Modeling adaptive biological systems

R.J. Bagley, J.D. Farmer, S.A. Kauffman[a], N.H. Packard[b], A.S. Perelson and I.M. Stadnyk[c]

Theoretical Division and Center for Nonlinear Studies, Los Alamos National Laboratory, Los Alamos, NM 87545,
[a]Department of Biochemistry and Biophysics, University of Pennsylvania School of Medicine, Philadelphia, PA 19104,
[b]Center For Complex Systems Research and Department of Physics, University of Illinois, Champaign-Urbana, IL 61826
and [c]Computer Science Department, University of Michigan, Ann Arbor, MI 48103 (U.S.A.)

During the evolution of many systems found in nature, both the system composition and the interactions between components will vary. Equating the dimension with the number of different components, a system which adds or deletes components belongs to a class of dynamical systems with a finite dimensional phase space of variable dimension. We present two models of biochemical systems with a variable phase space, a model of autocatalytic reaction networks in the prebiotic soup and a model of the idiotypic network of the immune system. Each model contains characteristic *meta-dynamical* rules for constructing equations of motion from component properties. The simulation of each model occurs on two levels. On one level, the equations of motion are integrated to determine the state of each component. On a second level, algorithms which approximate physical processes in the real system are employed to change the equations of motion. Models with meta-dynamical rules possess several advantages for the study of evolving systems. First, there are no explicit fitness functions to determine how the components of the model rank in terms of survivability. The success of any component is a function of its relationship to the rest of the system. A second advantage is that since the phase space representation of the system is always finite but continually changing, we can explore a potentially infinite phase space which would otherwise be inaccessible with finite computer resources. Third, the enlarged capacity of systems with meta-dynamics for variation allows us to conduct true evolution experiments. The modeling methods presented here can be applied to many real biological systems. In the two studies we present, we are investigating two apparent properties of adaptive networks. With the simulation of the prebiotic soup, we are most interested in how a chemical reaction network might emerge from an initial state of relative disorder. With the study of the immune system, we study the self-regulation of the network including its ability to distinguish between species which are part of the network and those which are not.

Kewords: Complex systems; Self-organization; Adaptation; Evolution; Modeling; Networks.

1. Meta-dynamics for adaptive networks

We present two models of biological systems, each representing a particular approach to investigating an outstanding problem of biology. The emergence of life from the prebiotic soup is studied with simulations of autocatalytic reaction networks of polymers. The network is autocatalytic because every reaction between polymers is catalyzed by a polymer in the network. The immune system is studied with a model of an idiotypic network. This model emphasizes that antibodies can themselves be identified as being anti-gens and thus be recognized by other antibodies. This type of self interaction may provide some degree of control of the immune response to a foreign antigen.

Models of biological systems vary greatly in their level of description. There are models of pattern formation, morphogenesis, and perception which describe the phenomenological behavior of systems but which do not attempt to explicitly describe each individual element of the system. Other models, such as those for a sequence of chemical reactions, usually quantitatively measure the time evolution of the important properties of each

element of the system. The models presented here combine these two approaches by exhibiting collective behavior which is determined by the interactions of the basic elements of the system.

The autocatalytic network model and the immune network model have the same basic design. Both have microscopic components that are coupled together. They are characterized by a structure variable, and a strength variable that represents the concentration of the component in a well stirred aqueous environment. All concentrations vary in time due to a coupling determined by the structure variables, which serve the role of representing microscopic physical and chemical properties. In both the autocatalytic network and the immune network, the structure variable is a bit string representing the shape of a polymer.

These systems have two different types of temporal variation. The first is adjustment of the concentration of each component in response to the couplings. This adjustment is implemented with a set of ordinary differential equations for the concentration variables. The second is an innovative process, the addition of new components into the system. The addition of new components implies that the set of ordinary differential equations must change, both by adding new equations, and by adding new couplings to the old equations. Thus the dimension of the state space of the concentration dynamics can grow with time. Innovation generally does not happen completely randomly, but in accordance with a rule that specifies how the new structure variables are formed, and how they are coupled to pre-existing ones. Since this rule builds the ordinary differential equations that govern the dynamics of the concentration variables, we call it the *meta-dynamics* of the system.

Each system contains a network of interactions, and the topological structure of the network may be represented as a *graph* with components represented by *vertices* and interactions between components with *edges*.

Each network will have a characteristic topology which can be studied with the methods of graph theory (Harary, 1969; Bollobas, 1979). We attain a more complete description of the system by attributing dynamical information to the graph, with the concentration and structure variables placed at the vertices of the graph, and coupling strengths placed at each edge. Adjustment changes the concentration variables at each vertex, and innovation changes the topology of the graph by adding new vertices or new edges between vertices. Macroscopic, collective behavior is observed in the overall time evolution of the system caused by both adjustment and innovation.

In our discussion of autocatalytic reaction networks, we include an explicit rendering of the data structures used in our simulations to represent the networks. The methods for ascribing data structures to components and types of interactions, and for recording the relationships between data structures are similar in both models described in this paper. Therefore, we do not include a parallel discussion for the immune system.

2. Autocatalytic reaction networks

2.1. A definition

In a single autocatalytic reaction, one of the reactants also catalyzes the reaction. Though a catalyst does not affect the equilibrium concentration of the reactants, a catalyst does serve to lower the activation energy barrier of the rate limiting steps of the reaction (Atkins, 1978). Since the rate of a reaction is dependent upon the concentration and effectiveness of its catalyst, catalysts can be said to regulate reactions.

An autocatalytic reaction *network* contains molecules and reactions between molecules, with each reaction catalyzed by some molecule in the *network*. The catalyst for any reaction is rarely a reactant in that reaction, and thus there are few simple autocatalytic reactions in the network. Instead, the catalysts are at arbitrary locations in the

reaction graph. Given that each reaction may have one or several catalysts and that the influence of the reactions a molecule catalyzes upon the molecule itself is indirect and variable, the internal regulation of the reactions of the network by the catalysts of the network is much more complex than that of a simple autocatalytic reaction.

In our studies of autocatalytic reaction networks, the molecules are linear polymers. The allowed reactions consist of the condensation of two polymers to form one larger polymer, or the cleavage of one polymer to form two smaller polymer products. All reactions in the network are constrained to be reversible, so that each reaction represents both a cleavage and a condensation. We assume that the effect of all catalysts is simply to reduce the activation barrier between reactants, so that a polymer which catalyzes the (forward) condensation reaction must also catalyze the (reverse) cleavage reaction. In Eqn. (1), A, B, C and E each represent a polymer. In the forward reaction, A and B react to form C, while in the backward reaction, C splits into A and B. Both reactions are catalyzed by E.

$$A + B \overset{E}{\rightleftharpoons} C \qquad (1)$$

Autocatalytic reaction networks are potentially adaptive systems, and have interesting features as dynamical systems and as networks. In the following section, we discuss the relevance of studies of autocatalytic reaction networks to the problem of the origin of life. We proceed to a description of the model and its simulation, combined with a presentation of the data types which we employ in our code to represent the network during a numerical simulation. We conclude the presentation of autocatalytic reaction networks with a discussion of some investigations and future work.

2.2. One step in the emergence of life: a short review

A standard assumption among evolutionary scientists is that life initiated in an aqueous environment frequently termed the *prebiotic soup*. This soup is thought to have been populated by monomers and very small polymers which were produced through a series of spontaneous reactions. How did we make the transition from the simple soup of monomers to the complicated metabolisms and self replicating structures found in contemporary organisms?

We believe that this transition occurred through a series of steps involving increases in the complexity of the catalytic interactions in the soup, a boot-strapping procedure in which both the polymers and their interactions become increasingly more sophisticated. These autocatalytic sets 'replicate' themselves without templating, by assimilating elements of a food set; they implicitly store information; they can evolve. While they are clearly very different from contemporary living organisms, they share enough of their properties that we will often refer to them as *proto-organisms*.

In contemporary organisms large polymers act as catalysts, templates, components of metabolisms, and structural elements, among other roles. Thus we know with hindsight that the development of a mechanism for the synthesis of large polymers is an essential step in the evolution of life. But from the point of view of a successfully emerging proto-organism, polymers are important for two reasons. First, polymers have properties which allow them to perform functions which small molecules cannot do, or can do only poorly. Second, these functions can be performed on other polymers, so that a network of polymers operating upon each other establishes a new level of organization and collective behavior. We want to provide a plausible mechanism for the emergence of a polymer network, as well as explore the collective behavior of the network.

In developing a mechanism for synthesis of large polymers, the proto-organism must respond to many constraints. Of the most important of these constraints is the need for

energy. Like the spontaneous synthesis of monomers, the synthesis of large polymers in concentrations far from expected equilibrium values requires activation energy. The synthesis of monomers is thought to have made use of energy directly, such as from lightning or a thermal bath, but the gross production of polymers requires that an energy gradient be tapped and energy stored in a form usable by the proto-organism. The proto-organism must therefore develop a primitive metabolism which will evolve with the organism as energy usage becomes even more sophisticated.

A possible precursor to a primitive metabolism is polyphosphate, a molecule which releases energy when a monomer phosphate is cleaved from the parent molecule. R. Fox (R. Fox, 1988) has suggested several energy gradients in the prebiotic environment that would have led to the formation of energy retentive polymers such as polyphosphates. Molecules such as ATP currently play this role. Modern experiments have verified that a polymer can be activated by transferring a monomer phosphate from a polyphosphate to an end site of the polymer (Paecht-Horowitz and Katchalsky, 1973). The activated end of the polymer then proceeds to ligate to other monomers or polymers at a much more favorable rate than would be the case for non-activated polymers. Such a scheme may have led to the formation of polymer networks, though this mechanism for distributing energy to the network cannot be said to be metabolic until the formation of the energy retentive polyphosphate is controlled in some way by the polymer network. An example of the development of a metabolism would be if the synthesis of polyphosphates and the activation of a polymer by a polyphosphate was specifically catalyzed by a polymer of the network.

