Rediscovering Computational Autopoiesis

Barry McMullin
Santa Fe Institute
1399 Hyde Park Road
Santa Fe, NM 87501
USA

mcmullin@eeng.dcu.ie
http://www.eeng.dcu.ie/~mcmullin

Francisco J. Varela
LENA—CNRS URA 654
Hopital del la Salpetriere
47, Blvd. de l'Hopital
75651 Paris cedex 13
France.

fv@ccr.jussieu.fr

Abstract

This paper summarises some initial empirical results from a new computer model (artificial chemistry) which exhibits spontaneous emergence and persistence of *autopoietic* organisation. The model is based on a system originally presented by Varela, Maturana and Uribe [11]. In carrying out this re-implementation it was found that an additional interaction (chain-based bond inhibition), not documented in the original description by Varela et al., is critical to the realisation of the autopoietic phenomena. This required interaction was re-discovered only following careful examination of (unpublished) source code for an early version of the original model. The purpose of the paper is thus twofold: firstly to identify and discuss this previously undocumented, but essential, interaction; and secondly to argue, on the basis of this particular case, for the importance of exploiting the emerging technologies which support publication of completely detailed software models (in addition, of course, to conventional publication of summary experimental results).

Keywords: Autopoiesis, Artificial Life, Artificial Chemistry, Origin of Life.

1 Introduction

The concept of *autopoiesis* [3, 9] occupies a distinctive position in the entire field of biology as one of the very few substantive attempts to give an integrated characterisation of the nature of living systems which is clearly separate from a mere listing of arbitrary "properties" (such as metabolism, growth, reproduction etc.). The concept was originated some twenty-five years ago, by Humberto Maturana and Francisco Varela [10]. Its influence since then has been diverse and sustained—see, for example, [13, 7].[1]

The first widely distributed, and thus seminal, description of the concept of autopoiesis was that of [11], which was illustrated with a computer model of a "minimal" example. Experimental data from this model showed both the spontaneous formation and ongoing repair of an autopoietic system embedded in a two dimensional, discrete space. This was accompanied by a qualitative description of the artificial chemistry realised by the model, and a more detailed algorithmic account of the simulation program.

This computer model has been extremely influential in providing a relatively simple, graphic, exemplar of the concept of autopoiesis. It demonstrated that the idea of autopoietic organisation, although subtle and abstract, could be instantiated in a relatively simple, and concrete, system.

However: a recent reappraisal of the original presentation of this computer model has revealed significant flaws—flaws which, if they were left uncorrected, might tend to undermine its role as a concrete example of autopoiesis.

A number of technical difficulties with even interpreting the original algorithm, and apparent discrepancies between the algorithm and the experimental data, have been discussed in a previously published working paper [5]. That paper also incorporates, as an appendix, the `FORTRAN-IV` code of a version of the original program used by Varela et al. Careful study of this code has now allowed the identification of an additional interaction, present in the code, but omitted from all published descriptions of the model.

[1] An excellent, comprehensive, bibliography of the literature on autopoiesis is maintained by Randall Whitaker at:

`http://www.informatik.umu.se/~rwhit/AT.html`

In this paper we present experimental results from a completely new implementation of the qualitative chemistry described by Varela et al. which suggests that this additional interaction is, indeed, critical to the realisation of the autopoietic phenomena; and that, conversely, provided this additional interaction is included, the autopoietic phenomena are not dependent on any particular details of the original program or algorithm, but may be expected in *any* system sharing the same qualitative chemistry.

2 The Original Qualitative Chemistry

The chemistry takes places in a discrete, two dimensional, space. Each position in the space is either empty or occupied by a single particle. Particles generally move in random walks in the space. There are three distinct particle types, engaging in three distinct reactions:

- Production: Two substrate (`S`) particles may react, in the presence of a catalyst (`K`) particle to form a link (`L`) particle.
- Bonding: `L` particles may bond to other `L` particles. Each `L` particle can form (at most) two bonds, thus allowing the formation of indefinitely long chains, which may close to form membranes. Bonded `L` particles become immobile.
- Disintegration: An `L` particle may spontaneously disintegrate, yielding two `S` particles. When this occurs any bonds associated with the `L` particle are destroyed also.

Chains of `L` particles are permeable to `S` particles but impermeable to `K` and `L` particles. Thus a closed chain, or membrane, which encloses `K` or `L` particles effectively traps such particles.

3 The Phenomena

The basic autopoietic phenomenon predicted for this system is the possibility of realising dynamic cell-like structures which, on an ongoing basis, produce the conditions for their own maintenance. Such a system would consist of a closed chain (membrane) of `L` particles enclosing one or more `K` particles. Because `S` particles can permeate through the membrane, there can be ongoing production of `L` particles. Since these cannot escape from the membrane, this will result in the build up of a relatively high concentration of `L` particles. On an ongoing basis, the membrane will rupture as a result of disintegration of component `L` particles. Because of the high concentration of `L` particles inside the membrane, there should be a high probability that one of these will drift to the rupture site and effect a repair, *before* the `K` particle(s) escape, thus re-establishing precisely the conditions allowing the build up of that high concentration of `L` particles.

A secondary phenomenon which *may* arise is the spontaneous establishment of an autopoietic system from a randomised initial arrangement of the particles.

Clearly, the issue of spontaneous formation does not arise unless the system actually supports autopoietic organisation. In this sense the phenomenon of autopoietic organisation is *logically* prior to spontaneous formation (though chronologically following from it). For this reason, the phenomenon of spontaneous formation will *not* be considered further in this paper. Instead, in all the experiments reported, a putatively autopoietic entity will be artificially introduced into the system; the question at issue will be whether this entity succeeds in realising the autopoietic reaction network already described.

