

Size Selection and Adaptive Evolution in an Artificial Chemistry

Michael W. Lucht*
University of Tasmania

Abstract A simple artificial chemistry for the Squirm3 artificial environment, consisting of replicators that produce quasi-universal enzymes, is presented. The aim of this system is twofold: first, to demonstrate the survival of extracellular replicators despite the presence of faster-replicating parasites; second, to observe the evolution of adaptively useful enzymes. Accomplishing these goals will underpin future attempts to attain open-ended and/or creative evolution in the Squirm3 environment. The first aim is achieved by attaching enzymes to their replicators. Our software implementation demonstrates replicators with 10 bases prospering in the presence of parasites with zero bases. To accomplish the second aim, a process for creating selection pressure toward longer molecules is introduced. The evolution and subsequent dominance of a replicator that produces an adaptively useful enzyme is demonstrated experimentally. Finally, we comment on the crucial role played by neutral evolution and discuss the biological significance of our results.

Keywords

Evolution, self-organization, emergence, molecular simulation, origin of life, complex systems

I Introduction

There is no agreement on exactly how life first began on Earth. Most widely accepted is the RNA world hypothesis, which holds that life began with the RNA molecule [13]. Certain kinds of RNA molecules, called ribozymes, have the crucial property of being able to act as both replicator (like DNA) and catalyst (like proteins). Other researchers consider it more likely that RNA had simpler predecessors [27, 28]. A more exotic possibility is *clay theory*, which holds that clay crystals bootstrapped organic life [4].

Deciding which of these theories is correct presents considerable difficulties. Ideally, one would like to observe the first steps of life in situ. Unfortunately, there is little hope of doing so on Earth, because the environmental conditions have changed; for example, the presence of oxygen in the atmosphere. Additionally, as already commented by Darwin, living organisms tend to devour the complex biological molecules that might have been their ancestors [29]. Finally, the emergence of life might be an extraordinarily rare event, making it unlikely to be observed in the wild.

In vitro research, on the other hand, has offered glimpses into the early development of life. Famously, in 1953 the Miller-Urey experiment [25] was able to create amino acids, the building blocks of proteins, from a “soup” of simple molecules. In 1967, Mills et al. [26] studied self-replicating RNA outside a living cell. Under environmental conditions favoring RNA that replicated fastest, Mills et al. observed that the RNA lost 83% of its length, demonstrating that the elimination of dispensable

* School of Computing and Information Systems, University of Tasmania, Locked Bag 1359, Launceston, TAS 7250, Australia.
E-mail: mwlucht@postoffice.utas.edu.au

sequences conferred a selective advantage. However, the RNA required external help in the form of enzymes in order to replicate. In 2001, Johnston et al. generated ribozymes with general RNA polymerization activity [19]. Unfortunately the ribozyme's efficiency was insufficient for self-sustaining RNA replication. In 2009, Lincoln and Joyce demonstrated self-sustainable RNA replication based on two ribozymes catalyzing each other's synthesis [22]. However, the component parts for synthesis were four oligonucleotides, not the individual nucleotides themselves.

Despite these advances, it is evident that reconstructing the steps of biological evolution in a test tube, from an initial replicator to the first simple cell, remains a long way off. Even if it were to be achieved, it would still leave some fundamental questions about the origin of life unanswered. For instance, does evolution require organic chemistry to "run on" to produce the bountiful diversity of life that developed within our biosphere, or is evolution an information-processing phenomenon with organic chemistry only one out of many possible representations? And—if it were to be found to be an information-processing phenomenon—what are the essential properties that such a system must possess?

Artificial life (ALife) provides us with the means to explore these kinds of questions. ALife systems may be broadly categorized into systems that contain an explicit definition of what constitutes an individual organism within the system, and those that do not. In Geb [7], for example, each individual is issued its own neural network, and is endowed with an inbuilt set of actions (reproduce, fight, turn, move). In Geb, continuous adaptation results because of a dynamic fitness landscape. Given any population of individuals, evolution can always find individuals with slightly different behaviors that confer on them a selective advantage.

This kind of ALife system comes with inbuilt limitations. The explicit definition of what constitutes an individual organism constrains the system's evolutionary potential. For example, no matter how long Geb runs, organisms will never be able to learn how to "hide." The explicit definition of the individual also makes it unsuitable for studying abiogenesis.

Generally speaking, what constitutes an individual organism is less explicitly defined in an artificial chemistry (AC). Informally, Dittrich et al. [8] described an AC as a system similar to a real chemical system, but man-made, and argue that it is the "right stuff" for the study of early evolution. Formally, they defined an AC as a triple (S, R, A) , where the first parameter, S , stands for the set of all possible molecules, composed of atoms; R is a set of collision rules between molecules; and A is the reactor algorithm, which specifies how R will be applied to all the molecules in the system. Dittrich et al. also provided a comprehensive review of the different categories of ACs that have been developed. One of these is lattice molecular systems (LMSs), which are ACs consisting of a regular lattice. Each lattice site can either be empty or hold a single atom. Molecules are formed by forming bonds between atoms.

Several highly visual LMSs have been developed by Tim J. Hutton in a series of articles beginning in 2002. All of them take place in the Squirm3 artificial environment, which employs a greatly simplified real-world physics and chemistry. Squirm3 chemistry consists of reaction rules between two or three atoms, describing what bonds are formed or broken when atoms come in contact. Squirm3 does not contain any explicit notion of what constitutes an individual. Individual organisms simply arise as a result of interactions between atoms.

What are the advantages of working in the Squirm3 artificial environment, compared to working with more abstract ACs like Tierra [30] or even more abstract ACs like lambda calculus? Funes [11] commented about the difficulty of recognizing complexity in purely symbolic domains, such as Tierra, and recommended using complex, reality-based domains instead. The difficulty of appreciating the significance of evolved behaviors has also been noted in Geb [7].

In contrast, although Squirm3 greatly simplifies real-world chemistry, it nevertheless bears a close resemblance to it. Evolutionary innovations that occur in Squirm3 are likely to be recognized, because they can be expected to have biological parallels. For instance, membranes in Squirm3 are easily recognized as such [17, 18]. In contrast, it is not clear whether the concept of membrane even exists in Tierra. Organisms in Tierra and other abstract ACs may evolve survival mechanisms without real-world analogues. Even when identified, their significance might be difficult to interpret.

1.1 Evolution Types

This section defines and discusses some of the different kinds of evolution that we shall be referring to in this article.

Adaptive evolution adapts individual organisms to their habitat. Adapted organisms can be expected to persist for a long time, and even dominate the population. Adaptation can be either to the external environment, or to competing organisms. We also find it useful to differentiate between uphill and downhill adaptive evolution. *Uphill adaptive evolution* is adaptive evolution that increases the size of the genome, while *downhill adaptive evolution* is adaptive evolution that decreases the size of the genome. Since we desire to commence evolution with the minimum replicator (a replicator with the shortest possible genome), downhill adaptive evolution would clearly accomplish little except to preserve the dominance of the minimum replicator. It is uphill adaptive evolution that is the precondition for the more interesting types of evolution below. However, even uphill adaptive evolution need not necessarily result in organisms that are more complex.

Complexity evolution is adaptive evolution that results in organisms becoming more complex. Unfortunately, the concept of complexity is difficult to define [9]. While it is tempting to focus only on the complexity of individual organisms, the importance of the complexity of the biosystem as a whole should not be overlooked [21].

Creative evolution, as defined by Taylor [34], means that individual organisms can interact with the environment and other organisms with few restrictions. Moreover, creative evolution also signifies the evolution of new mechanisms for sensing and interacting with the environment. Taylor suggests the evolution of vision, photosynthesis, flying, and multicellular organisms as examples. The evolution of a membrane, we surmise, could be another example. Creative evolution may be seen as an extreme case of complexity evolution. Taylor's notion of creative evolution bears much similarity to the major transitions in biological evolution identified by Maynard-Smith and Szathmáry [24]. For a somewhat philosophical discussion of creativity in evolution, see Bentley [3].

Open-ended evolution, also defined by Taylor [34], occurs when new, adaptively successful individuals keep appearing in the population indefinitely. It is adaptive evolution that never ceases. Bedau et al. [2] developed a statistical test (the ALife test) for open-ended evolution in terms of unbounded evolutionary activity metrics. Maley [23] raised problems with the ALife test, as did Channon [5, 6], who refined the test and showed that the artificial life system Geb [7] can pass it. Stout and Spector [33] validated the ALife test (with Channon's improvements). While the ALife test is a useful tool for evaluating and comparing ALife systems, one has to be aware that not even the meaning of "open-ended evolution" is universally agreed upon. For example, Ruiz-Mirazo et al. [31] list a lack of a predetermined upper bound of organizational complexity as a core feature of open-ended evolution. Since it does not appear essential for adaptively successful individuals in Taylor's definition to be more complex, the two definitions seem to differ. In this article, we abide by Taylor's definition for open-ended evolution.

With these definitions, we can say that creative evolution may occur in the absence of open-ended evolution. Fundamental novelty may arise initially, such as a new sense or the appearance of a membrane, but it then ceases.

