Evolvable Self-Replicating Molecules in an Artificial Chemistry

Tim J. Hutton ,
Biomedical Informatics Unit,
Eastman Dental Institute, University College London,
256 Grays Inn Road, London WC1X 8LD, UK.
T.Hutton@eastman.ucl.ac.uk

Abstract

This paper gives details of Squirm3, a new artificial environment based on a simple physics and chemistry that supports self-replicating molecules somewhat similar to DNA. The self-replicators emerge spontaneously from a random soup given the right conditions. Interactions between the replicators can result in mutated versions that can out-perform their parents. We show how artificial chemistries such as this one can be implemented as a cellular automaton. We concur with [9] that artificial chemistries are a good medium in which to study early evolution.

Keywords: evolution, self-organization, emergence, molecular simulation, origin of life, complex systems.

1 Introduction

This work addresses the second of Bedau at el.'s [2] open challenges for artificial life research: 'Achieve the transition to life in an artificial chemistry in silico'. While a definition and quantitative test for life is still out of reach, a working definition in this context was suggested: 'a self-reproducing organizational form constructing itself in a simple environment and capable of evolution'.

One motivation for simulating evolution was described by McMullin [17]: that, unlike our typical experience with engineering where complexity necessarily decreases as machines produce other machines, evolution provides a mechanism whereby complexity can increase with successive generations. This *evolutionary growth of complexity* has never been conclusively demonstrated in a simulated system and is the ultimate goal of our work.

Maley [15] captured the direction of our research: 'The richness of biological life has never been replicated in our artificial models of evolution'. He gave four rather abstract requirements for systems to display *open-ended evolution*, where entities of increasing complexity and diversity are produced without limit.

Taylor [36, 37, 38] suggested some practical features necessary for evolution that is both open-ended and potentially *creative*, these include that the 'proto-DNA' should be a) an indefinite hereditary replicator, and b) fully embedded in a world with c) rich interactions.

McMullin [17] also discussed why one minimum requirement for evolution seems to be a self-replicator, suggesting that von Neumann saw it as the special case between degenerative production and increasing-complexity production, being able to produce something of the same level of complexity as itself. In von Neumann's design for a self-reproducing machine [41] (only recently implemented in full [24]) this was achieved by connecting a universal computer to a universal constructor in a cellular automaton (CA) environment. A tape of instructions guides the construction of a copy of the whole machine, including the tape itself. Certainly such a machine is capable in theory not only of replicating itself but also of creating machines of greater complexity that are still self-replicating, if the correct changes occurred to the contents of the tape.

This guaranteed evolutionary capacity is appealing and impressive but it *appears* that von Neumann's design for a self-replicator capable of the evolutionary growth of complexity is more complex than it need be - note that biological organisms contain neither a universal computer nor a universal constructor. Following this line of thought, much smaller CA replicators have been created, starting with Langton's famous loops in 1984: [13, 3, 28, 39, 23, 6, 7, 30, 31]. These automata are not in general capable of universal computation (with the notable exception of [23]) or universal construction but can self-replicate repeatedly, and in some cases are capable of limited evolution either through artificial selection [7] or natural selection caused by the competition for space [31]. The spontaneous appearance of evolvable self-replicators from a 'primordial soup' has been demonstrated [6].

Self-replication is possible in media other than a CA. Computer viruses are strings of machine code that attempt to copy themselves onto other computers by hiding in executable files that are distributed. Once a costly annoyance, computer viruses now have to contend with a plethora of anti-virus software preventing their successful replication. The study of replicating computer programs is also a respectable pursuit, with famous systems such as Tierra [26], Avida [1] and Amoeba [21, 22] descendents of A.K.Dewdney's Core War [8]. The computer programs in such systems are capable of universal computation and can evolve into more efficient forms. In this medium also, the spontaneous emergence of self-replicators has been demonstrated [21].

A third medium in which we can create self-replicators is the chemistry of our own Earth. Advances in molecular biology have made possible the synthesis of evolvable RNA-based replicators [40, 27, 42]. Combining such molecules with membrane-like lipid bilayers [14, 35] would let us see what the very first living organisms on our planet might have been like. This research is tremendously exciting because we know that our chemistry contains everything required for the staggering increase in complexity from RNA molecules to human beings, although we are unclear exactly what it is about our chemistry or the molecules involved that made this possible.

In this paper we argue that a fourth medium might be worth exploring to answer this question and the wider question of evolvability in general: artificial chemistries (ACs). We first discuss some advantages of ACs for studying evolution and evolvability, and then demonstrate one way that the minimal requirement for evolution (evolvable self-replicators) can be created.

The term artificial chemistry has been used as a loose metaphor to describe any medium in which artificial life may be simulated, including CA-based systems and computer programs. We believe this usage to be confusing since there is a specific type of simulation that better deserves the term - those that are similar to a real chemical system in that they involve molecules and collisions between them. Dittrich et al. [9] recently reviewed the state of the art in artificial chemistries and gave the following definition: an AC is a triple (S, R, A) where S is the set of all possible molecules, R is a set of collision rules and A is an algorithm describing the domain and how the rules are applied to the molecules inside (the physics). Neither the molecules nor the collisions in an AC need to be explicit and the AC can either be closely modelled on real chemistry or can abstract away from it. The domain can either be 'well-stirred' (any molecule can collide with any other with equal probability) or can have an imposed topology restricting collisions to be in some sense local.

A traditional CA can be seen as a special case of an AC, one where the molecules lie on every point of a regular grid and are not permitted to move, having 'collisions' only with their immediate neighbours. By comparison, an AC that permits movement will witness the collision of many different molecules in the course of the simulation - we believe that this allows for richer interactions than is possible in a CA without at least a) some mechanism to make the normal self-replicating operation of the cells robust to local variations and b) some mechanism for permitting local variation to be sampled at appropriate points.

Rich interactions are important for our goal because without them the replicators cannot grow in complexity. Consider that the reduction in entropy caused by information becoming stored in a genome is only possible because the information is useful to help the genome survive in its environment - the information is constantly kept up-to-date through many and complex interactions with competitors and with a varied habitat. Any mis-copying or deletion of important information in the genome causes the creature concerned to be less 'fit' in exploiting its surroundings and its genome to be copied less successfully. Thus in theory a genome can expand and be maintained over many generations in the face of mutations but only if all aspects of it are continually tested by the environment.

In the context of early evolution, Szathmáry stated [34]: "The potential now lies in modelling chemical organization and evolution *in abstracto*." Dittrich et al. [9] also argued that ACs are "the right stuff" for simulating prebiotic and biochemical evolution. However, until now the minimal requirement for evolution (the evolvable self-replicator) has not been shown in an AC.

