

This is a repository copy of *Functional grouping analysis of varying reactor types in the Spiky-RBN AChem*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/121298/>

Version: Published Version

Conference or Workshop Item:

Krastev, Mihail Svilenov, Sebald, Angelika Anne-Marie orcid.org/0000-0001-7966-7438 and Stepney, Susan orcid.org/0000-0003-3146-5401 (2017) Functional grouping analysis of varying reactor types in the Spiky-RBN AChem. In: European Conference on Artificial Life 2017, 04-08 Sep 2017.

Reuse

This article is distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs (CC BY-NC-ND) licence. This licence only allows you to download this work and share it with others as long as you credit the authors, but you can't change the article in any way or use it commercially. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Functional grouping analysis of varying reactor types in the Spiky-RBN AChem

Mihail Krastev^{1,3}, Angelika Sebald^{2,3}, Susan Stepney^{1,3}

¹Department of Computer Science, University of York, UK

²Department of Chemistry, University of York, UK

³York Centre for Complex Systems Analysis
mk599@york.ac.uk

Abstract

We explore the effects that different reactor types have on Spiky-RBN AChem systems, looking at mass conserving and flow reactors. To assist in analysing the behaviour we introduce an activity measure based on possible system state changes as a result of changes in particle properties. This leads to a discussion on approaches to engineering complex systems towards specific goals.

INTRODUCTION

The Spiky-RBN Chemistry (Krastev et al., 2016) is a type of subsymbolic Artificial Chemistry (ssAChem) designed such that particle properties and behaviour are derived from its underlying system dynamics. The dynamic behaviour of the constituent random boolean networks (RBN) (Kauffman, 1969; Kauffman et al., 2003) allows us to define the properties of our atomic particles. The RBN's nodes are partitioned into "spikes", which form the basic linking units. Each spike has a linking property derived from the behaviour of the underlying RBN. The linking property has an integer value and can be positive or negative. Figure 1 shows a representation of an RBN partitioned into spikes. A reaction occurs when the linking properties of the reactants are compatible, and results in a composite particle. This composite particle is a composite RBN built from the reactant RBNs, and its new spike properties are based on the new composite RBN dynamics. See Figure 1 for an overview, and see Krastev et al. (2016) for the details of the various algorithms involved.

Krastev et al. (2016) use a heuristic algorithm to search for suitable sets of 20 atomic Spiky RBNs, from the vast space of possible RBNs. These sets are chosen to be reactive and to produce a rich variety of composite particles, in terms of particle size, structure and composition.

The reactor rules used during that search are geared towards an exploration of the possible reaction space rather than any analysis of the dynamics. The reactor contains one of each atomic particle and each unique composite particle found so far; there are equal concentrations of all particles found so far, and there is always enough material for further




Figure 1: An atomic and composite particle in the Spiky-RBN system.

reactions. The original 20 atomic particles quickly become a small portion of the overall reactor contents, which biases the exploration towards larger composite particle reactions.

Such reactor rules are not very realistic. A mass-conserving reactor, in which concentrations change as a result of reactions, is closer to reality. From a system's perspective, it also provides pressures on the types of particles that can be seen, and could lead to more diversity in the smaller composite particles generated. Reactant concentrations play an important role in reaction dynamics, so it is interesting to compare different reactor rules and their effects on observed system dynamics. Here we explore how different types of mass-conserving reactors affect the dynamics of our system.

Analysing and quantifying the effects of these different rules on the observed systems can be challenging. In AChems with defined reaction rules there is often a desired behaviour that the rules aim to produce, so analysis can be geared towards finding that behaviour. With the Spiky-RBN Chemistry, the aim is to explore the dynamics and activity of ssAChems and how they are affected by different rules.

One approach to analysis is quantifying the complexity of the system. There are multiple definitions of complexity, and from there, multiple approaches to measuring it. Within our system we could look at the complexity of the RBNs (Ciencias del Espacio, 2000; Wang et al., 2011), or of the composite particle geometries (Rashevsky, 1955; Karreman, 1955). We are interested in the reaction networks our system produces. There are metrics for chemical reactions specifically (Karreman, 1955; Bonchev and Trinajstić,

1977), and for graph complexity (Dehmer and Mowshowitz, 2011). However, we want our analysis tooling not only to provide a numeric quantity, but also to allow us to analyse and reason about the system behaviour. Therefore, we move away from these more generic complexity measures, and define our own metric, which is designed to accommodate the reaction processes of our Spiky-RBN system, while still being based on core concepts of complexity.

