

Tractable models of self-sustaining autocatalytic networks

Mike Steel^{*,†} and Wim Hordijk^{*,‡}

[†]*Biomathematics Research Centre, University of Canterbury, Christchurch, New Zealand*

[‡]*Institute for Advanced Study, University of Amsterdam, The Netherlands*

E-mail: mike.steel@canterbury.ac.nz; wim@WorldWideWanderings.net

Abstract

Self-sustaining autocatalytic networks play a central role in living systems, from metabolism at the origin of life, simple RNA networks, and the modern cell, to ecology and cognition. A collectively autocatalytic network that can be sustained from an ambient food set is also referred to more formally as a ‘Reflexively Autocatalytic F-generated’ (RAF) set. In this paper, we first investigate a simplified setting for studying RAFs, which are nevertheless relevant to real biochemistry and allows for a more exact mathematical analysis based on graph-theoretic concepts. This, in turn, allows for the development of efficient (polynomial-time) algorithms for questions that are computationally NP-hard in the general RAF setting. We then show how this simplified setting for RAF systems leads naturally to a more general notion of RAFs that are ‘generative’ (they can be built up from simpler RAFs) and for which efficient algorithms carry over to this more general setting. Finally, we show how classical RAF theory can be extended to deal with ensembles of catalysts as well as the assignment of rates to reactions according to which catalysts (or combinations of catalysts) are available.

Introduction

A central property of the chemistry of living systems is that they combine two basic features: (i) the ability to survive on an ambient food source, and (ii) each biochemical reaction in the system requires only reactants and a catalyst that are provided by other reactions in the system (or are present in the food set). The notion of a self-sustaining ‘collectively autocatalytic set’ tries to capture these basic features formally, and was pioneered by Stuart Kauffman,^{1,2} who investigated a simple binary polymer model to address questions that relate to the origin of life. The notion of a collectively autocatalytic set was subsequently formalised more precisely as ‘Reflexively Auto-catalytic and F-generated’ (RAF) sets (defined shortly) and explored by others.³ RAFs are related to other notions such as Rosen’s (M;R) systems,⁴ and ‘organisations’ in Chemical Organisation Theory.⁵ The application of RAFs has expanded beyond toy polymer models to analyse both real living systems (e.g. the metabolic network of *Escherichia coli*⁶) and simple autocatalytic RNA systems that have recently been generated in laboratory studies.⁷

The generality of RAF theory also means that a ‘reaction’ need not refer specifically to a chemical reaction, but to any process in which ‘items’ are combined and transformed into new ‘items’, and where similar ‘items’ that are not used up in the process facilitate (or ‘catalyse’) the process. This has led to application of RAF theory to processes beyond biochemistry, including biodiversity,⁸ cognitive psychology,⁹ and (more speculatively) economics.³

In this paper, we show how RAF theory can be developed further to:

- provide an exact and tractable characterisation of RAFs and subRAFs when reactants involve just food molecules;
- extend this last concept to general catalytic reaction networks by defining a new type of RAF (‘generative’) which couples realism with tractability; and
- include reaction rates into RAF theory and show that an optimal RAF can be calculated in polynomial time.

We begin with some definitions.

Catalytic reaction systems and RAFs

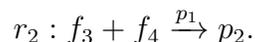
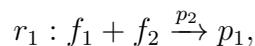
A *catalytic reaction system* (CRS) consists of a set X of ‘molecule types’, a set \mathcal{R} of ‘reactions’, an assignment C describing which molecule types catalyse which reactions, and a subset F of X consisting of a ‘food set’ of basic building block molecule types freely available from the environment. Here, a ‘reaction’ refers to a process that takes one or more molecule types (the ‘reactants’) as input and produces one or more molecule types as output (‘products’). C can be viewed as a subset of $X \times \mathcal{R}$.

Definitions (RAF, maxRAF).

Given a CRS $\mathcal{Q} = (X, \mathcal{R}, C, F)$, a subset \mathcal{R}' of \mathcal{R} is said to be a *RAF* for \mathcal{Q} if \mathcal{R}' is nonempty and satisfies the following two conditions.

- *Reflexively autocatalytic* (RA): each reaction $r \in \mathcal{R}'$ is catalysed by at least one molecule type that is either present in the food set or is generated by another reaction in \mathcal{R}' .
- *Food-generated* (F): for each reaction $r \in \mathcal{R}'$, each reactant of r is either present in the food set F or can be generated by a sequence of reactions from \mathcal{R}' , each of which has each of its reactants present either in the food set or as a product of an earlier reaction in the sequence.

In other words, a RAF is a subset of reactions that is both self-sustaining (from the food set) and collectively autocatalytic. In forming a RAF from the food set, some reactions may initially need to proceed uncatalysed (and thereby at a lower rate) but once formed every reaction in the RAF will be catalysed. A simple example of a RAF is the pair of reactions



with food set $\{f_1, f_2, f_3, f_4\}$, as shown in Fig. 1. Note that either r_1 or r_2 must first proceed uncatalysed, but once one reaction has occurred, the system continues with both reactions catalysed.

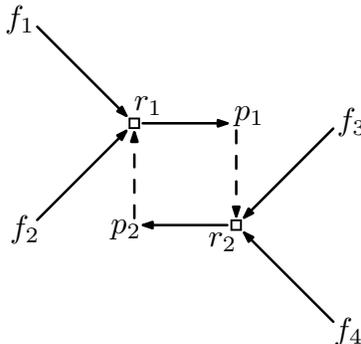


Figure 1: A simple RAF involving two reactions, r_1 and r_2 and a food set $F = \{f_1, f_2, f_3, f_4\}$, with catalysation arcs shown as dashed arrows.

The food-generated condition (F) can also be formalised as follows: For an arbitrary subset \mathcal{R}' of \mathcal{R} let $\text{cl}_{\mathcal{R}'}(F)$ be the (unique) minimal subset W of X that contains F and has the property that if $r \in \mathcal{R}'$ and all the reactants of r are in W then the product(s) of r are also in W . The (F) condition now becomes the statement that each reactant of each reaction in \mathcal{R}' is present in $\text{cl}_{\mathcal{R}'}(F)$. Note also that, assuming the (F) condition holds, the (RA) condition becomes equivalent to the stronger condition that each reaction $r \in \mathcal{R}'$ is catalysed by at least one molecule type that is present in $\text{cl}_{\mathcal{R}'}(F)$.

Definitions (subRAFs, irrRAFs, closed RAFs, minimal closed RAFs, CAFs)

If a CRS \mathcal{Q} has a RAF, then it has a unique maximal RAF, which contains any other RAFs for \mathcal{Q} . This maximal RAF for \mathcal{Q} is referred to as the *maxRAF* of \mathcal{Q} (denoted $\text{maxRAF}(\mathcal{Q})$) and it can be constructed by an algorithm that runs in polynomial time in the size of \mathcal{Q} .¹⁰ The maxRAF may contain one or more proper subsets of reactions that are themselves RAFs for \mathcal{Q} , in which case we call any such subset a *subRAF* of the maxRAF.

An RAF \mathcal{R}' is an *irreducible* RAF (irrRAF) if it contains no subRAF. In other words, removing any single reaction from an irrRAF \mathcal{R}' gives a set of reactions that does not contain a RAF for \mathcal{Q} .

Given any subset \mathcal{R}' of reactions from \mathcal{R} , the *closure* of \mathcal{R}' , denoted $\overline{\mathcal{R}'}$ is the (unique) minimal subset \mathcal{R}'' of \mathcal{R} that contains \mathcal{R}' and satisfies the property that if a reaction r from \mathcal{R} has each of its reactants and at least one catalyst present in the food set or as a product of a reaction from \mathcal{R}'' then r is in \mathcal{R}'' . We say that a RAF \mathcal{R}' is *closed* if it is equal to its closure (i.e. $\mathcal{R}' = \overline{\mathcal{R}'}$). In particular, the maxRAF is always closed (closed RAFs are the type of RAF that is most closely related to, but still different from, organisations in Chemical Organisation Theory⁵).

A *minimal closed RAF* for a CRS \mathcal{Q} is a closed RAF \mathcal{R}' for \mathcal{Q} that does not contain any other closed RAF for \mathcal{Q} as a strict subset. Any closed irrRAF is a minimal closed RAF but a minimal closed RAF need not be an irrRAF.

Given a CRS $\mathcal{Q} = (X, \mathcal{R}, C, F)$, a stronger notion than a RAF is that of a *constructively autocatalytic F-generated* (CAF) set for \mathcal{Q} . This is a nonempty subset \mathcal{R}' of \mathcal{R} for which the reactions in \mathcal{R}' can be ordered in such a way that for each reaction r in \mathcal{R}' , each reactant and at least one catalyst of r has the property that it is either produced by an earlier reaction from \mathcal{R}' or is present in the food set. In other words, a CAF is like a RAF with the extra requirement that no spontaneous (uncatalysed) reactions are required for its formation (i.e. the catalyst needs to be already present when it is first needed). The RAF in Fig. 1 fails to be a CAF.

