

Bondable Cellular Automata

Mark Hatcher¹, Wolfgang Banzhaf¹ and Tina Yu¹

¹Memorial University of Newfoundland, St John's, NL A1B 3X5, Canada

mhatcher@mun.ca, banzhaf@mun.ca, tinayu@mun.ca

Abstract

We present the Bondable Cellular Automata model, which uses simple 1-dimensional, binary cellular automata as the base atomic elements of an artificial chemistry. Reactions are dependent upon an emergent, 'resolution independent' observable, measurable for individual or composite cellular automata structures. We discuss the rationale behind our choice of observable, 'mean polarity', and behind the choice of a bonding mechanism based on this observable. From simple experimentation we observe that using cellular automata as the underlying dynamical system coupled with mean polarity as the reaction success criterion shows potential to support sustainable emergent behaviour.

Introduction

The general model for an artificial chemistry consists of the {S (material), R (reaction rule set), A (algorithm)} triplet, with R applied to S according to A (Dittrich et al, 2001). Typically R is hand-coded and applies to a single level of structural hierarchy, while S is composed of atomic types with little or no internal dynamic. This provides ease of analysis at the expense of flexibility.

This work applies a recent, alternate approach of 'sub-symbolic' artificial chemistry, described in RBN-World (Faulconbridge et al, 2009), and in (Faulconbridge et al, 2010), where reactions can apply at any level, with reaction success based upon an internal dynamic of colliding bodies.

Here cellular automata (CA) are used to provide the internal dynamic. By using CA as base atoms it is possible to construct 'Bondable CA' (BCA) systems where the application of R is dependent upon an emergent, possibly 'resolution independent' observable of individual (atomic) or composite (molecular) CA structures. The bonded CA within composite structures are able to exchange state information, introducing new dynamics to the CA, and potentially leading to emergent behaviour and structure in the chemical system.

Why use Internal Dynamics?

If we allow each body (be it an atom or composite molecule) in an artificial chemistry to possess an internal dynamical system, then we can allow the reactions which occur between bodies to affect the configuration of their systems. Moreover we can allow the reverse: for the configuration of the bodies' systems to affect their ability to react. This way a feedback loop is formed.

This can be achieved by forming reaction rules that are based upon the value of an observable of each body's internal dynamical system. If the chosen observable is measurable for any body of any internal structure or size, be it a single atom or complex molecule, then potentially any two bodies can react with each other, even if their size and structure differ. This allows composite bodies of arbitrary size and structure to be constructed, and allows their dynamical systems to couple and interact. As they grow these composite bodies will take the form of an increasing hierarchy of systems within systems within systems, all interacting with each other.

Further, in the BCA model we have chosen to use an observable that reflects change in the configuration of a body's dynamical system as it occurs, whether or not this change has been caused by reaction with another body. So two bodies that meet the criterion to react with each other at one instance in time might not do so at another instance because of interim change in the values of each body's observable; and vice versa. Similarly, when two bodies react to form a larger, single body, the interaction between their dynamical systems will cause changes to each over time, and might lead to structural instability. If, according to the values of their observables, subcomponents of the single body no longer meet the criterion to remain bonded then decomposition of the single body will occur. Such decomposition will have a knock-on effect upon the internal dynamic of the remaining body, which in turn may cause further, future decomposition, and so on. Thus we have introduced an element of spontaneity to the reactions that take place, allowing them to occur well after or even in the absence of collision between bodies.

So using bodies that possess internal dynamical systems and basing reaction rules upon a suitable observable of those systems allow a rich set of reaction types to take place, between bodies unbounded by size, thus providing a sophisticated platform upon which we can model and explore multi-layered dynamical systems and how they interact.

Why use Cellular Automata?

The Cellular Automaton is an ideal underlying system for a sub-symbolic chemistry. A tenet of artificial life research is that complex, interesting behaviour may arise from a deceptively simple mechanism, and the cellular automaton is a deceptively simple dynamical system; deceptive because intuition would suggest that from simple rules must emerge simple outcomes. Yet we knew in 1966 that CA with large, intricately constructed transition rules were capable of

universal computation (von Neumann, 1966), and we know now that 1-dimensional CA are capable of performing universal computation (Cook, 2004), or of generating pseudo-random numbers that pass all current tests for randomness (Wolfram, 2002). It is even speculated that CA-type processes are at work in nature, such as the colour patterns generated in the spiral of the Cone Snail, or in the formation of snowflakes (Wolfram, 2002).

Further, a strand of research that allows the CA's transition rule to be altered during iteration has allowed computations to be performed that are not possible with 'simple', standard 1-dimensional CA (Mitchell et al, 1993, 1997), (Kano and Wu, 2003), as opposed to those CA which require a partition of cells to store the 'program' to be executed, such as (Cook, 2004). Other work has explored the capability of 2-dimensional CA to perform computations, including universal computation (Sapin et al, 2007).

