RBN-World: The Hunt for a Rich AChem

Adam Faulconbridge, Susan Stepney, Julian F. Miller and Leo Caves

YCCSA, University of York, Heslington, YO10 5DD, UK
asf500@yccsa.org

Abstract

An Artificial Chemistry (AChem) is a set of components and interactions that result in a composable system. Ideally, the system is rich, and results in rich higher-order emergent properties. We present a methodology for discovering interesting AChems through a series of tests that probe elementary low-level properties. In doing so, we assume that these elementary properties are a necessary, but not sufficient, basis for higher-order emergent properties, such as autocatalytic sets and hypercycles. The test strategy is applied to RBN-World, a sub-symbolic chemistry. This results in identifying a number of new and interesting RBN-World chemistries that appear richer than our original parameterisation.

Introduction

One approach towards the goal of Artificial Life (ALife) has been Artificial Chemistry (AChem), particularly for the origins of life. Unlike many ALife approaches, life-like properties are not explicitly designed in, but emerge from the dynamics of the system. AChems have been applied in other contexts [16, 12] however here we focus on their role as approach to the study of composable systems capable of exhibiting rich higher-order emergent behaviour.

In its most basic form, an AChem is a collection of molecules and reactions that describe transformations between groups of molecules, and an algorithm which determines how the reactions are applied over time [2]. There are a large number of possible AChem designs (relating to the nature of the components, interactions and reactions) each with a potentially large parameter space. Moreover, some examples of emergent systems (Boids [14], Conway's Game of Life [8], etc) only exhibit emergence at a small subset of possible parameters. This motivates the need to develop strategies to search the parameter spaces of AChems to find those regions that exhibit rich emergence.

Here we describe a set of tests suitable for any AChem and apply those tests to filtering 200 alternatives of an AChem — RBN-World [7].

Desired high-level properties

Determining how to evaluate different AChems is a difficult task. The overall goal when developing an AChem for ALife is an emergent system capable of open-ended evolution. The metric for this is unclear; some suggestions include Chemical Organization Theory [1] and Granger causality [15]; however, searching for interesting chemistries using metrics such as these would not be computationally tractable over the large search space of alternative chemistries. Several mid-level properties have been previously suggested as important in the emergence of rich evolutionary characteristics; in the context of artificial chemistry, three of particular relevance are *autocatalytic sets* [11], *hypercycles* [4, 6, 5] and *heteropolymers* or co-polymers [13]. Desirable characteristics of artificial chemistries have been suggested before [17] however, these are design specifications rather than emergent properties.

Autocatalytic Sets An autocatalyst is a molecular species that catalyses its own production. Autocatalytic sets are two or more molecular species where one or more reactions producing each member of the set is catalysed by itself or another member of the set [11]. The members of an autocatalytic set may be, but do not have to be, autocatalysts themselves. In addition, autocatalytic sets may overlap with individual molecular species belonging to more than one set. Autocatalytic sets are thought to be important to the emergence of life because of their characteristic growth; as long as substrate is available, the members of an autocatalytic set will continue to increase in concentration.

Hypercycles Hypercycles are a collection of coupled self-replicative units and thought to be important as a higher-order organization [4, 6, 5] — many biochemical metabolic processes are hypercycles for example.

Heteropolymers Polymers are molecules composed of repeated subunits. Heteropolymers are molecules composed of non-identical subunits, such as DNA or proteins which both have a repeating backbone structure with different side-groups attached to it. The important feature of heteropolymers is their capacity for information storage encoded into the ordering of the subunits.

Desired low-level properties

Searching for autocatalytic sets, hypercycles and/or heteropolymers would be an useful step towards finding artificial chemistries with sufficiently rich emergent properties. However, this is still too computationally intensive to be useful as an initial step. We suggest that the space of possible chemistries can be first reduced by selecting for specific features thought to be required by the higher goal; towards that end, the features being examined should be low-level and computationally tractable.

In order to hunt for rich AChems, we specify tests for low-level properties that we believe are necessary (but possibly insufficient) 'stepping stones' to higher-order emergent behaviour. The tests can be structured, and as a result chemistries that fail the lowest-level tests are not considered for the intermediate tests thus allowing subsequent searches to focus on interesting subspaces.