Ono and Ikegami [18, 19, 20] demonstrated one kind of self-replicating structure in an AC - membranes that can grow and split. In the system described the membranes are of very limited evolvability since they lack in information storage capability. A similar model is the Lattice Molecular Automaton (LMA) [16, 25] which is a more direct simulation of real-world physics and chemistry, with forces and energies propagated between the lattice positions. Amphiphilic polymers - those with hydrophobic ('water-hating') monomers at one end and hydrophilic ('water-loving') monomers at the other - self-organise when immersed in water into recognisable higher-level structures (micelles) capable of growth and separation. In many

ways such models of simple proto-membrane-forming particles are complementary to Squirm3, we discuss this further later in the paper.

JohnnyVon [33] is a simulation of replicating strands of codons in a continuous 2D space. In some ways it is similar to Squirm3 in that it supports a simple method of template-based catalysis. Its physics is different, JohnnyVon has forces acting on the codons that depend on their state. One drawback is that it is much slower to run and this may restrict its usefulness for simulating evolution, a process which intrinsically requires many interacting entities of reasonable complexity.

In this paper we first give details of the Squirm3 system (section 2) and then in section 3 we examine some of its properties. In section 4 we show how the artificial chemistry simulation can be implemented as a cellular automaton with special update rules.

2 System Description

The system we shall describe developed from attempts to realise a representation for artificial creatures that was both computationally efficient and flexible in shape. Mass-spring models were rejected because of the potentially high computational cost of simulating the physics of multi-body systems. Traditional cellular automata were not used because they do not easily permit creatures to interact with each other, and it is felt that these rich interactions are essential for driving creative forces in evolution. From a purely systems design point of view artificial chemistries seem to offer the right mix of flexibility of form and a rich spectrum of possibilities regarding construction.

2.1 Components

The main component in the Squirm3 world is the *atom*. The word is not intended to imply an exact correspondence with real-world atoms but is instead used because they are the smallest unit in our simulation and indivisible. Each atom has a *type* $\in \{\mathrm{a}, \mathrm{b}, \mathrm{c}, \mathrm{d}, \mathrm{e}, \mathrm{f}\}$ and a *state* $\in \{0, 1, 2, \ldots\}$. The state of an atom is subject to change but its type is fixed.

A *bond* is a physical connection between two atoms. Bonds are formed and broken by *reactions*. Two or more atoms connected by bonds form a *molecule*.

2.2 Physics

In this section we specify the physics of Squirm3. A different set of physical rules with a similar chemistry may well give the same overall result but we give details here as an 'existence proof', to show that the behaviour described is experimentally replicable.

The physics of our virtual world is based on a two-dimensional grid as with cellular automata, each square of which is either empty or occupied by an atom. Each atom occupies exactly one square.

The atoms move around at random - at each simulation time step each atom either moves to an empty neighbouring square or stays still. A move to a square is valid if the atom remains within the neighbourhood of every atom it is bonded with. For this two-dimensional implementation we have used the 8 (Moore) neighbourhood.

An atom can react with any other atom inside its immediate von Neumann neighbourhood (4-neighbourhood). Each reaction is specified by the types and states of the atoms involved and whether they are bonded or not.

2.3 Chemistry

The first reaction (R1) is represented below:

$$\mathrm{e8} + \mathrm{e0} \rightarrow \mathrm{e4e3} \tag{1}$$

This notation is intended to be similar to that used in traditional chemistry (compare for example [9, 10]). The numbers, however, do not refer to the number of atoms of that type that are present (as a subscript would) but to the state of the atom. Reaction R1 indicates that when an atom of type e with state 8 physically encounters an atom of type e and state 0 then a bond is formed between them and their states change to 4 and 3 respectively.

All of the reactions are listed in Table 1. Some involve x and y, these are variables, standing for any of the other states $\mathrm{a} - \mathrm{f}$. Thus reaction R2 indicates that atoms of any type that are in states 4 and 1 and are connected will react to adopt states 2 and 5 respectively, and remain connected. Within each reaction the value of each variable remains fixed, thus reaction R3 describes how two atoms of the same type with states 5 and 0 form a bond.

R1:	$\mathrm{e8} + \mathrm{e0} \rightarrow \mathrm{e4e3}$
R2:	$x4y1 \rightarrow x2y5$
R3:	$x5 + x0 \rightarrow x7x6$
R4:	$x3 + y6 \rightarrow x2y3$
R5:	$x7y3 \rightarrow x4y3$
R6:	$\mathrm{f4f3} \rightarrow \mathrm{f8} + \mathrm{f8}$
R7:	$x2y8 \rightarrow x9y1$
R8:	$x9y9 \rightarrow x8 + y8$

Table 1: The reactions needed for self-replication.

Unlike traditional chemical formulae the molecules (which may be composed of many atoms) are not completely described in the reactions, which only concern the two atoms that are next to each other. For example, an e8 atom will react with an e0 atom whether either is bonded with other atoms or not. Without this, we would have to include all possible molecules in the reaction space, and this would limit the potential complexity of the molecules since the reaction set is necessarily finite.

These reactions could be written out in full without the variables, there would be 188 of them. With six atom types and ten states, there are $60 \times 60 \times 2 = 7200$ possible inputs (when two atoms are next to each other, bonded or unbonded), the other $7200 - 188 = 7012$ of them are null reactions (do not cause any change).

Figure 1 demonstrates how an example chain of four bonded atoms $\mathrm{e8} - \mathrm{a1} - \mathrm{b1} - \mathrm{f1}$ will replicate itself, drawing on a random soup of atoms in state 0. Reactions

R1 to R5 duplicate the chain by letting atoms of the correct type bond in the right place at the right time. Reactions R6 to R8 split the chain from its copy, leaving each in the original starting configuration.

The reactions are such that *any* string of a1's, b1's, c1's and d1's with an e8 at one end and an f1 at the other will replicate when immersed in a sufficiently large random soup of atoms in state 0. This is by design, since it becomes simple to show that there exists a path of small transitions between viable self-replicators that leads from the simplest to one of arbitrary complexity. The atoms a − d are a potential genetic code, information that is repeatedly copied from generation to generation.

We have made provision for 4 bases a − d but this is purely arbitrary. With n bases the addition of states e and f means that there are $200(n+2)^2$ possible reactions of which $5n^2 + 21n + 4$ are required for replication in the scheme set out above.

3 Observations

In this section we examine the results of three experiments; simulations of the physics and chemistry with different settings.

Experiment 1: Repeated self-replication and simple evolvability

For the first experiment we set up the world with the reactions given in Table 1, starting with one chain of e8-a1-b1-c1-f1 in a random soup of atoms in state 0 (Fig. 2(a)). The atoms are drawn as filled squares, with bonded atoms joined by lines. For clarity we use only a small world (20×20) and only 75 extra atoms.