Thus simple CA, when allowed to interact, have the potential to produce a wider range of behaviours than in isolation. Placing them within the framework of an artificial chemistry, with the continual 'composition, decomposition, re-composition' processes of combinations of CA, allows this potential to be explored in a more open-ended, emergent manner. Figure 1 shows how linking two circular, 1-dimensional CA between just one cell of each can lead to large changes in configuration within a short time, and in automaton B's case interferes with its otherwise short and simple cycle of just 8 configurations.

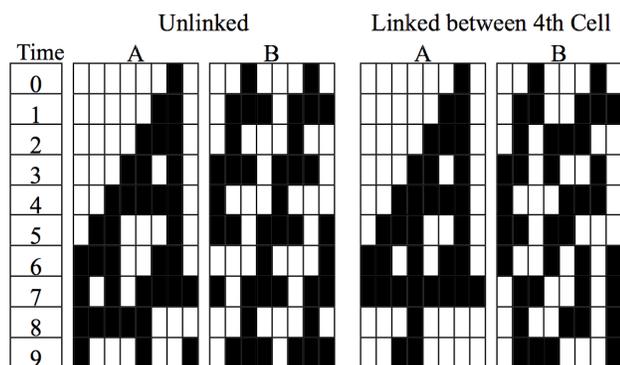


Figure 1: The impact of linking two circular, 1-dimensional cellular automata between one cell of each. The linked cells see each other's neighbourhood when updating. Note that the simple, cyclic configuration of unlinked automaton B becomes disturbed through interaction, and that both CAs' configurations are affected.

This paper describes the BCA model from different perspectives: from the perspective of the individual CA cells, from the perspective of the CA, the *atom*, and from the perspective of the *molecules*, composed of many CA. It describes the *reaction mechanism* that allows composition, or *bonding*, and decomposition, or *unbonding*, to occur between *bodies*, be they atoms or molecules. It provides an example of collision leading to bonding then subsequent, spontaneous decomposition. It describes and discusses the reasoning

behind and the impact of using *mean polarity* as the underlying observable upon which reactions are based. Finally conclusions are drawn, which will steer the direction of future work.

The BCA Model

The Bondable Cellular Automata model is an artificial chemistry that uses 1-dimensional, binary CA for its base (atomic) elements. These atoms bond to form molecules and molecules further bond to form larger molecules of arbitrary size. Adopting the approach in (Faulconbridge et al, 2009), the reaction rules between bodies (whether they be singular atoms or composite molecules) are not explicitly defined for each type of body, but instead reaction success is based upon the comparison of the value of a single observable for each body; an observable that is based upon the internal CA configurations yet can be measured for any constructible body.

Model Perspectives

Since BCA is an artificial chemistry based upon cellular automata it is useful to describe and observe it from different perspectives.

Sub-atomic Level. BCA can be viewed as a collection of interacting cells. Each cell updates its binary state each iteration, according to the collective state of its perceived neighbourhood of other cells and its assigned transition rule. Figure 2 illustrates the cell's perspective.



Figure 2: The central (green) cell perceives its neighbourhood as its left and right neighbours, which for rule-width 2 is the 2 cells either side (in blue), but BCA allows a cell to perceive the neighbourhood of a cell in another CA, allowing information to flow between them.

Atomic Level. BCA can be viewed as a collection of bondable atoms. Each atom is a circular CA and bonding causes cells in one atom to link to cells in another atom, as shown in Figure 3. Atoms with positive polarity can bond to atoms with negative polarity while atoms with the same polarity cannot bond and atoms with zero polarity are always inert; this is the 'bonding criterion'. At any instant in time an atom's *polarity* is defined as:

$$\text{polarity} = \text{count}(\text{cells in state '1'}) - \text{count}(\text{cells in state '0'})$$

So, in Figure 3 the upper atom has 5 cells in state '1' and 7 cells in state '0', giving it a polarity of $5 - 7 = -2$ (polarity sign: *negative*). Similarly the lower atom has polarity of $12 - 0 = +12$ (polarity sign: *positive*). Hence these atoms can bond.

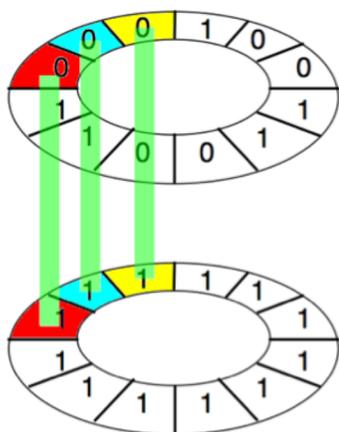


Figure 3: Two bonded atoms are linked between pairs of cells (colour-coded accordingly). Each cell sees its partner's neighbourhood when updating.