Conversely, while open-ended evolution will continually produce individuals that differ from their ancestors, it need not be particularly creative, or even result in an increase in complexity. The evolution that occurs may just be "more of the same." For example, although Geb has been shown to pass Channon's test for open-ended evolution, the "few restrictions" requirement for creative evolution would seem to limit Geb's potential for that kind of evolution, since Geb organisms are limited to a small number of predefined actions.

1.2 Review of Squirm3 Artificial Chemistries

The Squirm3 artificial environment was introduced to the ALife community in 2002 [14]. It being an LMS, the reaction vessel in Squirm3 is a 2D grid in which each square can either be occupied by a single atom or be empty. The movement of the atoms in the grid amounts to Brownian motion.

Atoms are labeled by atom type, which remains constant, and atom state, which may change. Atoms can also form bonds with each other, thereby creating molecules. The reactions are described by reaction rules. For example, the reaction $e_0 + e_1 \rightarrow e_5 e_2$ specifies that if an atom of type e in state 0 comes into contact with an atom of type e in state 1, and the two atoms do *not* already have a bond between them, then they will form a bond, and the first e atom will be in state 5 and the second in state 2.

The AC described in [14] consists of eight reactions, all of them between two atoms. To easily refer to this AC, we name it H-8. Let us also define the following: Atoms without any bonds to other atoms are *free atoms*. Free atoms in state 0 are *food atoms*, since they are the raw material out of which replicators make copies.

Since the space in the 2D grid is finite, and since Squirm3 replicators have not reached the level of sophistication necessary to attack each other for food, evolution in the grid will halt as soon as all the food atoms are used up. To prevent this, a reaper process is needed to periodically prune the replicators in the grid, and to insert a fresh supply of food atoms. Hutton therefore introduced a *half-flood*, which periodically empties half the grid and fills the empty area with a random mixture of food atoms. This flood type, with minor modification, is also used in all his subsequent articles.

With H-8, it was shown that:

- a. When a suitable molecule is inserted into a *soup* of food atoms, it spontaneously replicates.
- b. Mutations occur without the need for any specific mutation reaction, due to the interaction between different molecules.
- c. Seeding the grid with longer molecules results in the evolution of shorter molecules. This is a direct consequence of the half-flood, because shorter molecules in Squirm3 tend to replicate faster. The same phenomenon was also observed in vitro [26], where downhill adaptive evolution of RNA occurs in an environment favoring the fastest-replicating RNA.
- d. When a small random factor is applied to a soup of food atoms (analogous to cosmic rays), replicators arise spontaneously. This is largely thanks to the fact that the minimum replicator ef consists of just two atoms.

In [15] H-8 was extended to 15 reactions (we hence refer to this AC as H-15). H-15 introduces reactions between three atoms. This prevents molecules from interfering with each other's replication reactions, avoiding spontaneous mutations as observed in H-8. While spontaneous mutations might seem like a desirable trait for an AC, they deprive the experimenter of the ability to control the mutation rate.

Rather than setting all free atoms to state 0 (making them food atoms), some of the free atoms were set to states that made them useless as raw materials for molecular replication. To make them palatable, reactions were introduced that allow atoms of types a and b that were part of a replicator to act as enzymes, converting free atoms in non-palatable states into food atoms. For these extremely simple enzymes, even their position in the replicator is irrelevant.

In one experiment, the grid was divided into three sections, with food atoms in one section, and two different varieties of non-food free atoms in the remaining two. By seeding the grid with the minimum replicator, creating a periodic half-flood, allowing some mixing between the sections, and adding a mutation reaction, evolutionary growth from ef to $eabf$ and $ebaf$ was observed. The catalytic utility of these simple enzymes counterbalanced the pressure, due to the half-floods, toward shorter molecules. Thus, it was shown that uphill adaptive evolution is possible in a Squirm3 chemistry. H-15 was also an early example of how attaching enzymes to replicators can serve as an alternative to membranes, but this approach was not followed up in subsequent articles.

Hutton's next Squirm3-based AC [16] had 13 reactions (so we call it H-13). While the reactions for replication in H-13 and H-15 are essentially identical, H-13 introduces quasi-universal enzymes capable of catalyzing any two-atom reaction not involving an enzyme. An enzyme in H-13 is a free atom of type d with its state specifying the reaction it catalyzes. This enzyme state is programmed by

the sequence of atoms of type a , b , and c in the replicator, which is interpreted in base 3. H-13 hence weakens the analogy between atoms in the AC and physical atoms, as individual physical atoms cannot be enzymes. In contrast, the analogy between the a -, b -, and c -type atoms in the replicator and DNA nucleobases is strong, and therefore one may refer to these atoms as *bases*.

For experimental purposes, one of the reactions vital for replication in H-13 was disabled. For replication to occur, the presence of an enzyme that catalyzes the missing reaction was therefore essential. Starting with a population of molecules that produce this enzyme (the *host*), mutations resulted in the emergence of parasites. Parasites are replicators that do not produce the essential enzyme and are thus only able to replicate by using the enzymes produced by the host. Eventually the parasites drove the host into extinction. With the essential enzyme no longer being produced, the parasites became extinct soon after. Thus, the experiment always resulted in global extinction.

Hutton concluded that the global extinctions in H-13 were due to the replicators being unable to benefit exclusively from the production of their enzyme. In [17] this problem was overcome by means of an AC consisting of 28 reactions (which we call H-28), which traps the useful enzyme inside a membrane. The replicators of this AC more closely resemble entire cells than strings of DNA. Not surprisingly, replication in H-28 is significantly more complex than in H-13 and mimics the process of biological cell division.

Again, one of the AC reactions was disabled, and the grid was seeded with a replicator that produces the enzyme essential for replication to occur. It was demonstrated that not only was the cell-like replicator capable of indefinite reproduction (as long as food atoms were available), but also the membrane successfully protected the replicators against parasites. However, no evolution was observed.

An alternative implementation of a membrane AC for Squirm3 was presented in [18]. In this implementation, unlike H-28, all the reactions needed for replication are two-atom reactions. (Unwanted mutations, as observed in H-8, are prevented by the membrane.) This has the advantage of making it possible, at least in principle, for the replication reactions to be completely described within the replicator itself. The change does, however, increase the reactions in the AC to 41. (We call it H-41.) It also increases the number of atom states needed for self-reproduction from 18 (in H-28) to 38.

Again, it was confirmed that membranes protect the replicators from parasites. Interestingly, a few neutral mutations were observed, lengthening the genome without altering the produced enzyme. Due to the workings of chance, occasionally such mutants usurped the original species. Finally, Hutton showed how longer replicator sequences are able to outcompete shorter sequences in situations where the longer sequence allows the cell to utilize a larger food supply.

While the cells' reproduction in H-41 (as well as H-28) is almost uncanny in its mimicry of biological cell division, it is a consequence of 41 (or 28) cleverly hand-coded reactions. It remains to be demonstrated how such complex behavior can arise spontaneously. However, it is a useful existence proof, showing that such complex "biological" behavior is possible in Squirm3.

Based on this review of Squirm3 literature, what are the weaknesses that our article will address?

None of the Squirm3 ACs, except H-15, have demonstrated any uphill adaptive evolution. And in H-15, the grid had to be constructed in such a way that each additional base provided the organism with an immediate adaptive advantage. This approach is unsuitable for evolving quasi-universal enzymes, because only the complete and correct sequence of many bases will provide an adaptive advantage. Yet in the presence of half-floods, any increase in the length of the genome, other things being the same, is penalized. Consequently, the sequences of bases required for adaptively useful quasi-universal enzymes always had to be hand-coded. In this article, we will show how they may evolve.

The other limitation of Squirm3 ACs is the global extinction problem reported in H-13. Although H-28 and H-41 offered membranes as the solution, this raises the problem of how the relatively complex cell-like organisms could have evolved. The alternative approach of attaching the enzyme to the replicator was shown to work in H-15, but only for replicators with two bases and extremely basic enzymes. We will show that this approach can be extended to work with a quasi-universal enzyme of ten bases (which in our AC is the maximum number of bases required to catalyze any two-atom reaction not involving enzymes).

By addressing the global-extinction problem and showing how quasi-universal enzymes can evolve, our article lays the groundwork for future work toward creative and/or open-ended evolution in the Squirm3 environment.

2 System Description

2.1 Overview

The replicators of our AC are chains of atoms with an atom of type e at one end, an atom of type f at the other, and bases in between (atoms of types a , b , c , and d). We adopt the quasi-universal enzyme first introduced in H-13, whose state is programmed by the sequence of bases. Enzymes are atoms of type d with a state greater than or equal to 10 (e.g., $d444112$). The enzyme state defines a specific two-atom reaction catalyzed by the enzyme (e.g., $e0 + e1 \rightarrow e5 e2$). The mapping between enzyme state and catalyzed reaction is arbitrary, allowing it to be adjusted to suit the needs of a particular experiment.

The 14 two- and three-atom reactions of our AC are defined in Figure 1. The AC includes a number of innovations. Most significantly, the enzyme does not float free after it has been programmed, as in H-13, but remains attached to the replicator that produces it.