The structure of RAFs in elementary catalytic reactions systems

Let CRS $\mathcal{Q} = (X, \mathcal{R}, C, F)$. We say that \mathcal{Q} is *elementary* if it satisfies the following condition:

- Each reaction r in \mathcal{R} has all its reactants in F .

An elementary CRS is a very special type of CRS; however it has arisen both in applications to real experimental chemical systems^{7,11} and in theoretical models.¹² The RAF shown

in Fig. 1 is a simple example of an elementary CRS. It is possible to extend the definition of elementary CRS to also allow for reversible reactions, by requiring only one side of the reaction to contain molecule types that are exclusively from F .

In this section, we show that elementary RAFs have sufficient structure to allow a very concise classification of their RAFs, closed subRAFs, irrRAFs, and ‘uninhibited’ closed RAFs (a notion described below), something which is problematic in general. We then extend this analysis to more complicated types of RAFs in the next section.

Our analysis in this section relies heavily on some key notions from graph theory, so we begin by recalling some concepts from that area.

Definitions

In this paper, all graphs will be finite. Given a directed graph $\mathcal{D} = (V, A)$, recall that a *strongly connected component* of \mathcal{D} is a maximal subset W of V with the property that for any vertices u, v in W , there is a path from u to v and a path from v to u .

It is a classical result that for any directed graph $\mathcal{D} = (V, A)$, the vertex set V can be partitioned into strongly connected components. This, in turn, induces a directed graph structure, called the *condensation (digraph) of \mathcal{D}* , which we will denote by \mathcal{D}^* . In this directed graph, the vertex set is the collection of strongly connected components of \mathcal{D} and there is an arc (U, V) in \mathcal{D}^* if there is an arc (u, v) in \mathcal{D} with $u \in U$ and $v \in V$. By definition, \mathcal{D}^* is an acyclic directed graph. Moreover, the task of partitioning V into strongly connected components and constructing the graph \mathcal{D}^* can both be carried out in polynomial time.¹³ Note that the strongly connected component containing v will consist just of v if v is not part of a cycle involving another vertex.

We now introduce some further definitions. Given a directed graph $\mathcal{D} = (V, A)$:

- We say that a strongly connected component S of \mathcal{D} is a *core* if either $|S| > 1$ or $|S| = 1$ (say $S = \{r\}$) and there is an arc from r to itself. Note that \mathcal{D} has a core if and only if \mathcal{D} has a directed cycle.

- A *chordless cycle* in a directed graph $\mathcal{D} = (V, A)$ is a subset U of vertices of \mathcal{D} for which the induced graph $\mathcal{D}|U$ is a directed cycle (here $\mathcal{D}|U = (U, A')$ where the arc set A' for $\mathcal{D}|U$ is given by $A' = \{(u, v) \in A : u, v \in U\}$). Note that if $|U| = 1$, this means that there is an arc from the vertex in U to itself.
- A vertex v in V is *reachable* from some subset S of V if there is a directed path from some vertex in S to v . More generally, a subset U of V is reachable from S if there is some vertex $v \in U$ that is reachable from S .

First main result

The following theorem provides graph-theoretic characterisations of RAFs, irrRAFs, closed RAFs, and minimal closed RAFs within any elementary CRS.

Given any CRS, \mathcal{Q} , consider the directed graph $\mathcal{D}_{\mathcal{Q}}$ with vertex set \mathcal{R} and with an arc (r, r') if a product of reaction r is a catalyst of reaction r' . In addition, for any reaction r that has a catalyst in F , we add the arc (r, r) (i.e. a loop) into $\mathcal{D}_{\mathcal{Q}}$ if this arc is not already present; this step is just a formal strategy to allow the results to be stated more succinctly, and does not necessarily mean that a product of r is an actual catalyst of r .

We can now state our first main result, the proof of which is provided in the Appendix.

Theorem 1. *Let \mathcal{Q} be an elementary CRS. Then:*

- (i) *\mathcal{Q} has a RAF if and only if $\mathcal{D}_{\mathcal{Q}}$ has a directed cycle, and this holds if and only if $\mathcal{D}_{\mathcal{Q}}$ contains a chordless directed cycle. The RAFs of \mathcal{Q} correspond to the subsets \mathcal{R}' of \mathcal{R} for which the induced directed graph $\mathcal{D}_{\mathcal{Q}}|\mathcal{R}'$ has the property that each vertex has in-degree at least 1.*
- (ii) *The irrRAFs of \mathcal{Q} are the chordless cycles in $\mathcal{D}_{\mathcal{Q}}$. The closed irrRAFs of \mathcal{Q} are chordless cycles from which no other vertex of $\mathcal{D}_{\mathcal{Q}}$ is reachable. The smallest RAFs of \mathcal{Q} are the shortest directed cycles in $\mathcal{D}_{\mathcal{Q}}$.*

- (iii) *The closed RAFs of \mathcal{Q} are the subsets of \mathcal{R} obtained by taking the union of any one or more cores of $\mathcal{D}_{\mathcal{Q}}$ and adding in all the reactions in \mathcal{R} that are reachable from this union.*
- (iv) *Each minimal closed RAF of \mathcal{Q} is obtained by taking any core C of $\mathcal{D}_{\mathcal{Q}}$ for which no other core of $\mathcal{D}_{\mathcal{Q}}$ is reachable from C , and adding in all reactions in \mathcal{R} that are reachable from C .*
- (v) *The number of minimal closed RAFs of \mathcal{Q} is at most the number of cores in $\mathcal{D}_{\mathcal{Q}}$, and thus it is bounded above by $|\text{maxRAF}(\mathcal{Q})|$. These can all be found and listed in polynomial time in $|\mathcal{Q}|$.*
- (vi) *The question of whether or not a given RAF for \mathcal{Q} (e.g. the maxRAF) contains a closed RAF as a strict subset can be solved in polynomial-time.*

Figs. 2, 3 and 4 illustrate Parts (i)–(iv) of Theorem 1. Some of these examples are based on reaction networks that come from actual experimental RAF sets.

Remarks:

- Parts (ii)–(vi) of Theorem 1 hold even when \mathcal{Q} is not elementary, provided that $\mathcal{Q}' = (X, \mathcal{R}', C, F)$ is elementary where \mathcal{R}' is the maxRAF of \mathcal{Q} .
- Although cores share no reactions in common, it is quite possible for minimal closed RAFs to share reactions in common.
- The last sentence of Part (ii) implies that the size of the smallest RAF is equal to the length of the shortest directed cycle in $\mathcal{D}_{\mathcal{Q}}$ and this can be found in polynomial time in $|\mathcal{Q}|$ (by a depth-first-search or network flow techniques). This is in contrast to the problem of finding the size of a smallest RAF in a general CRS, which has been shown to be NP-hard.¹⁶
- An important extension of the RAF concept allows for molecules to inhibit reactions as well. In this case, Theorem 1(v) provides the following result.

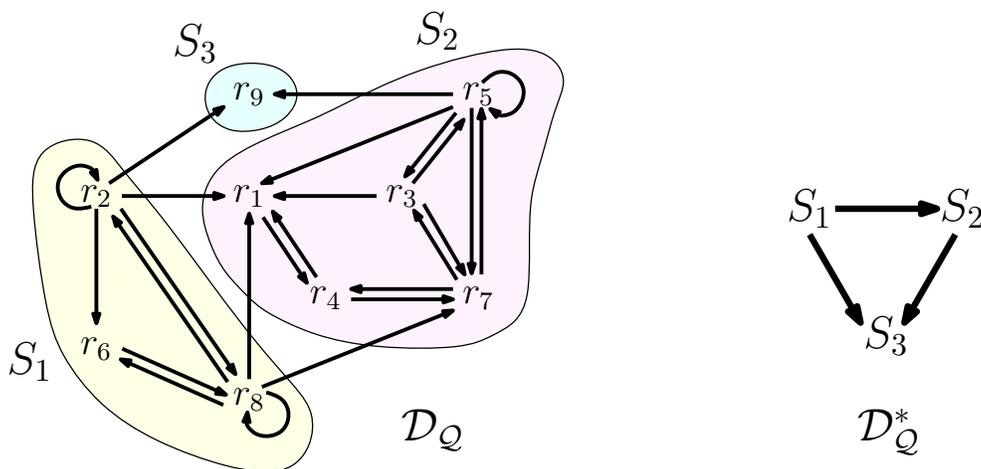


Figure 2: The directed graph \mathcal{D}_Q for an elementary CRS \mathcal{Q} (adapted from an experimental system¹¹) that has three strongly connected components (S_1, S_2, S_3), of which S_1 and S_2 are cores. The associated (acyclic) condensation digraph \mathcal{D}_Q^* is shown on the right. The unique minimal closed RAF is $S_2 \cup S_3$; the other closed RAF is the full set itself, namely $S_1 \cup S_2 \cup S_3$. The reactions subsets $S_1, S_1 \cup S_3, S_1 \cup S_2, S_1 \cup S_3$, and S_2 are all RAFs but not closed RAFs. A computer-based search finds 305 RAFs altogether. There are six chordless cycles in this CRS, which correspond to the six irrRAF: $\{r_2\}, \{r_5\}, \{r_8\}, \{r_1, r_4\}, \{r_4, r_7\}$ and $\{r_3, r_7\}$. Note that this representation of the CRS is in terms of the molecules produced by reactions that have reactants in the food set. However, each reaction produces a single (and unique) product so we can identify the product with the reaction in this example.

