SimSoup: An Artificial Chemistry Model for Investigation of the Evolution of Metabolic Networks

Chris Gordon-Smith

c.gordonsmith@gmail.com,
WWW home page: http://www.simsoup.info

Abstract. The mechanism for evolution in the first lifeforms is a key question that must be addressed by any explanation of the Origin of Life. The SimSoup artificial chemistry model is described, and it is shown how catalytic reactions can be represented as 'compound interactions'. A possible mechanism for inheritance in metabolic networks is outlined using graphical notation developed for SimSoup. Preliminary results from computer simulations are presented; it is demonstrated that a SimSoup network has multiple persistent states and that transitions between these states can occur as a result of perturbations or random fluctuations. It is argued that this may have relevance to understanding the mechanism of evolution in early lifeforms.

1 Introduction

Among the many theories of the origin of life, two major viewpoints can be identified; I refer to these as the Genetic View and the Metabolic View.[1] These can be briefly described as follows:-

- **Genetic View:** Template replicating molecules or crystals were crucial for the origin of life, and have from the outset been the carriers of inherited information that makes evolution possible
- **Metabolic View:** The first living entities were metabolic systems, and they evolved by exploring the possibilities for new kinds of metabolic network.

The well known RNA World theory is an example of a Genetic View theory. Cairns-Smith's clay crystal theory [1] can also be regarded as essentially genetic in nature. Examples of Metabolic View theories are those discussed by Oparin [8], Kauffman [6], Dyson [2], Segré et al. [9], [10], [11], Jain and Krishna [5], and Krishna [7].

A key challenge for Metabolic View theories is to explain how an essentially network oriented system can evolve. The conditions for inheritance in such networks have been modelled by Segré et al. in [11]. A model for the appearance

[1] Another important categorisation of viewpoints is the distinction made by Wächtershäuser [13] between heterotrophic and autotrophic theories.

of highly non-random organisations in chemical networks on the prebiotic earth has been investigated by Jain and Krishna in [5], and by Krishna in [7].

These models work by representing the catalytic influence of each molecular species on the formation of other molecular species. The modelled species themselves are assumed to be formed from 'food' reactants that have constant concentrations. Krishna [7] identifies some potential drawbacks of this assumption.

2 Typographical Convention

In this paper, terms with a specific meaning in the SimSoup model are capitalized. The word 'Molecule' therefore refers to an object in the SimSoup model, whereas the word 'molecule' has the meaning generally used in chemistry.

3 The SimSoup Artificial Chemistry Model

SimSoup is an artificial chemistry model that includes an abstract representation of the static and dynamic properties of a chemical network. The model is designed to enable the properties of such networks to be explored, particularly with a view to investigating evolutionary properties.

SimSoup does not assume a constant concentration of 'food' molecules. However, in its current form it does make the following key simplifications:-

- The detailed structure of molecules is not represented
- Interactions are assumed to take place in a well stirred reactor
- The energy of molecular species is not represented
- There is a fixed number of molecular species.

The following sub-sections describe the logical structure and behaviour of the model in terms of the model entities and the mechanisms by which they interact.

3.1 The Static Model

In SimSoup, the laws of Chemistry are represented by a set of Molecule Types and Interaction Types that set a framework within which dynamic behaviour can take place.

Molecule Types and Interaction Types: A **Molecule Type** is analogous to a molecular species in real chemistry. It has a mass, and can be associated with Interaction Types in which individual Molecules participate as Reactants or Products. In addition, SimSoup Molecule Types can represent unstable intermediates that are not normally referred to as molecules by chemists.

An **Interaction Type** is a category of interaction between Molecule Types. Each Interaction Type has either one or two Reactant Molecule Types, and

either one or two Product Molecule Types. Different Interaction Types can be combined to produce Compound Interactions (see section 3.3). In SimSoup there are three forms of Interaction Type, representing the simplest interactions. These are shown in Table 1.

Form of Interaction Type	**Interaction Scheme**	**Interaction Rate**
Construction	$A + B \rightarrow C$	kab
Transformation	$D \rightarrow E$	kd
Fission	$F \rightarrow G + H$	kf

Table 1. The three forms of Interaction Type in SimSoup

The table shows each form of Interaction Type and its Interaction Scheme. For example, the Interaction Scheme for a Construction is $A + B \rightarrow C$. This signifies that when a Construction occurs, it consumes one Molecule of each of types A and B, and produces a single Molecule of type C.

Each Interaction Type has a Rate Constant k that determines the Interaction Rate (see section 3.2). In addition, each Interaction Type must conserve mass. For example, in the Fission $F \rightarrow G + H$, the mass of Molecule of Type F must be the same as the sum of the masses of G and H.[2]

As a model simplification, Constructions with both Reactants of the same Molecule Type and Fissions with both Products of the same Molecule Type are ruled out.