Another problem is that the production of polymers of the emergent polymer network must be favored over other polymers which might be produced via uncatalyzed collisions between molecules in the prebiotic soup. Linear biopolymers such as proteins and nucleic acids are identified by a sequence of component monomers. If there are n different kinds of monomer available, then there are $\sum_{l=1}^{L} n^l$ possible combinations of monomers up to a maximum size polymer of L monomers. For example, $n = 20$ and $L = 100$ results in $\sim 10^{130}$ possible molecules. Given that there is some finite amount of energy and material in the environment, the alternatives are either to make all possible polymers in very small concentration, or to make only a few of those possible in high concentration. The latter possibility is the most favorable to the development of the network, especially if the polymers of the network are particularly good catalysts for the reactions of the network.

This problem of how to focus resources is partially overcome by the fact that the reaction affinity of a monomer is a function of the type of monomer which is the co-reactant. For example, the experimental thermalization of a pool of different amino acids in equivalent concentrations resulted in the synthesis of a hexamer protein in a concentration 10^4 times in excess of what it would be if all possible polymers were equally favorable (S. Fox and Wang, 1968). Thus, some concatenations of monomers are naturally favored over others. An additional means of favoring the synthesis of the polymers of the network is to catalyze the reactions of the network. Since catalysts increase the rate of a reaction, catalyzed reactions makes quicker use of the available resources. Hence, the reactants of catalyzed reactions are more robust than those of uncatalyzed reactions.

Two principal considerations in devising a consistent theory for the emergence of polymer networks are (1) the choice of mechanism for synthesis and (2) the choice of material. In the biochemistry of contemporary life, the transcription of nucleic acids is achieved via templating upon nucleic acids, a process catalyzed by proteins. Meanwhile, proteins are synthesized via ligation reactions moderated by a complicated translation apparatus of nucleic acids and proteins. The monomeric

composition of a protein is a direct reflection of the monomeric composition of the parent nucleic acid, and it is the genotypic information stored in the nucleic acid which varies under selective pressure. But without very sophisticated catalysts replication via templating is error prone, and the information stored in a template would be lost after only a few generations (Radman and Wagner, 1988; Orgel, 1963). In present life, catalysis with proteins and templating with nucleic acids are mutually sustaining mechanisms. In a prebiotic soup where there are initially no polymers to act as templates or as catalysts, we must at least explain how either templates or catalysts emerged independently.

One proposal is that the first polymer network was composed of nucleic acids which replicated via templating without catalysts. However, only oligonucleotides have been produced in templating reactions without catalysts, and not in prebiotic conditions (Joyce, 1987). If catalysts are required, we can imagine that one of many different proteins somehow spontaneously generated from an existing template might have been able to catalyze its own formation, thereby forming an autocatalytic loop (Loomis, 1988). It would be at this point that the sequence of bases which represent a nucleic acid can be said to store information, and the relationship between the composition of the nucleic acid and the protein is established. But what would serve as the first template? Clays have been suggested as possible templates for the synthesis of polymers in the prebiotic environment (Cairns-Smith, 1982). However, there is as yet little evidence that the local structure of the clay can be reflected in the structure of a product polymer after synthesis by templating, so that clays at best would aid the generation of templates of random composition. Experiments have successfully demonstrated polymerization on clay surfaces (Paecht-Horowitz et al., 1970), but once formed, the polymers do not readily detach from the clay. At this point, the need for catalysts and the synthesis of the first templates are problems which make the template first theory implausible. Much like the onetime pre-eminence of anaerobic bacteria among life forms, there must have been a period when a mechanism other than the direct templating from polymers or clay surfaces was the fundamental process sustaining robust polymer populations.

An alternative hypothesis is that polymer populations emerged and were sustained via catalyzed condensation reactions (Calvin, 1969; Dyson, 1985; Farmer et al., 1986a; Kauffman, 1986; Watson et al., 1988). One such proposal postulates that the initial components of the prebiotic environment, monomers and perhaps dimers, could weakly catalyze condensation reactions among themselves to produce small polymers (Farmer et al., 1986a; Kauffman, 1986). Another proposal considers the low catalytic activity of monomers and dimers, and suggests that polymers up to a size which would allow a given polymer to present a spatial configuration to a potential substrate must have formed without the catalytic help of smaller polymers (R. Fox, 1988). In either case, catalytic activity may have been provided by surfaces in the environment. The products of these initial condensations would then catalyze reactions producing larger polymers, leading to an emergent network of catalyzed reactions and the production of still better catalysts. A further assumption is that the specificity of a catalyst for some substrate is strong enough to ensure a significant increase in the reaction rate. Conversely, the specificity of a catalyst for a substrate must also be weak enough so that one polymer can effectively catalyze more than one reaction, and similarly one reaction can have a plurality of effective catalysts. This latter condition for the catalytic specificity results in a reaction network which is highly connected, a property which may be crucial to the success of the network.

Assuming that the polymers first appeared in a catalytic reaction network, we next consider possible materials. Of course, a material which might eventually act as both a catalyst

and as a template would resolve many difficulties in tracing the co-evolution of templates and catalysts. Experiments have shown that ribosomal ribonucleic acid (RNA) in *Tetrahymena* (Cech, 1986a,b) and viroid RNA (Keese and Symons, 1985; Buzayan et al., 1986; Hutchins et al., 1986) have catalytic properties. In light of these experiments, it is possible that the first autocatalytic reaction network was composed of RNA molecules with latent templating capabilities. An immediate problem with the choice of RNA as the seminal polymeric material is that the shortest RNAs which have shown catalytic properties is 19 bases long (Uhlenbeck, 1987), so we still must explain how such a relatively complex molecule was generated without catalysts.

Another possible choice for the first polymeric material is proteins, which are easier to synthesize than nucleic acids. For example, the synthesis of a nucleic acid monomer, a nucleotide, requires three dehydration reactions to bind the component sugar, base (purine or pyrimidine), and phosphate. This is thermodynamically less favorable than the formation of many types of the amino acids which combine to form proteins. Some amino acids form readily in prebiotic conditions (Miller and Urey, 1959). Also, sugars are less stable in solution than amino acids, and bases are difficult to synthesize under prebiotic conditions (Miller and Orgel, 1974). Another argument in favor of proteins is that even small proteins and amino acids demonstrate at least some catalytic activity, and thus it is possible that an autocatalytic reaction network developed from a set of molecules of minimal complexity. Polypeptides formed from the catalytic reaction network might then catalyze the formation of nucleotides and nucleic acids. With this scenario, the catalytic reaction network provides a scaffolding upon which the templating mechanism develops.

There are other considerations which favor proteins. Experiments have shown that proteinoids, multi-branched polymerizations of amino acids, spontaneously self-assemble into microspheres (S. Fox et al., 1959). R. Fox (1988) has proposed that a primitive vesicle would protect important chemicals not in abundance and may have provided the interface for the coupling of oxidation-reduction energy to the synthesis of polyphosphates. A microsphere made of protein polymers produced by the reactions of an autocatalytic reaction network is a plausible candidate for a primitive cell membrane.

The issue of what material constituted the first autocatalytic reaction networks can only be settled by experiment or much more detailed computer simulations than those presented here. Our model is sufficiently general that we can simulate either proteins or nucleic acids by simply changing a few parameters. More important to our point of view, the use of autocatalytic reaction networks to model the emergence of a proto-organism emphasizes two points. First, catalysis or something like catalysis is required to favor the continued production and hence survival of a small set of molecules. Second, the function of catalysis is assumed by the molecules of the network, which implies that the network exerts a sophisticated control over the reactions which produce the polymers of the network. This self-regulation by the network is ultimately related to the way in which the network adapts and exhibits collective behavior.

2.3 Description of the model: assumptions and approximations

For a polymer network to be successful, the concentrations of the polymers of the network must be much greater than any other polymers in the environment, and the reactions of the network must win the competition for resources. As we have discussed, the increase in the rates of the reactions afforded by catalysis achieves this end. We assume that catalyzed reactions are so much faster than spontaneous reactions that spontaneous reactions can be neglected. At some point in the development of the network, the resources and catalysts needed to

sustain a reaction or set of reactions may be redirected to a newly created portion of the network. The rate of the vestigial reactions of the network will approach that of spontaneous reactions. In order to keep the dimensions of the system as small as possible, we delete reactions and polymers from the network when they have effectively died out.

Presumably catalysts have evolved to become more effective than the primitive catalysts of the first polymer networks (Koshland, 1987). Contemporary catalysts have also become more specific, meaning that a particular catalyst will often act upon only one kind of substrate. Specificity and effectiveness are often directly related. While inferior effectiveness is a liability for primitive catalysts, a low degree of specificity implies that they may catalyze several reactions. Conversely, several catalysts can speed the rate of one reaction and in concert may insure the robustness of the reaction and its products.

The rules by which reactions are created and catalysts are assigned to reactions are an important part of the model, for these rules reflect the importance of characteristics of the polymers of the network upon the evolution of the network. Since it is too difficult to predict the chemical properties of a polymer based upon its monomeric composition, we must make approximations in the assignment of reactions and catalysts. The simplest metadynamical rule creates a condensation reaction by selecting two chemicals at random, and conversely a cleavage is created by selecting a polymer and the hydrolysis site on that polymer randomly. Similarly, a catalyst for a particular reaction is randomly chosen from the chemical species available in the system. This simple rule of selecting reactants and catalysts randomly obviously does not even consider the properties of the polymers. An improvement is to implement several varieties of a match rule which assesses the catalytic potential of a polymer based upon the quality of the match between the substrate and a sub-sequence of the catalyst which will act as the binding site. Conversely,

reactants are chosen which will match best with available catalysts. The match rule method for assigning the values for the binding affinity of the various monomers for each other is ad hoc but reflects the asymmetrical preferences of nature. With the random rule for assigning catalysts, the efficiency of the catalyst is an arbitrary value within a defined range. With the match rule, the efficiency is chosen as a function of the quality of *match* to the substrate. Since the effectiveness of a catalyst is often correlated with size, a further improvement is to make the catalytic efficiency a function of the length of the catalyst.

An important attribute of autocatalytic reaction networks with match rules is that the effect of a variation in the monomeric composition of a polymer on catalytic efficiency can be measured. In addition, an approximation of the complementarity between the shapes of catalysts large enough to form a three-dimensional surface and substrates is registered. We can model natural variability by spontaneously introducing a new polymer into the network, and observing its effect upon the structure of the network. A novel polymer may have the important effect of seeding several new reactions, which in turn may create new polymers causing a "chain reaction" in its effect.