As the simulation runs, the molecule replicates as expected by random interactions with the soup of atoms in state 0. Figure 2(b) shows the state of the world after 3544 iterations, 11 copies of the molecule are visible. Eight of the molecules are in the original form (e-a-b-c-f) and are unable to replicate further because there aren't any e0's left. The other three are in the process of replicating but are stuck halfway because there aren't any b0's left.

One side-effect of the chemical replication mechanism is that the concentration of free atoms in the soup changes. Of the free atoms in Fig. 2(b) there is a higher proportion of type d than before, since they are not being used for the replication. This kind of natural side-effect is one of the attractive properties of working with ACs - knock-on effects can create different ecological niches.

Another interesting property of ACs is that they are intrinsically 'dirty', in the sense that the reactions will happen whenever the right components come together, not necessarily when the designer of the AC had intended. Achieving control in such a situation can be difficult, there are many opportunities for the 'wrong' atoms to bond together. In Squirm3, this dirtiness manifested itself spectacularly and unexpectedly - with the same experimental settings as specified above we sometimes see not just copies of the original replicator but also different ones!

This occurs because there is a flaw in the reaction set concerning reaction R4 which connects any two atoms in states 3 and 6. In the normal replication process,

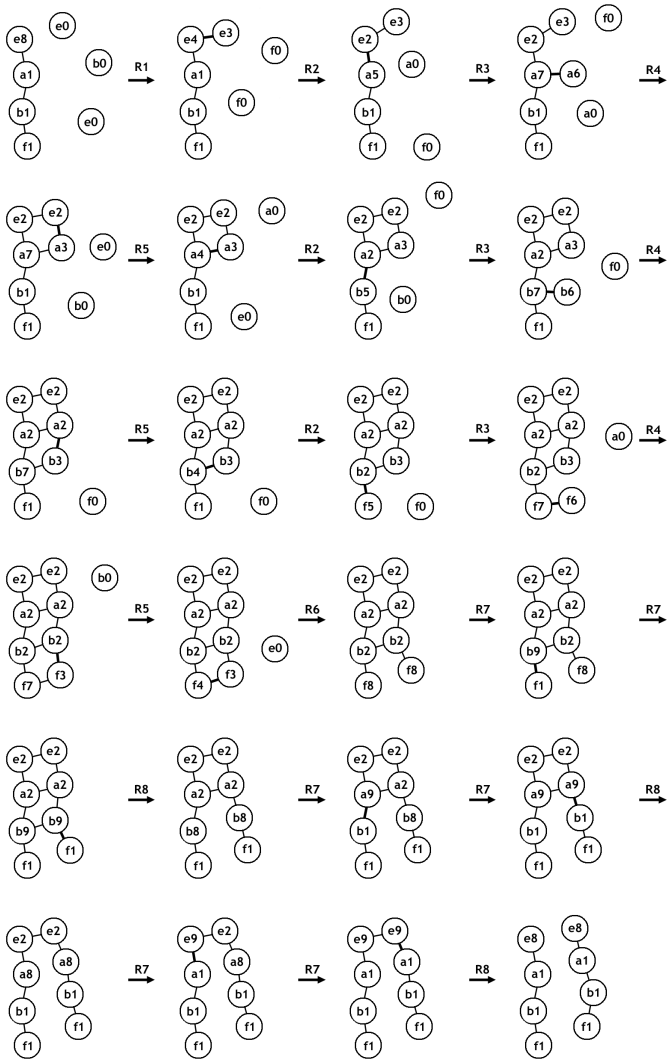

Figure 1: The replication sequence in Squirm3, illustrated on a short molecule e8 − a1 − b1 − f1. Any molecule composed of a string of a1's, b1's, c1's and d1's with an e8 at one end and an f1 at the other will replicate itself in a similar way when immersed in a random soup of atoms in state 0. Note that at each step the reaction shown is the only one that can apply.

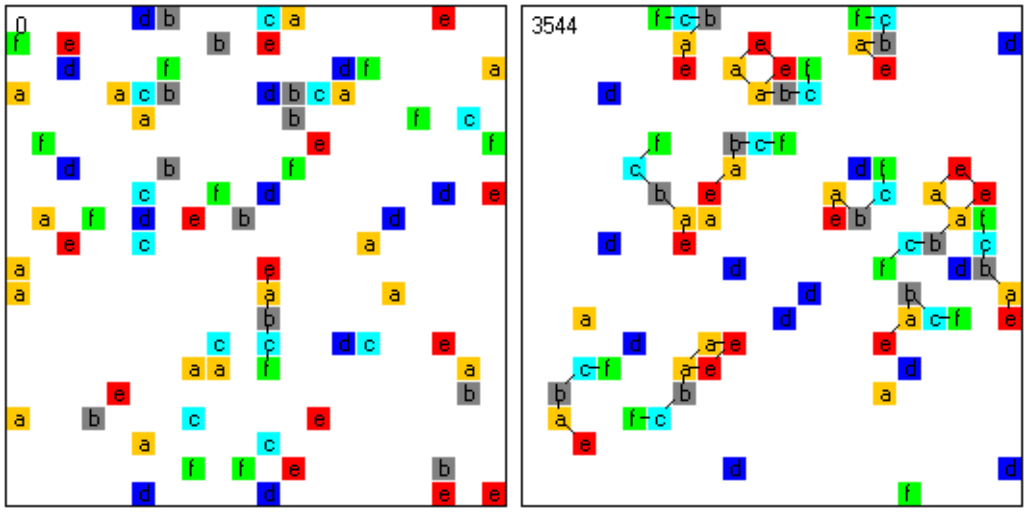

(a) (b)

Figure 2: Two snapshots of the running simulation. The starting point (a) contains one chain of atoms. After 3544 iterations (b) the chain has replicated several times, using the raw material surrounding it.

illustrated in Fig. 1, these atoms are part of the same molecule and their proximity means that they tend to quickly react with each other. After several rounds of replication, however, the possibility arises that the atoms in states 3 and 6 actually belong to different molecules. Often this results in a molecule that no longer has the ability to replicate but occasionally further reactions allow it to split into molecules that still have the necessary arrangement for self-replication but that are not exact copies of the originals.

In a different run of the same experiment we can see where this has happened. Fig. 3(a) shows where after 435 iterations two identical molecules have become attached with each other while replicating. Later in the run (Fig. 3(a)) they manage to separate again but the four resulting molecules are not identical. Two of them have the same structure (e-a-b-c-f) but the third has the form e-a-b-f and the fourth is e-c-a-b-c-f. There has been crossover of the genetic sequences, resulting in a base 'c' moving from one molecule to another. It must be noted that the increase in length caused by this is not yet an increase in complexity, since the bases a - d have no additional effects - any combination of them is equally fit.