In the current version of SimSoup, all Molecule Types and Interaction Types must be specified at the outset. The possibility of extending the model to enable a more 'open ended' approach is discussed briefly in section 6.

Graphical Representation of the Static Model: The motivation for the SimSoup model is to enable investigation of networks. A graphical representation of the static model that emphasises its network characteristics is therefore useful.

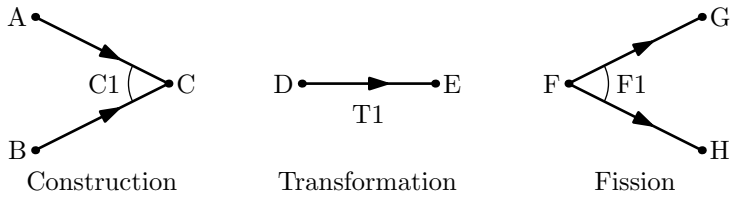

Fig. 1. Diagrammatic representation of the three forms of SimSoup Interaction Type

[2] Without this restriction, cycles that grow indefinitely can occur.

Figure 1 shows such representations of the three forms of SimSoup Interaction Type. Each distinct Molecule Type is represented by a point in a plane. An Interaction Type is represented by drawing lines joining the Reactant and Product Molecule Types. Arrows on the lines indicate the direction of Interactions from Reactant to Product. The arcs on the Construction and Fission indicate that the two lines joined by the arc form part of the same Interaction Type. The labels by the arcs, or near the centre of the single line in the case of the Transformation, are a identifiers for the Interaction Types.

The notation can be used to represent a network of Interaction Types. Figure 2 shows an example in which the Product C of Construction C1 is the Reactant for Transformation T1, and the Product D of T1 is the Reactant for Fission F1. The overall result is that a Molecule of type A can combine with a Molecule of type B, and this can result in the final production of Molecules of types E and F, via the intermediates C and D.

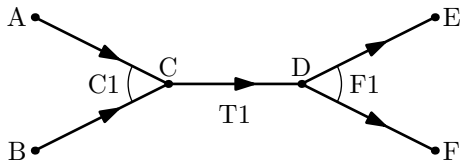

Fig. 2. Diagrammatic representation of a simple network of Interaction Types consisting of Construction C1, Transformation T1, and Fission F1.

3.2 The Dynamic Model

The static model establishes network of Interaction Types that can result in actual Interactions taking place when Molecules are introduced.

The Reactor, Molecules and Realisable Interaction Types: The SimSoup **Reactor** is envisaged as an enclosed space within which Molecules exist and interact. In the model discussed in this paper, this space is assumed to be three dimensional, with Molecules of each type being evenly distributed throughout the space. An extension of the model to cover situations in which Interactions take place on a surface as envisaged by Wächtershäuser [14] is possible.[3]

Each **Molecule** in the Reactor is an instance of a Molecule Type. At any point in time, there are a number of **Realisable Interaction Types**. An Interaction Type is Realisable (possible) if there are sufficient Reactants for one

[3] The assumption of a well stirred Reactor would not apply for a model of surface metabolism, and so the model would need to be extended to represent the locations of Molecules on the surface.

Interaction. For example, the Construction $X + Y \rightarrow Z$ is possible if there is at least one Molecule of type X and at least one Molecule of type Y.

Interaction Rate: The rate per unit volume at which Interactions of a particular type occur is dependent on the Rate Constant k and the concentration(s) of the Reactant(s). The concentration of a Molecule Type is the number of Molecules of that type present, divided by the volume of the Reactor. Each Interaction Type's Rate Constant is a randomly selected constant value.

The Interaction Rates per unit volume for different Interaction Types are shown in Table 1 in section 3.1. For a Construction $A + B \rightarrow C$ the Interaction Rate *kab* is proportional to the product of the concentrations of the Reactants. This is consistent with the behaviour of *second order* reactions in chemistry. The reason for this behaviour is that the frequency of collisions between molecules of two different species in a solution (or a gas) is proportional to the product of the concentrations.

Transformations and Fissions are analogous to *first order* reactions in chemistry. Only a single Molecule is required for these Interactions to occur and there is no requirement for a collision. The Interaction Rate is simply proportional to the concentration of the (single) Reactant Molecule Type.

The three forms of Interaction Types described in section 3.1 do not in themselves cover all of the types of reaction that can take place in chemistry. For example, they do not represent reactions of the form $A + B \rightarrow C + D$. In addition, they do not represent catalytic reactions such as $A + X \rightarrow B + C + X$, in which the catalyst, X, is necessary for the reaction to proceed but is not itself either consumed or produced in the reaction.