An Activity Measure for Spiky-RBN

The dynamics of our system can be described as a reaction network: a directed graph where nodes are particles in the system, and observed reactions are directed edges connecting reactants to products. Defining where a complexity measure should be maximum can be difficult. We can state that our measure should tend to zero when the behaviour is ‘uninteresting’, by which we mean when the system is either ordered or random in terms of reaction possibility and particle properties. The particle properties (the spikes in our system) are the only thing that particles can observe about each other, and the main thing that dictates what reactions can take place. We therefore think of these spikes as a descriptor of the system. This does not fully describe our system, since it is capable of degeneracy: two different RBNs could have the same spikes produced by different underlying dynamics. In this case their properties would be the same, however their different dynamics could become visible after they have reacted and formed different composites.

A purely random system has no structure for defining which reactions occur, or the resulting properties of the products. Every possible reaction attempt would be successful, and any two particles could create a composite linked between any of their spikes. Additionally, there would be no connection between the property of the product and the properties of the reactants that formed it. In effect we would see a huge number of possible unique particles with almost no internal structure. Eventually the system would tend to one single composite containing all atomic particles.

A fully ordered system can take two forms. The simplest is a completely inert system where no reactions are possible at all. Alternatively, some reactions are possible, but no reaction changes particle properties—composites have identical spikes as their reactants—and therefore no reaction changes the behaviour in the system. In this case the set of possible reactions never changes. In practice this would be a system which produces only “polymers”, particles with a rigid internal structure which are only ever capable of the same reactions they are a product of.

“Interesting” systems lay in between these extremes of order and randomness. Some reactions maintain the behaviour of particles and composites, while other reactions result in changes in behaviour causing products to have different spikes from their reactants, therefore changing the possible set of reactions. An open ended system is one

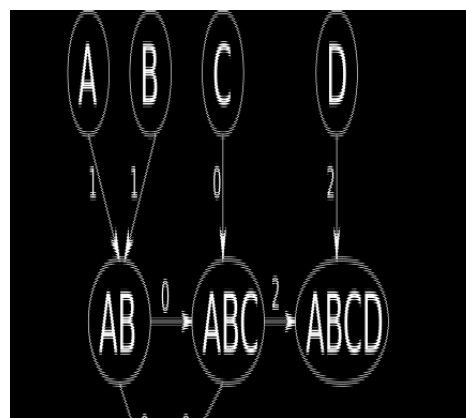


Figure 2: Set of reactions expressed as a reaction graph (top), and a functional grouping (bottom). Edges represent reactions labelled with how many spikes were changed by the reaction.

where a comparatively small but still unlimited number of reactions cause changes in behaviour. The occurrence of such reactions would open up new possible interactions to the system, possibly leading to new particle types. Our metric is designed to measure such behaviours in terms of spike properties.

Functional Grouping

We start from a complete reaction network, a graph describing all reactions that have occurred in the reaction vessel we are analysing.

We reduce this full reaction network graph to include only those reactions which result in a change to spikes, where the spikes of the product are different from the (non-linked) spikes of the reactants. This focusses on products that have new linking properties. We do this by merging particle nodes where reactions do not change spikes. Figure 2 shows an example of this grouping, reactions that produce *ABC* and *ABCAB* do not cause a change in spikes therefore we merge the reactants with the products.

These groupings represent sets of particles that share a common root structure and for which there exist reactions that do not change the reactant’s spikes.

If we describe the state of the system as the set of possible reactions that the system can undergo, reactions within a group do not change the system state. Edges between groups represent reactions that change unbound spikes. A change of unbound spikes means a change in possible reactions and therefore a change in system state.

This method of ‘functional grouping’ also allows us to reason about the bonding spikes themselves. If we observe that most successful reactions involving a given spike result in no change to unbound spikes we can think of that bonding spike as “frozen”. Alternatively if any link on a given spike changes the properties of the other unbound spikes then we

can consider that bonding spike as “active”. We can then reason about the likelihood of reactions causing changes to the system based on what spikes are involved in the link and how frozen or active those spikes are.

In most cases a reaction will change only some spikes, not all of them. This allows us not only to label spikes as frozen or active, but more precisely say how perturbing they are to the system, how many other spikes they change. We label edges on the reaction network with how many spikes are changed as a result of the reaction. Any change in a spike will result in a different set of reactions being possible, therefore tracking how much an individual spike changes gives no additional information, so we do not track that.

