

AUTOCATALYTIC REPLICATION OF POLYMERS

J. Doyne FARMER

Center for Nonlinear Studies, MS B258, Los Alamos National Laboratory, Los Alamos, NM 87545, USA

Stuart A. KAUFFMAN

Department of Biochemistry and Biophysics, School of Medicine, University of Pennsylvania, Philadelphia, PA 19104, USA

and

Norman H. PACKARD

The Institute for Advanced Study, Princeton, NJ 08540, USA

We construct a simplified model for the chemistry of molecules such as polypeptides or single stranded nucleic acids, whose reactions can be restricted to catalyzed cleavage and condensation. We use this model to study the spontaneous emergence of autocatalytic sets from an initial set of simple building blocks, for example short strands of amino acids or nucleotides. When the initial set exceeds a critical diversity, autocatalytic reactions generate large molecular species in abundance. Our results suggest that the critical diversity is not very large. Autocatalytic sets formed in this way can be regarded as primitive connected metabolisms, in which particular species are selected if their chemical properties are advantageous for the metabolism. Such autocatalytic sets may have played a crucial role in the origin of life, providing a bridge from simple molecular species to complex proteins and nucleic acids. Many of our results are experimentally testable.

1. Introduction

The origin of life poses fundamental problems concerning the prebiotic origins of molecular species. Present life forms contain a rich set of chemical constituents which regenerate themselves through processes of replication, transcription, and translation, relying heavily on the collaboration of proteins and nucleic acids. Although contemporary life may have evolved from earlier, simpler forms, there seems to be a minimum level of complexity below which life based on templating is not possible. How, then, were the conditions needed to achieve life based on templating ever reached?

The experiments of Miller and Urey [1] and others indicate that it is possible to form amino acids and other small metabolites from simpler constituents in a fairly direct way, but it is a long

way from there to the complex proteins and nucleic acids needed for contemporary life. The possibility that we study here is that the origin of life came about through the evolution of autocatalytic sets of polypeptides and/or single stranded RNA. By *autocatalytic set* we mean that each member is the product of at least one reaction catalyzed by at least one other member. Our central thesis here is that templating is not required to achieve an autocatalytic set. Instead, simple polymers can catalyze the formation of each other, generating autocatalytic sets that evolve in time to create complex chemical species whose properties are tuned for effective collaboration with each other. The system thus bootstraps itself from a simple initial state to a sophisticated autocatalytic set, which might be regarded as a precursor life form. The sets provide the rich substrate of raw materials and coupled catalytic relationships needed for

the origin of contemporary life. The importance of autocatalytic properties has also been discussed by Calvin [2], Eigen [3], Eigen and Schuster [4, 5], Kauffman [6, 7], Rossler [8, 9, 10], and Dyson [11, 12].

The molecular species comprising our autocatalytic sets may be any polymers that can undergo catalyzed cleavage and condensation reactions. Given a sufficiently diverse supply of simple monomers and small polymers as raw materials, catalyzed reactions between them can form more complex molecular species. The new species in turn serve as substrates and catalysts for further reactions, increasing the complexity of the system until an autocatalytic set is created. Catalyzed pathways from the "food set" maintain the members of the autocatalytic set. The result is a metabolism converting simple molecules into complex ones. Different catalytic pathways within the network compete with each other, enhancing the concentration of some species at the expense of others. The set thus evolves in time, creating a rich but focused collection of molecular species whose cooperative chemical properties make them more fit than others. Such autocatalytic sets might serve as a bridge from simple building blocks such as amino acids and nucleotides to contemporary life based on a genetic code.

Experiments have shown that ligation (end condensation) and cleavage reactions among polypeptides may be catalyzed by other polypeptides [13, 14, 15]. More complicated reactions may also be catalyzed, e.g. transpeptidation reactions, where pieces of two polypeptides switch places. These more complicated reactions may, however, be regarded as compositions of ligation and cleavage. Thus, there is good experimental evidence indicating that the type of catalyzed reactions we study here occur for polypeptides.

In addition, recent experiments have shown that certain single stranded RNA sequences can be reproduced autocatalytically. It has been conjectured that this mechanism might be extended to include arbitrary sequences [16]. If this were true,

then single strands of RNA would be capable of catalyzing each other in precisely the fashion that we model here. Our model is general, describing any family of polymers undergoing catalyzed cleavage and condensation reactions.

What we do here can be viewed as a preliminary attempt to simulate and understand the qualitative behavior of a model chemistry. Simple constituents forming a food set are pumped into a stirred tank which is allowed to overflow. Our results demonstrate that what happens is critically dependent on the properties of the food set. Below a critical complexity very little happens. Above this critical complexity, however, a chain reaction is triggered that generates a rich autocatalytic set.

The model chemistry we construct here also presents an interesting problem in dynamical systems theory. For a fixed set of chemical species we model the dynamics in terms of ordinary differential equations describing the change in concentration of each species. The set of chemical species can change in time, however, so that the equations themselves are dynamic.

Rossler first stressed the qualitative dynamics of autocatalytic networks that grow in complexity [8, 9, 10]. Even though the system is always finite at any given time, it potentially explores an infinite dimensional space. The system evolves by changing its equation of motion. These changes can take place over a very long time scale. Such systems provide a novel type of dynamical system which has received very little attention.

The work described here has practical motivations as well. Current molecular cloning techniques now make it possible to generate very large numbers of novel DNA molecules in expression vectors, hence to clone very large numbers of novel genes, and study their RNA or protein products for catalytic or other activities. Such experiments are now underway, and may ultimately lead to actual construction of autocatalytic polymer systems, possibly with useful commercial properties. Our work may be viewed as a feasibility study for this process. Our results suggest

that there is a minimum critical complexity of the food set and catalytic properties which must be achieved if such efforts are to succeed.

Basic aspects of the model discussed here have already been introduced elsewhere [6, 7].

2. Description of the model

Real chemical reactions among large sets of peptides or oligonucleotides are too complicated to have any hope of simulating in detail because of limitations on both our current knowledge of biochemistry and of current computer technology. Any simulation must use a drastically simplified model of chemical reactions. The goal in constructing such models is to preserve the qualitative features that are most important in generating robust reproduction of what is typically seen in experiments.

The polymers under study are assumed to be one dimensional chains, represented by one dimensional strings over an n letter alphabet, of the form s_1, s_2, \dots . The letters s may represent different amino acids, or alternatively, different nucleotides. The strings are oriented from left to right. In proteins this corresponds to the orientation from the amino terminal on one end to the carboxy terminal on the other end, and for single stranded RNA it corresponds to the asymmetry between the 5' and 3' ends. The polymers are allowed to have any length up to a given maximum.* Two types of reactions are allowed: End joining reactions (or ligation), i.e., gluing two polymers together, and cleavage reactions, i.e., splitting a larger polymer into two pieces. Since these reactions are reversible, the presence of a given end joining reaction implies that the corresponding cleavage reaction is also present. Other types of reactions, such as exchange, can be represented as combinations of these two.

*We used 1000 as the maximum length in our simulations, but this has very little effect on the results.

