RBN-World
A Sub-Symbolic Artificial Chemistry

Adam Faulconbridge[1], Susan Stepney[2], Julian F. Miller[3], and Leo Caves[4]

[1] YCCSA, University of York, Heslington, YO10 5DD, UK
asf500@york.ac.uk
[2] susan@cs.york.ac.uk
[3] jfm7@ohm.york.ac.uk
[4] lsdc1@york.ac.uk

Abstract. We describe a composable dynamical system that uses the emergent properties of coupled random Boolean networks (RBNs) as a basis for a sub-symbolic artificial chemistry. The approach shows potential for open-ended emergent properties and may lead to a foundation for artificial life.

1 Introduction

Artificial chemistries [1] (AChem, AC) have been used for investigations into the emergence and/or early development of biological phenomena in an abiotic environment with arguable success at generating systems with multiple levels of emergence. We believe this is because previous artificial chemistries have used symbolic approaches. We propose a sub-symbolic [2] approach based on composable random Boolean networks (RBNs) to produce a system within which self-organizing multi-level structures could emerge.

In symbolic representations each "atom" has no internal structure. In a sub-symbolic representation, the "atoms of meaning" are emergent properties of complex dynamics. An example useage of sub-symbolism is neural networks in the field AI; the learned information emerges from the network structure and the weights of the links, rather than being explicitly encoded in a fixed set of symbols. Sub-symbolic representations allow new, unforeseen, "atoms of meaning" to emerge from the developing system.

2 Sub-symbolic Artificial Chemistry

Artificial chemistries are analogous to real-world chemistry in that indivisible building blocks (atoms) bond together to produce larger structures (molecules). However, real-world atoms have internal structure (e.g. electron shells) that is not incorporated in an artificial chemistry based on symbols. We propose using a sub-symbolic representation to account for this feature.

A sub-symbolic representation suitable for an artificial chemistry should exhibit the following features:

– Deterministic and computationally tractable
– Emergent characteristics
– Composability to enable sub-symbolic representations of molecular structures can be constructed
– Upward and downward causation so that low-level changes have the potential to disrupt higher-level structures and *vice versa*

Composability of a rich sub-symbolic representation allows molecular structures, such as functional groups or polymers, to potentially be more than the sum of their parts. In this fashion we hope that analogies to biological structures may emerge: a protein is one entity but it is composite of amino acids, each of which is composite of several functional groups, which are themselves composite of multiple atoms.

Sub-symbolic composability allows reactions between novel structures to occur without the need to specify additional reaction rules. This is important for evolution within an artificial chemistry as it potentially enables open-ended development.

Decomposability is also a desirable feature of a sub-symbolic representation. By allowing interactions at multiple levels of structure, lower-level changes have the potential to alter higher-level structures (e.g. the breakage of bonds where catalysts separate from their products).

There are many possible sub-symbolic representations, and many possible artificial chemistries using them. We have made some arbitrary choices for the representation and reaction rule in order demonstrate proof-of-concept. The system described below is only one example of many possible artificial chemistries using this sub-symbolic framework. In addition, the example reactions represent only a tiny sample of possible reactions within this chemistry.

3 RBN-World: chemistry

RBNs

Random Boolean networks [3–5] (RBNs) are our system of choice for a sub-symbolic artificial chemistry due to their rich dynamical structure[5]. In this work, we use a reaction rule based upon matching cyclelengths and composition of RBNs.

An RBN consists 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 its input nodes to its state at the next timestep. We use $k = 2$ for all RBNs described here. Function and initial state of each node are chosen at random.

[5] RBNs were originally devised as simple models of the genetic regulatory network within a cell. Subsequent work using RBNs has continued this theme and focused on reflecting biological networks. However, here we use RBNs as composable dynamical systems with emergent properties.