Reactions can only take place between molecules which are actually present in the system. The standard use of continuous equations to describe the concentrations of molecules poses a problem, since in principle it invokes an infinite number of polymer species, even though most have infinitesimally small concentrations. Yet physical systems have finite resources, so we establish a *threshold* which is the effective concentration corresponding to only one polymer molecule in the system. If a polymer species is to be a candidate for participation in new reactions as a reactant or as a catalyst, the concentration of the species must be above threshold. The number of new reactions which is assigned at every update of the reaction graph is

parameterized by a value called the *probability of catalysis*. Roughly speaking, new reactions will be created only if the concentrations of a fraction of previously created polymers have recently exceeded the threshold concentration.

Another simple assumption is that the forward rate constants for all condensation reactions are the same; similarly, the reverse rate constants for all hydrolysis reactions are equivalent. As was cited, reactions between different monomers under the same conditions in nature do not proceed equally. A better assumption regarding rate constants is to strongly favor the rates of some reactions based upon the properties of the monomers which are about to form or hydrolyze a bond with *affinity* rules. R. Fox suggests that the successful emergence of life relied upon the coincidence that kinetically favored polymers were also good catalysts (R. Fox, 1988). If true, this argument would justify attributing superior catalytic capability to the polymers which are produced in dominant concentrations.

A final approximation is the method by which we account for the amount of reactants and catalysts which are not available to react because of involvement in ongoing reactions. The rate of a reaction is dependent upon the concentration of available catalysts and reactants, so catalysts and reactants which are in intermediate complexes cannot contribute to the rate of a reaction. If all of the catalyst for a particular reaction is saturated with substate, then the unbinding of this intermediate complex becomes the rate limiting step for that reaction. We do not attempt to model all of the possible intermediates for each reaction; our simplest scheme for modeling the intermediates of each reaction would so increase the dimension of our system that we would be unable to simulate large networks. Instead, each reaction results in a product which is in a bound state, which dissociates at a prescribed rate, releasing the product and catalyst. This method requires that we model the time evolution of the amount of each poly-

mer which is free and which is in a bound state, and thus only doubles the dimension of the system.

2.4. A sample simulation: equations of motion and data structures

We simulate the development of an autocatalytic reaction network in a chemostat. The simulation is begun by initializing the chemostat with a *food set* of monomers and small polymers less than or equal to some length r. The chemostat is presumed to be well stirred, so that the time evolution of the concentration of each chemical species is computed by an ordinary differential equation. Each different monomer is represented as a letter of the alphabet $\{a, b, c, ...\}$, and polymers are represented as strings of letters, i.e. $\{cbaabac\}$. The polymers of the network are oriented, like proteins $(CO_2^- \rightarrow NH_3^+)$ and nucleic acids $(3' \rightarrow 5')$, so that baa and aab represent different polymers.

In all living systems, complicated structures are being built up from basic elements at the same time that they are also degrading. We assume that there is a supply of basic food elements which flows into the system. We simulate the flux of mass through the system by supplying the chemostat with the food set at a rate τ, while depleting all species in the chemostat as a linear function of concentration $(-\delta x_i)$. As a matter of convenience, we choose our initial conditions to conserve the total mass of the chemostat. (Even if we do not do this, the system will achieve a constant mass level at steady state.) Letting x_i^0 be the initial concentration for the ith member of the food set

$$\tau_i = \delta x_i^0 \tag{2}$$

x_i^0 is the initial concentration for the ith member of the food set.

Living systems also contain metabolisms which convert energy from the environment into a chemically usable energy source. We employ a simple scheme suggested by R. Fox (1982) to simulate a primitive metabolism.

Energy from a primitive energy gradient is stored in the form of polyphosphates. Initially, we supply only monophosphate (p) which polymerizes to form diphosphate (p_2), a reaction driven by an external source of energy (γ). Each bond of a polyphosphate molecule stores enough energy to activate a polymer of the catalytic reaction network. Once activated, polymers condense at a rate competitive with hydrolysis.

During a numerical experiment, the details of a catalytic reaction network are stored in several types of data structures. The identifying characteristics of each polymer is stored in a data structure of type **Poly**:

```
struct Poly {
    struct React *react;    /*points to linked
                            list of reactions the
                            polymer participates
                            in*/
    struct Cat *cat;        /*points to linked
                            list of reactions this
                            polymer catalyzes*/
    double conc;/           /*concentration of
                            polymer*/
    char *str;              /*representation of
                            polymer as a string
                            of letters*/
};
```

Poly contains str, a pointer to a character string which represents its monomeric composition. The real variable conc is the concentration of the polymer.

The details of every reaction between polymers of the network are stored in a data structure of type **React**:

```
struct React {
    struct Poly *c,         /*points to the data
                            structure for the
                            polymer which
                            cleaves to form A
                            and B*/
    struct Poly *a, *b;     /*points to the data
                            structure for the
                            polymers which
```

condense to form C*/
```
    struct Cat *cat;        /*points to the list
                            of catalysts for this
                            reaction*/
    double flux;            /* mass flux through
                            reaction node in
                            direction of
                            condensation*/
};
```

React contains c, a, and b which are pointers to the polymers which participate in the reaction, and *flux*, a measure of the net mass flux through the reaction node.

In a graphical representation, **Poly** and **React** represent the nodes of the graph (see Fig. 1). There are also two kinds of connections between these nodes, reactive links and catalytic links. Since a reaction may have more than one catalyst, and a polymer may catalyze more than one reaction, information about catalysts are stored in a linked list of type **Cat**:

```
struct Cat {
    struct Poly *poly       /*poly which is the
                            catalyst*/
    struct React *react;    /*reaction which is
                            catalyzed*/
    double vel;             /*efficiency of
                            catalyst*/
    struct Cat *next;       /*next in linked list
                            of catalysts*/
};
```

Cat contains a pointer to the polymer which is the catalyst *poly*, a pointer to the catalyzed reaction *react*, and a the value for the efficiency of the catalyst for that particular reaction. The variable *next* points to the next catalyst in the list. Returning to our definition of the data types **Poly** and **React**, both structures contain pointers to Cat. However, the lists of data type **Cat** referred to by **Poly** and **React** are different. In the case of **Poly**, we list all of the reactions catalyzed by a particular polymer. With **React**, we list all of
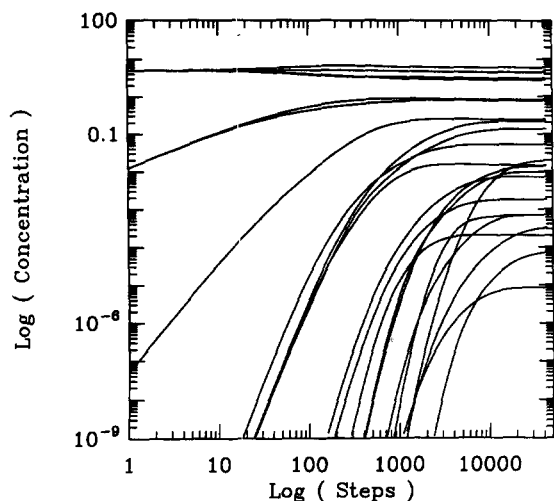
Fig. 1. A graph of a reaction network at steady state. This network resulted from a simulation which began with a food set of two different monomers represented by a and b. Each reaction is represented by an arrow drawn from the cleavage products and pointing to the condensate of the reaction. Though the catalytic relationships are not explicitly drawn, each reaction is catalyzed by at least one polymer in the network.

the polymers which catalyze a particular reaction.

At every integration step, the composite data structure is traversed and the equations of motion are reconstructed by assigning the appropriate term for each interaction to the equations of the participants. The data structure is accessed with *polylist,* which points to a linked list of data structures of type *Poly.*

After each integration step, the rules for creating reactions are applied to the polymers in the chemostat. Time series of the concentrations of the polymers of a sample simulation are shown in Fig. 2.

We have written our software in C, a structured programming language (Kernighan and Ritchie, 1988). C has several features which simplify the task of modeling meta-dynamical

Fig. 2. The time evolution of the concentrations of each species exhibiting an approach to a steady state. New polymers which are introduced with the growth of the graph emerge from the lower axis.

systems, including user-defined data types, dynamic memory allocation and the referencing of variables with pointers.

An improvement in the choice of language for our code would be an object oriented language (Cox, 1986; Stroustrup, 1986), where procedures and functions which operate on a particular data type are included in the definition of that data type. Code modification is much easier since any changes to the definition of a data type will not affect the rest of the program. In our studies of meta-dynamical systems, modifications to the simulation usually require changes in the definitions of data types. When such a change has been made with a program written in C, Fortran, or many other standard languages, we often must revise every procedure which referenced the data type. As with our simulation of autocatalytic networks, many details of a complex system under study are not known. And since the systems are usually not amenable to analysis, ideas can only be tested by rewriting and simulating the program. With object oriented languages, the process of investigating many possibilities for the system via code modification is simpler.

TABLE 1

A schematic representation of each reaction between chemical species of the network followed by the differential term each reaction generates. The respective term is subtracted from the equation of motion for each species on the left side of the reaction arrow, and added to the equation of motion for each species on the right side. x_i is the concentration for species i. A and B represent polymers which condense to form C; E is the catalyst for the reaction. A bar indicates polymers in the bound state. The rate constants are k_f for polymerization, k_i for hydrolysis, k_u for the dissociation of bound complex, k_f for polymerization of phosphate, k_r for hydrolysis, and k_u for the activation of a polymer. A $*$ indicates that the polymer is activated. k_f^* is the rate constant for the condensation of an activated polymer with another polymer. v is the effectiveness of each catalyst for a particular reaction. h represents water.