Experiment 2: Evolution - adaptation to the environment

To replenish the environment and allow replication to continue we rerun the simulation with an additional effect. Every T time steps we remove all the atoms in one half of the area and refill it with random atoms in state 0. Figure 4 shows how this 'flood' removes any molecules that were in that half and inserts fresh 'raw

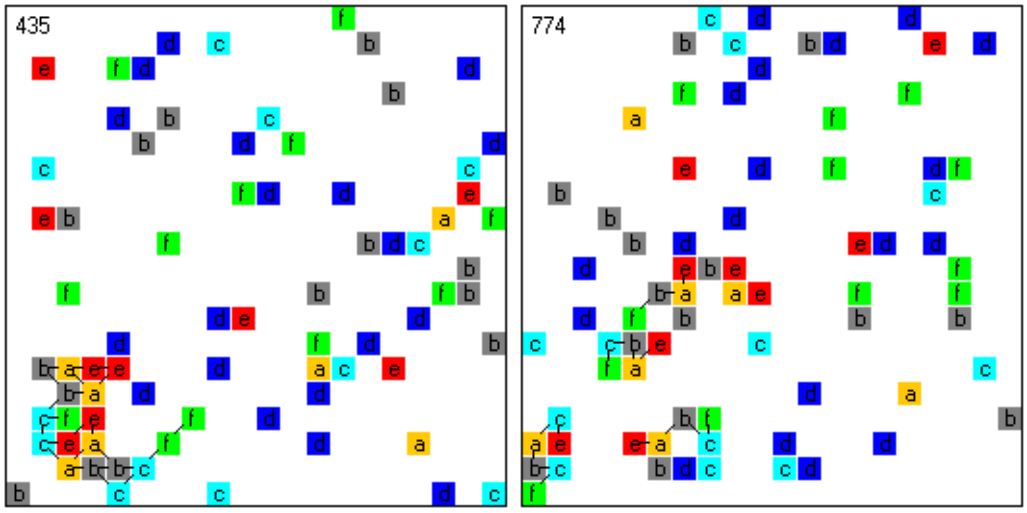

(a) (b)

Figure 3: In another run of experiment 1 we see two molecules (originally e-a-b-c-f) become tangled while replicating (a). When they manage to separate (b) the four resulting molecules are different.

material', that can be used by the remaining replicators. We alternate the flood between different halves of the area to clean out non-replicating molecules - ones that have got stuck against an edge or tangled with each other. In the 'warm pools' of a prebiotic world such partial mixing with larger pools or the sea would be commonplace.

In this experiment the selection pressure is in favour of replicators that replicate quickly - those molecules that are shorter tend to do better than longer ones because they are more likely to spread to cover the entire area before the next flood and thus survive.

We run the simulation with a flood period of $T = 2000$. Figure 5 shows 9000, 15800 and 34500 time steps in the second experiment, with a visible reduction in the length of the prevalent replicators. In the third image the world is dominated by replicators that are the shortest possible.

Figure 6 shows a plot of the replicator size as the simulation progresses. These results are from a 100×100 world with $T = 20000$, starting with a replicator that was seven atoms in length. The size was computed by counting the linked atoms after the reaction e9e9 $\rightarrow$ e8 + e8 occurred.

After optimal replicators become prevalent we see no further evolution. The replicators are not capable of evolving features that might be of advantage in this competitive situation, such as the ability to take other replicators apart to use as raw material, the chemistry does not yet have the capacity for this. Thus while this experiment does show that the replicators are capable of evolving it is in some aspects disappointing. It shows that Squirm3 is not complete, that we need more

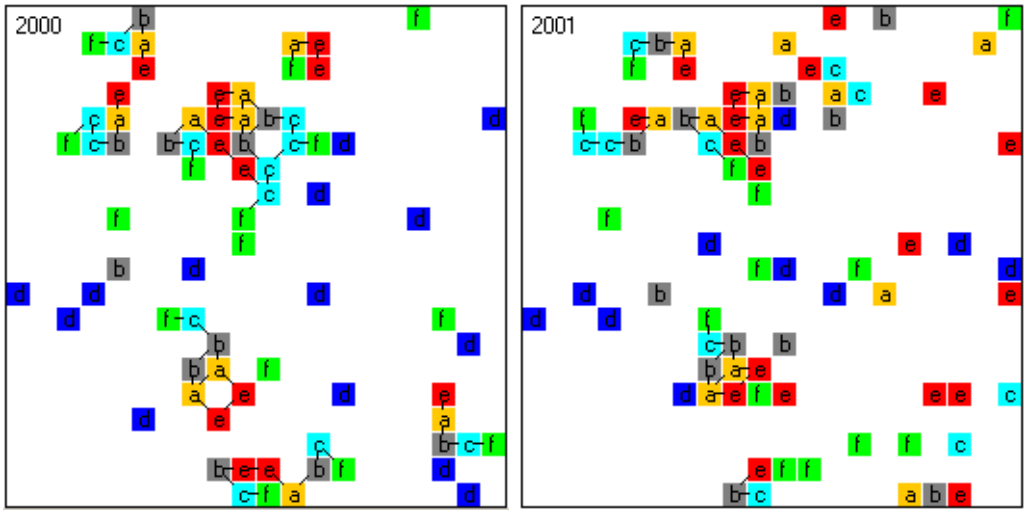

Figure 4: Before and after the flood. In the right image the right-hand half of the area has been replenished with atoms in state 0.

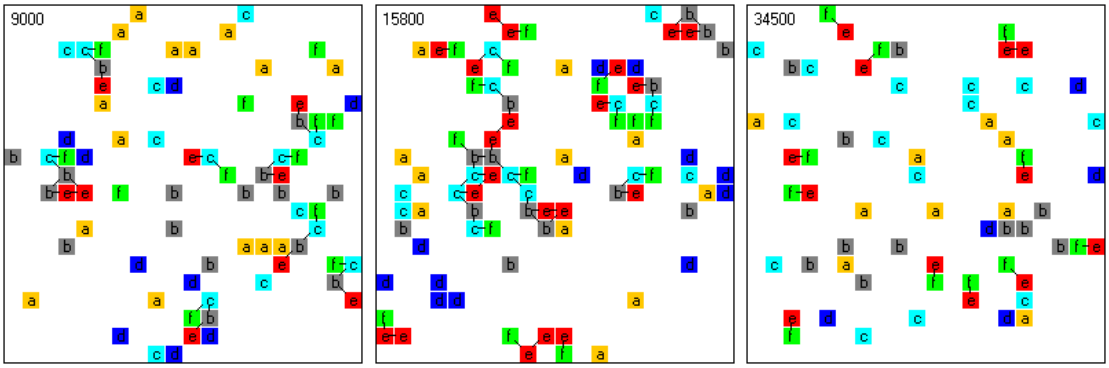

Figure 5: Three stages in experiment 2. Shorter replicators take over and come to dominate the world.

