Abstract

We study networks that are a generalization of replicator (or Lotka-Volterra) equations. They model the dynamics of a population of object types whose binary interactions determine the specific type of interaction product. We show that the system always reduces its dimension to a subset that contains production pathways for all of its members. The network equation can be rewritten at a level of collectives in terms of two basic interaction patterns: replicator sets and cyclic transformation pathways among sets. Although the system contains well-known cases that exhibit very complicated dynamics, the generic behavior of randomly generated systems is found (numerically) to be extremely robust: convergence to a globally stable rest point. It is easy to tailor networks that display replicator interactions where the replicators are entire self-sustaining subsystems, rather than structureless units. A numerical scan of random systems highlights the special properties of elementary replicators: they reduce the effective interconnectedness of the system, resulting in enhanced competition, and strong correlations between the concentrations.

1 Introduction

Replicators are entities that are copied during interactions with other entities. Replicator equations, or Lotka-Volterra equations, are commonly used to describe the population dynamics of replicators [24, 20]. In these models, replicators are usually assumed to be objects without internal structure, and the copying process is subsumed into a single reaction event.

A variety of systems from computer science, chemistry, biology, ecology, and economics, deal with objects that are not replicators themselves, but rather interact with other objects to produce further objects not necessarily of the same type. Here, we focus on such reaction networks in which the products can be viewed as being a function of both interacting objects.

An example of the types of reaction networks considered here comes from chemistry, where two molecules react to produce new molecules whose nature is typically not determined by chance, but by the reactive properties of the educts. As a corollary, a system could be a replicator as a whole without any individual object of the system being a replicator. This requires that every object has some production pathway involving only objects of the same system. Such a scenario may have occurred in an early phase of prebiotic molecular evolution, preceding the emergence of individual replicators [12]. Models of such an autocatalysis at the level of collectives of biopolymers have been proposed [5, 2].

The production of specific objects involves the specification of a mapping that allows one to determine the product given the interaction partners. This mapping can be encoded in the interaction partners themselves, as in chemistry, or it can be specified by a look-up table. The former case allows the study of systems with a potential infinity of types, at the expense of a stochastic dynamics [7] or a so-called meta-dynamics [5]. The latter case, in contrast, can be cast in terms of ordinary differential equations, with the interaction matrix serving as the look-up table. Such matrices may be either structured or random.

In this paper we analyze the system using the ODE approach. This approach allows us to form a simple specification of the system. Using this specification, we are able to address the issue of the generic behavior present in an ensemble of randomly generated systems. Moreover, by biasing the generation of interaction matrices, we can investigate the transition of the system to special cases like replicator dynamics, as well as the organizational stability of the system with respect to interaction modifications. Our approach is related to studies of replicator systems with mutation. In such systems, a replicator is not copied correctly, but gives rise to a variant type [25]. Such work uses a perturbation approach in which the reference state is given by the replicator equation and the mutation field is treated as a perturbation.

In section 2 we define our reaction network equation—a simple generalization of replicator equations. Section 3 surveys some of the special cases contained in the equation. Though it has a simple form, the network equation is analytically difficult to work with; This is in part due to the equation's generality. In Section 4 we present a few rigorous results. Section 5 addresses the problem of finding structure in network systems by grouping together individual objects in to sets, and rephrasing the equations at the level of these collectives. Section 6 summarizes our experience with numerical integration of a large number of both random and non-random networks.

2 The Reaction Network

We consider a system S of n types of objects, where two objects i and j interact to produce one or more types of objects $k_1, \ldots, k_l$. The reaction products are assumed to be again in S, and the interaction partners i and j are retained. Hence i and j play the role of catalysts. Any material needed to build the products is buffered, and therefore does not explicitly enter the kinetic equations. Furthermore, the system is placed into a continuously stirred tank reactor with an unspecific dilution flux, $\Phi(t)$, that keeps the total number of particles constant. In the following we switch to relative concentrations, $0 \leq x_i(t) \leq 1, i = 1, \ldots, n$. The state space of our model is the concentration n-simplex. The rate equations arise through mass action kinetics, and in the deterministic setting considered here they become:

$$\dot{x}_k = \sum_{i=1}^{n} \sum_{j=1}^{n} \alpha_{ij}^{k} x_i x_j - x_k \Phi(t), \quad k = 1, \ldots, n, \tag{1}$$

with second order rate constants α_{ij}^k for the reaction $i + j \longrightarrow i + j + k$ and with $\Phi(t)$ such that $\sum_k x_k(t) = 1$. Throughout the paper we will refer to equation (1) as the "catalytic network equation." In the next section we will show that this equation contains many of the well known models of prebiotic evolution and population genetics as special cases.

There is a natural way of splitting the n^3 coefficients, α_{ij}^k, into rate constants, a_{ij}, and transmission coefficients, t_{ij}^k, by defining

$$a_{ij} = \sum_{k=1}^{n} \alpha_{ij}^{k} \tag{2}$$

$$\alpha_{ij}^{k} = a_{ij} \cdot t_{ij}^{k}. \tag{3}$$

The a_{ij} gives the fraction of reactive collisions between i and j, and t_{ij}^k denotes the relative frequency of the reaction product k. The $\sum_{k=1}^{n} t_{ij}^k$ is equal to 1 for any reactive pair (i, j), and is equal to 0 for those pairs that do not interact at all. To keep the sum of all concentrations at unity, the dilution flux $\Phi(t)$ equals the total production rate of the system:

$$\Phi(t) = \sum_{r,s,t=1}^{n} \alpha_{rs}^{t} x_r x_s = \sum_{r,s=1}^{n} a_{rs} x_r x_s. \tag{4}$$

In population genetics, the use of difference equations is more common then the use of differential equations. The most general second order model is

$$x_k' = \frac{1}{u} \cdot \sum_{i,j=1}^{n} t_{ij}^{k} w_{ij} x_i x_j, \tag{5}$$

where x_k denotes the relative abundance of a haplotype k in the population, w_{ij} is the fitness of a diploid genotype composed of the haplotypes i and j, t_{ij}^k is the probability that a gamete produced by a parent with genotype (i, j) has haplotype k, and u is a normalization factor that keeps the variables x_k normalized [1]. Equation (5) is precisely the discrete time analogue of our catalytic network equation. Population geneticists restrict the transmission coefficients, t_{ij}^k, to reflect the genetic setting. Here, we do not consider such restrictions.