All nodes in the network simultaneously update their state at time t based on the states of their inputs at time $t-1$. The state of an RBN is the collection of states of the nodes. All RBNs have cyclic behaviour, returning to a previous state after sufficient timesteps. The number of timesteps on a cycle is the *cyclelength*. The distribution of cyclelengths is highly skewed with median $\sqrt{n}$, with a long tail to a theoretical (but rarely seen) maximum value of 2^n. This means that discovering a RBNs cyclelength is computationally tractable.

RBNs exhibit sensitivity to noise perturbations [3], i.e. a change in one node may change the behaviour of the entire network, or may do nothing. This gives a structured richness to the system that is rarely found in combination with such simplicity. RBNs have a vast number of possibilities, yet they have a number of emergent properties (e.g. cyclelength) with complex many-to-one mappings.

For use in a chemistry, RBNs need to be be combined and fragmented. Therefore some modifications have to be made to traditional RBNs which are described below.

Some further definitions are needed for an artificial chemistry; *atoms*, *bonds*, *reactions* and *molecules*. Due to space limitations, we can not cover all of the system in depth and therefore only outline the important features.

Atoms

We define an additional feature of RBNs; *bonding sites* to make bRBNs (bonding random Boolean networks). Bonding sites (b) are one or more additional nodes that are each taken as an input by one ordinary node chosen at random (a single ordinary node has at most one input from bonding sites). Bonding sites do not have any inputs; their state is determined by whether they are "filled" or "empty". See figure 1 for two example bRBNs. For the preliminary work described here, $b = 2$ for all bRBNs.

The coupling of bRBNs through bonding sites mean that a reaction can change one input to a single node. Due to the sensitivity of the dynamics of RBNs, the change of state of bonding site on formation of a bond can have a wide range of effects (or none).

Bonds

A bond links two bRBNs. There can be multiple bonds between the same pair of bRBNs. Each bond requires one bonding site within the pair of bRBNs to become "filled", and each "filled" bonding site is associated with only one bond. In the chemistry described here, we require that the two bRBNs linked by a bond must have equal cyclelengths both before and after bonding.

Reactions

To form a bond, we require that the two bRBNs have equal cyclelengths both when the bonding sites are "empty" and when the bonding sites are "filled". We do not require the cyclelength when the bonding site is "empty" to be equal to the cyclelength when the bonding site is "filled". Example structures before and after a reaction are shown in figures 1 and 2 respectively (summarized in figure 3).

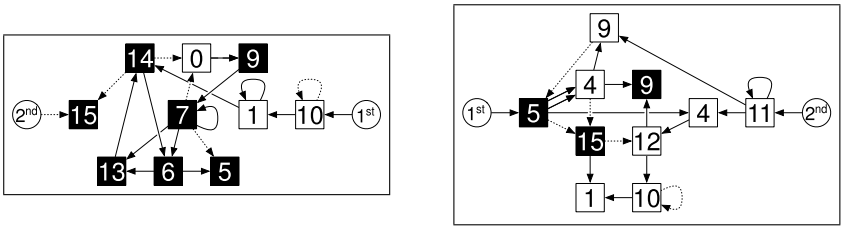

Fig. 1. Two example bRBNs ($n = 10$, $k = 2$, $b = 2$). Numbers are Boolean functions, colour indicates state at this timestep. Edges indicate where outputs are connected to; dashed lines indicate inputs that are always ignored. White circles represent "empty" bonding sites with their bonding order pre-specified.