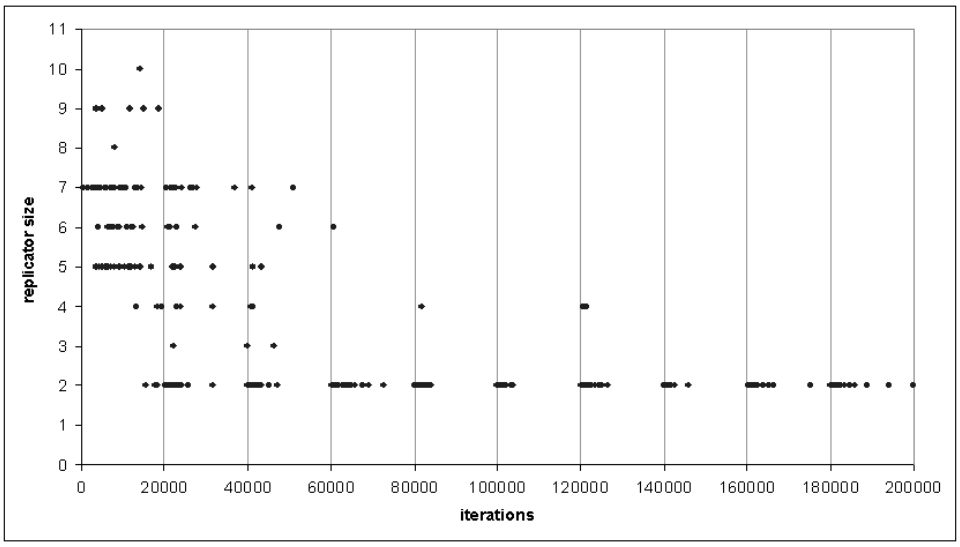

Figure 6: The replicator size decreases with successive generations because smaller replicators have an advantage.

reactions to permit the growth of complexity. We will return to the question of how this might be achieved later in the paper.

Experiment 3: Spontaneous appearance of self-replicators

For our third and final experiment we will demonstrate how the self-replicators can arise spontaneously from a soup of atoms. We initialise a 100×100 world with atoms of random type all with state 0. None of the reactions given in Table 1 are applicable, so without additional effects nothing would happen. We introduce, however, an occasional 'cosmic ray' with some low probability. The effect of a strike on an atom is to perturb its state, leaving its bonds and its type unchanged but its state as a random value. It turns out that this minimal intervention is sufficient to permit self-replicators to emerge spontaneously.

For this experiment we set $p_{\text{cosmic}} = 0.00001$ per atom per timestep, ie. at every timestep for a given atom there is a 0.00001 probability that its state will be randomised. After a few thousand iterations clumps of atoms become visible. These clumps can form, for example, when one atom in the soup has its state changed to 5. When it next encounters an atom of the same type in state 0 then reaction R3 applies, turning them into a 7-6 molecule. Any wandering atom in state 3 that happens to encounter the atom in state 6 will join to it by reaction R4, and so on. Such molecules may react occasionally but they do not replicate because they do not usually happen to have the correct structure.

Eventually, however, a molecule is created by these random events that has the structure needed to replicate itself: e8 $\{x1\}^*$ f1. After this point the composition of the world changes rapidly. The raw material that is used in the construction of the replicators vanishes rapidly as they multiply. We include a flood as in experiment 2 with period $T = 10000$ to clear out any persistent clumps of inactive atoms.

Figure 7 shows how the reaction rate increases suddenly when the first replicators appear, reflecting their capacity to act as catalysts.

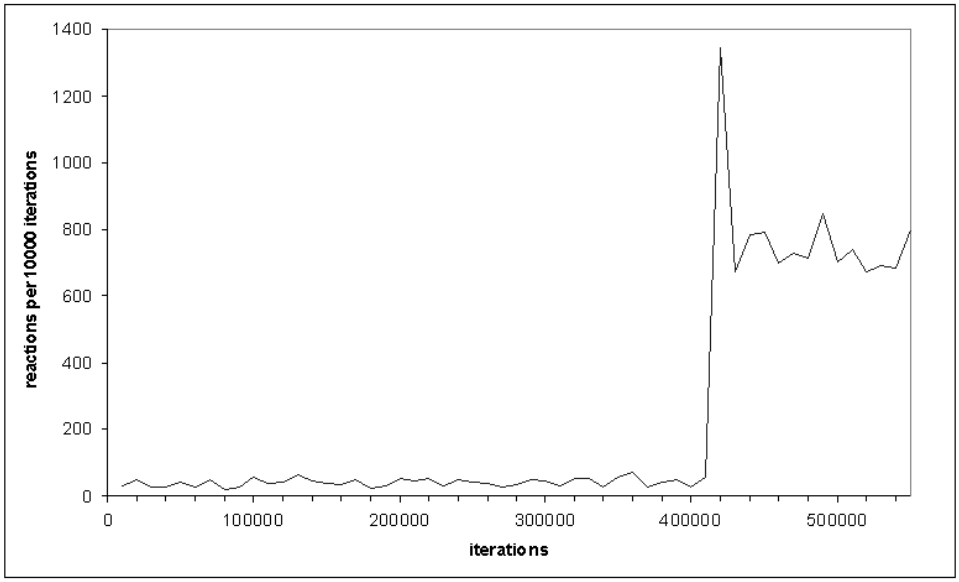

Figure 7: The reaction rate over time in experiment 3. When the first self-replicators appear (at around 400000 iterations) the average reaction rate jumps to a higher level.

4 Implementation in a CA

Cellular automata (CAs) are attractive for several reasons. Firstly, since the input to each cell update comes only from the cells nearby, the operation of the CA is relatively easy to follow. In a simulation with action-at-a-distance the cause of some change might not be easily ascertainable, it could be confusing to watch.

Computationally, CAs are attractive because their cost scales with the increase of the area being simulated, independently of the number of 'live' cells. Physics simulations involving rigid-body motion requiring collision detection, by contrast, can incur high computational costs as the number of interacting entities increases.

The form of simple artificial chemistry used in this 2D implementation of Squirm3 can be implemented in a CA. To achieve this we have to make use of two different neighbourhoods, first for the reaction step (the chemistry) and secondly for the movement step (the physics).

Reactions occur when two atoms are directly next to each other - when one is in the 4 (von Neumann) neighbourhood of the other. To avoid missing a possible reaction for a given atom we require the pairs of squares shown in Fig. 8 to be considered.

Since these updates only involve the two atoms in the squares, the pair checking can be run in parallel. This leads us to the four neighbourhood steps for the reactions phase illustrated on a 5×5 world in Fig. 9.