3 Special Cases

The catalytic network equation has a variety of important special cases that arise by imposing constraints on the transmission coefficients. Such cases have been widely analyzed. In this section we briefly make these connections.

3.1 Replicator equation

When the interaction between i and j yields a copy of i or j, then we obtain the second order replicator equation in n species on the simplex [11], which has been shown to be topologically equivalent to the $(n-1)$-dimensional Lotka-Volterra equation on $\mathbb{R}_{+}^{n-1}$ [10].

The special choice of transmission coefficients, $t_{ij}^k = 0$ iff $i \neq k \vee j \neq k$, and the abbreviation $b_{kj} = t_{kj}^k a_{kj} + t_{jk}^k a_{jk}$ yields the replicator equation

$$\dot{x}_k = x_k \left[\sum_{j=1}^{n} b_{kj} x_j - x_k \sum_{i,j=1}^{n} b_{ij} x_i x_j \right]. \tag{6}$$

The replicator equation (6) has further interesting special cases, some of which are: (1) Fisher's selection equation [6] with $b_{ij} = b_{ji}$, (2) Schloegl's model of independent replication with $b_{ij} = c_i \cdot \delta_{ij}$, and (3) the hypercycle [4], where $b_{ij} = c_i \cdot \delta_{i,j-1}$ (indices modulo n).

3.2 Complementary replication

In complementary replication the interaction between i and j is still a copying process, but the product is the complementary template of i or j, as, for example, occurs in RNA replication. Let k^* denote the complement of k, with b_{kj} as before, the network equation becomes:

$$\dot{x}_{k^*} = x_k \sum_{j=1}^{n} b_{kj} x_j - x_{k^*} \sum_{ij} b_{ij} x_i x_j. \tag{7}$$

This system of equations has recently been studied by Stadler [26].

3.3 Mutation replication equation

The template directed replication of i, catalyzed by j, may have a limited accuracy. In this case a particular product mutant k will be produced with a probability expressed by the transmission coefficients t_{ij}^{k}. In the usual template directed DNA or RNA replication this probability only depends on the template to be copied, but not on the catalyst. The rate constants α_{ij}^{k} factorize accordingly into the replication rate, a_{ij}, and the mutant frequencies, $Q_{ki} = t_{ij}^{k}$ for all j.

$$\dot{x}_k = \sum_{i=1}^{n} \sum_{j=1}^{n} Q_{ki} a_{ij} x_i x_j - x_k \Phi(t), \quad k = 1, \ldots, n. \tag{8}$$

This equation has the same form as our network equation. The product k, however, is generated by noise, rather than by a specific action involving both i and j. This is expressed by the fact that the transmission coefficient Q_{ki} does not depend on j. Equation (8) has been recently investigated [28].

Fertility selection and selection-recombination equations [11] are further special cases of the network equation (1).

3.4 Turing Gas

A recent model [7, 8] considers a system in which objects encode symbolic transformations that operate on other such objects yielding new objects representing further transformations. The universe of objects and their transformations is given by the axioms of a calculus known as the λ-calculus (e.g. see [21]), in which there is no syntactical distinction between a function and its arguments.

The dynamical system results when such objects "react" according to mass action kinetics. Upon collision one object acts as a function that is applied to the other object acting as argument, resulting in a new object (the value). In the Turing gas model the overall dynamics follows the network equation (1), but is implemented as a stochastic process, because the system need not be closed with respect to the applicative behavior of its component objects: new types of objects may appear as a result of interactions occurring within the system. To model instances of the Turing gas with a deterministic set of equations we require that the system be closed with respect to the interactions occurring in it.

Because the application, $\circ$, of a map to an argument yields a unique value, the transmission coefficients are now given by

$$t^h_{gf} = \begin{cases} 1 & \text{if } h = f \circ g \\ 0 & \text{otherwise,} \end{cases} \tag{9}$$

with, for example, the $a_{ij} = 1, \forall\ i, j$. Additional constraints on the interactions may be introduced by setting particular a_{ij} to zero.

Section 5 and some numerical experiments reported in section 6 refer to the Turing gas case.

4 Properties

First we observe that the network equation is invariant to a symmetric or asymmetric choice of the rate coefficients α^k_{ij}. For any given k consider the $n \times n$ matrices $A^{(k)} = (\alpha^k_{ij})$, whose nonzero entries are those interactions that result in k. Let $\tilde{A}^{(k)} = (A^{(k)} + A^{(k)^T})/2$ be the symmetrized version of $A^{(k)}$. Rewriting the catalytic network equation in terms of quadratic forms in $A^{(k)}$ gives

$$\dot{x}_k = (\mathbf{x}A^{(k)}\mathbf{x}) - x_k \sum_t (\mathbf{x}A^{(t)}\mathbf{x}) = (\mathbf{x}\tilde{A}^{(k)}\mathbf{x}) - x_k \sum_t (\mathbf{x}\tilde{A}^{(t)}\mathbf{x}). \tag{10}$$

Next we note that the $x_k(t)$ remain nonnegative for all $t > 0$, provided all symmetrized coefficients $\tilde{\alpha}^k_{ij}$ are nonnegative for $k \neq i$ and $k \neq j$. This means that if degradation reactions for k are to be considered, they have to involve k as a reactant. In the following we will not consider degradation reactions at all. The total concentration is only constrained by the unspecific dilution flux $\Phi(t)$.

Nothing is changed by adding constant and linear terms to the network equation. They can always be absorbed into the second order terms.

By a straightforward application of Hofbauer's transformation [10], equation (1) can be shown to be topologically equivalent on the simplex (except for the face $\{x_n = 0\}$) with the following $n-1$ dimensional system of equations on $\mathbb{R}^{n-1}_+$:

$$\dot{y}_k = \omega_k + \sum_{j=1}^{n-1} \beta_{kj} y_j + \sum_{i,j=1}^{n-1} \gamma^k_{ij} y_i y_j - y_k \cdot \Theta(y), \tag{11}$$

with coefficients ($1 \leq i,j,k \leq n-1$):

$$\omega_k = \alpha^k_{nn}, \quad \beta_{kj} = \alpha^k_{nj} + \alpha^k_{jn}, \quad \gamma^k_{ij} = \alpha^k_{ij}$$
$$\Theta(y) = \alpha^n_{nn} + \sum_{j=1}^{n-1} (\alpha^n_{nj} + \alpha^n_{jn}) y_j + \sum_{j=1}^{n-1} \alpha^n_{ij} y_i y_j \tag{12}$$

A suitable choice of the above coefficients yields the direct Lotka-Volterra analogue of our catalytic network equation:

$$\dot{y}_k = \sum_{i,j}^{n} \gamma^k_{ij} y_i y_j - d_k \cdot y_k, \tag{13}$$

where the d_k are species specific degradation rates.