Such reactions cannot in general be modelled by simple rate laws of the kind described above. These more complex reactions occur in a series of steps, which chemists call **elementary reactions**. The rate of an elementary reaction depends on its **molecularity**, which is the number of molecules coming together to react. In a **unimolecular reaction**, a single molecule shakes itself apart or rearranges its atoms into a new formation. Examples of this are the isomerisation of cyclopropane to propene, and the decomposition of N_2O_2 to two molecules of NO. In a **bimolecular reaction**, two molecules collide and interact in some way. An example is the reaction $H + Br_2 \rightarrow HBr + Br$.

The rate of a unimolecular reaction is proportional to the concentration of reactant present, while the rate of a bimolecular reaction is proportional to the product of the concentrations of the reactants.

SimSoup Transformations and Fissions are unimolecular, while Constructions are bimolecular. Termolecular elementary reactions in which three molecules come together are possible, but occur infrequently due to the low probability of such collisions; they are therefore not represented in SimSoup.

Where a set of elementary reactions results in an overall reaction, the rate of the overall reaction will depend on the concentrations of the reactants at each stage. Some of the reactants involved may be short lived **intermediates**; species that are not consumed or produced by the overall reaction, but which

play a role in one of the elementary reactions. In general it is not possible to produce an equation describing the overall rate of a set of elementary reactions without taking account of the intermediates.

Operation of the Model: In contrast with models based on differential equations, SimSoup models each Molecule as a separate object, and Molecule objects are added to and removed from the Reactor as Interactions take place. The model proceeds on an iterative 'timestepping' basis: the actions for one timestep are executed, the time is incremented and the actions for the next timestep are then executed and so on. The basic operation of the model is as follows:-[4]

- Execute actions for current time:-
  - For each Realisable Interaction Type:-
    - Determine the number of Interactions of this type that are to occur in the current timestep, calculated according to the Interaction Rate formulae in Table 1. Under certain circumstances it is possible for the calculated number to exceed the available number of Molecules of a Reactant (for example in the case of a Construction where one Reactant has a very high concentration and the other has a very low concentration). In such cases, limit the number of Interactions to $\frac{1}{10}$ of the number of the 'scarce' Reactant Molecules available, or if the number available is 10 or less, set the number of Interactions such that there is a $\frac{1}{10}$ probability that all the Molecules of the scarce Reactant are removed, and a $\frac{9}{10}$ probability that no Interactions occur at this timestep
    - Add the Product Molecules and remove the Reactant Molecules for each Interaction
  - Add the 'food' Molecules (see section 5.1)
  - Add perturbation Molecules if required at this timestep (see section 5.1)
  - Remove the 'leakage' outflow Molecules (see section 5.1)
  - Increment Time
- Go back to start and repeat for new timestep

Maintaining Realisable Interaction Types: Whenever a Molecule is added to or removed from the Reactor, a check is made to determine whether any Interaction Types have as a result become Realisable, or have stopped being Realisable. For example, suppose the Reactor contains 15 Molecules of type A, and one Molecule of type B. If the Molecule of Type B is removed for any reason (eg a Transformation with scheme $B \rightarrow C$), then any Realisable Interaction Types with B as a Reactant are no longer Realisable, and so will not be considered at the next timestep.

[4] The model has some similarities to the Gillespie algorithm described in [3] and [4]. A key difference is that the Gillespie algorithm proceeds from one molecular interaction to the next, calculating the time and type of the next reaction at each step, whereas SimSoup proceeds in fixed timesteps, calculating the number of Interactions of each type at each timestep.

3.3 Compound Interactions

In SimSoup, a set of elementary reactions that combine to produce an overall reaction is modelled as a set of Interaction Types that are connected in a way that enables **Compound Interactions** to take place.

The simple network of Figure 2 that was discussed in section 3.1 shows a Construction, Transformation and Fission combining to form a Compound Interaction with the overall scheme $A + B \rightarrow E + F$. The intermediates are C and D, and the overall rate depends on the concentrations of these intermediates as well as of A and B. If the network in Figure 2 is considered in isolation, then the overall dynamic behaviour of the system is determined by the concentrations of A and B and the rate constants for the three Interaction Types. The concentrations of the intermediates in this case are not independent variables. However, in the more general case, the intermediates may also be participating in other Interaction Types, and the behaviour of the system cannot be understood without taking account of these.

For this reason SimSoup does not explicitely represent 'Compound Interaction Types'. Instead, Compound Interactions are simply the result of the simple Interaction Types combining in particular ways.

Catalysis in SimSoup: In SimSoup, catalytic reactions are represented as Compound Interactions that occur as a result of a set of linked Interactions.