4 The SCL Program

The newly developed program is called SCL (for Substrate-Catalyst-Link) [6]. This has been implemented using the SWARM[2] simulation system, developed at the Santa Fe Institute[3].

A conscious decision was taken that SCL would *not* be based on the *algorithm* originally published by Varela et al., but should rather reflect an independent implementation of the same *qualitative* chemistry. This stemmed partly from the previously documented problems with the original algorithm [5]; but it also reflected a desire to test the robustness of the autopoietic phenomena—i.e. are they perhaps reliant on some artifact of the original program and/or algorithm, or are they robust outcomes from the given qualitative chemistry.

The version of SCL used to generate the results described in this paper (v0.05.01) differs in minor ways from that described in the previously published documentation (v0.04) [6]. The complete source code relating to SCL v0.05.01 is available as:

```
ftp://ftp.santafe.edu/pub/swarm/
users-contrib/anarchy/scl-0.05.01.tar.gz
```

The SCL data files relating to the experiments described here are available in:

```
ftp://ftp.santafe.edu/pub/swarm/
users-contrib/anarchy/scl-data00.tar.gz
```

A key to the display of the three particle types in SCL is provided in figure 1.

5 Experimental Protocol

This paper will describe the results from two separate experiments with SCL. In each case the same experimental protocol was followed. Each experiment consisted of 5 runs of SCL. The initial configuration was identical in

[2] `http://www.santafe.edu/projects/swarm`
[3] `http://www.santafe.edu`

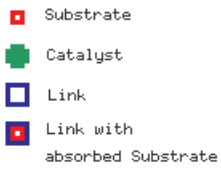

Figure 1: Key to Particle Types.

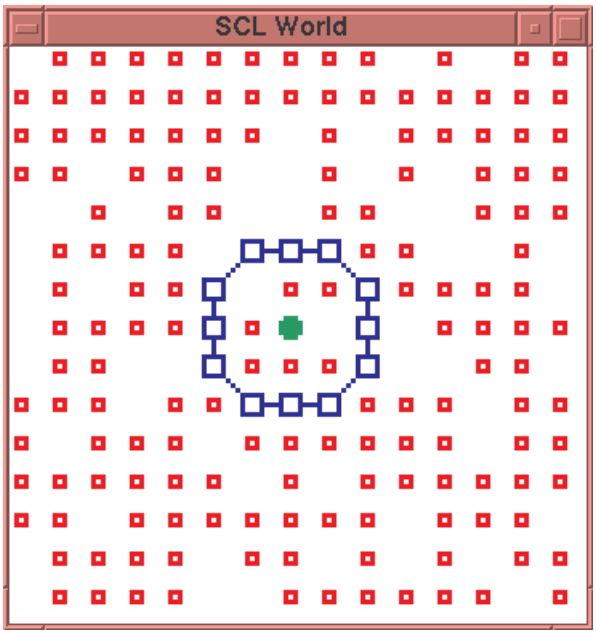

Figure 2: Initial Configuration.

all runs, and is illustrated in figure 2. This comprises a single artificially constructed cell-like entity, being a closed membrane of `L` particles enclosing a single `K` particle. This is embedded in a 15×15 toroidal space. The five runs in each experiment differed only in the initial state of the underlying pseudo-random number generator. These five distinct initial states are specified in the files `run1.stt` through `run5.stt` in the distributed data file archive (`scl-data00.tar.gz`).

The two experiments differed only in that experiment 1 implemented just the reactions of the qualitative chemistry described in section 2, whereas experiment 2 incorporated the newly rediscovered chain-based bond inhibition interaction (to be discussed in section 7).

SCL supports a variety of parameters controlling reaction rates, mobility parameters, etc. Apart from the parameter controlling the additional interaction just mentioned, these parameters were held constant across all runs and both experiments. These parameter sets for experiments 1 and 2 are specified in the files `exp1.prm` and `exp2.prm` respectively, in the distributed data file archive.

The `disintegrationProbability` parameter was set at 0.001 in all cases.[4] This is the probability that any given `L` particle will disintegrate per unit time. The membrane in the initial configuration is composed of 12 `L` particles. It follows that the expected time to first rupture of the initial membrane is given by:

$$\tau = \frac{1}{1-(1-P_d)^{12}} \simeq 84$$

6 Experiment 1

6.1 Run 1-1

As expected, `S` particles initially permeate through the membrane and, under the influence of the `K` particle, production of `L` particles starts. However, instead of these `L` particles remaining mobile, trapped within the membrane, in readiness to repair any rupture, they begin to spontaneously bond to *each other*. Given that bonded `L` particles are specified to be immobile, this means that such particles are *not* available to drift to a rupture site. The screenshot of figure 3 was taken at time 110. The membrane has not yet suffered any decay. However, the interior of the membrane is now completely clogged with bonded—and thus immobile—`L` particles. Only two open positions remain inside the membrane, one occupied by the `K` particle. Since the production reaction requires two `S` particles adjacent to each other and to the `K` particle, there is no longer any available site for further production within the membrane, and further production of `L` particles is impossible. It follows that, whenever the membrane does eventually rupture, there will be no mobile `L` particles available to effect a repair.

In fact, the membrane suffers a double rupture at times 234 and 235, yielding the configuration shown in figure 4. The chain which had previously been formed inside the membrane now becomes spliced to one side of the rupture site, forming a folded chain. This no longer encloses the `K` particle. Indeed, should the folded chain become closed, the `K` particle would necessarily be *outside* it. Thus, the initial, putatively autopoietic, entity has clearly now irreversibly degenerated, without having undergone even a single episode of self-repair.