Corollary 1. *When inhibition is also allowed in an elementary CRS \mathcal{Q} , it is possible to determine in polynomial time whether \mathcal{Q} contains a closed RAF \mathcal{R}' for which no reaction is inhibited by any molecule type produced by \mathcal{R}' .*

Proof. There is a closed RAF for \mathcal{Q} that has no inhibition if and only if there is a minimal closed RAF for \mathcal{Q} that has no inhibition. By Part (v) of Theorem 1, there are at most $|\max\text{RAF}(\mathcal{Q})|$ minimal closed RAFs for an elementary CRS \mathcal{Q} , and these can all be checked in polynomial time to determine if any of them have the property that no reaction is inhibited by any molecule type produced by the reactions in the set. \square

- Part (v) of Theorem 1 raises the question of whether this result might apply with the restriction that \mathcal{Q} is elementary. In other words, is the number of minimal closed RAFs in a (general, nonelementary) CRS bounded polynomially in the size of \mathcal{Q} ? The answer turns out to be ‘no’ as the following example shows.

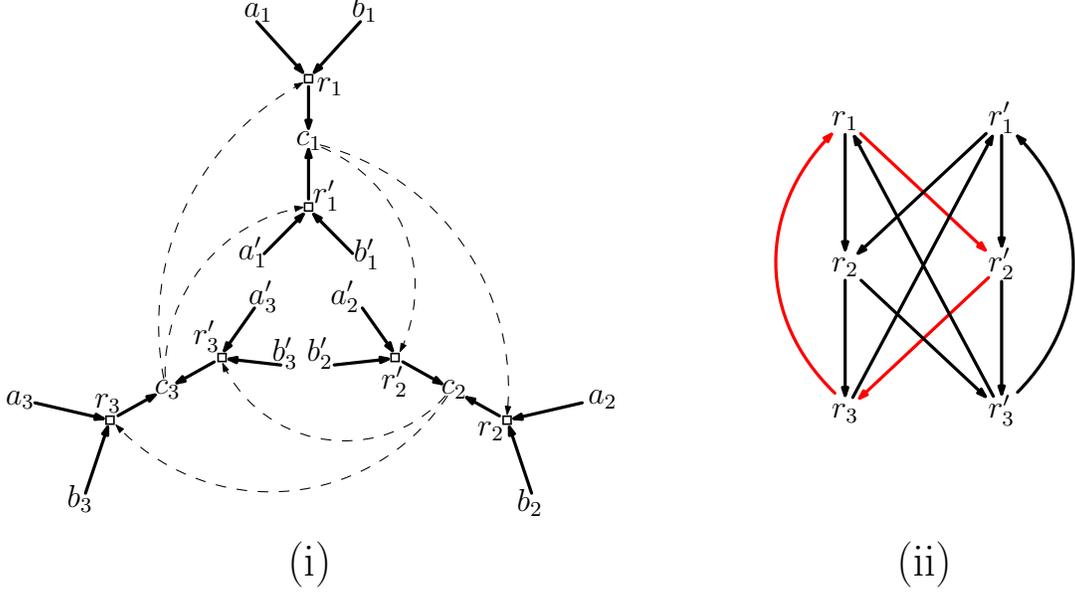


Figure 3: (i) An elementary CRS (with food set F equal to the 12 elements labelled a_i, a'_i, b_i, b'_i for $i = 1, 2, 3$) that has eight irrRAFs, each of which has size 3 (this example can be extended to produce an elementary CRS with $2n$ reactions and 2^n irrRAFs¹⁴). These irrRAFs correspond to the eight chordless cycles in the graph \mathcal{D}_Q shown in (ii), with one of these chordless cycles indicated by the red arcs. None of these irrRAFs is closed. There are 27 RAFs for Q in total.

Consider the CRS $\mathcal{Q}_k := (X, \mathcal{R}, C, F)$ where

$$X = \{f, x, \theta\} \cup \{x_1, x'_1, \dots, x_k, x'_k\} \cup \{\theta_1, \dots, \theta_k\}, F = \{f\},$$

and for $[k] = \{1, 2, \dots, k\}$, the reaction set is:

$$\mathcal{R} = \{r_x, r_\theta\} \cup \{r_i : i \in [k]\} \cup \{r'_i, i \in [k]\} \cup \{\bar{r}_i : i \in [k]\} \cup \{\bar{r}'_i : i \in [k]\},$$

where these reactions are described as follows (with catalysts indicated above the arrows):

$$r_x : f \xrightarrow{\theta} x,$$

$$r_\theta : \theta_1 + \theta_2 + \dots + \theta_k \xrightarrow{\theta} \theta,$$

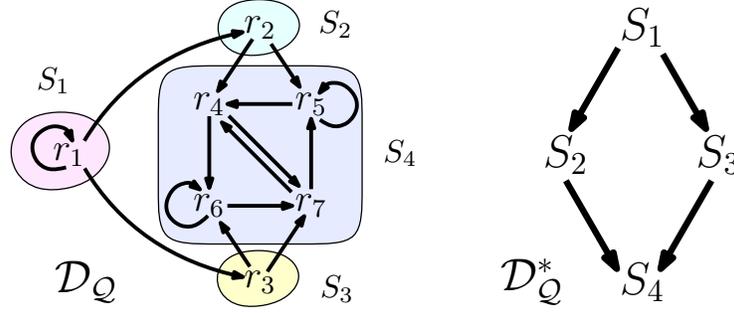


Figure 4: The directed graph $\mathcal{D}_{\mathcal{Q}}$ for an elementary CRS (from an experimental system¹⁵), shown on the left, has 67 RAFs and two closed RAFs (the whole set and $\{r_4, r_5, r_6, r_7\}$). The strongly connected components of $\mathcal{D}_{\mathcal{Q}}$ are $\{r_1\}$, $\{r_2\}$, $\{r_3\}$ and $\{r_4, r_5, r_6, r_7\}$, two of which are cores (namely, $\{r_1\}$ and $\{r_4, r_5, r_6, r_7\}$). The associated condensation digraph $\mathcal{D}_{\mathcal{Q}}^*$ is shown on the right. For this RAF, there are four irrRAF, namely $\{r_1\}$, $\{r_5\}$, $\{r_6\}$, and $\{r_4, r_7\}$.

and for all $i \in [k]$:

$$r_i : x \xrightarrow{x_i} x_i, \quad r'_i : x \xrightarrow{x'_i} x'_i,$$

$$\bar{r}_i : x_i \xrightarrow{\theta_i} \theta_i, \quad \bar{r}'_i : x'_i \xrightarrow{\theta_i} \theta_i.$$

Thus, \mathcal{Q}_k has a food set of size 1, a reaction set of size $4k + 2$, and $3k + 3$ molecule types. Fig. 5 provides a graphical representation of \mathcal{Q}_3 .

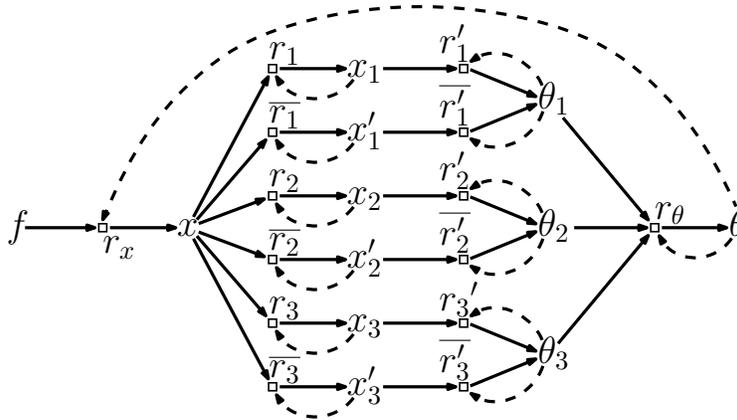


Figure 5: The CRS \mathcal{Q}_3 .

Proposition 1. *The minimal closed RAFs of \mathcal{Q}_k coincide with the irrRAF for \mathcal{Q}_k , and there are 2^k of them. More precisely, \mathcal{R}' is a minimal closed RAF of \mathcal{Q}_k if and only if \mathcal{R}' contains r_x and r_θ and for each $i \in [k]$, \mathcal{R}' contains either (i) r_i and \bar{r}_i but*

neither r'_i nor \bar{r}'_i , or (ii) r'_i and \bar{r}'_i but neither r_i nor \bar{r}_i .