For simplicity, at this stage we consider only catalyzed reactions. Uncatalyzed reactions may ultimately play an important role in allowing "tunnelling" to new species, but in view of the enormous discrepancy in reaction rates, it seems plausible to ignore them for the moment. Thus in order for a reaction to occur at all, it must have at least one catalyst present in the system. Our reactions are of the form.



where c is the concatenation of a and b , e is the enzyme catalyzing the reaction, and h represents water, which plays an important role in setting the equilibrium concentrations.

As illustrated in fig. 1, the list of species together with their reactions can be visualized as a graph, with arrows connecting the two cleavage products a and b and the condensate c to a nodal point, corresponding to the reaction. The network of catalysts can be thought of as a superimposed graph connecting the catalysts to the reaction nodes. These graphs can change in time as new species are created or old species are eliminated from the system. Assignment of chemical properties amounts to giving a procedure for determining the reaction and catalysis graphs, together with associated kinetic properties such as reaction velocities, equilibrium constants, etc. In the following section we address the problem of constructing reaction graphs, deferring the problem of establishing the chemical kinetics to section 5.

3. Graph theoretic results

Producing a chemically reasonable way to assign reaction graphs is the most difficult aspect of this problem, and the place where the greatest assumptions must be made. The complication comes about because of the enormous difficulty in predicting the chemical properties of a polymer directly from its chemical composition. This is

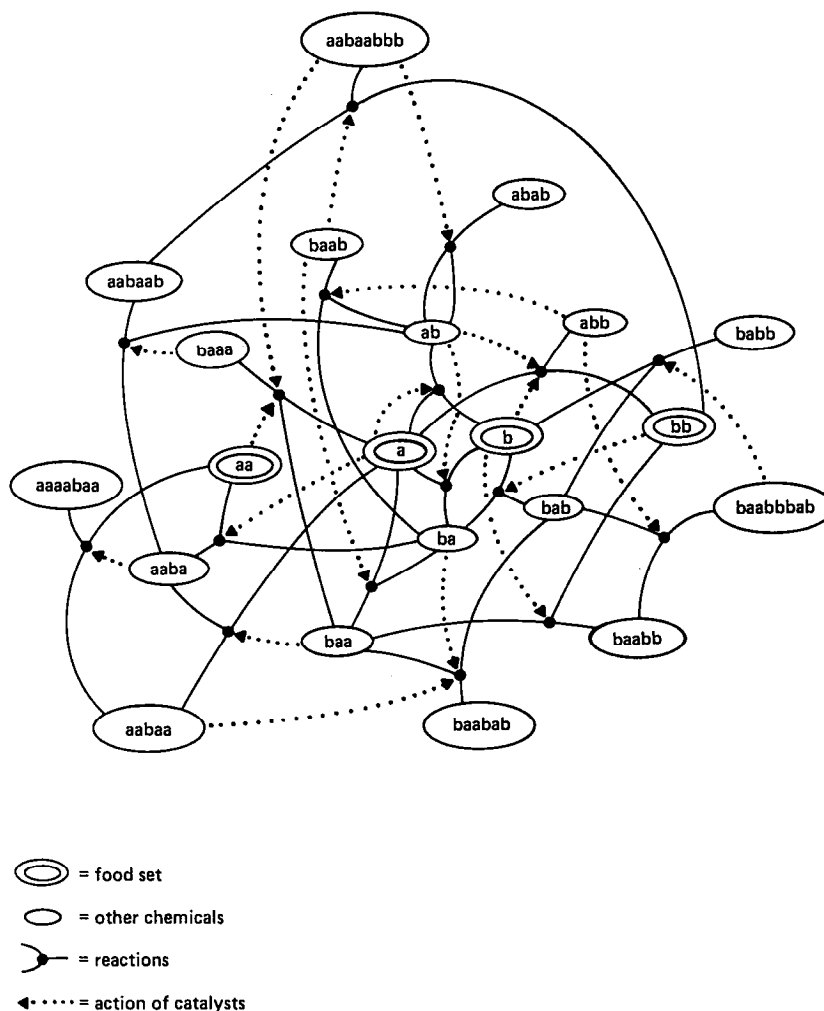


Fig. 1. A typical example of a graph that might describe an autocatalytic set. The reactions are represented by nodes connecting cleavage products with the corresponding condensate. Dotted lines indicate catalysis pathways, and point from the catalyst to the reaction being catalyzed.

particularly true of polypeptides. Although two polymers might have quite similar composition, differences in conformation and structure can cause their chemical properties to be quite different. One approach is to assume that this is so complicated that the reactions catalyzed by each polymer can be assigned at random. This method has unrealistic aspects, since it is clear that if an enzyme catalyzes a given reaction, in spite of all

the complication it is still likely to catalyze similar reactions. To explore the extent to which the basic results of the model depend on the details of this assumption, we are investigating other approaches, for example match rules [17]. This will be presented in a future paper. We ultimately hope that the salient properties of the behavior can be characterized by a few basic parameters, and will otherwise be independent of the details.

Our rule for random assignment of reactions is implemented as follows: For a given starting list of molecular species, we compute the maximum number of allowed condensation and cleavage reactions by counting the number of distinguishable combinations of string concatenations and string cleavages. The number of reactions that we actually assign is obtained by multiplying the number of allowed reactions of each type by a probability P . To assign condensation reactions, we chose two molecules at random, while for cleavage reactions we chose a molecule and a cleavage point at random. In both cases enzymes are chosen at random from the set of species currently present. Kinetic properties are also assigned at random, as discussed in the next section.

Assignment of reactions can be viewed as a dynamical process. We initialize the system by choosing a starting list, called the “firing disk”, typically chosen to be all possible strings shorter than a given length L . Reactions within the firing disk are assigned as described above. Condensation reactions may generate new species outside the firing disk, thereby expanding the list. The introduction of new species creates new reaction possibilities; to take these into account, on the next time step we count the number of combinatorial possibilities involving the new species. Multiplying by P gives the number of new reactions. This process is repeated on subsequent time steps. As long as new species are created on each step the graph continues to grow; otherwise growth stops.

One of the immediate questions that we wish to address is, “How many species do we need in the food set in order to generate an autocatalytic network?” Whether or not a graph grows indefinitely depends strongly on P . For sufficiently small P , it is virtually guaranteed to stop growing, and the graph is *subcritical*. For sufficiently large P , it is virtually guaranteed to grow without bound, and it is *supercritical*. The value of P on the boundary between subcritical and supercritical behavior is the *critical* value.

Although the reaction graphs studied here are strongly non-isotropic, an idea of the behavior that we should expect as P is varied can be obtained through a comparison with results from the theory of isotropic random graphs. Let R be the ratio of the number of edges (arrows) to the number of nodes. When R is small most of the nodes are isolated, i.e., they do not connect with any edges. For $R \approx 0.5$, most of the nodes become connected into one large cluster, and for $R \geq 1$ cycles of all lengths begin to form. Since R is proportional to P , we expect that as P increases the reaction graph will undergo a transition from a sparsely connected system with many isolated components to a richly connected system with one large component predominating.

For any given probability P , there is always a critical size for the firing disk so that the ratio R becomes large enough that the graph is highly connected, and thus supercritical. As the length of a polymer increases, the number of possible ways to cleave it also increases. As a result, the number of possible reactions that it participates in goes up. As already shown [7], in the limit of large sets the ratio of the number of reactions to the number of species goes as L , the maximum size present in the firing disk. As a result, for sufficiently large firing disks the reaction graph is supercritical.