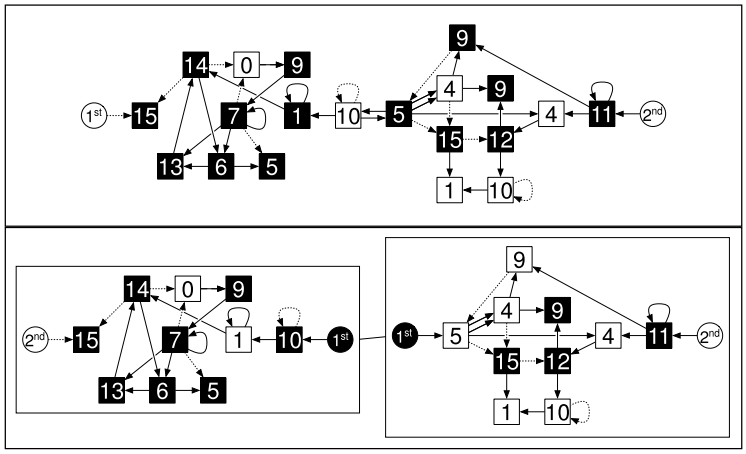

Fig. 2. Example RBN-molecule constructed from RBN-atoms in figure 1. Above is the composite bRBN, and below the component bRBNs. Black circles represent "filled" bonding sites.

If a bond is not formed, it is attempted again with any higher-level structures the pair of bRBNs are part of. This iteration of attempting bonding and retrying for higher-level structures continues until either a bond is formed or there are no more higher structures. See figure 4.

Molecules

bRBNs that are linked by one or more bonds can be expressed as a composite bRBN. The composite bRBN structure is the same as the structure of the component bRBNs, except that inputs from "filled" bonding sites are replaced with direct reciprocated inputs (e.g. figure 2). Non-composite bRBNs are *RBN-atoms*, and a composite bRBN (at any level) is a *RBN-molecule*. RBN-molecules may form bonds in the same manner as RBN-atoms to make higher-level composite structures. In this representation we track multiple levels of structure in order to allow decomposition events. Note that a node can be in different states at different levels of the structural hierarchy

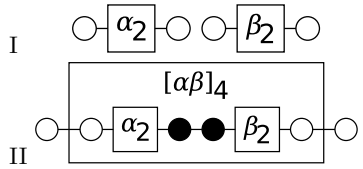

Fig. 3. Abstract representation of figures 1 and 2; RBN-atoms (I) and the RBN-molecule (II). Squares represent the RBN abstracted to a letter and the number shows the current cyclelength. Square brackets denote "is built of" to show that in $[\alpha\beta]_4$ the subscript refers to the combined bRBN rather than just β. The internal structure of a composite bRBN can be similarly expressed, e.g. $[\alpha_2\beta_2]_4$. Note that all RBN-atoms should have square brackets, e.g. $[\alpha]_2$, but for brevity they are omitted for single atoms. The 'lollipops' represent bonding sites; white when "empty" and black when "filled". This reaction can be expressed in symbolic form as $\alpha_2 + \beta_2 \rightarrow [\alpha_2\beta_2]_4$.

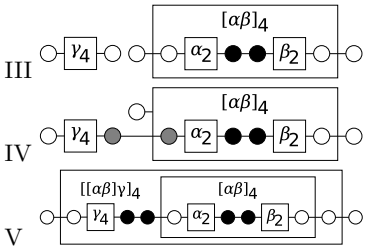

Fig. 4. Example formation of a multi-level structure. The RBN-molecule $[\alpha\beta]_4$ from figure 3 reacts with RBN-atom γ_4. Step III shows the initial condition, and step IV shows the attempted bonding between γ_4 and α_2. The cyclelengths do not match, so they do not bond. The bonding attempt is repeated between γ_4 and $[\alpha\beta]_4$. The cyclelengths do match and step IV shows the forming bond (indicated by the grey 'lollipops'). The final structure is shown in step V. The nested boxes show that γ_4 is bonded to $[\alpha\beta]_4$ rather than α_2. The reaction can be expressed in symbolic form as $[\alpha_2\beta_2]_4 + \gamma_4 \rightarrow [\gamma_4[\alpha_2\beta_2]_4]_4$.

Effects of Bonding

There are two direct consequences to the formation of a bond:

1. The process of bonding changes a bonding site in each linked bRBN from "empty" to "filled". This changes one input to one node, which can potentially lead to a change in cyclelength.
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 single bRBN.

Additional bonds can be formed as long as the requirements for bonding can be satisfied. Preliminary investigations suggest that complicated structures with multiple levels do form: this requires further investigation to characterise fully.

