Richard J. Bagley and J. Doyne Farmer
Complex Systems Group, Theoretical Division, and Center for Nonlinear Studies, Los Alamos National Laboratory, Los Alamos, NM 87545 and Santa Fe Institute, 1120 Canyon Road, Santa Fe, NM 87501

Spontaneous Emergence of a Metabolism

Networks of catalyzed reactions with nonlinear feedback have been proposed to play an important role in the origin of life. We investigate this possibility in a polymer chemistry with catalyzed cleavage and condensation reactions, studying the properties of a well-stirred reactor driven away from equilibrium by the flow of mass. Near equilibrium the distribution of material is uninteresting; it favors short polymers but is otherwise homogeneous. However, under appropriate non-equilibrium conditions, the situation changes radically: The nonlinear feedback of the reaction network focuses the material of the system into a few specific polymer species, whose concentrations can be orders of magnitude above the background. Like a metabolism, the network of catalytic reactions "digests" the material of its environment, incorporating it into its own form. For this reason we call it an *autocatalytic metabolism*. We vary the diet of an autocatalytic metabolism, and demonstrate that under some variations it persists almost unchanged, while in other cases it dies. We argue that the dynamical stability of autocatalytic metabolisms gives them regenerative properties that allow them to repair themselves and to propagate through time.

Artificial Life II, SFI Studies in the Sciences of Complexity, vol. X, edited by C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen, Addison-Wesley, 1991

CONTENTS

1. MOTIVATION

1.1 SETTING THE STAGE FOR AN ORIGIN OF LIFE

When Miller and Urey discovered that amino acids could be formed under conditions that might be similar to those of the prebiotic earth,[28] the spontaneous synthesis of proteins seemed just around the corner. However, this turns out to be much more difficult than the spontaneous synthesis of individual amino acids. Equilibrium conditions tend to favor dissociation, and generate a concentration profile that is fairly uniform. Except for occasional fluctuations, for long polymers the population of any given molecular species is typically zero. The population distribution of polymers is homogeneous, nonspecific, and uninteresting. This is in contrast to living organisms, which have high concentrations of a few *specific* polymer species.[1]

Contemporary organisms achieve specificity through a codependent relationship between templates and enzymes. Proteins and nucleic acids synthesize each other through a replication mechanism in which none of the components synthesizes itself. Even for the simplest organisms, this process is highly complex. There seems to be a minimum level of complexity below which a replicating machine based on proteins and nucleic acids simply cannot function. While it is easy to understand how such a replicating machine perpetuates itself, it is difficult to understand how the necessary initial conditions ever arose on their own. The probability that both enzymes and templates could be created through a statistical fluctuation is effectively nil. This suggests that other processes preceded contemporary life.

The idea that enzymatic activity might have set the stage for the origin of life was developed by Oparin, who suggested that coacervates may have played a major role.[33] Early experiments unsuccessfully attempted to use clays and other materials as nonspecific catalysts for polymerization.[34] Calvin studied several different scenarios through which catalytic activity could provide a selection mechanism, even without self-replication.[5,6] In 1971 Rössler,[37,38] Eigen,[10] and Kauffman[23] developed this idea further. In particular, Rössler envisioned a form of chemical evolution similar to that studied here. He emphasized the importance of specific catalysts which catalyze only a small fraction of all possible reactions. Along these same lines, Kauffman[24] later modeled the problem in terms of random graphs, and showed that under reasonable assumptions the probability of catalytic closure is quite high.[2] The random graph model was developed into a kinetic model that could be simulated on a computer by Farmer et al.[11] This line of investigation, which attempts to find possible precursors facilitating the emergence of life should be contrasted with

[1] An exception is provided by the experiments of Sidney Fox, who by heating a mixture of amino acides demonstrated the formation of polypeptides, called proteinoids.[20,21] The structure of proteinoids is not random; some subsequences, such as certain hexapeptides, occur much more frequently than others. In contrast, our goal is to increase the concentration of entire molecules, so that it is several orders of magnitude above equilibrium.

[2] Another toy model investigating the possibility that a metabolism might have spontaneously emerged without a replicator is due to Dyson.[8]

other work that addresses the (also very interesting) question of the early evolution of life once replication has already begun.[1,9,10,39]

In this paper we study the behavior of a network of catalyzed chemical reactions, along the lines laid out by Farmer et al.[11] We make several enhancements of the model and analytically study a few simple cases to gain better intuition about the dynamics. We also improve the simulation so that it is several orders of magnitude faster. This allows us to simulate the kinetics of a complicated reaction network in a matter of seconds. As a result, we are able to widely explore the parameter space and answer many of the questions originally raised in earlier papers.

Our main result is that, under appropriate conditions, a catalytic reaction network can focus most the material of its environment into a few chemical species. For this to happen the system must be driven the appropriate distance from equilibrium, polymerization must be favored, and it must have diverse kinetic parameters. Favoring polymerization may require the addition of energy, for example, through pyrophosphates energized by light. In spite of these restrictions, there is a wide range of parameters in which the material of the system is focused into only a few species, which dominate over the background.

Focusing radically alters the material composition of the environment. The species that emerge reinforce each others' production and largely take over the reaction vessel, excluding other possibilities. Since this behavior is analogous to that of a metabolism, we call the resulting set of species and reactions an *autocatalytic metabolism*. Under appropriate conditions autocatalytic metabolisms can evolve out of a simple, undifferentiated initial state, generating a sequence of complex, highly differentiated, final states. Like contemporary organisms, these final states are composed of a highly focused, specific set of long polymers. While the autocatalytic metabolism does not replicate itself in the usual sense, it propagates itself by taking over any medium with suitable properties, sustaining itself as long as the appropriate conditions are met. Furthermore, it may generate a lineage of related autocatalytic metabolisms.

The model that we study here applies to any system in which polymers can catalyze the formation of other polymers through cleavage and condensation reactions. If the basic building blocks are amino acids, then the polymers are called polypeptides or proteins. Such reactions are common among proteins, forming the basis for many of the functions of living organisms. If the basic building blocks are nucleotides, then the polymers are called nucleic acids (RNA or DNA). It is well known that polypeptides possess a large repertoire of catalytic activities of this type; the recent discovery of specific catalysis reactions in RNA suggests that nucleic acids may also possess the necessary properties.[7] Whether the polymers are polypeptides or nucleic acids changes the parameters but not the basic form of the model.

Even if the model we discuss here has nothing to do with the actual origin of life on earth, it might provide a *possible* origin of life in the laboratory. Although at present we cannot predict the outcome of experiments in detail, we can make qualitative predictions that provide broader experimental guidelines. The accumulation of more experimental knowledge can be used to determine the unknown parameters

of our model, which in turn should sharpen its predictive value for experiments. Although many important experimental details are still unknown, and some important questions await further study, our numerical simulations suggest that it may be possible to synthesize autocatalytic metabolisms in the laboratory.

1.2 NONTRIVIAL DISSIPATIVE STRUCTURES

The problem of the emergence of life is embedded in the broader problem of understanding self-organizing phenomena, which from the point of view of a physicist may be more interesting anyway. Many non-living systems exhibit self-organizing properties, albeit much weaker than those of living systems. Is there a sharp distinction between living and nonliving systems? Or can there exist levels of organization that are between those of present living and non-living systems? Can evolution and other self-organizing properties of living systems be viewed as manifestations of a general law that describes the tendencies of matter to organize itself?

There are many simple examples that have been cited as instances of self-organizing phenomena in nature. For example, when a fluid is heated from below, under appropriate conditions, patterns of convection cells form. The macroscopic structure of these cells is internally generated by the system itself, and is not apparent in its initial conditions. Such patterns are often called *dissipative structures*, because they occur when energy flows into a system and then is dissipated.[32]

Several researchers have asserted that life is a non-equilibrium phenomenon, associated with dissipative structures.[32] This is certainly true, but it is a very weak statement. While deviation from equilibrium is a necessary condition for life, it is far from sufficient. Driving a system away from equilibrium does not necessarily cause the emergence of order—in fact, it often has precisely the opposite effect. A central question that must be addressed in a theory of self-organizing phenomena is: Why do some non-equilibrium situations foster the spontaneous emergence of organization, while others do not?

There is a big gap between the dissipative structures of simple non-living systems, such as patterns in fluid convection, and the much richer dissipative structures associated with living systems. The model discussed here is intended to bridge this gap, at least to some extent, by showing the possibility for dissipative structures that are intermediate in complexity between living and non-living systems. Autocatalytic metabolisms are more complex than convection patterns, in that they propagate specific information through time. One autocatalytic metabolism can seed the formation of another, similar metabolism. The autocatalytic metabolisms of this model can be viewed as proto-life forms, since they have a metabolism, they evolve and store information, and they reproduce (although more continuously and with less fidelity than contemporary organisms). They are also dynamically stable, and so capable of self-repair. They, thus, have many of the essential properties of living systems, albeit in a much less sophisticated form.