Reaction		Rate
$A^* + B + E$	$\rightarrow$ $CE + h + p$	$k_f^* v x_E x_{A^*} x_B$
$A + B + E$	$\rightarrow$ $CE + h$	$k_f v x_E x_A x_B$
CE	$\rightarrow$ $C + E$	$k_u x_C x_E$
$C + E + h$	$\rightarrow$ ABE	$k_i h v x_E x_C$
ABE	$\rightarrow$ $A + B + E$	$k_u x_A x_B x_E$
$2p$	$\rightarrow$ p_2	$k_f y p^2$
$p_2 + h$	$\rightarrow$ $2p$	$k_r p_2 h$
$A + p_2$	$\rightarrow$ $A^* + p + h$	$k_A x_A p_2$
$A^* + h$	$\rightarrow$ $A + p$	$k_i x_A h$

2.5. Pursuing the model: thermodynamics, the graph, and collective behavior

For fixed reaction graphs, simulations have shown that the reactions of the network attain a steady state. One of the first things to study are factors which affect the steady state concentrations of the polymers of the network. The influence of some factors are easy to explain. For example, dehydrating the chemostat favors condensation reactions by mass action, since the polymerization of amino acids and nucleotides release a water molecule for every bond formed. There are other relatively more complex factors which will influence the concentrations of polymers, and here we focus on five of them:

(1) the supply of activation energy mediated by polyphosphates;

(2) mass flux, which includes the influx

(driving) of the food set and the efflux (dissipation) of all species;

    (3) autocatalysis;

    (4) saturation of catalysts;

    (5) the topology of the catalytic subgraph.

One thing which we do not demonstrate directly is that the combination of mass flux and autocatalysis allows the reactions of the network to dominate over uncatalyzed spontaneous reactions in the competition for food and the continuous reproduction of structures. In a closed environment without mass flux, both catalyzed and uncatalyzed reactions will eventually achieve thermodynamic equilibrium. In our system open to mass flux, the system can only achieve a nonequilibrium state, and the enhancement of the rate of a reaction due to catalysis allows the autocatalytic network to make quick use of the resources available. Yet the stronger the influx of the food set, the more the mass in the system is concentrated in smaller polymers. In this case, catalysis only serves to restore the concentrations of polymers to an equilibrium distribution (Fig. 3). The formation of large polymers in abundance is effected with the inclusion of activation energy provided by polyphosphates (Fig. 4).

The saturation of polymers in bound complexes detracts from the action of polyphosphates, since bound polymers cannot be activated (Fig. 4). Saturation also favors small molecules simply because the portion of a polymer which is in a bound complex is not available for reaction or catalysis (Fig. 3). We demonstrate that extremely high levels of mass flux and saturation can retard the production of large polymers.

Though it is energetics which plays the largest role in shifting mass from smaller to larger polymers, the topology of the catalytic subgraph also has a secondary effect upon the distribution of mass within the network. Figure 5 shows the results of two simulations, each with the same reaction subgraph, but with different catalysts assigned to the reactions.
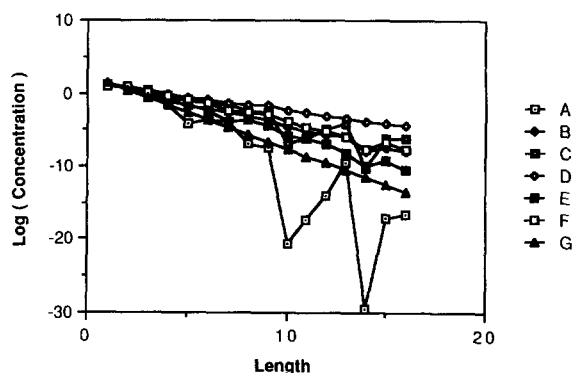
We are performing experiments on reac-

Fig. 3. A comparison of several versions of our equations of motion for the reaction network of Fig. 1 without phosphates. Shown are curves of the steady state concentrations of the polymer populations with respect to the length of the polymers. The system which generated E includes mass flux without catalysis, C includes mass flux with catalysis, and D and F are the equilibrium results without mass flux with and without catalysis. We include the equilibrium result G for the reversible condensation of bivalent monomers in which all possible reactions are considered (Flory, 1953). Comparing G with curves D and F demonstrates how the configuration of a system determines equilibrium. Curves A and B are results from versions of the model which include bound complexes.
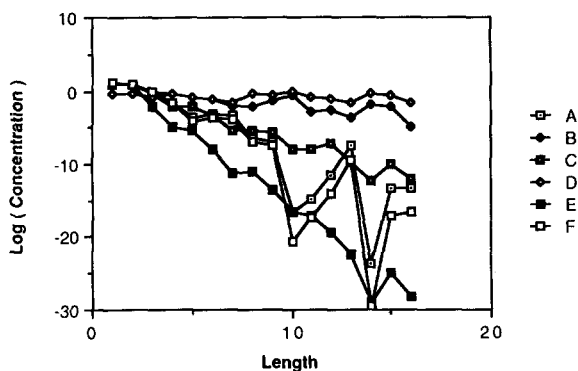

Fig. 4. A set of results for systems identical to Fig. 3 but for the inclusion of polyphosphates. Polymers activated with a phosphate condense to form larger polymers at a rate greater than the rate of hydrolysis and the rate of condensation with deactivated polymers. In comparison with Fig. 3 we identify the overall shift in mass towards larger polymers due to phosphorylation with curves B, D, and F.
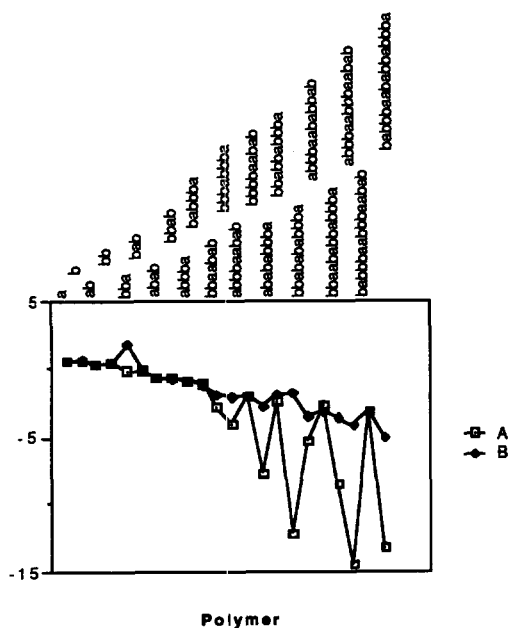
Fig. 5. Curves A and B plot the distribution of polymer concentrations at steady state for the same reaction network (Fig. 1), but with different catalytic connectivities.

tion networks which include a large number of spontaneous reactions to ascertain the minimum strength of catalysis among the components required to give the network of catalyzed reactions a sufficient advantage over uncatalyzed reactions in the competition for food. One difficulty of simulating spontaneous reactions with our present scheme is that the products of such reactions are low, and continuous equations are unrealistic when dealing with concentrations near or below the threshold. In response to this problem, we are developing a *molecular simulator* which simulates reactions as individual events between species with integer concentrations.

Quite often the supply of resources to a system is variable. If such a variation dips below the level of tolerance for a system, the system must adapt to the modified environment if it is to survive. By varying the mass and energy flux into our model chemostat, we can simulate environmental change and

measure the adaptive capacity of autocatalytic reaction networks. For example, if a network has developed several reaction pathways which lead back to the formation of a food component, then the network should remain robust in response to a reduction of the influx of that food component. Thus, a network which is being fed the food components a and ab, and contains the reaction $a + b \leftrightarrow ab$, should remain robust if a is no longer provided, since it could replenish a by hydrolyzing ab. This response is that of a metabolism because it is converting the resources available in the environment to a form useful to the network. It may also be true that this second state could be achieved only after the network had first developed with a steady supply of a, and which would indicate the presence of *hysteresis* or path-dependent states in the dynamics of the system.

In our description of the rules which we employ to assign catalysts to reactions, match rules were proposed which would include some notion of the complementarity of catalyst and substrate. However, a *strict* match rule would also be unrealistic, since it implies a continuous change in function with respect to a small change in the sequence of the catalyst. In fact, the properties of many contemporary proteins are completely changed with the variation of one residue. Thus we are exploring rules intermediate to the strict match and the random assignment rules to better simulate the real distribution of function with respect to composition. The most important aspect of matching catalysts to substrates as some function of catalyst composition is that the polymers of the network have characteristics additional to concentration and length. The sequence of monomers has meaning, retaining information which defines its usefulness as a catalyst. Thus, although autocatalytic sets do not store information in the same direct fashion that templating structures do, there is implicit information storage simply in the list of chemical species present in the soup.

Since we do not have exact knowledge

regarding the composition of the prebiotic soup, it is important to determine the minimal requirements for the emergence of an autocatalytic reaction network. Given a distribution of match strengths which reflect the energy of binding between catalyst and substrate, what is the minimum level of complexity required of the food set so as to provide enough catalysts for the initiation of the emergent network? Initially we perform our experiments without dynamics, a regime which can be analytically explored. The result of such investigations is a minimum set of criteria which is required for the connectedness of an autocatalytic network to be established. Methods from the study of random graphs (Palmer, 1985; Bollobas, 1985) are proving useful in the study of the transition to greater degrees of connectivity in the graph during the course of its growth. The inclusion of dynamics in these simulations will allow us to extend purely graphical results into a physical domain.

We are also searching for useful measurements of the graph which are predictive of network behavior. One such approach compares the topology of the catalytic reaction graph to the dynamics of the network. The topology is quantified by a number of different indices which capture some aspect of the connectivity of the graph (Bollobas, 1979; Rouvray, 1986; Trinajstic, 1983). Another approach is to measure the direction and strength of the net mass flux through a reaction node, generating a vector *flux field* on the reaction nodes of the graph. Since equilibrium systems are in *detailed balance* meaning zero net flux through each reaction node, the flux field would be another means of measuring the shift from equilibrium. It will also be apparent from the flux field which topological loops from the reaction graph are dynamical cycles, i.e. a sequence of reactions which are proceeding most favorably in one direction and which form a closed loop such as the Kreb's cycle of contemporary metabolisms.

Spatial heterogeneity within natural systems exists on several spatial scales. On one scale defined within a chemostat, the concentrations of each polymer may depend upon position. It is not tractable to study partial differential equations at the level of description in our model, and we think that our approximation that the chemostat is well stirred is viable for small enough enclosures. On a larger scale, we may consider that many autocatalytic networks emerged near to one another but experienced different environmental conditions. To study the interaction of a population of autocatalytic reaction networks, we are performing a series of numerical experiments with chemostats arranged on a lattice. Each chemostat may vary in the network it contains or in the environment it is in, and each chemostat is open to its immediate neighbors on the lattice. Networks from neighboring chemostats will compete for food and phosphates. One of our goals in such a study is to determine the better network topologies for a particular environment.