Bonding causes cells in one atom to link to cells in the other atoms. Cells can have only one link at a time. Linking causes a cell to view the corresponding cell's neighbourhood rather than its own when updating. Atoms can be bonded to many other atoms, limited only by the availability of unlinked cells. This leads to a rich, complete-graph structure for bonded atoms, as illustrated in Figure 4.

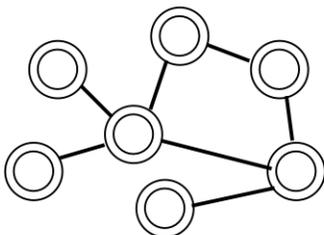


Figure 4: The underlying atomic structure of a molecule. Each atom is a circular CA, and each CA has its own transition rule.

Molecular Level. BCA can be viewed as an artificial chemistry, a collection of molecules that collide in pairs and bond.

Due to two body reactions each molecule consists of exactly two (sub) molecules or (conceptually) of a single atom. Figure 5 shows how this leads to a binary-tree structure for molecules, with each parent molecule containing two child molecules.

When molecules collide, if they meet the bonding criterion, then they bond. If they bond this may cause changes to the internal configuration of the CA, which in turn affects polarity, which may lead to sub-components unbonding, which is explained later in the Bonding Example section.

Molecules bond by linking pairs of atoms, one in each pair from each molecule, and how these pairs are chosen and linked is described in the next section. A key aspect of BCA is that the rich underlying atomic structure of molecules is hidden at this molecular level. This greatly simplifies the

description of molecular types and reactions. Also it is possible for two molecules with the same molecular identity to have different structure at the atomic level, thus allowing isomers to be modeled. Moreover, since a molecule's polarity is the emergent outcome of initial configuration, the transition rules and the topology of the underlying atomic bonds, it is possible for two molecules of the same type to possess opposing polarities and thus bond.

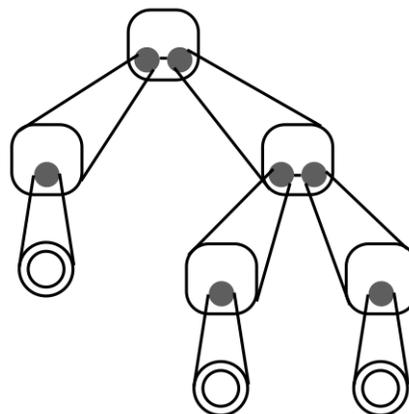


Figure 5: Molecular structure is nested, forming a binary tree. Key: Molecule (rounded square), Atom (outlined ring).

Material and Reactions

Atoms and Molecules

In BCA, 1-dimensional, unbounded (circular) binary CA form the atomic elements, their type identified by their transition rule. Atoms can bond to form molecules, and molecules can bond to atoms or other molecules to form ever-larger molecules.

Single atoms behave like standard CA, updating the state of each cell each iteration according to the application of the transition rule to the cell's neighbourhood.

When two atoms are bonded, pairs of cells between each atom are 'linked' together. When updating, linked cells still use their own transition rule but see their partner's neighbourhood instead of their own. Thus the cell states (the 'configuration') of paired atoms affect each other.

Collisions

BCA's topology is a well-mixed soup, consisting of *bodies*, which can be single (unbonded) atoms or composite molecules. At any point in time collision can randomly occur between any two bodies. Two bodies that collide are known as the reactants, and are tested to see if they will bond according to the bonding criterion.

Bonding Criterion

Reactants bond according to the mean polarity of each reactant's configuration.

Reactants with opposing polarity signs may bond (positive-to-negative or negative-to-positive) whilst any reactant with

neutral-signed polarity is considered inert and will not react with any other body.

We note that polarity is a scalable, ‘resolution independent’ observable in our model, since it can be measured for single atoms and for molecules of arbitrary size. Polarity is also an emergent property of the configuration of a body, which itself is an emergent property of the initial configuration and transition rule(s) of the underlying CA, and of the effect of the bonds between them. Hence polarity is an emergent property of the underlying dynamical systems and how they interact.

Bonding Mechanism

Atomic Level. Two atoms with opposing polarities bond by forming links between pairs of cells. Which cells link to form the ‘bond site’ is determined by comparison of each CA’s configuration at the time of collision. In keeping with the concept that ‘opposites attract’, the longest continuous run of currently unlinked state 1 cells is identified in the ‘positive’ atom, and the longest run of currently unlinked state 0 cells is identified in the ‘negative’ atom. The shorter of these runs determines the size of the bond site. Each cell in the shorter run is then linked to a cell in the longer run on a 1-to-1 basis until the shorter run is exhausted. Thus every cell has the potential to contribute to a bond site, and therefore to interact with cells in other CA, while the actual location and size of the bond site is an emergent outcome of the automata’s current configurations.