Let $S = \{1, 2, \ldots, n\}$ denote the set of object types in the system. A closed set A is defined to be a set of types such that no interaction of types in A produces types in $S \setminus A$, that is $\forall\ i,j \in A,\ k \in S \setminus A,\ \alpha^k_{ij} = 0$. Clearly, the subsimplex spanned by A is invariant. If S contains no closed subsets, then the system is permanent [11] – no type disappears (by definition) – and hence there is at least one interior rest point.

We refer to a system as active if $\liminf_{t\to\infty} \Phi(t) > 0$ for all initial conditions in the interior of the simplex. An active system always reduces its dimensions until it reaches a support in which all types have some production pathway. This is expressed in the following proposition, whose (trivial) proof we omit.

Proposition. (a) Any type m that is not produced disappears as $t \to \infty$.
(b) In an active system, k disappears if all types that produce it disappear.

5 Decomposition

5.1 A dynamical system on an algebraic structure

In this section we consider the case where the action of a particle of type j on a particle of type i results in a unique product k. This occurs in models like the Turing gas (section 3.4). More generally, we consider an algebraic structure, given by a finite set S of discrete objects together with a law of composition, $\circ$: $S \times S \to S$. Finiteness is required for an ODE treatment to be sensible.

Consider a population P on S, and denote by x_k the relative abundance of k in the population. The dynamics on P shall be governed by the network equation (1), with

$$\alpha^k_{ij} > 0 \Rightarrow k = j \circ i \tag{14}$$

The α_{ij}^k express dynamical constraints in the form of rate constants. For each $j \in S$ there are two sets $I(j)$ and $K(j)$ such that $\alpha_{ij}^k > 0$ iff $i \in I(j)$ and $k \in K(j)$. The object j then corresponds to a map from its domain $I(j) \subseteq S$ into its range $K(j) \subseteq S$. In particular, when some α_{ij}^k are zero, then the domain of j is a proper subset of S. We also say that this map "is the action of j on S."

Clearly, the case of a stochastic dynamics on an infinite set S allows the study of the interplay between organization emerging on the "semantic level," that is, the action of the maps on the current population, and organization at the syntactic level, that is, the regularities emerging simultaneously in the syntax of the maps [8]. To keep things simple we do not use any particular algebraic structure, but rather generate the α_{ij}^k at random. In some cases, we transcribe organizations obtained by runs with the Turing gas model into a set of α_{ij}^k and obtain systems with interesting population dynamics that would otherwise have an extremely low probability of being generated at random. This is the case in section 6.2.2 where we consider a "life-cycle" type of organization.

A vast number of numerical experiments (section 6.1) shows that the generic behavior of the network equation for randomly generated systems, is extremely robust and consists of only fixed points. In most cases these fixed points seem to be unique and globally stable. We will, however, explicitly engineer systems with more complicated dynamics (section 6.2).

In the following we show that the objects can always be grouped into sets such that at the set level the network equation for each set has either the form of a replicator equation or of a cyclic transformation among the sets. This result is more conceptually rather than analytically useful. In fact, in most cases a random network is a highly interconnected single replicator, albeit with a complicated internal structure. The decomposition does not tell us anything about the internal dynamics of the replicators. Nevertheless, it does suggest the possibility of coupling several random networks in a hypercyclic fashion (section 6.2.1).

5.2 Decomposition of the network equation

Definition. Let $\mathcal{P}$ be a partition of S such that each set $K \in \mathcal{P}$ is invariant under the action of each element $z \in S$, namely

$$\forall\, z \in S: \quad z \circ K = \{z \circ x | x \in K \cap I(z)\} \subseteq K. \tag{15}$$

Lemma. There is a unique partition $\mathcal{P}$ such that no $K \in \mathcal{P}$ can be further decomposed into action invariant subsets.
Proof. Let $\mathcal{P}_1 \cap \mathcal{P}_2$ be the collection of all pairwise intersections $K_i \cap K_j$, $K_i \in \mathcal{P}_1$, $K_j \in \mathcal{P}_2$. Clearly, $z \circ (K_i \cap K_j) \subseteq (K_i \cap K_j)$. The partitions fulfilling (15) form a partially ordered set with respect to the relation $\leq$, where $\mathcal{P}_1 \leq \mathcal{P}_2$ iff every set

$K \in \mathcal{P}_1$ is contained in some set $L \in \mathcal{P}_2$. Since S is finite this poset is also finite. The partition $\mathcal{P} = \bigcap \mathcal{P}_\ell$ is the unique least element of the poset.

In the following we will always refer to this finest partition $\mathcal{P}$. The above states that every set $K \in \mathcal{P}$ can be considered as behaving like a replicator, and that there is a unique decomposition into "smallest" replicator units. If $\mathcal{P}$ is nontrivial, that is $\mathcal{P} \neq \{S\}$, then the network equation is termed replicator decomposable. In fact, on the level of sets it can be rewritten in the form of a replicator equation.

Let $x_{k\in K}$ denote the concentration of the species k in set K. The network equation gives

$$\dot{x}_{k\in K} = \sum_{i\in K} \sum_{J\in\mathcal{P}} \sum_{j\in J} \alpha^{k\in K}_{i\in K, j\in J} x_{i\in K} x_{j\in J} - x_{k\in K}\Phi(t). \tag{16}$$

Let $x_K = \sum_{k\in K} x_{k\in K}$ and $y_{k\in K} = x_{k\in K}/x_K$. At the set level our model becomes

$$\dot{x}_K = x_K \left[\sum_J \mathcal{A}_{KJ}(t) x_J - \Phi(t) \right] \tag{17}$$

with the coupling coefficients between the replicator sets given by

$$\mathcal{A}_{KJ}(t) = \sum_{i\in K} \sum_{j\in J} \sum_{k\in K} \alpha^{k\in K}_{i\in K, j\in J} y_{i\in K} y_{j\in J}. \tag{18}$$

The coupling coefficients depend on time as well as on the initial conditions, and are bound by

$$\min_{j\in J, i\in K} \sum_{k\in K} \alpha^{k\in K}_{i\in K, j\in J} \leq \mathcal{A}_{KJ}(t) \leq \max_{j\in J, i\in K} \sum_{k\in K} \alpha^{k\in K}_{i\in K, j\in J}. \tag{19}$$

Now consider a set $K \in \mathcal{P}$. K may be further partitioned into a collection of subsets $\mathcal{Q}(K) = \{L_1, \ldots, L_r\}$ such that

$$\forall\ z \in S: \quad z \circ L_r \subseteq L_s \tag{20}$$

With the same argument as above there is a unique minimal partition with respect to the order relation $\leq$ on the poset of partitions fulfilling (20).

