The Blind Watchmaker's Workshop: three Artificial Chemistries in the context of Eigen's Paradox

Simon Hickinbotham[1], Adam Faulconbridge[2] and Adam Nellis[1]

[1]Department of Computer Science, [2]Department of Biology
York Centre for Complex Systems Analysis, University of York, Heslington, YO10 5DD, UK
sjh@cs.york.ac.uk www.yccsa.org

Abstract

We use Artificial Chemistries (ACs) as a way of addressing problems in Artificial Life (ALife) and evolution, by considering Eigen's paradox — small replicators with poor fidelity can not encode sufficient information to build a replicator with improved fidelity. We describe three AC case studies for different periods in the early evolution of the earth. From these, we discuss more general properties that are useful for ACs to possess for evolution, and compare our properties to those described by other authors.

We do not present a resolution of Eigen's paradox; rather we demonstrate a way of thinking about AC in the context of early evolution. Eigen's paradox is one key issue in this period. We use ACs as a model paradigm and from these we extract relevant properties that can be considered separately from the specific ACs that informed them; these properties can be used to inform design and analysis of future ACs.

Introduction

Artificial Chemistries (ACs) are a useful basis for experiments in Artificial life and evolution. Approaches to ACs in this area tend to emulate the 'central dogma' of biology, whereby information is encoded on macromolecules analogous to DNA, RNA, and proteins. This is a difficult modelling challenge due to the size of the molecules relative to their atomic constituents, and the complexity of the interactions between them. An alternative to this approach is to seek ACs that more closely resemble models of the early evolution of life on earth which do not have such a constrained linear flow of information. These stages may be easier to model due to their relative simplicity, and from these models, a set of properties can be derived that allow better models of the macromolecules of the central dogma of biology to be constructed. However, this pathway is not well understood in paleobiology and is therefore difficult to emulate. Recent work in paleobiology suggests that there were many different modes of evolution before the central dogma of biology became prevalent [25]. These modes exploit a more vague distinction between template (genotype-carrying) molecules and machine (phenotype) molecules. In this paper, we report work on ACs carried out separately by the three authors, that collectively emulate this period in the history of life.

One of the key problems an AC must handle is that any route from pre-biotic chemistry to the central dogma of biology must resolve *Eigen's paradox* [5]. This is Manfred Eigen's observation of the following cycle:

- Low-fidelity replicators are only able to preserve small genomes reliably.
- Small genomes limit the power of the phenotypes they express.
- So a small genome cannot encode a phenotype which contains a high-fidelity replicating mechanism

In essence, the poor copy fidelity of early genotypes could not encode the phenotype sufficiently accurately to preserve any improvements in copy fidelity.

We do not attempt to resolve Eigen's paradox here. Instead, we used the paradox as a challenge for AC design. This allows us to set ACs in a context and discuss their properties relative to this context. We argue for Goldberg's 'piecewise engineering' approach in the first instance [12] and take the view that a 'one size fits all' approach to AC design is not the most efficient way of approaching difficult problems. These problems are characterised by a system (such as chemistry, in the case of Eigen's paradox) that changes how it behaves as it develops through time. Before the resolution of Eigen's paradox, replicators were constrained in their size and therefore in their functionality; once the paradox has been resolved, this ceiling is lifted which allows for further evolution and adaptation, eventually leading to the central dogma of biology that we recognise today.

ACs can be used to produce Artificial Life (ALife) systems in which evolutionary features (such as reproduction or mutation) are not explicitly defined *a priori*. Instead, they are emergent properties of the system and as such are implicitly embedded:— they can be changed by the ALife system, rather than having to be pre-specified by a designer.

We investigate this by considering three different ACs which can represent the chemistry that existed before, after and during Eigen's paradox (figure 1). These chemistries come from recent work by the authors, developing ACs for three challenges: the origin of life [10]; the evolution of evolvability (*meta*-evolution) [21]; and as the basis for a self-maintaining genetic algorithm [16]. Note that the emphasis in these works is placed heavily on replication processes and do not consider the role of the container in the context of resolving Eigen's paradox. None of our chemistries currently model a cell membrane within the chemistry itself (but our chemistries do occupy a set volume and thus at least have the abstract concept of a container) although the emergence of membranes is linked to the emergence of replicators in models of the early earth. After describing these three chemistries, we discuss the properties they possess, how these relate to properties considered interesting by other authors [24] and how they relate to Eigen's paradox.