Proof. The ‘if’ direction in the second sentence is clear, since any such \mathcal{R}' is easily seen to be a closed subRAF, as well as being an irrRAF, and thus is a minimal closed RAF. For the ‘only if’ direction, a subset \mathcal{R}' of \mathcal{R} is a RAF of \mathcal{Q}_k precisely if the following two properties hold: (a) \mathcal{R}' contains r_x and r_θ , and (b) for each i , \mathcal{R}' contains either r_i and \bar{r}_i or r'_i and \bar{r}'_i (in order to generate θ_i , which is required by r_θ). Unless \mathcal{R}' satisfies the stronger condition (i) or (ii) (for each $i \in [k]$) listed in the statement of Proposition 1, \mathcal{R}' is not minimal. \square

Another question that Part (v) of Theorem 1 suggests is the following: does an elementary CRS always have at most a polynomial number of closed RAFs? Again, the answer is ‘no’, and the construction to show this is much simpler than the previous example. Consider the elementary CRS with $F = \{f_1, \dots, f_n\}$, $X = F \cup \{x_1, \dots, x_n\}$, together with the set \mathcal{R} of k catalysed reactions $r_i : f_i \xrightarrow{x_i} x_i$ for $i = 1, \dots, k$. This CRS has $2^k - 1$ closed RAFs, one for each nonempty subset of \mathcal{R} .

The probability of a RAF in an elementary CRS

Given an elementary CRS \mathcal{Q} , suppose that catalysis is assigned randomly as follows: each molecule type catalyses each given reaction in \mathcal{R} with a fixed probability p , independently across all pairs (x, r) of molecule type x and reaction r . The probability $p_{\mathcal{Q}}$ that \mathcal{Q} has a RAF is simply the probability that $\mathcal{D}_{\mathcal{Q}}$ has a directed cycle (by Theorem 1(i)).

In the case where each reaction in \mathcal{R} has just a single product, then the asymptotic behaviour of $p_{\mathcal{Q}}$ as $|\mathcal{R}| \rightarrow \infty$ is equivalent to the emergence of a directed cycle in a large random directed graph, which has been previously studied in the random graph literature.¹⁷

Here we provide a simple lower bound on $p_{\mathcal{Q}}$. Let $\lambda = p|\mathcal{R}|$ be the expected number of reactions that each molecule type catalyses. The following result gives a lower bound on $p_{\mathcal{Q}}$ that depends only on λ and which converges towards 1 as λ grows.

Proposition 2. $p_{\mathcal{Q}} \geq 1 - \left(1 - \frac{\lambda}{|\mathcal{R}|}\right)^{|\mathcal{R}|} \sim 1 - e^{-\lambda}$, where \sim denotes asymptotic equality as $|\mathcal{R}|$ grows.

Proof. Consider the probability p_r that a single reaction r has an arc to itself (such an event is sufficient but not necessary for $\mathcal{D}_{\mathcal{Q}}$ to contain a directed cycle). If r produces $m \geq 1$ products, we have $p_r = 1 - (1 - p)^m \geq pm \geq p = \lambda/|\mathcal{R}|$. The probability that no reaction has an arc to itself is therefore $\left(1 - \frac{\lambda}{|\mathcal{R}|}\right)^{|\mathcal{R}|}$. Since $(1 - x/n)^n \sim e^{-x}$, we obtain the result claimed. \square

Eigenvector analysis

A previous study¹² considered the dynamical aspects of an ‘autocatalytic set’ in a CRS, which is closely related to the notion of an RAF (our graph $\mathcal{D}_{\mathcal{Q}}$ differs from theirs in two respects, firstly the vertices here represent reactions rather than molecule types, and we also permit self-loops from a reaction to itself). We now present the analogues of these earlier dynamical findings in our setting (and formally, with proofs).

Given an elementary CRS \mathcal{Q} , let $A_{\mathcal{Q}}$ denote the adjacency matrix of the directed graph $\mathcal{D}_{\mathcal{Q}}$. Thus the rows and columns of $A_{\mathcal{Q}}$ are indexed by the reactions in \mathcal{R} in some given order, and the entry of $A_{\mathcal{Q}}$ corresponding to the pair (r, r') is 1 precisely if (r, r') is an arc of $\mathcal{D}_{\mathcal{Q}}$ and is zero otherwise. By Perron-Frobenius theory for non-negative matrices, $A_{\mathcal{Q}}$ has a non-negative real eigenvalue λ of maximal modulus (amongst all the eigenvalues) and if $\mathcal{D}_{\mathcal{Q}}$ is strongly-connected (i.e. $A_{\mathcal{Q}}$ is irreducible), then $A_{\mathcal{Q}}$ has a left (and a right) eigenvector with eigenvalue λ whose components are all positive. The following results are analogues of the former study¹² to our setting.

Proposition 3.

- (i) If \mathcal{Q} contains no RAF, then $\lambda = 0$.
- (ii) If \mathcal{Q} contains a RAF, then $\lambda \geq 1$.

(iii) If A_Q has an eigenvalue > 0 with an associated left eigenvector w , then the reactions r for which $w_r > 0$ forms a RAF for \mathcal{Q} .

Proof. Part (i) follows from Part (i) of Theorem 1, combined with the fact that the adjacency matrix A of an acyclic directed graph is nilpotent (i.e. specifically, A^{l+1} is the all-zero matrix when l is the length of a longest path in the directed graph) and thus all the eigenvalues of A are equal to zero.¹⁸ For Part (ii), if \mathcal{Q} contains a RAF, then \mathcal{D}_Q has a minimal (chordless) directed cycle (which could just be a loop on a vertex). Let w be the vector that has value 1 for each vertex in this minimal directed cycle and is zero otherwise. Then w is both a left and right eigenvector for A_Q with eigenvalue 1. For Part (iii), let $\mathcal{R}' = \{r \in \mathcal{R} : w_r > 0\}$. The condition $wA_Q = \lambda w$ translates as $\sum_{r \in \mathcal{R}} w_r A_{rr'} = \lambda w_{r'}$. Since the right-hand side is non-zero for each reaction $r' \in \mathcal{R}'$, it follows that $w_r A_{rr'} \neq 0$ for at least one reaction $r \in \mathcal{R}'$; In other words, each reaction in \mathcal{R}' is catalysed by the product of at least one reaction in \mathcal{R}' . Since \mathcal{Q} is elementary, this implies that \mathcal{R}' is a RAF. \square

To illustrate an application of Proposition 3, consider the the system of 9 reactions from Fig. 2. In this case, $\lambda \geq 1$ since the system contains a RAF (*cf.* Proposition 3(ii)). Regarding Part (iii), three of the eigenvalues of A_Q are strictly positive, and for the three corresponding left eigenvectors, one has strictly positive entries for the three reactions r_2, r_6, r_8 , which form the subRAF S_1 shown in Fig. 2. A second left eigenvector has strictly positive entries for the reactions $r_1, r_3, r_4, r_5, r_7, r_9$, and these form the minimal closed subRAF $S_2 \cup S_3$ shown in Fig. 2. The third left eigenvector has strictly positive entries for the reactions r_1, r_4, r_7 which forms a subRAF of S_2 .

We end this section by noting that Part (iii) of Proposition 3 does not hold if left eigenvectors are substituted for right ones. A counterexample is given by the elementary CRS for which A_Q is the 2×2 matrix with both rows equal to $[0, 1]$; in this case, A_Q has a principal eigenvalue of $+1$ but the associated right eigenvector is a column vector with strictly positive entries, but this does not correspond to a RAF for \mathcal{Q} .

Generative RAFs

We now introduce a new notion which describes how simple RAFs can develop into more complex ones in a progressive way. This section will build on, and apply the results concerning elementary CRSs, particularly Theorem 1.

Given a CRS $\mathcal{Q} = (X, \mathcal{R}, C, F)$ and a subset Y of X containing F , let $\mathcal{R}|Y$ be the subset of reactions in \mathcal{R} that have all their reactants in Y , and let

$$\mathcal{Q}|Y := (X, \mathcal{R}|Y, C, Y).$$

In other words, $\mathcal{Q}|Y$ is the CRS obtained from \mathcal{Q} by deleting each reaction from \mathcal{R} that does not have all its reactants in Y , and by expanding the food set to include all of Y .

Definition (generative RAF sequence): Given any CRS (i.e. not necessarily an elementary CRS) $\mathcal{Q} = (X, \mathcal{R}, C, F)$, a *generative RAF sequence* for \mathcal{Q} is a sequence $\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_k$ of subsets of \mathcal{R} that satisfies the following two properties:

- \mathcal{R}_1 is a closed RAF of $\mathcal{Q}|F$;
- for each $i > 1$, \mathcal{R}_i is a closed RAF of $\mathcal{Q}|Y_i$ where $Y_i = F \cup \pi(\mathcal{R}_{i-1})$, and where $\pi(\mathcal{R}_{i-1})$ refers to all molecule types that are produced by a reaction from \mathcal{R}_{i-1} .

\mathcal{R}_1 is thus a closed RAF of the elementary CRS $\mathcal{Q}|F$, and for each $i > 1$, we add the products of \mathcal{R}_{i-1} to the food set F of \mathcal{Q} and take \mathcal{R}_i to be a closed RAF of the resulting (induced) elementary CRS. In other words, the next closed RAF is built upon an enlarged food set generated by the previous closed RAFs in the sequence. The sets in any generative RAF sequences are nested, as the following result shows.