Synthesis is the formation of bonds is the lowest level property possible; however it is important not only that synthesis can occur in an AChem, but also that too much synthesis does not occur. If every molecule can bond with every other molecule, the chemistry is trivial and will not support rich dynamic higher-level properties.

Self-Synthesis is bonding between two identical atoms or molecules. As with synthesis, this is important for the formation of larger molecular structures but also should be able to occur between any two identical atoms/molecules.

Decomposition should also be possible, but not universal, within the AChem. Without the breakdown of larger molecules, many conceivable mechanisms for higher-level properties become impossible and the system may reach a steady state once all raw materials have been consumed.

Substitution is a potential emergent behaviour given that a particular AChem exhibits synthesis and decomposition. While arguably not important in itself, substitution represents the potential for relationships between more than one or two molecules.

Catalysis is another property of interest. We define catalysis as a series of reactions that do not consume the catalyst, yet the overall reaction would be slower (or not occur at all) without it.

RBN-World: Overview

RBN-World [7] is an AChem framework combining random Boolean networks (RBNs) [9, 10, 3] via bonding sites.

RBNs consist of n nodes synchronously updated in discrete timesteps. Each node in the RBN has a Boolean state, inputs from k nodes, and a Boolean function that maps the state of inputs to an updated state at the next timestep. The state of an RBN is the collection of states of all its nodes. All RBNs have cyclic behaviour, returning to a previous state after sufficient number (usually small) of timesteps.

To use RBNs in a chemistry some modifications have been made — we refer to the modified RBNs as bRBNs (bonding random Boolean networks). Important aspects of these are:

Atoms Within each RBN, there are one or more *bonding sites* (b); these are additional nodes that provide inputs to ordinary nodes. Bonding sites do not have any inputs, instead their state is determined by whether they are "bonded" or "unbonded".

Bonds A *bond* links two bRBNs, and there can be multiple bonds between the same pair of bRBNs. Each bond requires one "unbonded" site within each of the bRBN pair to become "bonded", and each "bonded" site is associated with only one bond.

Bonds are formed as a consequence of *reactions* when specific criteria are met. If a bond is not formed by a reaction, it is attempted again with any higher-level structures (e.g. molecules) that the pair of bRBNs are part of. This iteration of attempting bonding and re-trying for higher-level structures continues until either a bond is formed or there are no more higher structures.

Molecules bRBNs that are linked by bonds can be expressed as a composite bRBN. The composite bRBN's inputs and functions are the component bRBNs with inputs from "bonded" sites are replaced with direct inputs from the other "bonded" node. Non-composed bRBNs are *RBN-atoms*, and a composite bRBN is a *RBN-molecule*. A composite bRBN that is part of a larger composite structure is a *functional group* (by analogy with functional groups in chemistry, such as the amine group). RBN-molecules undergo reactions and form bonds in the same manner as RBN-atoms to make further higher-level composite structures. Note that an internal RBN node can be in different Boolean states at different levels of the structural hierarchy.

Bonding Consequences Forming a bond has two direct consequences:

1. The process of bonding changes a bonding site in each linked bRBN from "unbonded" to "bonded". This changes one input to one node, which can potentially lead to a change in the dynamic behaviour of the Boolean network.
2. The bRBNs linked by the bond form a new higher-level composite bRBN. If one of the participants of the bond was already a component in another bRBN, then the composite structures are combined into a larger composite bRBN.

In addition to the direct consequences, there are potential indirect consequences as well. The formation of a bond may

change the dynamics of either bRBN, which may cause the bonding requirements to be violated. When bonding criteria are no longer valid, bonds *break* and the associated bonding sites reverts to "unbonded". This also alters any higher-level composite structures, collapsing them if they are not distinct from their lower-level components.

Due to the combinatoric nature of Boolean networks, there are a vast number of possible bRBN-atoms. However, only a small subset will lead to the emergence of sufficiently rich properties – most of the chemistry that underpins life is consists of a restricted number of elements: Carbon, Hydrogen, Nitrogen, and Oxygen. Finding analogues of such highly composable elements (and implicitly their interactions) in a particular AChem is our task here.