Molecular Level. When two molecules with opposing polarities collide, they form bonds between pairs of atoms, in the manner described above. The molecules will attempt to form bonds between two pairs of atoms, but in practice might form a bond between only one pair, or even not be able to form any bond at all, as the process below explains.

Each molecule is polled for its atom with most positive polarity, and its atom with most negative polarity. Again, by the principle that opposites attract, the most positive atom in the first molecule is paired with the most negative atom in the second molecule (and respectively for the other pair). These atoms bond together at the atomic level in the manner described in the previous section.

Sometimes the chosen pairs of atoms cannot bond, because one or both atoms has no free, unlinked cells with which to form a bond. In this situation any ‘fully linked’ atom is overlooked and the molecule polled for the next most positive/negative atom as appropriate. This process will continue if necessary until either a bondable pair of atoms is found or no more candidates exist. In the latter situation this will lead to the molecules bonding between just one pair of atoms, or in the extreme case not bonding at all.

Allowing two pairs of atoms to possibly form the bond between molecules maintains consistency with the concept that polarity underpins the bonding mechanism, with the most-oppositely polarised atoms in each molecule being attracted to each other and attempting to bond. The key benefit of allowing more than one pair of atoms to bond between molecules is that it allows a rich graph structure to develop at the atomic level, illustrated in Figure 4. If only one pair of atoms were allowed to bond between molecules then this structure would be restricted to a tree, providing less

opportunity for interaction between the atoms within a molecule.

Unbonding Mechanism

Unbonding occurs spontaneously at the molecular level. If a body consists of more than a single atom, then every iteration of the system the bonds between the two sub-components that form a body are tested. The test is simple: if the two atoms which are actually bonded no longer attract, then they unbond. Unbonding removes all links between paired cells in each atom, and their CA no longer interact.

This unbonding will weaken the link, and hence interaction, between the two sub-components of the body, and if it was the last bond will lead to separation of the body into smaller bodies. Those bodies will then themselves be subject to potential spontaneous decomposition, and so on.

Bonding Example

Figure 6 illustrates by example the composition of two colliding molecules, and their subsequent decomposition into two new molecules.

Let A and B be molecules in the BCA system. Further, let A be composed of sub-molecules C and D, since this fact will become useful when describing the decomposition stage.

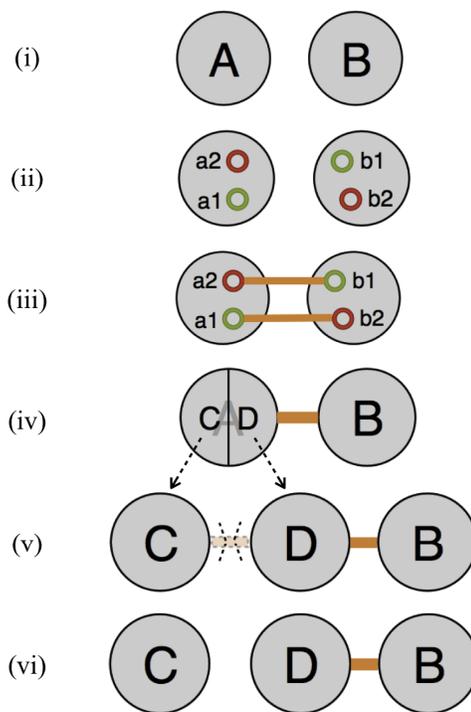


Figure 6: Example of the bonding of two molecules, A and B, and the subsequent, spontaneous decomposition of the resultant into two different molecules, C and D-B.

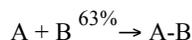
Composition. Suppose that A and B collide (Figure 6.i). They have opposing polarities and attempt to bond.

Let the most positive and most negative atoms in A be identified as a1 & a2 (respectively as b1 & b2 in B; see Figure

6.ii). Unlinked cells are available on each atom and so a1 bonds to b2 while a2 bonds to b1 (Figure 6.iii). The molecular equation for this reaction is simply:



A-B denotes composition. Since the reaction success of A + B is not guaranteed for every collision between the types, the current equation is insufficient. So during simulation our system records the percentage reaction success of collision between pairs of types and we can more accurately write the equation as:

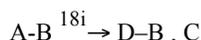


Note that this says nothing about the underlying atomic structure, thus allowing the aforementioned isomers to exist.

Decomposition. We now suppose that the formation of A-B leads to changes in the cell states of the underlying CA, through interaction between bonded atoms. These changes subsequently cause the bonded atoms in A's two sub-molecules, C & D (see Figure 6.iv), to no longer attract. Hence the bonds break and A decays (Figure 6.v).

This leads to the breakaway of C as a separate body, while D remains bonded to B, effectively forming a new molecule (Figure 6.vi).