Lemma 1. *If $\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_k$ is a generative sequence for \mathcal{Q} then $\mathcal{R}_i \subseteq \mathcal{R}_{i+1}$ for each $i \in \{1, \dots, k-1\}$.*

Proof. For each reaction $r \in \mathcal{R}$, let $\rho(r)$ denote the set of reactants of r . We prove Lemma 1 by induction on k . For $k = 2$, suppose that $r \in \mathcal{R}_1$. Then $\rho(r) \subseteq F$ and there exists some

molecule type $x \in F \cup \pi(\mathcal{R}_1)$ that catalyses r . Since \mathcal{R}_2 is a closed RAF, and since the reactants and at least one catalyst (namely x) are available in the enlarged food set for \mathcal{R}_2 , namely $Y_2 = F \cup \pi(\mathcal{R}_1)$, then $r \in \mathcal{R}_2$. Thus Lemma 1 holds for $k = 2$.

Suppose now that Lemma 1 holds for $k = m$ and that $\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_{m+1}$ is a generative sequence for \mathcal{Q} . We need to show that $\mathcal{R}_m \subseteq \mathcal{R}_{m+1}$. To this end, suppose that $r \in \mathcal{R}_m$. Then $\rho(r) \subseteq \pi(\mathcal{R}_{m-1})$ and there exists a molecule type $x \in F \cup \pi(\mathcal{R}_m)$ that catalyses r . Now $\mathcal{R}_{m-1} \subseteq \mathcal{R}_m$ by induction and so $\rho(r) \subseteq \pi(\mathcal{R}_m)$. Thus the reactants and at least one catalyst of r are in $Y_{m+1} = F \cup \pi(\mathcal{R}_m)$, and so, by the closure property, $r \in \mathcal{R}_{m+1}$. This establishes the induction step, and thereby the lemma. □

By Lemma 1, the sequence $\overline{\mathcal{R}}_1, \overline{\mathcal{R}}_2, \dots$, either terminates with the empty set at $\overline{\mathcal{R}}_1$ or it stabilises at some subset $\overline{\mathcal{R}}(\mathcal{Q}) := \cup_{i \geq 1} \overline{\mathcal{R}}_i$, which must equal $\overline{\mathcal{R}}_j$ for some sufficiently large value of j (with $j \leq |\mathcal{R}|$).

Definition (generative RAF):

Given a CRS \mathcal{Q} and a RAF \mathcal{R}' for \mathcal{Q} , we say that \mathcal{R}' is a *generative* RAF if there is a generative RAF sequence $\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_k$ for \mathcal{Q} with $\mathcal{R}_k = \mathcal{R}'$.

The motivation for considering this notion of RAF is two-fold. Firstly a generative RAF can be built up from simpler RAFs (starting with an elementary one) by generating the required catalysts at each step (i.e. some reactions may still need to proceed initially uncatalysed, but a catalyst for the reaction will be generated by some other reaction by the end of the same step). This avoids the possibility of long chains of reactions that need to proceed uncatalysed until a catalyst for the very first link in the chain is produced, which seems less biochemically plausible. A second motivation for considering generative RAFs is that they combine two further desirable properties: namely an emphasis on RAFs that are closed (i.e. all reactions that are able to proceed and for which a catalyst is available will proceed), and generative RAFs are sufficiently well structured that some questions can be

answered in polynomial time that are problematic for general RAFs (Theorem 2(iv) provides an explicit example).

Note that a generative RAF is necessarily, by definition, a closed RAF. A natural question at this point is the following: Is every closed RAF in a CRS generative? The answer to this is ‘no’ in general; for example, a CRS may have a maxRAF that requires too much ‘jumping ahead’ with catalysis (chains of initially spontaneous reactions) to be built up in this way. Fig. 6 shows an example of a generative RAF and a closed RAF that is not generative.

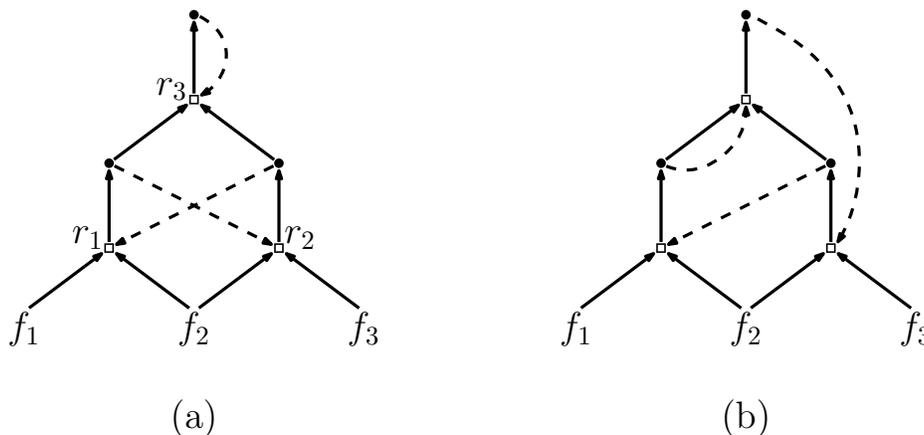
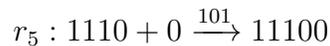
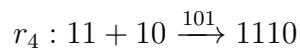
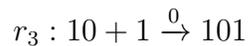
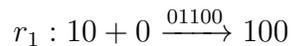


Figure 6: (a): A closed RAF that is generative (a generative RAF sequence starts with the elementary closed RAF $\{r_1, r_2\}$, and then adds r_3). (b): A different pattern of catalysis converts the three reactions into a RAF that is no longer generative. In both cases, the food set is $F = \{f_1, f_2, f_3\}$.

Another instructive example is the following maxRAF that arose in a study of the binary polymer model:³



where $F = \{0, 1, 00, 01, 10, 11\}$. This maxRAF contains six subRAFs, two of which are closed, namely, the full set of all five reactions, which is not generative, and the subset $\{r_3, r_4, r_5\}$, which is generative.

Lemma 2. *The closure of the union of any finite collection of generative RAFs is a generative RAF. Thus if \mathcal{Q} has a generative RAF, there is a unique maximal generative RAF for \mathcal{Q} that contains all other generative RAFs for \mathcal{Q} .*

Proof. Given a CRS \mathcal{Q} and generative RAFs \mathcal{R}' and \mathcal{R}'' , we first show that the closure of $\mathcal{R}' \cup \mathcal{R}''$ is a generative RAF. Let $(\mathcal{R}'_i, i = 1, \dots, k_1)$ and $(\mathcal{R}''_i, i = 1, \dots, k_2)$ be two generative RAF sequences, the first for \mathcal{R}' and the second for \mathcal{R}'' . Consider the following sequence $(\mathcal{R}_i, i = 1, \dots, k_1 + k_2)$ of subsets of \mathcal{R} : For $i = 1, \dots, k_1$, set $\mathcal{R}_i := \mathcal{R}'_i$, and for $i = k_1 + j$ (where $0 < j \leq k_2$), set \mathcal{R}_i to be the closure of $\mathcal{R}' \cup \mathcal{R}''_j$. For each i , \mathcal{R}_i is then a closed RAF for \mathcal{Q} , $\mathcal{R}_{k_1+k_2} = \mathcal{R}' \cup \mathcal{R}''$ and \mathcal{R}_i is a generative RAF sequence for $\mathcal{R}' \cup \mathcal{R}''$. The first sentence of Lemma 2 now follows by induction on the size of the collection of generative RAFs, and the second sentence follows by applying the result to the (finite) collection of all generative RAFs for \mathcal{Q} . \square

We now describe how to construct the maximal generative RAF described in Lemma 2 by a procedure that runs in polynomial time in the size of \mathcal{Q} .

Given a CRS $\mathcal{Q} = (X, \mathcal{R}, C, F)$, consider the following sequence $(\overline{\mathcal{R}}_i, i \geq 1)$ of subsets of \mathcal{R} . Let $\overline{\mathcal{R}}_1 = \text{maxRAF}(\mathcal{Q}|F)$, and for $i > 1$, let:

$$\overline{\mathcal{R}}_i := \text{maxRAF}(\mathcal{Q}|F \cup \pi(\overline{\mathcal{R}}_{i-1})).$$

Note that $\overline{\mathcal{R}}_1$ may be empty even if \mathcal{Q} has a RAF (as Fig. 6(b) shows), in which case, $\overline{\mathcal{R}}_i = \emptyset$ for all $i \geq 1$. However, if $\overline{\mathcal{R}}_1$ is nonempty, then $\overline{\mathcal{R}}_i$ forms an increasing nested sequence of RAFs, by Lemma 1.

We can now state the second main result of this paper, the proof of which is provided in the Appendix.

