

# Autocatalytic sets and chemical organizations: Modeling self-sustaining reaction networks at the origin of life

Wim Hordijk<sup>1</sup>, Mike Steel<sup>2</sup> and Peter Dittrich<sup>3</sup>

<sup>1</sup> Konrad Lorenz Institute for Evolution and Cognition Research,  
Klosterneuburg, Austria

<sup>2</sup> Biomathematics Research Centre, University of Canterbury, Christchurch,  
New Zealand

<sup>3</sup> Friedrich Schiller University, Jena, Germany

E-mail: [wim@WorldWideWanderings.net](mailto:wim@WorldWideWanderings.net)

**Abstract.** Two related but somewhat different approaches have been proposed to formalize the notion of a self-sustaining chemical reaction network. One is the notion of *collectively autocatalytic sets*, formalized as RAF theory, and the other is *chemical organization theory* (COT). Both formalisms have been argued to be relevant to the origin of life. RAF sets and chemical organizations are defined differently, but previously some relationships between the two have been shown. Here, we refine and explore these connections in more detail. In particular, we show that so-called *closed RAFs* are chemical organizations, but that the converse is not necessarily true. We then introduce and apply a procedure to show how chemical organizations can be used to find all closed RAFs within any chemical reaction system. We end with a discussion of why and how closed RAFs could be important in the context of the origin and early evolution of life.

*Keywords:* Autocatalytic sets, chemical organizations, origin of life

## 1. Introduction

Currently the dominant paradigm in origin of life research is the RNA world (Gilbert 1986, Joyce 2002). However, despite experimental progress towards the spontaneous formation of RNA (Powner et al. 2009, Hud et al. 2013, Patel et al. 2015), the RNA world hypothesis still has problems (Benner et al. 2012, Szostak 2012), and so far no one has been able to show that RNA can catalyze its own template-directed replication.

An alternative view is that the origin of life most likely involved a *network* of cooperating chemical elements, all mutually “helping” (i.e., catalyzing) each other’s formation from a basic food source (Nghe et al. 2015, Hordijk & Steel 2017). Such a cooperative network is more formally described as an *autocatalytic set* (see below). Autocatalytic sets have been studied extensively recently, and have been shown to exist in computational models as well as in real chemical and biological reaction networks (Hordijk 2013, Sousa et al. 2015, Hordijk & Steel 2017).

In previous work we have shown that autocatalytic sets often consist of a hierarchical structure of smaller and smaller autocatalytic subsets (Hordijk et al. 2012), which satisfies one of the main requirements for them to be evolvable (Vasas et al. 2012, Hordijk & Steel 2014). In fact, in theory there can be an exponentially large number of autocatalytic subsets within a given reaction network (Hordijk et al. 2012), and we showed that in practice one can indeed expect a large number of them (Steel et al. 2013).

However, in a dynamical sense it turns out that many of these autocatalytic subsets are only “transient”, in the sense that they would almost immediately become part of a larger subset where all reactions that *can* happen indeed *will* happen. This has previously been defined as a *closed* autocatalytic set (Smith et al. 2014). Even though there can, in principle, also be an exponentially large number of such closed autocatalytic subsets within a given reaction network, an unresolved question so far has been how many of them can be expected in practice in arbitrary reaction networks.

A concept related to autocatalytic sets is that of chemical organizations. A *chemical organization* is a similar but more general formalism describing closed and self-maintaining chemical reaction networks (Dittrich & Speroni di Fenizio 2007). Although this formalism does not explicitly include catalysis or a food source, these can be included implicitly (explained in detail below). In earlier work, some connections between autocatalytic sets and chemical organizations were already derived (Contreras et al. 2011, Steel et al. 2013, Smith et al. 2014).

Here, we refine and explore these connections in more detail. In particular, we show that the chemical organizations within an autocatalytic set include all its closed autocatalytic subsets. This corrects a minor error in Contreras et al. (2011), by showing that only *closed* autocatalytic subsets are chemical organizations. However, we also show that not *all* chemical organizations are necessarily closed autocatalytic sets, thus answering a question that was left open in Contreras et al. (2011). We then introduce a procedure for finding all closed autocatalytic subsets within a given autocatalytic network by calculating its chemical organizations. This procedure is illustrated by applying it to a simple computational model of polymer-based reaction networks, which provides insight into the expected number of closed autocatalytic subsets in practice. We end by discussing possible consequences for origin of life studies, in which closed autocatalytic subsets are likely to be of primary interest from a dynamical and evolutionary point of view.

## 2. Methods

The methods used in this paper are all theoretical and computational. In particular, they are based on the formalisms of autocatalytic sets and chemical organizations, which are then applied to a class of computational models known as the binary polymer model. These formalisms and models are reviewed briefly in the current section.

### 2.1. RAF theory

The concept of autocatalytic sets, as used here, was originally introduced by Kauffman (1971) (Kauffman 1986, Kauffman 1993), and subsequently formalized and further developed as RAF theory (Steel 2000, Hordijk & Steel 2004, Mossel & Steel 2005, Hordijk & Steel 2017). Here, we briefly review the basic definitions and results of RAF theory.

First, we define a *chemical reaction system* (CRS) as a tuple  $\mathcal{Q} = (X, \mathcal{R}, C)$  consisting of a set  $X$  of molecule types, a set  $\mathcal{R} \subset X \times X$  of reactions, and a catalysis set  $C \subseteq X \times \mathcal{R}$ . A reaction  $r \in \mathcal{R}$  is an ordered pair  $r = (A, B)$  where  $A, B \subset X$ ,  $A \cap B = \emptyset$ , representing the reactants  $A$  and products  $B$  of a chemical reaction, respectively. Note that in the RAF framework (and its graph theoretical realization), the stoichiometric coefficients are not required, so it is sufficient to represent a reaction by a pair of sets. A *catalyzation*  $(x, r) \in C$  indicates that molecule type  $x$  catalyzes reaction  $r$ .

We also explicitly consider the notion of a food set  $F \subset X$ , which is a subset of molecule types that are assumed to be directly available from the environment. Often, we will include  $F$  in the CRS description by writing  $\mathcal{Q} = (X, \mathcal{R}, C, F)$ .

The notion of catalysis plays a central role here. A *catalyst* is a molecule that significantly speeds up the rate at which a chemical reaction happens, without being “used up” in that reaction. Catalysis is ubiquitous in life (van Santen & Neurock 2006). The majority of organic reactions are catalyzed, and catalysts are essential in determining and regulating the functionality of the chemical reaction networks that support life.

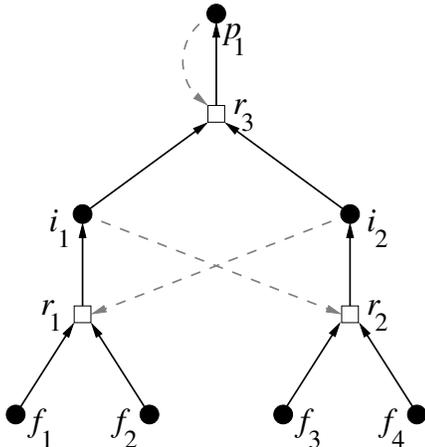
Given a set of molecules  $A \subseteq X$  and a set of reactions  $\mathcal{R}' \subseteq \mathcal{R}$ , we define the *closure*  $cl(A, \mathcal{R}') \subseteq X$  as the set of molecules obtained by adding to  $A$  all molecules that can be produced (in arbitrarily many steps) from  $A$  using reactions  $\mathcal{R}'$ .

An *autocatalytic set* (or *RAF set*, for short) is now defined as a subset  $\mathcal{R}' \subseteq \mathcal{R}$  of reactions that is:

- (i) *Reflexively Autocatalytic* (RA): each reaction  $r \in \mathcal{R}'$  is catalyzed by at least one molecule type that is either a product of  $\mathcal{R}'$  or is present in the food set  $F$ ; and
- (ii) *F-generated* (F): all reactants in  $\mathcal{R}'$  can be created from the food set  $F$  by using a series of reactions only from  $\mathcal{R}'$  itself; formally for all  $(A, B) \in \mathcal{R}'$ ,  $A \subseteq cl(F, \mathcal{R}')$ .