Finding a single chemistry to span these phases is much harder than finding different chemistries modelling each situation appropriately. The goal of our work in these three areas is to derive a new set of desired properties, to aid us in designing a series of ACs that together form an innovative artificial evolutionary platform. We are interested in finding which properties of ACs contribute to evolution and evolvability in general. Focusing on Eigen's paradox as an example of evolvability is a way in which we can tease out these properties.

The Context of Eigen's Paradox

A time-line of the beginnings of evolution on the early earth is shown in figure 1. This period is interesting to ALife researchers because it resolved Eigen's paradox [22], a key problem in evolution. The period begins with the 'late heavy bombardment' of the earth by debris from space as the solar system formed — only after this was the planet thought to be stable enough for life to prosper. Then come the well-known phases in the development of life on this planet, from the pre-biotic chemical 'soup' to the emergence of the central dogma of biology. The graphic in the middle of figure 1 illustrates the inheritance of genetic strategies over this period. Essentially, many different evolutionary strategies are prevalent, until the central dogma sweeps the planet as shown by the shaded region at the bottom of the graphic. Eigen's paradox is resolved before the emergence of replicator molecules that precede the central dogma of biology. The three chemistries forming the basis of the current contribution are shown to the right of the graphic in figure 1. These are described below.

From the perspective of the central dogma, Eigen's paradox is insoluble. It is not possible to construct a long genotype for an accurate copying phenotype from the basis of a short genotype that encodes an inaccurately-copying phenotype. And yet, the central dogma is common to all known life. Potential resolutions to Eigen's paradox are:

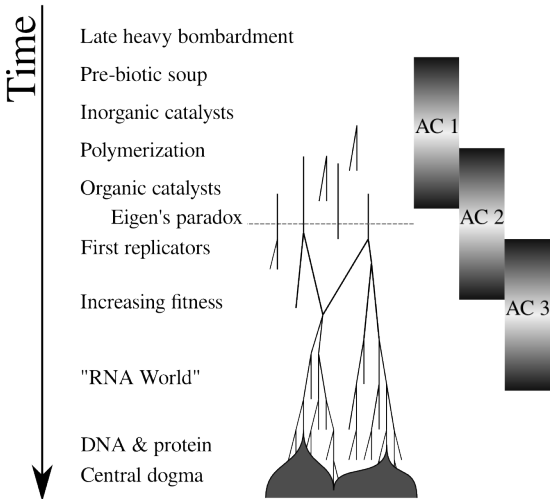

Figure 1: Timeline of the beginning of evolving systems. Events leading to the central dogma of biology are shown on the left. The resolution of Eigen's paradox is required for the emergence of competent replicators. The central graphic shows the myriad different evolutionary processes that are thought to have been prevalent before the central dogma. The three Artificial chemistries are shown on the right of the figure.

1. Stochastic processes throughout the planet over a billion years could ensure that, even though on average a short sequence does not copy well, given enough sequences, some might work well enough for long enough to encode a faithful genotype-copying arrangement.

2. Environment: there may have been local isolated environments where fidelity was higher and denaturation was reduced. If a long & accurate replicator could have arisen there, it could have spread to other locations; e.g. the presence of inorganic compounds such as clay crystals, could have aided replication [2].

3. The assumption that short sequences imply low fidelity is false. It may have been possible to construct some efficient copier from a short genome in some 'lost' chemistry. Alternatively, some collective property of the system does the job of forming an accurate template before the arrival of specialised template-carrying molecules.

