

# Bonding as an Emergent Phenomenon in an Abstract Artificial Chemistry

Dominique Chu<sup>1</sup> and Rune Vabø<sup>2</sup>

<sup>1</sup>School of Computer Science, University of Birmingham, UK

<sup>2</sup>Institute for Marine Research, Bergen, Norway

hsvdg@uib.no

## Abstract

In this article we describe an implementation of a very simple toy artificial chemistry originally proposed as a “Gedanken” model by Groß and McMullin. They predicted that this model would have a number of interesting properties including emergent bonding, and top-down constraints on bond-stability. We found that not all of the original claims could actually be verified.

## The Problem of Novelty in Artificial Life Systems

In a contribution to the Artificial Life 8 conference Groß and McMullin (Groß and McMullin, 2003) discussed the problem of novelty emerging in Artificial Life particularly focusing on agent-based models (ABM) (Casti, 1997; Holland, 1995). While acknowledging the usefulness of ABMs in general, they pointed out that ABMs are not particularly good at creating what they call perpetual novelty. As one of the main limitations for the emergence of novelty they identified what they called “meta-rules,” that is hard-coded procedures that define the scheduling of events in the model, the order in which agents are called, mutated, resources are supplied to the environment and so on. While necessary for the simulation those meta-rules are “hard-coded” and as such significantly constrain the possible behaviors displayed by the model.

Groß and McMullin discussed several ways to avoid the limitations of meta-rules. One possibility is to evolve/modify the meta-rules themselves at run time. Unfortunately this only looks like a solution, as this would require the specification of meta-meta-rules that specify in which way the meta-rules are to be changed. Given the insight that meta-rules are indispensable, Groß and McMullin decided that it would be better to think about ways to minimize the limitations posed by meta-rules, rather than to attempt to avoid them altogether.

In their article Groß and McMullin focused the discussion on the emergence of novelty in agent-based artificial chemistries (Dittrich et al., 2001). In such artificial chemistries the behavior of particles is determined in a

bottom-up fashion, meaning that the rules of behavior function at the level of the individual agent. While this is a plausible way to implement artificial chemistries, it is not entirely unproblematic from the point of view of emerging novelty. To see this consider the usual way to represent bonding in such models. A common way to do this is to implement an explicit bonding rule that specifies under which conditions in the simulation bonds are formed and break. Groß and McMullin see inherent limitations in this method to define bonding.

The “new” high level particle formed as a result of this kind of bonding is little more than a collection of its primitive components plus a few motion constraints which have been preconceived and explicitly pre-programmed. In this sense bonds—their formation behavior, rupture—are completely pre-specified[...]. [...]This seems distinctively unlike the formation of new molecular species in nature and it also restricts the potential for novelty creation in the model. (Groß and McMullin, 2003)

Groß and McMullin call this approach to bonding “bottom-up,” because whether or not a bond will be formed is determined exclusively by the inner and outer states of the particles that participate in the bonding. This is opposed to their suggested “top-down” approach to bonding, where the bond is an emergent result of the environmental conditions of the particles, while there are no explicit rules for bonding given. Top-down bonding, they conjecture, will lead to artificial chemistries with a much higher potential for emergent novelty.

In their paper they describe a simple “Gedanken” artificial chemistry as an example for a system that has top-down constraints on bonds. Admitting that this system is probably too simple to display perpetual novelty, they list a few key-features even this very simple model will possess:

- **Emergent Bonding:** Particles can form bonds with one another, although no explicit rules for bonding is specified.

Copyrighted Material




Figure 1: The 2 types of particles. Different gray-scales indicate the field-strength. The right side shows an *A*-particle, on the left is a *B*-particle. The particles themselves are not visible in this picture, but located in the center of the fields.

- **Variety:** The model supports a variety of compound configurations (isomers).
- **Configuration Dependence (Top Down Constraint):** Which particles of a compound can bond with yet another particle depends on specifics of the compound itself.

In this article we report some experimental results obtained from an implementation of this “Gedanken” artificial chemistry. Particularly, we are trying to establish whether or not the phenomena predicted by Groß and McMullin can indeed be observed.

### Description of the Model

The model is implemented in Java using the Repast (Repast, 2004) simulation libraries. Every particle type is characterized by the particular “force-field” it induces in its neighborhood (see figure 1 for an example of an agent type). The details of this field are defined by the user at run-time and remain fixed for each particle. In principle the model allows an arbitrary amount of different particle types to be defined at run time; for all simulations reported here we consider only 2 types. Each agent has 6 momentum components (one component for each possible direction in a hexagonal grid; see below) and 2 variables describing its position. Position and momentum variables are updated at each time step. The momentum components can take a maximal value of 100 and a minimal value of 0.