The critical transition is illustrated in fig. 2, where we plot the number of new species created as a function of time, using typical parameter values. For the subcritical case, shown in fig. 2a, the number of new species decays exponentially until the graph stops growing. As P is increased the decay time increases, until eventually it reaches a critical point where it goes to infinity. Fig. 2b illustrates a case near the critical point; N_t exhibits significant fluctuations, and on any given run the graph may either stop growing or grow without bound. If P is increased still more, the graph becomes supercritical, as illustrated in fig. 2c; after an initial period of decay, the graph grows at a rate that is faster than exponential. For parameter values sufficiently near critical, the be-

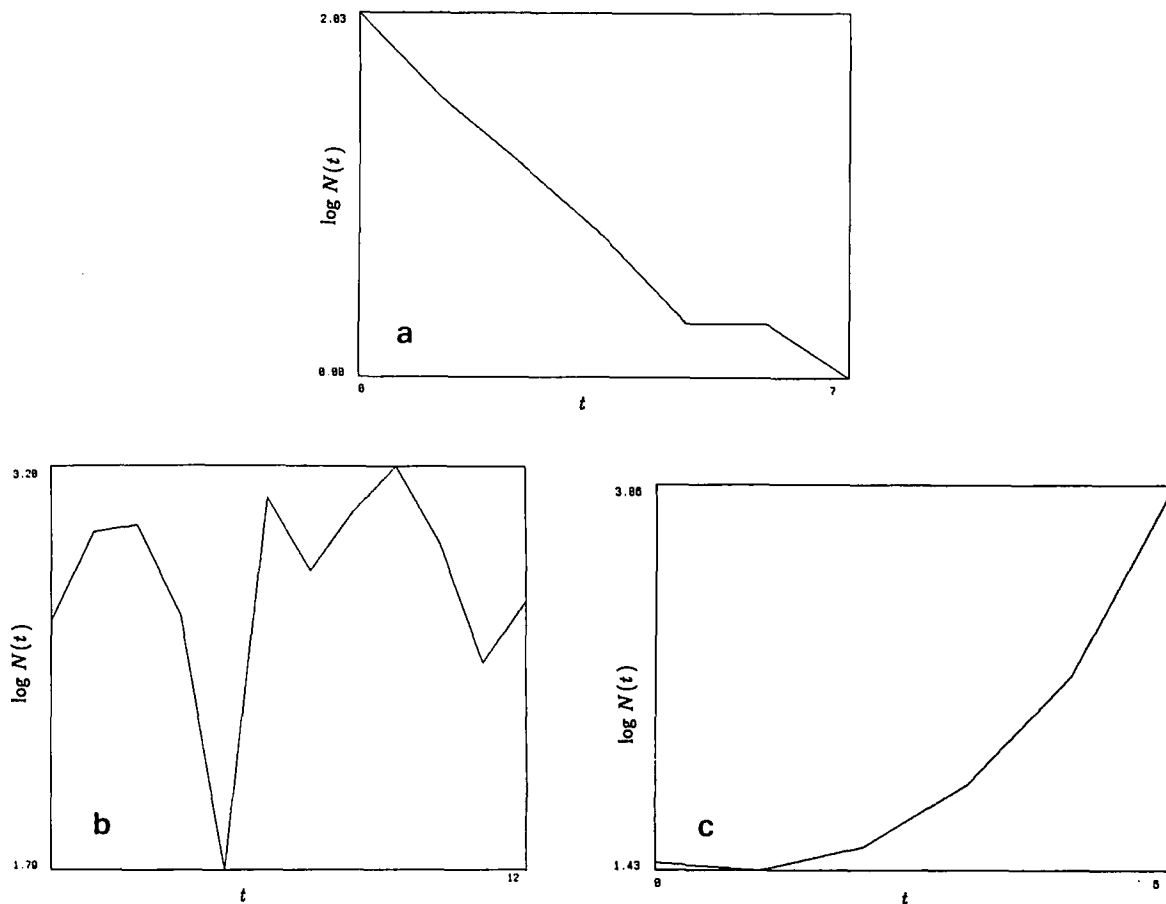


Fig. 2. The number of new molecular species created at each iteration during the generation of a random reaction graph. All cases illustrated used $B = 2$, and $L_f = 6$. (a) Below the supracritical point the graph decays. (b) Near the supracritical point, the graph size behaves erratically. (c) Above the supracritical point, the graph size increases without bound. Note the semi-log scale.

havior seen in these figures is typical*. Since the assignment of reactions depends to some extent on random roundoff, re-seeding the random number generator may result in slightly different behavior. It is even possible that on one run a graph is subcritical, while on another it is supracritical, illustrating that the critical point is only defined in

*Far from criticality, for low P the graph will immediately stop growing, while for very high P it will immediately start growing without any initial decay.

an average sense. Nonetheless, our results are generally quite well behaved, and the deviations from the mean behavior are not unreasonably large.

The critical surface can be estimated numerically using a simple trial and error procedure. For a given alphabet size B and a given firing disk radius L , values of P are found, P_{\max} and P_{\min} , so that P_{\max} is clearly supracritical and P_{\min} is clearly subcritical. The midpoint is tested; if supracritical, it replaces P_{\max} ; if subcritical, it replaces P_{\min} . This process continues until a value of P is found where the graph grows so slowly that it is effec-

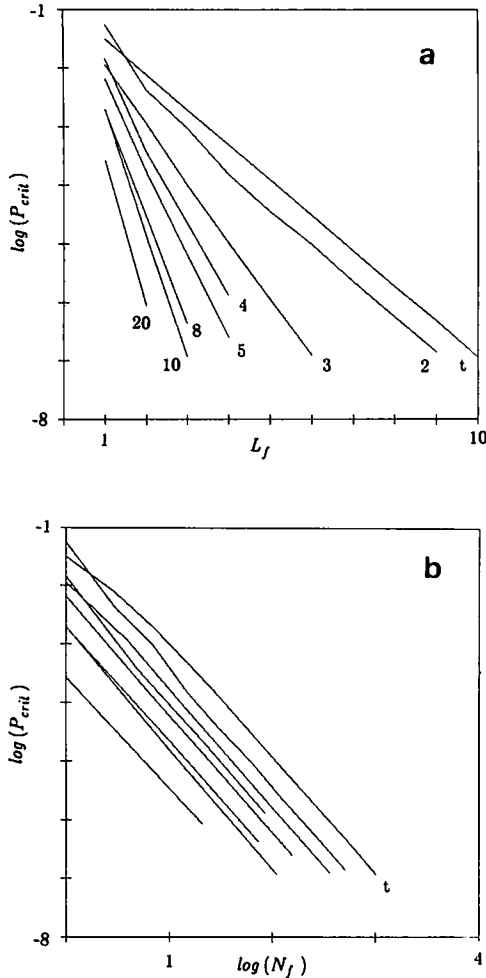


Fig. 3. (a) Dependence of P_{crit} , the probability of catalysis at the supracritical threshold, on L_f , the maximum length of the molecular species in the initial firing disk. The alphabets have different numbers of letters, as indicated. The line labeled t is from the theoretical results outlined in appendix A. (b) As in (a), except that the L_f is replaced by $\log(N_f)$, where N_f is the total number of species in the firing disk.

tively very close to critical. The results obtained with different values of B and L are given in fig. 3.