Activity Measures in Terms of Functional Grouping

We reason about the dynamics of our system in terms of the identified functional groups as follows. As the number of groups approaches the total number of unique particles, we describe the system as random, since each particle has properties different from the reactants that made it. Conversely, as the number of groups approaches one we have an ordered system, since all found particles share the same properties.

We plot the activity A of the system in terms of the ratio between the number of groups and total number of unique particles produced:

$$A = \frac{FG(t)}{U(t)} \quad (1)$$

Here $FG(t)$ is the number of functional groups after t bonding attempts and $U(t)$ is the number of unique particles after t attempts.

Clearly, $0 < A \leq 1$. As A tends towards 1 the system is more random, and as it tends towards 0 the system is more ordered. This can more clearly be observed as change in activity, ΔA , over a period of $t + n$ reaction attempts.

$$\Delta A_n^t = \frac{FG(t+n) - FG(t)}{U(t+n) - U(t)} \quad (2)$$

Mapping these values over the runtime of a reactor and comparing them across different reactor types provides us with a way to reason about the system dynamics. It also highlights moments of increased activity, which we can then explore in more detail.

Experimental Implementation

Krastev et al. (2016) found 20 different sets of 20 atoms that display more interesting behaviour than randomly chosen atom sets. Here we use those atom sets to explore four different reactor rule sets: a simple mass conserving example, two types of flow reactors and one hierarchical decanting reactor.

The mass conserving reactor is initialised with 200 instances of each atomic particle. A reaction attempt consists

of removing two particles at random, attempting a reaction, and returning to the reactor whatever the products of the reaction are: a new particle for a successful attempt, or the original reactants for an unsuccessful attempt. Since no reaction can create or destroy atoms, only change links, the reactor will conserve its mass.

A flow reactor proceeds in the same way, by selecting particles to attempt reactions and returning the products. Periodically, however, the flow reactor injects new atomic particles into the system. In order to maintain mass a random set of particles with equal total mass is removed from the reactor.

The first variant of the flow reactor injects those atomic particles that have been involved in bonding over the period. It ‘feeds’ reactions that have already been observed. This should encourage more composite particles of similar types and more closely explore a dominant particle type. The second variant injects random atomic particles. This should encourage diversity of smaller composite particle types.

The decanting reactor outflows the largest 40 unique particles that are generated through a run to a second reactor. This second reactor is populated with 100 copies of each of these outflow particles, and is then run. This can be thought of as decanting the heavier particles and allowing them to interact separately. This is a way to observe if larger composite particles have different activity from smaller particles.

Each reactor is run for 100,000 reaction attempts. For each of the 20 previously found atom sets, denoted reactor 0–19, we run each reactor experiment 10 times.

Results

Figures 3–6 show results for six of the reactors, the two with the highest number of unique particles (reactors 13 and 18), two with the lowest (8 and 0) and two with a median number (7 and 2). Table 1 shows the average number of unique particles found under each reactor type. A unique particle is defined as one with a structure or composition that has never previously been seen in the run. In Spiky RBN spikes are uniquely identifiable so structural uniqueness also includes particles that have the same geometric structure but the bonds are between different spikes. We consider reactors with a high number of unique particles to be more reactive and those with a smaller number more inert.

Figures 3–6 show the activity measure ΔA (Equation 2) for each reactor with an $n = 10,000$. Each boxplot represents the ΔA of the 10 repetitions of the reactor. So a low median would indicate activity over those 10,000 reactions was low (most reactions do not affect binding property), a high median represents high activity (most reactions do change binding property). The spread of each boxplot shows the difference between the 10 runs, high spread indicates the system is very dependant on the order of particles selected for binding, low spread the system is less dependant on the other of selection.

