Towards the Implementation of Evolving Autopoietic Artificial Agents

Barry McMullin
Barry.McMullin@rince.ie

Dominique Groß
Dominique.Gross@rince.ie

Presented at
6th European Conference on Artificial Life ECAL 2001
(September 10-14, 2001, University of Economics, Prague, Czech Republic.)

Dublin City University
Research Institute for Networks and Communications Engineering
Artificial Life Laboratory
Technical Report Number: bmcm-ecal-2001

Abstract

We report modifications to the SCL model(McMullin and Varela, 1997), an artificial chemistry in Swarm designed to support autopoietic agents. The aim of these modifications is to improve on the longevity of the agents and to implement growth. We demonstrate by means of two simulation runs that the improvements indeed have the desired effects.

1 Background: SCL

A computational model of autopoiesis was originally described by Varela et al. (1974). This was re-implemented by McMullin in Swarm[1] as the SCL system (McMullin and Varela, 1997; McMullin, 1997). SCL contains three different types of particles: substrate (S), catalyst (C) and links (L). In this model an autopoietic agent or organisation consists of a closed (self-repairing) chain of doubly bonded links which encloses a C particle. In order to qualify as an interesting instance of autopoiesis, its longevity has to exceed that of its components (the links) which are designed to be unstable. In the experiments presented in (McMullin and Varela, 1997) the minimal autopoietic organisation could be sustained for up to approximately 1800 time-steps, but was eventually irreparably corrupted by some decaying links which could not be replaced quickly enough.

2 New Features: SCL-GRO

We report a modified version of SCL, called SCL-GRO[2]. This aims at an improvement of the self-repair mechanism of the autopoietic agents (without trivializing the problem by, for example, turning off the decay of links), and the introduction of growth. The motivation for these improvements derives from an overall long-term goal of realising evolution of artificial autopoietic agents. The following additional features have been implemented.

- **Affinity:** Links have been given an affinity for chains of bonded links (closed or open); if a link is in the neighborhood of a chain it may drift along the chain but will not move away from it.
- **Smart repair:** Free links actively scan for decaying links in a chain, i.e., links which will decay as soon as there are enough free lattice sites around them to do so. If they find them, they replace these links.
- **Improved mobility of links:** Whereas the L particles in the original implementation were immobile once they were bonded, we allow bonded links to move as well (subject to the bond constraints).
- **Displace motion:** If there is a free link in the neighborhood of a doubly bonded link, then with some user-specified probability it triggers the breaking of one of the bonds of the bonded link, and "pushes" this (now singly-bonded) link in the direction away from the free link; the latter then moves into the newly vacated lattice position. In this configuration the free link can be integrated into the chain through a chain splicing event, which also restores the closure of the membrane.

In earlier versions of SCL a special feature, so-called bond inhibition had to be introduced in order to ensure that the free links inside the membrane did not bond and thus become unavailable for repair of the membrane. With growth now enabled this measure became ineffective in SCL-GRO; bond inhibition only works in the immediate neighborhood of the membrane (thus preserving the principle of locality!). With a larger membrane most of the free links inside will not be in such a neighborhood. This problem has been resolved by differentiating the bonding interaction into three cases, depending on the prior bonding states of the link particles. The distinct cases are for bonding between two free links (chain initiation), between a free link and a single bonded link (chain extension), and between two singly bonded links, possibly via an additional free link (chain splicing). The probability for the first of these (chain initiation) is then set to zero. This ensures that free links will not spontaneously bond to each other, but can still bond to pre-existing chains (in order to maintain—and grow—the membrane).

3 Results

We will now shortly summarise results from selected, but typical runs in SCL-GRO. The results given below hold approximately for all experiments we have performed.

Longevity: We turned off growth (by turning off the displace motion) in order to compare the longevity of the autopoietic agents in SCL-GRO with those in the original SCL. We seeded the SCL-GRO world with a medium sized autopoietic agent.

[1] `http://www.swarm.org/`

[2] `http://www.eeng.dcu.ie/~alife/src/scl-gro/`

In the course of the simulation the membrane undergoes some change and also inserts a few additional links. Despite continuous ruptures of the membranes, in the specific experiment described here the self-repair mechanism of the autopoietic agent remains effective until about time-step 5400, when a terminal rupture happens.

Taking into account that the probability of decay per time-step of a link is set to 0.002 and that the initial configuration consisted of 42 links, then the expected life-time of the membrane—in the absence of repair—is about 12 time steps. This is to be contrasted with the actual life-time of the membrane of over 5000 which is more than 400 times longer. The improvement over the original SCL model becomes clear if one compares this result to the analogous number of the earlier reported experiments (McMullin and Varela, 1997) where the autopoietic agent lived for about 22 times its expected (unrepaired) life-time (decay probability of 0.001, 12 links in the membrane and an absolute life time of no more than 1800 time-steps). It is clear that in experimental settings like this, quantitative results are only of limited value in the absence of more extensive coverage of the parameter space. However, there is a factor of 20 between the two results, which has some suggestive value.

Growth: Let us now turn on the displace motion and start the experiment again with an initial configuration, consisting of a small membrane enclosing a catalyst and four free links such as in (McMullin and Varela, 1997). Starting the simulation, the membrane immediately starts to deform into irregular shapes. Occasionally new links are inserted into the chain, which leads to a growth process. Inside the membrane the catalyst continuously produces new free links, which tend to drift outwards, ready to repair a potential fracture of the membrane. Once the agent has attained a certain size three functionally distinct parts or layers can be discerned. First there is the most inner layer, the "reactor", characterized through the presence of the catalyst and substrate which has drifted through the chain of links into the interior of the membrane. The middle-layer consists of free links, which were produced in the "reactor" and are "waiting" to repair a potential rupture. The third layer is the "membrane", which has already been described above. Once the membrane attains a certain size, it usually takes a rectangular shape; this is apparently partly due to the affinity between free links and the membrane and the consequent inhibition of its inward motion, and partly a reflection of the underlying lattice geometry.