Our chemistries explore the third possibility for resolution of the paradox. ACs for ALife could be used to find evolutionary mechanisms simpler than the central dogma of biology — this forms the central design objective of our ACs. It involves seeking simpler molecular machinery than DNA, RNA and protein, which will be easier to simulate computationally. However, by discarding the central dogma of bi-

ology, we have lost the ability to design replicators by looking at biology and attempting to copy what we see because these primitive replicators no longer exist on the Earth. We are faced with the task of designing from scratch an AC that can support recognisable evolution.

The paradox is related to ACs in two ways. Firstly, if we have an AC that cannot resolve this paradox, then the AC has a (small) maximum genome size that it can not overcome. If we want genomes larger than this size, then we must explicitly add in high-fidelity replicators. Secondly, the ACs may foster new theories about how Eigen's paradox can be resolved. We can design ACs to test these new theories.

Implementations

We now present a brief overview of the three chemistries referenced in figure 1. In its most basic form, an AC is defined as[4]:

- A set of molecules (both those present at a point in time and all possible molecules)
- Reactions that describe transformations between sets of molecules
- An algorithm which determines how the reactions are applied to the set of molecules present

A number of different ACs have been developed from this basis, without much consensus on which approach is 'best'. However, there have been a number of different properties and characteristics proposed as interesting features or requirements. ACs have also been applied in various other contexts [23, 20], but the power of ACs is limited if evolutionary processes are not implicit in the representation.

Our approach is to decompose the problem into three phases: emergence of self-replicators (AC1); evolution of evolvability (AC2); stable but primitive evolutionary system (AC3).

AC1: Emergence of Replicators

AC1 is an analogue of the pre-biotic soup in which early replicators emerged. It is designed as an source of open-ended chemical novelty and innovation, in which replicating molecular species may be initially formed. In this phase, replicators do not yet exist and therefore other processes and structures, such as autocatalytic sets [19] and hypercycles [6, 7, 8], are the focus of investigation.

One of the problems investigating the earliest phase of evolution is that there cannot be an assumption of a pre-existing replicating structure — it must be initially formed from other reactions. In order to achieve this, the chemistry must spontaneously generate sufficient novelty in order to describe templates and the molecular machinery to replicate them.

To implement an AC for this phase, we have developed a novel molecular representation classification, which we call

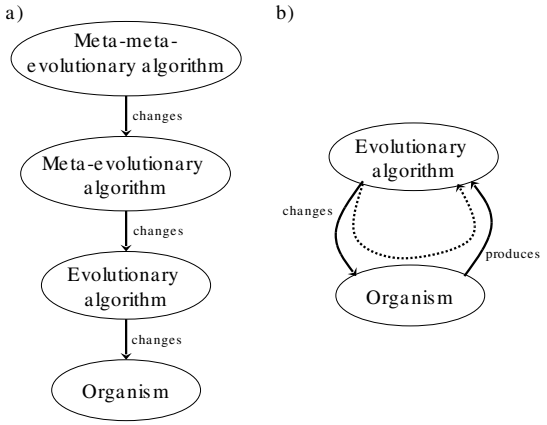

Figure 2: a) Naive meta-evolution suffers from the problem of how many meta-levels to use. b) Having the evolutionary algorithm as an emergent property of the organisms solves this problem. Evolution itself can choose how many levels of evolutionary algorithm to encode within the organism.

"sub-symbolic". Rather than reactants and products of reactions being defined in advance, they are determined by *bonding criteria* applied to *bonding properties* of the molecular species present; the bonding properties are themselves a emergent property of each atoms collection of sub-symbolic components. This means that for any molecule (either created within the system or provided by external input) all of its interactions can be generated dynamically.

Rather than try to specify a single AC that can achieve the emergence we seek, we have designed a framework within which many ACs can exist (RBN-World [10]). To find individual ACs that may achieve the goal of emergent replicators within this design space, we have developed a series of tests for desirable low-level properties. These form a set of 'stepping stones' that lead towards self-replicating systems. [9]

At the end of this phase, we anticipate a collection of molecules that form an autocatalytic set — production of every member of the set is catalysed by at least one member of the set. Taken as a cooperative collective, this forms a proto-organism capable of growth and replication.