These results agree quite well with a previous estimate of the critical value of P , based on combinatorial arguments [7]. In Appendix A this calculation, which was previously made only for the case where $B = 2$, is extended to arbitrary

values of B . The central result is that criticality occurs roughly when

$$P \approx B^{-2L}. \quad (2)$$

Since the number of molecular species present in the firing disk approaches B^L for B large, the implication of eq.(2) is that the critical point occurs when the firing disk has the order of the square root of $1/P$ species.

Fig. 3. illustrates the dependence of P_{crit} , the critical probability of catalysis on the initial food set. Fig. 3a shows $\log(P_{\text{crit}})$ as a function of L_f , the size of the firing disk, including a comparison with the theoretical curve obtained in appendix B for alphabet size $B = 2$. Fig. 3b shows a plot of $\log(P_{\text{crit}})$ as a function of $\log(N_f)$, where $N_f = \sum_{i=1}^{L_f} B^i$ is the number of species in the critical firing disk, for a variety of different values of B , illustrating that $\log(P_{\text{crit}})$ scales linearly with N_f , the number of different species in the initial set, independent of the alphabet size.

Another interesting property of the system is the distribution of the number of species with a given length. For the firing disk, the number of species with a given length L is B^L . The distribution function is therefore exponentially increasing with length, with a sharp cutoff. Figs 4a and 4b illustrate two successive iterations of the graph update procedure. The underlying form of this distribution is revealed by plotting on a semi-logarithmic scale, as shown in fig. 4c, which demonstrates that the tail is roughly exponential. Similar results are observed at other parameter values, providing the total number of different species is large enough to produce good statistics.

The "sharp edge" of the firing disk can cause a ringing phenomenon for small polymer lengths, as illustrated in fig. 4d. After the first iteration, the distribution is extended from the firing disk by a factor of two in length, with the basic shape repeated. The repetition occurs because (for $B = 2$) the longest length contains half of the total number of species, and the condensation reactions between the longest lengths dominate the next




Fig. 4. Distribution of polypeptides as a function of their length. (a) and (b) illustrate the action of one iteration of the graph update procedure. (c) Same as (b), but with a semi-log plot, indicating the exponential tail. (d) Length distribution for a system with only 2 symbols in the alphabet, which causes a ringing at short lengths.

step, causing a peak at $2L$. (This effect is even larger for $B > 2$, when the dichotomy between the edge and center of the disk is even more extreme.) Successive iterations repeat this process, generating more peaks. After enough iterations have taken place, however, the peaks begin to overlap, and the distribution develops a smooth tail, as illustrated in fig. 4c.

The important point of these results is that the number of simple species needed to generate a supracritical graph is not very high – with a twenty letter alphabet, for example, a firing disk containing all possible monomers and dimers is sufficient

to generate a supracritical graph with a probability of catalysis of only 10^{-8} . Of course, this depends significantly on the random assignment of catalytic properties. It seems likely, however, that use of alternate schemes such as matching will only improve matters.

P is difficult to estimate in real systems. Perhaps the best estimate can be obtained from studies of the immune system, where it is known that an antibody recognizes a random antigen with a probability of roughly 10^{-5} [18]. Since antigen antibody recognition is a reasonable model for enzyme ligand binding at a catalytic site, this gives

some feeling for the probability that a randomly chosen polypeptide catalyzes a given reaction. The estimate is dangerous because antibodies may constitute a highly evolved set, specialized for ligand binding, and because ligand binding is not the same as catalysis. Another indication can be obtained from the work of Fox [19], which suggests that weak catalytic properties in abiogenically derived protenoids are not rare.

4. Extending the model to include kinetics

The existence of a supracritical graph is not sufficient to guarantee the generation of an autocatalytic set. Once chemical kinetics are taken into account competition for resources limits growth. Differences in efficiency can produce drastic differentials in concentration, so that many species are effectively not present in the system at all. Thus, even though it complicates matters considerably, a consideration of kinetics is essential for a realistic assessment of the potential to generate autocatalytic sets.

To simulate the influx and efflux from a chemostat, we assume that the raw materials from the foodset are supplied at a constant rate. For convenience, unless otherwise noted we set the food set equal to the firing disk. To do this we initialize the system by setting the concentration of all species of length less than L_f to a nonzero value, and setting all others to zero, thereby defining the food set. We then drive the system by adding to the concentration of the members of the food set at a constant rate. All molecular species are removed with the same first order loss rate, as would be expected, for example, from letting the chemostat overflow at a constant rate. After the passage of sufficient time, with this type of driving and damping the set of species present in the system becomes fixed, and the concentration of each approaches a constant value.

The fact that the system has a finite mass imposes a minimum concentration, corresponding to a single molecule. For most purposes concentrations significantly below this level can be ignored. This is an important practical benefit, since it allows us to restrict the simulation to a finite number of different species. Physically the threshold sets the effective size of the system. If the chemostat has roughly the volume of a bacterium, for example, one molecule per bacterium corresponds to a concentration of about 10^{-9} molar. In our simulations we are often forced to use higher thresholds, which can be thought of as decreasing the size of the container. For very small systems it might be more appropriate to use a stochastic molecular model rather than continuous concentrations, but we doubt that this will change our results very much.

A simplifying assumption that we have made is that we have a well-stirred chemostat. This is not necessarily the case. Spatial structure might well develop in a primordial soup, which could considerably complicate the problem. In this case, the size parameter (which depends on the threshold) should be adjusted to the volume of the typical region over which mixing and chemical interactions occur. In more realistic simulations we should model enclosure in a lipid boundary (liposome) or a protenoid microsphere [19], with a radius of 1–20 microns (which corresponds roughly to the size of a bacterium). Diffusion across volumes in this size range is fast enough to ensure near well stirred conditions.

Updating of the reaction graph proceeds as described in the previous section, but with a few important differences. When a reaction generates a new species, it is not immediately added to the system as before, but is rather placed on a list of *candidates*. In solving the differential equations for concentration updating, candidates are treated the same as other species, but they are not allowed to catalyze new reactions until their concentration crosses the threshold, at which point they are fully admitted to the system. In the simulations re-

ported on here we have ignored the possibility that a species crosses the threshold and then drops below it again; we do not remove reactions when this takes place.

We now discuss the differential equations that we use to model the kinetics. A fully realistic model is intractable for our purposes, since it requires keeping track of many intermediates for each reaction, resulting in an enormous multiplication in both the number of variables and the number of equations. The common approximations that are normally employed for catalyzed reactions are of no help, since there is no clear separation of the molecular species into products, enzymes, and substrates: A given species can play all three of these roles in different reactions. The result is that for a given reaction none of the components are known to definitely be either saturated or unsaturated.