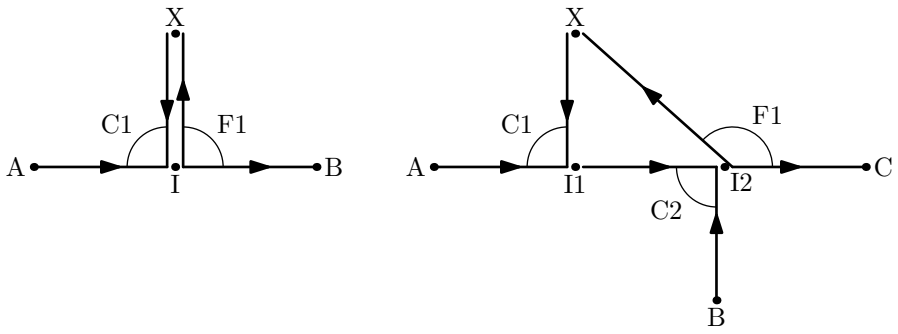

a) Catalytic Transformation b) Catalytic Construction

Fig. 3. Diagrammatic representation of two catalytic Compound Interactions

Figure 3 shows two examples of Compound Interactions involving a catalyst. Figure 3a shows a Catalytic Transformation in which a Molecule of type A is converted to a Molecule of type B with the assistance of a catalyst Molecule of type X. The Catalytic Transformation takes place as two separate Interactions, a Construction followed by a Fission. The Construction produces the intermediate

I, which then undergoes a Fission, releasing a Molecule of type B and the catalyst. The overall scheme for this can be written as:- $A \xrightarrow{x} B$.

Figure 3b shows a Catalytic Construction, in which Molecules of types A and B combine to form a Molecule of type C with the assistance of a catalyst X. In this case there are three separate Interactions (two Constructions and a Fission) and two intermediates (I1 and I2). The overall scheme for this can be written as $A + B \xrightarrow{x} C$.

Note that the concept of molecularity discussed in section 3.2 is not relevant for Compound Interactions; it is only relevant to the three simple forms of SimSoup Interaction Type. Thus, the Catalytic Transformation of Figure 3a does not have an overall molecularity; it is a (bimolecular) Construction followed by a (unimolecular) Fission. Similarly, the Catalytic Construction of Figure 3b is a (bimolecular) Construction followed by another (bimolecular) Construction followed by a (unimolecular) Fission.

4 A Mechanism for Inheritance in Metabolic Networks

This section uses the notation developed in section 3 to illustrate a mechanism for inheritance in metabolic networks. The idea that the prebiotic world could have included 'protocells' containing molecules that were able to reproduce by a process of incorporating 'food' molecules into a network of metabolism-like reactions and dividing when a certain size is reached has been put forward in a number of theories of the origin of life. An example is The Lipid World described in [9].

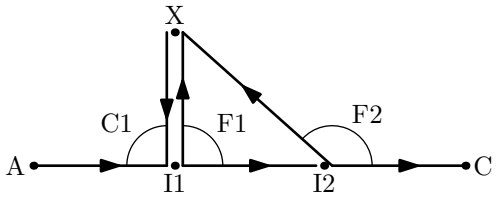

Fig. 4. A Compound Interaction that can 'remember'

A key question in regard to such theories is whether such protocells are capable of evolution. How can hereditary information be remembered and passed on to offspring by entities that have no alphabet based genetic material that can store information in the way that the DNA of contemporary organisms does?

Figure 4 shows a Compound Interaction that includes Construction C1, followed by two Fissions, F1 and F2. Suppose that a protocell with a network of Compound Interactions exists, and it is producing a surplus of A. Initially the protocell contains no Molecules of type X, and so the Compound Interaction of Figure 4 will not take place. Now suppose that a chance event occurs such that the protocell comes into contact with, and absorbs, a small quantity of X (perhaps a quantity of X splashes into the vicinity of the protocell). The protocell will immediately begin converting its surplus of A into Molecules of C and X.

Because one new Molecule of the catalyst X is produced for each Molecule of A that is consumed, an external supply of X is not needed to keep the Compound Interaction going. This means that the network of reactions in the protocell has flipped to a new state. The protocell now has a surplus of C, and a substantial quantity of the catalyst X that is required for the production of C. If it can incorporate some or all of the C molecules into the overall network in some way, then it may be able to consume 'food' molecules faster and reproduce more quickly. When the protocell does divide, it will pass on part of its stock of X. As a result of this, its offspring will also be able to continue to produce molecules of C whenever any molecules of A are available.