A formal definition of RAF sets is also provided in Hordijk & Steel (2004) and Hordijk et al. (2011), including an efficient (polynomial-time) algorithm for finding RAF sets in a general CRS, or determining there is none.

This RAF algorithm returns the unique *maximal* RAF set (maxRAF), which is the union of all RAF (sub)sets within a given CRS. If no RAF set exists within a CRS, the RAF algorithm returns the empty set. A maxRAF can often be decomposed into several smaller subsets which themselves are RAF sets, i.e., subRAFs (Hordijk et al. 2012)). If such a subRAF cannot be reduced any further without losing the RAF property, it is referred to as an *irreducible* RAF, or irrRAF (Hordijk & Steel 2004). A simple example is given in Figure 1.



**Figure 1.** A simple example of a maxRAF  $\mathcal{R}' = \{r_1, r_2, r_3\}$  that could exist within a larger CRS with food set  $F = \{f_1, f_2, f_3, f_4\}$ . Molecule types are represented by black dots and reactions by white boxes. Solid arrows indicate reactants going into and products coming out of a reaction, while dashed arrows indicate catalysis. This maxRAF contains the irreducible RAF  $\mathcal{R}_i = \{r_1, r_2\}$ .

RAF theory has been applied extensively to simple polymer-based models of chemical reaction networks (see below), showing that autocatalytic sets are highly likely to exist at modest (and realistic) catalysis rates, and under a wide variety of model assumptions (Hordijk & Steel 2004, Mossel & Steel 2005, Hordijk et al. 2011, Hordijk, Hasenclever, Gao, Mincheva & Hein 2014, Hordijk, Wills & Steel 2014, Smith et al. 2014). Moreover, these results show that generally many hierarchical levels of subRAFTs and irrRAFTs exist in such networks (Hordijk et al. 2012, Hordijk et al. 2015). The formal RAF framework has also been applied successfully to analyze real chemical and biological networks in terms of autocatalytic sets (Hordijk & Steel 2013, Sousa et al. 2015). Autocatalytic networks have even attracted attention in the physics community (Farmer et al. 1986, Stadler et al. 1993, Jain & Krishna 1998, Hordijk et al. 2010).

## 2.2. Closed RAFTs and CAFs

As we defined the closure of a set of molecules (relative to a set of reactions), we can similarly define the closure of a set of reactions (relative to a set of molecules), leading to the notion of a closed RAF. Given a CRS with a food set,  $\mathcal{Q} = (X, \mathcal{R}, C, F)$ , a (non-empty) subset  $\mathcal{R}'$  of  $\mathcal{R}$  is said to be a *closed* RAF if  $\mathcal{R}'$  is a RAF that satisfies the additional closure property (Smith et al. 2014):

- cp*:  $\mathcal{R}'$  contains each reaction  $r \in \mathcal{R}$  for which each reactant and at least one catalyst are either generated by another reaction from  $\mathcal{R}'$  or are part of the food set  $F$ .

In other words, a closed RAF is a RAF in which each reaction that *can* happen indeed *will* happen.

This notion of the *closure* of a RAF  $\mathcal{R}'$ , denoted  $\overline{\mathcal{R}'}$  was already defined in Smith et al. (2014), and  $\mathcal{R}'$  is a closed RAF if and only if  $\mathcal{R}' = \overline{\mathcal{R}'}$ . Briefly, the closure of a RAF  $\mathcal{R}'$  is obtained by starting with  $\mathcal{R}'$  and sequentially adding new reactions from  $\mathcal{R}$  if their reactants and at least one catalyst are either an element of the food

set or a product of a reaction in the set so-far constructed. The notion of closure is well-defined, i.e., it does not depend on the order in which reactions are added (Smith et al. 2014). The irreducible RAF  $\mathcal{R}_i = \{r_1, r_2\}$  in Figure 1 is also a closed RAF, as is the maxRAF  $\mathcal{R}' = \{r_1, r_2, r_3\}$ .

Note that a maxRAF in any CRS is always closed. Thus, any CRS that has a RAF has a unique maximal closed RAF (namely the maxRAF). However, as pointed out in Smith et al. (2014), although the union of RAFs is always a RAF, the union of closed RAFs need not be a closed RAF. Note also that, given a CRS  $\mathcal{Q}$  and food set  $F$ , a closed RAF is completely determined by the subset of molecules that participate in its reactions (Smith et al. 2014, Lemma 1).

Finally, a stronger notion of a RAF is a CAF (*constructively* autocatalytic and F-generated), which requires that the reactions can be ordered so that for every reaction in the sequence, each of its reactants and at least one catalyst are either produced by an earlier reaction in the sequence or are in the food set  $F$  (Mossel & Steel 2005). One can also define a *closed* CAF in an analogous way, namely as a CAF  $\mathcal{R}'$  that satisfies the additional closure property (*cp*) above.

### 2.3. Chemical organization theory

Chemical organization theory (COT) is an alternative formal framework for defining and studying closed and self-maintaining reaction networks. Following ideas by Fontana and Buss (Fontana 1992, Fontana & Buss 1994), COT has been developed to study the hierarchical structure (Speroni di Fenizio et al. 2000) and dynamics (Speroni Di Fenizio & Dittrich 2002) of artificial chemistries that employ catalytic reactions under a general dilution flow. Later, the theory was generalized to arbitrary reaction networks (Dittrich & Speroni di Fenizio 2007).

Note that chemical organizations do not require an explicit notion of catalysis or a food set. However, COT shares with RAF theory a precise axiomatic approach, where the objects of interest are sets of molecule types as opposed to sets of reactions. Nevertheless, as we show below, there is a direct correspondence between the two theories, and the mathematical representation of one theory can be readily converted into one of the other, and vice versa.

In COT, a *reaction network* is a pair  $(X, \mathcal{R})$  where  $X$  is a set of molecule types and  $\mathcal{R}$  is a set of reactions. A *reaction*  $r \in \mathcal{R}$  is a pair of multisets over  $X$ , representing the reactant and product side, respectively. Note that catalysis can be formally included by adding a catalyst  $x$  to both the reactants and the products of  $r$ , e.g.,  $a + b + x \rightarrow c + x$ . Note that this is excluded in RAF theory, where  $A$  and  $B$  are sets and  $A \cap B = \emptyset$ . The reactions  $\mathcal{R}$  can be represented by two matrices  $(L_{i,r})$  and  $(R_{i,r})$ , where  $L_{i,r}$  is the stoichiometric coefficient of molecule type  $i$  on the left hand side of reaction  $r \in \mathcal{R}$ , and  $R_{i,r}$  is the stoichiometric coefficient of  $i$  in reaction  $r \in \mathcal{R}$  on the right hand side.  $\mathbf{S} = (S_{i,r}) = (R_{i,r} - L_{i,r})$  is the stoichiometric matrix. Given  $A \subseteq X$ , we denote  $\mathcal{R}_A = \{r \in \mathcal{R} : \forall i \in X \text{ with } L_{i,r} > 0 \Rightarrow i \in A\}$  as the reactions that can “fire” given molecules  $A$ . Finally,  $\pi(r) = \{i \in X : R_{i,r} > 0\}$  is the set of product types of reaction  $r$ .

Given a reaction network  $(X, \mathcal{R})$  a *chemical organization* is a set of molecule types  $O \subseteq X$  that is:

- (i) *Closed*: none of the reactions that can be applied using only molecules from  $O$  generates any molecules that are not already in  $O$ , i.e., for all  $r \in \mathcal{R}_O : \pi(r) \subseteq O$  (equivalent to  $cl(O, \mathcal{R}) = O$ ); and

- (ii) *Self-maintaining*: all molecules in  $O$  that are consumed by reactions  $\mathcal{R}_O$  within  $O$  can be produced by reactions from  $\mathcal{R}_O$  at a non-negative rate while all reactions  $\mathcal{R}_O$  have a strictly positive flux; formally, there exists a flux vector  $\mathbf{v} = (v_r)$  with  $v_r > 0$  for  $r \in \mathcal{R}_O$  and  $v_r = 0$  otherwise, such that  $\mathbf{S}\mathbf{v} \geq 0$ .

