

Finding Chemical Organisations in Matter-Conserving AChems

Jonathan Young¹ and Simon Colton^{1,2}

¹EECS, Queen Mary University of London, London, UK

²SensiLab, Faculty of IT, Monash University, Melbourne, Australia

j.a.young@qmul.ac.uk

Abstract

Chemical Organisation Theory (COT) provides a way of understanding the evolution of collectively self-maintaining sets of molecular species. Adding conservation of matter to an Artificial Chemistry (AChem) can increase evolutionary activity, so it may be useful to understand the evolution of organisations under these conditions. We show how in a reaction network generated by a matter-conserving chemistry, every edge within an organisation must be a part of a cycle in the organisation's bipartite representation. A consequence of this fact is used to alter an existing algorithm to more efficiently discover the complete set of organisations. The altered algorithm is shown to be faster than the original when tested on reaction networks generated by the Spiky-RBN AChem. Overall, this paper contributes useful tools for analysing chemical evolution in matter-conserving chemistries.

Introduction

Artificial Chemistries (AChems) are commonly abstract computational models that are employed to better understand generic principles of life and its origin which are not contingent upon terrestrial conditions and history (Dittrich et al., 2001; Langton, 1989). Evolving self-maintaining or self-replicating networks of molecular species may have preceded the development of life (Kauffman, 1986; Hordijk et al., 2010). Chemical Organisation Theory (COT) can be used to identify subnetworks of a reaction network that do not create any molecular species outside the subnetwork and that are able to maintain their existence over time (Dittrich and Speroni di Fenizio, 2007; Centler et al., 2008). These subnetworks are called *organisations*. The concept of an organisation was originally introduced as part of a minimal theory of biological organisation (Fontana and Buss, 1994). Hence, COT is an appropriate means of analysing the evolution of chemical organisations in models of prebiotic evolution (Matsumaru et al., 2006b). Algorithms for finding the complete set of organisations in a reaction network can be expensive (Centler et al., 2008), but there are more efficient algorithms for reaction networks that are generated under specific conditions (Milreu et al., 2010; Speroni di Fenizio, 2015). For example, the set of organisations generated by

an AChem forms a lattice if every molecular species has a nonzero probability of disappearing from the reaction vessel (Speroni di Fenizio, 2015). An algorithm that takes advantage of this property can find the complete set of organisations much faster than a brute force approach (Speroni di Fenizio, 2015).

In some AChems, the total amount of matter is conserved over time. In this paper, *matter* refers to an indivisible unit which everything within the system is made from. We describe matter as being *conserved* in a chemistry if the total number of these units within the system cannot change over time and every reaction consumes the same amount of units that it produces. AChems have been shown to behave differently when matter-conservation is added. For example, conserving the total matter in the Stringmol (Hickinbotham et al., 2010) AChem can increase evolutionary activity (Hickinbotham and Stepney, 2015). The Spiky-RBN (Krastev et al., 2016) AChem is also found to behave differently in matter-conserving reaction vessels (Krastev et al., 2017). Milreu et al. (2010) hypothesise that more efficient algorithms for finding organisations may exist for matter-conserving chemistries. But to the best of our knowledge, no such algorithm has been proposed. Such an algorithm would make it easier to characterise the evolution of organisations under such conditions. This could help facilitate research, examples of which may include comparing the evolutionary behaviours of matter-conserving AChems, or comparing the behaviour of an AChem with and without matter-conservation.

Here we prove that in a reaction network generated by a matter-conserving chemistry (artificial or natural), every edge within an organisation must be a part of a cycle in the organisation's bipartite representation. A consequence of this fact is used to alter the constructive algorithm proposed by Centler et al. (2008) such that the complete set of organisations can be discovered more efficiently. We demonstrate that this altered algorithm is approximately nine times faster than the unaltered version when tested on reaction networks generated by the Spiky-RBN (Krastev et al., 2016) AChem.

In this paper, we discuss COT and introduce its core defin-

* March 2023: See erratum for this paper here: https://doi.org/10.1162/isal_x_00560

* June 2023: See erratum for this paper here: https://doi.org/10.1162/isal_x_00561

itions, then provide all relevant theorems before providing the altered version of the constructive algorithm, and then present experimental results and discuss conclusions.