To cope with this problem we have developed a new scheme for approximating the behavior of catalyzed reactions. This scheme makes one major approximation, namely, that the binding velocity is the same for all intermediates. This assumption is not accurate, but we do not think that it makes an important difference in our results. We allow the reaction velocities due to catalysis to vary from reaction to reaction, and since the binding velocity takes on meaning only in relation to the other velocities of the system, we do have the diversity in relative reaction rates which seems to be the qualitatively important factor.

With this assumption the concentration of each species can be separated into two parts, the *free* part, and the *bound* part, which is the sum of the concentrations of all intermediates in which a given species is bound. This doubles the number of variables, but this is a minor complication compared with keeping track of all intermediates separately. The resulting dynamical equations are derived in detail in Appendix B. The results can be summarized as follows: For every catalyzed reaction of the form of eq. (1), let x_a and x_b denote the free concentrations of the products a

and b, x_c denotes the free concentration of the condensate c, and x_e denotes the free concentration of the enzyme. Bound concentrations are denoted similarly as \bar{x}_a, \bar{x}_b , etc. This reaction has a forward velocity

$$v_f = v_0 k_f x_e x_a x_b \quad (3)$$

and a backward velocity

$$v_r = v_0 k_r x_e x_c h,$$

where k_f is the forward equilibrium constant, associated with the condensation reaction, k_r is the backward equilibrium constant, associated with the cleavage reaction, v_0 characterizes the increase in velocity caused by catalysis for a given reaction, and h is the concentration of water. Each reaction changes the free concentration of a, b, c, and e, according to the rules $\dot{x}_a = -v_f$, $\dot{x}_b = -v_f$, $\dot{x}_c = -v_r$, and $\dot{x}_e = -v_r - v_f$. When this is combined with the effects of driving and damping, the total change in concentration for any given species i is

$$\dot{x}_i = \sum [\text{reaction terms}] + k_u \bar{x}_i - k_d x_i + rH(L_f - L), \quad (4)$$

where k_u is a ‘‘dissociation constant’’ or ‘‘unbinding velocity’’, which is inversely proportional to the typical time that a complex remains bound in an intermediary, and k_d is the dissipation constant reflecting chemostat overflow. The last term is the driving term, which is nonzero only for members of the food set; l is the length of species i , L_f is the radius of the food set, r is the rate of concentration input, and H is the Heaviside function, $H(x) = 1$ for $x > 0$, $H(x) = 0$ for $x < 0$. The change in bound concentrations for each reaction can similarly be written as $\dot{\bar{x}}_a = v_r$, $\dot{\bar{x}}_b = v_r$, $\dot{\bar{x}}_c = v_f$, and $\dot{\bar{x}}_e = v_r + v_f$, so that the total change in the bound concentration for any given species can be written as:

$$\dot{\bar{x}}_i = \sum [\text{reaction terms}] - k_u \bar{x}_i - k_d \bar{x}_i.$$

In the simulations reported here, we set $h = 1$, $k_f = 1$, $k_r = 10$, $k_u = 1000$, and v_0 was randomly chosen in the range $10 < v_0 < 1000$. The first three values are reasonable approximate averages for the forward and reverse rates of peptide bond formation in an aqueous medium, where cleavage is favored [20].

5. Results incorporating kinetics

In this section we describe how the purely graph theoretical results are modified when kinetics is taken into account. Certain properties are obvious. For example: The fact that the system has a finite mass automatically limits the growth of the reaction graph. Even if the graph is strongly supracritical, there is a maximum size beyond which it simply cannot grow, due to the fact that resources are finite. Thus, in a strict sense there is no such thing as supracritical behavior when kinetics are considered. There does, however, remain a marked qualitative distinction between supracritical and subcritical behavior, in terms of the size to which the system grows (cf. fig. 5).

Kinetics introduces an important third parameter that influences the critical behavior, namely, the concentration threshold required for catalysis. The limiting cases are clear: As the threshold goes to zero, we return to the purely graph theoretic case. At the other extreme, regardless of how supracritical the purely graph theoretical case is, it is always possible to set the threshold high enough that no new reactions are ever catalyzed. In between these two extremes the number of species initially increases, but then levels off as it exhausts the capacity of its resources.

Fig. 5. gives an overview of the effect of the threshold concentration and the probability of catalysis on the total number of polymers formed from the food set. For zero threshold, the system reduces to the case of graph iteration considered above, where the critical probability of catalysis is known. Fig. 5 shows, for low thresholds (10^{-13}), that as the probability of catalysis increases past




Fig. 5. A contour plot of the logarithm of the asymptotic graph size asymptotic graph size as a function of P , the probability of catalysis and $\log(T)$, the logarithm of the concentration threshold. All of the simulations in this figure were made using a 2 letter alphabet, with a firing disk of radius 2 and with driving and damping constants both set at 10,000. Contours B through H represent graph sizes ranging from 7 to 10,000.

this value, the total number of polymers shows a sharp increase. For larger probabilities of catalysis, the system forms a large number of polymers. The finite number of polymers formed will reflect three factors: (i) Constant influx of polymers in the food set and first order loss of all molecular species assures that the system will relax to a steady state with constant total mass. (ii) The dissociation constant favoring cleavage over condensation in an aqueous environment implies that on average larger polymers are present in lower concentrations. (iii) Consequently, the existence of a threshold for catalysis limits the number of polymers that can come to function as catalysts, and therefore limits the growth of the graph. This effect is seen even for low values of the threshold, where as the probability of catalysis increases beyond the critical value for supracritical behavior, the rate of growth of the asymptotic number of polymers in the system slows. As the threshold increases, fewer polymers are able to function as catalysts. Therefore, the probability of catalysis

must increase to achieve the same asymptotic number of polymers.

The main conclusion to be drawn from these experiments is that the chemistry specified initially must have a minimum complexity for active autocatalytic behavior to take place, where complexity of the chemistry is determined by the size of the initial food set and by the probability of catalysis. There is a minimum size of the system below which autocatalytic behavior cannot take place, given a fixed probability of catalysis. Kinetics tends to inhibit the formation of a large reaction network, which decreases the extent of autocatalytic behavior. For a sufficiently large system, however, this is not a problem.

6. Future work

The results that we have reported here represent a first step toward an understanding of autocatalytic sets. A great deal remains to be done. To make more contact with experiments, and to investigate the robustness of the ideas presented here, we intend to build in a much higher degree of realism than exists in the present model. Planned enhancements include:

1) A string match model for assignment of catalytic properties, such as in ref. [17]. It is plausible that the probability of catalytic activity increases with polymer length. This feature is not present in the purely random model, but is easily incorporated using string matching. For realism, the details of the match procedure will need to be specialized for either polypeptide or RNA catalysis. The string match approach provides a good contrast to the random assignment scheme used here, and should give us a good feeling about the robustness of the behavior.

2) Inclusion of "high energy" compounds such as ATP, whose exergonic hydrolysis is coupled to endergonic peptide bond synthesis, thus driving synthesis of large peptides. In the absence of such high energy compounds, the concentrations of small and large polymers heavily favors small

polymers due to thermodynamic considerations reflected here in the higher rate of cleavage over condensation in an aqueous environment.

3) A more detailed study of the loop properties in the system.

4) A treatment of uncatalyzed as well as catalyzed reactions. Although uncatalyzed reactions proceed much more slowly, they may play a crucial role in seeding catalyzed reactions. This can dramatically lower the critical threshold.