In short, our protocell has adopted a new characteristic (the ability to produce C) as a result of a one-time perturbation that introduces a small quantity of X. Once this new characteristic is acquired it persists (ie it is 'remembered') and is passed on to offspring, even though the original external input of X stops.

This provides the basis for a population of protocells to evolve. If the 'C producing' protocells are able to reproduce faster than the 'non C producing' protocells, then they will come to dominate.

5 Exploration of an Artificial Chemistry using SimSoup

The SimSoup Model described in section 3 has been implemented as a computer simulation. Details of the simulation program are available at the SimSoup website [12]. Figures 5 to 7 show output from two preliminary runs of the simulation. The plots are colour coded, and also labelled to assist reading monochrome prints.

5.1 Simulation Scenarios

Scenario 1 used for the simulation run shown in Figures 5 and 6 is as follows. The static model defining the laws of chemistry has:-

- **200 Molecule Types**: Mass chosen in each case as a random integer from the range 1 - 5 inclusive
- **Interaction Types chosen to promote catalysis:** 310 Constructions and 310 Fissions are set-up in pairs, in such a way that the Product of each Construction forms the Reactant for the associated Fission so that many possibilities for Catalytic Transformations exist. The Reactants and Products

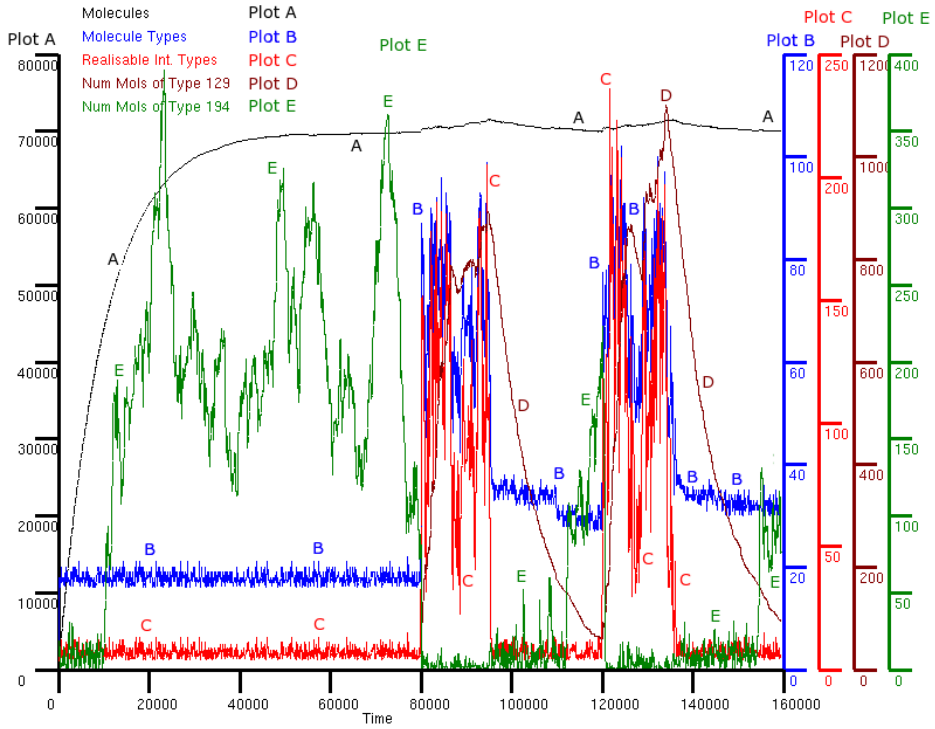

Fig. 5. SimSoup Output for Scenario 1: Plot A = Number of Molecules in Reactor, Plot B = Number of Molecule Types in Reactor, Plot C = Number of Realisable Interaction Types in Reactor, Plot D = Number of Molecules of type 129 in Reactor, Plot E = Number of Molecules of type 194 in Reactor.

are otherwise chosen at random. The Rate Constants for the Constructions and Fissions are chosen as random integers from the range 0 - 45 inclusive.

The simulation conditions driving the dynamic model are:-

- **Start with empty Reactor:** The Reactor initially contains no Molecules
- **'Food' supply:** At each timestep, three Molecules are introduced, one of each of three Molecule Types (the introduced Molecules are always of the same three types). This represents an inflow of 'food' Molecules.
- **A cap on growth:** At each timestep, a small number of Molecules are removed at random, with a removal probability of 0.0001 for each Molecule. This ensures that the total number of Molecules in the Reactor does not grow indefinitely as more 'food' Molecules are added
- **Perturb Reactor at regular intervals:** Starting at time 80,000, ten 'perturbation' Molecules of each of twenty randomly selected Molecule Type are introduced to the Reactor at intervals of 40,000 time units. Each time