AC2: Meta-Evolution

AC2 overlaps with AC1. AC2 is a meta-evolution phase in which speed and fidelity of replications increases as a loosely-replicating proto-entity becomes more capable of maintaining both its own fidelity and the fidelity of a larger reaction network [21]. The proto-entity will gradually evolve robust replication until it is widespread and prevalent.

AC2 implements an analogue of a traditional genetic algorithm (GA) in the same medium as the organisms themselves (figure 2). This requires the organisms and algorithm to be

implemented in a single representation, which a sufficiently rich AC can provide. We have identified the following requirements of an AC for meta-evolution:

- *template molecule(s)* that encode enzymes, including indirectly encoding the reactions that they can perform.
- *translation enzymes* that "read" the template molecule and construct the enzymes that are coded for.
- *replication enzymes* that can copy templates with some stochastic error so that mutations can occur.

We will encode initial examples of all of the above into template molecules within the system. This will allow meta-evolution to happen, because mutations occurring on the template molecule can cause the EA to change.

One part of evolving the EA is evolving the concept of mutation. We enable evolution of mutation because mutations can occur due to inexact copying of the template (mutation-on-copy). The replication enzymes are encoded on the template, and so the process of replication (and thus the process of mutation) can evolve under its own control.

The replication machines in this AC contain complex internal structure, and replication is a multi-step, character-by-character process. To replicate a template molecule, each character is replicated in turn by the following sequence of steps:

1. The next character from the template is read;
2. The replicator makes an internal representation of the next character;
3. Raw materials are picked up from the environment;
4. The raw materials are used to write the next character to the copy;
5. The replicator moves on to the next character on the template and the copy.

Because the copying process involves many steps, there are many ways in which is can go wrong. This means that many different types of mutation are possible, and also many different ways in which the replicator can evolve.

The replicators emerging from AC1 can be seen in AC2 as primitive and unstable with have low fidelity (high mutation). These will undergo metaevolution within AC2 to become the stable replicators of AC 3 exhibiting high fidelity (low mutation).

In relation to Eigen's paradox, this AC has a representation of replicating chemicals that can evolve their own copying fidelity. Therefore changes in the template and/or copy fidelity can be recorded over time and different conditions. This will enable examination of the conditions under which Eigen's paradox is resolvable and if it is inevitable.

AC3: "RNA world"

AC3 represents molecules that can copy with relatively high accuracy, even though there is not necessarily a distinction between template and machine.

AC 3 is called Stringmol . The Stringmol chemistry was developed to emulate molecular systems in such a manner that the binding and reactions between molecules could be varied using evolutionary approaches. In a nutshell, a molecule consists of a *sequence* along with a set of *flags* and *pointers* that allow the sequence to be executed as a program. Further details are available in [16] and [14]

There are two key features of the Stringmol system. The first is the *binding scheme*, which specifies the probability of two molecules joining together and creating a reaction. The second is the *mutation-reaction scheme*, which specifies how reactions occur under an environment of mutation, and determines what the products of the reaction are. Thus we have rules that handle the alignment of two strings of symbols (bound pair of molecules), and interprets the strings as a program and a data repository simultaneously.

Experiments with mutation in the Stringmol system have shown that a wide variety of phenomena can occur with no extenally-applied evolutionary pressure. In particular, we see the spontaneous emergence of autocatalytic sets from a basic replicase system [15].

Properties of Artificial Chemistries

It is useful to consider ACs in the light of the properties of ALife listed in [1]. ACs offer a route to generating "life" from the non-living by: A.2, exploring the transition to life *in silico*; A.3, discovering novel living organisations; A.4, determining how rules and symbols are generated from physical dynamics. Once a 'living' AC is constructed, then investigation can proceed, to: B.6, determine what is inevitable in open-ended evolution; B.7, explore evolutionary transitions (e.g. Eigen's paradox); B.8, provide the base layer of a hierarchical dynamical system; B.10, form the currency of an information processing theory for evolving systems. These ALife properties drive the properties of the underlying chemistry. One classification of desirable properties of an AC by Suzuki et al was published in [24] and is reproduced for convenience in table 1 alongside our summarised interpretations. We divide those ten properties into three groups: *molecule & reaction* properties, *membrane* properties and *mutation* properties.