3. The immune system[*]

3.1. Introduction

The immune system is our primary defense against pathogenic organisms and as such has evolved strategies for recognizing antigens, i.e. foreign cells and molecules. The basic strategy, called clonal selection, is interesting from a theoretical viewpoint since it uses a pseudo-random process to perform specific pattern recognition tasks. Further, the immune system is capable of learning and remembering patterns that it has recognized. To understand how learning, memory and pattern recognition is accomplished, we have developed an adaptive network model of the immune system using principles of *metadynamics*. As we discuss below, one consequence of the pattern recognition algorithm that the immune system uses is that certain novel components of the immune system such as antibody cannot be distinguished from

---

[*]Portions of this section have been taken almost verbatim from Perelson (1989).

antigen. Thus the immune system may recognize and respond to itself in a self-referential way. This leads to the idea, first presented by Jerne (1973,1974), of the immune system as a network of interacting cells and molecules, normally looking inward, and being perturbed by exogenous antigens. Our adaptive network model can be used to examine in a quantitative fashion the operation of immune networks. Before going into the details, we first review some of the features of the immune system.

The basic defense cells of the immune system are a class of white blood cells known as *lymphocytes*. B lymphocytes are the cells that make antibody molecules. The other class of lymphocytes, T lymphocytes, are involved in cell–cell interactions. Some T cells are responsible for killing tumors or virally infected cells, whereas others secrete factors that promote the growth and differentiation of B cells and T cells.

Both B and T cells have receptor molecules on their surfaces that can recognize antigen. For convenience we shall speak about lymphocytes being of different types, where the type of a lymphocyte refers to the specificity of its receptor. It is estimated that the B cell population within an animal such as a mouse contains approximately 10^7 different specificity types at any time. The T cell population may be slightly less diverse. The number of different receptors, i.e. lymphocyte types, that an animal could potentially make is unknown, but, based on what we know of the genetic processes involved, the number could easily exceed 10^{10}. When antigen enters the body it encounters a large number of lymphocytes. If some of these cells have receptors that recognize the antigen these cells may become stimulated, begin proliferating and secreting molecules such as antibody that can lead to the elimination of the antigen. A single cell that recognizes the antigen can grow into a large clone of cells all of which are capable of fighting the antigen. The cells that grow into clones are selected by the antigen, hence the name clonal selection.

Although clonal selection operates at the level of both B cells and T cells, here we restrict our discussion to B cells and the antibodies they secrete. In broad outline a similar story applies to T cells.

The receptors on a B cell are made by a random genetic process during the maturation of a B cell in the bone marrow. The process involves choosing at random various immunoglobulin gene segments from gene families and joining these segments together in an error prone way. The fully assembled receptors when expressed on the cell surface advertize the capability of the B cell. Many, perhaps most, B cells have receptors that are useless and detect nothing. Such cells probably die within a few days of leaving the bone marrow. However, on occasion a randomly made receptor is useful and detects something. The cells carrying this receptor are then amplified via clonal selection, and the receptor molecule, now known to be useful, is secreted as antibody. During the proliferation of the clone point mutations occur at unusually high frequency within the genes coding for the antibody molecule. Thus further refinements in specificity can occur during the course of an immune response.

For a few weeks after exposure to antigen the cells making antibody against the antigen are maintained within the body at an elevated level and can be thought of as carrying a short-term memory of the exposure to antigen. However, as is well known, memory of certain diseases can last many years or even a lifetime. The mechanisms of long-term memory storage are not fully understood. Some memory may be carried by "memory cells", cells that have responded to antigen by proliferation and then which enter a special quiescent state in which they may live for an unusually long time. Upon re-exposure to antigen these long-lived memory cells may then become restimulated and once again grow into a clone of cells actively secreting antibody. Memory may also be carried dynamically. Rather than totally terminating, an immune response may maintain itself in a

steady-state or dynamic state (e.g. limit cycles or aperiodic trajectories) in which stimulated lymphocytes are kept at a low population level via interactions with antigen or other antibodies.

For clonal selection to work as a means of generating protective antibodies the antibody repertoire needs to be complete, i.e. the population of B lymphocytes needs to be sufficiently large and diverse that essentially any antigen will be recognized by some lymphocyte or set of lymphocytes in the population. Quantitative calculations by Perelson and Oster (1979) indicate that if antibodies are truly made at random the repertoire will be complete in all immune systems with 10^5 or more elements. Interestingly, the smallest known immune system, that of a young tadpole, is estimated to have on the order of 10^5 different antibody types.

A consequence of the repertoire being complete is that the body can recognize its own antibodies as antigens. Antibodies are just protein molecules. A novel protein made by a stimulated B cell and secreted into the blood stream looks to the rest of the immune system just the same as a novel molecule made by a bacterium or some other pathogen. If the antibody concentration gets high enough, the system recognizes it and responds to it by making another antibody that can bind to it. The unique portions of an antibody molecule that the immune system uses for recognition purposes are called *idiotopes*. The antibodies raised against other antibodies are called *antiidiotypic antibodies.**

Because the immune system makes antiidiotypic antibodies, one can imagine a cut on your finger, which allows bacteria to enter your body, stimulating an immune response which then cascades via antibodies being made against antibodies into an event potentially involving all of the white blood cells in your body. Thus there are two extreme views of an immune response. In classical clonal

selection the antigen excites only those cells that can make antibody against the antigen. This, in general, is a rather small portion of the immune system; typically one in 10^5 cells respond to a given antigenic determinant. At the other extreme, the immune system is an idiotypic network and potentially the entire network is involved in the response to any antigen. How can we assess these two extreme models? Can strict clonal selection (no network interactions) quantitatively account for immune phenomena? If not, how can one approach understanding the role of immune networks? Here we pursue the answers to these questions by formulating what we believe are realistic models of idiotypic networks and examining their properties. We shall see that as a parameter describing the probability that one antibody recognizes another antibody changes, our model predicts a phase transition; on one side of the transition networks are rather sparse and responses should be dominated by single clones or a few interacting clones. On the other side of the transition the network is highly interconnected and thus can become fully engaged. (The specific dynamic rules governing the growth of clones will determine how much of the network is involved in any particular response.)

3.2. The model

The hallmark of the immune system is its diversity and the variability of its components. The diversity arises in many ways. Each cell or more precisely clone of cells has a particular specificity defined by its receptor type. A realistic immune system contains between 10^5 and 10^8 different clones at any time. Clones may only exist for some period of time. They die and are replaced by other newly generated clones. Further, when the cells of a clone are stimulated to grow by contact with antigen, they may mutate at unusually high rates generating great diversity within the clone itself.

Modeling clonal selection presents some unusual problems. Not only do clones appear,

---

*One can remember these names by thinking of the idiotopes as the *idio*syncratic parts of the antibody molecule.

disappear and change, but each clone while present in the system is potentially able to interact with any other clone. To investigate such a diverse, dynamically changing system with a standard population dynamic model is enormously difficult. Clearly, it is not feasible to measure all of the interactions between clones and then use that information to construct a set of differential equations describing each clone. In this situation using meta-dynamics is natural. One constructs a rule and then uses this rule to assemble the appropriate dynamical equations. To illustrate the method we construct a model of an immunological network in which antibodies recognize antigen and other antibodies.

We model the immune system as an open system. The system initially contains a set of randomly made antibody molecules. The antibodies in this model can be thought of as molecules which have the property of self reproduction (via B cells which are not explicitly in the model). Alternatively, one can think of the antibodies as receptors and thus as labels for B cell types. The model can be used in many different ways. One can mimic the formation of an idiotypic network during ontogeny by allowing these antibodies (B cells) to interact among themselves in the absence of antigen. This provides a realistic means of initializing the system. Interactions of the immune system with foreign antigens can be modeled by introducing antigens into the system at fixed or random times. The formation of new B cells in the bone marrow can be simulated by inputing new antibodies into the system. A set of differential equations that determines the concentration of each antibody and antigen in the system is constructed via a meta-dynamical rule. A loss term is included in the differential equations so that antibodies that are not amplified (via clonal selection) are ultimately lost from the system. To implement a means of deleting species we introduce a threshold. If an antibody or antigen concentration decreases below the threshold we assume the last molecule of that type has been eliminated. At this point the differential

equation for that species is removed from the system of equations and all interaction terms involving the "removed" species are eliminated from the other equations in the system. The converse procedure is used when a new species enters the system. An equation for the new species is installed in the set of differential equations and a meta-dynamic rule used to generate interactions between the new species and all existing species. We now turn our attention to the development of appropriate meta-dynamic rules.

Immunological recognition is governed by chemical interactions between molecules. For an antibody or B cell receptor to recognize an antigen it must bind it. Binding between antigen and antibodies generally involves short-range non-covalent interactions based on electrostatic charge, hydrogen binding, van der Waals interactions, etc. In order for the molecules to closely approach each other an appreciable portion of their surfaces need to be complementary In some cases, the complementary regions may be planar, while in others they more closely resemble a bump and a groove. Both shape and charge distributions, as well as the existence in the appropriate complementary positions of chemical groups that can form hydrogen bonds and interact in other ways, are properties of antigens and antibodies that are important in determining the interactions between these molecules. Our meta-dynamic approach therefore should take these factors as descriptive properties of molecules and then use them in formulating the dynamical equations. Unfortunately, features such as the tertiary structure of molecules are known in only a few cases and tertiary structure is still very difficult to predict a priori. To mimic the required shape, charge and amino acid complementarity between molecules, we represent antibodies and antigens by binary strings. Complementarity can then be defined by any of a number of rules, and the degree of complementarity can be quantified and used as a measure of the affinity of the antigen-antibody interaction. Figure 6 illustrates the simplest rule: two molecules
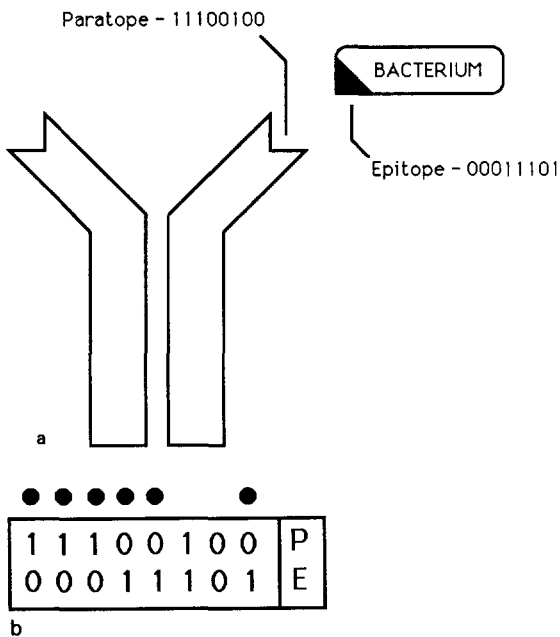
Fig. 6. (a) A schematic representation of the structure of an antibody and a typical antigen. The shape of the combining site of an antibody and the shape of an antigenic determinant (epitope) are represented by binary strings. (b) Complementarity can be assessed by aligning the antibody and antigen strings and then summing the number of positions at which a 1 is matched by 0.

are complementary if their binary string representations are complementary. In support of this rule there is evidence suggesting that if a peptide or protein is read from one strand of a double stranded DNA molecule and a "complementary peptide" is synthesized by reading from the complementary DNA strand, then the peptide and complementary peptide will bind specifically and with high affinity (Bost et al., 1985a,b; Smith et al., 1987; Shai et al., 1987). Further, in the case of the hormone ACTH, antibodies against ACTH and antibodies against the complementary peptide seem to be an idiotypic — antiidiotypic pair, leading to the speculation that idiotopes and antiidiotopes may represent complementary sequences in the hypervariable regions of such immunoglobulin pairs (Smith et al., 1987).