RBN-World: Alternative Chemistries

During the development of RBN-World, it became clear that a number of modelling decisions had to be made based on limited information; for example, the size of the bRBNs and bonding criteria between them. Also, pragmatically, a number of choices and assumptions were made without explicit consideration of alternatives. These choices may have impact on the emergent properties of the AChem.

To investigate the alternative chemistries, some of the choices have been explicitly defined in order to determine their effect upon the resulting AChem. It is worth noting that the decisions around which alternatives to study have themselves been made based on limited information from preliminary experiments and exploratory ideas.

Four different categories of alternatives have been identified with multiple options within those categories. As well as these separate alternatives, combinations of alternatives from different categories can also be investigated.

Bonding Property

One of the novel aspects of RBN-World is the use of properties of the underlying dynamical system to determine bonding. However, it is not clear which property would be most suitable and what effect different properties might have. Several alternatives are considered here, each with distinct distributions. See tables 1 and 2 for summary and example.

Cyclelength (c) is the number of different states the bRBN passes through between repeats. Cyclelength has a large but bounded asymmetric discrete distribution of values, with a median of approximately $\sqrt{n}$ for small values of k [9].

Flashing counts how many Boolean nodes change state during the cycle. RBNs typically have a 'frozen core' of static Boolean nodes, and flashing is the inverse of this. This can expressed as follows; let a state of 'true' have a value of 1 and a state of 'false' have a value of -1; N be the set of nodes in the bRBN; $s_{i,j}$ be the state of the i^{th} node at the j^{th} state of the repeating cycle. Then:

$$N_{i_{\text{flashing}}} = \begin{cases} 1 & \text{if } \left| \sum_{j=1}^{c} s_{i,j} \right| \neq c \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

$$N_{\text{flashing}} = \sum_{i \in N} N_{i_{\text{flashing}}} \tag{2}$$

Flashes is the total number of Boolean node state changes over the cycle. As at least one node must change state at each step around the cycle, this is related to the cyclelength and the flashing property. This can be expressed as:

$$N_{\text{flashes}} = \frac{1}{2} \sum_{i \in N} \sum_{j=1}^{c} \left| s_{i,j} - s_{i,j-1} \right| \tag{3}$$

Total is the sum of all Boolean node values at all time steps over the cycle. This is a property of the states of the bRBN rather than its dynamics and is related to the cyclelength property and the number of Boolean nodes.

$$N_{\text{tot}} = \sum_{i \in N} \sum_{j=1}^{c} s_{i,j} \tag{4}$$

Magnitude is the larger out of the total number of Boolean nodes at all time steps over the cycle that are in the 'true' state compared with the number that are in the 'false' state.

$$N_{\text{mag}_\text{T}} = \frac{1}{2} \sum_{i \in N} \sum_{j=1}^{c} (1 + s_{i,j}) \tag{5}$$

$$N_{\text{mag}_\text{F}} = \frac{1}{2} \sum_{i \in N} \sum_{j=1}^{c} (1 - s_{i,j}) \tag{6}$$

$$N_{\text{mag}} = \max\{N_{\text{mag}_\text{T}}, N_{\text{mag}_\text{F}}\} \tag{7}$$

Proportion is the proportion of nodes in state 'true' averaged over both cyclelength and number of Boolean nodes.

$$N_{\text{prop}} = \frac{N_{\text{mag}_\text{T}}}{n \times c} \tag{8}$$

Bonding Criteria

In addition to the bonding property, the bonding rule requires a comparison between the properties of two bRBNs for some criteria to be met. There are multiple possibilities to conduct this comparison, and this is another area for exploration.