The problem with half-floods is that they create selective pressure that favors the fastest-replicating molecules—usually the shortest. With respect to the reactor, our major innovation is the introduction of a new flood type that creates selective pressure toward longer molecules. We call it a *flow-flood*, because it simulates a fluidlike flow toward the right. Molecules and atoms are removed on the far right of the grid, while new food atoms are inserted on the far left. The only way in which molecules can survive the flood is by getting stuck on stationary *dirt* particles spaced throughout the grid. Since longer molecules are more likely to be trapped, this results in size selection in favor of longer molecules. A screenshot of a flow-flood in progress is presented in Figure 4, discussed in Section 2.4.

This concludes the brief overview of the system. The remainder of this section contains a detailed description of our AC and the flow-flood. Some readers may prefer to progress directly to experiment 1, returning to the minutiae later.

2.2 The Molecules

Following [14–18], our AC consists of atoms that possess both a type and a state. There are six atom types (a, b, c, d, e, f). Analogous to the atomic number of stable elements, the atom type cannot change. Atoms (except for atoms of type d) can be in one of eight possible states (0–7). Analogous to the excitation state of a physical atom, the state of an atom can change. In addition to the usual eight states, an atom of type d can also exist in state 10 and above, which indicates that it is an enzyme. Atoms can form bonds with each other, thereby forming molecules.

2.3 The Reactions

The reactions of our AC are shown in Figure 1. In order to facilitate comparison, the numbering of reactions R1 to R8 matches the corresponding reactions in [16].

Reactions R1 to R4 make a copy of the molecule, but bonds between atoms in the original and the copy remain. For that reason, reactions R5 to R8 are necessary to unzip the copy from the original, creating two independent molecules. They also set the atoms in the original and the copy to state 7, which triggers the enzyme-programming phase. An example of the replication process is shown in Figure 2.

In H-15, H-13, H-28, and H-41 the capturing of food atoms by the replicator proceeds one atom at a time in predetermined order. For example, until the molecule $e1 a1 c1 f1$ encounters an $e0$ atom to bond with its $e1$ atom, any $a0$, $c0$, and $f0$ atoms that float within range will be ignored. By the time the molecule finds an $e0$ atom, the $a0$, $c0$, and $f0$ food atoms encountered earlier might have moved far away or have been captured by another replicator.

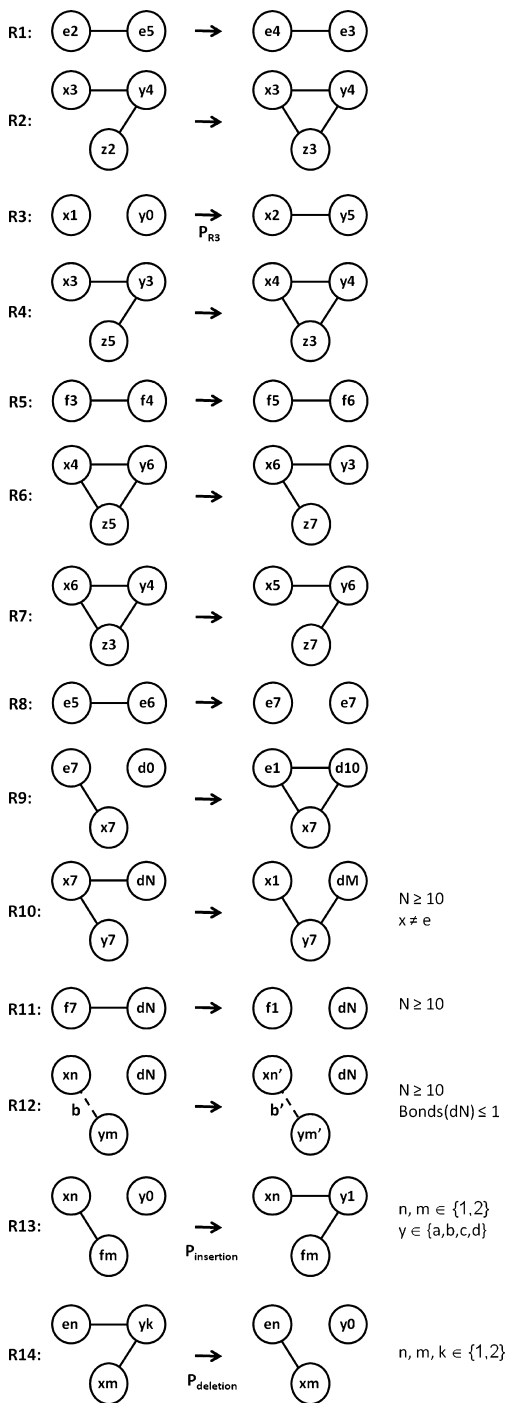


Figure 1. R1 to R8 are the reactions for self-replication. Reactions R9 to R11 program the enzyme. Reaction R12 carries out the enzyme action. Reactions R13, R14, and R3 respectively enable insertion, deletion, and point mutations. Unless otherwise specified, $x, y, z \in \{a, b, c, d\}$.

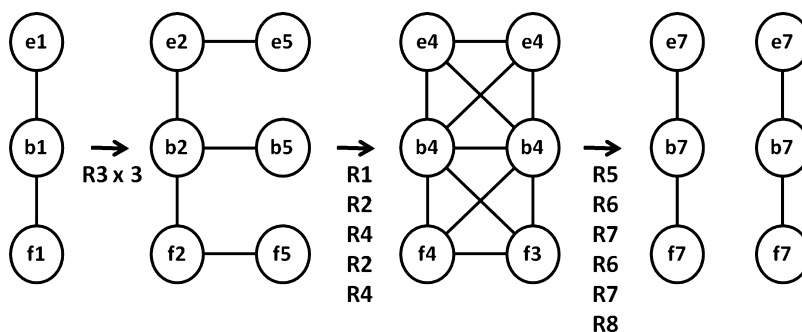


Figure 2. An example of a molecule replicating. $R3 \times 3$ indicates that $R3$ occurs three times. (The food atoms $e0$, $b0$, and $f0$ are not shown.) The order in which $R3$ occurs with respect to the atoms of type e , b , or f does not matter. On the other hand, the reactions shown under the two rightmost arrows must occur in the listed order.

This puts longer replicators at a considerable disadvantage compared to shorter ones. To reduce this disadvantage, our AC allows immediate bonding with any required food atoms during the replication phase of the molecule. Using the aforementioned example, if the $a1$ atom in the $e1 a1 e1 f1$ molecule were to come within reaction range of an $a0$ atom, there would be an immediate reaction binding the two atoms. This innovation, although in itself insufficient to make longer replicators competitive with shorter ones, nevertheless lessens the competitive disadvantage of longer replicators.

Mutations are handled by three different reactions, allowing their rates to be adjusted independently. An insertion mutation adds a base, a deletion mutation subtracts a base, and a point mutation swaps one base with another. Point mutations are implemented as part of the replication process via $R3$. P_{R3} is the probability that reaction $R3$ occurs per possible occurrence. When both atoms are of the same type (i.e., $x = y$), then $P_{R3} = 1$, that is, the reaction occurs with certainty. However, if $x \neq y$ and $x, y \in \{a, b, c, d\}$, then $P_{R3} = P_{\text{point mutation}}$ where $P_{\text{point mutation}}$ is a small probability that the wrong base gets copied. Handling point mutations during the copying process has the advantage that it makes the point-mutation rate independent of the total number of food atoms in the reactor.

Reactions $R9$ to $R11$ program the enzyme and set the states of the atoms in the molecule to 1, which triggers the next replication phase. In our AC, the enzyme that is programmed by the molecule does not float free as in H-13, nor is it surrounded by a membrane as in H-28 and H-41. Instead, the enzyme is attached to the molecule that has programmed it. One of the immediate advantages of this approach is that the reactions in our AC are of the same order of complexity as in H-13. The biggest single difference between our AC and H-13 is the manner in which the enzyme is programmed. Our process is shown in Figure 3. Enzyme programming begins via reaction $R9$ binding a food $d0$ atom to

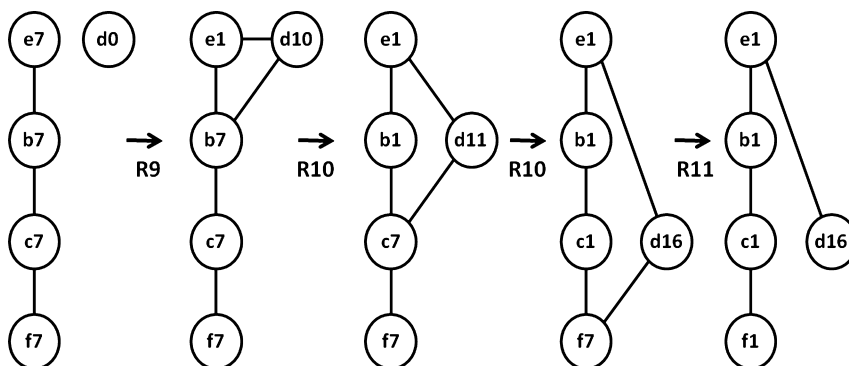


Figure 3. An example of a replicator with two bases programming an enzyme. A food atom of type d is attached to the replicator via $R9$ and becomes an enzyme to be programmed. A repeated application of $R10$ programs the enzyme, and $R11$ detaches the programmed enzyme from the f atom.

the replicator at its e end. The state of the d atom is set to 10, indicating that it is an enzyme. Reaction R10 enables the d atom to zip down the molecule, one base at a time (while remaining attached to the e atom), writing the base-4 number encoded by the sequence of bases to its state as it progresses. The condition $x \neq e$ in R10 is necessary to prevent an enzyme that has already been programmed once from zipping down the molecule again in subsequent enzyme programming phases. When the enzyme reaches the f atom, R11 will detach the enzyme from it. We are left with a fully programmed enzyme attached to the replicator's e atom.