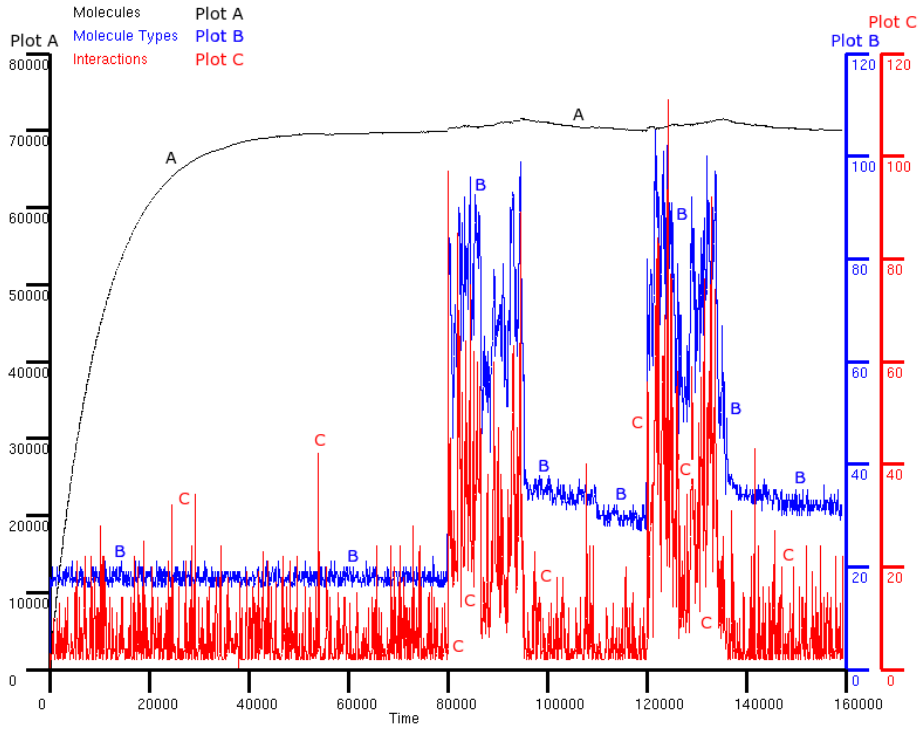

Fig. 6. SimSoup Output for Scenario 1: Plot A = Number of Molecules in Reactor, Plot B = Number of Molecule Types in Reactor, Plot C = Number of Interactions in Reactor at each timestep.

this is done, a different set of randomly selected Molecule Types is chosen. The starting time of 80,000 is chosen to allow the Reactor to reach a stable state before the series of perturbations starts. By this time the number of Molecules in the Reactor has reached almost 70,000.[5]

– **Volume:** The Reactor volume is 1000

Scenario 2 used for the simulation run shown in Figure 7 is the same as for scenario 1 described above, except that the Rate Constants are increased such that they are chosen as random integers from the range 0 - 48 inclusive.

[5] It is interesting to note that this value of 70,000 is substantially higher than the value that is reached when no Interactions are taking place. If no Interactions are occurring, then a steady state of about 30,000 Molecules is reached when the rate of removal of Molecules reaches the rate of supply of 'food' Molecules. The higher value that is reached under the scenario of Figure 5 is due to Fission Interactions consuming 'food' Molecules and replacing each one (and some of the Fission Products) with two Molecules, effectively multiplying the rate of supply. This multiplication cannot continue indefinitely because of the conservation of mass rule in SimSoup.