Chemical Organisation Theory

Minimal biological systems have been characterized as chemical systems that are capable of self-maintenance and reproduction. Fontana and Buss (1994) introduced the concept of a biological organisation as an operationally closed and self-maintaining chemical system. This work was further developed into COT, the goal of which is to identify subnetworks of a reaction network that do not create any molecular species outside the subnetwork and that are able to maintain themselves over time (Dittrich and Speroni di Fenizio, 2007; Centler et al., 2008). COT can be applied to any system in which entities can interact and form new entities (Dittrich and Speroni di Fenizio, 2007). In particular, It has been used to study models of the martian atmosphere (Centler and Dittrich, 2007), HIV infection (Matsumaru et al., 2006a), ecological systems (Veloz, 2020) and political systems (Dittrich and Winter, 2008). It has also been used to study the evolution of chemical organisations in an AChem model of prebiotic evolution (Matsumaru et al., 2006b) since self-maintenance may be a prerequisite for minimal biological systems (Fontana and Buss, 1994). COT has been shown to overlap with theories relating to autocatalytic sets of species. For example, a reflexively autocatalytic set that is generated by a food set and cannot generate species outside of the set is also an organisation (Hordijk et al., 2018). Autocatalytic sets that are organisations also contain catalytic loops or cycles (Contreras et al., 2011).

We will formally define what a chemical organisation is after some definitions are introduced for reaction networks generally. A reaction network $\langle \mathcal{M}, \mathcal{R} \rangle$ is defined as a set of reactions \mathcal{R} that occur among molecular species \mathcal{M} . Each reaction within \mathcal{R} is an ordered pair consisting of a set of *reactants* and a set of *products*. Therefore $\mathcal{R} \subseteq \mathcal{P}_M(\mathcal{M}) \times \mathcal{P}_M(\mathcal{M})$, where $\mathcal{P}_M(\mathcal{M})$ are all multisets that can be created from elements in \mathcal{M} and $X \times Y$ consists of all ordered pairs (x, y) that exist with $x \in X$ and $y \in Y$. Reactions are denoted $(A \rightarrow B) \in \mathcal{R}$ hereafter, where $A \in \mathcal{P}_M(\mathcal{M})$ is the multiset of reactant species and $B \in \mathcal{P}_M(\mathcal{M})$ is the multiset of product species. They are multisets because there can be multiple instances of the same element in a multiset, just like there can be multiple instances of the same species consumed or produced in a reaction. We also define a $|\mathcal{M}| \times |\mathcal{R}|$ stoichiometric matrix $S = (s_{i,j})$ where $|s_{i,j}|$ is the number of instances of species $i \in \mathcal{M}$ that are produced or consumed in reaction $j \in \mathcal{R}$. If $s_{i,j} > 0$, then species i is produced in the reaction, but if $s_{i,j} < 0$, it is consumed (Dittrich and Speroni di Fenizio, 2007; Centler et al., 2008). Given a subset of species $\mathcal{M}' \subseteq \mathcal{M}$, we say the reaction network *implied* by \mathcal{M}' is $R(\mathcal{M}') = \{(A \rightarrow B) \in \mathcal{R} \mid A \in \mathcal{P}_M(\mathcal{M}')\}$.

A chemical organisation is a set of species $\mathcal{O} \subseteq \mathcal{M}$ that is *closed* and *self-maintaining* as defined below. A set of species is considered closed if there are no reactions using species within the set that can react to form any species outside of the set (Dittrich and Speroni di Fenizio, 2007; Centler et al., 2008).

Definition 1. $\mathcal{O} \subseteq \mathcal{M}$ is *closed* if for all reactions $(A \rightarrow B) \in R(\mathcal{O})$ then $B \in \mathcal{P}_M(\mathcal{O})$.

A set of species is self-maintaining because all species within the set can sufficiently be produced such that none of them decays (Dittrich and Speroni di Fenizio, 2007; Centler et al., 2008).