Besides demonstrating the possibility for the spontaneous generation of autocatalytic metabolisms, one of our main purposes in this paper is to discover under

what conditions they can be expected to form. How does their formation depend on the parameters of the system, such as the flow of energy, or the inherent diversity of the underlying dynamics? Although our results are specific to this model, they nonetheless suggest several rules that may pertain to the more general problem of self-organization.

1.3 A SIMPLE MODEL FOR STUDYING EVOLUTION WITH AN EMERGENT NOTION OF FITNESS

In principle, it is possible to describe biological systems at a fundamental level in terms of their dynamics. At this level of description, "selection" is an emergent property of the dynamics. In practice, however, for most systems this is hopelessly intractable. As a result, studies of evolution are typically couched in terms of the fitness function, which is an empirical construct, disconnected from the laws of physics. Even so, in most systems the fitness function is known only in very special circumstances where all but a few relevant factors are neglected. In general, fitness is a complicated function of the external environment, which includes other organisms. As a result, most theoretical models for evolution make many *ad hoc* assumptions, postulating fitness functions that may be qualitatively different from those in the real world.

As pointed out by Eigen,[10] Rössler,[37] and others, chemistry provides an excellent forum for studying evolution. The laws of chemical kinetics are well understood, and make it possible to model the behavior of the system at a fundamental level. These laws determine population levels and therefore determine fitness. As in biological systems, fluctuations are always present, generating random variation. Thus, for chemical networks we can describe the fitness at the fundamental level of dynamics.

Even though autocatalytic metabolisms do not have templates or a genetic code, because of their specificity they are nonetheless capable of evolution. This is discussed in a companion paper.[4] Autocatalytic metabolisms, therefore, provide an interesting alternative for studying evolution in a chemical setting. It is also interesting to note that autocatalytic structures analogous to those we study here occur spontaneously in more abstract environments, as observed by Fontana[16,17] and Rasmussen et al.[36]

2. BACKGROUND

In this section we discuss some of the properties of catalyzed reaction networks, providing a background for the development of the simulation in Section 3. We discuss the reactions we are going to consider, and show why they are uninteresting at equilibrium. We then explain how the situation is altered as we move away from equilibrium, and how catalysis can play an important role in focusing the material

of the system into just a few chemical species. We define autocatalytic sets and the related notion of autocatalytic metabolisms.

2.1 SPONTANEOUS REACTIONS

We are interested in reversible polymerization reactions, in which two polymers either *condense* to form a single longer polymer, or a single polymer *cleaves* into two shorter polymers. Cleavage and condensation can be considered together as a single reversible reaction. The reaction in which polymers A and B join together to form C, giving off water, or equivalently, in which C hydrolyzes into A and B, can be written

$$A + B \rightleftharpoons C + H, \tag{1}$$

where H represents water.

Providing the concentrations are sufficiently high and the solution is well stirred, the law of mass action provides a good approximation of the kinetics. Let k_f be the rate constant for the forward reaction, $A+B \rightarrow C+H$, and k_r be the rate constant for the backward reaction $C + H \rightarrow A + B$. The rate equation for C is then

$$\dot{C} = \frac{dC}{dt} = k_f AB - k_r HC. \tag{2}$$

For convenience, whenever the meaning is unambiguous, we use the same symbol to represent both a polymer and its concentration. Similar equations apply for $\dot{A}$ and $\dot{B}$.

2.2 EQUILIBRIUM DISTRIBUTION OF POLYMERS

At equilibrium the concentrations of the polymers of a given length can be computed analytically using the classical theory of polycondensation reactions developed by Flory and Stockmayer.[15,40,44] For simplicity we assume that all the reactions have the same forward and backward rate constants, and that the reaction vessel is well stirred. Furthermore, we assume that the monomers are oriented so that each monomer has two sites, which we arbitrarily designate as the "+" site and "–" site.

We follow the treatment of Macken and Perelson.[27] Rather than solving for the concentrations of each polymer, it is more convenient to use an aggregate variable y, which is the concentration of free sites of a given kind (either "+" or "–"). We assume that the polymers are unbranched, and that they cannot form rings. For reactions of the form of Eq. (2),

$$\dot{y} = -k_f y^2 + k_r(m_0 - y)H, \tag{3}$$

where m_0 is the total concentration of monomers, which is equal to the concentration of free sites if nothing is bound. At steady state $\dot{y} = 0$ and the concentration of free sites is

$$y = \frac{(1 + 4\kappa y_0)^{1/2} - 1}{2\kappa}, \tag{4}$$

where $\kappa = k_f / H k_r$ is the equilibrium constant.

We now compute the concentration of polymers of length n. At equilibrium, let ρ be the probability for the formation of a bond. This is the ratio of bound sites to the total number of sites, i.e.,

$$\rho = \frac{m_0 - y}{y_0}. \tag{5}$$

Assume that each binding event is independent. For a given free site, the probability that it is attached to $n-1$ bonds followed by another free site is $\rho^{n-1}(1-\rho)$. Solving Eq. (5) for y shows that the concentration of free sites for a given value of ρ is $y = m_0(1-\rho)$. Thus, the concentration of polymers of length n is

$$x_n = m_0(1-\rho)^2 \rho^{n-1}. \tag{6}$$

At equilibrium, inserting Eq. (4) into Eq. (5) gives

$$\rho = 1 - \frac{(1 + 4\kappa m_0)^{1/2} - 1}{2\kappa m_0}. \tag{7}$$

Note that $\rho < 1$. Thus, Eq. (6) implies that the concentration of polymers of length n decreases exponentially with n, at a rate that depends only on the product of the equilibrium constant and the concentration of monomers. In a system with m distinct monomers, present initially at equal concentrations, the concentration of any particular polymer species of length n is further decreased by a factor of m^{-n}. For $m > 1$, even if $\rho \approx 1$, so that polymerization is favored, for large n the concentration of any particular species is quite small. For example, for polypeptides $m = 20$; the concentration of a polypeptide of length $n = 30$ is roughly 20^{-30} less than that of a monomer. For a container of finite size, this implies that, except for occasional fluctuations, most longer species are not present.

2.3 CATALYZED REACTIONS

The presence of a catalyst (enzyme) E can accelerate a reaction.

$$A + B \overset{E}{\rightleftharpoons} C + H. \tag{8}$$

At equilibrium the rate of the forward reaction equals that of the backward reaction, so that $\dot{A} = \dot{B} = \dot{C} = 0$. Catalysis speeds up the rate at which the system approaches equilibrium, but does not change the concentrations at equilibrium.

However, when the reaction is driven away from equilibrium, for example by externally supplying one of the participants in the reaction, catalysis can shift the steady state. This is the basis for the effect we study here.

Catalysis increases both the forward and backward rate constants by the same amount. This can be taken into account by defining a quantity ν that we call the *catalytic efficiency*. For fixed concentration of the reactants, the increase in the velocity of the reaction is proportional to the product of ν and the concentration E of the catalyst. The kinetic equation for C can be crudely approximated as

$$\dot{C} = (1 + \nu E)(k_f AB - k_r HC). \tag{9}$$

Similar equations apply to $\dot{A}$ and $\dot{B}$. When the catalytic efficiency $\nu = 0$, this reduces to the kinetic equation for a spontaneous reaction.

Note that, for a population of polymers, this reaction is just one reaction in a network of many. A given polymer may play the role of A in some reactions, and the role of C in others. To compute the rate of production of any given species, it is necessary to sum all the relevant reaction terms.

The approximation made in Eq. (9) neglects the effect of saturation, which comes about because the enzyme and the reactants are bound together for a finite time. During this time they cannot participate in new reactions, which lowers the effective reaction rates. If this is a dominant effect, so that most of the enzyme or product is bound at any given time, the reaction is *saturated*. To take this into account we do not use Eq. (9), but rather use a more accurate approximation. We keep track of the concentration of any given species x_i that is bound into complexes through an auxiliary variable $\overline{x}_i$, which is equal to the sum of the concentrations of all the complexes in which x_i is bound. To keep the simulation tractable, we assume that all complexes unbind at the same rate k_u. This approximation is described in more detail in the Appendix.

2.4 DRIVING FROM EQUILIBRIUM

To make anything interesting happen in a reaction network it must be driven away from equilibrium. In this model we investigate two different mechanisms. The first involves a flux of mass, and the second involves the formation of energetic pyrophosphate molecules, driven by light.

2.4.1 MASS FLOW We model a reaction vessel with a steady input flux of monomers or short polymers, and an output flux due either to diffusion or overflow of the reaction vessel. This might correspond to a prebiotic environment, or it might correspond to a chemostat in a laboratory experiment. The chemical species that are input are collectively called the *food set*. For simplicity, we assume an inflow rate δ of concentration per unit time, and an outflow that is proportional to concentration, with rate constant K.