New properties

Each of the three authors of this paper has independently developed ACs analogous to different stages in early evolution. We use these three 'case study' ACs to think about desirable properties of ACs in general.

In addition to the properties in table 1, there are some further properties we perceive to be desirable in an AC:

No.	Property	Interpretation	
1.	The symbols or symbol ingredients be conserved (or quasi-conserved) in each elementary reaction, at least with the aid of a higher-level manager.	Conservation of Mass	Molecules & reactions
2.	An unlimited amount of information be coded in a symbol or a sequence of symbols.	Molecules composed of atoms & bonds	
3.	Particular symbols that specify and activate reactions be present.	Catalysis	
4.	The translation relation from genotypes to phenotypes be specified as a phenotypic function.	Phenotypic gene expression	
5.	The information space be able to be partitioned by semi-permeable membranes, creating cellular compartments in the space.	Cells	Membrane
6.	The number of symbols in a cell can be freely changed by symbol transportation, or at least can be changed by a modification in the breeding operation.	Variable cell volume / concentration	
7.	Cellular compartments mingle with each other by some random process.	Cell movement	
8.	In-cell or between-cell signals be transmitted in the manner of symbol transportation.	Diffusion through membranes	
10.	Symbols be selectively transferred to specific target positions by particular activator symbols (strongly selective), or at least selectively transferred by symbol interaction rules (weakly selective).	Membrane pores & pumps	
9.	There be a possibility of symbols being changed or rearranged by some random process.	Spontaneous Mutation	Mutation

Table 1: The list of desirable AC properties from [24]. On the left is the original description, on the right is our summarised interpretation. NB: we classify property 10 as a membrane property along with 5-8 rather than a genome property with 9.

11. Novelty & innovation This is a property desired in evolutionary systems, and AC design should reflect this. If a new molecule is introduced to the chemistry, it should be able to interact with the other molecules present without requiring the AC to be changed. Furthermore, the AC should be able to generate novel molecules itself to allow innovative genetic architectures to emerge. This is related to Suzuki's properties #2: *Atoms and bonds* and #3: *Catalysis*, but rather than defining the function of molecules *a priori*, the possibility of novelty should be a general property of the molecular design. It is clear that ACs require this property in order to resolve Eigen's paradox, since without novelty there can be no transition between replicating systems. One can detect this property in absolute terms by asking whether it is possible to add a new molecular species to the system. If it is possible, one should then ask how *easy* it is to do so, and how easy it is for the system to *generate* new molecular species.

12. Range of Scales Although we do not think that all evolutionary phases should be supported by a single chemistry, we do think that chemistries should exhibit a wide range of scales — both spatially and temporally. Much of biology relies on reactions that proceed much slower than others, spanning several orders of magnitude in some cases. A large range of sizes of molecules are also present — from small metabolites consisting of a handful to atoms, to huge enzyme complexes with tens of thousands. Without such diversity, an AC would have limited scope for evolutionary exploration and therefore be restricted in terms of its potential behaviours and solutions to encountered problems.

A large range of spatio-temporal scales would also allow for smoother evolutionary slope climbing by gradual improvements once a solution has been found, for example with a faster rate or greater stability. Scale need not be measured in terms of size alone. Multi-scale representations are useful, because they offer a route to increase the efficiency of the system.

13. Dynamic environment History is littered with cases where an environmental change triggered an evolutionary breakthrough (punctuated equilibria [13]). There is also evidence that variation maintained by different environments can provide useful raw material for evolution, such as around deep-sea geothermal vents [11]. These dynamic environments can occur on many different scales; real-world biology varies from day/night cycles, to changing seasons and ice ages on a temporal scale and varies from micro-

environments between soil particles, through regional variations to continents (which themselves change over geological timescales). In order to utilise some of these dynamics, an AC should have parameters that can be varied (over time, space or both) to created different environments – analogous to temperature, pressure, pH, or other similar characteristics.