It is convenient to depict the relations between the sets of $\mathcal{Q}$ as a directed graph $\vec{\Gamma}$. Since $K \in \mathcal{P}$ the graph is connected, and each vertex has at most one outgoing edge. It follows that $\vec{\Gamma}$ contains either one circuit $C = L_1 \to L_2 \to \ldots \to L_r \to L_1$ or one endpoint $L_r \in \mathcal{Q}$.

At the level of sets, the equations for the members of $\mathcal{Q}(K)$ become

$$\dot{x}_{L_i} = x_{(\to L_i)} \left[\sum_{L_j} \mathcal{A}_{(\to L_i)L_j}(t) x_{L_j} - \Phi(t) \right] \tag{21}$$

where $(\rightarrow L_i)$ denotes the (possibly empty) parent set of L_i according to (20). L_r or C may have incident trees of sets, but the theorem of the previous section implies that only the types in C or L_r survive for $t \rightarrow \infty$.

One can thus uniquely decompose any S into a collection of replicator sets which in turn may be decomposed into a collection of sets reproducing themselves in a transitive cyclic fashion.

5.3 A cluster problem

Rarely does a randomly generated system decompose into distinct interacting replicator sets or into several sets that are transformed cyclically into each other. Nevertheless, the previous section suggests that grouping together several object types may result in a higher level description in terms of (possibly) fewer clusters at the expense of time dependent coefficients. Such a clustering could reveal a "natural" order not visible at the level of individuals. The problem is to find, according to some predefined criterion, a partition $\mathcal{P} = \{K_1, \ldots, K_n\}$ of the set of object types $\mathcal{S}$ such that

$$\forall\, k, i \neq j \quad K_i \cap K_j = \emptyset \quad \wedge \quad K_i \circ K_j \subseteq K_k \tag{22}$$

holds, where the action of a set A on a set B is defined as

$$A \circ B = \{x \,|\, x = i \circ j,\ (i,j) \in A \times B\}. \tag{23}$$

Both trivial partitions, $K_i = \{i\}$, $i = 1, \ldots, n$ as well as $K_1 = S$, comply with (22), but there is in general no unique partition in between. A sensible clustering criterion would, for instance, try to minimize connectivity of the interaction graph between clusters, and maximize it within clusters. A self maintaining set, that is, a replicator, by definition does not entertain transformation pathways to other sets. The decomposition into replicators is therefore optimal. In most cases, however, one is further interested in the internal structure of a replicator (which might be fairly complicated).

6 Numerical experiments

In this section we investigate numerically the effects of interconnectedness on the behavior of the network equation (1). Recall that this equation is not a Lotka-Volterra (or replicator) equation, but rather contains the latter as a special case. The study of random systems of the Lotka-Volterra type has received considerable attention in the past twenty years [9, 14, 15, 22, 18, 13, 19, 27].

6.1 Random Networks

At the level of individual object types, we can separate out the replicator part of the network equation:

$$\begin{aligned}
\dot{x}_k &= \mathcal{R}_k(x) + \mathcal{J}_k(x) \\
\mathcal{R}_k(x) &= x_k \left[\sum_{j=1}^{n} (\alpha_{kj}^k + \alpha_{jk}^k) x_j - \sum_{i,j=1}^{n} (\alpha_{ij}^i + \alpha_{ji}^i) x_i x_j \right] \\
\mathcal{J}_k(x) &= \sum_{i,j \neq k} \alpha_{ij}^k x_i x_j - x_k \sum_t \sum_{r,s \neq t} \alpha_{rs}^t x_r x_s
\end{aligned} \tag{24}$$

In the following experiments we begin by generating matrices at random that do not contain copying reactions thus $\mathcal{R}_k(x) \equiv 0$ for all k. For the sake of simplicity we consider the case where every interaction has a unique product (as in section 3.4, for instance). Next we lessen the connectivity of the reaction graph by setting randomly chosen rate constants to zero with probability p_{el}. The zeros in the interaction matrix represent "elastic" interactions. In addition, after generating an interaction matrix with a given p_{el}, we randomly introduce copy (replication) interactions, with probability p_{self} among the nonzero entries. If an interaction between i and j has been determined to be a copy action, then the product is chosen to be either i or j with equal probability. The relative magnitude of the replicator field $\mathcal{R}_k(x)$ is controlled by p_{self}.

For a system of n types we have n^3 entries of the interaction matrix, α_{ij}^k, and we only need to determine the (unique) product k for the n^2 entries corresponding to all pairs (i,j); All other entries are zero. For the transmission coefficients t_{ij}^k we have for every pair (i,j) and for any k:

$$t_{ij}^k = \begin{cases} 1 & \text{with probability} \quad \frac{1}{n-2}(1-p_{el})(1-p_{self}) \quad \text{for } k \neq i,j \\ 0 & \text{with probability} \quad \frac{1}{n} p_{el} \\ 1 & \text{with probability} \quad \frac{1}{2}(1-p_{el}) p_{self} \quad \text{for } k = i \text{ or } j \end{cases} \tag{25}$$

The rate constants are given by $\alpha_{ij}^k = a_{ij} \cdot t_{ij}^k$, with the a_{ij} drawn uniformly on $(0,1)$. In some instances we simply fixed $a_{ij} = 1$. We refer to this latter case as the "0/1-model."

In almost all of several hundred numerical integrations with $p_{el} = p_{self} = 0$ for n between 5 and 20 the system converged to a globally stable fixed point in the interior of the simplex. This behavior is reminiscent of a much more complicated model that studies a similar problem cast in terms of catalyzed ligation and cleavage reactions among polymers [2]. In a few cases our system reduced its dimension until it reached a set of types with no closed subsets. In all these cases the entire surviving network was a single replicator.