For an element of the food set, the kinetic equations are of the form

$$\frac{dx_i}{dt} = k_u \overline{x}_i + \sum(reaction\ terms) + \delta - K x_i, \tag{10}$$

$$\frac{d\overline{x}_i}{dt} = -k_u \overline{x}_i + \sum(reaction\ terms) - K \overline{x}_i. \tag{11}$$

The (*reaction terms*) are defined in the appendix in the discussion following Eq. (25). For a species outside the food set, the kinetic equations are of the same form, except that $\delta = 0$. There is a net flow of mass from the food set to the other elements of the system, which drives it away from equilibrium.

Because the reaction terms conserve mass, the total mass in the reaction vessel always goes to a fixed point, independent of initial concentrations. To see this, note that only the last two terms in Eq. (10) and the last term in Eq. (11) change the total mass. The total mass concentration is proportional to[3] $m = \sum n(i)x_i$, where $n(i)$ is the length of the ith species. Letting N_f be the number of elements in the food set, the rate of change of the total mass concentration is given by a simple differential equation,

$$\frac{dm}{dt} = N_f \delta - K m, \tag{12}$$

which has a global fixed point $m_0 = N_f/K\delta$. This means the initial mass is irrelevant anyway, and so for convenience in our simulations, we choose $m(0) = m_0$. Thus, there are effectively only two parameters relating to the flow of mass through the system, which can be δ and K, or equivalently δ and m_0.

Since δ and K are not intuitively easy to interpret, it is sometimes useful to quote results in terms of the *mean reaction number* r. This is defined as the mean number of times a given monomer participates in a reaction, on average, from the time it enters the vessel until the time it is flushed out. At equilibrium r is infinite, and when the other parameters are fixed, it decreases monotonically as δ increases.

2.4.2 PYROPHOSPHATES As we demonstrate in subsection 2.6, catalytic focusing requires conditions that favor polymerization. The tendency to polymerize can be enhanced by an appropriate input of energy. The mechanism that we investigate here involves pyrophosphate molecules, which play a role analogous to that of ATP in contemporary organisms. This mechanism is supported by early experiments.[35] The detailed sequence of reactions was suggested by Ron Fox,[4] and is illustrated in Table 1.

It proceeds as follows: Light causes the formation of pyrophosphate molecules (p_2), which is balanced by hydrolysis. When a pyrophosphate molecule binds to polymer A, it creates the energized form A^* and releases a phosphate atom in the process. A^* may hydrolyze, releasing the other phosphate atom, or it can bind to another polymer B (in the presence of the catalyst E). This occurs with a rate constant k_f^*, which is greater than the unenergized rate constant k_f. Thus, the addition of energy favors polymerization.

Based on simulations involving the full reaction scheme shown in Table 1, we found that when the concentration of pyrophosphate and the input of light are sufficiently high, the behavior is roughly equivalent to that obtained by simply using the equations given in the Appendix, with the effective forward rate constant equal to k_f^*. For convenience, in the numerical experiments described here we simply assume that the parameters quoted correspond to k_f^*, and use the simpler equations which do not involve pyrophosphate.

[3] To convert this to units of mass/volume there is a constant of proportionality, that depends on the mass of a monomer.

[4] Private communication.

TABLE 1 Pyrophosphate energizes and enhances polymerization. The first column lists the reactions and the second column the reaction rate. γ represents a photon (and in the column on the right represents the intensity of light), p a phosphate atom, and H water. A and B represent the polymers that condense to form C. E is the catalyst. A bar indicates a complex that is bound to the enzyme E. k_f is the rate constant for polymerization, k_r for hydrolysis, k_u for the dissociation of the bound complex, k_e for the polymerization of phosphate, and k_a for the activation of a polymer. Activated polymers are indicated with a "*" superscript. k_f^* is the rate constant for the condensation of an activated polymer with another polymer.

Reaction			Rate
$2p$	$\xrightarrow{\gamma}$	$p_2 + H$	$k_e \gamma p^2$
$p_2 + HC$	$\rightarrow$	$2p$	$k_r p_2 H$
$A + p_2$	$\rightarrow$	$A^* + p$	$k_a A p_2$
$A^* + H$	$\rightarrow$	$\overline{A + p}$	$k_r A^* H$
$A^* + B + E$	$\rightarrow$	$\overline{CE} + p$	$k_f^* \nu E A^* B$

2.5 REACTION GRAPH

Each distinct monomer can be assigned a character from a fixed alphabet, $a, b, c, \ldots$. A polymer can then be represented as a character string, for example ($acabbacbc\ldots$). We assume that the polymers are oriented, so that abc and cba are different strings. The topological structure of a network of reactions, each of the form of Eq. (8), can be represented as a polygraph with two types of nodes and two types of connections,[11] as illustrated in Figures 1 and 2. One type of node represents the

polymer species and is labeled by the corresponding string. The other type of node represents the catalyzed reaction and is labeled by a black dot. The polymers that participate in a reaction are connected to the corresponding reaction node by reaction links (black arrows), which point in the direction of condensation. Each polymer is connected to the reactions it catalyzes by a catalytic link (dotted line). Each reaction has at least four links: three reaction links, and one or more catalytic links.

2.6 CATALYTIC FOCUSING

Under appropriate conditions catalysis can focus most of the material of a reaction network into only a few species. The basic idea can be grasped by considering the simple reaction network

$$ba + H \rightleftharpoons a + b \overset{E}{\rightleftharpoons} ab + H,$$

as shown in Figure 1. Assume a and b are supplied at rate δ, and diffuse out of the container with rate constant K, as described in subsection 2.4. For simplicity, assume the concentrations of E and H are maintained at fixed values. Neglecting saturation, according to the approximation of Eq. (9), the rates of change of ab and ba are

$$\dot{[ab]} = \gamma(k_f[a][b] - k_r H[ab]) - K[ab], \tag{13}$$

$$\dot{[ba]} = k_f[a][b] - k_r H[ba] - K[ba], \tag{14}$$

where $[ab]$ is the concentration of polymer ab. Setting the derivatives to zero and using the mass conservation condition of Eq. (12) gives

$$\frac{[ab]}{[ba]} = \frac{1+\beta}{1+\frac{\beta}{\gamma}}. \tag{15}$$

$\beta = \delta/m_0 k_r H$ is a dimensionless parameter related to the deviation from equilibrium, where $m_0 = a(0) + b(0)$ is the total concentration of monomers. Note that $\beta \geq 0$. $\gamma = 1 + \nu E$ is a dimensionless parameter that characterizes the strength of catalysis. $\gamma \geq 1$, and $\gamma = 1$ corresponds to an uncatalyzed reaction.

Under what circumstances is the concentration of ab much greater than that of ba? At equilibrium $\beta = 0$ and the ratio of $[ab]$ to $[ba]$ is one. This ratio can become large only when $\beta \gg 0$, i.e., only when the system is driven well away from equilibrium. When $\beta \gg \gamma$ this ratio approaches γ; when $\beta \ll \gamma$ it approaches β. Thus, by varying γ and β the concentration of ab relative to ba can be made arbitrarily large.

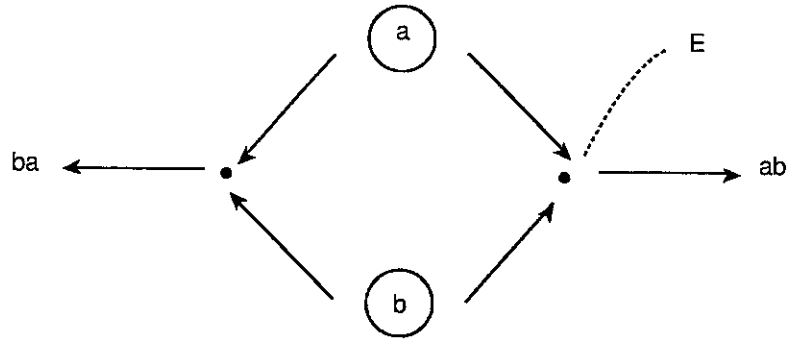

FIGURE 1 A simple network illustrating how steady-state concentration can be boosted by catalytic activity. a and b are driven at a fixed rate δ, and the enzyme E is maintained at a fixed concentration.

Note that the ability to focus comes about because the formation of one species is catalyzed, while that of the other is not. If all reactions were catalyzed equally, with equal kinetic parameters, there would be no focusing; the concentration of ab would equal that of ba. Focusing thus requires *specific catalysis*, in which some reactions are catalyzed more strongly than others.