Note that here the closure is always defined relative to the full reaction network  $\mathcal{R}$ . The self-maintaining condition formalizes the notion that molecules that can react will eventually react and that molecules can be regenerated as fast as they are used up (at least for some setting of the reaction rates as specified by  $\mathbf{v}$ ).

A reaction network in COT corresponds to a CRS in RAF theory, and is a flow system, since for each molecule  $x \in F$  in a CRS there is an inflow reaction  $\emptyset \rightarrow x$  in the corresponding reaction network in COT, and for each molecule  $x \in X$  in a CRS there is a decay reaction  $x \rightarrow \emptyset$  in the corresponding reaction network in COT (this will be derived in detail in Section 3.1). This implies the following properties (Dittrich & Speroni di Fenizio 2007): For any set  $A \subseteq X$ , we can generate uniquely a chemical organization  $O = G_O(A) = G_{SM}(G_{CL}(A))$ , where  $G_{CL}(A) = cl(A, \mathcal{R})$  is the smallest closed set containing  $A$  as subset and  $G_{SM}(A)$  the largest self-maintaining subset of  $A$ . All chemical organizations form an algebraic lattice, with union and intersection of two chemical organizations  $O_1, O_2$  computed by  $G_O(O_1 \cup O_2)$  and  $G_O(O_1 \cap O_2)$ , and  $G_O(\emptyset)$  as the smallest and  $G_O(X)$  as the largest chemical organization.

With  $G_O$  we can map any state  $\mathbf{x}$  of a dynamical flow system uniquely to a particular chemical organization  $G_O(\phi(\mathbf{x}))$ , with  $\phi(\mathbf{x})$  the molecules with strictly positive concentrations in  $\mathbf{x}$ . Thus we can track qualitative transitions in the space of chemical organizations. Given a “chemical” differential equation of the form  $\dot{\mathbf{x}}(t) = \mathbf{S}\mathbf{v}(\mathbf{x}(t))$  (with kinetic laws  $\mathbf{v}(\mathbf{x})$ ), basically all long term behaviors tend towards chemical organizations, i.e., given a limit set and any state of this set, the molecular types with positive concentrations form a chemical organization (Peter & Dittrich 2011). COT can be used to unravel the hierarchical structure of models (Centler & Dittrich 2007), follow the chemical evolution of constructive artificial chemistries, predict growth phenotypes (Kaleta et al. 2009), design chemical computers (Matsumaru et al. 2009), and coarse-grain chemical dynamics (Mu et al. 2016).

#### 2.4. The binary polymer model

Kauffman (1993) introduced an abstract model of chemical reaction networks, now known as the *binary polymer model*, which has been used extensively to study autocatalytic sets, e.g., (Bagley & Farmer 1992, Filisetti et al. 2012, Hordijk, Wills & Steel 2014). In this model, molecules are represented by bit strings that can be concatenated into longer ones or cut into smaller ones. Catalysis is then assigned at random.

More specifically, the binary polymer model generates a CRS as follows:

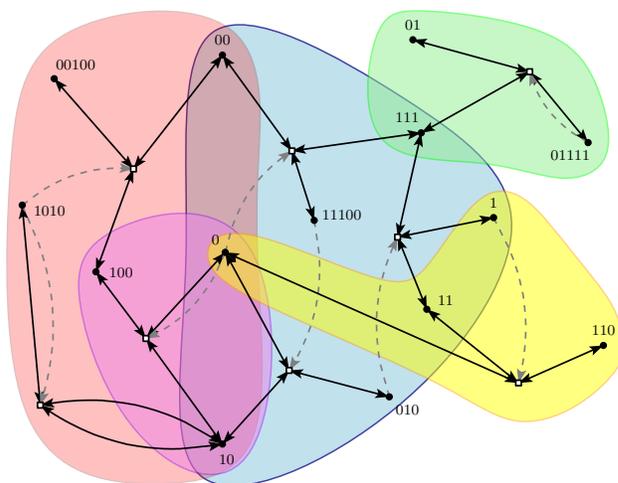
- The molecule set  $X$  consists of all bit strings up to (and including) a maximum length  $n$ :  $X = \{0, 1\}^{\leq n}$ .
- The food set  $F$  consists of all bit strings up to (and including) length  $t$ , where  $t \ll n$ . Here,  $t = 2$  is used:  $F = \{0, 1, 00, 01, 10, 11\}$ .
- The reaction set  $\mathcal{R}$  consists of two types of reactions:
  - (i) *Ligation*: concatenating two bit strings together into a longer one, e.g.  $00 + 111 \rightarrow 00111$ .
  - (ii) *Cleavage*: cutting a bit string into two smaller ones, e.g.  $010101 \rightarrow 01 + 0101$ .

All possible ligation and cleavage reactions between bit strings are included in  $\mathcal{R}$ , as long as none of the reactants or products violates the maximum molecule length  $n$ .

- The catalysis set  $C$  is constructed at random for each instance of the model, as follows. For each bit string  $x \in X$  and reaction  $r \in \mathcal{R}$ , the pair  $(x, r)$  is independently included in  $C$  with probability  $p$ .

This model has two main parameters: the maximum molecule length  $n$  and the probability of catalysis  $p$ . We keep the maximum length of food molecules fixed at  $t = 2$  throughout.

Note that in this model each ligation (forward) reaction has a corresponding cleavage (backward) reaction. In the model, if a molecule catalyzes a reaction, it catalyzes it in both the forward and the backward direction. Because of the random catalysis assignments, each instance of the model gives rise to a different reaction network, which may or may not contain one or more RAF (sub)sets and/or chemical organizations. Figure 2 shows a maxRAF as found by the RAF algorithm in an instance of the binary polymer model with  $n = 5$  and  $p = 0.0045$ . It also shows some of the RAF subsets that exist within the maxRAF as differently colored areas. Note that the top right subset, shaded green, is not in itself a RAF, but it can be added to the blue subRAF to create a larger RAF; such a set is called a ‘co-RAF’.



**Figure 2.** A maxRAF as found by the RAF algorithm in an instance of the binary polymer model with  $n = 5$  and  $p = 0.0045$ . Several RAF subsets within this maxRAF are indicated by colored shapes (the top right reaction, shaded green, is not itself a RAF but a co-RAF; see text).

### 3. Results

We now show how the concepts of closed RAFs and chemical organizations are formally related by deriving some new mathematical results. We then introduce a computational procedure involving the calculation of the chemical organizations within a (max)RAF to find all its closed subRAFs. Finally, we show the results of applying this procedure to instances of the binary polymer model.

### 3.1. Theoretical results

Previously it was shown that an F-generated set is always a chemical organization, but the reverse may not be true (Contreras et al. 2011, Steel et al. 2013). We now show that there is an even closer connection between RAFs and COT, in particular in the context of *closed* RAFs.

First note that, in general, a CRS  $\mathcal{Q}$  can have many closed RAFs but it has at most one closed CAF. To see this, recall that there is always a unique maximal closed RAF for  $\mathcal{Q}$  (namely the maxRAF). However,  $\mathcal{Q}$  may also have many smaller closed (sub)RAF. In fact, it is straightforward to construct an example of a CRS  $\mathcal{Q}$  for which the number of closed RAFs is exponential in the size of  $\mathcal{Q}$ . For example, if we let  $F = \{f_i, i \in [n]\}$ ,  $X = F \cup \{x_i, i \in [n]\}$ ,  $\mathcal{R} = \{r_i := f_i \rightarrow x_i, i \in [n]\}$ , and  $C = \{(x_i, r_i), i \in [n]\}$  (here  $[n] = \{1, \dots, n\}$ ), then the CRS  $\mathcal{Q} = (X, \mathcal{R}, C, F)$  has  $2^n - 1$  closed RAFs, corresponding to the non-empty subsets of  $\mathcal{R}$ .