The equation for decomposition is:



N_i represents the number of iterations the composite survived for, and the comma indicates separation. We can consider N_i as the reaction rate for a decomposing reaction. The full chain of events can be written as:



Some composites will decompose into the original two molecules that formed them, so reversible reactions can be supported by the model.

During the run of a simulation we can track the entire flow of compositions and decompositions for all molecules to derive the reaction network.

The Impact of Polarity

The choice of polarity as the basis for the bonding criterion followed a process of deduction and experimentation.

Table 1 lists the key, measurable properties of a CA and their suitability for the role. It was quickly identified that any candidate property for underpinning the bonding criterion would need to be not just resolution independent, but would also need to at least in part reflect the dynamic nature of a CA's configuration in order to be an emergent property leading to emergent behaviour.

Therefore rule width, dimension, number of cell states and size were discounted as too trivial to be useful since they remain constant or ignore the CA's configuration; they effectively reduce to static elemental types seen in the general artificial chemistry model. Likewise Transition Rule was

discarded for the above reason and further since it cannot be consistently defined for bonded structures.

	Resolution Independent?	Dynamic?
Rule width	Yes	No
Dimension (1d, 2d, ...)	Yes	No
No. of possible cell states	Yes	No
Size	Yes	Weakly
Transition rule	No	No
Configuration	Yes	Yes

Table 1: Summary of key measurable properties of CA.

This left measures based upon Cell Configuration, which fall into two broad categories: long-term measures and instant measures.

A good example of a long-term measure is cyclelength, as used for RBN-World. We can measure the cyclelength of a CA as the number of iterations required for the configuration to return to a previous state. One strength of using cyclelength is that it is an emergent outcome of bonding; as structures bond the cyclelengths of the sub-components and the combined structure can change. Also, since cyclelength is partly dependent upon other properties of CA such as transition rule and the current configuration, it could provide a valuable reflection of the nature (and specifically Wolfram classes (Wolfram, 1984)) of combined CA.

However one downside to using cyclelength is that its value for a particular body remains static until that body reacts with another body through collision. We believe that using an alternative observable, one whose value can change both because of and independent of reaction with other bodies, adds flexibility to the model since it allows spontaneity and uncertainty to what reactions occur and when.

The other downside to using cyclelength and similar measures is computational cost. Determining the cyclelength of a body in BCA requires direct simulation, since it is dependent upon not just the CA but also how they are bonded; in the worst case its time cost is the Cartesian product of the combined width of a body's CA.

Therefore initially BCA employed the instant measure of polarity at the moment of collision. Instant polarity is an emergent outcome of both internal configuration and the bonding mechanism, is computationally inexpensive (the cost is linear with respect to combined width) and allows for spontaneous unbonding. It also provides underlying consistency to the model: the bonding criterion, the location of bonds in the bonding mechanism (at molecular and atomic levels) and the unbonding mechanism can all be based upon this single characteristic.

Unfortunately during simulation the measure of instant polarity proved to be too stochastic for some transition rules, as Figure 7 demonstrates. The reaction success of colliding bodies became dependent upon the time of collision, which is randomly chosen.

Therefore the measure of mean polarity was implemented. This retains the benefits of instant polarity, including computational cost rising linearly with size, but, as seen in Figure 7, also smoothes the impact of large changes in

configuration between successive iterations. In our experimentation the mean is calculated from the instant that a body is formed, but alternative calculations, such as the mean for the last 100 cycles, could be used, allowing later changes in configuration to have greater impact upon it.

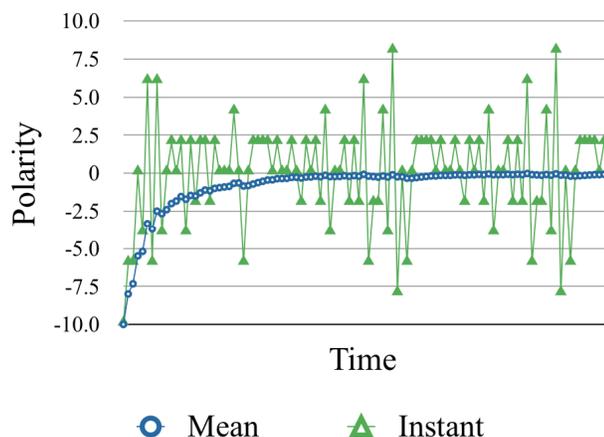


Figure 7: Polarity of a Rule 30 Cellular Automata over time. Whilst the snapshots of polarity change erratically between iterations, mean polarity smoothly settles to a steady value.

Experiments in Mean Polarity

Rule width	1
Width (number of cells)	12
Dimension	1d
Number of possible cell states	2 (binary)
Number of transition rules	256
Initial configuration of cell states	'000000000001'
<i>Number of iterations:</i>	
Isolated CA	4096
Paired System	100000

Table 2: Set-up for the Isolated and Bonded Pairs experiments.