An interesting consequence of R9–R11 is that a replicator will add an additional enzyme each time it goes through its replication–enzyme-programming phases. An old replicator will eventually be weighted down by a curly head of enzymes, impeding its movement and replication.

Our enzyme encoding scheme follows that of H-41. The value of the enzyme is given by the atom types of the bases, which are interpreted as a base-4 number with the highest digit attached to the e atom. The value of M in R10 is thus given by $M = 4(N - 10) + \text{val}(x) + 10$, where $\text{val}(a) = 0$, $\text{val}(b) = 1$, $\text{val}(c) = 2$, and $\text{val}(d) = 3$. For example, the molecule $ebcf$ produces the enzyme $d16$ (because $(bc)_4 + 10 = 12_4 + 10 = 16$).

Presently the reactions in our AC only allow a replicator to have a single enzyme (not counting multiple copies of the same enzyme). While this suffices for the experiments in this article, extending the AC to allow replicators to have multiple (different) enzymes will be necessary for future work.

Reaction R12 defines the enzyme action. The condition $\text{bonds}(dN) \leq 1$ ensures that the reaction only occurs when the enzyme has no or one bond with other atoms. This prevents enzymes that are in the middle of being programmed (with two bonds) from catalyzing reactions. The variables b and b' take on values of 0 and 1. A value of 0 indicates that no bond exists, while a value of 1 indicates that one does.

Our AC has 589,824 possible two-atom reactions not involving enzymes, including null reactions like $e4f3 \rightarrow e4f3$. To cover all possibilities, replicators with up to 10 bases are required (ef to $ecaddddddddf$).

Since the mapping between the enzyme state and the reaction it catalyzes is arbitrary, this article introduces an *enzyme-to-reaction mapping offset* to create a 1 : 1 circular mapping between the enzyme and the enzyme's action. With this offset, the values of n, m, b, n', m', b' in R12 are given by N , the state of the enzyme, via the relation

$$N - 10 = (2(2(6(6(8(8n + m) + n') + m') + x) + y) + b) + b' + \text{offset} \% 589824$$

where $\%$ is the remainder operation. For example, with $\text{offset} = 0$, $d10$ catalyzes $a0 + a0 \rightarrow a0 + a0$, $d11$ catalyzes $a0 + a0 \rightarrow a0 a0$, and $d589,833$ catalyzes $f7 f7 \rightarrow f7 f7$. With an enzyme-to-reaction mapping offset of 1, $d10$ catalyzes $a0 + a0 \rightarrow a0 a0$, and $d589,833$ catalyzes $a0 + a0 \rightarrow a0 + a0$. The offset makes it easy to adjust the catalytic action of an enzyme to whatever suits a particular experiment.

Finally, reactions R13 and R14 define insertion and deletion mutations, respectively. R13 inserts new atoms at the f end of the replicator with probability $P_{\text{insertion}}$ per possible occurrence. Likewise, R14 deletes atoms from the e end of the replicator with probability P_{deletion} per possible occurrence. Restricting insertions and deletions to the e and f ends of the replicator has the advantage that the likelihood of an insertion or deletion is independent of the number of bases in a replicator. This represents an improvement over H-15, H-13, H-28, and H-41, where insertions and deletions can occur anywhere along the genome; this disadvantages species with longer genomes, as a high mutation rate can cause a species to go extinct, due to individuals of that species mutating into different species.

2.4 The Reactor

As is typical in Squirm3, our reactor consists of a 2D grid of squares. Each time step consists of a reaction phase followed by a movement phase. Before every reaction phase, the order in which the

atoms are processed is shuffled randomly to avoid the possibility of ordering artifacts affecting the results.

During the reaction phase, only atoms that are in each other's Moore neighborhood are able to react. Thus, in order for two- or three-atom reactions to occur, the participant atoms have to occupy a 2×2 block of adjacent squares.

During the movement phase, each atom randomly moves into one of the eight squares of its Moore neighborhood, subject to the following constraints: (a) the atom may not move out of the 5×5 neighborhood minus corners of all atoms to which it is bonded; (b) an atom cannot move into a square already occupied by another atom. If movement without breaking either constraint is impossible, then the atom forfeits its move.

A flow-flood consists of a general flow of atoms from left to right. Atoms reaching the far right of the grid are removed and replaced by an equal number of food atoms on the far left of the grid. Given a large number of movement steps, this would obviously result in global extinction. To prevent this, we place non-reactive stationary dirt atoms throughout the grid, whose purpose it is to allow molecules to get stuck during the flow-flood. Significantly, longer molecules, such as replicators with many bases, are more likely to get stuck on these dirt atoms than shorter molecules. Once molecules are caught on the dirt, they tend to remain caught, so the exact number of flow-flood movement steps does not matter, as long as there are enough of them to eliminate the molecules that fail to be stuck. A screenshot of the grid during a flow-flood is shown in Figure 4.




Figure 4. The grid during a flow-flood. The large squares represent dirt particles, atoms are circles (filled circles if they have at least one bond, hollow circles if they have no bonds), and the lines represent bonds between atoms. The flow is from left to right. Long molecules get stuck on dirt like seaweed. Three molecules at the far right were not stopped by the dirt and will shortly disappear from the grid. After the flow-flood is over, the left and right grid walls are inserted back.

As with the greatly abstracted laws of chemistry in our AC, we do not attempt to accurately model the laws of fluid dynamics, but instead rely on simple rules to simulate the essential behavior. A movement step during a flow-flood is just like an ordinary movement step, with one extra constraint: Movement into one of the three cells in the Moore neighborhood to the left of the atom being moved is only allowed if there exists another atom, directly to the right of the atom being moved, that has not moved during the previous flow-flood movement step. This includes dirt atoms, which of course never move. This simple rule works well at keeping long molecules stuck, while at the same time shaking loose entangled bundles of shorter molecules.

The only other rule needed for flow-floods concerns the enzymes. A molecule that replicates multiple times may eventually have so many enzymes attached that it will be unable to move. A half-flood can remove them easily, but not so a flow-flood. We therefore introduce the rule that enzymes detach from their replicators during a flow-flood. Being free atoms, the enzymes will quickly drain away. For a physical analogy, one might say that the bond between an enzyme and its replicator is a weak bond, easily broken during the turbulence of a flow-flood.

3 Experiment I

3.1 Overview

This section presents three closely related experiments that investigate whether long extracellular replicators can prosper in an environment of half-floods, despite the emergence of parasites.

The key replicator in this series of experiments is *ebcdabcdabcf*, a molecule that produces enzyme *d444112*, which is set to catalyze $e0 + e1 \rightarrow e5 e2$. This reaction, because it is essential for replication to occur, is built into our AC under R3. One of the advantages of working with an AC is that we can easily turn off R3 for atoms of type *e*. Doing so, the only way for $e0 + e1 \rightarrow e5 e2$ to occur is via *d444112*.

Subsequently, we refer to *ebcdabcdabcf* as the *host species*, or just *host* for short. Replicators that do not produce enzyme *d444112* are parasites, because they are only able to replicate by freeloading off the host's enzyme. If the host goes extinct, the supply of *d444112* runs out, resulting in the extinction of all replicators.

Global extinction is of course exactly what happens in [16]. Experiment 1a is a control to replicate that result. The grid is seeded with the host, and mutations are set to occur. Although our reactions differ from H-13, the selective pressure toward shorter molecules is the same, due to the half-floods. Because our AC attaches the enzyme to the replicator, for the purpose of experiment 1a we add a temporary reaction to break the enzyme-replicator bond.

Since we expect experiment 1a to end in global extinction, experiment 1b checks whether attaching the enzyme to the replicator avoids that outcome. For this, and all subsequent experiments, the temporary reaction that breaks the enzyme-replicator bond is removed from the AC.

The purpose of experiment 1c is to increase our confidence in the result of experiment 1b. In addition to having the host compete against parasites arising from mutations, we also seed the grid with *ef*, the minimum replicator.

3.2 Experimental Settings


This subsection contains sufficient detail to enable the experiment to be replicated. Readers not interested in the minutiae are advised to skip to Section 3.3, referring back to this section later.