6.2 Run 1-2

On this run, the initial rupture of the membrane occurs relatively early, at time 31. Just one `L` particle has been produced within the membrane by this time. However,

[4] This is a factor of 10 smaller than the value originally suggested by Varela et al. This reflects the fact that the *maximum* rate of the production reaction is approximately this much slower in SCL than in the original model, so that 10 timesteps in SCL can be considered roughly comparable to one timestep in the original.

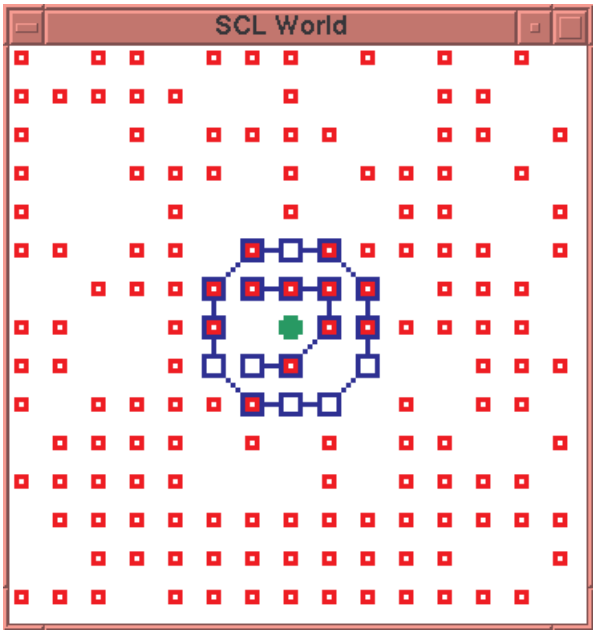

Figure 3: Experiment 1, Run 1, Time 110.

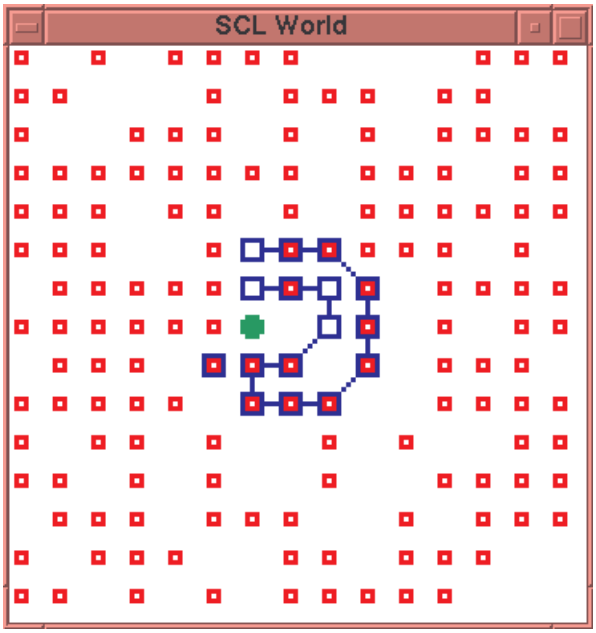

Figure 4: Experiment 1, Run 1, Time 235.

since this means the particle is still mobile, there is at least some possibility that it may drift to the rupture site and effect a repair. A second L particle is produced at time 63, thus improving the possibility of a repair. However, at time 69, these two L particles bond to each other, thus becoming immobile, and unavailable to drift

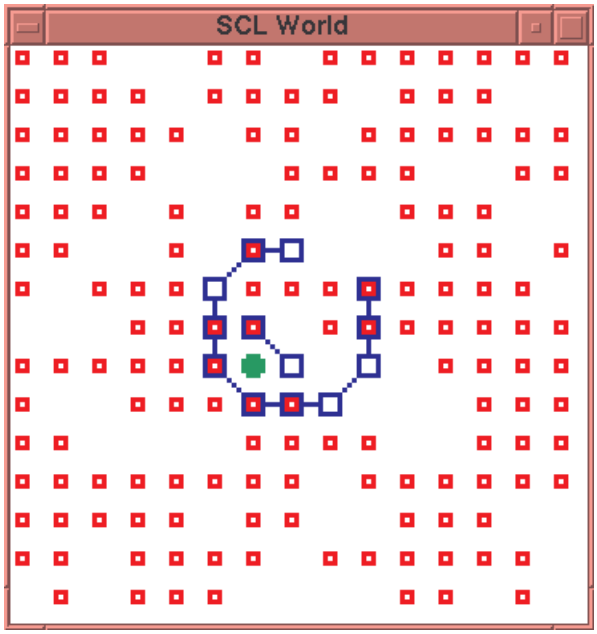

Figure 5: Experiment 1, Run 2, Time 069.

to the rupture site. Moreover, they are located in such a position that the K particle is blocked in all four directions. This configuration is illustrated in figure 5. Not only is the K particle now also effectively immobile, but, again, there is no space available adjacent to it to permit further production of new L particles. Thus, there is no possibility of repairing the existing rupture. As with run 1, the initial entity has clearly already irreversibly degenerated.

6.3 Runs 1-3 to 1-5

Given the descriptions of runs 1 and 2, only a brief description of the remaining runs is necessary. Precisely the same failure mechanism is again observed: the L particles produced within the membrane spontaneously bond to each other, thus becoming immobile, rather than remaining available to drift to a rupture site when it arises; the interior of the membrane becomes progressively clogged up, until there is no longer space available for further production. At this point, since no L particles are available to repair any rupture, and no more can be produced within the membrane, the original entity has effectively degenerated. In all three runs this occurs without even a single episode of successful repair of the membrane. The times at which this condition is reached are as follows:

Run	Time
3	282
4	126
5	165

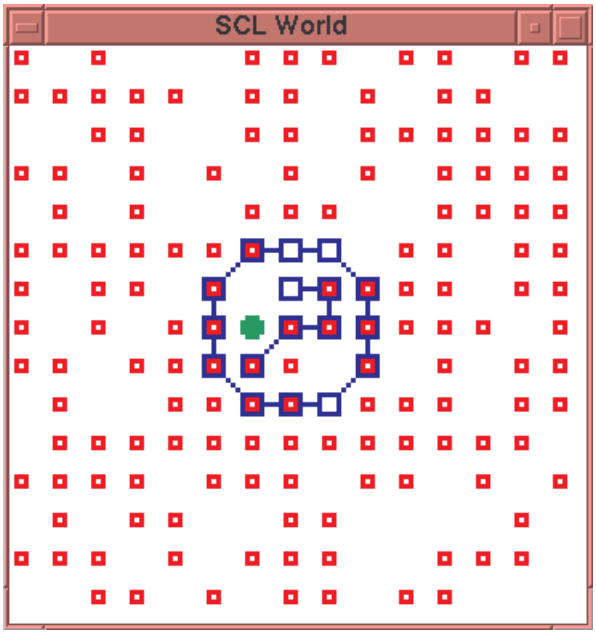

Figure 6: Experiment 1, Run 3, Time 282.

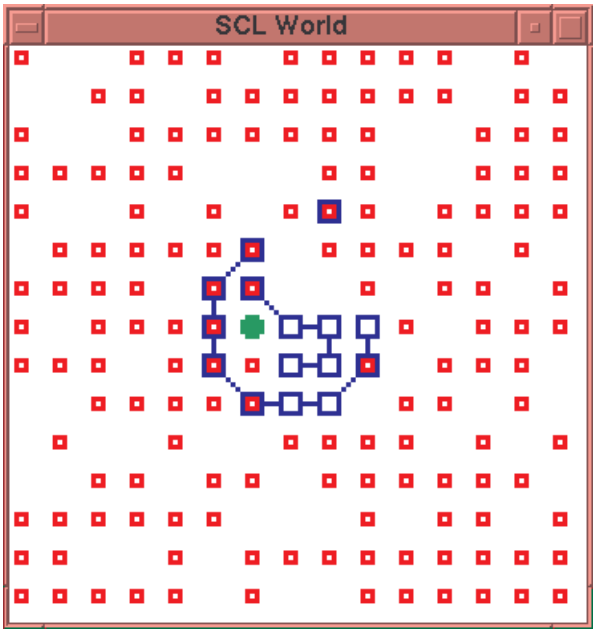

Figure 7: Experiment 1, Run 4, Time 126.

These terminal configurations for runs 3–5 are shown in figures 6, 7 and 8 respectively.

6.4 Discussion 1

In all five runs of experiment 1 a consistent *failure* of the autopoietic process was observed. This is due to

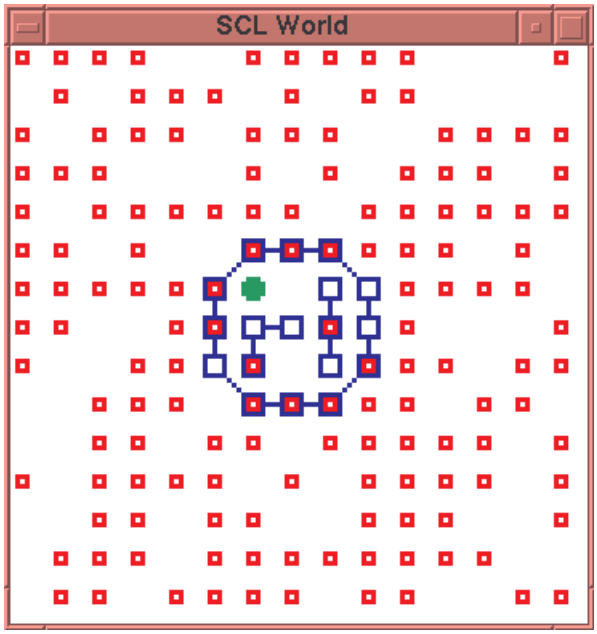

Figure 8: Experiment 1, Run 5, Time 165.

the spontaneous and premature bonding of the `L` particles produced within the membrane, thus making them immobile and unavailable to effect a repair to the membrane. With the benefit of these experimental results, it seems fairly obvious that this failure mode was already implicit in the *qualitative* chemistry described by Varela et al. [11]. It is evidently *not* dependent on any particular details of the implementation, nor on the specific parameters settings.

This conclusion is corroborated by the fact that this same failure mechanism has been observed in previous (unpublished) experiments with two other, *independent*, implementations of this reaction scheme [1, 8], and has also been previously reported by Lizana [2]. This class of failure seems also to have been recognised, at least implicitly, in the re-implementation(s) carried out by Zeleny [12, Figure 4].

Two attempted solutions to this failure mechanism were briefly investigated, before the preferred solution, to be discussed in section 7, was finally identified.

Firstly, the bonding reaction was separated into two cases: bonding between two free `L` particles, and bonding between a free `L` particle and an `L` particle already having one bond. The latter is the case of interest for membrane repair. These were controlled by separate rate parameters. This allowed the "spontaneous" bonding reaction to be made very slow. This should ensure that the free `L` particles formed within the membrane would not spontaneously bond with each other but would rather be

held in reserve for membrane repair.[5] However, this idea proved largely ineffective. The problem is that, once a rupture *does* occur it frequently happens that, instead of a single free L repairing the membrane, *all* of the free L particles become quickly incorporated into an inward spiraling chain fragment.

The second mechanism appears to have been independently suggested by both Zeleny [12] and Lizana [2]. This involves inhibiting bonding to a free L particle in some neighborhood of any K particle. In Lizana's case, this effect seems to have been limited to the *immediate* (Moore) neighborhood of a K particle, whereas Zeleny seems to have used arbitrarily large (and dynamically changing?) neighborhoods. The idea appears to be that the K particle(s) can establish zone(s) of bond inhibition around them. The membrane can then form (roughly) at the edge of such a zone. L particles within these zones will remain free, and ready to drift to a rupture site to effect a repair.