A key question about mean polarity is whether it would be *too* smooth a measure, essentially reducing in most cases to a static value over time. To answer this question simulations were run using CA with width 12, rule-width 1, providing 256 possible transition rules, or atomic types. Each type of CA was simulated in turn in isolation for 4096 iterations. This is the maximum theoretical cyclelength for a width 12 CA and thus allows a fair calculation of mean polarity over time for the individual CA. Each CA began with the same initial configuration of a single cell set to state '1', all others '0'.

Following this we attempted to bond every possible pair of CA in turn. If they bonded then the simulation was run for a further 10000 iterations, far short of the maximum theoretical 16.7 million iterations required to cover all possible cyclelengths, but in practice sufficient time to determine the long-term sign and magnitude of the pair's mean polarity. Table 2 records the set-up for both experiments. This is the first step into the reaction chemistry of the BCA model.

Results and Discussion

Isolated CA. For 141 out of 256 (56%) of types the sign of mean polarity changes during simulation for isolated atoms. Table 3 shows that although all types begin with negative polarity approximately a quarter of types finish with neutral mean polarity and are hence inert.

Positive	75 (29%)
Negative	115 (45%)
Neutral	66 (26%)

Table 3: Tally of final sign of mean polarity of isolated CA after 4096 iterations.

That such a sizeable proportion of CA types become inert raises concern. However the 'periodic table' of types (Figure 8 in the Appendix) shows that the vast majority of types that achieve neutral mean polarity take more than 256 iterations to do so, and the opportunity for those CA to bond with others and be 'rescued' from inertia remains open during that time.

The periodic table illustrates a strong correlation between the time taken to settle and the amount by which mean polarity changes. All CA begin with mean polarity of -10. Those that finish with that same value never change polarity. Those that finish with mean polarity close to -10 tend to settle within 8 iterations, and as final polarity drifts away from -10, so the time taken to reach the new value tends to increase. We see that many CA types reach high positive polarities, notably acquiring a change of sign, and relatively speaking take their time to do so. This is good since it demonstrates that the mean polarity measure is dynamic in the majority of cases for long enough to present a 'window of opportunity' for different reactions to occur.

Bonded Pairs.

Total possible unique pairs	32640
Of which bonded	8625 (25%)
<i>Change seen in mean polarity value:</i>	
Individual CA	14049 (81%)
Paired System	7921 (91%)
<i>Change seen in mean polarity sign:</i>	
Individual CA	5799 (34%)
Paired System	3938 (45%)
Both CA and paired system changed polarity	371 (4%)
Both CA changed polarity but paired system did not (the changes 'cancel out')	253 (3%)
One CA's polarity remained stable while the other's and the paired system's changed	402 (5%)

Table 4: Summary Data for the Bonded Pairs experiment.

Table 4 summarises the results of attempting to bond in pairs. The listed percentages for changes seen in polarity are calculated as proportions out of the total number of pairs that succeeded in bonding. We observe that in around 4 out of

every 5 of such cases the act of bonding re-introduces a dynamic element to the value of polarity of the individual CA, and in around 1 out of every 3 cases also leads to a sign change. We also see that in almost half of cases following bonding the paired system's polarity sign changes. This reinforces the view that mean polarity is an emergent outcome of both the underlying CA configurations and the bonding mechanism, not just because it aggregates the values for the subcomponents, but crucially because the subsequent interaction caused by bonding leads to changes in the subcomponents' values themselves.

In the 5% of cases where bonding causes a sign change for one of the CA, if this change were rapid then during a full simulation this would lead to rapid decay of the combined body, releasing both atoms near-instantly back into the soup. This provides the potential for bonding to cause the appearance of unstable molecules, leading to rapid chain reaction; a complement to the slower decay modeled in the unbonding mechanism. In such paired systems, the CA whose sign remains stable is also of interest, since it might have catalytic properties, causing change in the CA it bonds with whilst itself remaining unchanged in polarity sign. Seeking and identifying CA atoms with this property, and possibly even molecules too, is a further step in the research.

Additionally, we observed that bonding causes an overall drift away from neutrality for the bonded pairs, so the act of bonding leaves proportionally fewer inert bodies in the system than if CA were left to iterate in isolation.

These results suggest that rather than locking CA into inert structures with (near) static mean polarity as was feared, bonding could be a self-sustaining process, keeping the system active.

Conclusion and Future Work

The modeling and polarity experimentation suggest that BCA shows the potential to allow behaviour and structure to be emergent properties of both the bonding mechanism and the underlying CA configurations. More generally it allows us to study the results of the interaction that occurs between simple dynamical systems when they are placed within the framework of an artificial chemistry.