2.7 AUTOCATALYTIC SETS AND METABOLISMS

To achieve catalytic focusing the enzyme E must be maintained at high concentration. One way for the system to accomplish this by itself is through an autocatalytic reaction, in which one of the products catalyzes its own formation. A simple example is

$$A + B \overset{C}{\rightleftharpoons} C + H. \tag{16}$$

If we set $C = ab = E$ in reaction (13), then the enzyme is produced automatically, and the focusing maintains itself.

Simple autocatalytic reactions such as reaction (16) are obviously very special. A more common situation occurs when autocatalysis involves a cooperation between reactions, in which one species catalyzes the formation of another. An *autocatalytic set* is defined as a set of chemical species such that each member of the set is produced by at least one catalyzed reaction involving only members of the set. This notion was introduced by Calvin,[6] Eigen,[10] Kauffman,[23] and Rössler.[37] Since the reactions we are considering are reversible, a species can be produced either by cleavage or condensation. Thus reaction (16) is an autocatalytic set, and so is

$$A + B \overset{A}{\rightleftharpoons} C + H. \tag{17}$$

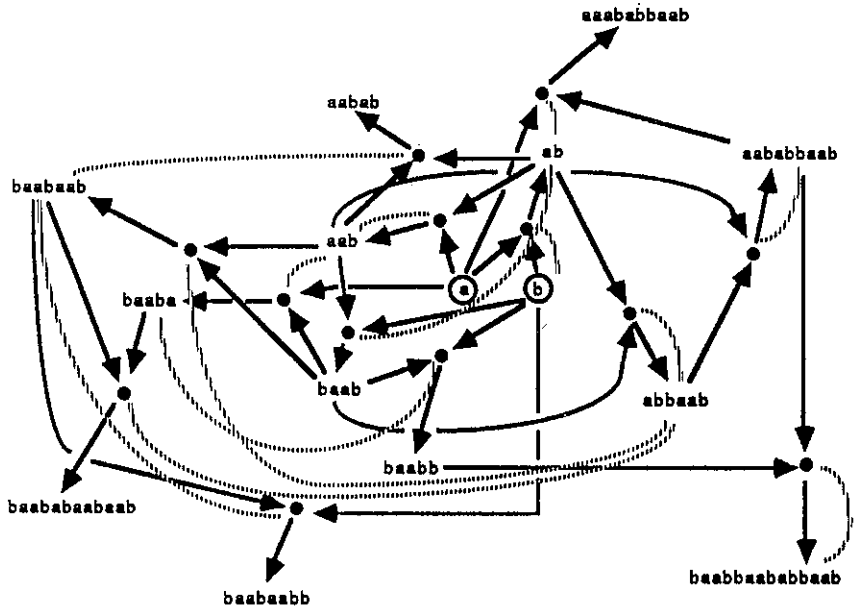

FIGURE 2 An autocatalytic network consisting of 15 species. The monomers a and b, which are circled, comprise the food set. Character strings represent the polymers. The fat dots represent reaction nodes. The arrows connecting the reactants to the nodes are reaction links and point in the direction of condensation. The broken lines connecting polymer species to reactions are catalytic links, indicating which species catalyze which reactions.

A more complicated (and more typical) autocatalytic set is shown in Figure 2. This network happens to have one catalytic link per reaction, a rather special property. We will use this reaction network, or variants with more catalytic links, for many of our numerical experiments. Note that autocatalytic sets may contain other autocatalytic sets as subsets.

Using the random graph model described in the next section, Kauffman[24] showed that for any given probability of catalysis, if the food set is sufficiently large, the resulting graph will almost always have an autocatalytic set. This result is encouraging, since it suggests that autocatalytic sets exist under fairly reasonable conditions. We must emphasize, however, that the presence of an autocatalytic set in a reaction network, in and of itself, does not imply that there will be any interesting departures from equilibrium behavior. To achieve catalytic focusing it is critical that the kinetic parameters are favorable. Thus the graph-theoretic notion of an autocatalytic set is a necessary but not a sufficient condition.

To make this distinction, we define an *autocatalytic metabolism* as an autocatalytic set whose concentrations make significant departures from the values they would have if none of the reactions were catalyzed. The phrase "significant departures" is subjective, and is admittedly rather vague. However, from an operational point of view, in our simulations we often see a clear distinction between autocatalytic sets that can function as metabolisms and those that cannot, as shown in Section 6.

3. SIMULATION

In principle the kinetic equations are all we need to know in order to simulate the behavior of a reaction network. In practice, however, there are two major problems: The first is that the kinetic parameters cannot be determined from first principles. To deal with this we construct an artificial chemistry, as discussed in subsection 3.1. The second problem is that there are an infinite number of possible reactions, and it is intractable to solve all of them; we must focus our computational resources on only the most relevant ones. Our method for doing this is discussed in subsection 3.2.

3.1 ARTIFICIAL CHEMISTRY

In a real chemical system the efficiencies and rate constants of the reactions depend on detailed properties of chemical composition, as well as on thermodynamic parameters such as temperature and pressure. While a computation of these constants from quantum mechanics and statistical mechanics is possible in principle, from a practical point of view, at this point in time it is hopelessly intractable.

To circumvent this problem, the approach introduced by Kauffman,[24] Farmer et al.,[11,12] and Bagley et al.[2] is to invent an artificial chemistry, a set of rules stating which catalyzed reactions occur, and with what strength. An artificial chemistry cannot reproduce the behavior of a real chemistry in detail, but it may reproduce many of the correct qualitative properties. An artificial chemistry can produce complex behavior, even though it is simple from a calculational point of view. By exploring different artificial chemistries, we can discover which properties cause significant changes in behavior, and which do not. We can begin with simple chemistries and move toward more complex chemistries, adding layers of realism as needed. The knowledge gained in this way can be useful in guiding experimental investigations of real systems, by pinpointing the essential quantities that need to be measured in experiments in order to make the model more realistic.

Since our primary interest is in understanding the effect of catalysis, we first address the problem of assigning a catalytic efficiency to each reaction. We do this using two different methods. In the first, we construct a completely disordered artificial chemistry, by assigning catalytic efficiencies at random, and in the second, we

construct a highly ordered artificial chemistry, assigning them with a string matching algorithm. The random method is more disordered than real chemistry, and the string matching method is more regular than real chemistry. From a qualitative point of view, we hope that real chemistry lies somewhere between these two extremes.

To be strictly correct, every possible reaction should be included in the reaction graph. However, in practice the reaction graph must be trimmed, so that computational resources are used only for the most essential reactions. We take advantage of the fact that the vast majority of reactions are catalyzed only weakly and can be treated essentially as spontaneous reactions. The graph represents only those reactions with sufficient catalytic efficiency to make them significantly different from the corresponding spontaneous reactions. Thus, when we refer to a "catalyzed reaction," we mean a "strongly catalyzed reaction," and when we refer to a spontaneous reaction, we mean a "weakly catalyzed reaction."

3.1.1 RANDOM ASSIGNMENT OF REACTIONS In some cases changing a single monomer can have a dramatic effect on the chemical properties of a polymer, either because it causes a drastic change in the configuration of the polymer or because it alters the properties of a critical site. If this were always the case, then chemistry would be random. For a random chemistry there is no correlation between chemical formulas and chemical properties. This is unrealistic. However, it does have the advantage of being easy to implement, and lies at one extreme in the space of all possible chemistries.

Following Kauffman,[24] we assume that out of all possible spontaneous reactions, only a fraction p are catalyzed with sufficient strength to be significantly different from spontaneous reactions. The set of reactions that is catalyzed is chosen at random. To see the basic idea, imagine creating a list of all possible catalyzed reactions. For m distinct monomers the number of species of length n is m^n, the number of possible spontaneous reactions is the order of m^{2n}, and the number of possible catalyzed reactions is the order of m^{3n}. The reactions that are strongly catalyzed can be determined by flipping a biased coin that returns heads with probability p. Reactions that receive heads are assigned a non-zero value of ν, and all others are assigned $\nu = 0$. The random rule generates an ensemble of possible chemistries, corresponding to all possible sequences of random choices.

Operationally the procedure described above would be very time consuming, since m^{3n} can be a very large number. It can be made much more efficient by decomposing the problem properly, focusing attention only on reactions involving species that are already present in the reaction vessel, and taking care to avoid double counting, as described by Farmer et al.[11]

To determine catalytic efficiencies, one simple possibility is to set $\nu =$ constant, so that all the reactions on the graph are catalyzed with the same efficiency. Another natural possibility is to choose the catalytic efficiencies at random according to a given probability distribution, for example by making the probability of a given efficiency uniform within given maximum and minimum values. We employ both of these in our simulations.