Following [16–18], we disable one of the reactions essential for replication to occur. Our choice is to disable R3, but only for atoms of type *e*, that is, $P_{R3} = 0$ when $x = y = e$. The reaction $e0 + e1 \rightarrow e5 e2$, and hence replication itself, can henceforth only proceed if there exists at least one enzyme that catalyzes $e0 + e1 \rightarrow e5 e2$.

We set the enzyme to reaction-mapping offset to 161,099, which results in $e0 + e1 \rightarrow e5 e2$ being catalyzed by *d444112*, which is the enzyme produced by *ebcdabcdabcf*. The replicator in this experiment thus has 10 bases, which is the maximum number of bases needed to catalyze any two-atom reaction not involving enzymes.

In all three experiments, a grid of size 120×120 is seeded with 15 copies of *ebcdabcedabcf*, randomly distributed, with the states of the atoms initialized to 7 so that the molecules start in their enzyme-producing phase. In addition to the initial replicators, the grid is filled with 1,200 food atoms. For every 40,000 time steps, a half-flood occurs, alternating between the top and bottom halves of the grid. The probabilities for deletion, insertion, and point mutation per possible occurrence are set to 0.00003, 0.0001, and 0.00005, respectively.

As mentioned in the preceding section, experiment 1a requires a temporary reaction (used in this experiment only) to dislodge the enzyme from the *e* atom:



The experimental settings for experiment 1b are largely identical to those of experiment 1a, with the crucial difference that enzymes remain attached to the replicators that produce them, that is, reaction R15 is removed from the AC. Also, because the deletion rate tends to be lower when the enzymes are attached, we increase the probability for deletion per possible occurrence to 0.0003. The probabilities for insertions and point mutations remain unchanged at 0.0001 and 0.00005.

In experiment 1c, we seed the grid with 15 instances of *ef* as well as 15 instances of *ebcdabcedabcf*, randomly scattering the 30 molecules in the grid. All other settings remain identical to experiment 1b. We have to ensure that the enzyme produced by *ef* does not influence the result of this experiment. With an enzyme-to-reaction mapping offset of 161,099, the enzyme *d10* programmed by *ef* catalyzes $e2 \ c1 \rightarrow e3 \ c6$. Since *ef* does not contain atoms of type *c*, *d10* neither helps nor hinders the survival chances of *ef*. Likewise, since in the molecule *ebcdabcedabcf* the atom bonded to *e* is of type *b*, not *c*, the enzyme *d10* does not directly affect the survival chances of the host. In short, the catalytic action of *d10* should not affect the outcome of this experiment for the two species that we are matching up, although it might affect other species that arise from mutations.

The replication counts reported on in the following section were derived by counting the occurrences of reaction R8—the final step of a molecule’s self-replication.

3.3 Results

Figure 5 shows the replications per flood interval (smoothed) against time (in units of floods) for a typical run of experiment 1a. Smoothing was done to improve the clarity of the graph, and involved taking the average over current, previous, and subsequent flood intervals. As expected, parasites drove *ebcdabcedabcf* (the host) into extinction, an outcome observed in every run of the experiment. Only the identity of the parasite(s) causing the extinction varies; they need not even be shorter than the host. Likewise, in every run that has been attempted, the victory for the parasites is short-lived. As the supply of *d444112* diminishes due to repeated half-floods, a global extinction results, confirming the result of [16].

Curiously, in some runs of experiment 1a, very short parasites emerged much sooner than would be expected, via the repeated occurrences of deletion mutations. This was due to enzymes that catalyze irregular mutations. For example, *ebcdabcedaacf*, just a single point mutation away from the host, produces enzyme *d444108*, which catalyzes $e0 + d1 \rightarrow e5 \ d2$. With this enzyme, a replicator may use an atom of type *e* instead of the base *d* in the copy. For example, if the original replicator is *eaaaabf*, and $e0 + d1 \rightarrow e5 \ d2$ occurs, then the copy will become *eaaaebf*. When *eaaaebf* replicates in turn, the replication process may start from the second *e* in the molecule, making the replicated copy *ebf*. This shows that even in a simple system like this, quasi-universal enzymes can generate surprises.

A typical result for experiment 1b is shown in Figure 6. That run took place for over 18 million time steps without the host going extinct. Experiment 1b thus demonstrates that attaching the enzyme to the replicator can prevent global extinction. The host’s lowest point was during the interval between flood 166 and flood 167, when it replicated only nine times. (It appears to be more in the graph, because of the smoothing.) This event can be traced to the emergence of parasite *ebcdabcaf*, which produces enzyme *d6946*, which catalyzes $f2 \ c2 \rightarrow f1 \ c6$, a reaction that sabotages the replication

replications per flood interval

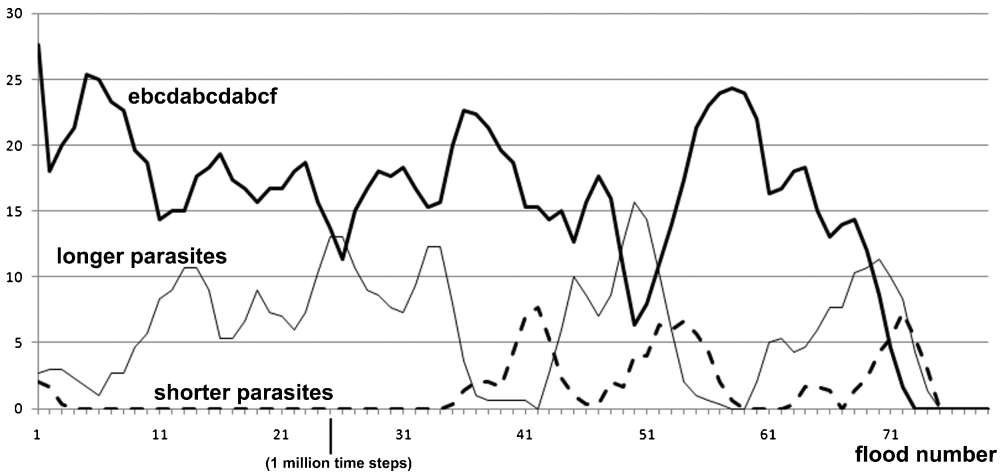


Figure 5. A typical result for experiment 1a, in which enzymes are unattached to replicators. The graph shows the replication count during each flood interval for the host (bold line), for parasites of the same length or longer than the host (thin line), and for parasites shorter than the host (broken line). All replicators produce enzymes, but only the host produces an enzyme critical for replication. As expected, the host becomes extinct (at flood 73), followed shortly by the extinction of the parasites.

of any molecule that has an atom of type *c* bound to an atom of type *f*, as does the host. However, unable to produce the crucial enzyme, this parasite did not persist for long.

Figure 7 shows a typical result for experiment 1c. The host prospered, despite having to compete against the minimum replicator. In fact, *ef* came close to extinction on several occasions, while the

replications per flood interval

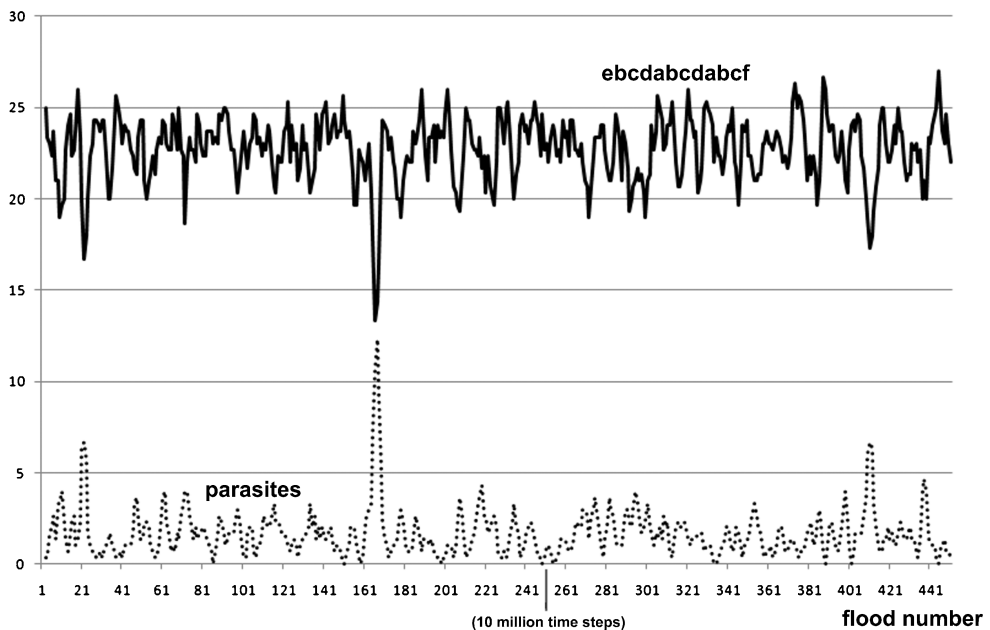


Figure 6. A typical result for experiment 1b, in which enzymes are attached to the replicators. The number of replications by the host (bold line) and the sum total of all parasite replications (dotted line) are plotted. Although mutations ensure a steady supply of parasites, they rarely pose a challenge to the dominance of the host, despite a nervous moment around flood 167.

replications per flood interval




Figure 7. A typical result for Experiment 1c, in which enzymes are attached to the replicators. Numbers of replications by the host (bold line), the minimum replicator (thin line), and all other parasites (broken line) are plotted. In this experiment, the host's adversary is the minimum replicator. The host was never in danger of extinction, while the minimum replicator came close on several occasions. Parasites emerging via mutations played an insignificant role in this experiment.

host never did. This visual impression is confirmed by the statistics. The mean numbers of replications per flood interval for the host and *ef* are similar (20.0 for the host versus 16.9 for *ef*), while the standard deviation for *ef* is twice that for the host (4.0 for the host versus 8.1 for *ef*).