One is tempted to conjecture that the network equation (1) with no closed subsets has one globally stable rest point $\mathbf{x} > \mathbf{0}$. Unfortunately, a counterexample is given by the life-cycle system of section 6.2.2. This system contains no closed subsets, yet produces a limit cycle. This example is rather contrived and is very unlikely to occur by chance. Moreover, it has a very sparse interaction matrix that requires $p_{el} \gg 0$. Also note that the addition of only one single arc at an appropriate place collapses its limit cycle into a rest point. At small p_{el}, however, the system is exceedingly likely to contain no closed subsets (therefore becoming permanent), as well as being highly interconnected. As long as $p_{self} < 1$ the field $\mathcal{J}_k(x)$ acts strongly like a mutation field [28], resulting in a flow that points into the interior of the simplex everywhere. This leads to attractors located near the center $(1/n, \ldots, 1/n)$.

The productivity, Φ (see equation (4)), is a global quantity that reflects the behavior of the system as long as we are dealing with stable rest points. In figure 1a we show the dependency on p_{el} and p_{self} of the average equilibrium productivity $\langle\Phi(\infty)\rangle$ for 10-dimensional systems. At each grid point we average over 20 interaction matrices chosen at random according to (25). In view of the results of many preliminary integrations, we tested only one initial condition chosen randomly for each instance. Figure 1b plots the average dimension of the support at equilibrium—the number of surviving types—under the same conditions as figure 1a. Figures 1c and 1d show the corresponding data for the 0/1-model.

Since we only know a few rigorous results about the network equation, it is rather difficult to unravel the effect of variations in p_{el} and p_{self} on the network's dynamical behavior. Thus, a more phenomenological approach may help.

$$\begin{aligned}
\langle\Phi\rangle &= \sum_{p=1}^{n}\sum_{q=1}^{n}\langle x_p x_q \sum_r \alpha_{pq}^r\rangle \approx \langle\sum_r \alpha_{pq}^r\rangle \sum_{p=1}^{n}\sum_{q=1}^{n}\langle x_p x_q\rangle && (26)\\
&\approx \langle a_{pq}\rangle\langle t_{pq}^r\rangle \cdot n^2 \cdot \overline{x_p x_q} \approx \langle a_{pq}\rangle\langle t_{pq}^r\rangle(n^2\overline{x_r}^2 + R) && (27)\\
&\approx a \cdot (1-p_{el}) \cdot (1+R). && (28)
\end{aligned}$$

Brackets denote ensemble averages, bars are system averages. Equation (26) is a mean field approximation for highly interconnected systems. Next we replace ensemble averages by system averages, and extract the correlation between concentrations, R, from the second moments. Because attractors are located near the center we set $\overline{x}_r \approx 1/n$. Where n is the dimension of the system at equilibrium, and a is 1/2 for uniformly distributed rates and 1 for the 0/1-model. The $(1-p_{el})$ term results from averaging the transmission coefficients in equation (25).

When all self-replicators are suppressed, $p_{self} = 0$, we observe a linear dependence of $\Phi(\infty)$ on p_{el}. This is what one would expect from (28) if there were no correlations, R. This leads to the assumption that correlations are introduced mostly by switching on replication interactions. Indeed, a plot of $R = \langle\Phi\rangle/[a\cdot(1-p_{el})] - 1$

(figure 1e) confirms that at low p_{self} the correlations, R, are almost independent of p_{el}. As p_{self} increases, however, low connectivity enhances the correlations introduced by individual replicators. In view of the brutal assumptions underlying (26)-(28) we are surprised to observe that the linear dependence $\Phi = a \cdot (1 - p_{el})$ holds even for very sparse systems produced at high p_{el}.

In figure 1b the average number of surviving types is seen to remain at or close to the initial dimension of the system, $n = 10$, until around $p_{el} = 0.6$ where the equilibrium dimension begins to drop quickly. The equilibrium set of types is not likely to contain closed subsets. (Equilibrium states with closed subsets can be found, at very low probabilities, for example, in replicator equations.) We will refer to the interaction graph as the directed graph with the set of types $S = \{1, 2, \ldots, n\}$ as its vertex set, and with edge set $E = \{i \to k \mid \alpha_{ij}^k > 0 \text{ for some } j\}$. The interaction graph is a representation of the transformation pathways. The decomposition of section 5.2 corresponds to the decomposition of the interaction graph into connected components.

A set of types that contains no closed subsets implies that the corresponding interaction graph is strongly connected. Strong connectivity, however, does not imply that the set of vertices contains no closed subsets. Nevertheless, it seems that randomly generated directed graphs that are strongly connected are also unlikely to contain closed subsets. In fact, for $p_{self} = 0$ the curve reflects the changes in the connectivity of the system. At $p_{el} \ll 1$ the graph contains $\mathcal{O}(n^2)$ directed edges. From the theory of random graphs [3] we know that if the number of edges reaches the threshold of $\mathcal{O}(n \log n)$ an undirected graph is connected with probability one as $n \to \infty$. We suspect something similar holds for strong connectedness in directed graphs.

The transition to dimension reducing behavior becomes sharper as the system size increases (not shown), as expected from random graph theory. This transition is not reflected in Φ.

On the other extreme, when p_{el} is fixed at 0, increasing p_{self} gradually turns off the field $\mathcal{J}_k(x)$ in (24) and switches on $\mathcal{R}_k(x)$: the network equation increasingly resembles a replicator equation. Since replicator equations have a variety of different dynamical behaviors [11] we stopped our statistics at $p_{self} = 0.95$. When reaction rates were chosen to be all degenerate at unity, we frequently hit fixed-point manifolds at $p_{self} > 0.8$, while random reaction rates produced basins of attraction for a variety of locally stable rest points. This explains the difference in figures 1b and 1d at high values of p_{self} and low values of p_{el}. Again, as p_{self} increases, the connectivity of the system decreases, the system contains more closed subsets that compete with each other according to replicator dynamics, and dimension reducing behavior suddenly begins in a fashion similar to the dependency on p_{el} along the $p_{self} = 0$ cut. The overall shape of the surface in figure 1b indicates that the effect of p_{self} is similar to that of p_{el} with respect to connectivity.

Figure 1a shows that a higher frequency of replicators tends to increase the overall productivity Φ at equilibrium. According to (28) the increase in Φ at constant p_{el} is attributed to the introduction of positive correlations among the types as p_{self} is increased and interactions become predominantly replications. The decrease of Φ with increasing p_{el}, at constant p_{self}, is expected from (28), and due to the decrease in connectivity. At higher p_{self} the linear dependence of the flow on connectivity breaks down, pointing to an additional dependence of the concentration correlations on the system's connectivity (see also figure 1e).