Equal is the simplest bonding criteria; form a bond where the value of bonding property is *equal* within 0.1% of the maximum possible range of values to allow for numerical error). This can be expressed as:

$$\frac{p(N_i) - p_{\text{min}}}{p_{\text{max}} - p_{\text{min}}} - \frac{p(N_j) - p_{\text{min}}}{p_{\text{max}} - p_{\text{min}}} = 0 \pm 0.001 \tag{9}$$

Measurement	Minimum	Maximum	Description
Cyclelength	1	2^n	Count of steps on cycle
Flashing	0	n	Count of nodes that change state
Flashes	0	$n \times c$	Count of changes of node states over
Total	$-n \times c$	$n \times c$	Sum of node states over cycle
Proportion	0	1	Proportion of node steps with a value of True on cycle
Magnitude	1	$n \times c$	Maximum count of node states with False/True on cycle

Table 1: Alternative bRBN bonding criteria properties. n is the number of nodes within the bRBN, c is the cyclelength of the bRBN.

	bRBN node			
	A	B	C	D
Cycle steps	F	T	F	F
	F	F	F	F
	T	T	F	F
	T	T	T	F
	T	F	T	F
	F	T	T	F

$c = 6$

$N_{\text{mag}} = 14$

$N_{\text{flashing}} = 1 + 1 + 1 + 0 = 3$

$N_{\text{mag}_\text{T}} = 3 + 4 + 3 + 0 = 10$

$N_{\text{flashes}} = \frac{4 + 8 + 4 + 0}{2} = 8$

$N_{\text{mag}_\text{F}} = 3 + 2 + 3 + 6 = 14$

$N_{\text{tot}} = 0 + 2 + 0 + -6 = -4$

$N_{\text{prop}} = \frac{10}{4 \times 6} = 0.417$

Table 2: Example bonding properties for a $n = 4$ bRBN. Although only one would be used for any specific AChem, here they are all displayed. The table indicates the states of the bRBN nodes at each sequential step on the cycle.

where N_i and N_j are the bRBNs involved in the bond, $p(x)$ is a function to calculate the bonding property of bRBN x, and p_{min} & p_{max} are the minimum and maximum possible bonding property values.

Similar is a relaxation of the equal criteria — i.e. within 5% of the maximum possible range of values.

$$\frac{p(N_i) - p_{\text{min}}}{p_{\text{max}} - p_{\text{min}}} - \frac{p(N_j) - p_{\text{min}}}{p_{\text{max}} - p_{\text{min}}} \leq 0.05 \tag{10}$$

Different is the inversion of similar.

$$\frac{p(N_i) - p_{\text{min}}}{p_{\text{max}} - p_{\text{min}}} - \frac{p(N_j) - p_{\text{min}}}{p_{\text{max}} - p_{\text{min}}} \geq 0.05 \tag{11}$$

Sum one (applicable only to proportion) allows the formation of bonds where the proportion property of the interacting molecules total to one (± 0.001 allowing for numerical error).

$$p(N_i) + p(N_j) = 1 \pm 0.001 \tag{12}$$

Sum Zero (applicable only to total) requires that the total property of the bRBNs sum to a value of zero (± 0.001).

$$p(N_i) + p(N_j) = 0 \pm 0.001 \tag{13}$$

Sum One and Sum Zero are applicable only to proportion and total bonding properties respectively as these are the only bonding properties that can meet these bonding criteria.

n	k	Bonding Property	Bonding Criteria
5	2	Equal	Cyclelength
10	3	Similar	Flashing
15		Difference	Flashes
20			Total
25			Magnitude
			Proportion
		Sum One	Proportion
		Sum Zero	Total

Table 3: Features of the 200 alternative AChems tested. Every chemistry must have one feature from each column. Horizontal lines cannot be crossed within the table when moving from one column to the next. For example, 5 – 2 – Equal – Cyclelength is valid, 20 – 2 – Sum One – Proportion is valid, but 5 – 2 – Sum One – Flashes is not valid.

Sizes of bRBNs

The number of nodes (n) within each bRBN-atom must be chosen. A range of values at intervals was investigated ($n \in \{5, 10, 15, 20, 25\}$ with the potential to expand this range if there appears to be a directional trend).

The size of a bRBN does not have much impact on the chemistry directly. However, it does alter the distribution of the bonding properties, and their responses to bond formation, which in turn affects the propensity for different types of reactions.