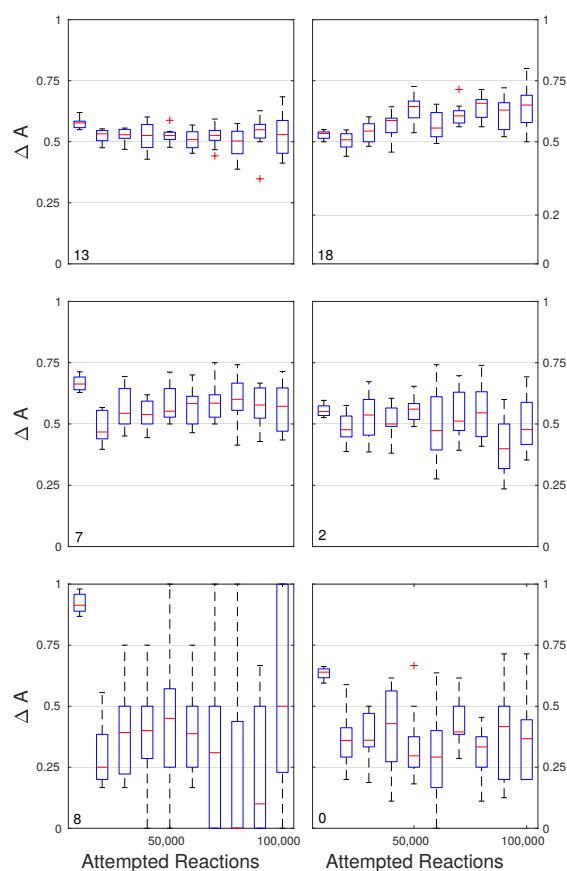


Figure 3: Mass Conserving: Plot shows the change in activity every 10,000 reactions attempts. The boxplots show the spread in behaviour over the 10 individual runs.

Mass Conserving

Figure 3 shows results for the mass conserving reactor. The reactors with low numbers of unique particles (reactor 8 and 0) show a much higher spread in activity between runs. This is likely because the overall activity of the system is much more reliant on which specific particles emerge. Reactor 8 on average creates only 93 unique particles over the run. The huge spread in activity is the result of precisely when specific particles emerge in each run. In the top two systems this is obscured since unique particles are more abundant.

The two top reactors show a difference in behaviour over time. Since the reactors are mass conserving, the later in the run the more likely it is that larger particles will be involved in reactions. For example the activity of reactor 13 is similar between different points in the run and averages around 0.5. This suggests that there is a balance between reactions that change binding property and those that do not, and that this balance is not affected by the contents of the reactor. Reactor 18 however shows an increase in activity value later in the run. This suggests that successful reactions later in the

Reactor	Mass Conserving	Food Flow	Food Random	Decanted
13	1445	2924	1478	1819
18	1286	2915	1302	2112
2	488	1117	486	1263
7	468	787	459	1156
0	182	251	180	307
8	93	181	97	293

Table 1: Average unique particles discovered

run are more likely to result in changes to binding property. Since the only difference between earlier and later in the run is the size of particles involved we can say that composite particles have different behaviours from the atomic particles in this reactor.

Flow Food

Figure 4 shows results for a flow reactor which inputs atomic particles that were part of previous observed reactions.

This reactor type resulted in substantially more unique particles being found in each reactor instance. Over the run the concentration of intern particles in the system decreases as they are replaced by reactive ones so the system finds more reactions overall.

As with the mass conserving example we see spread between runs increase as the reactors become more inert. Comparing Figure 3 and Figure 4 shows that the flow food reactor in general have a lower spread between runs. This is likely for the same reason, since food flow reactors are more reactive the system is less sensitive to exactly which particles are observed.

Flow food reactors also show a decrease in activity measure compared to the mass conserving counterparts (Figure 3). A decreased activity measure indicates that more reactions are found which do not change the binding property. This behaviour is likely caused by reactions that progressively add the same particle to a structure without changing its binding property. Because these reactions are successful the flow reactor keeps feeding them, increasing the concentration of reactant particles and making the reactions more likely.

Flow Random

Figure 5 shows results for a flow reactor which inputs random particles from the reactor atomic set.

Overall this reactor type shows extremely similar behaviour to the mass conserving reactor, both in terms of activity measure and number of unique particles produced. This is not surprising since random inflow of particles will maintain the concentrations of atomic particles in the system making it equivalent to no flow. The similarity indicates that the concentration of composite particles in the mass