5) If a set is indeed autocatalytic, there should be a sufficient number of connected pathways within it so that it can sustain itself as long as there are pathways to the food set. To test this, we intend to demonstrate that once the firing disk initiates growth, the autocatalytic set persists even after the food set is contracted down to a smaller number of elements.

6) A study of the evolutionary capacity of autocatalytic sets in competition with each other.

7. Summary

We have extended the previous analysis of autocatalytic sets [7] in several respects: Using a random scheme for assignment of catalytic properties, we have confirmed the existence of a critical transition from a (subcritical) graph with very few connections, to a (supercritical) graph which is highly connected. When it is subcritical, the graph growth rate decays exponentially; when it is supercritical, the growth rate may initially decay, but eventually becomes faster than exponential. After the passage of sufficient time, the supercritical graph apparently exhibits universal properties, such as an exponential tail in the frequency distribution of lengths. We have verified a previous estimate for the location of the critical surface in parameter space, and extended it to alphabets with more than two letters.

Based on a new scheme for dealing with saturation effects in networks of catalytic reactions, we have begun a numerical study of the effect of kinetics in autocatalytic sets. While kinetics trun-

cates supracritical behavior, the effects of graph supracriticality are evident in the size and richness of the reaction graph that develops with a concentration threshold.

Our results suggest that autocatalytic properties may have played a major role in supplying the complex chemical prerequisites needed for the origin of life. The scenario that we are suggesting proceeds as follows: When the concentration and variety of simple polypeptides exceeds a critical threshold, an autocatalytic network is triggered, causing the formation of a rich set of proteins. Similarly, short strands of nucleic acids might generate autocatalytic networks, catalyzing the formation of a rich set of RNA molecules. Eventually these two networks begin interacting with each other, making use of cross-catalytic properties to form an even richer set. Throughout this process, chemical kinetics naturally selects polymers with the most efficient properties for cooperating in the metabolic pathways of the autocatalytic network; the species which emerge from this competition are effectively the "fittest". Finally, in the presence of this rich set of highly evolved raw materials, the nucleic acids begin templating the formation of proteins and each other, creating the basis for life as we now know it.

Acknowledgements

We would like to thank Ron Fox and Brian Goodwyn for valuable discussions.

Appendix A

Supracritical firing disc scaling law

The aim of the appendix is to derive relations between the probability that an arbitrary peptide catalyses an arbitrary reaction, P , and the maximum length polymer maintained in an initial set of polymers of length $L = 1, 2, \dots, L_f$, such that that set of polymers generates an infinite catalysed reaction graph.

We derive results initially for $B = 2$ kinds of amino acid monomers, and restrict attention to end condensation and cleavage reactions, as defined in the text. The set of $2^{(L_f+1)}$ polymers is called the "firing disc". Since this initial set is maintained, the initial generation of novel peptides derives only from end condensation of preexisting peptides to form peptides of length $L_f + 1, L_f + 2, \dots, 2L_f$. Any specific peptide of length $L', L_f < L' \leq 2L_f$ can be formed in

$$Y = 2L_f + 1 - L' \quad (\text{A.1})$$

ways via end condensation. The number of such reactions among the polymers of length L' is

$$2^{L'} [2L_f + 1 - L']. \quad (\text{A.2})$$

Each of the possible new condensation reactions might be catalysed by any of the members of the initial firing disc, with probability P , thus the expected number of new condensation reactions which are catalysed, Q_1 is

$$Q_1 = P 2^{L_f+1} \sum_{L'=L_f+1}^{2L_f} 2^{L'} [2L_f + 1 - L']. \quad (\text{A.3})$$

If the total number of catalysed reactions given in (A.3) is small with respect to the possible number of new peptides to length $2L_f$, then Q_1 is a good estimate of the number of novel peptides whose formation is catalysed by members of the firing disc. The criteria we develop below for supracritical behavior based on this assumption is self-consistent in satisfying this constraint.

There are Q_1 novel peptides that are available to participate in new reactions either as substrates or catalysts. We discriminate 4 cases: (i) The Q_1 may catalyse reactions among members of the firing disc whose large product is of length L_f or less, hence within the firing disc. We ignore these since the firing disc is maintained. (ii) The Q_q can catalyse any of the condensation reactions among members of the firing disc as substrates, forming products of length $L_f + 1$ to $2L_f$. (iii) The Q_1 can participate as substrates in end condensation reactions on the left or right terminus of any of the $Q_1 + 2^{L_f+1}$ old and new members of the peptide set now in the catalysed reaction graph. These

novel condensation reactions might be catalysed by any of the old and new members of the reaction graph. (iv) The Q_1 are substrates for new cleavage reactions, catalysed by any of the old or new members of the reaction graph.

Case ii, Catalysis of old condensation reactions

The number of novel peptides formed by this mechanism on the second iteration is

$$Q_2 = Q_1^2 / 2^{L_r+1}. \quad (\text{A.4})$$

Case iii, New condensation reactions

The number of new peptides after the first iterate is Q_1 with lengths up to $2L_r$. On a second iterate, the maximum length achievable increases to $4L_r$. It is not necessary to keep track of the size distribution in detail, however. End condensation of any of the Q_1 can with one another, or on each end of the initial peptides, hence by this mechanism the number of new condensation reactions is:

$$Q_1^2 - Q_1 + 2Q_1 2^{L_r+1} > Q_1 [Q_1 + 2^{L_r+1}]. \quad (\text{A.5})$$

Each of these might be catalysed by any of the total number of peptides present, with probability P , hence the number of new peptides, Q_2 created by this mechanism is

$$Q_2 > P Q_1 [Q_1 + 2^{L_r+1}]^2. \quad (\text{A.6})$$

Case iv, New cleavage reactions

The number of new cleavage sites afforded by the Q_1 is the sum of the produce of the number of peptides of length L' , times their $L' - 1$ internal bonds. Each of these might be catalysed by any of the new or old polymers. This gives:

$$Q_2 = P 2^{L_r+1} \sum_{L'=L_r+1}^{2L_r} 2^{L'} (2L_r + 1 - L') (L' - 1). \quad (\text{A.7})$$

More simply, a minimum estimate of the number of internal sites available for cleavage among the Q_1 is $Q_1 L_r$, while the maximum $Q_1 (2L_r - 1)$. Since any of the old or new molecules might catalyse these cleavage reactions, the maximum and minimum estimates of Q_2 by this mechanism are:

$$Q_2 < 2 P L_r Q_1 Q_1 + 2^{L_r+1}, \quad (\text{A.8a})$$

$$Q_2 > P L_r Q_1 Q_1 + 2^{L_r+1}. \quad (\text{A.8b})$$

The total number of new peptides added to the catalysed reaction graph on the second iteration by these mechanisms, $Q_{2(\text{tot})}$ is greater than the sum of the right hand sides of eqs. (A.4) + (A.6) + (A.8b).