Theorem 2. *Suppose that $\mathcal{Q} = (X, \mathcal{R}, C, F)$ is a CRS.*

- (i) *If $\overline{\mathcal{R}}_1 \neq \emptyset$, then $\overline{\mathcal{R}}_k$ is a generative RAF for each $k \geq 1$. Moreover, $\overline{\mathcal{R}}(\mathcal{Q})$ is a generative RAF, and it contains all other generative RAFs for \mathcal{Q} .*
- (ii) *If \mathcal{R}' is a closed RAF for \mathcal{Q} then \mathcal{R}' is a generative RAF for \mathcal{Q} if and only if $\mathcal{R}' = \overline{\mathcal{R}}(\mathcal{Q}')$, where $\mathcal{Q}' = (X, \mathcal{R}', C, F)$.*
- (iii) *The construction of $\overline{\mathcal{R}}(\mathcal{Q})$ and the determination of whether an arbitrary closed RAF for \mathcal{Q} is generative can be determined in polynomial time in the size of \mathcal{Q} .*
- (iv) *There is a polynomial-time algorithm that can determine whether or not a generative RAF \mathcal{R}' contains another generative RAF as a strict subset.*

Remarks:

- If a CRS has a CAF (defined at the end of the Introduction), then the (unique) maximal CAF is generative. However, a generative RAF need not necessarily correspond to a maximal CAF.
- One could further insist that the closed RAFs in a generative RAF sequence are chosen to be minimal closed RAFs and/or non-inhibited, and/or the most dynamically favourable RAFs, as Theorem 1(v) ensures that these RAFs can be identified in polynomial time.
- Part (iv) of Theorem 2 provides an interesting contrast to the general RAF setting for which determining whether a closed RAF (e.g. the maxRAF) in an arbitrary CRS contains another closed RAF as a strict subset has unknown complexity.

RAFs with reaction rates

In this third and final section, we consider a further refinement of RAF theory, by explicitly incorporating reaction rates into the analysis. This conveniently addresses one shortcoming

implicit in the generative RAF definition from the last section – namely a generative RAF necessarily grows as a monotonically increasing nested system with the length of its associative generative RAF sequence (Lemma 1). However, once a sufficiently large generative RAF is established, one or more of its subRAFs may then become dynamically favoured if it is more ‘efficient’ (i.e. all its reactions proceed at higher reaction rates than the generative RAF it lies within), as we shortly illustrate with a simple example.

Suppose that we have a CRS $\mathcal{Q} = (X, \mathcal{R}, C, F)$ and a function $f : C \rightarrow \mathbb{R}^{\geq 0}$ that assigns a non-negative real number to each pair $(x, r) \in C$. The interpretation here is that $f(x, r)$ describes the *rate* at which reaction r proceeds when the catalyst x is present.

Given \mathcal{Q} and f , together with a RAF \mathcal{R}' for \mathcal{Q} , let:

$$\varphi(\mathcal{R}') = \min_{r \in \mathcal{R}'} \{ \max \{ f(x, r) : (x, r) \in C, x \in \text{cl}_{\mathcal{R}'}(F) \} \}$$

In other words, $\varphi(\mathcal{R}')$ is the rate of the slowest reaction in the RAF \mathcal{R}' under the most optimal choice of catalyst for each reaction in \mathcal{R}' amongst those catalysts that are present in $\text{cl}_{\mathcal{R}'}(F)$.

Example: Fig. 7 provides an example to illustrate the notions above. In this CRS the three reactions comprises a RAF, with a φ -value equal to 1. However there are three subRAFs, and one of these (namely $\{r_2, r_3\}$) has a higher φ -value. However, the less optimal closed subRAF $\{r_1, r_2\}$ is generative and likely to have formed before the optimal one; otherwise $\{r_2, r_3\}$ would require a chain of two reactions to occur uncatalysed (r_2 followed by r_3) before the catalysts for them become available. The closed RAF $\{r_1, r_2\}$ may then expand to $\{r_1, r_2, r_3\}$ before this second closed RAF is subsequently out-competed by its subRAF $\{r_2, r_3\}$, since the catalysed reactions in this subRAF run twice as fast as the reaction r_1 .

Our third main result in this paper shows that finding a RAF to maximise φ can be achieved by an algorithm that runs in polynomial time in the size of \mathcal{Q} . The proof is

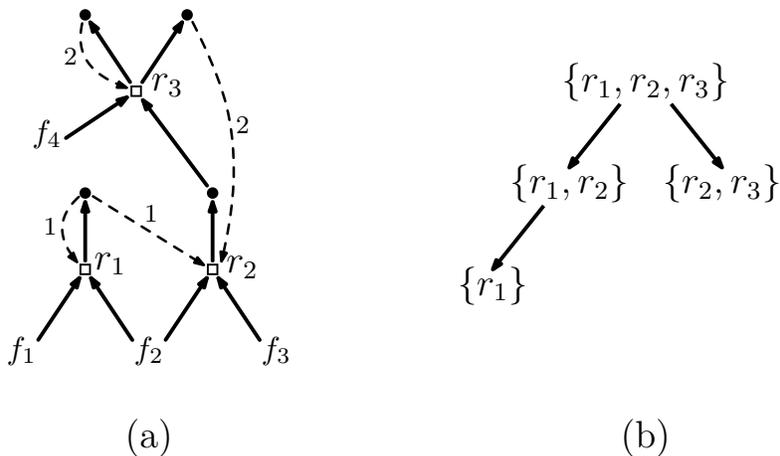


Figure 7: (a) A RAF in which the catalysis arcs have associated rates (namely, the values 1 and 2 as indicated). The poset consisting of the maxRAF and its three subRAFs (partially ordered by set inclusion) is shown by the Hasse diagram in (b). All four RAFs have φ -values of 1 except for the subRAF $\{r_2, r_3\}$, which has a φ -value of 2. This optimal RAF $\{r_2, r_3\}$ is not a generative RAF (whereas the other three RAFs are generative; indeed, $\{r_1\}$ and $\{r_1, r_2\}$ are elementary). Nevertheless, once the generative maxRAF $\{r_1, r_2, r_3\}$ has formed, $\{r_2, r_3\}$ can then emerge as the dominant sub-RAF.

provided in the Appendix.

Theorem 3. *There is a polynomial-time algorithm to construct a RAF with largest possible φ -value from any CRS \mathcal{Q} that contains RAF. Moreover, this constructed RAF is the maximal RAF with this φ -value.*

Remark: For the example in Fig. 7, we have the subRAFs $\mathcal{R}_1 = \{r_1\}$, $\mathcal{R}_2 = \{r_2, r_3\}$ with $\varphi(\mathcal{R}_1) < \varphi(\mathcal{R}_2)$. In this case, there is a path in the poset from \mathcal{R}_1 to \mathcal{R}_2 on which φ is non-decreasing (this path goes ‘up’ then ‘down’ in Fig. 7(b)). An interesting question might be to determine when this holds: in other words, from a sub-optimal RAF, can a more optimal RAF be reached by a chain of RAFs that, at each stage, either adds certain reactions or deletes one or more reactions, and so that the optimality score (as measured by φ) does not decrease?

Rates for ‘catalytic ensembles’

We can extend the results on rates in the previous section to accommodate the following feature: a reaction for which a combination of two or more catalysts is present may proceed at a rate that is higher than if just one catalyst is present.

We formalize this as follows. Recall that in a CRS $\mathcal{Q} = (X, \mathcal{R}, \mathcal{C}, F)$, the set \mathcal{C} represents the pattern of catalysis and is a subset of $X \times \mathcal{R}$. Thus $(x, r) \in \mathcal{C}$ means that x catalyses reaction r . Now suppose we wish to allow a combination (ensemble) of one or more molecules to act as a catalyst for a reaction. In this case, we can represent the CRS as a quadruple $\mathcal{Q} = (X, \mathcal{R}, \mathcal{C}, F)$ where $\mathcal{C} \subseteq (2^X - \emptyset) \times \mathcal{R}$ and where $(A, r) \in \mathcal{C}$ means that the ensemble of molecules in A acts as a (collective) catalyst for r , provided they are all present. We refer to \mathcal{Q} as a *generalised* CRS. The notions of RAF, subRAF, CAF, and so on, can be generalized naturally. For example, the RA condition for a subset \mathcal{R}' is that for each reaction r , there is a pair $(A, r) \in \mathcal{C}$ where each of the molecule types in A is in the closure of F relative to \mathcal{R}' .

Note that an ordinary CRS can be viewed as a special case of a generalised CRS by identifying (x, r) with the pair $(\{x\}, r)$. Note also that each reaction may have several ensembles of possible catalysts, and some (or all of these) may be just single molecule types.

Given a generalised CRS $\mathcal{Q} = (X, \mathcal{R}, \mathcal{C}, F)$ we can associate an ordinary CRS $\mathcal{Q}' = (X', \mathcal{R}', \mathcal{C}', F)$ to \mathcal{Q} as follows. Let

$$\mathcal{A}_C := \{A \subseteq 2^X - \emptyset : \exists r \in \mathcal{R} : (A, r) \in \mathcal{C}\};$$

(so \mathcal{A}_C is the collection of catalyst ensembles in \mathcal{Q}). For each $A \in \mathcal{A}_C$, let x_A be a new molecule type, and let r_A be the (formal) reaction $A \rightarrow x_A$. Now let

$$X' := X \dot{\cup} \{x_A : A \in \mathcal{A}_C\};$$

$$\mathcal{R}' := \mathcal{R} \dot{\cup} \{r_A : A \in \mathcal{A}_C\}; \text{ and}$$

$$C' := \{(x_A, r) : (A, r) \in \mathcal{C}\} \dot{\cup} \{(x_A, r_A) : A \in \mathcal{A}_C\}.$$

Note that $C' \subseteq X' \times \mathcal{R}'$.