Definition 2. A set of species $\mathcal{O} \subseteq \mathcal{M}$ is *self-maintaining* if there exists a vector of reaction rates (called a flux vector) $r = (r_1, \dots, r_n) \in \mathbb{R}_{\geq 0}^n$ such that the following conditions hold: (1) For every reaction $j \in R(\mathcal{O})$, its corresponding flux is $r_j > 0$. (2) For every reaction $j \in \mathcal{R} \setminus R(\mathcal{O})$, its corresponding flux is $r_j = 0$. (3) For every species $i \in \mathcal{O}$, the concentration change $(\mathcal{S}r)_i$ is nonnegative $(\mathcal{S}r)_i \geq 0$.

We define $\Pi(\mathcal{R}')$ with $\mathcal{R}' \subseteq \mathcal{R}$ as the set containing every flux vector $r \in \Pi(\mathcal{R}')$ that exists with a positive flux $r_j > 0$ for every reaction $j \in \mathcal{R}'$ and zero flux $r_j = 0$ for every reaction $j \in \mathcal{R} \setminus \mathcal{R}'$. Note that one organisation can be a subset of another, so organisations can have a hierarchical structure (Dittrich and Speroni di Fenizio, 2007; Centler et al., 2008). Such a hierarchical structure is considered a requirement for evolvability in autocatalytic sets and organisations, since it shows that they can combine and grow over time (Hordijk et al., 2012; Speroni di Fenizio, 2015; Hordijk et al., 2018). The following additional definitions are useful when discussing algorithms designed to find organisations within a reaction network (Dittrich and Speroni di Fenizio, 2007; Centler et al., 2008).

Definition 3. A species $i \in \mathcal{O}$ with $\mathcal{O} \subseteq \mathcal{M}$ is *consumed within* $R(\mathcal{O})$ if there exists a reaction $j \in R(\mathcal{O})$ with $s_{i,j} < 0$ and is *produced within* $R(\mathcal{O})$ if there exists a reaction $j \in R(\mathcal{O})$ with $s_{i,j} > 0$.

Definition 4. A set of species $\mathcal{O} \subseteq \mathcal{M}$ is called *semi-self-maintaining* if all species $s \in \mathcal{O}$ that are consumed within implied network $R(\mathcal{O})$ are also produced within it.

Note that a self-maintaining set of species is also semi-self-maintaining. A set of species that is closed and semi-self-maintaining is a *semi-organisation*.

Definition 5. A set of species $\mathcal{O} \subseteq \mathcal{M}$ is *reactive connected* if every species is a product or reactant of at least one reaction, and if for any two species $o_i \in \mathcal{O}$ and $o_j \in \mathcal{O}$ there exists a sequence of n species (o_1, \dots, o_n) with $o_k \in \mathcal{O}$ for $k = 1, \dots, n$ such that $o_i = o_1$, $o_j = o_n$ and o_k is directly connected to o_{k+1} for $k = 1, \dots, n-1$. Two species $o_l \in \mathcal{O}$ and $o_m \in \mathcal{O}$ are directly connected if there exists a reaction $(A \rightarrow B) \in R(\mathcal{O})$ such that $o_l \in A \cup B$ and $o_m \in A \cup B$.

Several algorithms have been proposed for finding organisations within any given reaction network (Centler et al., 2008; Milreu et al., 2010; Speroni di Fenizio, 2015). This static approach to finding organisations is complemented by a dynamic approach which can relate a reaction vessel's state at any point in time to an organisation (Dittrich and Speroni di Fenizio, 2007; Matsumaru et al., 2006b). Although it is possible to precisely determine the evolution of organisations through time, predictions about the evolution of organisations can still be made by attaining the complete set of organisations within a reaction network and determining their hierarchical structure (Speroni di Fenizio, 2015).

Organisations in Matter-conserving Chemistries

Milreu et al. (2010) define a reaction network as being *mass-consistent* if there exists a mass vector \mathbf{m} containing a mass $m_i > 0$ for each species i in the network such that $\mathbf{m}\mathcal{S} = 0$. In such a reaction network, there can be no *inflow* or *outflow* reactions. An inflow reaction ($\emptyset \rightarrow B$) produces products B without consuming any species, and an outflow reaction ($A \rightarrow \emptyset$) consumes reactants A without producing any species (Dittrich and Speroni di Fenizio, 2007). Here, the term *matter-conserving* is used to describe a mass-consistent reaction network or chemistry that generates such a network.