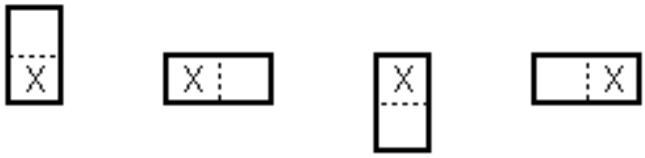

Figure 8: The 4 neighbours that need to be checked for possible reactions.

For the movement phase we can do something very similar. We can update the position of each atom in parallel as long as we ensure that there can be no conflict of two atoms moving to the same place - we have to 'lock' the local (Moore neighbourhood) atoms in place. We do this by using the neighbourhood update steps as illustrated in Fig. 10 on a 5×5 area. It is the atom (if any) in the centre of each square of nine that is considered for moving. If the square of nine overlaps the edge then the possible moves are reduced.

Note that in Fig. 10 each of the squares in the 5×5 area is checked exactly once. This update scheme does allow an atom to move more than once for each time step which we could disallow but within the random framework we have set up where reactions are rare compared to the number of timesteps this has very little effect.

Thus each time step update consists of four reaction update steps followed by nine movement update steps. This makes it unusual compared to traditional CAs which use a single update of the same neighbourhood at each time step. For a discussion of different neighbourhoods see Tim Tyler's web pages at `http://www.cell-auto.com`.

The CA implementation described could easily be extended to three or more dimensions, at each update step the immediate neighbours need to be checked for reaction candidates and then the movement neighbourhood needs to traversed. Similarly a simulation based on the collision of spheres in a continuous space (whether 2D or 3D) would exhibit the same properties, though the replication rates would inevitably be different.

5 Discussion

5.1 Self-Replication and Evolution

Artificial chemistries provide another medium in which self-replication is possible. While some ACs can be implemented as a CA (as this one can), the natural representation of a reaction-system is an abstract one, with various possible physical manifestations, and so justifies its own category.

Sayama [31] compared CA systems and computer programs in a table, which we reproduce here (Table 2) and augment with a third row. Sayama used this table to point out the key factors that differentiate between the categories. A system exhibiting competitiveness over mere self-replication requires both mortality and the spatial interaction between individuals. Competitive self-replicators will only be evolvable if they are robust to variation, and they will only be able to adapt to the actions of others if there is functional interaction between individuals. The

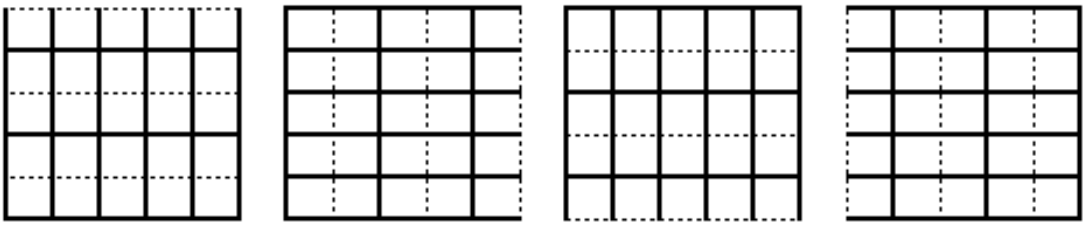

Figure 9: The 4 checks running in parallel for all atoms in the world. These 4 steps guarantee that every pair has been considered.

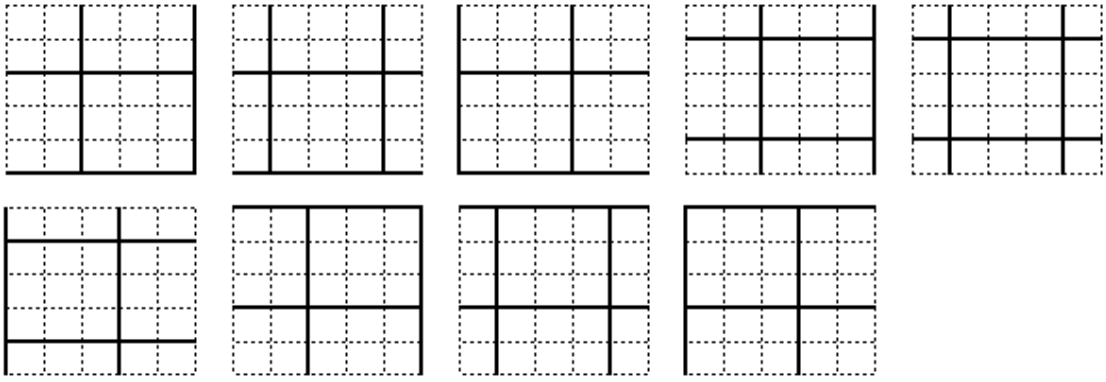

Figure 10: The 9 update steps for the movement of the atoms. In each step the atom (if any) in the centre of each square of nine can move to one of its empty 9-neighbour cells (if any).

fourth category is important because it is thought that the high selection pressure caused by an 'arms race' between two species is one source of the evolutionary drive to find ever more creative design solutions.

	Behaviour			
			Evolvable	
Media	**Self-reproductive**	**Competitive**	**Adaptive to physical environment**	**Adaptive to other individuals**
Computer programs	Computer viruses →	Core wars →		Tierra, Avida, Amoeba
Cellular automata	Langton's SR loops →	SDSR loops →	Evoloops →	?
Artificial chemistries	Ono and Ikegami's cell, LMA →		Squirm3 →	?

Table 2: A comparison of self-replicators in different categories.

One possibility to allow such interactions in Squirm3 is to include simple catalytic reactions, such as:

$$\mathrm{a1} + \mathrm{b0} \rightarrow \mathrm{a1} + \mathrm{b5} \tag{2}$$

Such reactions would not immediately interfere with the replication cycle but would affect the local chemistry in various ways. The example above would mean strings with a1 in them would cause b5's to be produced from b0's, which would then react with further b0's to make b3b6 pairs. Any strings with b1's in them would find it harder to replicate since there would be fewer of the required b0's available. This might yield a competitive advantage to those strings with a1 in them over strings with b1 in them.

Alternatively, reactions could be added that cause direct interactions between strings, such as:

$$\mathrm{c1} + \mathrm{d1} \rightarrow \mathrm{c1} + \mathrm{d9} \tag{3}$$

When a string with c1 in it encountered a d1 atom that was part of another replicating string it would cause that string to cease replicating, again this should give a competitive advantage to strings with c1 in them.

As defence against this kind of attack, a string could exist within a loop of non-interacting atoms, as shown in Fig. 11(a). As with real life cells, the membrane serves to protect the replicator from potentially damaging external reactions.