By contrast, a CRS has at most one closed CAF by virtue of the following result.

**Theorem 1.** *Any CRS  $\mathcal{Q} = (X, \mathcal{R}, C, F)$  has at most one closed CAF.*

*Proof.* Consider the following sequence of subsets of  $\mathcal{R}$ . Let  $\mathcal{R}_1$  be the set of reactions in  $\mathcal{R}$  that have all their reactants and at least one catalyst present in  $F$ . For  $i > 1$ , let  $\mathcal{R}_i$  be the reactions in  $\mathcal{R}_{i-1}$  along with those reactions  $r$  in  $\mathcal{R}$  for which every reactant of  $r$  and at least one catalyst of  $r$  are either present in  $F$  or are produced by a reaction in  $\mathcal{R}_{i-1}$ . Formally, for  $i > 1$ :

$$\mathcal{R}_i = \mathcal{R}_{i-1} \cup \{r \in \mathcal{R} : \rho(r) \subseteq \pi(\mathcal{R}_{i-1}) \cup F \text{ and } \exists x \in \pi(\mathcal{R}_{i-1}) \cup F \text{ with } (x, r) \in C\}.$$

Here  $\rho(r)$  denotes the set of reactants of  $r$ , and  $\pi(\mathcal{R}_{i-1})$  is the set of products of reactions in  $\mathcal{R}_{i-1}$ .

By definition,  $\mathcal{R}_i$  is an increasing sequence of subsets of  $\mathcal{R}$  (i.e.,  $\mathcal{R}_i \subseteq \mathcal{R}_{i+1}$  for all  $i$ ) and since  $\mathcal{R}$  is finite, either  $\mathcal{R}_k = \emptyset$  (which implies that  $\mathcal{R}_1 = \emptyset$ ), in which case  $\mathcal{Q}$  has no CAF, or some  $k$  exists for which  $\mathcal{R}_{k+1} = \mathcal{R}_k \neq \emptyset$  in which case  $\mathcal{R}_k$  is a maximal closed CAF for  $\mathcal{Q}$ . Moreover, we claim that  $\mathcal{R}_k$  is the *only* closed CAF for  $\mathcal{Q}$ , since if  $\mathcal{R}'$  is a closed CAF, then, by the closure property of  $\mathcal{R}'$ , we have that  $\mathcal{R}_1 \subseteq \mathcal{R}'$  and (by induction)  $\mathcal{R}_i \subseteq \mathcal{R}'$  also holds for all  $i > 1$ . Thus  $\mathcal{R}_k \subseteq \mathcal{R}'$ ; since  $\mathcal{R}_k$  is a maximal closed CAF for  $\mathcal{Q}$ , this implies that  $\mathcal{R}_k = \mathcal{R}'$ .  $\square$

We now derive a direct relationship between closed RAFs and chemical organizations. Let  $\mathcal{Q} = (X, \mathcal{R}_0, C, F)$  be a CRS (with food set  $F$ ) and let  $\mathcal{R}$  be any subset of  $\mathcal{R}_0$ . We first state the following definitions:

- Let  $\tilde{\mathcal{R}}$  denote  $\mathcal{R}$  in “expanded” format, where each reaction  $r = A \rightarrow B$  in  $\mathcal{R}$  that is catalyzed by molecules  $x_1, x_2, \dots, x_k \in X$  is replaced by the  $k$  reactions  $A + x_i \rightarrow B + x_i, i = 1, \dots, k$ , where the left- and right-hand sides should be viewed as multisets (i.e.  $x_i$  may already appear as a reactant or product, so another copy of it is added on each side). In other words, each molecule that catalyzes a reaction is explicitly included as both a reactant and a product for that reaction.
- Let

$$\tilde{\mathcal{R}}^\pm = \tilde{\mathcal{R}} \cup \mathcal{R}_F^+ \cup \mathcal{R}_X^-,$$

where

$$\mathcal{R}_F^+ := \{\emptyset \rightarrow f, \text{ for all } f \in F\} \text{ and } \mathcal{R}_X^- := \{x \rightarrow \emptyset, \text{ for all } x \in X\}.$$

In other words,  $\mathcal{R}_F^+$  is a set of “inflow” reactions that introduce food molecules, and  $\mathcal{R}_X^-$  is a set of “outflow” reactions that remove each molecule type.

- Let  $\mathcal{O}(\mathcal{Q})$  denote the set of chemical organizations that exist within the algebraic chemistry  $(X, \tilde{\mathcal{R}}^\pm)$ .
- For  $o \subseteq X$ , let  $\mathcal{R}(o)$  denote the set of reactions in  $\mathcal{R}$  that have all their reactants and at least one catalyst in  $o$ . In other words,  $\mathcal{R}(o)$  is the subset of reactions in  $\mathcal{R}$  that can happen in a catalyzed manner when the molecule set  $o$  is present.
- Finally, for a subset  $\mathcal{R}'$  of  $\mathcal{R}$ , let  $O(\mathcal{R}') := \pi(\mathcal{R}') \cup F$ , where  $\pi(\mathcal{R}')$  is the set of molecules in  $X$  appearing as a product of some reaction in  $\mathcal{R}'$ . In other words,  $O(\mathcal{R}')$  is the union of the food set and all molecules produced by all reactions in  $\mathcal{R}'$ .

The next Lemma will be helpful in our second theoretical result stated below.

**Lemma 1.**

- (i) *If a maxRAF  $\mathcal{R}$  in a CRS  $\mathcal{Q}$  contains no CAF, then there will be a chemical organization  $o \in \mathcal{O}(\mathcal{Q})$  consisting of only  $F$  which does not correspond to any closed RAF;*
- (ii) *The reverse is not necessarily true: If there is a chemical organization  $o \in \mathcal{O}(\mathcal{Q})$  consisting of only  $F$ , there may still be a corresponding closed RAF  $\mathcal{R}(o)$ .*

*Proof.*

*Part (i):* If  $\mathcal{R}$  contains no CAF, it means that there are no reactions that can proceed catalyzed when only molecules in  $F$  are present. In other words, in the corresponding reaction network  $\tilde{\mathcal{R}}^\pm$ , the only reactions that can happen are the inflow and outflow reactions involving just the food molecules. Thus, the set  $o = F$  is a chemical organization for  $\tilde{\mathcal{R}}^\pm$ : it is closed (no new molecules outside of  $F$  are generated) and self-maintaining (we can choose strictly positive flux vector entries  $v_i$  for the inflow reactions and  $v_o$  for the outflow reactions such that  $v_i \geq v_o$ ).

However, since there are no reactions that can proceed catalyzed when only molecules in  $F$  are present, the set  $\mathcal{R}(o)$  is empty, and thus does not correspond to any closed RAF.

*Part (ii):* A counter-example within the binary polymer model with  $t = 2$  is a maxRAF  $\mathcal{R}$  consisting of the following two reactions:



where  $\xrightarrow{00}$  indicates a reaction that is catalyzed by the molecule 00 (and similar for 11). These reactions only involve food molecules, and since  $\mathcal{R}$  is the maxRAF,  $o = F$  is a chemical organization for  $\tilde{\mathcal{R}}^\pm$ . So, there is a chemical organization  $o \in \mathcal{O}(\mathcal{Q})$  consisting of only  $F$ , but there is a CAF and a corresponding closed RAF  $\mathcal{R}(o)$  (both equal to the maxRAF).  $\square$

We can now fully state our second theoretical result.