In an organisation within a matter-conserving reaction network (MCRN), reactions cannot exist that will eventually deplete any species. A reaction network or subnetwork can be transformed into a bipartite graph to demonstrate that the network can only be an organisation if every edge within the directed graph is a part of a cycle. Let $\langle \mathcal{M}, \mathcal{R} \rangle$ be a MCRN and let $\mathcal{O} \subseteq \mathcal{M}$ be a closed set of species with implied reaction network $R(\mathcal{O}) \subseteq \mathcal{R}$.

Definition 6. The *bipartite representation* $G(\mathcal{O}) = (V, E)$ of \mathcal{O} is a directed graph with set of nodes V and set of edges E that are constructed using the following steps:

1. Add a node v_i to V for each species $i \in \mathcal{O}$.
2. Add a node v_j to V for each reaction $j \in R(\mathcal{O})$.
3. Add a directed edge to E starting at v_i and ending at v_j for each combination of species and reaction $(i, j) \in \mathcal{O} \times R(\mathcal{O})$ that exists with stoichiometric coefficient $s_{i,j} < 0$.
4. Add a directed edge to E starting at v_j and ending at v_i for each combination of species and reaction $(i, j) \in \mathcal{O} \times R(\mathcal{O})$ that exists with stoichiometric coefficient $s_{i,j} > 0$.

Let $G(\mathcal{O}) = (V, E)$ be the bipartite representation of \mathcal{O} . Note that all $(v_1, v_2) \in E$ either start from a reaction node $v_1 = v_{A \rightarrow B}$ and end at a species node $v_2 = v_i$, or start at a species node $v_1 = v_i$ and end at a reaction node $v_2 = v_{A \rightarrow B}$. Hereafter a directed edge is notated as an ordered pair of nodes (x, y) with the first element of the pair x being the

start of the edge and the second element y being the end of the edge. To demonstrate that every edge $(v_1, v_2) \in E$ must be a part of a cycle for \mathcal{O} to be an organisation, we must first provide several definitions.

Definition 7. A sequence of nodes (v_1, \dots, v_n) with $v_j \in V$ for $j = 1, \dots, n$ is a *path* if all nodes in the sequence are distinct and there exists an edge $(v_j, v_{j+1}) \in E$ for $j = 1, \dots, n-1$. Note that each node $v_j \in V$ could be a reaction node $v_{A \rightarrow B}$ or a species node v_i .

Definition 8. A sequence of nodes (v_1, \dots, v_n) with $v_j \in V$ for $j = 1, \dots, n$ is a *cycle* if nodes v_1, \dots, v_{n-1} are distinct, $v_1 = v_n$ and there exists an edge $(v_j, v_{j+1}) \in E$ for $j = 1, \dots, n-1$. Note that all cycles will have a minimum length of three because a reaction node or species node cannot be connected to itself.

Definition 9. A reaction $(A \rightarrow B) \in R(\mathcal{O})$ *contains* an edge $(v_1, v_2) \in E$ if the edge ends or starts at reaction node $v_{A \rightarrow B}$. Formally, $(A \rightarrow B)$ contains (v_1, v_2) if $v_1 = v_{A \rightarrow B}$ and therefore $v_2 = v_i$ with $i \in B$ or if $v_2 = v_{A \rightarrow B}$ and therefore $v_1 = v_i$ with $i \in A$.

Definition 10. A sequence (x_1, \dots, x_n) *contains* an ordered pair (y, z) if there exists an element of the sequence x_j such that $y = x_j$ and $z = x_{j+1}$.

Definition 11. A node $y \in V$ is *reachable* from node $x \in V$ if $y = x$ or there exists a path (v_1, \dots, v_n) with $v_1 = x$ and $v_2 = y$.

Definition 12. A node $y \in V$ is *conversely reachable* from node $x \in V$ if $y = x$ or there exists a path (v_1, \dots, v_n) with $v_1 = y$ and $v_2 = x$.

Let $V^{\leftarrow}(v)$ be the set of all nodes that are reachable from $v \in V$ and let $V^{\rightarrow}(v)$ be the set of all nodes that are conversely reachable from $v \in V$. To demonstrate that every edge $(v_1, v_2) \in E$ must be a part of a cycle, we must identify why the presence of an edge that is not part of a cycle prevents \mathcal{O} from being an organisation. A formal definition of such an edge is provided.