Connectivity of bRBNs

Previous work on RBNs [10] has shown that the number of inputs (k) each node has can have an impact on their properties. There is also an interplay with the Boolean function

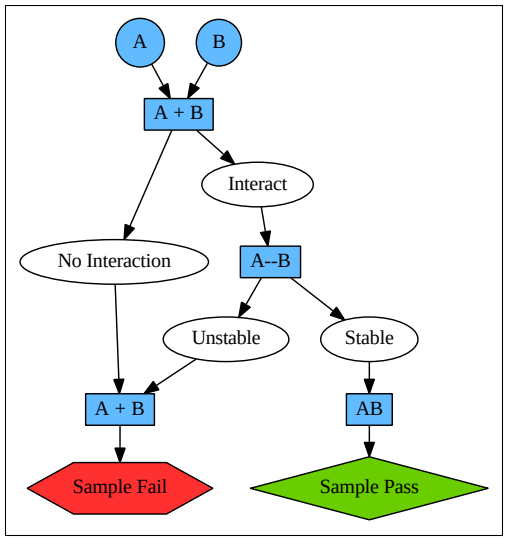

Figure 1: Schematic depiction of one sample for the 'synthesis' test its possible outcomes, and how those outcomes are interpreted as 'pass' or 'fail' for that sample. A & B are two sample atoms. A '+' symbol denotes separate atoms and '–' indicates a potential bond formation between two atoms. Adjacent atoms (e.g. AB) indicates that a bond has formed.

assigned to each node; certain functions can result in one or more inputs having no affect on the state of the node (canalisation) and more different Boolean functions are possible with more inputs.

As an initial assessment, we consider alternatives of two- and three-input bRBNs ($k \in \{2, 3\}$). In theory, any positive integer value equal to or less than the total number of nodes could be used. However, these are values known to be on the 'edge-of-chaos' — higher values are chaotic and lower values are statuc.

Combinations of Alternatives

The alternatives discussed above each change different, but potentially interlinked, aspects of the AChem. Different combinations of alternatives can be used, though some are mutually exclusive. Table 3 shows the possible combinations; in total there are 200 different AChems to be considered, each of which may have potentially different and interesting features.

Previous work [7] used $n = 10$ $k = 2$ with 'cyclelength' as the bonding property and 'equal' for the bonding criterion as an arbitrary initial choice from the 200 alternatives

Method

As discussed previously, there are a large number of potential alternative chemistries, and each of those has a very large number of potential elemental bRBNs.

Due to the vast number of possible bRBNs, exhaustively testing multiple chemistries is not feasible. Therefore, a random sampling approach is taken. In order for a chemistry to be have the potential for sufficiently rich properties, it is important that at the desired low-level behaviours are seen at least once. However, it is also important that the behaviours are not omnipresent — consider the synthesis test for example (described below); if every interaction resulted in the formation of a stable bond, it would rapidly coalesce into a single molecule and would therefore not exhibit sufficiently rich properties.

We do not seek to find the optimal subset of bRBNs in the optimal AChem; we are simply looking to remove those alternative AChems unlikely to exhibit sufficiently rich emergent properties.

Desired Behaviours

As well as the alternative chemistries, the tests for required low-level behaviours must also be defined. There is a natural structuring of prerequisites within the behaviours – decomposition can only occur if synthesis occurs for example. This can be used to increase the efficiency of the sampling.

Synthesis Synthesis is the lowest-level behaviour possible in an atom-based AChem. A pair of atoms is randomly sampled, the two atoms interact, and the outcome is recorded. RBN-World has a two-stage bonding process, and the bonding criteria must be met both at the start of the interaction and after bonding. If a stable bond can be formed, then the sample passes; if not, the sample fails (figure 1).

Self-synthesis The self-synthesis test the synthesis test between two copies of the same element. If a stable bond can be formed, then the sample passes; if not, the sample fails.