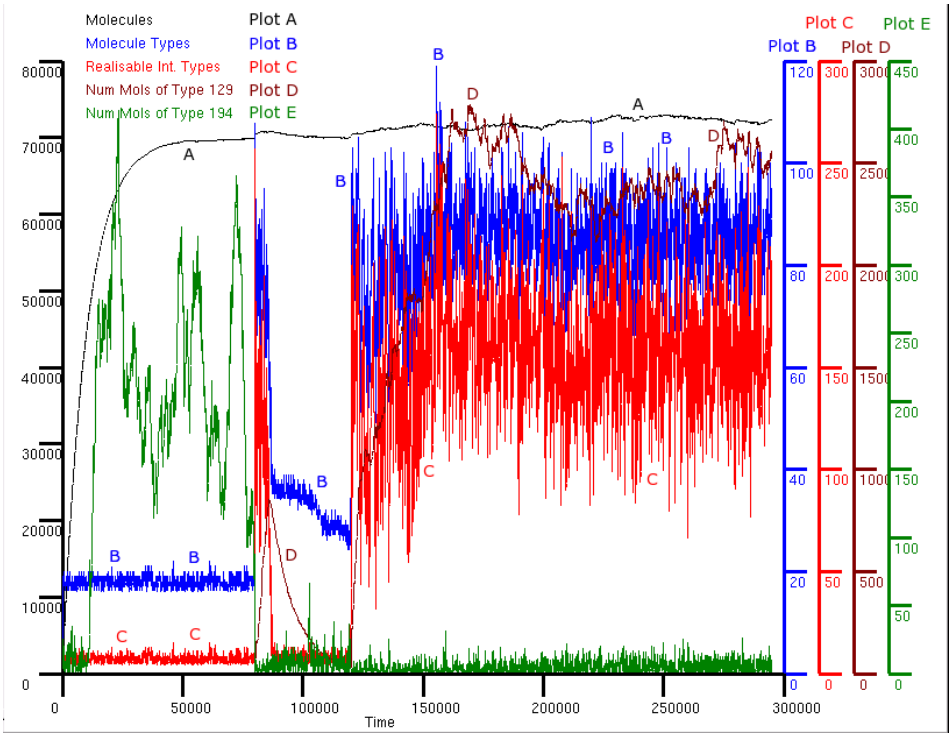

Fig. 7. SimSoup Output for Scenario 2: The simulation variables shown by plots A to E are the same as for Figure 5

5.2 Observations

The following observations on the output for scenario 1 in Figures 5 and 6 can be made:-

- **Initial interval of macroscopic stability:** During the period before the first perturbation at time 80,000, the number of Molecule Types and Realisable Interaction Types vary within narrow ranges (although the number of Molecules is still moving towards its stable value of 70,000)
- **Microscopic variability within macroscopic stability:** During this period, there is large variability in the number of Molecules of type 194. This is initially relatively stable, varying between 0 and 25. It then increases dramatically and varies between 100 and almost 400 during the period of macroscopic stability
- **Perturbations trigger more variable phases:** The first perturbation at time 80,000 triggers a substantially different phase during which the number of Molecule Type and the number of Realisable Interaction Types are

substantially higher. The subsequent perturbation at time 120,000 triggers a similar phase

- **Increased 'metabolic rate' during variable phases:** During the variable phases after each perturbation, the number of Interactions at each timestep increases substantially in synchronisation with the increase in Molecule Types
- **Catastrophic fallback:** The variable phases end abruptly with a dramatic reduction in Molecule Types and the number of Interactions per timestep
- **Slowly decaying behaviour after catastrophic fallback:** After each catastrophic fallback, the number of Molecule Types is substantially higher than during the initial stable phase, but decays slowly. This decay is much slower than the rate of removal of material from the Reactor. For example, for a period of 20,000 timesteps after the first catastrophic fallback, the number of Molecule Types falls from about 35 to about 30. During this period over 86% of the mass within the Reactor will have been removed by the 'leakage' outflow.
- **Return of microscopic variability during slowly decaying phases:** The variable behaviour of Molecule Type 194 disappears during the high metabolic rate periods, but returns during the slowly decaying phases

Other simulation runs show that the pattern in which a perturbation triggers a high metabolic rate with a subsequent catastrophic fallback continues up to time 360,000 and beyond.

The output for scenario 2 in Figure 7 shows that if the Rate Constants are increased slightly, the high metabolic rate behaviour reaches a state in which catastrophic fallback no longer occurs.

5.3 Interpretation

Perturbations or Random Fluctuations Trigger New Persistent States: The new patterns of behaviour that follow perturbation events often persist over long periods. For example, in Figure 5 the high metabolic rate phases last for about 15,000 timesteps before catastrophic fallback. In Figure 7 for scenario 2, the high metabolic rate behaviour persists for over 100,000 timesteps, and the plot suggests that it may continue indefinitely. The dramatic rise of the quantity of Molecule Type 194 at various times in Figure 5 is in each case the result of a random fluctuation, and in each case heralds a new state ('high and variable 194') that persists until the next perturbation.

The Network Remembers what it has Discovered: The persistence of new behaviour in the face of constant leakage of material can be regarded as a kind of memory. We can consider a new persistent state that is entered as a result of a perturbation or a random fluctuation to be a new kind of behaviour that the network has discovered. The fact that it can maintain the behaviour means that it 'remembers' what has been discovered. Where is the information stored? One suggestion is that it is held in the set of relative concentrations of the various Molecule Types in the Reactor. Information held in this way is described as **compositional information** by Segré et al. in [11].

Number of Persistent States A key question for further investigation is the number of persistent states. Figures 5 to 7 show that there are at least two macroscopic states: a low metabolic rate state and a high metabolic rate state. In addition, the behaviour of Molecule Type 194 in Figure 5 indicates that there are at least two variants of the low metabolic rate state.