The environment of the model is a 2d hexagonal discrete grid (with periodic boundary conditions). Each grid site is occupied by at most one particle at a given time. Every site is associated with a particular value (either zero, or a negative or positive integer) describing the field induced by an agent. If the force fields of agents overlap on one grid site, then the net value of the grid site is the sum of all contributions.

The movement rules of the model are highly schematic and no attempt is made to be physically realistic or even plausible: At each time step particles either stay where they are or move to a neighboring site (so maximum speed is one cell per time-step). In order to determine the direction of movement, agents check the field values at all 6 neighboring sites. If agents move in a direction

momentum-component is updated by adding the difference between the field values after and before movement to the momentum component. For example, if an agent had a momentum of one in the north direction, his current site has a field value of  $-2$  and the new site has a value of  $-1$ , then the updated component of the north component of the momentum is 2.

At every time step agents attempt to perform a movement that maximizes their momentum in the following way: In order to determine the direction of movement agents check the field values at all neighboring sites and calculate the momentum they would have after moving to this site. The direction chosen for the actual movement is the direction in which the momentum component *after* movement is maximal. If several momentum components are equal, a random one is chosen among the maximal ones. If all momentum components (before *or* after movement) are zero, no movement is performed. Particles do not feel their own force field.

An example might clarify this movement rule. Consider a particle that is on a site with a force-field of 2. Its north-component of the momentum is 5 while all other momentum components are 1. All neighboring field values surrounding it are 0 except for the north field value, which has a value of  $-1$ . Thus if the particle choose the northern direction its momentum component after the movement will be  $5-3=2$ . In all other movement directions the momentum after (and before) the movement will be 1. Since the north-direction will be maximal after the movement into the north direction, the north direction is chosen for movement.

Agents are subject to some random kicks which change (with a user defined probability) the value of a randomly selected momentum component by  $\pm 1$ . In absence of any fields, the agents thus perform a random walk over the grid driven by the random kicks.

### Emergent Bonding

The phenomena that can be observed in the model strongly depend on the choices of the force fields of the agents. We experimented with a number of values, but here we will report only results obtained with agents of the following type: Agent  $A = (-40, -35, -30, 15, 15, -30)$  and  $B = (-20, 0, 0, 0, 0, -20)$ , where the numbers in parenthesis denote the field values an agent induces in its first, second, third, ... neighborhood; the field values of those agents are depicted in figure 1. (At the moment the model supports up to 8 neighborhoods, but in the simulations reported here we will only use 6). The qualitatively results do not depend on the particle types as long as the particle types do support stable bonding; there are of course also particle types that show no interesting behavior (for example particles with no field). As anticipated by Groß and McMullin, the system does support emergent bonding between particles. Bonding between *A*-particles can be achieved by placing them in each others fourth neighborhood (the fifth




Figure 2: A closed chain of *A*-particles in its initial state.

one also works). The particles are then trapped between two “repulsive” fields effectively preventing them from coming too close to each other and from leaving each others’ neighborhood. The same arrangement, but with more flexibility for relative movement is achieved by placing *B*-particles into each others second, third, fourth or fifth neighborhood. We also found many other particle types that support bonding. A necessary condition for emergent bonding is that the particles have an inner and outer, strongly repulsive field—and sandwiched between those—a neutral or positive field.

Both *A* and *B*-particles can form strong and stable bonds with one another; note that due to the random kicks the particles receive even the most stable configurations might rupture.

### Chains

Both types of particles can not only engage in pairwise bonding, but can also be arranged into linear chains of the form  $\dots - A - A - A - A - A - A - \dots$ . For all practical purposes those chains can be made arbitrarily long. In the case of *A*-particles chains can be made as follows: Place two *A* particles into the fourth neighborhood on opposite sides of another *A*-particle to make an initial short chain. The chain can then be extended by adding *A* particles on either side of the initial chain. Simulations show that those chains have very good stability properties and can even be arranged into closed chains that typically remain bonded over very long times. Figure 2 shows an example of such a closed chain.

While chains (closed or open) are very stable, isomers might have very different stability properties. As an exam-

ple consider the linear chain of 12 *A*-particles. Like all linear chains of *A*-particles this configuration is stable. One possible isomer of this chain is a cross-shaped configuration, consisting of a chain of length 6 and 2 side chains of size 3 attached to the main chain between particle 3 and 4. We simulated this configuration 10 times and found that only in 2 of those runs the cross remained intact during the first 200 time steps. In the eight remaining runs one or several of the side chains broke away. Attaching only one side chain (such that the overall configuration looks like a “T”) improves the stability. Again we ran 10 simulations for the T-configuration and 10 for the inverted T-configuration (those cases are actually equivalent). In our experiments we observed that the configuration remained stable in 10 out of 10 runs and in 8 out of 10 runs respectively.