Decomposition This is the breaking of bonds, potentially leading to a molecule separating into two (or more) smaller molecules. In RBN-world this is triggered by an interaction between an bRBN molecule and another bRBN. In the decomposition test, samples of three atoms are taken and the first two attempt to form a stable bond. If they cannot form a stable bond, then that sample is ignored for determining pass/fail; this is a test for decomposition, not for synthesis. Once a stable molecule has been formed, it interacts with the third sample. This can have several possible outcomes; no interaction, formation of a larger molecule, or breakdown into two or three separate molecules. If it results in the bond between the first two sampled bRBNs breaking, then it is recorded as a pass; other outcomes are classed as fail (figure 2).

Substitution Similar to decomposition, the substitution test involves an interaction with a molecule that leads to replacement of part of the molecule with the reacting bRBN. The process is the same as the decomposition test, but the only valid outcome is a direct replacement of the second sampled bRBN with the third sampled bRBN, i.e. AC+B in figure 2.

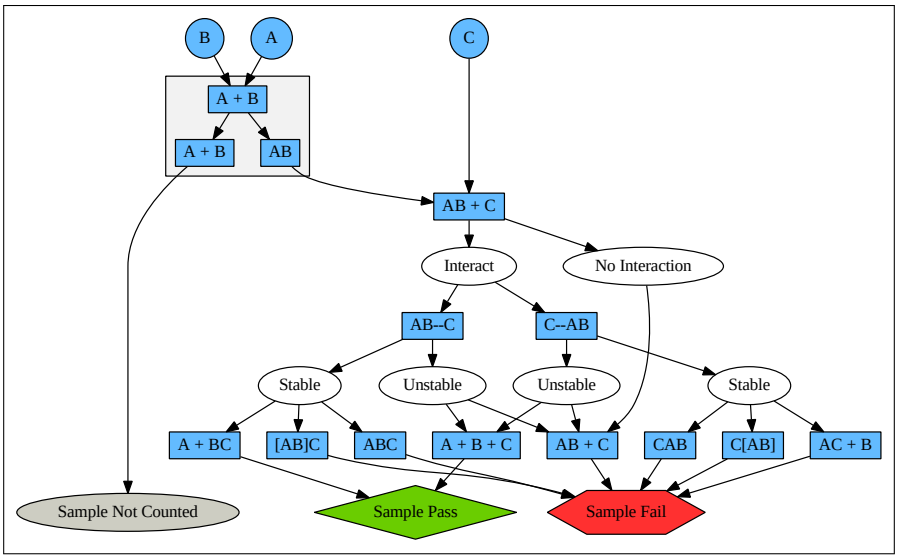

Figure 2: Schematic depiction of the 'decomposition' test. A requirement of the decomposition test is that synthesis must have first occurred, this part of the schematic is indicated in the highlighted subgraph (details removed for brevity).

Catalysis This is the highest-level property investigated here. Unlike the other desired properties, catalysis can take many forms. Any of the other tests could be repeated requiring the presence of a catalyst. For simplicity, we focus on catalysis of synthesis reactions.

The test proceeds as follows: as before, a sample of three bRBNs is taken and the first two attempt to form a stable bond. However, unlike decomposition or substitution tests, this time it is important that a stable bond does not form. If a bond does form, then the sample is not counted for pass/fail.

After that initial bond formation stage, the third bRBN in the sample attempts to form a bond with the first; this is analogues to interacting with a catalyst to form a temporary intermediate. If this does not form a stable bond, then again the sample is not counted for pass/fail.

The final step is to test that the second bRBN from the sample can substitute for the third bRBN. If this is the case, then the third bRBN has acted as a catalyst for the formation of the bond between the first and the second bRBN that would not occur directly (figure 3).

Results

The outcomes of testing the described alternative chemistries with 10,000 randomly generated samples of bRBNs is summarized in table 4 (testing took approx. 2 days on a 24 quad-CPU cluster). With each test a number of alternative AChems are ruled out; the chemistries that pass all tests are listed table 5.

Less than 5% of alternative chemistries pass all the tests. The n & k categories of alternatives have little or no influence on the low-level properties of the chemistry. The anomaly is $n = 25$, $k = 3$ with bonding property 'total' and a comparison of 'sum zero'; however, this may be due to sample size. Closer examination of this case shows that of $10,000$ samples in the decomposition test, $9,677$ were not counted (as the did not form a molecule that could break down) and none of the remaining 323 samples passed. In comparison, the $n = 20$ equivalent AChem where $9,382$ were not counted and 43 of the remaining 618 samples passed.