It is not unreasonable to suppose that there may in fact be many persistent states for a SimSoup network, each with a different composition of Molecule Types and a correspondingly different network of Interactions.

Inheritance and Evolution of Networks in SimSoup If a SimSoup network can indeed have many persistent states, then the ability mentioned above to remember a newly discovered state could form the basis for an evolutionary process. If the Molecules in the Reactor were to be divided into two groups at random and placed in separate Reactors (the analogy is with division of a protocell), then it is plausible to suppose that each would continue the same pattern of persistent behaviour. They would have both inherited similar compositional information, and this would be likely to lead to a similar set of Interactions occurring.

A key feature of this mechanism for passing on inherited information is that it can be expected to work even though the offspring do not inherit exactly the same compositional information. It has been seen that persistent states in the Reactor can continue even when there is substantial variability in the number of Realisable Interaction Types and the Number of Molecule Types present (eg from time 160,000 to 340,000 in Figure 7). The persistent state is therfore an attractor, so that when deviations occur the system is self correcting, provided these deviations are within certain limits.

The signficance of the above mechanism for the Origin of Life is that it suggests a way in which evolutionary processes could have gained a foothold without the need for complex template replicating molecules and the highly evolved enzymes needed for accurate template replication.

6 Conclusion and Prospects

It has been shown how an approach that explicitely simulates individual Molecules and the Interactions between them can be used to model the behaviour of a chemical network. The SimSoup model represents catalysis in a way that is analogous to the way catalytic reactions in chemistry involve the formation of a transient intermediate that temporarily incorporates the catalyst. The model has been shown to exhibit behaviour that includes some of the properties required for evolution, specifically the ability to discover and 'remember' new patterns of behaviour. There is scope for further investigation of these properties, particularly in regard to the factors affecting the number of persistent states and the degree of stability of these states.

The simulation approach can reflect the full structure of the network connecting all reaction participants (including substrate, catalysts, intermediate prod-

ucts and final products), and allows categories of behaviour to be represented that would be difficult to model using mathematical techniques. It also provides a framework for future extensions to enable more complex behaviour to be modelled.

Such extensions could include modifications such that new Molecule Types could be produced as the simulation runs. For example, rather than specifying all Molecule Types and Interaction Types as an input to the model, it would be possible to extend SimSoup such that an Interaction between two Molecule Types that had never previously come into contact would result in the creation of a new Construction and a new Product Molecule. This would enable the network of possible interactions to be extended in an open ended way as the model runs, increasing the possibilities for evolution. Another possible extension would be to add structure to Molecules to enable investigation of questions such as the evolution of complex molecular structures. A simple set of structure dependent rules could be introduced that would enable Molecules to bind to and split from one another with the possibility that they would evolve to develop 'machine-like' properties.

References

1. Cairns-Smith, A.G.: Genetic takeover. Cambridge University Press (1982)
2. Dyson, F.: Origin Of Life. Cambridge University Press (1999)
3. Gillespie, D.: General method for numerically simulating stochastic time evolution of coupled chemical-reactions. J. Comput. Phys. **22** (1976) 403-434
4. Gillespie, D.T.: Exact stochastic simulation of coupled chemical reactions. J. Phys. Chem. **81** (1977) 2340-2361
5. Jain, S., Krishna, S.: Autocatalytic Sets and the Growth of Complexity in an Evolutionary Model. http://arxiv.org/abs/adap-org/9809003
6. Kauffman, S.A.: The Origins Of Order. Oxford University Press (1993)
7. Krishna, S.: Formation And Destruction of Autocatalytic Sets in an Evolving Network Model. PhD Thesis, Indian Institute Of Science (2003), http://arxiv.org/abs/nlin.AO/0403050
8. Oparin, A.I.: The Origin Of Life On The Earth. Oliver And Boyd (1957)
9. Segré, D., Dafna, B., Deamer, D., Lancet, D.: The Lipid World. Origins Life Evol. Biosphere **31** (2001) 119-145
10. Segré, D., Lancet, D., Kedem, O., Pilpel, Y.: Graded Autocatalysis Replication Domain (GARD): kinetic analysis of self-replication in mutually catalytic sets. Origins Life Evol. Biosphere **28** (1998) 501-514
11. Segré, D., Shenhav, B., Kafri, R., Lancet, D.: The Molecular Roots of Compositional Inheritance J. Theor. Biol. **213** (2001) 481-491
12. SimSoup: Details of the SimSoup simulation program are available at http://www.simsoup.info/
13. Wächtershäuser,G.: Evolution of the first metabolic cycles. Proc.Natl. Acad. Sci. USA **87** (1990) 200-204
14. Wächtershäuser,G.: The Origin of Life and its Methodological Challenge. J. Theor. Biol. **187** (1997) 483-494