6.2 Special systems

In this section we briefly discuss a few systems with internal structures that do not usually arise in random matrices. According to the decomposition in section 5.2 we may build network equations with a set level structure given by well-known systems, for example, hypercycles. While typical hypercycle models operate with replicators that have no further internal structure, we may now couple entire networks.

6.2.1 Hypercyclically coupled networks

Figure 2 shows the form of the interaction matrix. Each diagonal block is a randomly generated network that behaves like a single replicator. The object types of block i interact additionally with those of block $i+1 \,(\mathrm{mod}\, n)$ to produce types in block $i+1\,(\mathrm{mod}\, n)$, thus providing an overall hypercyclic coupling between networks.

Figure 3 shows a projection of the limit cycle obtained for five randomly generated networks, each of which has dimension five. The limit cycle is a relaxation oscillation very similar to those present in simple replicator hypercycles [17]. The trajectory spends a long time near corners. One network therefore dominates the entire system for a certain period of time until its concentrations drops to almost zero, and the next network takes over.

It is interesting to note that the relative concentrations of the members within a block are now entrained in small oscillations (not shown), while as an isolated system each block exhibits a stable rest point.

The effect of varying the number of blocks in the system is essentially the same as varying the dimension of the hypercycle equation: stable rest points occur when the number of blocks is less than or equal to four, and limit cycles result for systems with more blocks. Varying the internal coupling within each block relative to the coupling among the blocks results in the same behavior as for the hypercycle case: increasing the internal coupling beyond a threshold lets the trajectory converge toward a heteroclinic orbit as in the May-Leonhard model [16].

For networks coupled in a replicator fashion, equation (17), we may expect chaotic behavior as observed in simple replicator equations [23].

Uncoupled networks mimic yet another special case of the replicator equation known as the Schlögl model. Competing independent self-replicators lead to the survival of one single type, depending on the initial condition. This "once for ever selection" is observed at the set level as well.

6.2.2 Life cycles

As opposed to a system of interacting replicators we now consider an organization that bears some resemblance to a biological life cycle. Objects of type k interact only with objects of the same type to produce objects of type $k+1$, and type n objects interact with each other to yield objects of type 1, thus closing the cycle. For simplicity we assume equal rate constants:

$$\dot{x}_k = x_{k-1}^2 - x_k \Phi(t). \tag{29}$$

It is easily seen that there is no rest point on the boundary, and that for any interior rest point the relation $x_k = (1/\Phi)x_{k-1}^2$ must hold. Iterating we obtain $x_k = x_k^{2n}/\Phi^{2^n-1}$, that is $\Phi = x_k$ for any k. Thus $\hat{\mathbf{x}} = (1/n, \ldots, 1/n)$ is the unique rest point of the system.

Theorem. The rest point $\hat{\mathbf{x}}$ is asymptotically stable for equation (29) if $n \leq 4$, and unstable otherwise.

Proof. The Jacobian at $\mathbf{x}$ is given by

$$J_{kl}(\mathbf{x}) = \frac{\partial \dot{x}_k}{\partial x_l} = 2x_{k-1}\delta_{k-1,l} - \delta_{kl}\Phi - 2x_k x_l \tag{30}$$

At the rest point $\hat{\mathbf{x}}$ this reduces to

$$J_{kl}(\hat{\mathbf{x}}) = \frac{1}{n}\left(2\delta_{k-1,l} - \delta_{kl} - \frac{2}{n}\right) \tag{31}$$

which is a circulant and therefore has eigenvalues

$$\lambda_k = -\frac{2}{n}\delta_{k0} + \frac{2}{n}e^{2\pi i k/n} - \frac{1}{n} \tag{32}$$

$\lambda_0 = -1/n$ corresponds to the unphysical direction. Thus $\hat{\mathbf{x}}$ is stable when

$$\cos(2\pi/n) < \frac{1}{2}. \tag{33}$$

Inequality (33) holds for $n < 6$, for $n = 6$ the rest point is nonhyperbolic, and for $n > 6$ it has eigenvalues with positive real part. This behavior is reminiscent of the hypercycle, with the exception that in the latter the transition from globally stable behavior to oscillatory behavior occurs at $n = 4$.

Figure 4 shows a projection of the limit cycle of equation (29) for $n = 10$. The trajectory is not of the relaxation oscillation type seen in the hypercycle case. The flow $\Phi(t)$, for example, remains (almost) constant along the trajectory. Our numerical experiments support the conjecture that equation (21) has a globally stable rest point for $n \leq 5$ and a globally stable limit cycle for $n \geq 7$ at the set level as well.

As mentioned previously, the behavior of the life-cycle equation is rather fragile with respect to the introduction of additional interactions. For instance, in the 10-dimensional example shown in figure 4 it suffices to set $\alpha^{10}_{4,5} = 1$, that is, to have type 5 producing 10 by acting on 4, and the system converges to a stable rest point. On the other hand, less strategically placed interactions, do not harm the limit cycle. A superposed counterrotating cycle with $\alpha^k_{k,k+1} = 1$ (hypercycle) or $\alpha^k_{k+1,k-1} = 1$ (no replicators) preserves the limit cycle. These structural perturbations, however, still have a purely cyclic organization.

7 Conclusions

In section 5.1 we alluded to a system of objects together with a law of composition that determines how two objects map into a third. Such algebraic structures may be viewed to some extent as abstractions of naturally occurring systems, as for instance chemistry. It is interesting to ask what happens if one defines a dynamical system on top of such an algebraic structure. A simple example has objects interacting with each other according to mass action kinetics, and thereby producing new objects in accordance with a law of composition [7]. In this paper we address the simplest instance of such a situation. The algebraic structure is defined here to be a finite look-up table. This look-up table may be generated at random or it may contain some prescribed structure. Mass action kinetics in a population constrained to a constant size in a continuously stirred tank reactor leads to a straightforward generalization of replicator equations, *cf* (1).

To generalize this system we developed the catalytic network equation. Our research indicates that the typical behavior of random networks is extremely robust. It consists of a globally stable rest point in the interior of the state space. Clearly, this is not the only behavior that can be found, since the network equation contains many systems known to exhibit complicated dynamics as special cases [25]. All these cases, however, are distinguished by some particular underlying algebraic structure, as, for instance, hypercycles. In fact, we have coupled entire subsystems in a hypercyclic fashion. The resulting global behavior is very similar to the elementary hypercycle, but the internal dynamics of the subsystems changed from rest points to oscillatory behavior. More generally, the networks allow us to study replicator dynamics of replicators endowed with some complex internal structure.