Other rules can also be used for determin-

ing complementarity. For example, since the strings represent molecules they need not be aligned when they interact. A match rule involving sequence shifting was discussed by Farmer et al. (1986b). Molecules generally do not interact over their entire length, but rather interactions tend to be localized. Stadynk (1987) uses a complementarity rule in which the number of adjacent complementary bits is important.

A graph can be used to represent the topology of an idiotypic network. As detailed in Section 1, a graph is generated by assigning each antibody and antigen type to a node and then representing complementarity between two molecules as a link joining their nodes (Fig. 7). By setting a threshold in the degree of complementarity required for the formation of a link different network graphs can be
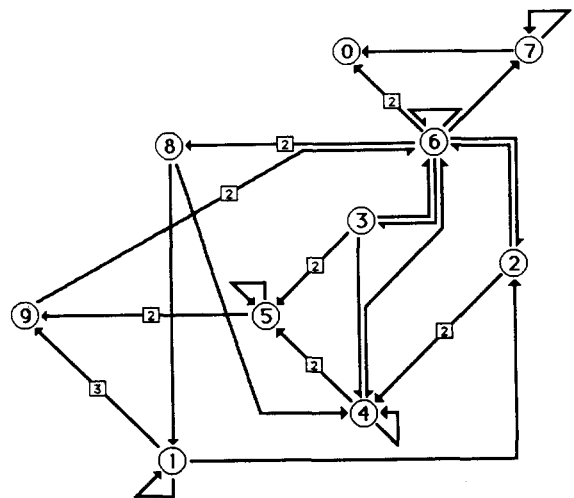

Fig. 7. A network graph generated usings strings of length 8. Here a link is drawn between two nodes if and only if at least 6 bits are complementary. Three levels of complementarity are noted. The highest level denoted by a 3 indicates that all 8 bits are complementary. There is only one such interaction. The next level denoted by a 2 indicates that 7 bits are complementary. Most of the interactions are level 1, (the label 1 is omitted for simplicity), corresponding to six bits being complementary. Links that are at complementary level 1 would be eliminated if the affinity criteria were changed to ≥ 7 bits complementary.

obtained, each graph representing interactions at a different minimal level of affinity.

We find there is some resemblance between our randomly generated networks and experimentally obtained networks (Fig. 8). However, since only a few networks have been mapped it may be premature to ascribe great significance to this resemblance.

In addition to a complementarity rule that generates the network topology one requires a set of meta-dynamical rules for generating the differential equations that determine the dynamics of the model. As in the case of the

complementarity rule many choices are possible. Unfortunately immunology has not progressed to the point where an agreed upon set of equations exist for describing the dynamics of the immune system. Various models have been developed by individual investigators, but none have been examined in sufficient detail by experimentalists to give one confidence in their predictions. In this domain of uncertainty our modeling has taken the form of a menu of different choices. A parameter in our simulation code allows one to choose the DeBoer model (DeBoer, 1988), the Hoffman
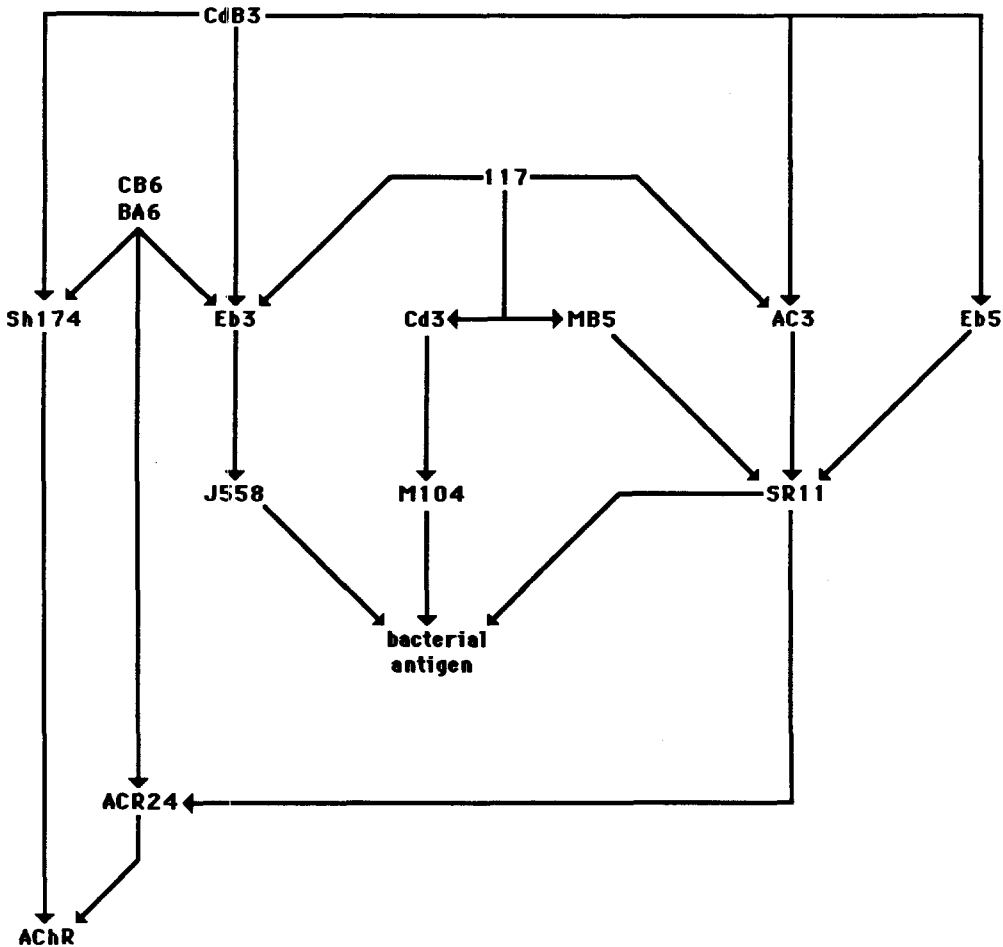

Fig. 8. Complementarities between antigens and antibodies mapped by Dwyer et al. (1986). This diagram shows a relationship between the immune response to the acetylcholine receptor (AChR) and dextran, a sugar found on the surface of some bacteria. The nodes are labeled with the names of monoclonal antibodies.

model (Hoffmann et al., 1988), the Segel—Perelson model (Segel and Perelson, 1988), the Farmer—Packard—Perelson model (Farmer et al., 1986b), or any of a number of immunologically detailed models that we have developed (Perelson et al., 1980; 1988). For purposes of illustration consider the following variation of the Farmer—Packard—Perelson model (Farmer et al., 1986b). Let x_i be the concentration of antibodies of type i and let y_k be the concentration of antigen of type k. Then

$$\frac{dx_i}{dt} = k_0 f$$

$$\left(\sum_j m_{ij} x_j x_i + \sum_k m_{ik} y_k x_i \right)$$

$$- k_1 g \left(\sum_j m_{ij} x_j x_i \right) - k_2 x_i$$

$$\frac{dy_k}{dt} = r_G(y_k) y_k - k_3 y_k - k_4 h \left(\sum_i m_{ik} x_i y_k \right)$$

The term proportional to k_0 represents the stimulation of production of antibodies of type i by their recognition of idiotopes on antibodies j or antigens of type k. The function $f(x)$ is a sigmoidal function, e.g.

$$f(x) = \frac{x^n}{(K + x^n)}$$

where n is an integer (typically 2 or 3). The saturating function is used here because the immune system can only produce antibodies at a limited rate, irrespective of the strength of stimulation. The match strengths m_{ij} and m_{ik} measure the degrees of complementarity between antibody and antibody and antibody and antigen, respectively. Molecules of type i may also be recognized as "antigens" by other antibodies in the system. Such recognition is assumed to lead to the elimination of antibodies of type at a rate proportional to k_1, but limited by a nonlinear saturating function $g(x)$. Antibodies have a limited lifetime in the body. The natural decay of antibody is modeled by the k_2 term.

Antigen is assumed to grow spontaneously with a density dependent growth rate $r_G(y_k)$. By choosing $r_G = r_0 (y_{max} - y_k)$ with r_0 and y_{max} constants, a logistic growth law is obtained. Antigen is assumed to die by natural causes at rate k_3 and to be eliminated by interaction with matching antibody at a rate proportional to k_4, but filtered through a saturating function $h(x)$.

In this model all interaction terms have been chosen to be bimolecular in accord with the law of mass-action. However, due to the limitation in rates of biological processes the quadratic terms are the arguments of other nonlinear functions that saturate. More complex interaction terms could be envisioned if one wished to include more of the biological and chemical details involved in the stimulation of antibody production by B cells. Thus, one might consider a more general form of this model

$$\frac{dx_i}{dt} = r_A(m_{ij} m_{ik} x_j x_j) x_i - k_2 x_i$$

$$\frac{dy_k}{dt} \quad r_G(y_k) y_k - k_3 y_k - k_4 \mu_{kill}(m_{ik} x_i y_k) y_k$$

where r_A, the antibody growth rate, is a function of all the other antibodies and antigens in the system and the strengths of their matches with antibody i, and μ_{kill} is the fraction of antigens of type k killed by interaction with matching antibodies. To be realistic both r_A and μ_{kill} need to be chosen as saturating functions. More explicit forms of these equations can be found in Perelson (1988) and Segel and Perelson (1988).