For the property and comparison alternatives, only those using 'proportion' as property and 'sum one' as the criterion or those using 'total' as the property and 'sum zero' as the criterion pass all tests. Whilst alternatives should be kept in mind, we now have evidence that these are options are more likely to be capable of rich emergent properties. As various

Test	# of AChems tested	# of AChems where all samples Passed or Not Counted	# of AChems where all samples Failed or Not Counted	# of AChems remaining
Synthesis	200	10	7	183
Self-Synthesis	183	110	53	20
Decomposition	183	0	6	177
Substitution	177	0	18	159
Catalysis	177	0	39	138

Table 4: Results from testing 10,000 samples from each of 200 alternative chemistries for low-level emergent behaviours. The prerequisite for decomposition and self-synthesis tests is synthesis. The prerequisite for substitution and catalysis tests is decomposition. See text for details.

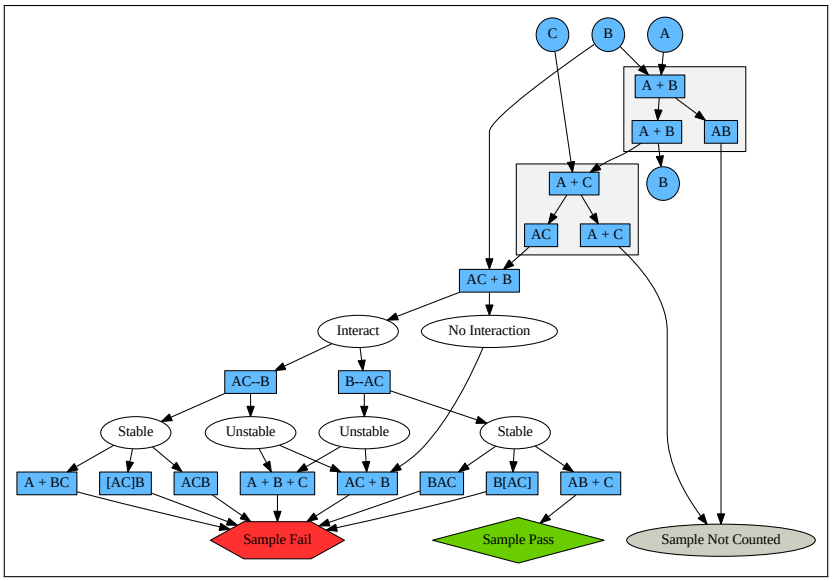

Figure 3: Schematic depiction of the 'catalysis' test. Requirements of the catalysis test are that A and B must *not* synthesise, and that C and A must synthesise; these are indicated by the highlighted subgraphs (details removed for brevity).

n	k	Measurement	Comparison
5	2	Proportion	Sum One
10	2	Proportion	Sum One
15	2	Proportion	Sum One
20	2	Proportion	Sum One
25	2	Proportion	Sum One
5	3	Proportion	Sum One
10	3	Proportion	Sum One
15	3	Proportion	Sum One
20	3	Proportion	Sum One
25	3	Proportion	Sum One
5	2	Total	Sum Zero
10	2	Total	Sum Zero
15	2	Total	Sum Zero
20	2	Total	Sum Zero
25	2	Total	Sum Zero
5	3	Total	Sum Zero
10	3	Total	Sum Zero
15	3	Total	Sum Zero
20	3	Total	Sum Zero

Table 5: The 19 alternative AChems that exhibit variation across all 5 low-level emergent behaviours tested.

different values of n and k were tested and did not affect which chemistries passed the tests, these values can be chosen based on other concerns, such computational tractability. One potential issue is that this work has only samples from atomic constituents; it is not guaranteed that molecular structures will also exhibit these behaviours. While various values of n were tested, molecular bRBNs of many atoms may not behave as an equivalent large bRBN atom due to the constrictions from reciprocal bonding sites between atoms.