We looked at the average behavior of the equilibrium productivity Φ as well as at the dimension of the support at equilibrium for a large number of randomly chosen interaction matrices. The averages were studied as a function of the density of elastic interactions and of the density of copy interactions. The former controls the sparseness of the interaction matrix, while the latter tunes the transition to pure replicator dynamics. We argue from these data that the main effect of individual replicators is to introduce strong positive correlations among the concentrations.

We showed that the systems always reduce their dimension to a self-maintaining subset of types. It is difficult to detect structure at the level of individuals in networks made up of a very large number of different object types. One may therefore group individuals to sets, and investigate the behavior of the system at the set level. We showed that a partition into replicator sets as well as sets transforming into each other in a cyclic fashion can always be made uniquely. However, in most randomly generated cases, the interconnections of production pathways are so dense that the entire system is just one single replicator. These replicators usually have an internal structure, whose decomposition into sets is no longer unique. This leads to an interesting clustering problem that deserves future attention.

A further line of inquiry seeks to understand how robust a particular interaction structure is with respect to modifications. We found a simple example for a cyclic interaction pattern of a non-replicator kind exhibiting limit cycles that can be destroyed by introducing a single additional interaction. All object types survive at the new rest point. At the same time there are several ways to structurally perturb the system by adding many interactions without affecting the limit cycle behavior. Hypercycles, in contrast, react quite differently to such modifications. Usually a short cut in a large hypercycle will lead to the extinction of some species, while preserving limit cycle behavior in the reduced system. Our simple system can be used to study issues of organizational stability that go beyond the mere addition of new interactions, but also include the addition of new types (that is, dimensions) and their interactions as well.

Acknowledgements

During part of this work P.F.S. was a postdoctoral fellow at the Max-Planck-Institut für biophysikalische Chemie in Göttingen. Thanks to Manfred Eigen and John McCaskill for stimulating discussions. We are particularly grateful to Josef Hofbauer, University of Vienna, for helpful comments and clarifications, and to Rob de Boer, University of Utrecht, for the source code of the ODE integration package GRIND from which we took the Rosenbrock-integrator. The Santa Fe Institute is sponsored by core funding from the U.S. Department of Energy (ER-FG05-88ER25054), the National Science Foundation (PHY-8714918), and the MacArthur Foundation.

References

[1] L. Altenberg and M.W. Feldman. Selection, generalized transmission and the evolution of modifier genes I: Reduction principle. *Genetics*, 117:559–572, 1987.

[2] R.J. Bagley and J.D. Farmer. Spontaneous emergence of a metabolism. In C.G. Langton, C. Taylor, J.D. Farmer, and S. Rasmussen, editors, *Artificial Life II*, pages 93–140. SFI Studies in the Sciences of Complexity, Addison-Wesley, Redwood City, 1991.

[3] B. Bollobás. *Random Graphs*. Academic Press, London, 1985.

[4] M. Eigen and P. Schuster. *The Hypercycle*. Springer Verlag, Berlin, 1979.

[5] J.D. Farmer, S.A. Kauffman, and N.H. Packard. Autocatalytic replication of polymers. *Physica D*, 22:50–67, 1986.

[6] R.A. Fisher. *The Genetical Theory of Selection*. Clarendon Press, Oxford, 1930.

[7] W. Fontana. Algorithmic chemistry. In C.G. Langton, C. Taylor, J.D. Farmer, and S. Rasmussen, editors, *Artificial Life II*, pages 159–210. SFI Studies in the Sciences of Complexity, Addison-Wesley, Redwood City, 1991.

[8] W. Fontana and L.W. Buss. A calculus of the ideas immanent in evolution. Technical report, Santa Fe Institute, 1992.

[9] M.R. Gardner and W.R. Ashby. Connectance of large dynamic cybernetic systems: critical value for stability. *Nature*, 228:778, 1970.

[10] J. Hofbauer. On the occurrence of limit cycles in Lotka-Volterra equations. *Nonlinear Analysis*, 5:1003–1007, 1981.

[11] J. Hofbauer and K. Sigmund. *The Theory of Evolution and Dynamical Systems*. Cambridge University Press, LMSST 7, Cambridge, 1988.

[12] S.A. Kauffman. Autocatalytic sets of proteins. *J. Theor. Biol.*, 119:1–24, 1986.

[13] L. Kindlmann and K. Rejmánek. Number of species at stable equilibrium of complex model ecosystems. *Ecol.Modelling*, 16:85–90, 1982.

[14] R.M. May. Will a large complex system be stable? *Nature*, 238:413–414, 1972.

[15] R.M. May. *Stability and Complexity of Model Ecosystems.* Princeton University Press, Princeton, 1973.

[16] R.M. May and W. Leonhard. Nonlinear aspects of competition between three species. *SIAM J.Appl.Math.*, 29:243–252, 1975.

[17] P.E. Phillipson and P. Schuster. Analytic solutions of coupled nonlinear rate equations. *J.Chem.Phys.*, 79:3807–3818, 1983.

[18] S.L. Pimm. *Food Webs.* Chapman and Hall, 1982.

[19] S.L. Pimm, J.H. Lawton, and J.E. Cohen. Food web patterns and their consequence. *Nature*, 350:669–674, 1991.

[20] S. Rasmussen. Toward a quantitative theory of the origin of life. In C. Langton, editor, *Artificial Life I*, pages 79–104. Addison-Wesley, Redwood City, 1988.

[21] G.E. Revesz. *Lambda-calculus, Combinators, and Functional Programming.* Cambridge University Press, Cambridge, 1988.

[22] A. Roberts. The stability of feasible ecosystems. *Nature*, 251:607–608, 1974.

[23] W. Schnabl, P.F. Stadler, C. Forst, and P. Schuster. Full characterization of a strange attractor: Chaotic dynamics in low dimensional replicator systems. *Physica D*, 48:65–90, 1991.

[24] P. Schuster and K. Sigmund. Replicator dynamics. *J.Theor.Biol.*, 100:533–538, 1983.

[25] P.F. Stadler. Selection, mutation, and catalysis. Universität Wien, PhD thesis, 1990.

[26] P.F. Stadler. Complementary replication. *Math.Biosc.*, 107:83–107, 1992.

[27] P.F. Stadler and R. Happel. The probability for permanence. *Math.Biosc.*, in press, 1992.