3.1.2 ASSIGNMENT OF REACTIONS BY STRING MATCHING The match rule provides an alternative artificial chemistry that is probably closer to real chemistry than the random rule discussed above. It lies at the opposite extreme—while the random rule is too disordered, the match rule is probably too ordered. For the match rule, changing a single monomer in a given polymer only causes a small change in its chemical properties. Two similar polymer strings always have similar chemical properties. The match rule assumes that the information contained in the string of a given polymer contains all the information needed to specify its chemical properties.

We roughly follow the approach used to model the immune system by Farmer et al.[12] For convenience, in this discussion we assume a two-letter alphabet consisting of a and b, although the rule is easily generalized to a larger alphabet. The two reactants A and B in Eq. (8) join together to form C. The character string corresponding to C is matched against that of enzyme E. There are several possible alignments; we require that the string E span the binding site between A and B. Each allowed alignment of E against C is given a match score according to the number of complementary matches, i.e., the number of cases where an "a" is paired against a "b." We then compute the probability P that a score as good or better would be obtained if the strings E and C were generated at random. We use this to define a quantity we call the *specificity* $s = 1/P$. The catalytic efficiency depends on the specificity through a function $\nu(s)$. We typically assume that high specificity corresponds to higher catalytic efficiency, and choose $\nu(s)$ to be linear. For a given choice of A, B, and E, the total catalytic efficiency is the sum of the efficiencies computed for each of the allowed alignments. For a more detailed description, see Bagley.[3]

The match rule assigns a catalytic efficiency to every possible catalyzed reaction. The reactions with catalytic efficiencies above a given threshold ν_0 are installed in the reaction graph. The match rule is completely deterministic, and generates a unique chemistry. The requirement that the specificity exceed a fixed threshold implies that very short polymers cannot participate in catalyzed reactions. Thus, the properties of the match rule are quite different from those of the random rule.

At this point we have studied the random rule more thoroughly than the match rule. We intend to present results using the match rule in a future paper.

3.1.3 OTHER KINETIC PARAMETERS The other relevant kinetic parameters that may vary from reaction to reaction are the forward reaction rate k_f, the backward reaction rate k_r, and the unbinding constant k_u. In order to use the approximation for the saturation problem described in the Appendix, it is necessary that k_u be the same for all reactions. Since we can vary the catalytic efficiency for each individual reaction, this is not a serious problem.

The rate constants k_f and k_r play an important role. At equilibrium diverse rate constants cause the polymers of a given length to have nonuniform concentrations. However, since to first approximation the rate constants only depend on the two monomers at the binding site, the resulting nonuniformities in the concentration profile are much more regular and less pronounced than those resulting from

catalytic focusing. For convenience we typically assume that k_f and k_r are the same for all reactions, although in some cases we also vary them randomly.

3.2 METADYNAMICS

The number of possible reactions is infinite. Of course, in reality only a finite number are important. Unfortunately, it is usually impossible to state in advance which reactions can be neglected. A metadynamical simulation attempts to solve this problem by restricting attention to a variable set of reactions and chemical species, and adding or deleting as needed.

The problem of determining the essential reactions is particularly severe for polymer chemistry, where the number of possible reactions grows exponentially with the length of the polymers. For $m = 20$ and $n = 10$, for example, there are already more than Avogadro's number of possible catalyzed reactions. Even if we knew the rate constants for every reaction, we could never hope to keep track of all of them in a computer simulation.

A real chemostat only contains a finite number of molecules and hence a finite number of species, with a finite number of possible reactions between them. Continuous differential equations fail to exploit this. Suppose, for example, that at $t = 0$ all the concentration is placed in a few species, and the concentrations of the rest are set to zero. According to the laws of continuous kinetics, for $t > 0$ there are generically *an infinite number of species* with non-zero concentrations. This unrealistic result comes from the approximation made in treating a discrete population of molecules as though it were continuous. This is equivalent to the assumption of a well-stirred reaction vessel with infinite volume. For the problem that we address here, the fact that the reaction vessel is finite is of critical importance.

One solution to this problem is to treat the kinetics as a random process, using integer populations and simulating molecular collisions through Monte Carlo techniques. This method is efficient when the populations are low, but is computationally inefficient when the populations are high. If the population is 10^6, for example, a differential equation solver might change the concentration by as much as one percent in a single step and still retain reasonable accuracy, while the same change with a stochastic simulator requires 10^4 steps.

The metadynamics approach offers an alternative that is computationally more efficient when there is a wide disparity in population size.[2,11,12] At any given time the reaction network is modeled by a finite set of continuous differential equations, representing the dominant reactions. The topological structure of this set of differential equations is represented as a graph containing only species that are either present in the reactor, or that can be produced by other species present in the reactor. As the concentrations change the dominant species and reactions may also change. The graph is changed to reflect this, which in turn changes the differential equations. The dynamics occur on two time scales, the faster time scale of the differential equations and the slower time scale for changing the graph. A typical example of a metadynamics simulation is shown in Figure 3.

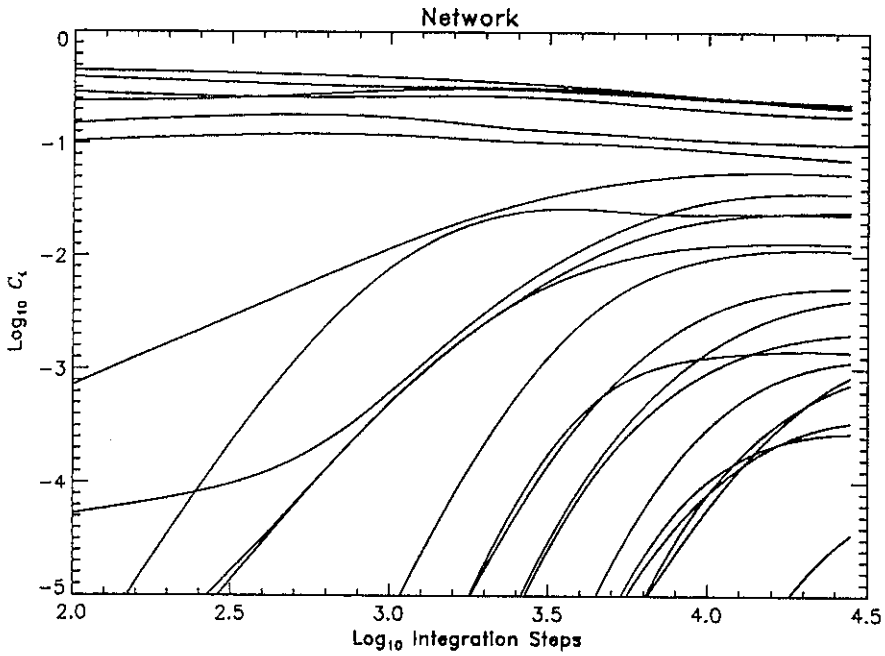

FIGURE 3 A typical metadynamics simulation. The logarithm of the concentration of each species is plotted against the logarithm of time. Initially only the four food set species a, b, ab and ba have non-zero concentration. Catalyzed reactions within the food set produce new species, which in turn have more catalyzed reactions. At the point where we begin the graph at $t = 100$, four new species have appeared, so there are a total of eight polymer species. Each new species can be seen appearing at the bottom of the graph as it crosses above the threshold. The system eventually approaches a steady state solution with 22 species above the threshold. The parameters are the same for all reactions, and are $k_f = 10^2$, $k_r = 10$, $\nu = 10^4$, $k_u = 10^4$, $H = 1$, $\delta = 10^2$, and $m_0 = 3$.

To take into account the fact that the container is finite, we impose a concentration threshold corresponding to the presence of a single molecule. If the concentration of a given species is significantly below this threshold, then that species is unlikely to be present. Therefore it cannot participate in reactions that produce other new species, and these reactions can be safely ignored. We only include reactions between species that are above threshold. They may produce new species, which rise above threshold; when this happens the new species are installed in the graph and allowed to catalyze new reactions. Similarly, species that were formerly above threshold may fall below it and be removed.

Since the number of species with concentrations above threshold at any given time is finite, the graph is also finite. Adjusting the threshold makes it possible to keep the graph from becoming unmanageably large—smaller reaction vessels have higher concentration thresholds, corresponding to smaller graphs. By making the reaction vessel sufficiently small, we can insure that the graph has less than a given number of elements. We typically strive to simulate a container that is roughly the size of a bacterium, although because of limitations in computer resources we are often forced to use smaller containers.