Of course the problem here is that the membrane is of fixed length and thus can only accommodate a certain number of replicators, soon there is no more room (Fig. 11(b)). Reactions that allowed the membrane itself to grow would assist the replicators inside.

Extending this style of fixed-link membrane to 3D, however, is a little bit difficult. A future chemistry that was a combination of Ono and Ikegami's (or Mayer

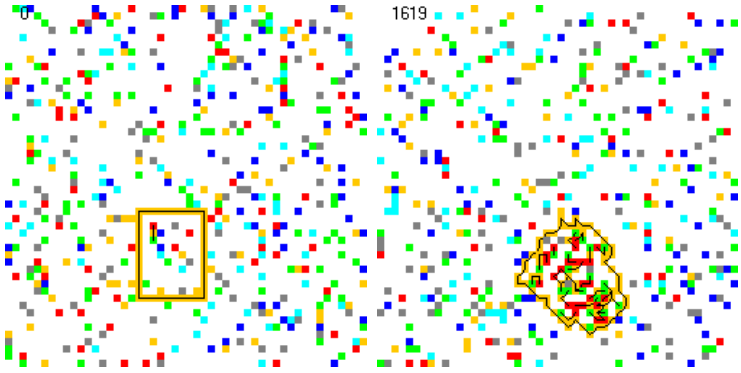

(a) (b)

Figure 11: If we start a replicator within a loop of non-interacting atoms (a) - here we've used a4's - then while single atoms can reach the inside to permit replication, the strings cannot get out. The strings are protected from any harmful external molecules but unless the membrane can grow the strings soon run out of room (b).

and Rasmussen's) and Squirm3 would be interesting, since flexible and extensible two- and three-dimensional membranes are exactly what Squirm3's replicating strings need to keep their reaction products to themselves and protect themselves from attack. If a string could encourage its surrounding membrane to grow by catalysing the right sort of reactions then it would give itself more room to replicate and would do better. The replicators and membranes could enter into a symbiotic relationship and eventually be indistinguishable from a single organism, as is thought to have happened in our world.

It is tempting to follow the route of adding reactions that permit more and more sophisticated behaviours. Much experimentation would be needed to find which reactions reliably caused the desired pattern of behaviour. However, such an approach would surely rob the simulation of the 'surprise!' factor of emergence [29] and wouldn't necessarily encourage the evolutionary growth of complexity.

So rather than carefully engineering reactions towards specific behavioural goals we need to consider how a wider range of actions can be implemented. If our replicating strings had some mechanism for producing different molecules ('proteins') as determined by the order of bases in their length then this would give them some of the powers of a universal constructor. We could then let the evolving strings work through the design space of that environment rather than us having to work through the design space of different chemistries. This is in fact the ultimate goal of creating evolutionary systems such as Squirm3.

5.2 Origin of Life

There is much debate about what form the first replicators on Earth would have taken. The RNA World theory [11, 12] proposes that self-replicating RNA molecules

sprung into existence from a soup of non-replicating molecules. By contrast, the Lipid World theory [32] suggests that in fact membranes came first, and that their innate abilities to grow and divide encouraged the development of more complicated structures. Another theory suggests that the crystals in certain clays acted as prebiotic replicators [4, 5], with their surrounding clays being the recipients of their properties and protecting them in return.

Artificial chemistries are good for testing theories about early replicators since they can be programmed with many different properties and run repeatedly. Their computational efficiency makes them effective for studying complex systems that only function in the presence of many interacting entities. This is especially true if the AC can be implemented in a CA - the 100×100 world experiments run at over 1000 iterations per second on a 1.7GHz PC when implemented in C++ (and the 20×20 ones at over 10000).

The Squirm3 system demonstrates one possible set of self-replicating molecules but by exploring other methods that work or don't work we might be able to form a model of what the chemical requirements are for evolutionary replication to gain a foothold.

6 Conclusions

We have described a novel system of reactions in an artificial chemistry environment in which self-replicating molecules exist and can spontaneously form under the right conditions. The molecules carry with them a sequence of atoms that aren't essential to the replication process but could form a set of instructions for influencing the local environment to improve the survival chances of the replicator. The evolution of interesting design solutions is the goal of the system but this is not yet achieved; the replicators can evolve but so far the only selection pressures encountered have been towards shorter molecules.

There is a world of difference between a syntactic description of the replication mechanism of Squirm3 and what actually happens:

$$\mathrm{e8}\,\{x1\}^*\,\mathrm{f1} + \{x0\}^* \rightarrow \mathrm{e8}\,\{x1\}^*\,\mathrm{f1} + \mathrm{e8}\,\{x1\}^*\,\mathrm{f1} \tag{4}$$

The effects of the local non-availability of raw material, the malicious interruption of replication, the use of a barrier to safeguard the replicator (Fig. 11) - none of these interesting factors can be captured by this abstract textual description of the replication process, it is the simulation of the individual atoms that gives us the rich interactions that we hope will lead to the evolutionary growth of complexity.

Squirm3 is implemented in Java and C++. The project is open-source and available online, so people can experiment freely:

`http://www.eastman.ucl.ac.uk/~thutton/Evolution/Squirm3`

Acknowledgments

Thanks to Daniel LaLiberte, Erwin Achermann and others on the Primordial Life mailing list for much helpful discussion and motivation. Three anonymous reviewers helped to clarify the important points of this paper.

References

[1] Adami, C., & Brown, C. T. (1994). Evolutionary learning in the 2D artificial life system "Avida". In R. Brooks & P. Maes (Eds.), *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems* (pp. 377–381). Cambridge, MA: MIT Press.

[2] Bedau, M. A., McCaskill, J. S., Packard, N. H., Rasmussen, S., Adami, C., Green, D. G., Ikegami, T., Kaneko, K., & Ray, T. S. (2000). Open problems in artificial life. *Artificial Life, 6*, 363–376.

[3] Byl, J. (1989). Self-reproduction in small cellular automata. *Physica D, 34*, 295–299.

[4] Cairns-Smith, A. G. (1982). *Genetic takeover and the mineral origins of life.* Cambridge: Cambridge University Press.

[5] Cairns-Smith, A. G. (1985). *Seven clues to the origin of life.* Cambridge: Cambridge University Press.

[6] Chou, H. H., & Reggia, J. A. (1997). Emergence of self-replicating structures in a cellular automata space. *Physica D, 110*, 252–276.

[7] Chou, H. H., & Reggia, J. A. (1998). Problem solved during artificial selection of self-replicating loops. *Physica D, 115*, 293–372.

[8] Dewdney, A. K. (1984). In the game called Core War hostile programs engage in a battle of bits. *Scientific American, 250*, 15–19.