Divergence criterion

Over successive iterations of the growth of the catalysed reaction graph, a sufficient criterion such that the firing disc is supracritical is that the sum of (A.4) + (A.6) + (A.8b) diverges. Since the sum is difficult to evaluate, a simpler criterion is that only one of (A.4), (A.6) or (A.8b) diverges. We show next that the rate of growth of new condensation reactions from the first to second iteration dominates over the remaining mechanisms and dominates the process over successive iterations, hence provides a satisfactory divergence criterion.

On the second iteration from (A.6) and (A.8a), the ratio of new peptides produced by new condensation reactions to new cleavage reactions is greater than $2^{L_r} / L_r$. This dominance derives from the fact that the number of possible new condensation reactions is roughly the product of the new times total molecules, with the number of potential cleavage reactions is a linear function of new peptide lengths. The criterion for divergence we shall derive based on use of condensation reactions, yields reasonably large critical values of L_r , for plausible probabilities of catalysis. Hence dominance of condensation is a self-consistent assumption with respect to this divergence criterion.

The ratio of new peptides due to new condensation reactions to catalysis of condensation reactions among the firing disc members is given by the ratio of (A.6) to (A.4) and yields:

$$\frac{P2^{L_t+1}(Q_1 + 2^{(L_t+1)})^2}{Q_1}. \quad (\text{A.9})$$

We shall show below that the criterion for divergence derived from condensation alone is consistent with condensation dominating graph growth with respect to this process as well, hence not only a sufficient criterion for divergence, but a good estimate of critical values.

For divergence based on new condensation reactions, it suffices to show that $Q_1 > 1$, $Q_2 > Q_1$ and $Q_2 > Q_1$ implies that $Q_n > Q_{n-1}$. Substituting in eq. (A.6) from eq. (A.3) for Q_1 , and letting the symbol Σ stand for the summation term in eq. (A.3)

$$Q_2 > P2^{2^{3L_t}} + 3(\Sigma + 2P\Sigma^2 + P^2\Sigma^3). \quad (\text{A.10})$$

Since $Q_1 = P2^{L_t+1}\Sigma$, $Q_2 > Q_1$ if

$$\frac{P^22^{3L_t+3}}{P2^{L_t+1}} = P2^{2L_t+2} > 1 \quad (\text{A.11})$$

or

$$2^{2L_t+2} > 1/P. \quad (\text{A.12})$$

Since the number of molecular species in the firing disc is 2^{L_t+1} , the equivalent sufficient criterion that $Q_2 > Q_1$ is that the number of molecules in the disc,

$$2^{L_t+1} > 1/P^{1/2}. \quad (\text{A.13})$$

Next we show that $Q_2 > Q_1$, and $Q_1 > 1$ implies $Q_n > Q_{n-1}$. From (6),

$$Q_2 > PQ_1(Q_1 + 2^{L_t+1})^2, \quad (\text{A.13a})$$

$$Q_3 > PQ_2(Q_2 + Q_1 + 2^{L_t+1})^2, \quad (\text{A.13b})$$

$$Q_n > PQ_{n-1}(Q_{n-1} + Q_{n-2} + \dots + Q_1 + 2^{L_t+1})^2. \quad (\text{A.13c})$$

If eqs. (A.13a–c) were equalities, then the squared term of each successive iteration is larger than the previous, hence if $Q_2 > Q_1 > 1$, then $Q_n > Q_{n-1}$. Deviation of Q_2 over the equality in eq. (A.13a) enters into Q_3 as a cubic term, hence augments Q_3 more than Q_2 , thus assuring that the series Q_n diverges.

A sufficient criterion for Q_1 to be greater than 1 when the divergence criterion in eq. (A.12) holds, is derived next. The expected number of novel peptides on the first iterate, Q_1 is given by eq. (A.3). The largest term in the sum in eq. (A.3) is 2^{2L_t} , hence, substituting this in place of the sum yields:

$$Q_1 > P2^{3L_t+1} = P2^{2L_t+2}2^{L_t-1}. \quad (\text{A.14})$$

Using the critical criterion of eq. (A.12) and substituting,

$$Q_1 > P2^{2L_t+2}2^{L_t-1} > (P2^{L_t-1})/P = 2^{L_t-1} \quad (\text{A.15})$$

or

$$Q_1 > 2^{L_t-1} \geq 1 \quad (\text{A.16})$$

for any positive L_t .

These results show that the criterion in eq. (A.12) or eq. (A.13), suffice for the growth of the catalysed reaction graph to diverge.

While eqs. (A.12) or (A.13) are sufficient conditions for a firing disc to be supraccritical, it remains to be considered whether they are good estimates of the critical maximal size polymer, L_c , maintained in the firing disc, (or number of peptide species), since processes due to catalysis of "old" condensation reactions (case ii) and catalysis of new cleavage reactions (case iv) were suppressed. Eq. (A.12) leads, for plausible probabilities of catalysis, P , on the order of 10^{-6} to critical polymer length, L_c , greater or equal to 9. The dominance of condensation over cleavage is given by the ratio of Q_2 in eq. (A.6) to Q_2 in (A.8b) and is over 50 fold for $L_c = 9$. Over successive itera-

tions of the graph, the dominance of new condensation over new cleavage increases.

Eq. (A.9) gives the dominance of new condensation reactions over "old" condensation reactions, case ii. Substituting for P from eq. (A.12) yields for the dominance:

$$Q_1/2^{L_t+1} + 2 + 2^{L_t+1}/Q_1 \quad (\text{A.17})$$

which has a minimum at 4. For $P = 10^{-6}$, and solving from eq. (A.3) for Q_1 , the dominance of new condensation is over 500-fold. Dominance grows over each iteration of the graph's growth. Thus, eqs. (A.12) and (A.13) are good criteria for supracriticality.

These criteria can be improved slightly as follows: In the marginally supracritical case, each iteration of the graph's growth, $(1, 2, \dots, n)$ adds an identical number Q_i new peptides to the catalysed reaction graph. After n iterations, the number of molecules in the graph is $nQ_1 + 2^{L_c+1}$. The latest increment, Q_n can condense on the left or right ends of all molecules from earlier iterations, and onto themselves as well, yielding

$$Q_n^2 - Q_n + 2Q_n[(n-1)Q_1 + 2^{L_c+1}] \approx 2Q_n[nQ_1 + 2^{L_c+1}]. \quad (\text{A.18})$$

This leads, after substitution, to a critical value of

$$2^{2L_c+2} > 1/2P. \quad (\text{A.19})$$

Finally, the criterion for supracritical discs can be generalized to $B > 2$ types of amino acids. As B grows larger, the total number of molecular species in a firing disc up to polymer length L_f converges to B^{L_f} . Substituting into eq. (A.3) and eqs. (A.10)–(A.12) leads to

$$B^{2L_c} > 1/P \quad (\text{A.20})$$

or slightly better, from eq. (A.19)

$$B^{2L_c} > 1/2P. \quad (\text{A.21})$$



Appendix B

Approximate equations of motion for catalyzed reactions

A central difficulty which must be addressed in any treatment of catalyzed reactions is saturation. Each reaction involves formation of several intermediate complexes which have a nonzero binding time. This effect can easily become strong enough so that most of a given species is bound up in complexes, saturating the reaction with a large slow down in the reaction rate. The total number of intermediate complexes is equal to the number of reactions times the number of complexes per reaction, which typically is much larger than the total number of species. A realistic simulation of the full system with all of the complexes requires an unrealistic amount of computer resources.