By allowing reactions to occur sometime after or even in absence of a collision BCA is also able to model useful chemical concepts such as variable decay rates, spontaneous reactions derived from internal configuration (rather than due to external trauma), isomers, and catalytic behaviour. Thus we believe that it is a worthy candidate system for the study of emergence.

The experimentation suggests that mean polarity could be an ideal resolution independent observable on which to base reaction success, possibly leading to a positive cycle in system behaviour, where reactions lead to changes in the internal structure of bodies leading to the potential for further reactions. However to assess this further experimentation is needed, including full simulation runs where numerous molecule of many types are present, and are allowed to react to form much larger bodies.

Mean polarity is only one of many possible ways to determine reaction success. Other candidate observables based upon CA configuration exist and can be explored too,

including the possibility of basing reaction success upon a family of measures, or on higher moments. For whichever set of observables we select, we need to strike a balance.

The experimentation using instant polarity has shown that using short-term measures which are based on a computationally inexpensive snapshot of the configuration can lead to essentially stochastic behaviour. Conversely, other long-term measures, such as cyclelength, are less sensitive to short term configuration changes but have increasingly large computational overhead as larger molecules appear, which impinges upon the scalability of experimental simulation. So other candidates will need to be able to display the balance that the use of mean polarity so far achieves, in tempering stochastic influence whilst keeping the computational overhead low.

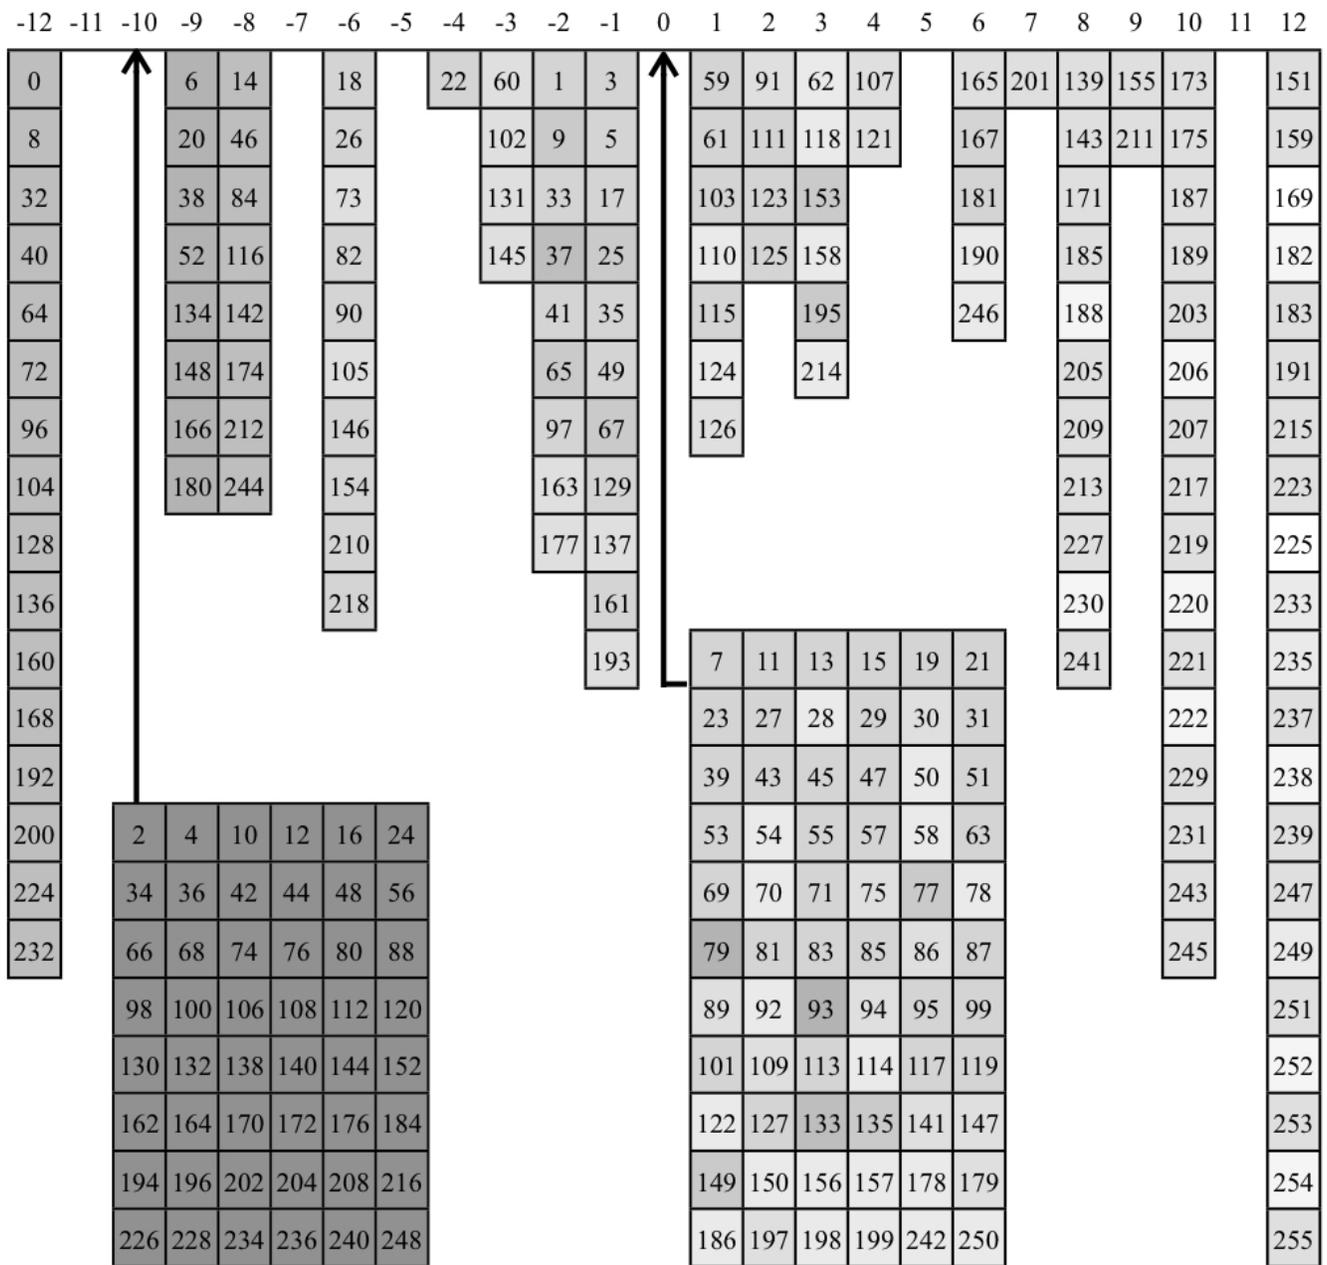
In the experimentation so far the rate of iteration for the CA has been identical to the rate of iteration for collision. So every time the system performs a collision (or a set of simultaneous collisions) it also updates the configuration of every CA. This need not be so, since we can instead allow the CA to operate in a different time frame and iterate them an arbitrary number of times between each collision.