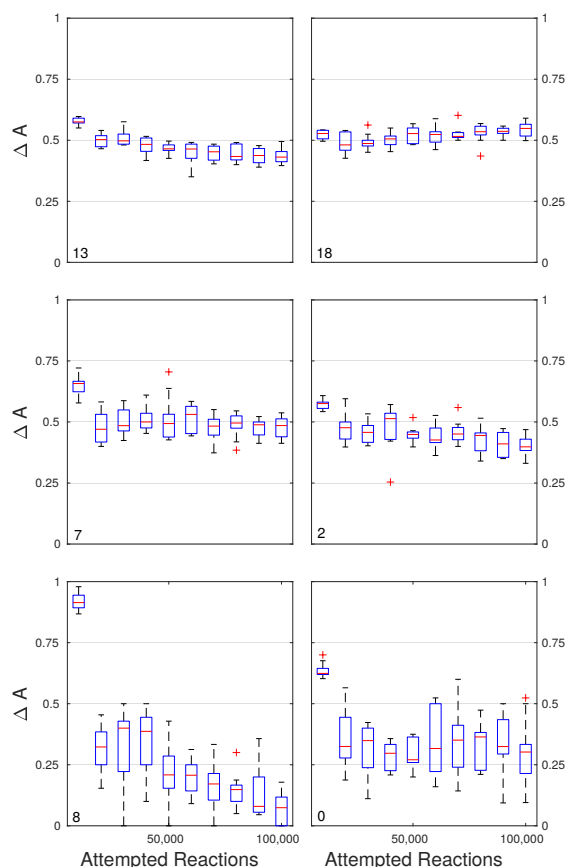


Figure 4: Flow Food: Plots showing the change in activity for each flow food reactor every 10,000 reactions attempts.

conserving reactor is relatively low. If randomly replacing particles does not effect behaviour that means that mostly atomic particles are being replaced by other atomic particles. This suggests that the mass conserving reactor never reaches a point where its behaviour is dictated by the reactivity of composite particles.

Decanted Reactor

Figure 6 shows results for the decanted reactors. Each decanted reactor is populated with 100 copies of the 40 largest particles discovered in the Mass Conserving Run. The decanted reactors are then run under mass conserving conditions for 100,000 reaction attempts. While each reactor instance is initialised with the same total number of composite particles, the total mass of each reactor instance varies depending on the total mass of the 40 largest particles discovered.

Table 1 shows that the decanted reactors are even more reactive then the atomic ones. In all cases substantially more reactions finding unique particles. This trend is not necessary uniform with reactor 18 showing a larger increase then

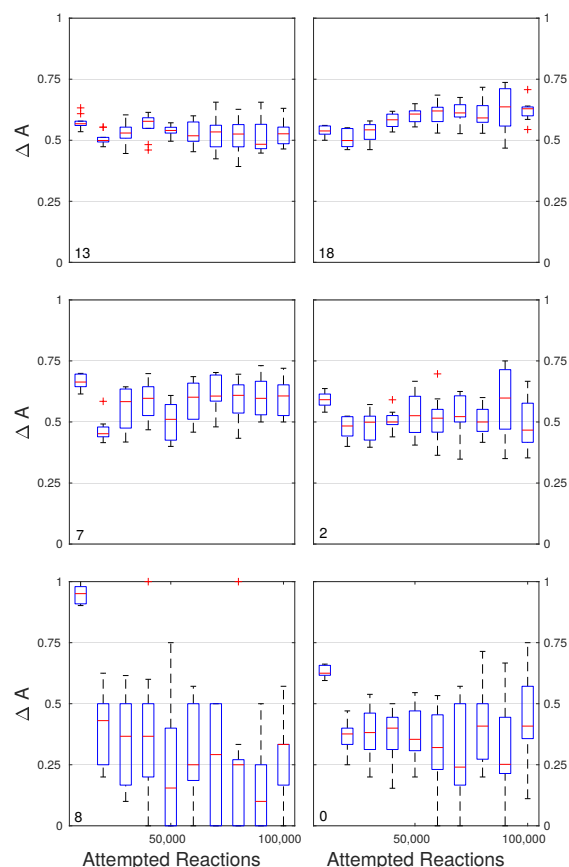


Figure 5: Flow Random: Plots showing the change in activity every 10,000 reactions attempts.

reactor 13, the same is true of reactors 8 and 0. This suggests that the behaviour of the reactor differs based on concentration of composites.

This change in behaviour is also observed when comparing Figure 3 and Figure 6. For example reactors 8 sees a very large decrease in activity measure. The majority of reactions in the decanted reactor now result in no change to the binding property. This explains the substantial increase in unique particles found since reactions of the form $X-X + X \rightarrow X-X-X$ are more likely not to change the binding property, meaning more X particles can be added to find new longer sequences.

Conversely reactor 0 shows an increase in activity measure. This means the system will find it harder to simply combine particles to create new larger composites since each reaction is more likely to change the binding property.