3.4 Discussion

Experiments 1b and 1c demonstrate that replicators that maintain bonds with their enzymes are able to survive the emergence of parasites. Because of the enforced proximity of the enzyme to its replicator, the replicator is statistically more likely to benefit from the enzyme than is a parasite. This solves the global-extinction problem observed in [16] and replicated in experiment 1a.

Although H-28 and H-41 have also solved this problem, their solution relied on membranes. Enabling membranes in Squirm3 requires a relatively large number of AC reactions and a nontrivial number of atoms in the minimum replicator. This raises the question of how such a cell-like entity could have arisen in the first place in Squirm3—a question with an obvious biological parallel. The comparative simplicity of the reactions and minimum replicator in our AC makes it a more believable evolutionary starting point.

We have also shown that quasi-universal enzymes are capable of producing interesting behaviors, even in simple experiments such as this. In experiment 1a (and to a lesser extent experiment 1b) we observed the emergence of enzymes that “invent” their own mutation operations. And in experiment 1b we detected an enzyme that directly attacked a base sequence in the host.

This provides hope that, as the potential for complexity is increased, for example by introducing replicators capable of producing more than one kind of enzyme, further interesting behaviors will emerge—for instance, adaptation to competing individuals, something not previously observed in the Squirm3 environment.

One of the advantages of working in the Squirm3 environment is the ease with which one may draw biological parallels. It is therefore only natural to speculate on what biological counterparts the replicators of our AC might have. The obvious candidates are extracellular ribozymes. Since the

strand of nucleotides in a ribozyme acts as both replicator and enzyme, the two functions are inseparable—just as in our replicator.

However, ribozymes produce enzymes via molecular folding, while Squirm3 enzymes are implemented as individual atoms. Consequently, the replicators of our AC are able to catalyze their own replication, while ribozymes either replicate or catalyze—but not both simultaneously. In Lincoln and Joyce’s experiment [22], for example, ribozymes had to catalyze each other’s synthesis. Based on the H-13 result, one would expect this separation of enzyme and replicator to lead to the emergence of parasites and (eventually) global extinction. The fact that this did not happen in [22] is probably due to the RNA being assembled out of coarse building blocks—four oligonucleotides. Provided that RNA constructed out of three or fewer of these oligonucleotides is not viable, this presents a formidable barrier to parasite formation.

The closest biological counterparts of the replicators in our AC might be polyribosomes, with constituent ribozymes that are able to take turns with their replicative and catalytic actions.

4 Experiment 2

4.1 Overview

In previous Squirm3 ACs, the sequences of bases necessary to produce adaptively useful quasi-universal enzymes were hand-coded. Ditto for experiment 1. The evolution of adaptively useful quasi-universal enzymes was never demonstrated in the Squirm3 environment.

This lack of uphill adaptive evolution can be traced to the half-floods, as this flood type provides strong selective pressure toward the minimum replicator. We now investigate whether flow-floods produce a different outcome.

Specifically, we want the enzyme that catalyzes $e0 + e1 \rightarrow e5 e2$ to evolve. In experiment 2 replication must be possible even in the absence of this enzyme, because initially it will not be present in the environment. Therefore, we enable $e0 + e1 \rightarrow e5 e2$ in our AC, but at a very low probability per possible occurrence. This has parallels in real-world chemistry. While catalysts increase reaction rates (often by a huge factor), the reaction may also occur, at much lower probability, without a catalyst.

We adjust our enzyme-reaction mapping so that $e0 + e1 \rightarrow e5 e2$ is catalyzed by *d238*, an enzyme produced by a molecule with only four bases: *edcbaf*. This makes it easier for evolution to find a solution within a reasonable running time of our program. The grid is seeded with multiple instances of the minimum replicator *ef*. Experiment 2a aims to discover whether flow-floods and mutations result in the evolution of *edcbaf*.

However, might half-floods not also produce *edcbaf*, given enough time? Experiments 2b and 2c are controls to examine this question. Both use half-floods instead of flow-floods. In experiment 2b, the grid is seeded with copies of *ef*, while in Experiment 2c it is seeded with copies of both *ef* and *edcbaf*. Except for the flood settings, all other parameters are kept identical to experiment 2a.

4.2 Experimental Settings

This section contains sufficient detail to enable the experiment to be replicated. Readers not interested in the minutiae at this point are again advised to skip to the results.

In experiment 2a, a grid of size 120×120 is seeded with 800 food atoms. The probabilities for deletion, insertion, and point mutation per possible occurrence are set to 0.0002, 0.0002, and 0.0005, respectively. A flow-flood is set to occur every 20,000 time steps.

Occasionally, due to chemical reactions catalyzed by the enzymes, dense clumps of particles may form that flow-floods are unable to remove. To take care of any such debris, a half-flood is set to occur very infrequently, every 35 flow-floods.

We place six *walls* of dirt, oriented vertically, in the grid. Each wall consists of 2×2 blocks of dirt atoms, with six empty squares between blocks, as shown in Figure 4. Apart from the food atoms, the grid is seeded with 80 instances of *ef*. The large number of seed molecules is required because short molecules are quite vulnerable to elimination by flow-floods. Alternatively, one may seed the grid with a single copy of *ef* and delay the occurrence of the first flow-flood until the replicator has time to be established.

The goal of this experiment is for an enzyme that catalyzes $e0 + e1 \rightarrow e5 e2$ to evolve. Setting the enzyme's reaction-mapping offset to 15,149 results in $e0 + e1 \rightarrow e5 e2$ being catalyzed by $d238$, which is an enzyme produced by a molecule with four bases: $edcbaf$. Because our run will start with zero copies of $d238$, it must be possible for replication to proceed without it. We therefore set the probability of $e0 + e1 \rightarrow e5 e2$ per possible occurrence to 0.003. This suffices for molecules to replicate, albeit slowly.

Incidentally, making the reactions of an AC probabilistic might be a way to help AC organisms to outgrow the built-in reactions of the AC. By giving each of the built-in reactions only a small probability of acting, it makes it easier for replicators to bypass them and find superior reactions (enabled by their enzymes).

For experiments 2b and 2c, we turn off the flow-flood and remove the dirt particles. A half-flood is instead set to occur every 50,000 steps. The flood interval is increased, because a half-flood requires more time for molecules to mix in the grid than a flow-flood. In experiment 2b the grid is seeded with 80 instances of ef , while in experiment 2c it is seeded with 10 instances of ef and 10 instances of $edcbaf$.

4.3 Results

Experiment 2a has always resulted in the eventual evolution and dominance of either $edcbaf$ or $eadcbaf$. Both replicators produce the adaptively advantageous $d238$ enzyme, since in our encoding scheme adding an a at the beginning of the genome does not change the enzyme produced. In keeping with our previous terminology, we shall refer to these two species as the hosts (shorthand for "host species"), and all other species as parasites (keeping in mind that parasites in experiment 2 are capable of replication without the hosts, although less efficiently.)

Figure 8a shows that, prior to flood 643, the proportion of short versus long parasites fluctuated wildly. For example, long parasites decreased from being the most prolific replicators at the 300th flood, to zero replications at the 380th flood, back to being the most prolific at the 530th flood. Clearly, the population of molecules at a single instant of time fails to provide an accurate picture of the long-term average size distribution of the molecules in the grid.

At flood 643, $eadcbaf$ evolved, followed shortly by $edcbaf$, which usurped $eadcbaf$ in its dominance of the population. Once dominant, the hosts remain entrenched in their position in every run of the experiment that has been attempted. However, the length of the run required for a host to arise can vary considerably.

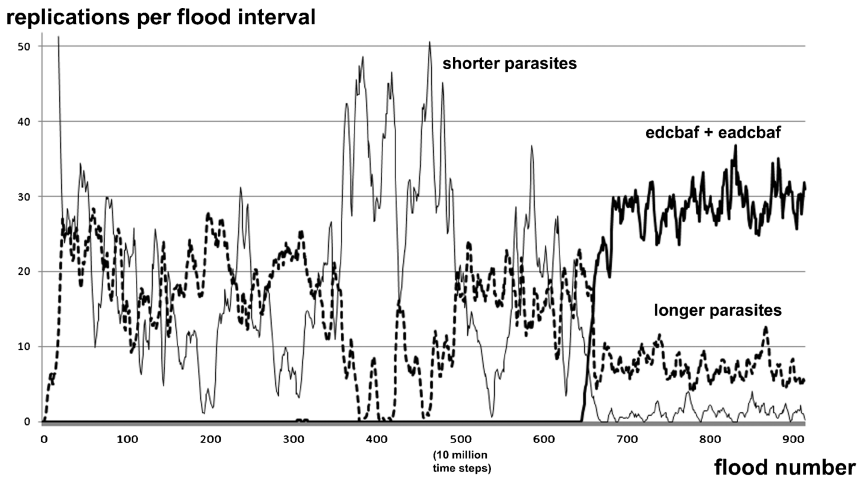
The evolution of the hosts is also evident in Figure 8b. In this scatterplot, successful species tend to form horizontal lines, as they replicate throughout a continuous range of flood intervals. Interestingly, $edcbaf$ first replicated around the 305th flood, but went extinct soon after. Small populations of an adaptively fit species may die out because of simple bad luck.