**Theorem 2.** *Given a CRS  $\mathcal{Q} = (X, \mathcal{R}, C, F)$  and its corresponding reaction network  $RN = (X, \tilde{\mathcal{R}}^\pm)$ , the following hold:*

- (i) *If  $\mathcal{R}'$  is a RAF for  $\mathcal{Q}$  then the corresponding set of molecules  $O(\mathcal{R}')$  is self-maintaining for  $RN$ ;*

- (ii) If, additionally, the RAF  $\mathcal{R}'$  is closed then the corresponding set of molecules  $O(\mathcal{R}')$  is a chemical organization for RN;
- (iii) The reverse of (ii) is not true in general, i.e., in RN there can be a chemical organization that does not correspond to any closed RAF in  $\mathcal{Q}$ .

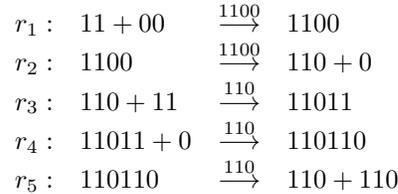
*Proof.*

*Part (i):* The argument that  $O(\mathcal{R}')$  is self-maintaining involves a similar approach to that of Lemma 4.1(i) of Steel et al. (2013), which, in turn, is related to Lemma 2 of Contreras et al. (2011).

Since  $\mathcal{R}'$  is F-generated, Lemma 3.1 of Steel et al. (2013) shows that  $\mathcal{R}'$  has a linear ordering  $r_1, r_2, \dots, r_k$  so that each reactant of  $r_i$  is either an element of  $F$  or is a product of some reaction  $r_j$  where  $1 \leq j < i$ . Now we can order the reactions in  $\tilde{\mathcal{R}}^\pm$  so that the reactions  $\mathcal{R}_F^\pm$  come first, followed by the expansions of  $r_1$ , then the expansions of  $r_2$  and so on, up to the expansions of  $r_k$ , finally followed by the reactions  $R_{\bar{X}}$ . Next, consider the corresponding stoichiometric matrix  $\mathbf{S}$ . The first non-zero element in each row of  $\mathbf{S}$  is  $+1$ , by the property of the linear ordering described for  $r_1, \dots, r_k$ , and noting that any catalyst  $x$  contributes  $+1 + (-1) = 0$  to any entry of  $\mathbf{S}$ . Finally, for any real matrix with this last property, there is a strictly positive column vector  $\mathbf{v}$  for which  $\mathbf{S}\mathbf{v} > \mathbf{0}$ , since if  $\mathbf{S}$  has  $c$  columns and if the largest absolute value of any negative entry of  $\mathbf{S}$  is  $b$ , then we can take  $\mathbf{v}$  to be the strictly positive vector that has its  $i^{\text{th}}$  coordinate given by  $v_{c-i} = (b+1)^i$  for  $i = 0, \dots, c-1$ .

*Part (ii):* Since the RAF  $\mathcal{R}'$  is closed,  $O(\mathcal{R}')$  also satisfies the closure property required for a chemical organization. Part (ii) now follows from Part (i).

*Part (iii):* We show this by example. Consider the following set of reactions  $\mathcal{R}$ , which could be a subset of an instance of the binary polymer model with  $n = 6$  and  $t = 2$ :



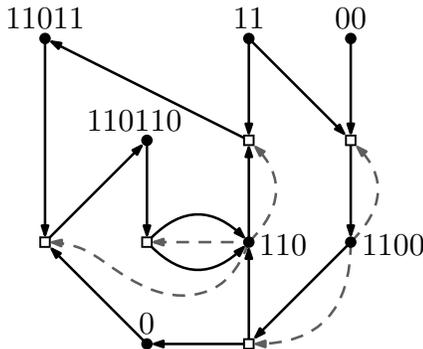
If we assume that these are the only catalyzed reactions, then  $\mathcal{R}$  is the maxRAF of this model instance. Figure 3 shows this maxRAF in a graphical way.

It turns out that  $(X, \tilde{\mathcal{R}}^\pm)$  contains three chemical organizations:

$$\mathcal{O}(\mathcal{Q}) = \{F, F \cup \{110, 11011, 110110\}, F \cup \{110, 1100, 11011, 110110\}\}.$$

First, consider the chemical organization  $o = F$ . Note that the maxRAF  $\mathcal{R}$  contains no CAF, since none of the catalysts are available yet when only molecules in  $F$  are present. So, according to Lemma 1(i),  $F$  is indeed a chemical organization for  $\tilde{\mathcal{R}}^\pm$ . However, according to the same Lemma, this chemical organization does not correspond to a closed RAF. This illustrates one specific case in which a chemical organization can fail to correspond to a closed RAF.

Next consider the second chemical organization listed above:  $o = F \cup \{110, 11011, 110110\}$ . It can be easily checked by hand that  $o$  is closed for  $\tilde{\mathcal{R}}^\pm$ . That  $o$  is also self-maintaining for  $\tilde{\mathcal{R}}^\pm$  can be seen by selecting flux vector entries  $v_3, v_4, v_5$  for the corresponding (expanded) reactions  $r_3, r_4, r_5$  as follows:  $2v_5 > v_3 > v_4 > v_5$ . The remaining entries of  $\mathbf{v}$  can then be chosen to be strictly positive so that  $\mathbf{S}\mathbf{v} \geq \mathbf{0}$ .



**Figure 3.** A maxRAF that could exist in an instance of the binary polymer model with  $n = 6$  and  $t = 2$  ( $F = \{0, 1, 00, 01, 10, 11\}$ ).

However,  $\mathcal{R}(o) = \{r_3, r_4, r_5\}$  is not F-generated (each of the three reactions has at least one non-food molecule as its reactants), and is therefore not a closed RAF of  $\mathcal{R}$ . This illustrates a second case in which a chemical organization can fail to correspond to a closed RAF.

The third chemical organization listed above,  $o = F \cup \{110, 1100, 11011, 110110\}$ , gives rise to an  $\mathcal{R}(o)$  that is equal to the maxRAF, and is thus a chemical organization that does correspond to a closed RAF.  $\square$

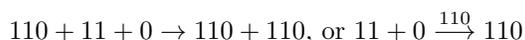
Part (ii) of Theorem 2 implies that the closed RAFs in a given (max)RAF  $\mathcal{R}$  correspond to chemical organizations. However, as Part (iii) shows, the converse need not always hold, even in the binary polymer model. Essentially, this is because the self-maintaining property of a chemical organization is an algebraic condition based on certain linear inequalities holding for a freely varying positive (flux) vector, and thus has a reasonable degree of freedom. On the other hand, a RAF  $\mathcal{R}$  has to be F-generated, which translates into a more strict combinatorial condition regarding the ability to order the reactions in  $\mathcal{R}$  in a certain way, i.e., for each reaction in the sequence each of its reactants (but not necessarily any catalyst) is either a product of an earlier reaction or is present in the food set.

The reason we need to include the inflow and outflow reactions in  $\tilde{\mathcal{R}}^\pm$  to make the correspondence between closed RAFs and chemical organizations work is because food molecules are implicitly assumed to be available from the environment in the RAF framework, without necessarily needing to be produced by the reaction network itself. The inflow reactions make this assumption explicit in the chemical organization context. Furthermore, if the outflow reactions were not included, there may exist chemical organizations containing molecules that are neither produced nor used up in any reaction. In a RAF this is not allowed, as all molecules need to be produced, through a series of reactions from the RAF set itself, from the food molecules. Therefore, inclusion of the outflow reactions forces all molecules to be produced at a strictly positive rate, rather than merely a non-negative rate (as is normally assumed for a chemical organization).

However, as Part (iii) of Theorem 2 shows, there may be chemical organizations for  $\tilde{\mathcal{R}}^\pm$  that do not correspond to closed RAFs in  $\mathcal{R}$ . First, any RAF  $\mathcal{R}$  that does not contain a CAF will have a chemical organization  $o = F$  consisting of only the

food molecules, which would correspond to an “empty” closed RAF  $\mathcal{R}(o)$ . However, as Lemma 1(ii) shows, we cannot simply discard any chemical organization that consists of  $F$  only. But since there can be at most one such chemical organization in any reaction network, it is easy to check if such a chemical organization actually corresponds to a closed RAF or not.