A similar trend can be seen between reactors 18 and 13. The decanted reactor 18 shows a much larger decrease in activity measure corresponding with a much larger increase in unique particles found. Meanwhile reactor 13 shows a smaller decrease in activity measure and a smaller increase

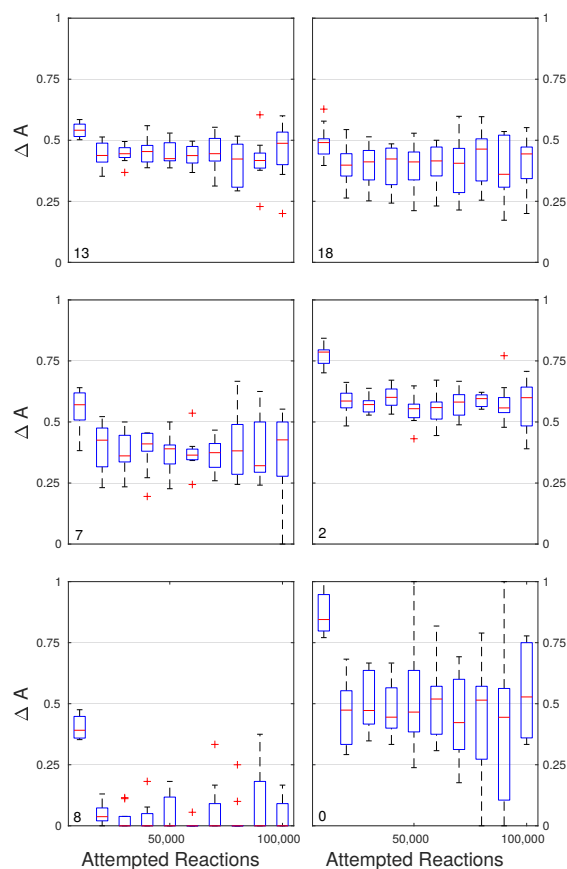


Figure 6: Decant: Plots showing the change in activity every 10,000 reactions attempts.

in unique particles found.

Overall this suggests that the main driving force behind new composites being discovered are polymer-like reactions, which are less likely to change the binding property of the product, allowing for further addition. However, this observation cannot be generalised. Reactors 2 and 7 show very similar increases in unique particles, yet the activity measure of reactor 7 decreases, while increasing for reactor 2.

Some reactors also show a difference in activity measure spread between Figure 6 and Figure 3. An increase in spread between decanted runs (such as in reactor 18) suggests that the system is more sensitive to the specific composite particles used to populate the run. Alternatively decreased spread between runs (observed in reactor 2) implies that most composites have very similar behaviour, so the system is less sensitive to which ones are picked to populate the reactor.

Summary of Results

The results presented here illustrate some of the behavioural richness that the Spiky-RBN system can exhibit.

We see that different particle sets exhibit varying be-

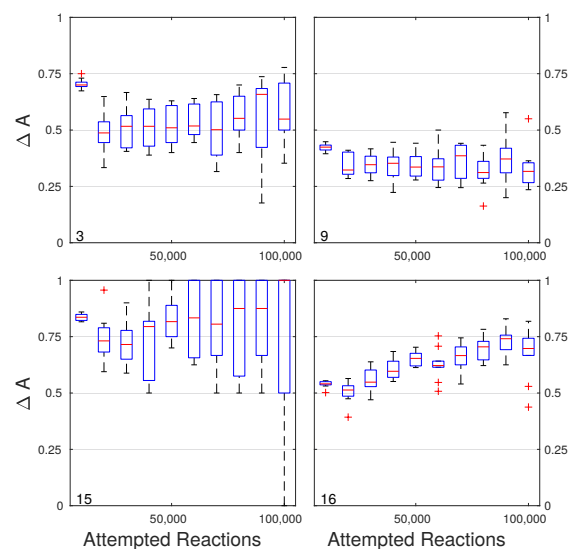


Figure 7: Other mass conserving reactors with interesting behaviour.

haviours over system runs and react in different ways to changes in reactor rules.

Figure 7 shows the range of behaviours exhibited by some of the other reactors that warrant in depth investigation. Reactor 3 shows growth in spread between runs over time possibly indicating diverging behaviour depending on which particles emerge. Reactor 15 shows very high chaos suggesting most reactions change the binding property and reactors 16 and 9 show strong trends towards order and chaos respectively.

Our measure for activity highlights possible behavioural differences in a way that can be related back to underlying mechanisms in the system. This is important if we hope to reason about the behaviours and the system as a whole.

The analysis methods used in this paper do obscure many details about system behaviour, for example reaction dynamics and particle composition. However, the functional grouping on which the activity measure is based preserves some of that information. Further analysis on the functional groupings can highlight other aspects of system behaviour.

Our reactor results for Spiky-RBN illustrate how the configuration and initial state of a given AChem (here, varying flow schemes under the same reaction rules) have a profound effect on the resulting reactions.