The traditional approximation schemes for dealing with saturation effects are not applicable here. These schemes typically assume that one of the components of the reaction, such as the enzyme, is the limiting factor, and that all of the other components are present in abundance. For our problem, however, any given species may serve as an enzyme or a product at the same time; the reaction can be limited by saturation of any of its components. A new approach is needed.

To develop an approximation scheme we consider the forward and backward reactions separately, and assume that each of them forms only one intermediate complex. A given reaction in the reaction network is of the form





where h represents water, c is the concatenation of the polymers a and b , and e is an enzyme. k_u is a phenomenological dissociation constant that depends on the average amount of time that the complexes remain bound.

To write explicit equations of motion for these reactions it is convenient to consider intermediate variables that represent the concentration of each chemical bound in complexes. Letting \bar{a} denote the amount of a bound up in $\bar{a}\bar{b}\bar{e}$, and using similar notation for the other strings, we obtain



The subscripts 1 and 2 are placed beside e to indicate that they come from different complexes. Since (B.1b) and (B.1d) are linear, however, it is possible to replace these with their sum: $e = e_1 + e_2$, and just write one equation where it is understood that \bar{e} means all the enzyme, including that bound in different reactions. Note that this reduction of the bound enzyme dynamics to one equation requires k_u to be the same for all reactions. Our representation of the reactions then becomes



We assume for the moment that the concentration of water remains constant.

To write down the equations of motion for the reactions listed in eq. (B.2) assume there is a matrix v_{abce} whose entries are equal to a rate constant that is zero unless the reaction is catalyzed. An arbitrary string s can (and typically

does) enter into more than one reaction at once, playing different roles in different reactions, i.e., it can act either as a condensate (c), a cleavage product (a or b), or an enzyme (e). The equations of motion must contain terms corresponding to each of these roles, which are determined by the index position of v_{abce} . The position of each index determines each role in the basic equation. Assume N species are present. Denote the concentration of the i th species as s_i for the free part, and \bar{x}_i for the bound part. The equations of motion for the bound part are then:

$$\begin{aligned} \dot{\bar{x}}_i = & \sum_{j,k,l} \left[k_f v_{jkil} x_j x_k x_l \right. \\ & (\text{forward reaction, } i \text{ in role of c}) \\ & + k_f v_{jkli} x_j x_k x_l \\ & (\text{forward reaction, } i \text{ in role of e}) \\ & + k_r v_{ijkil} x_k x_l h \\ & (\text{backward reaction, } i \text{ in role of a}) \\ & + k_r v_{jikil} x_k x_l h \\ & (\text{backward reaction, } i \text{ in role of b}) \\ & \left. + k_r v_{jkli} x_l x_i h \right] \\ & (\text{backward reaction, } i \text{ in role of e}) \\ & - k_u \bar{x}_i \\ & (\text{dissociation of } i \text{ complex}) \end{aligned}$$

and the equations of motion for the free part are:

$$\begin{aligned} \dot{x}_i = & - \sum_{j,k,l} \left[k_f v_{ijkil} x_i x_j x_l \right. \\ & (\text{forward reaction, } i \text{ in role of a}) \\ & + k_f v_{jikil} x_i x_j x_l \\ & (\text{forward reaction, } i \text{ in role of b}) \\ & + k_f v_{jkli} x_j x_k x_l \\ & (\text{forward reaction, } i \text{ in role of e}) \\ & + k_r v_{jkil} x_i x_l h \\ & (\text{backward reaction, } i \text{ in role of c}) \\ & \left. + k_r v_{jkli} x_l x_i h \right] \\ & (\text{backward reaction, } i \text{ in role of e}) \\ & + k_u x_i \\ & (\text{dissociation of } i \text{ complex}). \end{aligned}$$

v_{ijkl} is always very sparse; in practice the values are kept in a linked list. Furthermore, these terms can be matched up exactly, since the gain of the terms on the right side is equal to the loss of those on the left. The equations of motion may be written more compactly as:

$$\begin{aligned}\dot{\bar{x}}_i &= \sum_{j,k,l} \left\{ k_f (v_{jkil}x_l + v_{jkli}x_i)x_jx_k \right. \\ &\quad \left. + k_r h [(v_{ijkl} + v_{jikl})x_kx_l + v_{jkli}x_lx_i] \right\} - k_u \bar{x}_i, \\ \dot{x}_i &= - \sum_{j,k,l} \left\{ k_f x_i x_j [(v_{ijkl} + v_{jikl})x_l + v_{jkli}x_k] \right. \\ &\quad \left. + k_r h x_j x_l (v_{jkil} + v_{jkli}) \right\} + k_u \bar{x}_i.\end{aligned}$$

References

- [1] S.L. Miller and H.C. Urey, *Science* 130 (1959) 245.
- [2] Calvin, *Chemical Evolution* (Oxford Univ. Press, Oxford, New York, 1969).
- [3] M. Eigen, *Naturwissenschaften* 58 (1971) 465.
- [4] M. Eigen and P. Schuster, *Naturwissenschaften* 64 (1977) 541.
- [5] M. Eigen and P. Schuster, *the Hypercycle* (Springer, Berlin, Heidelberg, New York, 1979).
- [6] S. Kauffman, *J. Cybernetics* 1 (1971) 71.
- [7] S. Kauffman, *J. Theoret. Bio.* 119 (1986) 1–24.
- [8] O. Rossler, "A System Theoretic Model of Biogenesis," *Z. Naturforsch. B26b.* (1971) 741–746.
- [9] O. Rossler, "Chemical Automata In Homogeneous and Reaction-Diffusion Kinetics," *Notes in Biomath. B4* (1974) 399–418.
- [10] O. Rossler, "Deductive Prebiology," in *Molecular Evolution and the Prebiological Paradigm*, K.L. Roling ed. (Plenum, New York, 1983).
- [11] F. Dyson, "A Model for the Origin of Life," *J. Mol. Evol.* 118 (1982) 344–350.
- [12] F. Dyson, *Origins of Life* (Cambridge Univ. Press, London, 1985).
- [13] M.S. Silver and S.L.T. James, *Biochemistry* 20 (1981) 3177.
- [14] M.S. Silver and S.L.T. James, *Biochemistry* 20 (1981) 3183.
- [15] Y. Levin, A. Berger and E. Katchalsky, *Biochem. J.* 63 (1956) 308.
- [16] A.J. Zaugg and T.R. Cech, *Science* 231 (Feb. 1986) 470–475. T.R. Cech, *Proc. Nat. Acad. Sci.* 83 (1986) 4360.
- [17] J.D. Farmer, N.H. Packard and A.S. Perelson, "The Immune System, Adaptation, and Machine Learning," *Physica* 22D (1986) 187 (these proceedings).
- [18] A.S. Perelson and G.F. Oster, "Theoretical Studies of Clonal Selection," *J. Theo. Bio.* 81 (1981) 645–670.
- [19] S.W. Fox and K. Dose, *Molecular Evolution and the Origin of Life* (Freeman, San Fransisco, 1972).
- [20] M. Dixon and E.C. Webb, *Enzymes* (Academic Press, New York, 1960).