Both Zeleney and Lizana apparently got this mechanism to give somewhat satisfactory results. The mechanism has been investigated to only a limited extent with SCL. Specifically, the use of indefinitely large inhibition zones (as suggested by Zeleny) has *not* been pursued, since it violates an objective that the model should rely only on local (Moore neighborhood) interactions. With this (self-imposed) restriction, the results have generally been mediocre. Two counteracting effects have been noticed. Firstly, even within a relatively small membrane such as illustrated in figure 2, the K particle may transiently drift away from the central position; if free L particles also drift away from this position, then they may still be able to spontaneously bond and become immobile. Even though the K particle may drift back into their vicinity, it is now too late—the bonding has already occurred.[6] Secondly, if a rupture occurs in the neighborhood of the K particle it is now very difficult to effect a repair, even if a free L particle should drift into an appropriate position; worse still, this is precisely the situation in which swift repair is most important, lest the K particle should escape completely. These problems can be overcome, to an extent, by making the K particle immobile (in the center of the cavity). While Lizana's description is not fully detailed, it seems that this may be what she indeed did. The mobility of K particles also seems to have been severely constrained in a number of Zeleny's experiments. In our view, this significantly reduces the generality and interest of the model, and must be considered an unsatisfactory solution.

[5] Of course, this would make spontaneous formation of an initial membrane much less probable; but that issue was deferred.

[6] Presumably, the K particle should not be assigned an effect of rupturing these bonds again, because they cannot be distinguished from the bonds making up the membrane; on the other hand, this does seem to have been a mechanism actually used by Zeleny in some experiments [12, Figure 6].

7 Chain-based Bond Inhibition

By far the most troubling aspect of the results discussed above is that they are not consistent with the experimental results originally presented by Varela et al. [11].

In particular, a careful examination of those original results suggests that the model must have had some, unspecified, mechanism to overcome or preempt the class of failure now described here. However, given that the work was done over 25 years ago, it seemed that it would be extremely difficult to gain much further insight into this problem. The current author who was involved in the original work (Varela), no longer had any clear recollection of what additional mechanism was present in the model to account for this discrepancy.

Fortunately, a printout of an early version of the original simulation program, coded in FORTRAN-IV, has recently been rediscovered, and has now been incorporated in a published technical report [5]. As discussed in more detail in that report, this program has been rekeyed, and it has been possible to execute it again. This did not reproduce the precise results of the original publication; most likely it was not exactly the same version, and, in any case, the original pseudo-random number source is no longer available, so the precise execution trajectory is bound to be different. However, this did suffice to show that the program did, indeed, exhibit some mechanism whereby free links, confined within the membrane, tended *not* to spontaneously bond to each other. This motivated a detailed reanalysis of the program code, which finally resulted in the identification of a previously un-reported interaction—*chain-based bond inhibition.*

This is an interaction whereby bonding is inhibited to any free L particle which is in the immediate vicinity of another L particle which is doubly bonded. In effect then, a free L particle cannot form a bond as long as it is alongside (as opposed to at the end of) an existing chain of L particles; but it *can* form bond(s) when it is at the end of a chain; and, especially, when it is positioned at a site where a chain has broken (i.e. a rupture site).[7]

The next section reviews experimental results from SCL when the chain-based bond inhibition reaction is enabled.

8 Experiment 2

8.1 Run 2-1

Between time 0 and time 226 the initial membrane suffers two ruptures which are repaired with no change of

[7] This interaction has previously been outlined in the SCL documentation [6]. However, there is an error, or ambiguity, in that earlier description, in that it suggests that bond inhibition applies to both free and singly bonded L particles. In fact, it applies only to free L particles. Applying it to singly bonded L particles would actually *prevent* membrane repair from taking place.

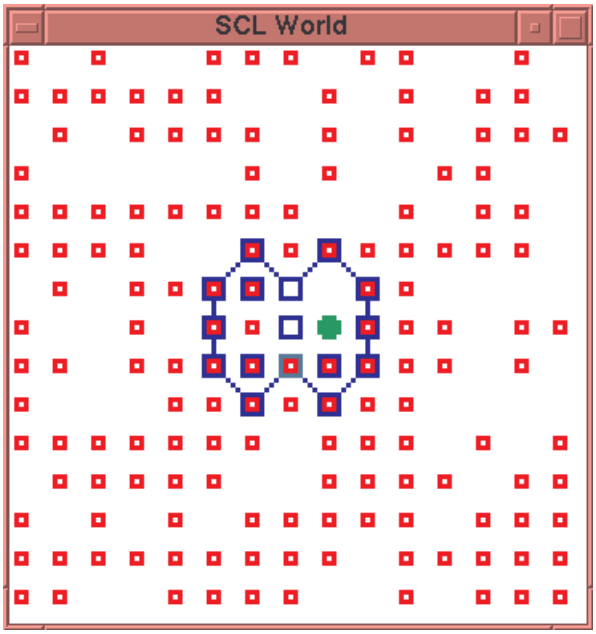

Figure 9: Experiment 2, Run 1, Time 444.