Dynamic environments allows a system to fully explore a chemistry, particularly if the rate of mutation varies. If the system can resolve Eigen's paradox locally within one environment, it can improve there and then spread to other environments — even if it could not evolve in those other environments directly.

14. Redundancy & degeneracy Successful evolutionary systems often contain *neutral mutation*. In an AC, this can be characterised by redundancy — multiple molecules that participate in equivalent reactions. However, neutral mutation is rarely completely neutral; it may have small side-effects. Degeneracy in an AC captures this by allowing two molecules to be equivalent for some reactions, but not for others.

In relation to Eigen's Paradox, redundancy and/or degeneracy can help by allowing multiple molecules to fulfil the same roles in the system. If one or more of these are lost through mutation, then the others may be able to partially or fully compensate. Techniques for measuring redundancy and degeneracy should be applicable to the AC, and give a feel for the expressive power of the system.

15. Emergent complex properties The reactions a molecular species participates in should be based on its structure, with similar molecules participating in similar reactions. However, there should be variation in this mapping such that while similar molecules in general have similar interactions, some similar molecules have very different interactions. This will allow an evolutionary landscape where gradual change generally occurs, yet there are some large changes in some regions. Combined with appropriate evolutionary pressures, this will lead to an efficient evolutionary engine.

16. Unified molecular representation There should be no 'special privileges' for template molecules — the property of holding genetic instructions should be an emergent property of the AC. This does not mean they have to be constructed from the same materials as other aspects of the chemistry, only that they should obey the same constraints and rules. In addition, if explicit membranes are used, they should also be represented without 'special privileges'.

The advantage of a unified molecular representation is that any part of the system can potentially interact with by any other part. This allows wider-ranging evolutionary changes and potentially highly innovative solutions to meta-evolutionary problems. It also means that the 'best' implementation of template molecules (or membranes) does not need to be hard-wired into the system beforehand — the system can be bootstrapped with an implementation that works and go on to optimise this itself.

17. Stochasticity Deterministic interactions between agents are a potential barrier to novel behaviour, and stochasticity can help smooth evolutionary changes by sampling the search space of possible alternatives. This leads to more efficient evolution when there are a large number of possible improvements.

18. Emergent mutation rates The replication mechanisms should enable the rate of error-on-copy to be modified. This allows the evolution of evolvability. A system that can reduce its own mutation rate in this manner can resolve Eigen's paradox by allowing larger templates to mutate less and so be more stable. But since the mechanism of genotype-encoding is changeable, the rate at which error accumulates cannot be set as an individual system-level parameter. Rather, the manifestation of error emerges from the reaction mechanism of the AC.

Mapping properties to three chemistries

Our three chemistries conform to the new properties listed in the previous section, thought no one chemistry contains all of them, but do not conform to some of the properties listed [24]. Below we show where our chemistries fit into Suzuki's and our own framework and the implications of those design decisions.

AC 1: Emergence of Replicators This AC analogue has a number of key properties within it. AC 1 implements #1: *conservation of mass* and #2: *atoms and bonds* of Suzuki's properties. Properties #3: *catalysis* and #4: *phenotypic gene expression* are deliberately not implemented in advance but are sought as emergent properties of the system. Our new property #11: *novelty & innovation* is the most important for this problem as we rely on novelty in order for replicators to emerge. Property #16: *unified molecular representation* is also key as we do not define what molecules fulfil which functions of the evolution of the system. #15: *Emergent complex properties* is another property that this systems is designed to exhibit, and is fundamental for the problem we are attempting to address.

Some properties we deliberately do not attempt to include in this AC. #18: *Emergent mutation* and Suzuki's #9: *spontaneous mutation* are not applicable to this phase, as there is not an explicit genome to be mutated; mutation-on-copy may appear as an emergent phenomenon however.