After the 643rd flood, the parasites that replicated most regularly are all a single mutation away from $edcbaf$, suggesting that they are not successful species in their own right, but are sustained via regular mutations from $edcbaf$. In three of these species, the d atom in $edcbaf$ is replaced by an atom of another type, while in the fourth the d atom is removed altogether. This is not a surprise; d atoms, which are also used as enzymes, tend to be scarcer than the other bases.


As for our controls, in experiment 2c, $edcbaf$ became dominant and eventually caused the extinction of ef . In experiment 2b, the dominant species remained ef , but longer molecules did occasionally arise. Figure 9 compares the total numbers of replications by number of bases in the replicator in experiments 2a and 2b.

4.4 Discussion

Experiment 2a demonstrates the advantage of flow-floods over half-floods. Unlike half-floods, flow-floods create selective pressure toward longer molecules. This provides a pool of replicators producing a variety of enzymes. When a sequence of bases that produces a useful enzyme is stumbled on by chance, it quickly comes to dominate the population. We have thus demonstrated uphill adaptive evolution of extracellular replicators.



(a)



(b)

Figure 8. (a) Typical run of experiment 2a. The number replications per flood interval (smoothed) is plotted against time (in units of floods). Smoothing involved taking the average replication count of the surrounding five floods. Replications by the hosts (bold line), shorter parasites (zero to three bases, thin line), and longer parasites (four to seven bases, broken line) are shown. The behavior of the system can be divided into two eras, with flood 643 as the dividing line. Before the 643rd flood, neutral evolution caused the space of replicators with up to seven bases to be explored. This changed after flood 643, when *eadcbaf* and *edcbaf* came to dominate the population. At the start of their first replication, there were over 80 replications by *ef* (cropped in the diagram). Subsequently, the great majority of shorter parasites possessed two or three bases, and longer parasites four or five bases. Parasites with eight bases and above were not observed in this run. (b) A scatterplot of species ID against time (in units of floods) for the same data set as in (a). It provides a view of the emergence of species over time. Species IDs are assigned sequentially to molecules in the order of their first replication. The size of the dot indicates the number of replications during the flood interval. A small dot denotes 1 to 4, a medium dot 5 to 9, and a large dot 10 or more replications. For those species that have been labeled, the initial *e* and final *f* atom have been omitted. Species 701, which is *eadcbaf*, was usurped by species 547, which is *edcbaf*.

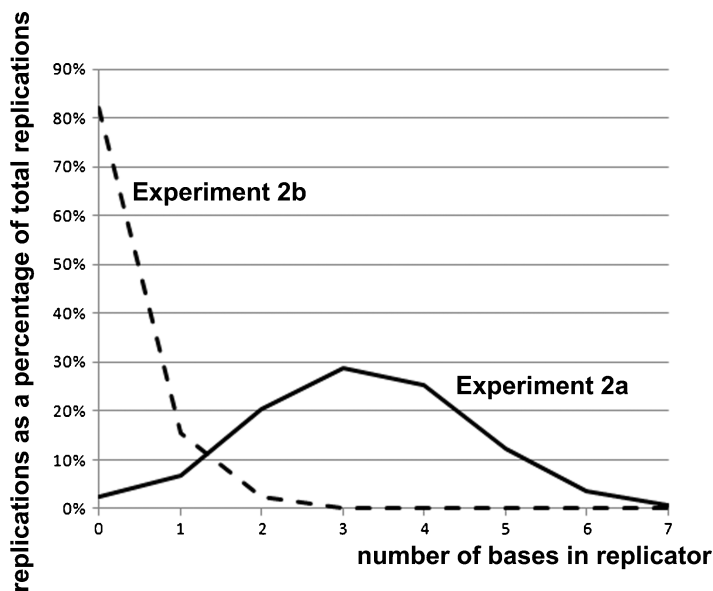


Figure 9. Percentage of total number of replications versus number of bases. Experiments 2a and 2b are compared. While experiment 2a peaked for molecules with three bases (with four bases close behind), experiment 2b has its peak at zero bases and declines exponentially. The count for experiment 2a was taken at flood 643, because we are interested in the distribution *before* the hosts begin their dominance.

From experiment 2c, we know that *if* the species *edcbaf* were to arise in the half-flood scenario, it too would come to dominate the population. Does that make the flow-flood superfluous? The broken line in Figure 9, easily recognized as exponential decay, provides the answer. While it is conceivable that a lucky sequence of insertion mutations could cause the emergence of *edcbaf* in an environment of half-floods, finding the right bases this way is three orders of magnitude less efficient than with flow-floods. (In experiment 2a, 25% of all replications are by molecules with four bases, compared to 0.02% in experiment 2b.) Although the decay constant of the exponential decay curve can be reduced by increasing the insertion mutation rate, this creates new problems, as it becomes increasingly likely that adaptively successful individuals will mutate before individuals of that species come to dominate the population.

While it is still barely possible for a replicator with four bases to arise under half-floods, the probability for an enzyme with 10 bases to evolve this way becomes infinitesimal. In contrast, flow-floods can be adjusted to select for even longer molecules (e.g., by increasing the dirt size.)

Clearly, flow-floods represent an important innovation for the Squirm3 artificial environment. However, does it have any broader significance? To answer this question, we first need to emphasize that the primary significance of the flow-flood lies not in the specific mechanism employed to implement it, but in *what* it does. It affords an environmental mechanism for providing selective pressure toward a longer genome that is unrelated to the information content of that genome. It is fascinating to speculate on what processes on early Earth might have done likewise.

Russell et al. [32] proposed that the RNA world began in inorganic compartments near hydrothermal vents. Baaske et al. [1] showed how the accumulation of biologically significant molecules in such compartments increases exponentially with the size of the molecule. Since it is speculated that early evolution occurred in these pockets of greatest concentration, this mechanism could have acted as a sieve, keeping longer strands of RNA trapped in these hatcheries for life while expelling shorter RNA, much like what occurs in the flow-flood.

A very different possibility arises from the chemoton, a model of a minimal protocell invented by Ganti [12]. Such a protocell is modeled in terms of its membrane, RNA templates, and metabolism. Ganti proposed that the RNA templates in such a protocell, rather than carrying useful genetic information, initially had the role of absorbing the waste products of the metabolism. Fernando and Di Paolo [10]

determined the existence of an optimum RNA template length for a specific chemoton model. Any longer or shorter than this optimum length, and the chemoton replicates more slowly. This brings to mind experiment 2a, in which the highest replicative success is achieved by replicators with two, three, and four bases.

Whether for these or any other reasons, once we have a mechanism that selects for a longer genome unrelated to the information content therein, it allows neutral evolution to take over. Experiment 2a thus provides an ALife demonstration of the importance of the neutral theory of evolution [20]. In terms of the stages defined by Wagner [36], the era before flood 643 represents a neutralist regime, and after flood 643 a selection regime.

5 Conclusions

We demonstrated uphill adaptive evolution in the Squirm3 environment to the level required for a replicator to program quasi-universal enzymes. Creative or open-ended evolution has yet to be demonstrated. However, since uphill adaptive evolution is a prerequisite for creative and open-ended evolution to occur, this article represents important progress.

We have also answered the question posed in [16], as to how extracellular RNA molecules could have evolved in complexity, despite having no choice but to share the fruits of their catalytic labor.

One of the advantages of Squirm3 is its close analogy with a real biology. This works in both directions, with Hutton's cells in H-28 and H-41 inspired by biological cells, and the results of our article allowing us to speculate on the early RNA world. Refer to the discussion sections on the two experiments for details.

Looking toward the future, one intriguing target for the demonstration of creative evolution is the evolution of a membrane, because Hutton already demonstrated that membranes are adaptively useful in the Squirm3 environment. We therefore present the following challenge to the ALife community: to develop a system in which, starting with a soup of free atoms and a simple "bootstrap" chemistry, a cell-like creature *similar to the one in H-41* evolves.

It must be emphasized that the challenge lies not so much in the membrane itself, as in the manner in which it arises. Merely producing a membrane within a lattice molecular system was, after all, achieved back in 1974 by Varela et al. [35]. Our challenge involves the evolution of a cell-like organism that contains a genome that programs enzymes that catalyze the reactions necessary for a membrane to form. This organism must furthermore be capable of self-replication.