3.3. Phase-transitions in idiotypic networks

Because of the uncertainty in the exact form of the dynamical equations of an immune network some of the most interesting results from our simulation model of the immune system have to do with the topology of idiotypic networks. Besides being able to generate realistic looking idiotypic networks (Figs. 7 and 8), we have discovered what appears to

be a phase transition in the structure of the network. In order to explain this phase transition we need to introduce some terminology.

When an antigen is injected into an animal the set of antibodies that the animal raises against the antigen are denoted Ab_1. These first level antibodies are the usual antibodies of a clonal selection theory. If the animal generates an antiidiotypic response, then it will view Ab_1 antibodies as antigens and raise a set of antibodies, denoted Ab_2, that are to some degree complementary to members of the set Ab_1. Continuing in this way, Ab_i antibodies can lead to the generation of Ab_{i+1} antibodies. In experiments, Ab_1, Ab_2, Ab_3 and Ab_4 antibodies have been found (Bona and Pernis, 1984). Whether an animal makes antibodies, under natural circumstances, at level higher than Ab_4 is not known. In order to examine this question and the more general question of the structure of idiotypic networks, we have studied the topology of idiotypic networks constructed according to our complementarity match rule. We construct a system containing N antibodies each n bits long, and a single n bit antigen. We then determine according to our match rule all antibodies that match the antigen and label these Ab_1. For each Ab_1 antibody we then determine the matches with the remaining antibodies in the system. Any antibodies that match Ab_1 antibodies are placed in Ab_2. We continue in this manner, assigning an antibody to layer i if it matches an antibody in layer $i - 1$, and if it has not already been assigned to some previous layer. This latter condition ensures that each antibody is assigned to a unique layer and generates a network graph that is a rooted tree, the antigen being the root. Figure 9 is an example of a small network represented in this way.

There are various properties of networks represented as rooted trees that should be noted. First, all N antibodies in the system need not appear in the diagram. For example, if no antibody matches the antigen the diagram will contain only the antigen root. There is nothing special about the antigen, the
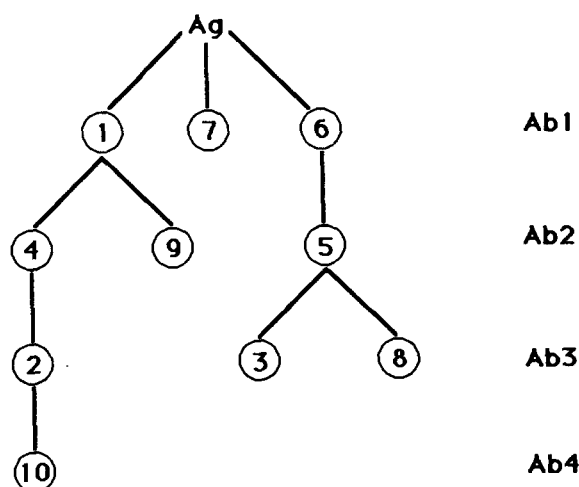

Fig. 9. A network of 10 antibodies drawn as a rooted tree.

diagram can "die out" at any layer because of a lack of matches. Thus, beginning with any root all of the antibodies in the system need not be assigned to an idiotypic level. Second, the number of layers in the tree is variable. Some trees will have many layers, others few.

Both the probability of the tree lacking some of the antibodies in the system and the number of layers in the tree depend on the probability that two strings match and hence on the match rule that we are using. For example, consider a system with 100 antibodies each represented by a 32 bit string (i.e. N = 100 and n = 32). Further, assume that our match rule is such that some threshold number of bits must be complementary for us to score two molecules as being complementary. Denote this threshold number n_{good}. If we choose n_{good} = 32, then we are essentially assured that the tree will be trivial, containing only the antigen. The probability of a match is $1/2^{32} \approx 2^{-10}$ and hence the system would need to contain 5×10^9 antibodies in order to expect even a single match.

Consider what happens as one varies n_{good} or equivalently the probability of a match. For small values of n_{good} all molecules in the system should match. In fact, using the simple

match rule shown in Fig. 6, we expect that on average half of the bits will match since 0 and 1 are chosen with equal probability. If all molecules match, then all of the antibodies in the system will be in layer 1. At values of n_{good} larger than $n/2$ there should be some matching, but not total matching. Thus one would expect trees with increasing numbers of layers. As n_{good} approached n, matches become rare and the number of layers in the tree and the number of antibodies in the tree will decrease. As reasoned above, when n_{good} reaches n we expect no antibodies in the tree when $N << 2^n$. Thus, we expect that a graph of the maximum layer reached in the tree versus n_{good} to approximate a smooth curve starting at 1, going through a maximum, and ending at 0. Surprisingly, we find that although the curve has these general characteristics it seems to approach a curve with a singularity as N is increased. Curves for $N = 100$ and 500 are given in Fig. 10. The behavior shown in the figure is typical for a system with a phase transition. There seems to be critical value of n_{good} at which the number of layers in an idiotypic network rises very sharply. For the system with 500 antibodies 29 layers are encountered at the critical value
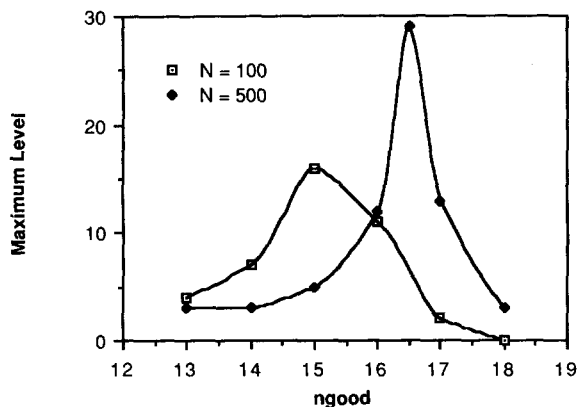

Fig. 10. Average of 100 Monte Carlo runs for 20 bit antibodies. In each run a network graph was generated, and the maximum level reached recorded. In the graph for N = 500 a set of runs was done with n = 32 to mimic a connectance that could be attained when n_{good} = 16.5.

of n_{good}. For systems with the diversity of the human immune system, i.e. containing 10^7 antibody types, we imagine that at the critical point hundreds or even thousands of layers may be present. We are currently developing a model, using percolation theory, from which we hope to predict the scaling behavior of this phase-transition (Perelson, 1989).

Why is the existence of this phase-transition interesting? Among immunologists the relevance of idiotypic networks to the functioning of the immune system is controversial. Some immunologists believe that idiotypic networks are an epiphenomenon and of no functional relevance, whereas other immunologists believe that idiotypic networks are the core of immune system, accounting for all of the normal activity of the immune system in times of health and controlling the system in times of disease. What we see in Fig. 10 is that outside of the phase transition region signals propagate only a few levels into the idiotypic network. When n_{good} larger than its critical value, this is due to the network being very sparse, with few antibodies at each level, and being composed of many unconnected components (Perelson, 1989). If the immune system normally operates in this part of the parameter regime, then some idiotypic antiidiotypic interactions will be present but the topology of the network will prevent a cascade of antibodies against antibodies from occurring. In the phase transition region or in the region where n_{good} is less than its critical value, the network is highly connected (Perelson, 1989) with many antibodies at each level, and the possibility exists for signals to propagate very deeply into the idiotypic network. In this regime network interactions have the ability to dominate any response. We have argued elsewhere (Perelson, 1989) that the immune system has parameters which place it in this part of the parameter regime. However, we must not forget that from topology alone one can not predict the dynamics of the immune response. It is the dynamical equations that will determine for any set of antibody and antigen con-

centrations and kinetic parameters whether or not deep penetration of the network will in fact occur.

In some autoimmune diseases, such as systemic lupus erythematosus (Shoenfeld, 1988), one observes many different types of anti-self antibodies. One might speculate that normally the dynamical interactions in the immune system prevent deep penetration and that network interactions are not as pronounced as topological considerations alone would lead us to believe. However, in autoimmune disease the kinetic parameters of the immune system may change in such a way as to allow all of the topologically possible interactions. If this occurs, many antibodies against antibodies will be raised, and some of these antibodies may cross-react with self components and contribute to the symptoms of autoimmune disease. Network interactions already have been implicated in causing myasthenia gravis (Dwyer et al., 1986).

4. Discussion

The autocatalytic reaction network is a useful abstraction to study for several important problems. First, we hope to understand the relevance of the connectivity and the hierarchical organization of the network to collective and specifically adaptive behavior. Second, though the networks we are studying are theoretical constructs, the improvement of experimental techniques may provide the opportunity to produce novel reaction networks in vitro. For such studies, the exact yields of each component will be of lesser importance than the behavior of the network. Our models are precursors to a potentially useful tool for experimentalists who would like to test a network design before implementation. Finally, we are trying to shed light on the origin of life. The polymers of the autocatalytic reaction network collectively demonstrate some of the properties which an organism must have. Autocatalytic reaction networks may have filled an evolutionary niche between the relative disorder of the prebiotic soup and the appearance of an organism. In our study of proto-organisms, we want to understand the means as well as the order of appearance of various submechanisms within the proto-organism.

Much remains to be done in order to understand the operation of the immune system. We believe that simulation models, such as the one we have developed using meta-dynamics, will be valuable tools. Although we have emphasized idiotypic networks, meta-dynamical models apply equally well to the study of non-network models, i.e. models that deal with interactions among cell subpopulations such as B cells, T helper and T suppressor cells. Meta-dynamical models are particularly useful in cases in which one wants to deal with mutation or other sources of variation among individual clones, with clones being created or destroyed.

