Open-endedness and thermodynamic reversibility in algebraic chemistry

Nathaniel Virgo

Earth-Life Science Institute (ELSI), Tokyo, Japan
nathanielvirgo+alifexv@gmail.com

Abstract

The lambda calculus chemistry model of Fontana and Buss is a classic in the literature on open-ended evolution. However, it lacks any notion of thermodynamic constraints, which in turn means that there is nothing that must be "used up" in order for replication to occur. By adding reverse reactions to a similar model, we are able to implement thermodynamically reasonable kinetics, although the system still lacks mass conservation. Here we outline our implementation and present some very preliminary results from initial investigations. In particular, we find that (i) the reverse reactions mean that nothing can be permanently lost from the system, enhancing its ability to continue generating novelty, and (ii) reverse reactions seem to create an implicit selection pressure against simple copy operations, removing the need to explicitly remove them from the system as in the original model. We believe that this approach will lead to a better understanding of the role that thermodynamic considerations must play in understanding the origin of life.

Considered as a chemical system, life has some unusual properties. It is not a sparse, deterministic "clean" chemistry of the kind usually sought by synthetic chemists, but neither is it entirely the same as the classic "messy" prebiotic chemistries such as Miller-Urey synthesis, which produce a huge combinatorial explosion of products. Life's chemistry has available to it a huge combinatorial space of possible products — every possible sequence of amino acids or nucleotides could in principle be produced, not to mention the huge space of possible organic molecules whose formation can be catalysed by them — but it restricts itself to only a comparatively tiny subset, restricted in large part to those that are in some way useful for its continued propagation.

Together, the complementary properties of unboundedness and self-restriction give life the ability to modulate its own composition, a prerequisite for forming novel chemical mechanisms in response to evolutionary needs. We are interested in understanding what is needed for this particular flavour of "open-endedness" to emerge in abiotic chemistry.

A classic artificial chemistry model that approaches this idea is the lambda calculus chemistry of Fontana and Buss (1994a,b). In this model, "molecules" are lambda calculus expressions (i.e. functions in a simple but Turing-complete language), and reactions are formed from the application of one expression to another. A key result was that the system settled into more or less complex "organisations", which, roughly, are sets of species that collectively re-create one another through reactions. However, while the system is capable of generating a wide variety of different organisations, each one is a fixed point of its dynamics, and hence it does not continue to develop complexity indefinitely.

However, the Fontana and Buss model lacks thermodynamic considerations. We have recently shown that even in a simple polymerisation chemistry, the addition of reverse reactions can lead to the self-organisation of complex autocatalytic mechanisms (Virgo et al., 2016). We wish to understand the effect of reversibility on much more complicated reaction networks. With this in mind we augment Fontana and Buss' model with reverse reactions.

Adding reversibility seems to change the dynamics markedly. Information can no longer be irretrievably lost from the system, so the system cannot reach a "dead end" where new molecules can no longer be produced. We suspect that this kind of reasoning will lead to new insights into open-ended chemical evolution and the origins of life.

Methods

The original lambda chemistry works as follows. There is a "reactor" containing a constant number of lambda expressions. At each step, two random expressions are drawn from the reactor, and one is given to the other as an input to produce a third. Then all three reactions are placed back into the reactor. Another random expression is then removed from the reactor in order to keep the population constant. The reaction part has the form

$$\mathrm{A} + \mathrm{B} \longrightarrow \mathrm{A} + \mathrm{B} + \mathrm{C}, \tag{1}$$

where $C = (A\,B)$ is the result of applying A to B and then beta reducing to normal form. It is of course possible for A, B and/or C to be different instances of the same species.

As a reaction this is somewhat odd, since it produces a product without destroying the reactants. (Though one could think of it as representing the reversible binding of a ligand B

to a catalyst A, which then catalyses the formation of C from a substrate that is not explicitly modelled, followed by the disassociation of A and B.) Nevertheless, we stick with this scheme and modify it by making Equation 1 into a reversible reaction,

$$\mathrm{A} + \mathrm{B} \rightleftharpoons \mathrm{A} + \mathrm{B} + \mathrm{C}. \quad (2)$$

This removes the need to remove random molecules from the reactor, since molecules can now be removed by the reverse reaction.