In our simulations the system always approaches a unique fixed point. This appears to be true independent of the way we model the reactions, i.e., whether or not we account for saturation or whether or not we include pyrophosphates. This suggests that there is a Lyapunov function for catalyzed kinetic equations in this class.[5]

We are able to speed up our simulations by several orders of magnitude by assuming the existence of a unique fixed point, which implies that the steady-state deterministic equations can be solved algebraically. We call the fixed point for any given set of differential equations a *dynamical fixed point*. The metadynamical simulation is simplified as follows: We find the dynamical fixed point of the current set of differential equations. We then examine it and check to see whether any species have moved above or below threshold in comparison to the previous dynamical fixed point. If so, we change the differential equations accordingly and find a new dynamical fixed point, and repeat the process. Eventually there are no changes compared to threshold and this procedure stops. We call the final dynamical fixed point the *metadynamical fixed point*. By reducing the metadynamics to a sequence of algebraic operations, on a typical workstation a simulation such as that of Figure 3 is compressed into a few seconds.

Note that although the deterministic simulation described above always reaches a metadynamical fixed point, once we reincorporate the stochastic effects of spontaneous reactions, the system may hop between many different metadynamical fixed points and the long-term dynamics become quite interesting. This is the basis for our claim that autocatalytic metabolisms can evolve, and is discussed in more detail in a companion paper.[4]

3.3 THE BACKGROUND OF UNCATALYZED REACTIONS

Restricting attention solely to the reaction graph is a good assumption as long as the reactions on the graph dominate over everything else. This is not always the case. For example, spontaneous reactions always dominate near equilibrium. To study the competition between catalyzed and spontaneous reactions, the spontaneous reactions must be taken into account, at least as an aggregate.

[5] These equations are reversible, which makes them different from many other autocatalytic equations that are known to display limit cycles, chaos, and hysteresis.

For the purposes of this discussion it is convenient to divide the chemical species into those of the catalytic network and those of the background, which are produced only by spontaneous reactions. This is illustrated in Figure 4. The *shadow* is a special subset of the background, consisting of species that can be directly produced by reactions involving only members of the reaction network. In a situation where the

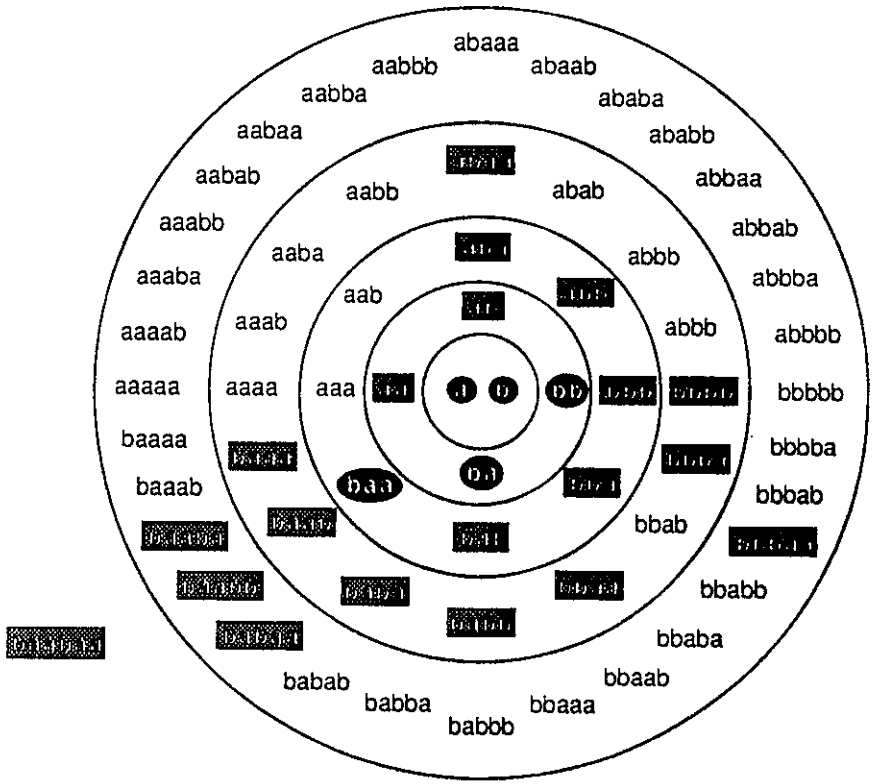

FIGURE 4 The "mandala" of polymer species. The *catalytic network*, shown with white letters on black ovals, includes the food set, but also consists of other species produced by catalytic reactions. There must be a continuous path in the corresponding graph of catalyzed reactions from members of the food set to each member of the catalyzed network; furthermore, the concentration must be above a concentration threshold corresponding to the presence of a single molecule. The *background* consists of everything that is not in the catalytic network. The *shadow*, shown by white letters on a grey background, is a special subset of the background, consisting of species that can be produced by a spontaneous reaction involving only themselves and members of the catalytic network.

concentrations of the reaction network are high, the shadow plays a special role because it is maintained at concentrations above the rest of the background. The shadow is in this sense similar to Eigen's quasi-species.[10]

For clarity we will first discuss the problem of modeling spontaneous reactions alone, without catalyzed reactions, and then return to discuss the case when spontaneous reactions are in competition with catalyzed reactions.

The difficulty of modeling spontaneous reactions comes about because there are so many of them. To make the problem tractable, we assume that all reactions have the same forward and backward rate constants k_f and k_r. The problem is then equivalent to that of modeling the spontaneous reactions in a system with only one distinct type of monomer ($m = 1$). We can lump together all polymers of a given length n, adding together their concentrations to get a combined concentration s_n. The allowed reactions are of the form

$$s_i + s_j \rightleftharpoons s_k + H. \tag{18}$$

A reaction is only possible if $i + j = k$. The contributions to the kinetics are now described in the order that the corresponding terms appear in Eq. (19): s_k receives contributions from condensation reactions between smaller species and from cleavages of longer species. It loses to condensations between polymers of length k and polymers of other lengths, as well as to cleavages of itself. If there are polymers of length k that are part of the food set, s_k will also gain from external driving. The resulting rate equation for s_k is

$$\frac{ds_k}{dt} = k_f \sum_{i+j=k} s_i s_j - k_r \sum_{n>k} (n-k+1)s_n - k_f s_k \sum_{j=1} m^j s_j - k_r(k-1)s_k + n_k \delta s_k - K s_k, \tag{19}$$

where $\delta = 0$ if k is outside the food set, and n_k is the number of elements of length k in the food set. For an alphabet of m monomers, with the assumption of uniform reaction rates, the concentration computed above is divided evenly among all the species present. It is

$$a_k = \frac{s_k}{m^k}. \tag{20}$$

We now outline our approach for treating the competition between spontaneous and catalyzed reactions: Let s_k now represent the sum of the concentrations of the polymer species of length k, excluding the catalytic network. The dynamics of the background due to its interactions with itself can still be taken into account by equations of the form of Eq. (19). However, the concentration for a single species is now defined to be

$$a_k = \frac{s_k}{(m^k - l_k)}, \tag{21}$$

where l_k is the number of elements of the catalytic network of length k. The coupling to the autocatalytic set can be taken into account by carefully counting all of the interactions. The members of the catalytic network contribute to the spontaneous background through their condensation and cleavage reactions. Similarly, they receive concentration from the spontaneous background. The only approximation necessary comes from the assumption that all the elements of the background have the same concentration. This is not strictly true. In particular, the concentrations in the shadow are typically above those of the rest of the background. However, this is a second-order effect, and we feel that in most cases our simulations are approximately correct. The problem of coupling the catalytic network to the spontaneous background is discussed in more detail by Bagley.[3]

4. EMERGENCE

In this section we present simulations of catalytic reaction networks under several different conditions. We begin by demonstrating that, under appropriate circumstances, autocatalytic metabolisms emerge at concentrations that are several orders of magnitude higher than those of the background. We explore the parameter dependence and show that there is a range of parameter values where autocatalytic metabolisms thrive.

4.1 DEPARTURE FROM EQUILIBRIUM

In Figure 5 we show a sequence of three simulations in which we drive the system further away from equilibrium by increasing the parameter δ. In each case we let the reaction vessel approach its steady-state behavior, and then plot the concentration of all the polymers in the vessel of length less than twelve. In Figure 5(a) $\delta = 0.01$, the mean reaction number $r \approx 210,000$, and the system is nearly at equilibrium. The observed behavior is in quantitative agreement with the predictions of subsection 2.2. The concentration falls off exponentially with length, all species of a given length have the same concentration, and there is no interesting structure in the concentration profile.

In Figure 5(b) $\delta = 10^5$, which corresponds to a mean reaction number of approximately 33. The concentrations of the catalyzed reaction network are orders of magnitude above those of the background, and most of the mass of the system is concentrated in the autocatalytic metabolism.

Finally, in Figure 5(c) we show the case when $\delta = 10^{7.5}$, corresponding to a mean reaction number of roughly 0.5. The dominance of the autocatalytic metabolism is less evident—the flow of matter through the system is so high that the reaction network has a much smaller effect on the composition of the system. This is because the flow of mass through the system is so large that there is no time for a given species to react before it is flushed from the system.