Definition 13. An edge $(v_1, v_2) \in E$ is an *acyclic edge* if there is no cycle within $G(\mathcal{O})$ that contains (v_1, v_2) .

Lemma 1. Given any acyclic edge $(v_1, v_2) \in E$, the set of nodes conversely reachable from v_1 do not intersect with the set of nodes reachable from v_2 , otherwise a path would exist from v_2 to v_1 which would mean that the edge is part of a cycle. Formally, $V^{\rightarrow}(v_1) \cap V^{\leftarrow}(v_2) = \emptyset$.

We can partition the nodes of a bipartite representation $G(\mathcal{O}) = (V, E)$ containing an acyclic edge into two sets of nodes (V_s, V_t) such that there is no edge in E starting from a node in V_t that ends at a node in V_s . This will be used to demonstrate that the set of species represented by nodes within V_s cannot collectively maintain themselves.

Definition 14. An *acyclic cut* partitions V into two subsets (V_s, V_t) based on an acyclic edge $(e_1, e_2) \in E$ such that

V_t is all nodes reachable from node e_2 and V_s is all other existing nodes. Formally, $V_t = V^{\leftarrow}(e_2)$ and $V_s = V \setminus V_t$.

Let (V_s, V_t) be an acyclic cut based on acyclic edge $(e_1, e_2) \in E$ and let E_c be the set of all edges $(v_1, v_2) \in E$ with $v_1 \in V_s$ and $v_2 \in V_t$. Note that $(e_1, e_2) \in E_c$ because $e_1 \in V_s$ and $e_2 \in V_t$.

Lemma 2. All edges in E_c are also acyclic edges.

Proof. All edges $(v_1, v_2) \in E_c$ are acyclic edges because $v_1 \in V_s$ and $v_2 \in V_t$. If (v_1, v_2) is part of a cycle then v_1 would be reachable from $v_2 \in V_t$ and therefore $v_1 \in V_t$ because V_t contains all nodes reachable from v_2 . \square

Lemma 3. There is no edge in E starting from a node in V_t that ends at a node in V_s .

Proof. There cannot exist an edge $(v_1, v_2) \in E$ with $v_1 \in V_t$ and $v_2 \in V_s$ because if $v_1 \in V_t$ then $v_2 \in V_t$ since v_2 is reachable from v_1 which is reachable from e_2 and $V_t = V^{\leftarrow}(e_2)$. \square

Lemma 4. A reaction $j \in R(\mathcal{O})$ cannot contain both an edge in E_c that starts at v_j and an edge in E_c that ends at v_j .

Proof. For all edges $(v_1, v_2) \in E_c$, $v_1 \in V_s$ and $v_2 \in V_t$. So a reaction $j \in R(\mathcal{O})$ cannot contain an edge $(v_1, v_2) \in E_c$ with $v_1 = v_j$ and an edge $(w_1, w_2) \in E_c$ with $w_2 = v_j$ because this would mean $v_j \in V_t$ and $v_j \in V_s$ which is impossible because $V_t \cap V_s = \emptyset$. \square

Let $V(X)$ be the set of species nodes for all species $i \in X$ with $X \subseteq \mathcal{O}$ and let $\mathcal{O}(Y)$ be the set of species for all nodes $v_i \in Y$ with $Y \subseteq V$ and $i \in \mathcal{O}$. Formally, $V(X) = \{v_i \mid i \in X\}$ and $\mathcal{O}(Y) = \{i \in \mathcal{O} \mid v_i \in Y\}$. Let $m(Y, j)$ be the net change in mass for reaction $j \in R(\mathcal{O})$ and all species $i \in \mathcal{O}(Y)$ with $Y \subseteq V$ and let $m(Y, r)$ be the net change in mass for flux vector r and all species $i \in \mathcal{O}(Y)$ with $Y \subseteq V$. Formally, $m(Y, j) = \sum_{i \in \mathcal{O}(Y)} s_{i,j} m_i$ and $m(Y, r) = \sum_{i \in \mathcal{O}(Y)} (Sr)_i m_i$. Also, let $U(X)$ be the set of distinct elements within multiset X . We can now show that there is a net loss of matter for all species $i \in \mathcal{O}(V_s)$ for all reactions involving these species.