AC 2: Meta-Evolution The purpose of this AC is to investigate a rich mutation scheme, in particular #18: *emergent mutation* This is done by an enzyme-driven copying process with both #14: *redundancy and degeneracy* and #17: *stochastic* properties. This AC will display the emergent

complex property of meta-evolution when the copying machine is both encoded on the template being copied (which requires a #16: *unified molecular representation*) and situated in a #13: *dynamic environment* to provide a changing evolutionary pressure.

Relating this to Suzuki's properties, exploring #9: *spontaneous mutation* is also part of the purpose of this chemistry. In order for there to be a template to copy, this chemistry must satisfy #2: *atoms and bonds*. To be able to encode enzymes, we must satisfy #3: *catalysis*. The translation machines described above satisfy #4: *phenotypic expression* as they are both encoded and represented within the chemistry. To make evolution happen, this chemistry will enforce #1: *conservation of mass* through the atomic structures, which imposes additional restrictions upon the potential evolutionary solutions. As with AC 1 above, this chemistry is not especially concerned with membranes, and so properties #6, #7, #8 and #10 are not applicable to this chemistry. However, property #5: *containers* is satisfied in that membranes are implemented as simple containers, but their only function is to keep enzymes close to the templates they are acting on. There is no direct cell-cell interaction.

AC 3: "RNA world" Relating this AC to the the molecular and mutation-reaction properties described in table 1 [24] indicates that #1: *conservation of mass*, #2: *atoms and bonds*, #3: *catalysis*, and #4: *phenotypic gene expression* are all applicable to Stringmol . The mutation-reaction framework is more complicated however. In Stringmol mutation only occurs as new molecules are constructed, not spontaneously as specified by Suzuki et al. Mutations occur during the selective copy of symbols during a reaction of a particular type. This mimics biology more closely and can potentially be built into the AC to implement the meta-evolution described in AC 2.

Although this deviates from Suzuki *et al.*'s specification, mutation still occurs and it's rate can be controlled in a similar manner to the 'spontaneous' mutation in described (a 'cosmic ray rate'). Stringmol system allows reliable replication to be specified, but has a set mutation rate that allows adaptation to occur. These are the conditions in an 'RNA-world' which the Stringmol system was designed to emulate, and which has the capability to produce innovative responses.

Turning to the remainder of our new properties, #14: *Redundancy & degeneracy* are properties of this system, as well as #17: *stochasticity* due to the variable binding affinities. There is also the possibility for #11: *novelty & innovation* in terms of novel sequences with novel behaviours. Interestingly, the baseline mutation scheme allows a richer suite of macro-mutations to arise, with dramatic changes in the inter-molecular dynamics of the replication process. Stringmol therefore possesses our new property #18: *Emergent mutation rates*.

Conclusion

AC designs have to trade off between being rich enough to exhibit interesting behaviours and being simple enough to be computationally tractable. To address this, we develop abstractions with two goals: 1, to make the rich behaviour computationally tractable, and 2, to discover which properties underlie the richness. When using ACs to address evolutionary problems, the goals become further complicated. For example, in real chemistry the problems and solutions regarding survival of the organism have changed over time — the first forms of life were very different to modern populations of multi-cellular organisms. We use Eigen's paradox as an example of applying ACs to a evolutionary problem. We are not aiming to provide a resolution of Eigen's paradox: we provide a way of thinking about problems in which the properties and behaviours of the chemistry change over time (before, during and after the paradox).

In this work we have not looked at properties involving membranes and other spatial characteristics (#5: *cells with membranes*, #6: *variable cell volume / concentration*, #7: *cell movement*, #8: *diffusion through membranes*, and #10: *membrane pores & pumps* from Suzuki *et al.*). This is because these properties are predominantly under the control of the 'kinetics' used for any particular implementation of an AC. In our experiences, the kinetics component of the model can often be interchanged between different ACs depending on the features under investigation and available computational resources. For example, previous work on membranes in an AC [17, 18, 3], whilst clearly demonstrating interesting behaviours, poses computational challenges when used for investigations of evolution and novelty.