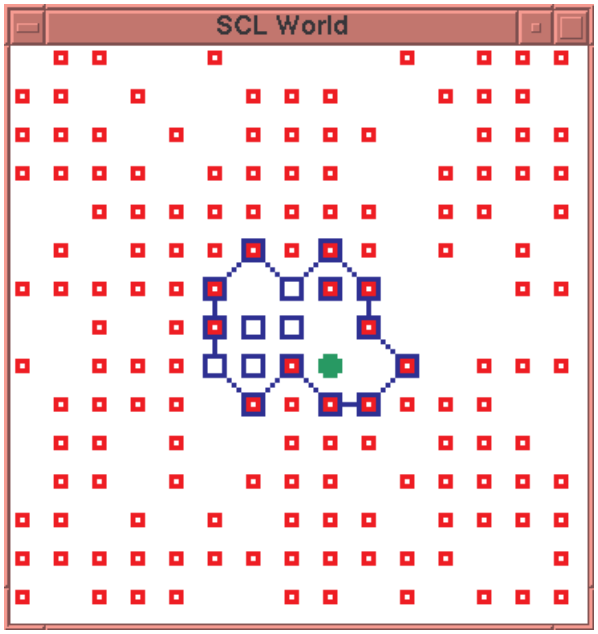

Figure 10: Experiment 2, Run 1, Time 1310.

membrane morphology. Between time 227 and 444 there are four rupture and repair episodes, yielding the new membrane morphology shown in figure 9.

This new morphology appears to be relatively robust. The entity persists in this morphology up to time 1250, in the course of which there are 5 more rupture and (successful) repair episodes. Between time 1250 and time 1310 there are two rupture and repair episodes yielding the new membrane morphology shown in figure 10. The entity survives in this morphology, through two more rupture and repair episodes until time 1741. There are then two ruptures in quick succession, at times 1742 and 1745. At time 1746 the membrane fragments, and partially spirals into the cavity, as shown in figure 11. It is then no longer possible to recover the closed membrane through any simple process of self repair.

8.2 Run 2-2

Between time 0 and time 133 the initial membrane suffers three ruptures which are repaired with no change of membrane morphology. A rupture at time 134 is repaired at time 137, yielding the new membrane morphology shown in figure 12. The entity persists in this morphology up to time 452, in the course of which there is one more rupture and (successful) repair episode. A further rupture at time 453 is eventually repaired at time 555; but in the interim, a second rupture at time 542 leads to a partial spiral into the cavity, as with run 1, and again it is then no longer possible to recover a closed

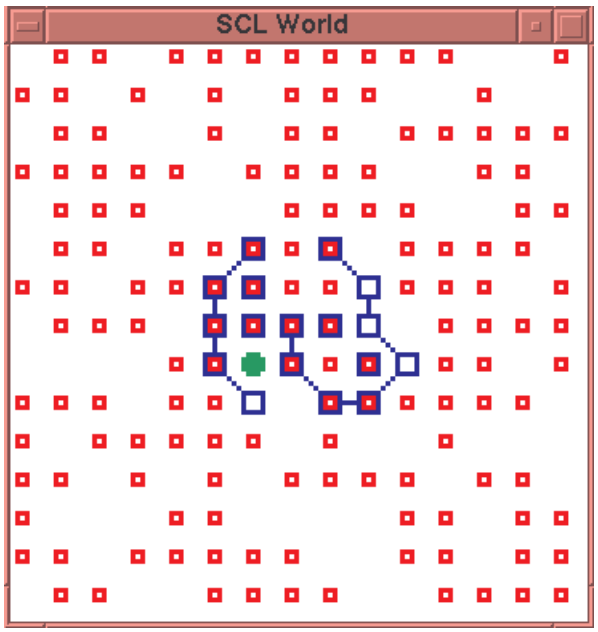

Figure 11: Experiment 2, Run 1, Time 1746.

membrane through any simple process of self repair.

8.3 Run 2-3

In this run there are two very early ruptures (times 6 and 13), before there has been time for an effective build up in the concentration of free `L` particles. The `K` particle

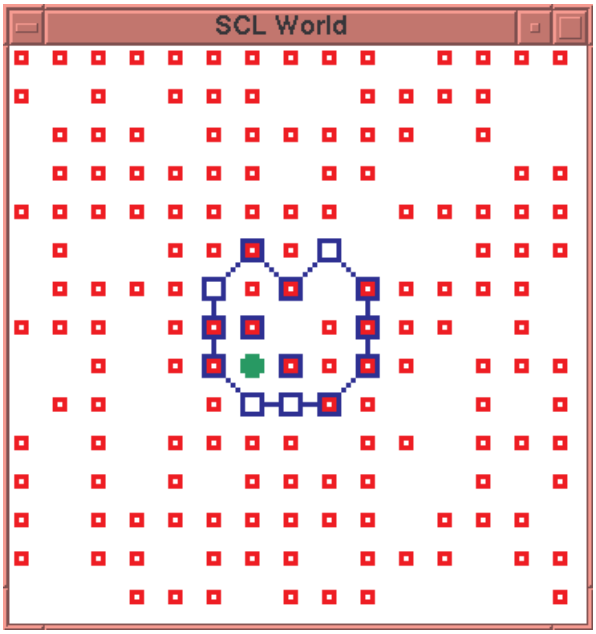

Figure 12: Experiment 2, Run 2, Time 137.