Theorem 5. For any reaction $j \in R(\mathcal{O})$ containing an edge in E_c , $m(V_s, j) < 0$ and $m(V_t, j) > 0$.

Proof. According to lemma 4, a reaction $j \in R(\mathcal{O})$ can contain edges in E_c that end at v_j if there are no edges in E_c that start at v_j . Conversely, j can contain edges in E_c that start at v_j if there are no edges in E_c that end at v_j . So the theorem must be proven in both scenarios. There are no other scenarios in which j can contain an edge in E_c .

Consider the first scenario in which a reaction $(A \rightarrow B) \in R(\mathcal{O})$ contains at least one edge $(v_i, v_{A \rightarrow B}) \in E_c$ with $i \in A$. For every species $i \in A$ with $(v_i, v_{A \rightarrow B}) \in E_c$, node $v_i \in V_s$ because edges in E_c are acyclic (see lemma

2) and $s_{i,A \rightarrow B} < 0$ (see definition 6). For every species $i \in B$, $s_{i,A \rightarrow B} > 0$ (see definition 6) and node $v_i \in V_t$ because v_i is reachable from node $v_{A \rightarrow B} \in V_t$. Hence, $m(V_s, A \rightarrow B) < 0$ and therefore $m(V_t, A \rightarrow B) > 0$ because all species produced by the reaction are in $\mathcal{O}(V_t)$, at least one species consumed by the reaction is in $\mathcal{O}(V_s)$ and in a MCRN $mS = 0$.

Consider the second scenario in which a reaction $(A \rightarrow B) \in R(\mathcal{O})$ contains at least one edge $(v_{A \rightarrow B}, v_i) \in E_c$ with $i \in B$. For every species $i \in B$ with $(v_{A \rightarrow B}, v_i) \in E_c$, node $v_i \in V_t$ because edges in E_c are acyclic (see lemma 2) and $s_{i,A \rightarrow B} > 0$ (see definition 6). For every species $i \in A$, $s_{i,A \rightarrow B} < 0$ (see definition 6) and node $v_i \in V_s$ because v_i is conversely reachable from node $v_{A \rightarrow B} \in V_s$. Hence, $m(V_t, A \rightarrow B) > 0$ and therefore $m(V_s, A \rightarrow B) < 0$ because all species consumed by the reaction are in $\mathcal{O}(V_s)$, at least one species produced by the reaction is in $\mathcal{O}(V_t)$ and in a MCRN $mS = 0$. \square

Theorem 6. In $R(\mathcal{O})$ there is a net loss of matter for all species $i \in \mathcal{O}(V_s)$. Formally, $m(V_s, r) < 0$ for any flux vector $r \in \Pi(R(\mathcal{O}))$.

Proof. We have already proven in theorem 5 that $m(V_s, A \rightarrow B) < 0$ and $m(V_t, A \rightarrow B) > 0$ for any $(A \rightarrow B) \in R(\mathcal{O})$ containing an edge in E_c . This means that $U(A) \cap \mathcal{O}(V_s) \neq \emptyset$ because there must be at least one species $i \in \mathcal{O}(V_s)$ with $s_{i,A \rightarrow B} < 0$ which makes i a reactant and $U(B) \cap \mathcal{O}(V_t) \neq \emptyset$ because there must be at least one species $i \in \mathcal{O}(V_t)$ with $s_{i,A \rightarrow B} > 0$ which makes i a product. Hence, $\mathcal{O}(V_s) \neq \emptyset$ and $\mathcal{O}(V_t) \neq \emptyset$ because E_c contains at least (e_1, e_2) which acyclic cut (V_s, V_t) is based on, and there exists at least one reaction containing an edge in E_c since all edges start or end at a reaction node.