Reverse reactions have the form A + B + C $\longrightarrow$ A + B. To implement this, we draw three random molecules from the reactor. We then calculate $(A\,B)$ as before, but this time we do not add the result to the reactor but instead check whether it is syntactially equal to C. If it is then this is indeed the reverse of a reaction A + B $\longrightarrow$ A + B + C, and so we remove C from the reactor. Otherwise we simply leave all three species in the reactor. It should be noted that a molecule will not necessarily be removed by the same reaction that created it, since there may be many choices of A and B such that $(A\,B) = C$.

We do not keep the reactor's population N constant, but instead allow it to vary according to the kinetics of the forward and reverse reactions. According to mass action kinetics, the forward reactions should happen at a rate $R_1 = k_1 N(N-1)$ since the two reactants are drawn without replacement, and trimolecular collisions, which may or may not result in reactions, should occur at a rate $R_2 = k_2 N(N-1)(N-2)$. Since the latter rate grows faster with N than the former the population tends to saturate, though the steady state concentration depends on the composition in the reactor.

As an important implementation detail, we used combinator calculus rather than lambda calculus as our underlying "physics". Combinator calculus is a close cousin of lambda calculus that shares with it the feature that expressions are Turing-complete functions that act upon other expressions, but it has the advantage of being easier to implement. We used the S and K combinators as a basis. We limit execution time to 1000 beta reduction steps, and we do not limit the size of expressions.

Preliminary results

At the time of writing we have only preliminary results from initial runs, so we cannot give any statistically significant results, but we can make some interesting observations.

Below are shown two expressions from two different runs of the simulation. These were both started with an initial population $\{S, S, S, K, K, K\}$ and run for 100000 time steps with $k_1 = k_2 = 1$, where each time step is a single molecular collision event. This resulted in populations of around 270 at the end of each run. We display expressions in an indentation-based format, in which $((S\,K)\,K)$ is written as $\begin{smallmatrix}\mathtt{SK}\\ \mathtt{K}\end{smallmatrix}$, whereas $(S\,(K\,K))$ is written simply as SKK. This makes it easier to discern patterns.

(i)

```
SSSSSK
    K
  S
 SSSSSK
    K
   KSKK
     SSK
      K
SSSK
 SSSSSK
      K
   S
   SSSSSK
      K
     KSKK
       SSK
        K
```

(ii)

```
SSSK
  S
KSKK
  KSKK
    KSKK
      KSKK
        SSK
          KKSK
              S
        SKK
         SKK
```

It can be seen that each of these molecules contains repeating motifs, and that these motifs are different between the two molecules. Each of the final populations contained a high diversity of products, so it will take some time to determine whether there are significant differences in the types of molecule produced by different runs. (There must be some differences, however, because execution time can differ markedly between runs.)

We believe that this diversity of end products is due to the information-preserving character of our system. No matter what state the reactor's population reaches, it is always in principle possible for it to retrace its steps and return to its initial state, from which any possible molecule can be reached. Thus there is no way for motifs or subexpressions to be permanently lost from the system, although they can become statistically very unlikely to recur.

In the Fontana and Buss model, the dynamics are dominated by "level 0" organisations, which consist of programs that simply copy their inputs, often converging to just a single species that copies itself. It was only when such simple organisations were explicitly prohibited that more complex ones arose.

However, in our system there appears to be a natural selection pressure against simple copy operations. For example, if a program Q were to simply ignore its input and produce a copy of itself as output, this program would be very susceptible to removal by reverse reactions.To avoid this negative selection pressure, programs must produce outputs that depend on their inputs; their outputs are then less likely to be destroyed by reverse reactions.

Identity functions seem to be uncommon in our system, even as subexpressions. The shortest way to write the identity function in combinatory calculus is $((SK)x)$ for any subexpression x, and these strings occur only rarely in the output. (We have not checked whether more complicated forms of the identity function are produced.) Moreover, the reactions generated by our system seem only rarely to be of the form A + B $\rightleftharpoons$ A + B + B or A + B $\rightleftharpoons$ A + B + A,

though the former does sometimes occur.