An analogy with autocatalytic networks may be helpful in our pursuit of models of the immune system that incorporate both T cells and B cells. Helper T cells secrete factors that facilitate the growth of B cells that have interacted with antigen or antiidiotypic antibody. Thus, one can view a helper T cell as a catalyst, enhancing the reaction

$$\text{antigen} + \text{B cell antibody} \qquad (3)$$

Here, like the model for autocatalytic networks with match rules, the T cell is assigned to a specific reaction as a catalyst depending on the sequence of its string representation. T cells help B cells in an antigen specific manner. Thus a T cell must have a receptor that matches the antigen in order to act as a catalyst for interactions involving that antigen. Viewing T cells as catalysts is an exciting modeling possibility that we have not yet pursued. Our experiences with autocatalytic network models suggests a rich set of dynamic phenomena may be revealed in such models.

In light of this analogy with the immune system, we see that it is useful to think of catalysis as a type of nonlinear control over a

process. In this broader sense, autocatalytic networks may share general properties with other kinds of systems. The types of "catalysis" which we use to distinguish one system from another will be dependent upon the rules of interaction for each system.

Models with meta-dynamics provide new opportunities for study. In an abstract sense, meta-dynamical systems provide the ability to study potentially infinite systems with finite, tractable models. More concretely, meta-dynamical systems are a digital laboratory for the study of evolution and adaptation. An important feature of these systems is that predefined fitness functions are not needed to evaluate the value of a response to environmental variation. Like real systems, the result of the interaction between system and environment in meta-dynamical systems is evaluated naturally. Instead of fitness functions, these models include low-level rules of interaction. Higher level behavior emerges from these low-level rules.

Both the autocatalytic reaction network and the immune system are examples of adaptive systems. In the autocatalytic reaction network, better catalysts may supercede existing catalysts, and the properties of a new polymer may enhance the robustness of the metabolism, or its efficiency at consuming and incorporating elements of the food set. Similarly, the immune system maintains within the body lymphocytes that detect foreign antigens while allowing lymphocytes that detect nothing to die and exit from the system. Both systems have the fundamental ability to change in directions which respond to the environment.

The meta-dynamical approach to these systems provides a viable means of simulating them, in a manner that naturally mimics their adaptive properties. A great deal remains to be done with this approach, both for the problems we have discussed, and in other areas, such as adaptive games, ecological models, and neural networks with plasticity in their synapses. We hope to stimulate further work in these directions.

References

Atkins, P.W., 1978, Physical Chemistry (W.H. Freeman, San Francisco, CA).

Buzayan, J.M., Gerlach, W.L. and Bruening, G., 1986, Nonenzymatic cleavage and ligation of RNAs complementary to a plant virus satellite RNA. Nature 323, 349–353.

Bollobas, B., 1979, Graph Theory: An Introductory Course (Springer-Verlag, New York).

Bollobas, B., 1985, Random Graphs (Springer-Verlag, New York).

Bona, C.A. and Pernis, B., 1984, Idiotypic Networks (Raven Press, New York).

Bost, K.L., Smith, E.M. and Blalock, J.E., 1985a, Regions of complementarity between the messenger RNAs for epidermal growth factor, transferrin, interleukin-2 and their respective receptors. Biochem. Biophys. Res. Commun. 128, 1373–1380.

Bost, K.L., Smith, E.M. and Blalock, J.E., 1985b, Similarity between the corticotropin (ACTH) receptor and a peptide encoded by an RNA that is complementary to ACTH mRNA. Proc. Natl. Acad. Sci. USA 82, 1372–1375.

Cairns-Smith, A.G., 1982, Genetic Takeover and the Mineral Origins of Life (Cambridge University Press, New York).

Calvin, M., 1969, Chemical Evolution (Oxford University Press, Oxford).

Cech, T.R., 1986, RNA as an enzyme. Sci. Am. 255, 64–75.

Cech, T.R. and Bass, B.L., 1986, Biological catalysis by RNA. Annu. Rev. Biochem. 55, 599–629.

Cox, B., 1986, Object Oriented Programming (Addison-Wesley, Reading, MA).

DeBoer, R.J., 1988, Symmetric idiotypic networks: connectance and switching, stability and supression, in: Theoretical Immunology, Part II, A.S. Perelson (ed.) (Addison-Wesley, Reading, MA) pp. 265–289.

Dwyer, D.S., Vakil, M. and Kearney, J.J., 1986, Idiotypic network connectivity and a possible cause of myasthenia gravis. J. Exp. Med. 164, 1310–1318.

Dyson, F., 1985, Origins of Life (Cambridge University Press, New York).

Farmer, J.D., Kauffman, S.A. and Packard, N.H., 1986a, Autocatalytic replication of polymers. Physica D 22, 50–67.

Farmer, J.D., Packard, N.H. and Perelson, A.S., 1986b, The immune system, adaptation, and machine learning. Physica D 22, 187–204.

Flory, P.J., 1953, Principles of Polymer Chemistry (Cornell University Press, Ithaca, NY).

Fox, R.F., 1982, Biological Energy Transduction: The Uroboros (John Wiley, New York).

Fox, R.F., 1988, Energy and the Evolution of Life (W.H. Freeman, New York).

Fox, S.W. and Wang, C.T., 1968, Science 160, 547.

Fox, S.W., Harada, K. and Kendrick, J., 1959, Production of sperules from proteinoids and hot water. Science 129, 1221−1223.

Harary, F., 1969, Graph Theory (Addison-Wesley, Reading, MA).

Hoffmann, G.W., Kion, T.A., Forsyth, R.B., Saga, K.G. and Cooper-Willis, A., 1988, The n-dimensional network, in: Theoretical Immunology, Part II, A.S. Perelson (ed.) (Addison-Wesley, Reading, MA) pp. 219−319.

Hutchins, C.J., Rathjen, P.D., Forster, A.C. and Symons, R.H., 1986, Self-cleavage of plus and minus RNA transcripts of avocado sunblotch viroid. Nucleic Acid Res. 14, 3627−3640.

Jerne, N.K., 1973, The immune system. Sci. Am. 229(1), 52−60.

Jerne, N.K., 1974, Towards a network theory of the immune system. Ann. Immunol. Inst. Pasteur 125 C, 373−389.

Joyce, G., 1987, Nonenzymatic template-directed synthesis of informational macromolecules. Cold Spring Harbor Symp. Quant. Biol. LXII, 41−51.

Kauffman, S.A., 1986, Autocatalytic sets of proteins. J. Theor. Biol. 119, 1−24.

Keese, P. and Symons, R.H., 1985, Domains in viroids: evidence of intermolecular RNA rearrangements and their contribution to viroid evolution. Proc. Natl. Acad. Sci. 82, 4582−4586.

Kernighan, B.W. and Ritchie, D.M., 1988, The C Programming Language, 2nd edn. (Prentice Hall, Englewood Cliffs, NJ).

Koshland, Jr., D.E., 1987, Evolution of catalytic function. Cold Spring Harbor Symp. Quant. Biol. LXII, 1−7.

Loomis, W.F., 1988, Four Billion Years: An Essay on the Evolution of Genes and Organisms (Sinauer Associates, Inc., Sunderland, MA).

Miller, S.L. and Orgel, L.E., 1974, The Origins of Life on Earth. Concepts of Modern Biology (Prentice Hall, Englewood Cliffs, NJ).

Miller, S.L. and Urey, H.C., 1959, Science 130, 245.

Orgel, L.E., 1963, Proc. Natl. Acad. Sci. 49, 517.

Palmer, E.M., 1985, Graphical Evolution: An Introduction to the Theory of Random Graphs (John Wiley, New York).

Paecht-Horowitz, M. and Katchalsky, A., 1973, Synthesis of amino acyl-adenylates under prebiotic conditions. J. Molec. Evolution 2, 91−98.

Paecht-Horowitz, M., Berger, J. and Katchalsky, A., 1970, Prebiotic synthesis of polypeptides by heterogeneous polycondensation of amino-acid adenylates. Nature 228(5272), 636−639.

Perelson, A.S. (ed.), 1988, Toward a realistic model of the immune system, in: Theoretical Immunology, Part II (Addison-Wesley, Reading, MA) pp. 377−401.

Perelson, A.S., 1989, Immune networks: a topological view, in: Theoretical Models for Cell to Cell Signalling, A. Goldbeter (ed.) (Academic Press, New York).

Perelson, A.S. and Oster, G.F., 1979, Theoretical studies of clonal selection: minimal antibody repertoire size and reliability of self-nonself discrimination. J. Theoret. Biol. 81, 645−670.

Perelson A.S., Goldstein, B. and Rocklin, S., 1980, Optimal strategies in immunology. III. The IgM-IgG switch. J. Math. Biol. 209−256.

Radman, M.R. and Wagner, R., 1988, The high fidelity of DNA duplication. Sci. Am. 259(2), 40−46.

Rouvray, D.H., 1986, Predicting chemistry from topology. Sci. Am. 255, 40−47.

Segel, L.A. and Perelson, A.S., 1988, Computations in shape space: a new approach to immune network theory, in: Theoretical Immunology, Part II, A.S. Perelson (ed.) (Addison-Wesley, Reading, MA) pp. 321−343.

Shai, Y., Flashner, M. and Chaiken, I.M., 1987, Anti-sense peptide recognition of sense peptides: direct quantitative characterization with the ribonuclease s-peptide system using analytical high-performance affinity chromatography. Biochemistry 26,669−675.

Shoenfeld, Y., 1988, Analyses of the idiotypes and ligand binding characteristics of human monoclonal autoantibodies to DNA: do we better understand systemic lupus erythematosus? in: Perspectives on Autoimmunity, I.R. Cohen (ed.) (CRC Press, Boca Raton, FL) pp. 135−141.

Smith, L.R., Bost, K.L. and Blalock, J.E., 1987, Generation of idiotypic and anti-idiotypic antibodies by immunization with peptides encoded by complementary RNA: a possible molecular basis for the network theory. J. Immunol. 138, 7−9.

Stadnyk, I., 1987, Schema recombination in pattern recognition problems, in: Genetic Algorithms and their Applications, J.J. Grefenstette (ed.) (Lawrence Erlbaum, Hillsdale, NJ).

Stroustrup, B., 1986, The C++ Programming Language (Addison-Wesley, Reading, MA).

Trinajstic, N., 1983, Chemical Graph Theory (CRC Press, Boca Raton, FL).

Uhlenbeck, O.C., 1987, Nature 328, 598−600.

Watson, J.D., Hopkins, N.D., Roberts, J.W., Steitz, J.A. and Weiner, A.M., 1988, Molecular Biology of the Gene, Vol.2, 4th edn. (Benjamin/Cummings, Reading, MA)