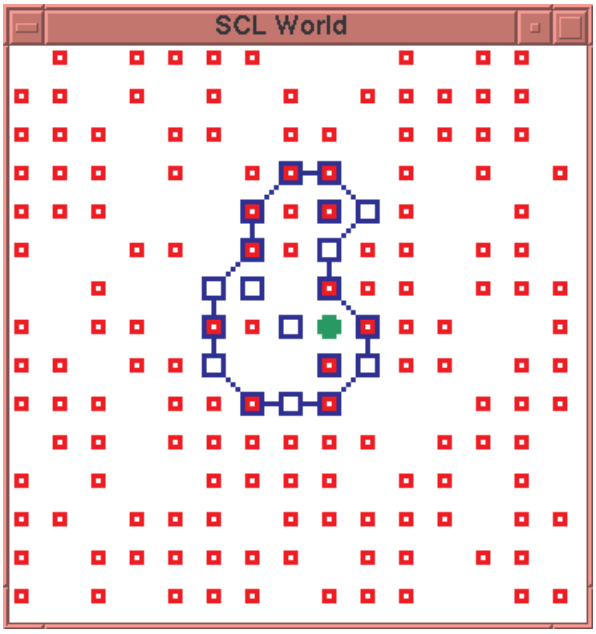

Figure 13: Experiment 2, Run 3, Time 148.

almost escapes immediately, but, instead, an extension of the membrane forms around it. There is a further independent rupture at time 139, but at time 148 a closed membrane reforms with the new morphology shown in figure 13

A further rupture at time 171 results again in an inward spiral and it is then no longer possible to recover a closed membrane through any simple process of self repair.

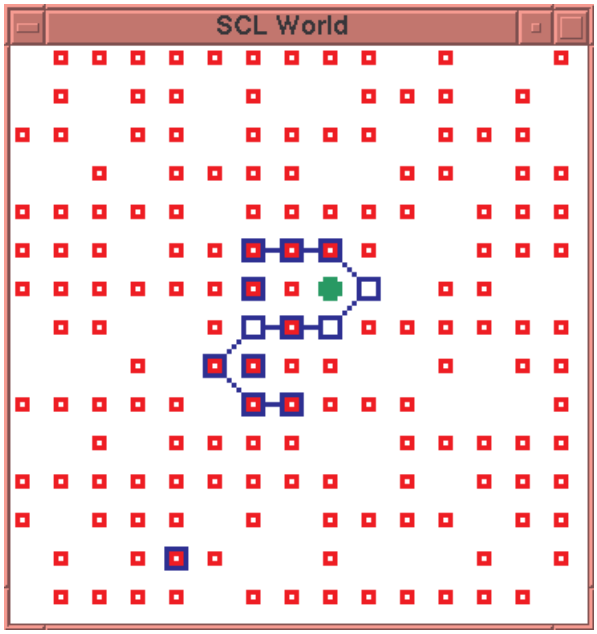

Figure 14: Experiment 2, Run 4, Time 1508.

8.4 Run 2-4

Following three rupture and repair episodes, at time 199 the entity forms into the same morphology encountered in run 1 (figure 9)—albeit, now rotated through 90°. This morphology again appears relatively robust, persisting from time 199 to time 1437, through 12 rupture and self-repair episodes. Between times 1438 and 1500 there are *four* additional ruptures. Of these, one is successfully repaired, but the overall damage to the membrane is now too great, and by time 1508 it has degenerated into a single curved chain as shown in figure 14. Again, it is then no longer possible to recover a closed membrane through any simple process of self repair.

8.5 Run 2-5

In this run there are two early ruptures at time 22 and 119, severely damaging the original membrane. An inward spiral forms. Coincidentally, another rupture allows the inward spiral to close forming a "new" membrane at time 245, with the morphology shown in figure 15. However, the cavity is now linear and thus does not afford *any* reaction sites for production of new `L` particles. This entity is therefore not capable of re-establishing the autopoietic reaction network.

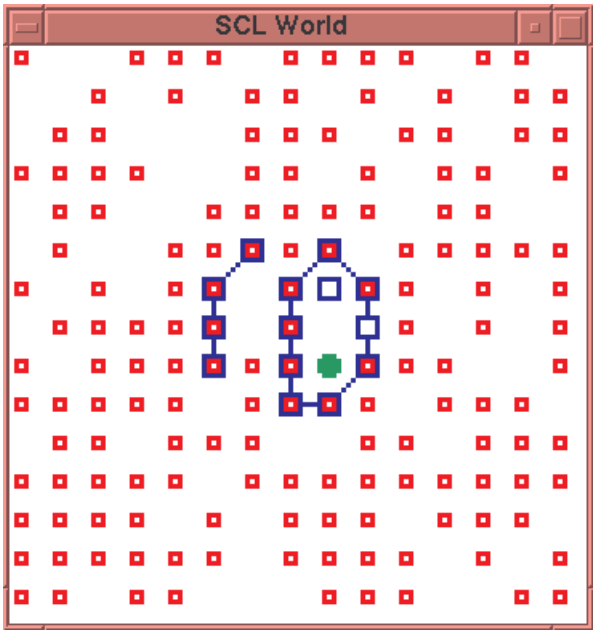

Figure 15: Experiment 2, Run 5, Time 245.

8.6 Discussion 2

There is substantial variation among the five runs comprising experiment 2. In runs 3 and 5, the initial entity effectively fails completely to establish a closed, autopoietic, reaction network. However, in runs 1, 2 and 4, an autopoietic reaction network *is* established, and a succession of successful repair episodes occurs. In runs 1 and 4 a morphology becomes established which is apparently particularly robust, persisting in each case for approximately 1000 time steps of the model.

The work reported here has not involved any extensive or comprehensive investigation of variations in the various reaction rate and mobility parameters available in the SCL model. It might well be possible to find combinations of parameter settings in which the establishment and maintenance of autopoietic reaction networks is more robust, and the autopoietic entities would thus be more stable and longer lived. However, the basic results of experiment 2 clearly show that this model *can* exhibit persistent, self repairing, autopoietic reaction networks, in the form originally described by Varela et al. [11]

> *Given that the only difference between experiment 1 and experiment 2 is the (re-)introduction of the chain-based bond inhibition interaction, it seems reasonable to conclude that this phenomenon of computational autopoiesis relies critically on the presence of this interaction.*

9 Conclusion

The primary conclusion from the work described here is that the original report of computational modelling of autopoiesis [11] was flawed, in that it failed to identify the chain-based bond inhibition interaction as being present and, indeed, as being an essential requirement for the achievement of the described autopoietic phenomena.