Second, note that in the counter-example given in the proof of Theorem 2, the three reactions  $(r_3, r_4, r_5)$  form what is known as an *autocatalytic cycle*. Reaction  $r_3$  starts with one 110 molecule plus a food molecule (11) to create 11011. Reaction  $r_4$  then takes this newly produced molecule plus another food molecule (0) to create 110110. Finally, reaction  $r_5$  splits this longer molecule into two copies of 110. In other words, this sequence of three reactions starts with one copy of 110, adds two food molecules, and then produces two copies of 110, which can be written in short-hand as:



In other words, the cycle of three reactions “reduces” to a single autocatalytic reaction. Hence the name *autocatalytic cycle*.

However, this is different from an *autocatalytic set*, even though each of the three reactions in the autocatalytic cycle is catalyzed by a molecule from the set itself (in this case, 110 for each reaction). In fact, the reaction set  $\{r_3, r_4, r_5\}$  is not F-generated and is thus not an autocatalytic set, but it does correspond to a chemical organization, as shown in the proof above. In conclusion, an autocatalytic *set* may contain a subset that forms an autocatalytic *cycle* that corresponds to a chemical organization, but not to a (closed) RAF.

There may even be other specific situations that give rise to a chemical organization but not a closed RAF. However, we argue that these cases are likely to be rare. Indeed, as the computational results below show, they do not seem to occur often (or at all) in random instances of the binary polymer model.

Assuming these exceptions are indeed rare, Theorem 2 suggests that calculating the chemical organizations in  $\tilde{\mathcal{R}}^\pm$  provides a reliable method for finding the closed RAFs in a (max)RAF  $\mathcal{R}$ . We therefore propose the following computational procedure:

**Procedure findClosedRAFs:**

Given a CRS  $\mathcal{Q} = (X, \mathcal{R}_0, C, F)$  with food set  $F \subset X$ , let  $\mathcal{R}$  be a (max)RAF for  $\mathcal{Q}$ . Let  $\pi(\mathcal{R})$  be the set of products of reactions in  $\mathcal{R}$ .

- (i) Convert  $\mathcal{R}$  into the expanded reaction set  $\tilde{\mathcal{R}}$  by replacing each  $r \in \mathcal{R}$  and  $x \in \pi(\mathcal{R}) \cup F$  for which  $(x, r) \in C$  by a reaction in which  $x$  is added to the left and right side. In other words, if a reaction  $r$  has more than one catalyst in  $\pi(\mathcal{R}) \cup F$ , it will be represented multiple times in  $\tilde{\mathcal{R}}$ , once for each of its catalysts.
- (ii) Convert  $\tilde{\mathcal{R}}$  into  $\tilde{\mathcal{R}}^\pm$  by adding the reactions in  $\mathcal{R}_F^+$  and  $\mathcal{R}_X^-$ .
- (iii) Calculate the chemical organizations  $\mathcal{O}(\mathcal{Q})$  in  $\tilde{\mathcal{R}}^\pm$ .
- (iv) For each chemical organization  $o \in \mathcal{O}(\mathcal{Q})$  reconstruct the corresponding set  $\mathcal{R}(o) = \{r \in \mathcal{R} | (\text{all reactants and at least one catalyst of } r) \in o\}$ .
- (v) For each such  $\mathcal{R}(o)$  check that it is a RAF set.

This procedure provides an exact algorithm for finding all closed RAFs, by way of calculating all chemical organizations within a (max)RAF  $\mathcal{R}$ . Given that there was

no known algorithm yet to enumerate all closed RAFs, or even to determine whether a maxRAF contains another (smaller) closed subRAF, the above procedure can be very helpful in providing more insight into the (so far) unresolved question of how many closed RAFs can be expected to exist within arbitrary reaction networks. As already argued earlier, it is exactly these closed RAFs (corresponding to chemical organizations) that are of interest from a dynamical point of view, and they play an important role in the potential evolvability of autocatalytic sets (see the discussion below).

Note, however, that calculating all chemical organizations (Step 3) is, in general, an intractable problem, as the number of chemical organizations (and also closed RAFs) can grow exponentially with the size of the reaction network. This algorithm is therefore limited in practice to relatively small reaction networks (a few hundred reactions, roughly), but it is guaranteed to find *all* closed RAFs.

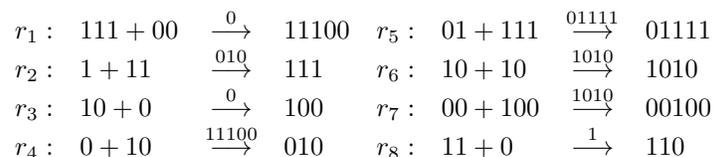
The last step, checking whether each chemical organization corresponds to a RAF set (and thus a closed RAF set), is necessary due to Theorem 2(iii). However, the standard RAF algorithm can be used for this and is therefore an efficient operation, especially compared to the computationally costly step of calculating all chemical organizations.

An illustrative example of applying this computational procedure to instances of the binary polymer model is provided next.

### 3.2. Computational results

Recall Figure 2, which shows an example of a maxRAF  $\mathcal{R}$  as found by the RAF algorithm in an instance of the binary polymer model with  $n = 5$  and  $p = 0.0045$ . This maxRAF consists of eight reactions. Some of the subRAFs within the maxRAF are indicated by colored shapes. When only food molecules (bit strings of lengths one and two) are present, only the purple and yellow subsets can immediately come into existence, as they use only food molecules as their reactants and catalysts. In other words, the union of the purple and yellow subsets forms the CAF for this network. However, the red and blue subsets each need at least one spontaneous (uncatalyzed) reaction to happen before they can come into existence. Finally, once the blue subset exists, it can be extended with the green subset, which, by itself, is not a proper RAF but can form a larger RAF together with the blue subset. The green extension also requires a spontaneous reaction to come into existence dynamically. A short movie showing these various subsets coming into existence at different times as a result of (stochastic) spontaneous reaction events is available online at Hordijk (2016).

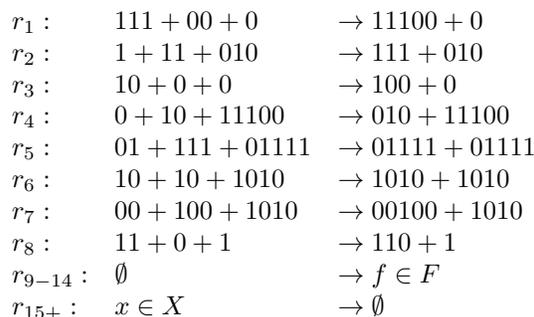
The eight reactions making up this particular maxRAF  $\mathcal{R}$  are as follows:



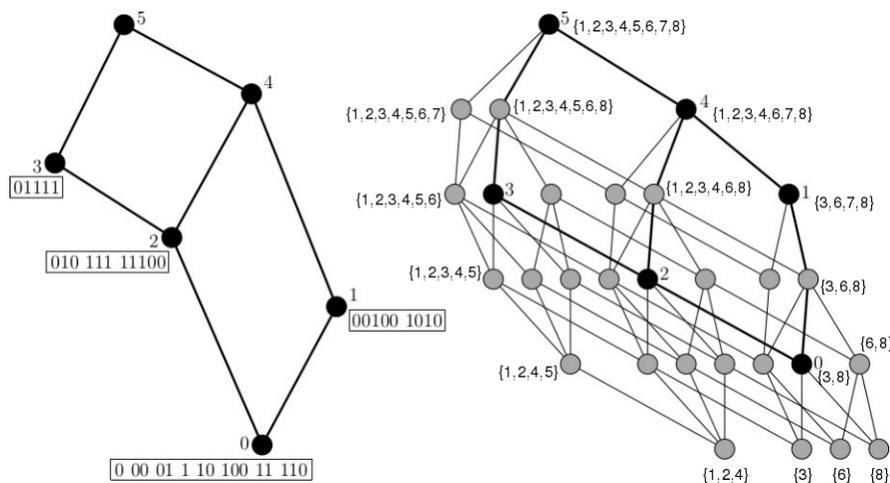
Note that for the calculation of the chemical organizations below, we only consider the ligation (forward) reactions here, to keep the example tractable.