Bonds can be destroyed as well as created. A bond is broken whenever its two linked bRBNs no longer have equal cyclelengths. The circumstances for this depends on the details of the bRBNs participating in the bond.

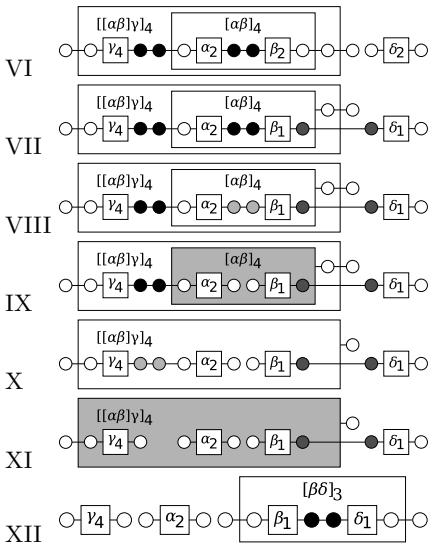

Fig. 5. Continuing from figure 4, δ_2 reacts with β_2. The first stage is to fill in a bonding site on δ_2 and β_2 (shown in dark grey in steps VII – XI above). This changes the cyclelength of β_2 to 1 (see step VII). As β_1 cyclelength is now different to α_2, the bond between them is removed (shown in grey in step VIII). This empties a bonding site in β_1 and α_2. The breakdown of the α_2 — β_2 bond also means $[\alpha\beta]_4$ no longer exists (shown in light grey in step IX) and therefore the γ_4 — $[\alpha\beta]_4$ bond and $[\gamma[\alpha\beta]]_4$ molecule no longer exist (shown in light grey in step X and step XI respectively). The final state at the end of the reaction is shown in step XII above. This reaction can be expressed in symbolic form as $[\gamma[\alpha\beta]]_4 + \delta_2 \rightarrow \gamma_4 + \alpha_2 + [\beta_1\delta_1]_3$.

An example of a reaction that leads to breaking bonds and the decomposition of RBN-molecules can be seen in figure 5. The molecule that has been built up by previous reactions in figures 3 and 4 goes on to react with another RBN-atom. This causes a change in cyclelength which triggers a cascade of bond breakage and structure fragmentation. Processes like these contribute to the rich complex dynamics of this artificial chemistry.

4 Discussion

Composite bRBNs are not identical to conventional RBNs: in addition to the presence of bonding sites, composite bRBNs have a distinctive topology due to the restricted connectivity between the component bRBNs. Here we demonstrate that this does not adversely impact the complex dynamics that we are interested in.

First, we show that the addition of bonding sites to RBNs (and thus the fixing of an input to a node) does not drastically change the distribution of cyclelengths ($\sqrt{n}$ median and long-tail). The results of examining 1,000 RBNs and RBN-atoms over $5 < n < 5000$ are shown in figure 6.

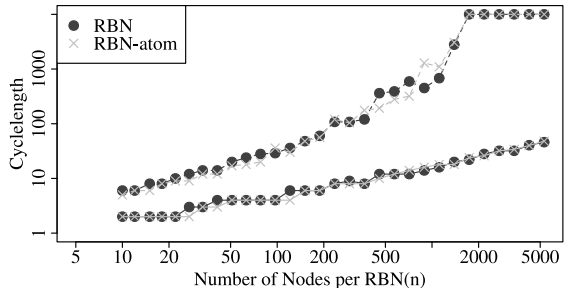

Fig. 6. Distribution of cyclelength in RBNs and RBN-atoms with n nodes. Solid lines are medians, dashed lines the 90^{th} percentiles; cyclelength was capped at 10,000.