In the next section we discuss how this might be exploited in engineering emergent systems.

Discussion

Figure 8 depicts a layered structure that can apply to complex systems in general; here it is instantiated with an in-

stances of our Spiky-RBN components. The bottom two (green) layers represent the explicitly implemented core components of the system, defining the “physics” of interactions, the specific environmental conditions, and the initial conditions. The top (pink) layer represents the observable data as gathered through analysis tooling for the particular instance.

Between these top and bottom layers is a (blue) layer of implicit “interaction channels”. These channels represent the emergent ways in which (bottom layer) core components can directly or indirectly influence each other, thereby giving rise to (top layer) observable emergent behaviours. The channels capture the emergent interactions between the core components, and will vary depending on the specific implementation of these components. In turn, observed emergent behaviour is a product of these channels. These channels are represented in their own layer in order to highlight that they are not explicitly encoded by the bottom two layers, and that the emergent behaviours observed depend on what interaction channels are present through the runtime.

For example, in natural chemistry the core properties of atoms give rise to a large number of such interaction channels. Different types of bonds, restrictions due to geometry, and endogenous environmental effects (such as exothermic or endothermic reaction effects) are all examples of emergent interaction channels. These channels are closely linked because they all derive from the same core properties.

We use this layered structure to discuss engineering emergent systems towards completing tasks. Reasoning about the behaviour of complex and emergent systems is inherently challenging. Emergence implies that observed behaviours are not readily apparent from the underlying components and their interaction rules. Interaction channels present themselves only when running the system, and are difficult to identify and define beforehand. For a system to show emergence and open-endedness, there has to be a rich set of interaction channels. To provide such channels, the system must be able to create new channels for itself, and exploit them to produce new behaviours. Overall, a relatively small set of core interactions has to provide a large rich set of possible interaction channels. These channels will be closely linked, since they all derive from the same small set of core components.

Because the system must generate its own closely linked channels in order to show emergence, it can be difficult to reason about it until after it has run. In order to reason about the effects of a specific interaction channel, one would have to observe what happens to the system when only that channel is changed. The only way to change an interaction channel is to change one of the core interactions from which it emerges. But because all the interaction channels depend on the core interactions, there is no way to guarantee that any changes made will affect only the channel under investigation. So in turn, any change in observed emergent behaviour

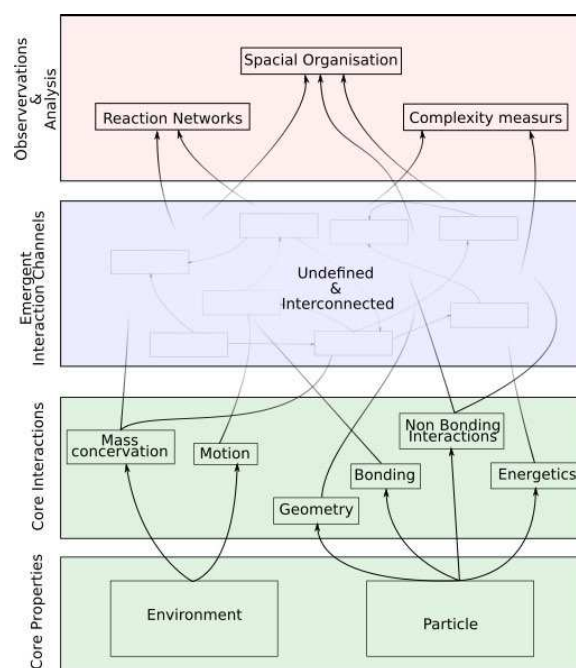


Figure 8: A representation of an instance of a complex system. The lower (green) layers represent explicitly defined components. The middle (blue) layer represents the emergent interaction channels. The top (pink) layer represents a snapshot of the running instance, containing information that has been selected for observation.

cannot in general be attributed solely to the one channel being investigated. This is a general issue with any complex system, artificial or emergent.

In a top-down approach to engineering such systems the developer needs to understand not only the lower two implemented layers (Figure 8), but also the interaction channels that emerge from them. Attempting to engineer specific emergent behaviour is complicated by the need to predict or define the emergent interaction channels. The restrictions that would have to be placed on the interaction channels in order to get specific emergent behaviour would likely rob the system of its ability to adapt and show novel emergence.

An alternative approach is to focus on the lower layer components (Figure 8) comprising the core properties and interactions. Any constraint that is maintained by all core interactions will also be maintained by the emergent interaction channels. That constraint will therefore be present in the observed emergent behaviours, provided the system is closed and has no external stimulus (so interaction channels are derived only from core mechanics). Here constraints on the core properties could be used to engineer the system towards a specific goal. Those constraints could be engineered through environmental configurations and initial conditions.