If a reaction $j \in R(\mathcal{O})$ does not contain an edge in E_c then j must either exclusively contain edges $(v_1, v_2) \in E$ with $\{v_1, v_2\} \subset V_s$ or exclusively contain edges $(v_1, v_2) \in E$ with $\{v_1, v_2\} \subset V_t$ since there cannot exist an edge $(v_1, v_2) \in E$ with $v_1 \in V_t$ and $v_2 \in V_s$ (see lemma 3). The equality $\sum_{i \in \mathcal{O}(V_s)} s_{i,j} r_j m_i = 0$ holds for any flux $r_j > 0$ and for all reactions $j \in R(\mathcal{O})$ that exclusively contain edges $(v_1, v_2) \in E$ with $\{v_1, v_2\} \subset V_s$. This equality also holds for any flux $r_j > 0$ and for all reactions $j \in R(\mathcal{O})$ that exclusively contain edges $(v_1, v_2) \in E$ with $\{v_1, v_2\} \subset V_t$ because $\mathcal{O}(V_t) \cap \mathcal{O}(V_s) = \emptyset$. The inequality $\sum_{i \in \mathcal{O}(V_s)} s_{i,j} r_j m_i < 0$ with $r_j > 0$ holds for all reactions $j \in R(\mathcal{O})$ that contain at least one edge in E_c (of which there is at least one). Hence, $m(V_s, r) < 0$ for any flux vector $r \in \Pi(R(\mathcal{O}))$. \square

Lastly, we prove that an organisation cannot contain an acyclic edge.

Theorem 7. An organisation within a MCRN cannot contain an edge within its bipartite representation that is not part of a cycle.

Proof. If a vector is taken from \mathcal{S} for a species $i \in \mathcal{O}$ and multiplied by some positive scalar $x > 0$, then the sign of the concentration change will not be altered. Formally, $\text{sgn}(x(\mathcal{S}r)_i) = \text{sgn}((\mathcal{S}r)_i)$ where $\text{sgn}(y) = 1$ if $y > 0$, $\text{sgn}(y) = 0$ if $y = 0$ and $\text{sgn}(y) = -1$ if $y < 0$. Hence, if $(\mathcal{S}r)_i \geq 0$ then $x(\mathcal{S}r)_i \geq 0$ and if $(\mathcal{S}r)_i < 0$ then $x(\mathcal{S}r)_i < 0$. If \mathcal{O} is a closed set of species with a bipartite representation $G(\mathcal{O}) = (V, E)$ containing an acyclic edge $(e_1, e_2) \in E$, then there exists an acyclic cut (V_s, V_t) based on (e_1, e_2) . According to theorem 6, there is a net loss of matter $m(V_s, r) < 0$ for any flux vector $r \in \Pi(R(\mathcal{O}))$. If $m(V_s, r) < 0$ for any such flux vector then there always exists at least one species $i \in \mathcal{O}$ with $(\mathcal{S}r)_i m_i < 0$ and therefore $(\mathcal{S}r)_i < 0$, because $\text{sgn}((\mathcal{S}r)_i m_i) = \text{sgn}((\mathcal{S}r)_i)$. The existence of a negative concentration change $(\mathcal{S}r)_i < 0$ for a species i within closed set \mathcal{O} means that \mathcal{O} cannot be self-maintaining. Hence, any closed set \mathcal{O} cannot be an organisation if its bipartite representation contains an edge that is not part of a cycle. \square

A decomposition theorem of organisations (Veloz et al., 2011; Veloz and Razeto-Barry, 2017) provides a way of partitioning a reaction network or subnetwork into modules to determine if it is an organisation. Matter-conservation has implications for this theorem. The *potential fragile-circuit* of a set of species \mathcal{O} is defined as all species within \mathcal{O} apart from *non-reactive*, *overproduced* or *catalyst* species (Veloz et al., 2011; Veloz and Razeto-Barry, 2017). A species is non-reactive if it is neither a reactant nor a product of any reaction in $R(\mathcal{O})$. A species $i \in \mathcal{O}$ is overproduced by flux vector $r \in \Pi(R(\mathcal{O}))$ if $(\mathcal{S}r)_i > 0$. A species $i \in \mathcal{O}$ is a catalyst if $s_{i,j} = 0$ for every reaction $j \in R(\mathcal{O})$ and $i \in A$ for at least one reaction $(A \rightarrow B) \in R(\mathcal{O})$. The potential fragile-circuit of \mathcal{O} is $\mathcal{O} - (N \cup C \cup F)$ where $N \subseteq \mathcal{O}$, $C \subseteq \mathcal{O}$ and $F \subseteq \mathcal{O}$ is the set of non-reactive, catalyst and overproduced species, respectively. An organisation within a MCRN has a net production of zero as it cannot contain an overproduced species.