This maxRAF actually contains a total of 29 subRAFs, which form a partially ordered set (poset), shown as a Hasse diagram in Figure 4 (right). The question is now which ones of these 29 subsets are also closed RAFs. To answer this, we calculate the chemical organizations within the maxRAF.

First, the maxRAF  $\mathcal{R}$  is converted to the corresponding reaction set  $\tilde{\mathcal{R}}^\pm$  as stated in the `findClosedRAFs` procedure:



Next, we calculate the chemical organizations within  $\tilde{\mathcal{R}}^\pm$ . It turns out this reaction set contains six chemical organizations, also forming a poset as shown in Figure 4 (left). The molecules contained in each chemical organization are shown next to the corresponding nodes in the Hasse diagram, with only the “new” molecules listed (i.e., those molecules that are not already in any of a chemical organization’s subsets). For example, the chemical organization labeled “3” contains molecule 01111 in addition to the molecules already listed for the chemical organizations labeled “0” and “2” (its subsets).



**Figure 4.** The maxRAF of Figure 2 contains six chemical organizations (left). These all correspond to closed RAFs, which are situated in the poset of 29 subRAFs (right). For some of the 29 subRAFs, it is indicated which reactions they consist of. For example, the label  $\{3, 6, 7, 8\}$  means the subRAF consists of the reaction set  $\{r_3, r_6, r_7, r_8\}$ .

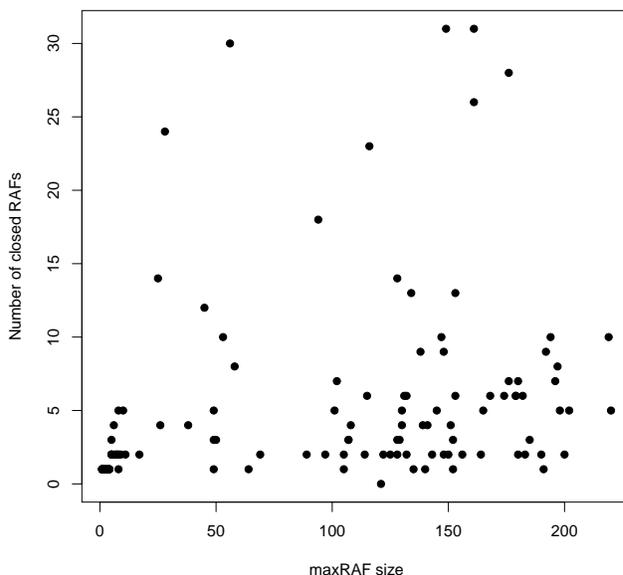
Finally, we check whether the reconstructed reaction sets  $\mathcal{R}(o)$  of these six chemical organizations are indeed RAF sets, which they all are. Therefore, all six chemical organizations  $o$  correspond to closed RAFs  $\mathcal{R}(o)$ . Figure 4 (right) shows how these six closed RAFs are situated within the poset of subRAFs. These closed RAFs correspond to the following reaction subsets and combinations of colored shapes in Figure 2:

0: $\{r_3, r_8\}$	$\{P, Y\}$
1: $\{r_3, r_6, r_7, r_8\}$	$\{P, Y, R\}$
2: $\{r_1, r_2, r_3, r_4, r_8\}$	$\{P, Y, B\}$
3: $\{r_1, r_2, r_3, r_4, r_5, r_8\}$	$\{P, Y, B, G\}$
4: $\{r_1, r_2, r_3, r_4, r_6, r_7, r_8\}$	$\{P, Y, R, B\}$
5: $\{r_1, r_2, r_3, r_4, r_5, r_6, r_7, r_8\}$	$\{P, Y, R, B, G\}$

where P=purple, Y=yellow, R=red, B=blue, and G=green.

This example explicitly shows how the procedure works. To get more insight into the expected number of closed RAFs within arbitrary maxRAFs, Figure 5 shows the number of closed RAFs found within 120 different maxRAFs of sizes up to around 200 reactions in the binary polymer model, using  $n = 6$  and  $p = 0.0050$ , and only forward (ligation) reactions. As this graph shows, these maxRAFs contain up to around 30 closed RAFs.

For maxRAFs of around 200 reactions, calculating the poset of subRAFs is already highly intractable. Even maxRAFs of around 20 reactions can have more than 1000 subRAFs. But calculating the chemical organizations (of which there are fewer) is still doable, with the largest instances from the current sample requiring up to one hour of computing time on a 1.50-GHz CPU. Unfortunately, though, for  $n = 7$  in the binary polymer model (with maxRAF sizes averaging close to 400 reactions), calculating the chemical organizations becomes too time-consuming as well.



**Figure 5.** The number of closed RAFs vs. the size of the maxRAF.

Finally, to confirm that the counter-examples as shown in the proof of Theorem 2 indeed seem to be rare, there are only 27 instances, out of the 120, where the number of closed RAFs is one less than the number of chemical organizations. These are actually all instances where there is no CAF in the maxRAF, which means that

there is one chemical organization consisting of just the food set  $F$  which does not correspond to a closed RAF (which is easy to check). In other words, cases like the autocatalytic cycle as illustrated in the proof of Theorem 2 have not occurred in the current sample. Thus, the proposed computational procedure indeed seems to provide a reliable method for finding all closed RAFs within a (max)RAF.

#### 4. Discussion

We have shown a direct and formal relationship between closed RAF sets and chemical organizations, refining and extending earlier established relationships between COT and RAFs (Contreras et al. 2011, Steel et al. 2013, Smith et al. 2014). This relationship is considered to be relevant in the context of the origin of life in the following way. As was shown elsewhere, the existence of subRAFs is one of the main requirements for autocatalytic sets to be evolvable (Vasas et al. 2012, Hordijk & Steel 2014). However, not every subRAF will necessarily be realized or persist in a dynamical sense. In particular, many subRAFs will only exist as a “transient” state, quickly becoming part of a larger subset where all reactions that currently *can* happen indeed *will* happen. Figure 2 showed an example where there are only six closed subsets out of a larger set of 29 subRAFs. Only these six closed subRAFs are dynamically “stable”, while the other subRAFs are “transient”. Thus, these closed RAFs are the equivalent of the viable cores (plus their periphery) of Vasas et al. (2012), possibly serving as units of variation and selection in early evolution via the possibility of having different combinations of closed RAFs existing inside different compartments (or “protocells”).

However, an important remaining question was how many *closed* subRAFs can be expected to exist within arbitrary chemical reaction networks. In other words, how much potential variation is there for such early evolution to work with? Here we have introduced and applied an exact computational procedure, based on calculating the chemical organizations within a maxRAF, to answer this question, thus confirming a suggestion made by Contreras et al. (2011). Even in a simple polymer model where molecules are represented by bit strings up to length six, we can already expect to see up to around 30 closed subRAFs. We expect this number to grow with increasing network sizes, although unfortunately the introduced computational procedure quickly becomes intractable for such larger sizes.

In conclusion, it seems that we can expect a large enough diversity of closed subRAFs to exist within arbitrary reaction networks (or, more precisely, within arbitrary maxRAFs) to satisfy a basic but important requirement for autocatalytic sets to be evolvable. In earlier work, we had already shown that autocatalytic sets (or maxRAFs) have a high probability of existing, also for moderate (and chemically realistic) levels of catalysis (Hordijk & Steel 2004, Hordijk & Steel 2017). Thus, together these results could have significant implications for a possible origin and further evolution of life as a collection of mutually catalytic molecules, i.e., a cooperative reaction network, rather than one or a few individually self-replicating RNA sequences.