In other words, \mathcal{Q}' is obtained from \mathcal{Q} by replacing each catalytic ensemble A by a new molecule type x_A and adding in the reaction $r_A : A \rightarrow x_A$ catalysed by x_A . The proof of the following lemma is straightforward.

Lemma 3. *A generalised CRS \mathcal{Q} has a RAF if and only if the associated ordinary CRS \mathcal{Q}' has a RAF that contains at least one reaction from \mathcal{R} . Moreover, in this case, the RAFs of \mathcal{Q} correspond to the nonempty intersections of RAFs of \mathcal{Q}' with \mathcal{R} .*

Now suppose that we have a generalised CRS $\mathcal{Q} = (X, \mathcal{R}, \mathcal{C}, F)$ and a function $f : \mathcal{C} \rightarrow \mathbb{R}^{\geq 0}$. The interpretation here is that $f(A, r)$ describes the *rate* at which reaction r proceeds when the catalyst ensemble A is present.

Given a RAF \mathcal{R}' for \mathcal{Q} , let:

$$\varphi(\mathcal{R}') := \min_{r \in \mathcal{R}'} \{ \max \{ f(A, r) : (A, r) \in \mathcal{C}, A \subseteq \text{cl}_{\mathcal{R}'}(F) \} \}$$

In other words, $\varphi(\mathcal{R}')$ is the rate of the slowest reaction in the RAF \mathcal{R}' under the most optimal choice of catalyst ensemble for each reaction in \mathcal{R}' amongst catalyst ensembles that are subsets of $\text{cl}_{\mathcal{R}'}(F)$.

Lemma 3 now provides the following corollary of Theorem 3.

Corollary 2. *There is a polynomial-time algorithm to construct a RAF for \mathcal{Q} with largest possible φ -value from any CRS \mathcal{Q} that contains a RAF. Moreover, this constructed RAF is the maximal RAF for \mathcal{Q} with this φ -value.*

Concluding comments

In this paper, we have considered special types of RAFs that allow for exact yet tractable mathematical and algorithmic analysis, and which also incorporate additional biochemical

realism (restricting the depth of uncatalysed reactions chains in generative RAFs and allowing reaction rates).

We first considered the special setting of ‘elementary’ systems in which all reactions (or at least those present in the maxRAF) have all their reactants present in the food set. This allows for the structure of the collection of RAFs, irrRAF, and closed subRAF to be explicitly described graph-theoretically. As a result, some problems that are computationally intractable in the general CRS setting turn out to be polynomial-time for an elementary CRS. For example, one can efficiently find the smallest RAFs in an elementary CRS, which is an NP-hard problem in general.¹⁶ Also, the number of minimal closed subRAF in an elementary CRS is linear in the size of the set of reactions (for a general CRS, they can be exponential in number). For future work, it may be of interest to determine if there are polynomial-time algorithms that can answer the following questions for an elementary CRS: (i) What is the size of the largest irrRAF? (ii) If inhibition is allowed, then is there a RAF that has no inhibition?

The concept of an ‘elementary’ CRS is an all-or-nothing notion. One way to extend the results above could be to define the notion of ‘level’, whereby a CRS has level k if the length of the longest path from the food set to any reaction product goes through at most k reactions (an elementary CRS thus has level 1). We have not explored this further here but instead, we consider the related alternative notion of a generative RAF. Briefly, a generative RAF allows a RAF to form by effectively enlarging its ‘food set’ with products of reactions, so that each step only requires catalysts that are either present or produced by reactions in the RAF at that stage. Although generative RAFs are more complex than elementary ones, their close connection to elementary RAFs (in a stratified way) allows for a more tractable analysis than for general RAFs. Moreover, unlike elementary RAFs, no special assumption is required on the underlying CRS; generative RAFs are just a special type of RAF that can be generated in a certain sequential fashion in any CRS.

In the final section, we considered the impact of rates of RAFs (which need not be

generative), and particularly the algorithmic question of finding a RAF that maximises the rates of its slowest reaction. Not only is this problem solvable in the size of the CRS, but it can also be extended to the slightly more general setting of allowing ‘catalytic ensembles’. The introduction of rates allows for the study of how a population of different closed subRAFs might evolve over time, in which primitive subRAF are replaced (out-competed) by efficient ones that rely on new catalysts in place of more primitive ones. We hope to explore this further in future work.

Acknowledgments

WH thanks the Institute for Advanced Study, Amsterdam, for financial support in the form of a fellowship.

References

- (1) Kauffman, S. A. *J. Theor. Biol.* **1986**, *19*, 1–24.
- (2) Kauffman, S. A. *The Origins of Order*; Oxford University Press, 1993.
- (3) Hordijk, W.; Steel, M. *Biosystems* **2017**, *152*, 1–10.
- (4) Jaramillo, S.; Honorato-Zimmer, R.; Pereira, U.; Contreras, D.; Reynaert, B.; Hernández, V.; Soto-Andrade, J.; Cárdenas, M. L.; Cornish-Bowden, A.; Letelier, J. C. (M,R) systems and RAF sets: common ideas, tools and projections. Proceedings of the ALIFEXII Conference, Odense Denmark, August 2010. 2010; pp 94–100.
- (5) Hordijk, W.; Steel, M.; Dittrich, P. *New J. Phys.* **2017**, DOI:10.1088/1367-2630/aa9fcd.
- (6) Sousa, F. L.; Hordijk, W.; Steel, M.; Martin, W. *J. Syst. Chem.* **2015**, *6*, 4.

- (7) Vaidya, N.; Manapat, M. L.; A., C. I.; Xulvi-Brunet, R.; Hayden, N., E. J. Lehman *Nature* **2012**, *491*, 72–77.
- (8) Gatti, R. C.; Hordijk, W.; Kauffman, S. *Ecol. Model.* **2017**, *346*, 70–76.
- (9) Gabora, L.; Steel, M. *J. Theor. Biol.* **2017**, *63*, 617–638.
- (10) Hordijk, W.; Steel, M. *J. Theor. Biol.* **2004**, *227*, 451–461.
- (11) Ashkenasy, G.; Jegasia, R.; Yadav, M.; Ghadiri, M. R. *Proc. Natl. Acad. Sci. USA* **2004**, *101*, 10872–10877.
- (12) Jain, S.; Krishna, S. *Phys. Rev. Lett.* **1998**, *81*, 5684–5687.
- (13) Tarjan, R. E. *SIAM J. Comput.* **1972**, *1*, 146–160.
- (14) Hordijk, W.; Steel, M.; Kauffman, S. *Acta Biotheor.* **2012**, *60*, 379–392.
- (15) Hordijk, W.; Steel, M. *J. Syst. Chem.* **2013**, *4*, 3.
- (16) Steel, M.; Hordijk, W.; Smith, J. *J. Theor. Biol.* **2013**, *332*, 96–107.
- (17) Bollobas, B.; Rasmussen, S. *Discrete Math.* **1989**, *75*, 55–68.
- (18) Cvetković, D. M.; Doob, M.; Sachs, H. *Spectra of Graphs, third ed.*; Barth, Heidelberg, 1995.

Appendix: Proofs of Theorems 1–3

Proof of Theorem 1

A key observation throughout is that in an elementary CRS \mathcal{Q} , any nonempty subset \mathcal{R}' of reactions automatically satisfies the F -generated property, so \mathcal{R}' forms a RAF for \mathcal{Q} if and only if \mathcal{R}' satisfies the reflexively autocatalytic (RA) property. By the manner in which $\mathcal{D}_{\mathcal{Q}}$

is constructed, the RA property means that the induced subgraph $\mathcal{D}_{\mathcal{Q}}|\mathcal{R}'$ has the property that each vertex has in-degree at least 1.

In particular, \mathcal{R} has a RAF if and only if $\mathcal{D}_{\mathcal{Q}}$ has a directed cycle. The ‘if’ direction of this claim is clear. For the ‘only if’ direction, suppose that \mathcal{R}' is a RAF and $r \in \mathcal{R}'$. By the assumption that each vertex in $\mathcal{D}_{\mathcal{Q}}$ has in-degree at least 1, there is a directed walk of length k (for any $k \geq 1$) involving vertices in \mathcal{R}' and ending in r . Since \mathcal{R}' is finite if we take $k > |\mathcal{R}'|$ then two vertices on this directed walk must coincide and the resulting subwalk between this vertex to itself gives a directed cycle in $\mathcal{D}_{\mathcal{Q}}$. Moreover, $\mathcal{D}_{\mathcal{Q}}|\mathcal{R}'$ contains a directed cycle if and only if this sub-digraph contains a chordless cycle (again, the ‘if’ direction is clear and the ‘only if’ direction follows by the finiteness of \mathcal{R} , so shortening each directed cycle by following a chord leads to a sequence of cycles of decreasing length that eventually terminates on a chordless cycle). This establishes Part (i).