[28] P.F. Stadler and P. Schuster. Mutation in autocatalytic reaction networks – an analysis based on perturbation theory. *J.Math.Biol.*, 30:597–631, 1992.

Figures and captions

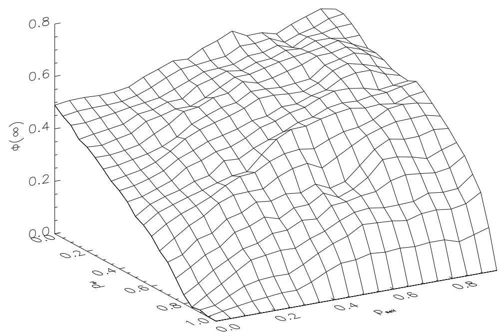

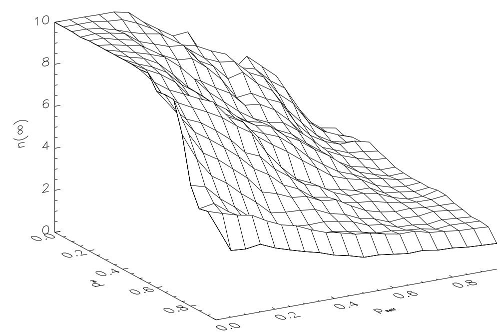

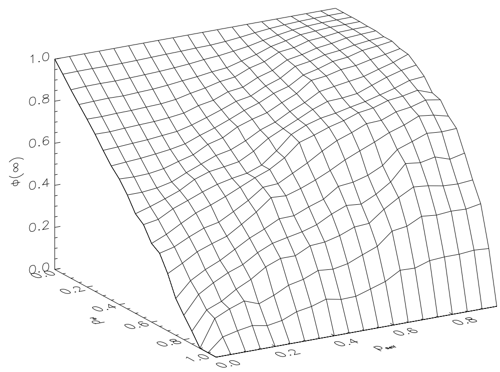

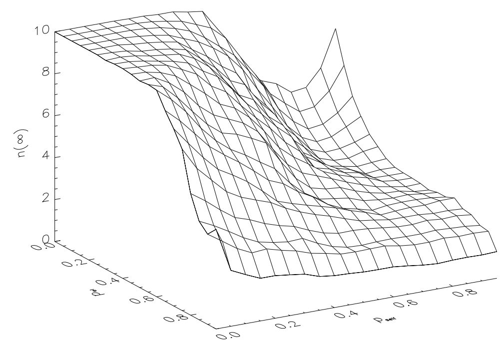

Figure 1.

a The equilibrium productivity $\Phi(\infty)$ is shown as a function of the density of elastic interactions (zero entries), p_{el}, and of the density of individual replicators, p_{self}, in a 10-dimensional system. Each grid point is an average over 20 random interaction matrices. The transmission coefficients were chosen randomly according to equation (25), the rate constants a_{ij} are uniform in $(0, 1)$.

b As in figure 1a, but the property shown is the dimension of the support at equilibrium.

c As in figure 1a, but rate constants are all degenerate, $a_{ij} = 1$ (0/1 model).

d As in figure 1b, but rate constants are all degenerate, $a_{ij} = 1$.

e The correlation, R, in equation (28) is determined from the measured average Φ of figure 1a, and plotted as a function of p_{el} and p_{self}.

Figure 2.
The interaction matrix of a system consisting of five self-maintaining networks (dark diagonal blocks) coupled in a hypercyclic fashion. Types interacting within each diagonal block i as well as in conjunction with types of block $i - 1 \bmod n$ (light off-diagonal blocks) produce types in block i. No single interaction is a copying process ($p_{self} = 0$).

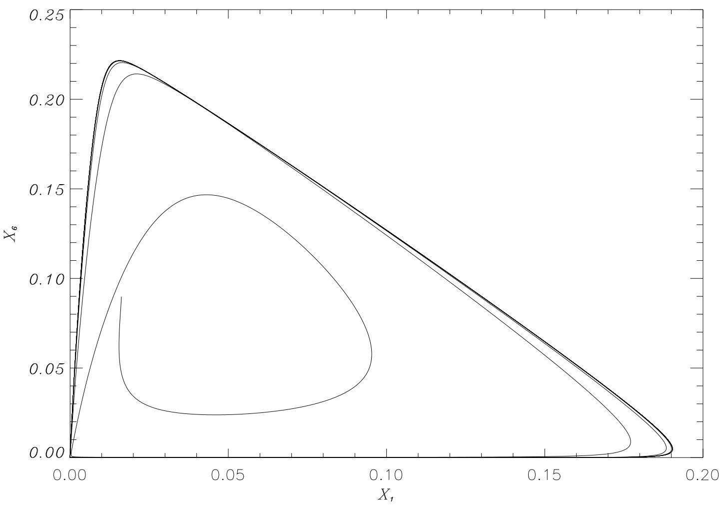

Figure 3.
The plot shows the relaxation oscillation of a 25-dimensional system consisting of five 5-dimensional self-maintaining random networks coupled as a hypercycle.

The form of the interaction matrix is shown in figure 2. The coordinates are relative concentrations of two types in different networks.

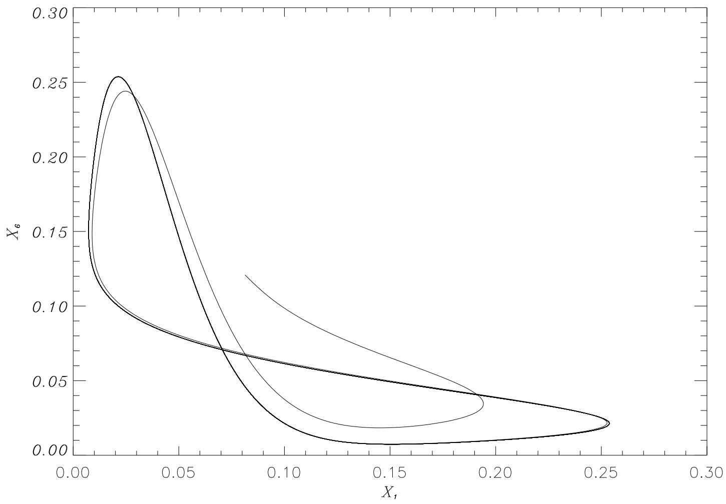

Figure 4.
The plot shows the limit cycle obtained from a 10-dimensional instance of the "life-cycle" equation (29).