These simulations demonstrate the ambiguity of the word "non-equilibrium" in this context. On one hand, the parameter δ is a control parameter that can be used to drive the system away from equilibrium. On the other hand, as demonstrated in Figure 5, increasing δ does not necessarily make the *physical properties* of the

system deviate further from those at equilibrium. In terms of the non-uniformity of the concentration profile, $\delta = 10^5$ produces a larger deviation from equilibrium properties than $\delta = 10^{7.5}$. The deviation from equilibrium properties is at a maximum when the driving away from equilibrium is finite. A flow of energy is needed to move the system away from the structureless equilibrium state, but too large a flow of energy again results in a structureless state.

We have explored several quantitative measures of the deviation of the physical properties of the system from those at equilibrium. One of these is the steady-state slope Λ of the concentration profile. In Figure 6 we plot the concentrations of the species in the network as a function of their length and compare them with the background, for a favorable set of parameters where the system exhibits an autocatalytic metabolism.[6] The logarithm of concentration as a function of length in Figure 6 gives roughly lines of slope Λ, indicating that the concentrations decrease roughly exponentially with length. To measure Λ we use a least-mean-squares fit. The longer polymers of the network are at much higher concentrations than those of the background. Consequently, for the network Λ is much larger (less negative) than that it is for the background. Λ thus provides a measure of the deviation from the equilibrium profile.

In Figure 7 we plot Λ as a function of the mass flow δ. For small values of δ the system is near equilibrium, and Λ for the network is nearly equal to Λ for the background. As δ increases Λ increases, reaches a maximum, and then decreases. Near the maximum, Λ approaches zero, indicating the concentration profile is almost flat.

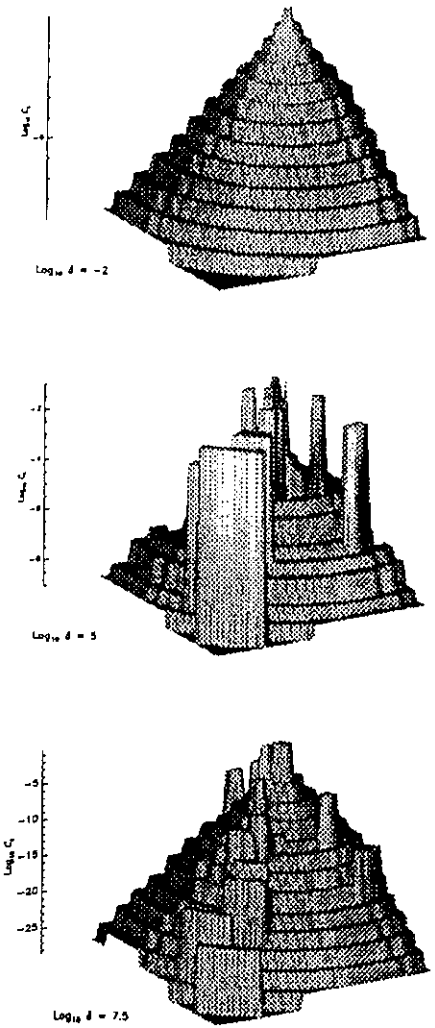

FIGURE 5 Concentration "landscapes" at different displacements from equilibrium. The two horizontal axes correspond to the polymer species, arranged concentrically with the shortest polymers in the center, as in Figure 4. The vertical axis corresponds to the logarithm (base 10) of the concentration of each species. The food set consists of the polymers a and b, using a variation of the network of Figure 2 with 118 catalytic links. In (a) $\delta = .01$, and the system is near equilibrium; the concentration falls off exponentially with length but is otherwise featureless. In (b) $\delta = 10^5$, which gives the behavior that is most distinct from that at equilibrium; in this case the concentration of the autocatalytic set is many orders of magnitude above the corresponding equilibrium values. Finally, in (c) $\delta = 10^{7.5}$, and the system is so far from equilibrium that it loses most of its interesting structure. (Note the change of scale in comparison with (b).) The other parameters are $k_f = 6.49 \times 10^2$, $k_r = 2.50$, $\nu = 8.97 \times 10^5$, $k_u = 5.00 \times 10^4$, and $m_0 = 2.0$.

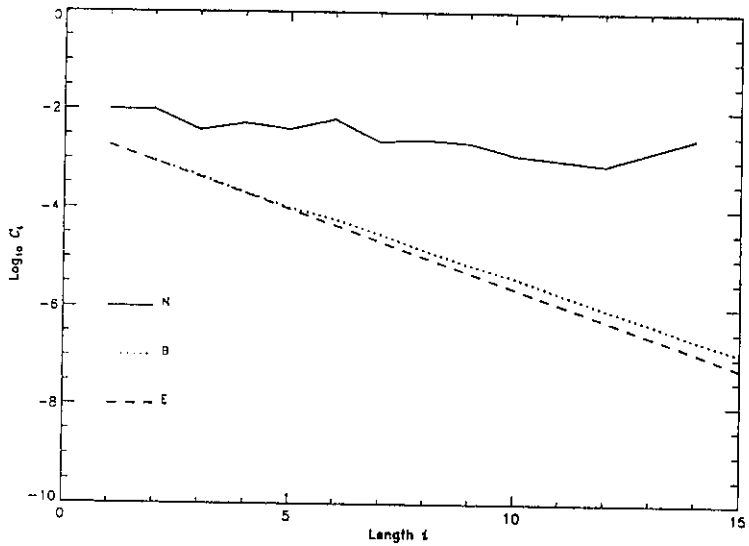

FIGURE 6 Mean concentration as a function of polymer length at steady state. The solid line (N) corresponds to the network, the dotted line (B) corresponds to the background, and the dashed line (E) corresponds to the equilibrium concentrations. The simulation is performed for the network of Figure 2 with parameters $k_f = 6.49 \times 10^2$, $k_r = 2.50$, $\nu = 8.97 \times 10^5$, $k_u = 5.00 \times 10^4$, $\delta = 1.79 \times 10^1$, and $m_0 = 2.0$.

[6]Note that this network has different parameters than those of Figure 5.

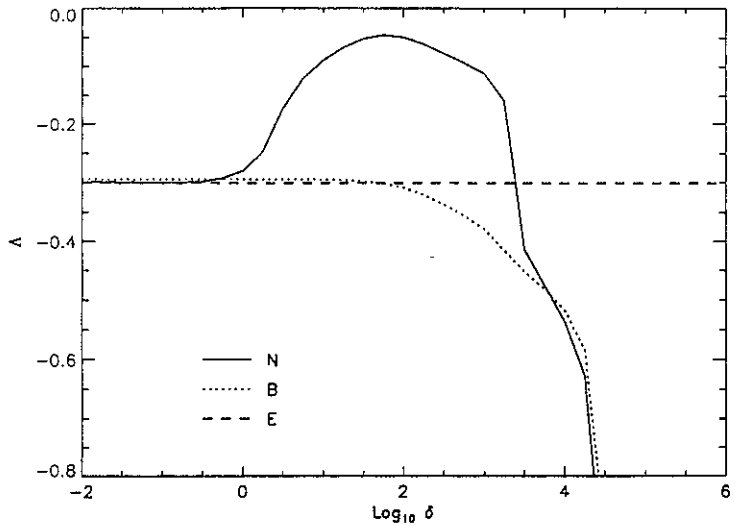

FIGURE 7 The slope of the concentration profile vs. the mass flux δ. (See Figure 6 for an example of a concentration profile.) The network (N) corresponds to the solid curve, the background (B) to the dotted curve, and equilibrium (E) to the dashed curve. Parameters are the same as those for Figure 6, except that δ is varied.

The mass concentrated in the metabolism provides another natural measure of the deviation of its properties relative to those at equilibrium. For example, in Figure 8 we plot the fraction of the mass in the background, the food set, and the catalytic network (subtracting the food set), as a function of δ. There is a central regime where the majority of the mass of the system is concentrated in the autocatalytic metabolism. Note that this regime overlaps with the regime where Λ is large. However, the two are somewhat skewed; Λ peaks at roughly $\delta = 10^2$, while the mass peaks when $\delta > 10^3$.

We feel that the need for a balanced energy flow for "interesting behavior" reflects a general principle. Another possible example is the fact that life evolved on Earth and not on Mercury or Pluto.