For Part (ii), a subset \mathcal{R}' of \mathcal{R} has the property that $\mathcal{D}_{\mathcal{Q}}|\mathcal{R}'$ is a chordless cycle, which implies (by Part (i)) that \mathcal{R}' is a RAF. It is also an irrRAF; otherwise, the cycle would have a chord. Conversely, if \mathcal{R}' has the property that $\mathcal{D}_{\mathcal{Q}}|\mathcal{R}'$ is not a chordless cycle, then either $\mathcal{D}_{\mathcal{Q}}$ does not contain a cycle (in which case it is not a RAF) or it contains a cycle which either has a chord or has other vertices reachable from it, in which case \mathcal{R}' is not an irrRAF. This establishes the first sentence of Part (ii). The arguments for the second and third sentences follow similar lines.

For Part (iii), it is clear that the union of one or more cores is a RAF; however, the resulting set of reactions is closed if and only if all reactions that are reachable from that set are also included.

For Part (iv), suppose that a core c' is reachable from another core c (by definition, c is not reachable from c'). Any closed RAF \mathcal{R}' that contains both c and c' is thus not minimal, since we could delete c' and all the reactions that are reachable from c' but not from c and obtain a strict subset of \mathcal{R}' that is also a closed RAF. On the other hand, if \mathcal{R}' has the property described in Part (iv), then it is a closed RAF by Part (iii) and it is also minimal,

since any closed RAF must contain at least one core, alongside all the reactions that are reachable from it.

Part (v) follows from Part (iv), since each minimal closed RAF is associated with exactly one core, and since cores are strongly connected components of $\mathcal{D}_{\mathcal{Q}}$ these cores are vertex-disjoint (i.e. two cores share no reaction). Consequently, the number of cores is bounded above by the number of reactions in the maxRAF of $\mathcal{D}_{\mathcal{Q}}$. Moreover, finding the strongly connected components of any digraph can be done in polynomial time in the size of the digraph.¹³ Each of these strongly connected components can then be tested in polynomial time to determine if it is a core; if so, one can then determine in polynomial time which other vertices are reachable from it. Thus the minimal closed RAFs can be listed in polynomial time in the size of \mathcal{Q} .

Part (vi) follows from Part (v) since \mathcal{Q} contains a closed subRAF if and only if it contains a minimal closed subRAF. □

Proof of Theorem 2

For Part (i), the proof of the first claim is by induction on k . For the base case $k = 1$, if $\overline{\mathcal{R}}_1 \neq \emptyset$, then $\overline{\mathcal{R}}_1$ is the maxRAF of $\mathcal{Q}|F$ and so is a closed RAF of \mathcal{Q} . Now suppose that $\overline{\mathcal{R}}_{k-1}$ is a generative RAF. In this case, $\overline{\mathcal{R}}_k$ is the maxRAF of $\mathcal{Q}|F \cup \pi(\mathcal{R}_{k-1})$ and so $\overline{\mathcal{R}}_k$ is a closed RAF of this CRS. For the second part, observe that if $(\mathcal{R}'_i, i \geq 1)$ is a generative RAF sequence for a generative RAF \mathcal{R}' , then, by induction on k , $\mathcal{R}'_k \subseteq \overline{\mathcal{R}}_k$, since $\overline{\mathcal{R}}_k$ is constructed by taking the maxRAF of $\overline{\mathcal{R}}_{k-1}$, so it contains any other closed RAF of $\overline{\mathcal{R}}_{k-1}$.

For Part (ii), observe that $\overline{\mathcal{R}}(\mathcal{Q}')$ is, by definition, a subset of \mathcal{R}' . Since $\overline{\mathcal{R}}(\mathcal{Q}')$ is generative (by Part (i)), the ‘if’ direction is established. Conversely, suppose that \mathcal{R}' is generative for \mathcal{Q} . Then \mathcal{R}' is also generative for \mathcal{Q}' . Since $\overline{\mathcal{R}}(\mathcal{Q}') \subseteq \mathcal{R}'$, and since $\overline{\mathcal{R}}(\mathcal{Q}')$ is the maximal generative RAF for \mathcal{Q} , it follows that $\overline{\mathcal{R}}(\mathcal{Q}') = \mathcal{R}'$.

For Part (iii), the proof of the first claim follows from the fact that the maxRAF can be computed in polynomial time in the size of the CRS;¹⁰ the second claim now follows from

Part (ii).

For Part (iv), the algorithm first constructs the generative RAF sequence $\overline{\mathcal{R}}_1, \overline{\mathcal{R}}_2, \dots$ to obtain $\overline{\mathcal{R}}(\mathcal{Q}') = \mathcal{R}'$ (where the last equality is from Part (ii)). One then sequentially applies Part (vi) of Theorem 1 to determine whether there is a closed RAF strictly within \mathcal{R}_1 , and if not whether there is within \mathcal{R}_2 , and so on. If no such closed (and strict subset) RAF is detected then \mathcal{R}' contains no other generative RAF as a strict subset, otherwise any such subset found is a generative RAF that lies strictly inside \mathcal{R}' . \square

Proof of Theorem 3

Let $\mathcal{L} = \{f(x, r) : (x, r) \in C\}$, and let $M = \max \mathcal{L}$. Consider the CRS $\mathcal{Q}' = (X', \mathcal{R}^*, C^*, F)$ obtained from \mathcal{Q} by first deleting any uncatalysed reaction and then replacing each reaction r that is catalysed by (say) $k \geq 1$ molecule types with k distinct copies of this reaction r_1, \dots, r_k , each of which is catalysed by a different one of the k molecule types. Thus each reaction r in \mathcal{R}^* is catalysed by exactly one molecule type, which we will denote as $x(r)$. For the associated catalysis set $C^* = \{(x(r), r) : r \in \mathcal{R}^*\}$, let f' be the rate function induced by f (i.e. if $r \in \mathcal{R}$ is replaced by $r_1, \dots, r_k \in \mathcal{R}^*$ then $f'(x(r_i), r_i) := f(x(r), r)$). For each $\ell \in \mathcal{L}$ let:

$$\mathcal{R}_\ell^* = \{r \in \mathcal{R}^* : f'(x(r), r) \geq \ell\}.$$

In other words, \mathcal{R}_ℓ^* is the set of catalyst-reaction pairs $(x(r), r)$ where the rate of reaction r when catalysed by the molecule type $x(r)$ is at least ℓ (as specified by the rate function f).

Now, let $\tilde{\mathcal{R}}$ be the maxRAF of \mathcal{R}_ℓ^* for the largest value of $\ell \in \mathcal{L}$ for which $\text{maxRAF}(\mathcal{R}_\ell^*)$ is nonempty. This set is well-defined, since $\mathcal{R}^* = \mathcal{R}_\ell^*$ when $\ell = 0$, and because \mathcal{R} (and thereby \mathcal{R}^*) is assumed to have a RAF. Notice that $\tilde{\mathcal{R}}$ can be efficiently determined, by starting at $\ell = M$ and decreasing ℓ through the (at most $|\mathcal{L}| \leq |C|$) possible values it can take until a nonempty maxRAF first appears (alternatively, one could start at $\ell = 0$ and increase ℓ until the last value for which a nonempty maxRAF is present).

Claim: $\tilde{\mathcal{R}}$ is a RAF that has the largest possible φ -value of any RAF for \mathcal{Q}' , and $\tilde{\mathcal{R}}$ contains any other RAF for \mathcal{Q}' with this maximal φ -value.

To establish this claim, suppose that $\tilde{\mathcal{R}} = \text{maxRAF}(\mathcal{R}_\ell)$ for $\ell = t$ and that $\text{maxRAF}(\mathcal{R}_\ell) = \emptyset$ for $\ell > t$ (i.e. t is the largest value of ℓ in \mathcal{L} for which \mathcal{R}_ℓ has a (nonempty) maxRAF). For each reaction r in $\tilde{\mathcal{R}}$, we then have $f'(x(r), r) \geq t$, and for at least one reaction r in $\tilde{\mathcal{R}}$, $f'(x(r), r) = t$ (otherwise, a larger value of ℓ would support a maxRAF). It follows that $\varphi(\tilde{\mathcal{R}}) = t$. Now if \mathcal{R}' is any other RAF for \mathcal{Q}' , let t' be the minimal value of $f'(x(r), r)$ over all choices of $r \in \mathcal{R}'$. Then $t' \leq t$ otherwise, \mathcal{R}_ℓ would have a nonempty maxRAF for a value $\ell = t'$ that is greater than t , contradicting the maximality of t . Thus $\mathcal{R}' \subseteq \mathcal{R}_t^*$ and so

$$\mathcal{R}' = \text{maxRAF}(\mathcal{R}') \subseteq \text{maxRAF}(\mathcal{R}_t^*) = \tilde{\mathcal{R}},$$

which shows that $\tilde{\mathcal{R}}$ contains any other RAF with this maximal value.

This establishes the above Claim, and thereby Theorem 3 for \mathcal{Q}' . However, the subset of reactions of \mathcal{R} whose copies are present in $\tilde{\mathcal{R}}$ provides a RAF for \mathcal{Q} that has the largest possible φ -value (namely t) and which contains any other RAF for \mathcal{Q} with this φ -value. \square