The reaction rules are chosen such that one would expect the system to reach a Gibbs equilibrium eventually. However, we suspect that after 100000 time steps our systems are still very far from this. With the parameters we use, once the system reaches a population of a few hundred the algorithm spends most of its time testing potential reverse reactions and rejecting them, so that the dynamics slow down markedly. Moreover, it might be that if run long enough the expressions would get larger indefinitely, in which case the system need not ever reach an equilibrium distribution. In future work we intend to add thermodynamic driving forces, e.g. via input and output flows or via temperature cycling, so that the system can attain non-equilibrium steady states.

Discussion

We have implemented an artificial chemistry model that is closely related to the lambda chemistry of Fontana and Buss, but which has the property of thermodynamic reversibility. This guarantees that information cannot be irreversibly lost, and data from some initial runs suggest that this leads to quite different dynamics from those of the original model.

In particular, the original model forms "organisations" that are "closed" in the sense that they cannot generate species outside of the organisation. In contrast, the information-preserving nature of our model means that if a species can be produced from the initial conditions then it can always in principle be produced from molecules present at later times, so closure can only occur at a statistical level. The lack of closure in this sense seems positive to us, since the system can never become fully closed off against the possibility of a new innovation. This hints at a deep relationship between open-endedness and thermodynamic reversibility, which might have some explanatory power in understanding why the open-endedness of the natural world has so far been absent in computational models.

There are many ways in which our model could be modified in future work. Perhaps the most important is the addition of a driving force, to keep the system away from thermodynamic equilibrium. Intuitively, remaining away from equilibrium should be necessary in order for complex behaviour to occur. Moreover, our original goal was to create a complex abstract chemistry in which something must be "used up" in order for replication to occur. To do this, one requires thermodynamic constraints together with a driving force. For the biosphere, the main driving forces are sunlight and geochemical gradients.

In order to approach the goal of having a complex algebraic chemistry with an explicit "energy source", it would be much more satisfying to have a reaction scheme where reactions consume reactants as well as producing products. The difficulty is that this must be done in such a way that information about the original reactants is not irretrievably lost. This would necessitate replacing combinator calculus with a more restricted calculus that preserves information. For example, if a calculus exists that is r-Turing complete in the sense of Axelsen and Glück (2011) (meaning that it is only able to evaluate computable invertible partial functions, but is able to compute all such functions), then it would be possible to implement reaction schemes along the lines of A + B $\rightleftharpoons$ A + C, resembling reversible catalysis. However, to our knowledge such a calculus has not yet been constructed.

More broadly, there are a wide variety of models within the emerging field of artificial chemistry, ranging from the very abstract to quite concrete models that are increasingly close to real chemistry. (See Banzhaf and Yamamoto, 2015, for a recent survey.) Many of these could be augmented with reverse reactions, and our preliminary results suggest that this may be a fertile ground for further research.

As a final comment, it is interesting that despite being quite far from a model of biological reproduction, it seems helpful to think of our system in terms of a "selection pressure" against simple copy operations. We find this interesting because for natural selection to occur one needs not just replication, but also variation and selection. Traditionally, studies in the origins of life have focused on the former more than the latter two, but if selection has important effects in complex chemical systems even in the absence of informational replicators, then perhaps a "selection first" paradigm should be added to the pantheon of approaches to the origins of life.

Acknoweldgements.

I thank Eric Smith, Norman Packard, Uwe Tangen, Nicholas Guttenberg, Stuart Bartlett and Olaf Wittkowski for discussions that stimulated and enhanced this work. This work was supported by the ELSI Origins Network (EON), which is supported by a grant from the John Templeton Foundation.

References

Axelsen, H. B. and Glück, R. (2011). What do reversible programs compute? In Hoffman, M., editor, *Foundations of Software Science and Computational Structures: 14th International Conference, FOSSACS 2011*, pages 42–56. Springer.

Banzhaf, W. and Yamamoto, L. (2015). *Artificial Chemistries*. MIT Press.

Fontana, W. and Buss, L. W. (1994a). "the arrival of the fittest": Towards a theory of biological organization. *Bulletin of Mathematical Biology*, 56(1):1–64.

Fontana, W. and Buss, L. W. (1994b). What would be conserved if "the tape were played twice"? *Proceedings of the National Academy of Sciences*, 91:757–761.

Virgo, N., Ikegami, T., and McGregor, S. (2016). Complex autocatalysis in simple chemistries. *Artificial Life*, 22(2):138–152.