A likely reason for the instability of the cross configuration is that particles 3 and 4 in the main chain each have to bond to 4 other particles. This interpretation is suggested by separate experiments with only five particles that are arranged in a linear chain of 3 particles and one particle attached to each side in the middle of the chain. In this configuration the central particle has to bond to 4 other particles which does not seem to be stable. We conclude that *A*-particles can form up to three bonds with other *A*-particles. If a particle has more than three bonds, then all bonds tend to be unstable, which leads to the disintegration of the compound.

The strength of the bonds depend on the context: Pairs of bonded *B* particles in isolation are extremely stable configurations; so are chains of *A*-particles. We performed 10 simulations with an *A*-particle chain of length 6 and (spatially separated) a pair of bonded *B*-particles. As expected both did not disintegrate in 10 out of 10 runs. However, if the bonded *B* particles are both placed into the 4-neighborhood of the middle particle of the *A* chain, then the compound did not remain stable. This configuration dissolved in 10 out of 10 simulation runs. It is particularly interesting to note that in all those cases the bonds between the *B* particles broke. Thus while in isolation inter-*B* bonds are extremely stable, within a compound of *A* particles those bonds are extremely unstable.

### Encapsulation

While performing our experiments we discovered another property that quite naturally emerges from certain configurations in the models, however without having been predicted by Groß and McMullin: Encapsulation of particles into larger wholes. Loosely, by encapsulation we mean that a compound of interacting particles starts to behave as a whole. There are two aspects to this phenomenon.

Firstly, the dynamics of compounds of particles can be understood to operate on two time-scales. Structures of bonded particles (as for example chains or closed chains) change their internal configuration on a fast time scale. On a




Figure 3: Two rings can collide with one another, without dissolving. Note that the model has periodic boundary conditions, so a part of the left chain is on the right side of the picture.

much slower time-scale the compound as a whole will drift through the environment; this overall movement is of course a result of the individual particles' motion, but at the same time imposes a constraint on it. This is best explained by an example: Consider a closed chain of *A*-particles. The chain exists because the constituent particles are trapped between the repellent force-fields of their neighbors. The individual particles see themselves surrounded by repulsive fields; consequently, the number of available degrees of freedom available to them is reduced. While free particles can move in any direction, the bonded ones are restricted to stay within the attractive fields of their neighbors. If a particle is bonded then its possible degrees of freedom will very much depend on where its neighbors move; at the same time, the neighbors' movement is itself limited by their neighbors' movement. In this sense the local movement of one particle in a compound is influenced by the global state of the compound as a whole.

Secondly, compounds remain intact even if they collide with other objects (i.e. other compounds or particles; see figure 3). Again this can be illustrated by a closed *A*-particle chain. If we put two or more chains in the environment, then those chains might collide with one another, however without being destroyed. The chains will not exchange particles with one another and particles of the chain will not form bonds with one another (which they would do in the free state). This second aspect of encapsulation is not only important when two chains interact, but also when various parts of the chain interact with one another. In simulations chains tend to fold onto themselves (see figure 4), bringing different parts of the chain into contact. Chains are only stable structures because the various parts of the chain do not react with one another when brought into contact.

In the particular case of *A*-particle chains (open and closed) the encapsulation property is a consequence of the outer negative field. Because of the particular choices of the fields of *A*-particles, properly bonded chains will be surrounded by




Figure 4: Chains tend to fold onto themselves. To the left we see a folded chain. To the right another folded chain without the grid or the fields represented. Those and similar chains have life times of the order of hundred-thousand time-steps.

a strongly repulsive field at each side of the chain. This outer field is particularly strong because it is a result of the overlay of the repulsive fields of the neighbors in the chain. This functions as an efficient barrier against most incoming particles and also allows chains to fold onto themselves without being dissolved by themselves.

## Discussion

In the present model, there was no explicit rule for bonding, yet a simple chemistry with characteristic bonding properties (i.e. arity and stability of bonds) did emerge. It is instructive to contrast bonding in the present model with artificial chemistries that rely on explicitly defined bonding rules, such as McMullin's SCL model (B. McMullin, 1997). In the SCL model (as in most similar artificial chemistries) establishing bonds between particles is a rather involved process requiring a degree of coordination between particles: At every time step for each particle the following things need to be checked:

- Check whether the particle is already engaging in the maximum number of allowed bonds.
- Check whether or not a suitable bonding partner is in the neighborhood
- Establish communication with potential bonding partners in order to agree on whether or not to form a bond.
- Finally, keep a record of all bonding partners in order to
  - ... coordinate motion
  - ... communicate possible breakage of the bond.