Given the lapse of time since the original publication, it is now difficult to suggest any definitive explanation as to how this interaction, actually present in the program code, came to be overlooked in the qualitative and algorithmic descriptions. However, as described elsewhere [10], the work was carried out during a difficult and turbulent time in Chile, and, further, there was a considerable time interval between the actual experiments and eventual publication. These factors together probably provide an adequate explanation for the oversight.

It should be emphasised that the substantive point of this paper is to correct the historical record. This is clearly relevant for anyone who wishes to reproduce, or *extend*, the phenomena of the original model. However, this correction does not add to, or modify, the original conceptual foundation of autopoiesis in any significant way.

In any case, the work described here also raises a more general question about the publication of computationally based ALife research. A key feature of scientific publication is that it should facilitate independent critical testing of whatever phenomena are presented. In this particular case, the defect in the original reporting (*not* a defect in the original model!) was uncovered only when a copy of the original program code was rediscovered by chance. At the time of the original publication, the technological facilities were not generally available to support easy distribution or access to accompanying code—but this is no longer the case. We would suggest therefore that as a general principle, published reports on computer models of ALife should be accompanied by access to the program code for the models on the World Wide Web.

Bare access to program code is, of course, of limited value in itself. Effective critical review would require that it should be "reasonably" feasible that others in the community be able to *execute* (and, indeed, modify) this code. This suggests a need for some degree of standardisation, where that is possible. The Swarm simulation system, with its open licensing for scientific research, offers a candidate platform for such standardisation. Indeed, this was a key reason for adopting Swarm in the development of the SCL system [6]. Our experience of using Swarm in this application suggests that it can provide a stable, efficient, and portable basis for wide dissemination of this kind of ALife research.

Acknowledgments

Initial work on re-implementing a computational model of autopoiesis was carried out by a former student of McMullin, Hyder Aswad. We are also grateful to Francoise Jullien and John Mingers for making available their independent re-implementations of the model of [11]. Ms. Marita Prandoni and Ms. Peggy Jones provided invaluable assistance in translating original documents from Spanish. All errors in the current paper remain, of course, our own responsibility.

This paper was written while McMullin was a visiting researcher (1996/97) with the Swarm Project Group at the Santa Fe Institute. This visit was made possible by generous support from Dublin City University and the Swarm Project. Swarm has been supported by Grant No. N00014-94-1-G014 from the Department of the Navy, Naval Research Laboratory, acting in cooperation with the Advanced Research Projects Agency. Swarm has also benefited from earlier support from The Carol O'Donnell Foundation and from Mr. and Mrs. Michael Grantham.

References

[1] Francois Jullien and Barry McMullin. FRJ's Simple Autopoiesis Program, 1995. Program source in Pascal, for MS-DOS platform.
`ftp://ftp.eeng.dcu.ie/pub/autonomy/frj_sap`

[2] Rosa Lizana. Modelo bidimensional de reproduccion como fractura de una unidad autopoietica. Master's thesis, Universidad De Chile, Facultad de Ciencias, Departamento de Biologia, 1981. In Spanish.

[3] Humberto R. Maturana and Francisco J. Varela. Autopoiesis: The organization of the living. In *Autopoiesis and Cognition: The Realization of the Living* [4], pages 59–138.

[4] Humberto R. Maturana and Francisco J. Varela. *Autopoiesis and Cognition: The Realization of the Living*, volume 42. D. Reidel Publishing Company, Dordrecht, Holland, 1980.

[5] Barry McMullin. Computational autopoiesis: The original algorithm. Working Paper `97-01-001`, Santa Fe Institute, Santa Fe, NM 87501, USA, January 1997.
`http://www.santafe.edu/sfi/publications/Working-Papers/97-01-001/`

[6] Barry McMullin. SCL: An artificial chemistry in *Swarm*. Working Paper `97-01-002`, Santa Fe Institute, Santa Fe, NM 87501, USA, January 1997.
`http://www.santafe.edu/sfi/publications/Working-Papers/97-01-002/`

[7] John Mingers. *Self-Producing Systems: Implications and Applications of Autopoiesis*. Plenum Publishing, New York, 1994.

[8] John Mingers and Barry McMullin. JM's Simple Autopoiesis Program, 1997. Program source in Pascal, for MS-DOS platform.
`ftp://ftp.eeng.dcu.ie/pub/autonomy/jm_sap`

[9] Francisco J. Varela. *Principles of Biological Autonomy*. North-Holland, New York, 1979.

[10] Francisco J. Varela. The early days of autopoiesis: Heinz and Chile. *Systems Research*, 13(3):407–416, 1996.

[11] Francisco J. Varela, Humberto R. Maturana, and R. Uribe. Autopoiesis: The organization of living systems, its characterization and a model. *BioSystems*, 5:187–196, 1974.

[12] Milan Zeleny. Self-organization of living systems: A formal model of autopoiesis. *International Journal of General Systems*, 4:13–28, 1977.

[13] Milan Zeleny, editor. *Autopoiesis: A Theory of Living Organization*, volume 3 of *North Holland Series in General Systems Research*. North Holland, New York, 1981.

Revision History

The first draft of this paper was published in February 1997, simultaneously as Working Paper No. `97-02-012` of the Santa Fe Institute and as Technical Report No. `bmcm9703` of the Dublin City University ALife Laboratory. This is a revised version, presented at ECAL-97, Brighton, UK July 1997, proceedings published by The MIT Press, edited by Phil Husbands and Inman Harvey.

Retrieval

The resources comprising this paper are electronically retrievable, in various formats, via the World Wide Web, from:

`http://www.eeng.dcu.ie/~alife/bmcm-ecal97/`

Copyright