[9] Dittrich, P., Ziegler, J., & Banzhaf, W. (2001). Artificial chemistries - a review. *Artificial Life, 7*, 225–275.

[10] Downing, K. L. (2001). Modular designer chemistries for artificial life. In L. Spector, E. D. Goodman, A. Wu, W. B. Langdon, H.-M. Voigt, M. Gen, S. Sen, M. Dorigo, S. Pezeshk, M. H. Garzon, & E. Burke (Eds.), *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001)* (pp. 845–852). San Francisco, CA: Morgan Kaufmann.

[11] Gilbert, W. (1986). The RNA world. *Nature, 319*, 618.

[12] Joyce, G. F. (2002). The antiquity of RNA-based evolution. *Nature, 418*, 214–221.

[13] Langton, C. G. (1984). Self-reproduction in cellular automata. *Physica D, 10*, 135–144.

[14] Luisi, P. L., Walde, P., & Oberholzer, Th. (1994). Enzymatic RNA synthesis in self-reproducing vesicles: an approach to the construction of a minimal cell. *Ber. Bunsenges. Phys. Chem., 98*, 1160–1165.

[15] Maley, C. C. (1999). Four steps toward open-ended evolution. In W. Banzhaf, J. Daida, A. E. Eiben, M. H. Garzon, V. Honavar, M. Jakiela, & R. E. Smith (Eds.), *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO99)* (pp. 1336–1343). Orlando, FL: Morgan Kaufmann.

[16] Mayer, B., & Rasmussen, S. (2000). Dynamics and simulation of micellar self-reproduction. *International Journal of Modern Physics C, 11*, 809–826.

[17] McMullin, B. (2000). John von Neumann and the evolutionary growth of complexity: Looking backwards, looking forwards... In M. A. Bedau, J. S. McCaskill, N. H. Packard, & S. Rasmussen (Eds.), *Artificial Life VII: Proceedings of the seventh international conference* (pp. 467–476). Cambridge, MA: MIT Press.

[18] Ono, N., & Ikegami, T. (1999). Model of self-replicating cell capable of self-maintenance. In D. Floreano, J. -D. Nicoud, & F. Mondada (Eds.), *Proceedings of the Fifth European Conference on Artificial Life (ECAL99)* (pp. 399–406). Lausanne, Switzerland: Springer.

[19] Ono, N., & Ikegami, T. (2000). Self-maintenance and self-reproduction in an abstract cell model. *Journal of Theoretical Biology, 206*, 243–253.

[20] Ono, N., & Ikegami, T. (2001). Artificial chemistry: Computational studies on the emergence of self-reproducing units. In *Proceedings of the Sixth European Conference on Artificial Life (ECAL01)* (pp. 186–195). Prague, Czech Republic: Springer.

[21] Pargellis, A. N. (1996). The evolution of self-replicating computer organisms. *Physica D, 98*, 111–127.

[22] Pargellis, A. N. (2001). Digital life behavior in the Amoeba world. *Artificial Life, 7*, 63–75.

[23] Perrier, J. -Y., Sipper, M., & Zahnd, J. (1996). Toward a viable, self-reproducing universal computer. *Physica D, 97*, 335–352.

[24] Pesavento, U. (1995). An implementation of von Neumann's self-reproducing machine. *Artificial Life, 2*, 337–354.

[25] Rasmussen, S., Baas, N. A., Mayer, B., Nilssen, M., & Olesen, M. W. (2002). Ansatz for dynamical hierarchies. *Artificial Life, 7*, 329–353.

[26] Ray, T. S. (1991). An approach to the synthesis of life. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial Life II: Proceedings of the Workshop on Artificial Life (Vol. X of SFI Studies in the Sciences of Complexity)* (pp. 371–408). Redwood City, CA: Addison-Wesley.

[27] Rebek Jr, J. (1994). Synthetic self-replicating molecules. *Scientific American, 271*, 34–40.

[28] Reggia, J. A., Armentrout, S. L., Chou, H. H., & Peng, Y. (1993). Simple systems that exhibit template-directed replication. *Science, 259*, 1282–1287.

[29] Ronald, E. M. A., Sipper, M., & Capcarrère, M. S. (1999). Design, observation, surprise! A test of emergence. *Artificial Life, 5*, 225–239.

[30] Sayama, H. (1998). Introduction of structural dissolution into Langton's self-reproducing loop. In C. Adami, R. K. Belew, H. Kitano, & C. E. Taylor (Eds.), *Artificial Life VI: Proceedings of the Sixth International Conference on Artificial Life* (pp. 114–122). Cambridge, MA: MIT Press.

[31] Sayama, H. (1999). A new structurally dissolvable self-reproducing loop evolving in a simple cellular automata space. *Artificial Life, 5*, 343–365.

[32] Segré, D., & Lancet, D. (2000). Composing life. *EMBO Reports, 1*, 217–222.

[33] Smith, A., Turney, P., & Ewaschuk, R. (2002). JohnnyVon: Self-replicating automata in continuous two-dimensional space. Technical Report ERB-1099, National Research Council Canada.

[34] Szathmáry, E. (1997). The first two billion years. *Nature, 387*, 662–663.

[35] Szostak, J. W., Bartel, D. P., & Luisi, P. L. (2001). Synthesizing life. *Nature, 409*, 387–390.

[36] Taylor, T. (1999). Creativity in evolution: Individuals, interactions and environment. In P. Bentley & D. Corne (Eds.), *Proceedings of the AISB'99 Symposium on Creative Evolutionary Systems, The Society for the Study of Artificial*

Intelligence and Simulation of Behaviour. Edinburgh: Morgan Kaufman.

[37] Taylor, T. (1999). *From artificial evolution to artificial life*. PhD thesis, Division of Informatics, University of Edinburgh.

[38] Taylor, T. (1999). On self-reproduction and evolvability. In D. Floreano, J. - D. Nicoud, & F. Mondada (Eds.), *Proceedings of the Fifth European Conference on Artificial Life (ECAL99)* (pp. 94–103). Lausanne, Switzerland: Springer.

[39] Tempesti, G. (1995). A new self-reproducing cellular automaton capable of construction and computation. In F. Morán, A. Moreno, J. J. Merelo, & P. Chacón (Eds.), *ECAL'95: Third European Conference on Artificial Life, LNCS929* (pp. 555–563). Heidelberg: Springer-Verlag.

[40] von Kiedrowski, G. (1986). A self-replicating hexadeoxynucleotide. *Angew. Chem. Int. Ed. Engl., 25*, 932–935.

[41] von Neumann, J. (1966). *The theory of self-reproducing automata, A. Burks (Ed.)*. Urbana, IL: University of Illinois Press.

[42] Wright, M. C., & Joyce, G. F. (1997). Continuous in vitro evolution of catalytic function. *Science, 276*, 614–617.