The emergent behaviours of the system are then the steps the system takes in order to reach and maintain those con-

straints. These emergent behaviours become a by-product of the system, and not the goal. The system can be thought of as a tool that gives not only a possible solution to the encoded problem but, through its emergent behaviours, gives an idea of how that solution was reached.

A top down approach to analysis of the system can then be taken (pink box of Figure 8). This analysis provides a snapshot of the instance, and is not necessarily capable of fully describing the interaction channels that lead to the observed behaviours. It could however provide an idea of what algorithm the system used to find a solution. This information could in turn be used to suggest what further constraints and initial conditions are likely to produce good results, without having to resort to changing the core “physics” of the system. For example, an evolutionary algorithm could be used to explore configurations and initial conditions to generate better instances of the system.

Conclusion and Future Work

In this paper we explore the range of dynamic behaviours our system can exhibit by looking at different instantiations of the model. We observe that there is a richness in the dynamics which can be reached without changing the core “physics” of the system.

A parallel can be drawn to Game of Life. Initial patterns define different instances of Game of Life much like the choice of atomic particles and environmental constraints define instances of Spiky-RBN. System dynamics vary substantially depending on the instance. Some Game of Life instances reach a “dead” state, others a dynamic final state; there are instances that show self replication and even Turing completeness. Even complicated instances rely on basic emergent elements, for example glider gun patterns are used to propagate information in more complicated instances. Game of Life is however very fragile: minimal differences in instantiation can completely destroy complex dynamics.

Symbolic approaches to AChems are at the other end of the spectrum. Dynamic behaviour is explicitly encoded into the “physics” (the set of predefined reactions). As such most of the instantiations of the system will behave similarly. While this is useful in guaranteeing that behaviour will be as expected, it inhibits the ability for new mechanisms to emerge.

Sub-symbolic AChems aim to be in the middle of this spectrum. By not predefining the reactions, the system can be richer than symbolic Artificial Chemistries in terms of emergent behaviour. Meanwhile mechanisms in the reaction algorithm constrain the emergent properties to only those that maintain predefined conditions, giving some level of control.

Future work will focus on engineering instances of Spiky-RBN. By identifying desirable properties and exploring the emergent reaction channels that give rise to those properties we aim to build instances with specific behaviour, for exam-

ple, ones generating composite particles with regular structures or generating reaction networks with cycles. This can be thought of as a top-down process, leaving the “physics” of the system fixed, and exploring the behaviours of instances guided by experimentation. This process would be akin to identifying gliders in Game of Life, and exploring how they arose, in order to build glider guns.

As Artificial Chemistries mature and we move from an exploratory to an exploitative phase, methods like this will aid us in creating specific behaviours.

Acknowledgment

The authors would like to thank Leo Caves for his constructive comments. Krastev is funded by an EPSRC DTA PhD studentship.

References

- Bonchev, D. and Trinajstić, N. (1977). Information theory, distance matrix, and molecular branching. *The Journal of Chemical Physics*, 67(10):4517–4533.
- Ciencias del Espacio, I. (2000). Measuring mutual information in random boolean networks. *Complex systems*, 12:241–252.
- Dehmer, M. and Mowshowitz, A. (2011). A history of graph entropy measures. *Information Sciences*, 181(1):57–78.
- Karremans, G. (1955). Topological information content and chemical reactions. *The bulletin of mathematical biophysics*, 17(4):279–285.
- Kauffman, S., Peterson, C., Samuelsson, B., and Troein, C. (2003). Random boolean network models and the yeast transcriptional network. *Proceedings of the National Academy of Sciences*, 100(25):14796–14799.
- Kauffman, S. A. (1969). Metabolic stability and epigenesis in randomly constructed genetic nets. *Journal of Theoretical Biology*, 22(3):437–467.
- Krastev, M., Sebald, A., and Stepney, S. (2016). Emergent bonding properties in the spiky rbn achem. In *Proceedings of the Artificial Life Conference 2016*, pages 600–607. MIT Press.
- Rashevsky, N. (1955). Life, information theory, and topology. *The bulletin of mathematical biophysics*, 17(3):229–235.
- Wang, X. R., Lizier, J. T., and Prokopenko, M. (2011). Fisher information at the edge of chaos in random boolean networks. *Artificial life*, 17(4):315–329.