Dependent on the chosen disintegration probability of the links, the chain sooner or later suffers a terminal rupture. Especially if the chain is already quite large, then locally there might not be enough repair material, i.e., free links, available to effectively perform a repair operation. This effect seems to be a cause for the terminal rupture in this run, where one clearly sees the lack of free links in the neighborhood of where the membrane is broken (not shown). Under this setting the terminal rupture of the membrane occurs at time step 1168. Various experiments (not reported here) suggest that the coarse qualitative features of the behaviour of the model remain invariant under a significant variation of the link-decay probability.

4 Discussion and Conclusion

The SCL-GRO model, as described in this article, clearly adds to the qualitative phenomenology of the original version of SCL in that it considerably extends the longevity of the autopoietic agents and supports their systematic growth. These might be useful steps in the direction of realising reproduction of autopoietic agents by growth and fission.

References

McMullin, B. (1997), Modelling Autopoesis: Harder than it Might Seem!, Technical Report `bmcm9704`, School of Electronic Engineering, Dublin City University, Dublin 9, Ireland.
`http://www.eeng.dcu.ie/~alife/bmcm9704/`

McMullin, B. and Varela, F. J. (1997), Rediscovering Computational Autopoiesis, *in* P. Husbands and I. Harvey, eds, 'Proceeedings of the Fourth European Conference on Artificial Life', Series: Complex Adaptive Systems, MIT Press, Cambridge.
`http://www.eeng.dcu.ie/~alife/bmcm-ecal97/`

Varela, F. J., Maturana, H. R. and Uribe, R. (1974), 'Autopoiesis: The Organization of Living Systems, its Characterization and a Model', *BioSystems* **5**, 187–196.

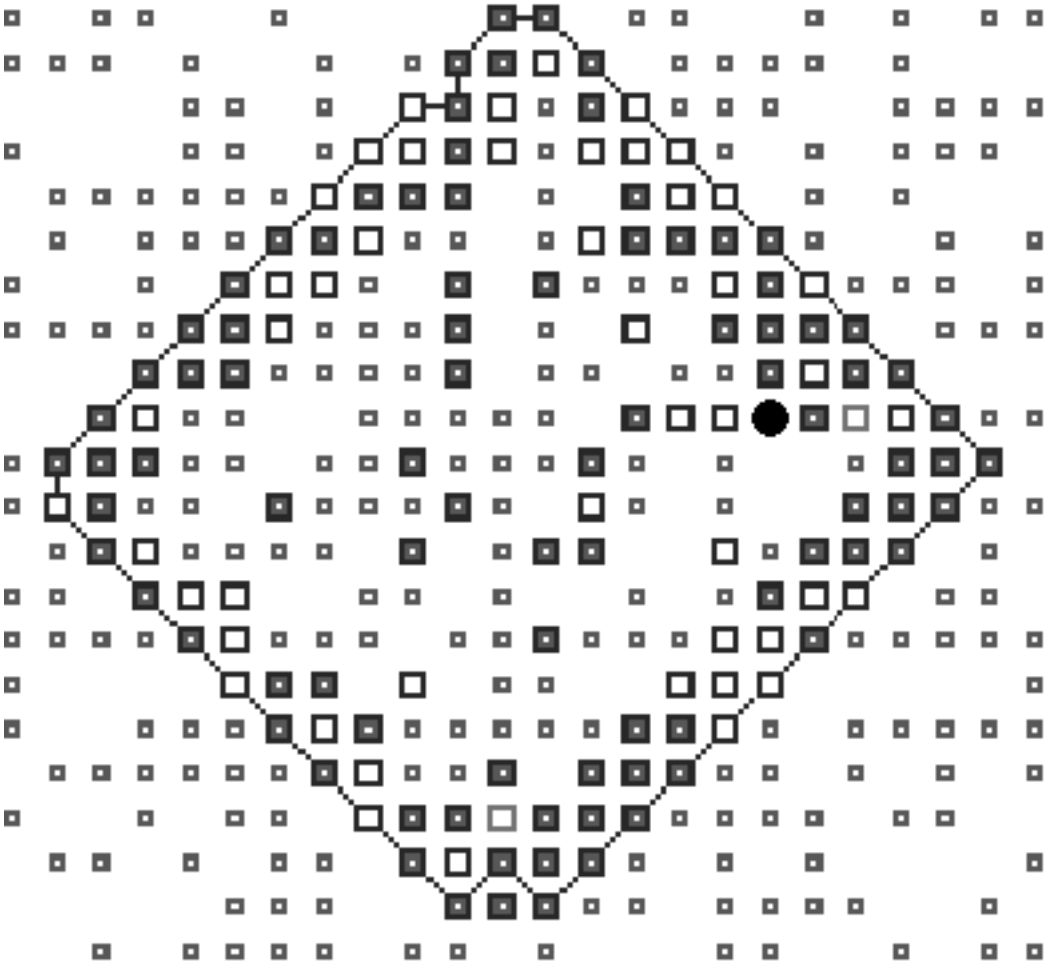

Figure 1: This is a screenshot from the experiment with growth turned off: after 3130 time-steps the autopoietic agent is still intact. The small squares are S particles; the large and large dark squares are L and L^+ particles (without and with absorbed substrates respectively); the filled circle is the catalyst. Three distinct layers of the agent are clearly visible.