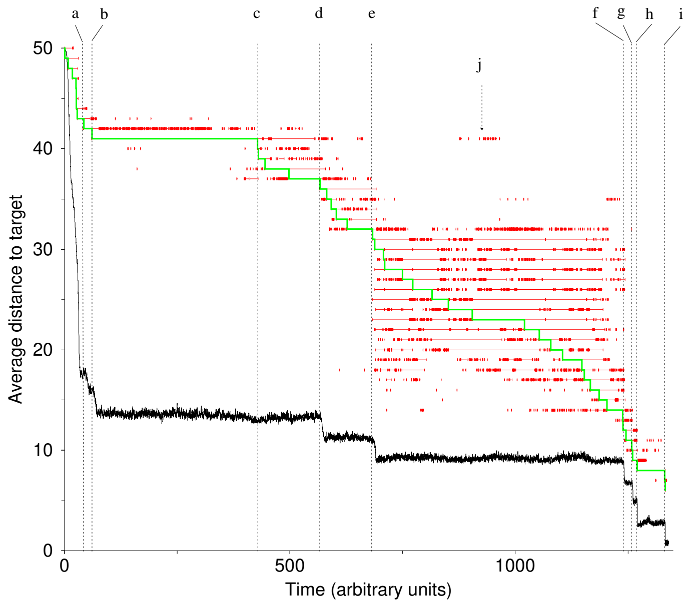
a
b
c
d
e
f
g
h
i
j
Average distance to target
0
10
20
30
40
50
0
500
1000
Time (arbitrary units)

(A)

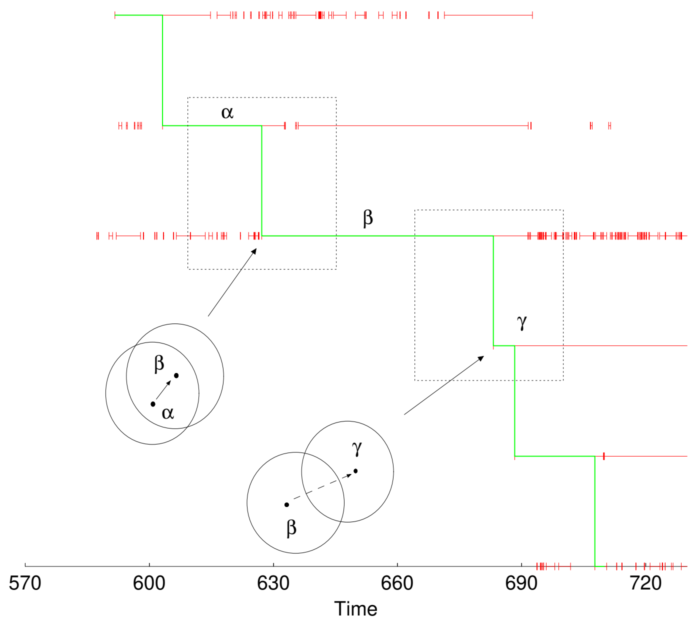
α
β
γ
570
600
630
660
690
720
Time

(B)

Figure 1

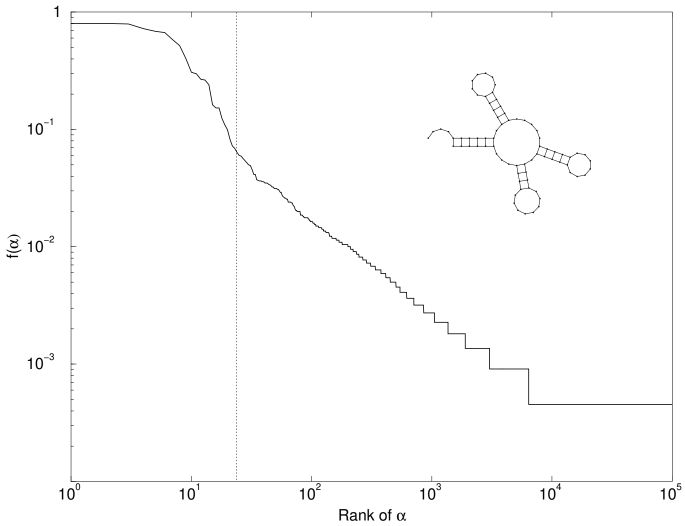

(A)

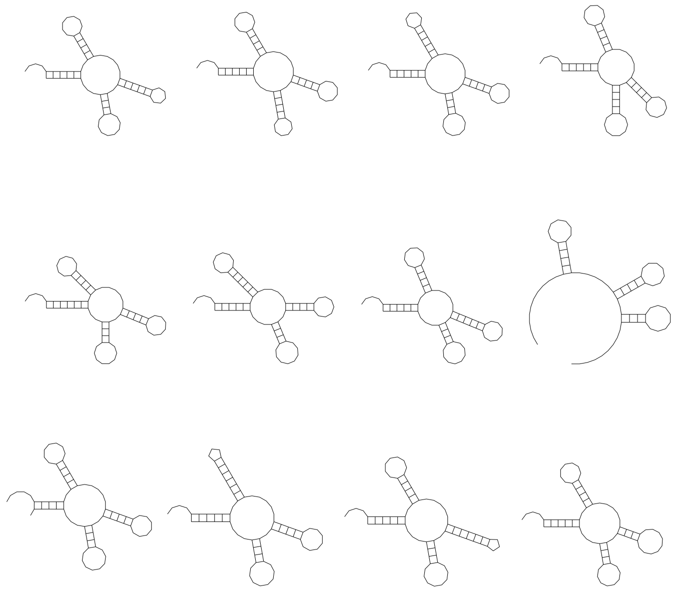

(B)