By considering specific ACs for three phases of evolution in the context of Eigen's paradox, we have concentrated on the properties needed for each phase. In all of these ACs, sub-symbolic atomic representations are useful because they preclude the need to create a set of reaction rules whenever a novel molecular species is produced, and so provide an appropriate platform for evolution to discover and preserve novel solutions which confer some benefit on the system. Effectively, using the sub-symbolic representation provides many properties for 'free'; #1: *conservation of mass*, #2: *atoms and bonds* and #3: *catalysis* from Suzuki's properties as well as #11: *novelty & innovation* and #16: *unified molecular representation* from our additional properties.

We have presented eight new properties in addition to the ten given in [24]. We have used Eigen's paradox as a context to map these properties onto our ACs to demonstrate how they can be used in the design and evaluation process. The resulting set of principles can be used for the design of a more generally applicable set of ACs.

Acknowledgements

Simon Hickinbotham and Adam Nellis are funded by the Plazzmid project, EPSRC grant EP/F031033/1. Adam

Faulconbridge is funded by BBSRC.

References

[1] Bedau, M. A., McCaskill, J. S., Packard, N. H., Rasmussen, S., Adami, C., Green, D. G., Ikegami, T., Kaneko, K., and Ray, T. S. (2000). Open problems in artificial life. *Artificial Life*, 6(4):363–376.

[2] Cairns-Smith, A. (1982). *Genetic takeover and the mineral origins of life*. Cambridge University Press.

[3] di Fenizio, P. S., Dittrich, P., and Banzhaf, W. (2001). Spontaneous formation of proto-cells in an universal artificial chemistry on a planar graph. *Lecture notes in computer science*, pages 206–215.

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

[5] Eigen, M. (1971). Self-organization of matter and the evolution of biological macromolecules. *Naturwissenschaften*, 58(10):465–523.

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

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

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

[9] Faulconbridge, A., Stepney, S., Miller, J. F., and Caves, L. (2010). RBN-world: The hunt for a rich achem. In *Proc. ALife XII*. MIT Press.

[10] Faulconbridge, A., Stepney, S., Miller, J. F., and Caves, L. (in press). RBN-world: A sub-symbolic artificial chemistry. In *ECAL 2009 (LNCS 5777)*. Springer.

[11] Gibson, R. N., Atkinson, R. J. A., and Gordon, J. D. M. (2001). *Oceanography and Marine Biology: An Annual Review*. CRC Press.

[12] Goldberg, D. E. (2002). *The Design of Innovation: Lessons from and for Competent Genetic Algorithms*. Kluwer Academic Publishers, Norwell, MA, USA.

[13] Gould, S. J. and Eldredge, N. (1977). Punctuated equilibria: the tempo and mode of evolution reconsidered. *Paleobiology*, 3(2):115–151.

[14] Hickinbotham, S., Clark, E., Nellis, A., Pay, M., Stepney, S., Clarke, T., and Young, P. (2009). An abstract metabolism of molecular microprograms. Technical Report (in preparation), University of York.

[15] Hickinbotham, S., Clark, E., Stepney, S., Clarke, T., Nellis, A., Pay, M., and Young, P. (2010). Diversity from a monoculture: effects of mutation-on-copy in a string-based artificial chemistry. In *Proc. ALife XII*. MIT Press.

[16] Hickinbotham, S., Clark, E., Stepney, S., Clarke, T., Nellis, A., Pay, M., and Young, P. (in press). Molecular microprograms. In *ECAL 2009 (LNCS 5777)*. Springer.

[17] Hutton, T. (2004). Making membranes in artificial chemistries. In *Workshop and Tutorial Proceedings, Artificial Life IX*, pages 71–76. Citeseer.

[18] Ikegami, T. (2009). Shapes and Self-Movement in Protocell Systems. *Artificial Life*, 15(1):59–70.

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

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

[21] Nellis, A. and Stepney, S. (2010). Meta Evolution (in preparation).

[22] Smith, J. M. and Szathmary, E. (1997). *The Major Transitions in Evolution*. New York: Oxford University Press.

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

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

[25] Vetsigian, K., Woese, C., and Goldenfeld, N. (2006). Collective evolution and the genetic code. *Proceedings of the National Academy Science*, 103:10696–10701.