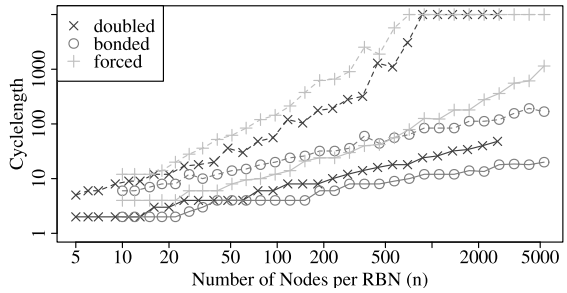

Fig. 7. Distribution of cyclelength for RBN-molecules composite of two RBN-atoms for n nodes. Solid lines are medians, dashed lines the 90^{th} percentiles; cyclelength was capped at 10,000. "Bonded" refers to atoms that are joined based on matching cyclelengths, "forced" are atoms joined regardless of cyclelength, and "doubled" refers to a single RBN-atom with $2n$ nodes.

Second, we show that RBN-molecules have a similarly shaped distribution of cyclelengths to RBN-atoms. The precise distribution is influenced by the details of the bonding scheme: the requirement for equal cyclelengths, and the topology of composite bRBNs. Therefore, we compare higher-level bRBNs formed in three different ways: bonding between two bRBNs with matching cyclelengths, forced bonding between two bRBNs without any requirements, and a single bRBN with twice the number of nodes, $2n$.

The results of examining 1,000 bRBNs at 30 different values where $5 < n < 5000$ are shown in figure 7. All three bonding schemes result in a broadening distribution with increasing n, though the rate of increase varies. Forced bonding has the steepest increase; this is most likely due to a "lowest common multiple" effect rather than the bond itself. If bond itself has no effect on cyclelength, then the composite structure must have a cyclelength equal to the lowest common multiple of its component bRBNs cyclelengths. If the two bRBNs have the same cyclelength, then the composite structure must also have that cyclelength.

These bonding schemes produce composite structures with long-tailed distributions of cyclelengths. This shows that bRBN-molecules maintain the in-

teresting dynamical properties of RBNs, and thus provide a basis for future higher-order emergence.

5 Future work

The artificial chemistry described here is a first step in exploring the emergent properties of composable discrete dynamical systems. We note that this framework allows for the specification of whole classes of new artificial chemistries. Some ideas for future work include:

- Varying n, b, and k.
- Limiting Boolean function sets, e.g. no fixed functions
- Characteristics other than cyclelength for bonding
- Requirements other than matching for bonding
- Locating bonding sites by emergent dynamical features (such as node activity) rather than pre-specifying them
- Using other dynamical systems, eg cellular automata (CAs), as atoms
- Identifying a small subset of networks as 'elements'; selected by, for example, a genetic algorithm
- Adding a measure of bond strength, allowing stronger bonds to replace weaker ones
- Introducing spatial aspects
- Introducing thermodynamics and / or entropy as implicit or explicit measures

We have introduced an artificial chemistry based on composable dynamical systems which offers the prospect of rich emergent properties with the potential for open-ended behaviour. A key aspect of our approach is the composition of sub-symbolic components into hierarchical structures, eschewing the need for additional externally imposed rules and / or symbols at each level of organisation. Here an illustrative RBN-based artificial chemistry has been used for proof-of-concept, but other dynamical systems and interaction schemes are possible. We propose that sub-symbolic composable systems provide a framework for the open-ended evolution of artificial life with emergent features.

References

1. Dittrich, P., Ziegler, J., Banzhaf, W.: Artificial chemistries — a review. Artificial Life **7**(3) (2001) 225–275
2. Smolensky, P.: On the proper treatment of connectionism. Behavioral and Brain Sciences **11**(1) (1988) 1–74
3. Kauffman, S.A.: Metabolic stability and epigenesis in randomly constructed genetic nets. Journal of Theoretical Biology **22**(3) (1969) 437–467
4. Kauffman, S.A.: The Origins of Order. Oxford univ. Press (1993)
5. Drossel, B.: Random Boolean Networks. In:. Volume 1 of Reviews of Nonlinear Dynamics and Complexity. Wiley-VCH (2008) 69–110