Further work will examine the reaction networks formed by full simulation of the model, assess the impact of allowing the CA configurations to iterate at a different rate to the collisions and assess the model's viability for application to the modeling of other domains.

References

- Cook, M. (2004). Universality in Cellular Automata. *Complex Systems*, 15(1):1-40.
- Dittrich, P., Ziegler, J. and Banzhaf, W. (2001). Artificial Chemistries – a Review. *Artificial Life*, 7(3):225-275.
- Faulconbridge, A., Stepney, S., Miller, J. and Caves, L. (2009). RBN-World: Sub-symbolic artificial chemistry. In *Proceedings of the Tenth European Conference on Artificial Life*, 10:371-378.
- Faulconbridge, A., Stepney, S., Miller, J. and Caves, L. (2010). RBN-World: The Hunt for a Rich Artificial chemistry. In *Proceedings of the Twelfth International Conference on the Synthesis and Simulation of Living Systems*, 12:285-292.
- Kanoh, H. and Wu, Y. (2003). Evolutionary Design of Rule Changing Cellular Automata. *Lecture Notes In Computer Science: Knowledge-Based Intelligent Information and Engineering Systems: 7th International Conference*, 1:258-264.
- Mitchell, M., Hraber, P. and Crutchfield, J. (1993). Revisiting the Edge of Chaos: Evolving Cellular Automata to Perform Computations. *Complex Systems*, 7:89-130.
- Mitchell, M., Hraber, P. and Crutchfield, J. (1997). Evolving Cellular Automata to Perform Computations. *Handbook of Evolutionary Computation*, Oxford University Press.
- von Neumann, J. (1966). *Theory of Self-Reproducing Automata*. University of Illinois, Urbana.
- Sapin, E., Bailleux, O., Chabrier, J-J. and Collet, P. (2007). Demonstration of the Universality of a New Cellular Automaton. *International Journal Of Unconventional Computing*, 3(2):79-104.

Appendix



Settled by:



Figure 8. A 'periodic table' for the transition rules in BCA. The value at which mean polarity settles for each transition rule in the Isolated CA experiment, with a starting configuration of a single bit set to '1'. The scale across the top shows the mean polarity. The number in each box indicates a transition rule. The greyscale shading shows how many iterations it takes for the value of mean polarity to settle to its final value. While individual shades may be hard to discern the general trend is clear.