Conclusions

We have presented simple tests of an AChem that can be used to restrict the design space to non-trivial chemistries. This is important, as for many AChems there are a large number of alternatives that should be considered – for RBN-World we have only examined a small fraction of possible alternatives. It has also been shown that our initial arbitrary choice of parameters did not pass these tests [7]. This is an important consideration as the processes that lead to the design of an AChem are typically opaque to the community.

A filtering metric provides a useful testing approach that does not require computationally expensive and/or exhaustive testing of molecules and/or reactions. It is also interesting to see that some AChems fail because all tested samples interactions failed, but some chemistries fail because all tested sample interactions passed; the presence of variation is a requirement for emergent properties.

Future work

Two specific alternative parameterisations of RBN-World have been identified as containing interesting atoms; future work can now be focused onto searching for specific small sets of elements within these chemistries that give rise to the high-level desired properties discussed earlier — autocatalytic sets, hypercycles and heteropolymers. These have not been tested for in the experiments described here due to the small samples from each chemistry that were being examined.

In addition, the low-level tests will be refined further. One example is that here only atoms were tested and there is no guarantee that these properties are also applicable for larger structures. As we can now remove the trivial, uninterest-

ing cases, computational effort will be concentrated on those non-trivial cases, in the hunt for rich AChems.

Acknowledgements

Thanks to Jess Wardman for constructive feedback.
This work was funded by BBSRC.

References

[1] Dittrich, P. and di Fenizio, P. S. (2007). Chemical organisation theory. *Bulletin of mathematical biology*, 69(4):1199–1231.

[2] Dittrich, P., Ziegler, J., and Banzhaf, W. (2001). Artificial chemistries-a review. *Artificial Life*, 7(3):225–275.

[3] Drossel, B. (2007). Random Boolean Networks. *Cybernetics and Systems*, 12(1):103–121.

[4] Eigen, M. and Schuster, P. (1977). The Hypercycle: Part A. *Naturwissenschaften*, 64(11):541–565.

[5] Eigen, M. and Schuster, P. (1978a). The Hypercycle: Part B. *Naturwissenschaften*, 65(1):7–41.

[6] Eigen, M. and Schuster, P. (1978b). The Hypercycle: Part C. *Naturwissenschaften*, 65(7):341–369.

[7] Faulconbridge, A., Stepney, S., Miller, J. F., and Caves, L. (2009). RBN-World: A Sub-Symbolic Artificial Chemistry. In *Proceedings of the tenth European Conference on Artificial Life*. Springer LNCS5777.

[8] Gardner, M. (1970). The fantastic combinations of John Conway's new solitaire game "life". *Scientific American*, 223:6–10.

[9] Kauffman, S. A. (1969). Metabolic stability and epigenesis in randomly constructed genetic nets. *Journal of theoretical biology*, 22(3):437–467.

[10] Kauffman, S. A. (1993). *The Origins of Order: Self-Organization and Selection in Evolution*.

[11] Kauffman, S. A. and Farmer, J. D. (1986). Autocatalytic sets of proteins. *Origins of Life and Evolution of the Biosphere*, 16(3-4):446–447.

[12] Meyer, T., Yamamoto, L., and Tschudin, C. (2008). An Artificial Chemistry for Networking. *Bio-Inspired Computing and Communication*, 5151:45–57.

[13] Odian, G. (2004). *Principles of Polymerization*. John Wiley & Sons, Inc., Hoboken, NJ, USA.

[14] Reynolds, C. W. (1987). Flocks, herds and schools: A distributed behavioral model. *Proceedings of the 14th annual conference on Computer graphics and interactive techniques*.

[15] Seth, A. (2010). Measuring Autonomy and Emergence via Granger Causality. *Artificial Life*, 18(Early Access):118.

[16] Straatman, B., White, R., and Banzhaf, W. (2008). An Artificial Chemistry-based Model of Economies. *Artificial Life*, 11:592.

[17] Suzuki, H., Ono, N., and Yuta, K. (2003). Several necessary conditions for the evolution of complex forms of life in an artificial environment. *Artificial Life*, 9(2):153–174.