4.2 DEPENDENCE ON PARAMETERS

How special are the parameters for which autocatalytic metabolisms occur? To answer this question as quantitatively as possible, we have systematically explored the parameter space, testing for the presence of autocatalytic metabolisms. We used two measures of the dominance of the autocatalytic set. One of these is Λ, and

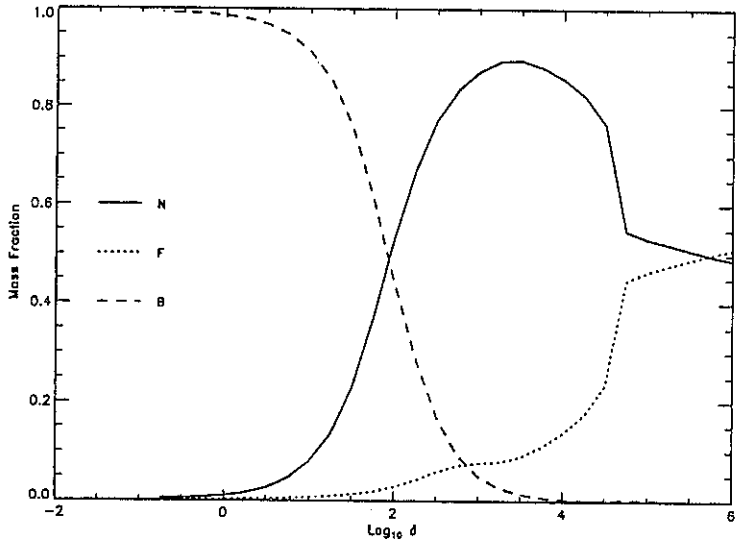

FIGURE 8 The fraction of the total mass for the network (N), the background (B), and the food set (F). (The fraction for the network excludes the food set.) The simulation is the same as that of Figure 7.

the other is the mass ratio $R = N/(B + F)$, where N is the mass of the network (neglecting the food set), B is the mass of the spontaneous background, and F is the mass of the food set.

In Figure 9 we show the behavior of Λ under variations of δ and ν. Note that Λ remains near zero for a broad range of parameter values. For comparison, in Figure 10 we plot the mass ratio R as a function of the same parameters, but for a network with fewer catalytic links. The behavior is more sharply peaked, but there is a broad range in which the autocatalytic set contains the majority of mass in the reaction vessel.

In Figure 11 we show the behavior of Λ under variations of ν and the unbinding rate constant k_u. This figure illustrates how distinct the behavior of the autocatalytic metabolism is from that of the background. In one regime, roughly corresponding to lower values of ν and lower values of k_u, Λ behaves just as it does for the spontaneous background. It is more negative and forms a relatively flat surface as a function of parameters. The other regime, which corresponds to larger values of the two parameters, has higher values of Λ. The transition from one regime to the other is quite sharp.

Finally, to illustrate the effect of varying the forward rate constant k_f, in Figure 12 we show the effect of varying k_f and δ. Once again we see a broad parameter

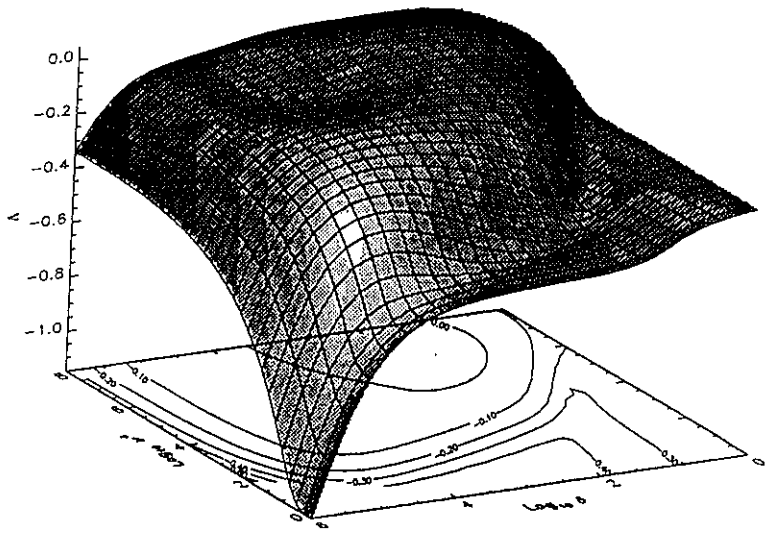

FIGURE 9 Λ vs. $\log_{10}\delta$ and $\log_{10}\nu$ for the variation of the network of Figure 2 with 118 catalytic links. The parameters are $k_f = 3.02\times10^4$, $k_r = 2.70$, $k_u = 7.11\times10^4$, and $m_0 = 2.0$.

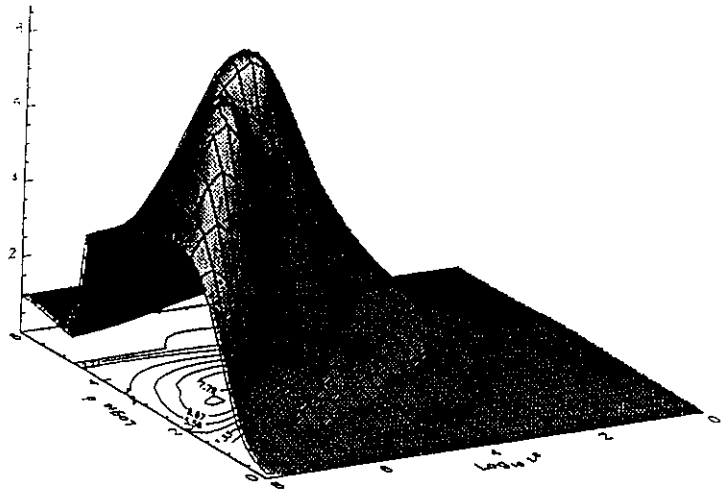

FIGURE 10 The mass ratio $R = N/(B+F)$ vs. $\log_{10}\delta$ and $\log_{10}\nu$ for the network of Figure 2. $k_f = 6.49 \times 10^2$, $k_r = 2.50$, $k_u = 5.00 \times 10^4$, $m_0 = 2.0$.

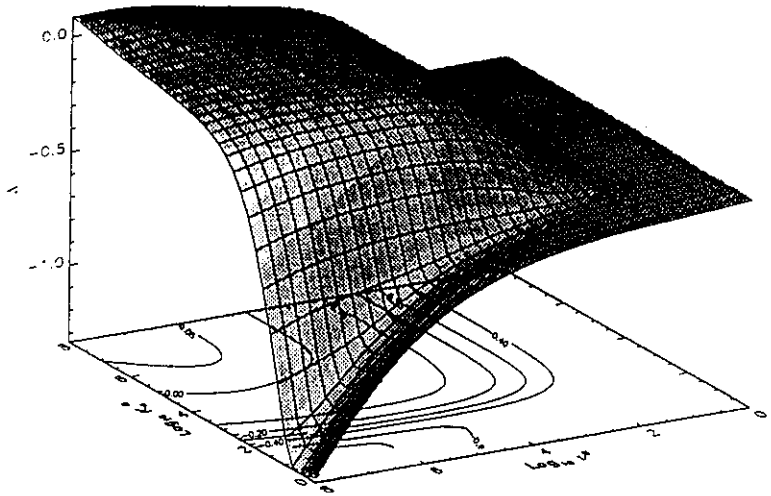

FIGURE 11 Λ vs. $\log_{10}\nu$ and $\log_{10}k_u$, for the variation of the network of Figure 2 with 118 catalytic links. $k_f = 3.02 \times 10^4$, $k_r = 2.70$, $\delta = 1.41 \times 10^2$, and $m_0 = 2.0$.

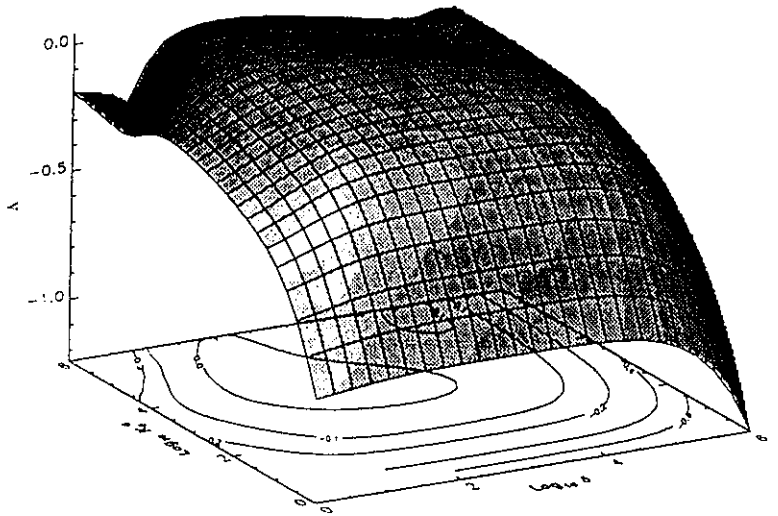

FIGURE 12 Λ vs. $\log_{10}\delta$ and $\log_{10}k_f$ for the variation of the network of Figure 2 with 118 catalytic links. $k_r = 2.70$, $\nu = 5.26 \times 10^5$, $k_u = 7.11 \times 10^4$, and $m_0 = 2.0$.