At present, the practical problems of meeting this challenge—at least on personal computers—seems staggering. Hutton estimated the number of bases required for a cell to direct its entire replication sequence at 500 [18]. However, the situation might not be as dire as all that. Tierra was able to reduce Ray's initial organism consisting of 80 instructions down to 22 instructions [30]. Likewise, we can hope that evolution will discover a simpler membrane mechanism than the one employed in H-41.

The evolution of a membrane might also be facilitated by intermediate stages—simpler phenotypes that also afford their replicators survival advantages. In the context of the flow-flood, it is easy to imagine an example. If a certain atom type existed in abundance, for example a , then a tail made from a atoms would endow its possessor with greater likelihood of getting stuck on dirt during a flow-flood. This tail might later become the membrane.

Ultimately, by its very nature, one cannot predict where creative evolution might lead. While continuing this research, it will be vital not to overlook the possible emergence of other unexpected survival mechanisms.

6 Further Material

The source code for the Java-based software used in this article, called BTL, as well as the raw data from the experiments, is available at <http://www.cis.utas.edu.au/users/mwllucht/BTL.html>.

Acknowledgments

It gives me great pleasure to thank Tim Hutton for his valuable advice, and Robert Rowe for his help with the biochemical passages. Finally, I would like to thank my two anonymous reviewers, whose comprehensive comments have done much to strengthen this article.

References

1. Baaske, P., Weinert, F. M., Duhr, S., Lemke, K. H., Russell, M. J., & Braun, D. (2007). Extreme accumulation of nucleotides in simulated hydrothermal pore systems. *Proceedings of the National Academy of Sciences of the United States of America*, *104*(22), 9346–9351.
2. Bedau, M. A., Snyder, E., & Packard, N. H. (1998). A classification of long-term evolutionary dynamics. In C. Adami, R. Belew, H. Kitano, & C. Taylor (Eds.), *Proceedings of Artificial Life VI* (pp. 228–237). Cambridge, MA: MIT Press.
3. Bentley, P. J. (1999). Is evolution creative? In P. J. Bentley & D. W. Corne (Eds.), *Proceedings of the AISB'99 Symposium on Creative Evolutionary Systems* (pp. 28–34). Society for the Study of Artificial Intelligence and Simulation of Behaviour.
4. Cairns-Smith, A. G. (1985). *Seven clues to the origin of life*. Cambridge, UK: Cambridge University Press.
5. Channon, A. (2002). Improving and still passing the ALife test: Component-normalized activity statistics classify evolution in Geb as unbounded. In R. Standish, M. Bedau, & H. Abbass (Eds.), *Proceedings of Artificial Life VIII* (pp. 173–181). Cambridge, MA: MIT Press.
6. Channon, A. (2006). Unbounded evolutionary dynamics in a system of agents that actively process and transform their environment. *Genetic Programming and Evolvable Machines*, *7*(3), 253–281.
7. Channon, A. D., & Damper, R. I. (1998). Perpetuating evolutionary emergence. In R. Pfeifer, B. Blumberg, J. A. Meyer, & S. W. Wilson (Eds.), *From Animals to Animats 5: Proceedings of the Fifth International Conference on Simulation of Adaptive Behavior* (pp. 534–539). Cambridge, MA: MIT Press.
8. Dittrich, P., Ziegler, J., & Banzhaf, W. (2001). Artificial chemistries—A review. *Artificial Life*, *7*(3), 225–275.
9. Edmonds, B. (1999). What is complexity?—The philosophy of complexity per se with application to some examples in evolution. In F. Heylighen & D. Aerts (Eds.), *The Evolution of Complexity* (pp. 1–18). Dordrecht, The Netherlands: Kluwer.
10. Fernando, C., & Di Paolo, E. (2004). The chemoton: A model for the origin of long RNA templates. In J. Pollack, M. Bedau, P. Husbands, T. Ikegami, & R. Watson (Eds.), *Artificial Life IX: Proceedings of the Ninth International Conference on the Simulation and Synthesis of Life* (pp. 1–8). Cambridge, MA: MIT Press.
11. Funes, P. (2001). *Evolution of complexity in real-world domains*. Doctoral dissertation. Available at <http://citeserx.ist.psu.edu/viewdoc/download?doi=10.1.1.69.5055&rep=rep1&type=pdf> (accessed June 2011). Brandeis University, Waltham, MA.
12. Ganti, T. (1971). *The principle of life* (in Hungarian). Budapest: Gondolat.
13. Gilbert, W. (1986). The RNA world. *Nature*, *319*, 618.
14. Hutton, T. J. (2002). Evolvable self-replicating molecules in an artificial chemistry. *Artificial Life*, *8*(4), 341–356.
15. Hutton, T. J. (2003). Simulating evolution's first steps. In W. Banzhaf, T. Christaller, P. Dittrich, J. T. Kim, & J. Ziegler (Eds.), *Proceedings of the Seventh European Conference on Artificial Life* (pp. 51–58).
16. Hutton, T. J. (2003). Information-replicating molecules with programmable enzymes. In *Proceedings of the Sixth International Conference on Humans and Computers* (pp. 170–175). Aizu-Wakamatsu: University of Aizu.
17. Hutton, T. J. (2004). A functional self-reproducing cell in a two-dimensional artificial chemistry. In J. Pollack, M. Bedau, P. Husbands, T. Ikegami, & R. A. Watson (Eds.), *Proceedings of Artificial Life IX* (pp. 444–449).
18. Hutton, T. J. (2007). Evolvable self-reproducing cells in a two-dimensional artificial chemistry. *Artificial Life*, *13*(1), 11–30.
19. Johnston, W. K., Unrau, P. J., Lawrence, M. S., Glasner, M. E., & Bartel, D. P. (2001). RNA-catalyzed RNA polymerization: Accurate and general RNA-templated primer extension. *Science*, *292*(5520), 1319–1325.
20. Kimura, M. (1983). *The neutral theory of molecular evolution*. Cambridge, UK: Cambridge University Press.

21. Korb, K. B., & Dorin, A. (2011). Evolution unbound: Releasing the arrow of complexity. *Biology and Philosophy*, 26(3), 317–338.
22. Lincoln, T. A., & Joyce, G. F. (2009). Self-sustained replication of an RNA enzyme. *Science*, 323(5918), 1229–1232.
23. Maley, C. C. (1999). Four steps toward open-ended evolution. In W. Banzhaf, J. Daida, A. Eiben, M. Garzon, V. Honavar, M. Jakiela, & R. Smith (Eds.), *Proceedings of the Genetic and Evolutionary Computation Conference* (pp. 1336–1343).
24. Maynard-Smith, J., & Szathmáry, E. (1995). *The major transitions in evolution*. Oxford, UK: Oxford University Press.
25. Miller, S. L. (1953). Production of amino acids under possible primitive earth conditions. *Science*, 117(3046), 528–529.
26. Mills, D. R., Peterson, R. L., & Spiegelman, S. (1967). An extracellular Darwinian experiment with a self-duplicating nucleic acid molecule. *Proceedings of the National Academy of Sciences of the United States of America*, 58(1), 217–224.
27. Nelson, K. E., Levy, M., & Miller, S. L. (2000). Peptide nucleic acids rather than RNA may have been the first genetic molecule. *Proceedings of the National Academy of Sciences of the United States of America*, 97(8), 3868–3871.
28. Orgel, L. (2000). A simpler nucleic acid. *Science*, 290(5495), 1306–1307.
29. Peretó, J., Bada, J. L., & Lazcano, A. (2009). Charles Darwin and the origin of life. *Origins of Life and Evolution of Biospheres*, 39(5), 395–406.
30. Ray, T. S. (1991). Evolution and optimization of digital organisms. In K. R. Billingsley, E. Derohanes, & H. Brown (Eds.), *Scientific excellence in supercomputing: The IBM 1990 Contest prize papers* (pp. 489–531). Athens, GA: The Baldwin Press.
31. Ruiz-Mirazo, K., Umerez, J., & Moreno, A. (2008). Enabling conditions for “open-ended evolution.” *Biology and Philosophy*, 23(1), 67–85.
32. Russell, M. J., Daniel, R. M., Hall, A. J., & Sherringham, J. A. (1994). A hydrothermally precipitated catalytic iron sulphide membrane as a first step toward life. *Journal of Molecular Evolution*, 39(3), 231–243.
33. Stout, A., & Spector, L. (2005). Validation of evolutionary activity metrics for long-term evolutionary dynamics. In *Proceedings of the Conference on Genetic and Evolutionary Computation* (pp. 137–142). New York: ACM Press.
34. Taylor, T. (2001). Creativity in evolution: Individuals, interactions and environment. In P. J. Bentley & D. W. Corne (Eds.), *Creative evolutionary systems* (pp. 79–108). San Mateo, CA: Morgan Kaufmann.
35. Varela, F. J., Maturana, H. R., & Uribe, R. (1974). Autopoiesis: The organization of living systems, its characterization and a model. *BioSystems*, 5(4), 187–196.
36. Wagner, A. (2008). Neutralism and selectionism: A network-based reconciliation. *Nature Reviews Genetics*, 9(12), 965–974.