This contrasts well with the present model where all an agent needs to do in order to calculate its own movement is to look at the force fields in its immediate neighborhood. The particles themselves do not "know" whether or not they are bonded. In fact there is no explicit notion of bonding in

the model. Yet still this model shows emergent “chemical” properties similar to the bonding of particles in the SCL model.

We could also verify that the model supports isomeric variants of compounds. Particles can not only form chains, but a variety of different configurations. When it comes to top-down constraints, the evidence is more ambiguous. In order to show that there are clear top-down constraints, it would be necessary to show that, all local conditions equal, an effect or interaction is only possible given some global configurations but not others.

One could interpret the loss of stability of the cross-shaped configurations relative to the “T” shaped configurations as a top-down constraint in this sense. However, closer examination shows that this difference in stability is a result of the fact that A-particles do not support more than 3 stable bonds, thus a local effect. The next best evidence is that of the disintegration of the, otherwise stable, B – B chain, when placed into the neighborhood of a 6 long A-particle chain. In order to really prove the point of the configuration dependence of the bond stability, it would have been necessary to show a configuration in which the B – B compound is stable despite being in the immediate neighborhood of a A-particle chain. We could not find such a configuration. We therefore conclude that the case for top-down constraints is at best weak.

A very interesting property that is supported by the model, but has not been predicted by Groß and McMullin, is encapsulation: Some compounds naturally form a unit in that they can interact with other particles or compounds (for example by colliding) yet retain their structural integrity. It should be noted that the property of encapsulation *per se* is very common in artificial chemistries. As an example one might think of the membrane in McMullin’s SCL model (B. McMullin, 1997; McMullin and Groß, 2001); this membrane is a structure that is very similar to the closed chain in the present model. However, the crucial difference is that in the SCL model the encapsulation property is a direct result of the bonding rules of the membrane particles. Membrane particles can bond with at most 2 other particles; once they have two active bonds, as all particles in a closed chain, they will not engage in any further interactions with other particles. In the SCL-model encapsulation was thus explicitly programmed into the system, while in the present model encapsulation is an emergent result of the way in which particles interact.

## Conclusion

Groß and McMullin conjectured that their proposed artificial chemistry would have emergent bonding, support isomers and demonstrate top-down constraints. We could readily verify the first two properties, but found only weak evidence for the third one. However, we also found an effect that has not been predicted by Groß and McMullin, namely

emergent encapsulation.

We think that the modeling strategy proposed by Groß and McMullin is a fruitful one and can be a useful complement to existing artificial chemistries that operate on a higher level. Particularly, top-down constraints would indeed be an interesting property to have in artificial chemistries. The modification of local properties by non-local configurations seems to be an important property of real chemistry, particularly in the realm of bio-chemical processes. It seems, though, that the artificial chemistry proposed by Groß and McMullin does not display top-down constraints. Whether or not a simple extension to the model can overcome this limitation is unclear at present.

For future extensions of the model it would be desirable to have an even richer phenomenology without substantially increasing the complexity of the basic rules of the model. Particularly it would be desirable to have emergent catalysts. In the current model we have only observed rudimentary forms of catalytic activity when two parts of a molecule push two unbounded particles together. However, this is a very primitive and unreliable effect. Extensions will most likely require continuous space/higher dimensions. Work in this direction is currently being undertaken.

## References

- B. McMullin (1997). *Scl: An artificial chemistry in swarm*. Dublin City University, School of Electronic Engineering, Technical Report: bmcm9702 and Santa Fe Institute Working Paper: 97-01-002.
- Casti, J. (1997). *Would-Be Worlds*. John Wiley & Sons, New York.
- Dittrich, P., Ziegler, J., and Banzhaf, W. (2001). Artificial Chemistries—A Review. *Artificial Life*, 7:225–275.
- Groß, D. and McMullin, B. (2003). The creation of novelty in artificial chemistries. In *Proceedings of the eighth international conference on artificial life*, pages 400–409. MIT Press.
- Holland, J. (1995). *Hidden Order*. Addison-Wesley, Reading.
- McMullin, B. and Groß, D. (2001). Towards the implementation of evolving autopoietic artificial agents. In *Proceedings of the 6th European Conference on Advances in Artificial Life*, pages 440–443. Springer-Verlag.
- Repast (2004). <http://repast.sourceforge.net/>.