Figure 2

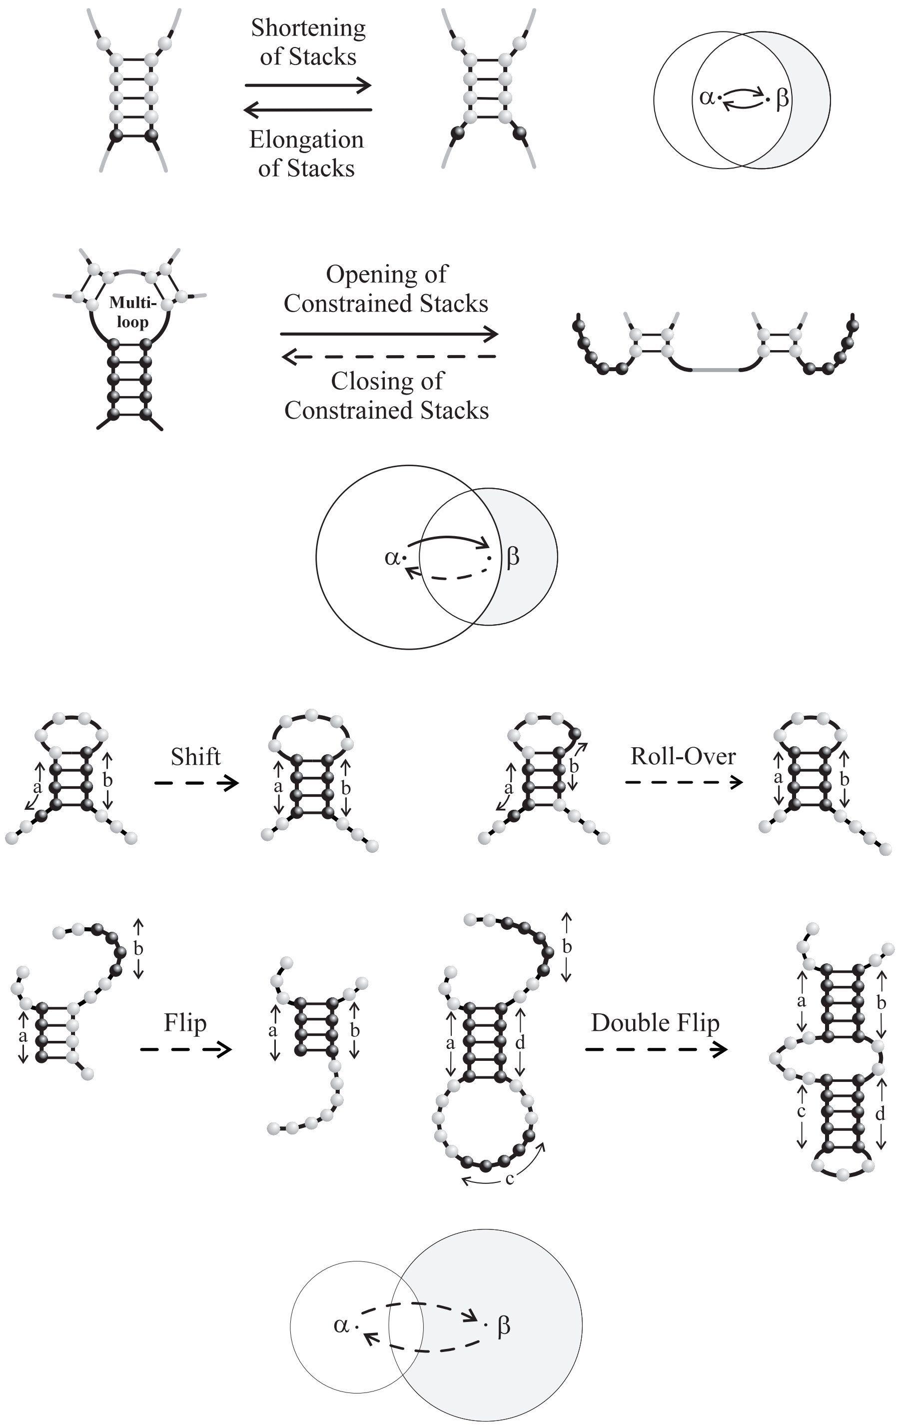
Shortening
of Stacks
Elongation
of Stacks
α
β
Multi-
loop
Opening of
Constrained Stacks
Closing of
Constrained Stacks
α
β
Shift
Roll-Over
a
b
Flip
Double Flip
a
b
c
d
α
β

Figure 3