Note that one of the main reasons the introduced procedure works is that it explicitly looks for chemical organizations within the maxRAF  $\mathcal{R}$  (or, more precisely, its expanded version  $\tilde{\mathcal{R}}^\pm$ ), rather than in the full reaction network  $\mathcal{R}_0$  in which the maxRAF  $\mathcal{R}$  exists. When asking whether a set of molecules  $X$  is a chemical organization or not, it is important to specify the reaction set  $\mathcal{R}$  that is implied. For example, in Contreras et al. (2011) it is stated that “all RAF sets are chemical

organizations”. However, this is only true relative to the reaction set  $\mathcal{R}'$  that consists of only the reactions in the given RAF (sub)set. It is (in general) not true relative to the maxRAF  $\mathcal{R}$  that a RAF subset  $\mathcal{R}'$  might be part of, or relative to the entire reaction network  $\mathcal{R}_0$  in which the RAF set exists.

Here we have made the relationship between chemical organizations and RAF sets more precise, by showing that all *closed* RAF (sub)sets are chemical organizations (relative to the maxRAF). Furthermore, we have answered an open question from Contreras et al. (2011): “whether all chemical organizations are F-generated is an ambiguous matter”. We have shown here, by counter-example, that not all chemical organizations are necessarily closed RAF sets, as they may not be F-generated. However, the exceptions can either be easily detected (such as chemical organizations that consist of only food molecules, when there is no CAF) or they seem to be rare (such as autocatalytic cycles, which may not be F-generated).

This precise correspondence between closed RAFs and chemical organizations also allows a direct application of COT to RAF sets and their dynamics. Thus we can track qualitative transitions in the space of subRAFs, unravel the hierarchical structure and follow the chemical evolution of closed RAFs, predict growth phenotypes, and coarse-grain the overall RAF dynamics.

Some computational questions still remain. For example, is it possible to obtain an estimate of the number of closed subRAFs for larger reaction networks or, alternatively, obtain a random but representative sample of closed subRAFs? Another interesting question is whether or not there is a polynomial-time algorithm for the following basic decision problem: Given a CRS  $\mathcal{Q} = (X, \mathcal{R}_0, C, F)$  having a maxRAF  $\mathcal{R}$ , does  $\mathcal{R}$  contain a closed RAF as a strict subset? We hope to address some of these questions in future work.

## Acknowledgments

WH thanks the KLI Klosterneuburg ([www.kli.ac.at](http://www.kli.ac.at)) for financial support in the form of a fellowship. The computational results presented in this paper were obtained using a Raspberry Pi ([www.raspberrypi.org](http://www.raspberrypi.org)).

## References

- Bagley R J & Farmer J D 1992 in C. G Langton, C Taylor, J. D Farmer & S Rasmussen, eds, ‘Artificial Life II’ Addison-Wesley Redwood City, CA pp. 93–140.
- Benner S A, Kim H J & Yang Z 2012 *Cold Spring Harbour Perspectives in Biology* **4**, a003541.
- Centler F & Dittrich P 2007 *Planetary Space Science* **55**(4), 413–428.
- Contreras D A, Pereira U, Hernández V, Reynaert B & Letelier J C 2011 in ‘Advances in Artificial Life, ECAL 2011: Proceedings of the Eleventh European Conference on the Synthesis and Simulation of Living Systems’ MIT Press pp. 176–183.
- Dittrich P & Speroni di Fenizio P 2007 *Bulletin of Mathematical Biology* **69**(4), 1199–1231.
- Farmer J D, Kauffman S A & Packard N H 1986 *Physica D* **22**, 50–67.
- Filisetti A, Graudenzi A, Serra R, Villani M, Füchslin R M, Packard N, Kauffman S A & Poli I 2012 *Theory in biosciences* **131**(2), 85–93.
- Fontana W 1992 in C. G Langton, C Taylor, J. D Farmer & S Rasmussen, eds, ‘Artificial Life II’ Addison-Wesley Redwood City, CA pp. 159–210.
- Fontana W & Buss L W 1994 *Bulletin of Mathematical Biology* **56**, 1–64.
- Gilbert W 1986 *Nature* **319**, 618.
- Hordijk W 2013 *BioScience* **63**(11), 877–881.
- Hordijk W 2016 ‘Autocatalytic sets arising’ <http://www.WorldWideWanderings.net/blog/?p=561>.
- Hordijk W, Hasenclever L, Gao J, Mincheva D & Hein J 2014 *Natural Computing* **13**(3), 287–296.
- Hordijk W, Hein J & Steel M 2010 *Entropy* **12**(7), 1733–1742.

- Hordijk W, Kauffman S A & Steel M 2011 *International Journal of Molecular Sciences* **12**(5), 3085–3101.
- Hordijk W, Smith J I & Steel M 2015 *Algorithms for Molecular Biology* **10**, 15.
- Hordijk W & Steel M 2004 *Journal of Theoretical Biology* **227**(4), 451–461.
- Hordijk W & Steel M 2013 *Journal of Systems Chemistry* **4**, 3.
- Hordijk W & Steel M 2014 *Origins of Life and Evolution of Biospheres* **44**(2), 111–124.
- Hordijk W & Steel M 2017 *BioSystems* **152**, 1–10.
- Hordijk W, Steel M & Kauffman S 2012 *Acta Biotheoretica* **60**(4), 379–392.
- Hordijk W, Wills P R & Steel M 2014 *Bulletin of Mathematical Biology* **76**(1), 201–224.
- Hud N V, Cafferty B J, Krishnamurthy R & Williams L D 2013 *Chemistry & Biology* **20**(4), 466–474.
- Jain S & Krishna S 1998 *Physical Review Letters* **81**, 5684.
- Joyce G F 2002 *Nature* **418**, 214–221.
- Kaleta C, Richter S & Dittrich P 2009 *Bioinformatics* **25**(15), 1915–1922.
- Kauffman S A 1971 *Journal of Cybernetics* **1**(1), 71–96.
- Kauffman S A 1986 *Journal of Theoretical Biology* **119**, 1–24.
- Kauffman S A 1993 *The Origins of Order* Oxford University Press.
- Matsumaru N, Hinze T & Dittrich P 2009 *International Journal of Nanotechnology and Molecular Computation* **1**(4), 1–19.
- Mossel E & Steel M 2005 *Journal of Theoretical Biology* **233**(3), 327–336.
- Mu C, Dittrich P, Parker D & Rowe J E 2016 in E Bartocci, P Lio & N Paoletti, eds, ‘Computational Methods in Systems Biology’ Vol. 9859 of *LNBI* Springer pp. 232–251.
- Nghe P, Hordijk W, Kauffman S A, Walker S I, Schmidt F J, Kemble H, Yeates J A M & Lehman N 2015 *Molecular BioSystems* **11**, 3206–3217.
- Patel B H, Percivalle C, Ritson D J, Duffy C D & Sutherland J D 2015 *Nature Chemistry* **7**, 301–307.
- Peter S & Dittrich P 2011 *Advances in Complex Systems* **14**(01), 77–96.
- Powner M W, Gerland B & Sutherland J D 2009 *Nature* **459**, 239–242.
- Smith J, Steel M & Hordijk W 2014 *Journal of Systems Chemistry* **5**, 2.
- Sousa F L, Hordijk W, Steel M & Martin W F 2015 *Journal of Systems Chemistry* **6**, 4.
- Speroni Di Fenizio P & Dittrich P 2002 *The Journal of Three Dimensional Images* **16**(4), 160–163.
- Speroni di Fenizio P, Dittrich P, Ziegler J & Banzhaf W 2000 in ‘German Workshop on Artificial Life (GWAL 2000), in print’ Bayreuth, 5.-7. April, 2000.
- Stadler P F, Fontana W & Miller J H 1993 *Physica D* **63**, 378–392.
- Steel M 2000 *Applied Mathematics Letters* **3**, 91–95.
- Steel M, Hordijk W & Smith J 2013 *Journal of Theoretical Biology* **332**, 96–107.
- Szostak J W 2012 *Journal of Systems Chemistry* **3**, 2.
- van Santen R A & Neurock M 2006 *Molecular Heterogeneous Catalysis* Wiley-VCH.
- Vasas V, Fernando C, Santos M, Kauffman S & Sathmary E 2012 *Biology Direct* **7**, 1.