Lemma 8. *If \mathcal{O} is an organisation within a MCRN then $(\mathcal{S}r)_i = 0$ for every species $i \in \mathcal{O}$ and any flux vector $r \in \Pi(R(\mathcal{O}))$.*

Proof. The equality $\sum_{i \in \mathcal{O}} (\mathcal{S}r)_i m_i = 0$ holds for any flux $r \in \Pi(R(\mathcal{O}))$ because a MCRN is mass-consistent. This means that if $(\mathcal{S}r)_i m_i > 0$ for some species $i \in \mathcal{O}$, then there must be at least one other species $k \in \mathcal{O}$ with $(\mathcal{S}r)_k m_k < 0$. Since $\text{sgn}((\mathcal{S}r)_l m_l) = \text{sgn}((\mathcal{S}r)_l)$ for any species $l \in \mathcal{O}$, if $(\mathcal{S}r)_i > 0$ for some species $i \in \mathcal{O}$ there must be at least one other species $k \in \mathcal{O}$ with $(\mathcal{S}r)_k < 0$ which means \mathcal{O} cannot be an organisation because it is not self-maintaining. Hence, if \mathcal{O} is an organisation then $(\mathcal{S}r)_i = 0$ for every species $i \in \mathcal{O}$. \square

One consequence of lemma 8 is that aside from non-reactive species and catalysts, all other species in \mathcal{O} con-

stitute the potential fragile-circuit of \mathcal{O} . If \mathcal{O} is an organisation, every species within the potential fragile-circuit must be both consumed and produced within $R(\mathcal{O})$ (see lemma 4 in Veloz et al., 2011). Note that in a MCRN, an acyclic edge can exist within \mathcal{O} 's bipartite representation which prevents \mathcal{O} from being self-maintaining even if every species within \mathcal{O} 's potential fragile-circuit is both consumed and produced within $R(\mathcal{O})$. For example, in the following network

$$\mathcal{O} = \{a, b, c, d\} \quad (1)$$

$$R(\mathcal{O}) = \{a \rightarrow b, b \rightarrow a, b \rightarrow c, c \rightarrow d, d \rightarrow c\} \quad (2)$$

every species within \mathcal{O} is both consumed and produced within $R(\mathcal{O})$, but reaction $b \rightarrow c$ induces an acyclic edge. A potential fragile-circuit can be decomposed into many independent smaller fragile-circuits which are linked by catalysts or overproduced species (Veloz et al., 2011; Veloz and Razeto-Barry, 2017). Another consequence of lemma 8 is that in a MCRN, these smaller circuits can only be linked by catalysts if the larger potential fragile-circuit is within an organisation.

Algorithm Optimisation

We propose an altered version of an algorithm given in Center et al. (2008) for computing reactive connected organisations of a given reaction network. The original algorithm is known as the *constructive approach*, and uses the reaction network's structure to find chemical organisations more efficiently than simply testing all 2^n combinations of n species. We propose a version of the algorithm for MCRNs that more efficiently computes the organisations by taking advantage of the fact that an organisation within a MCRN cannot contain an acyclic edge in its bipartite representation. The algorithm consists of four main functions which are shown as pseudo-code in Figs. 1-4. The algorithm first finds all semi-organisations in the reaction network and then determines which ones are organisations.

We will first redefine the concepts of semi-self-maintenance and semi-organisation for MCRNs. The bipartite representation of $\mathcal{O} \subseteq \mathcal{M}$ contains an edge that is not part of a cycle if \mathcal{O} contains a species that is only consumed or only produced within $R(\mathcal{O})$ (see definitions 3 and 6). According to theorem 7, \mathcal{O} cannot be an organisation if its bipartite representation contains an edge that is not part of a cycle. Hence, semi-self-maintenance can be redefined as follows:

Definition 15. *Within a MCRN $\langle \mathcal{M}, \mathcal{R} \rangle$, a set of species $\mathcal{O} \subseteq \mathcal{M}$ is called **semi-self-maintaining** if all species $s \in \mathcal{O}$ that are consumed within implied network $R(\mathcal{O})$ are also produced within it and if all species $s \in \mathcal{O}$ that are produced within implied network $R(\mathcal{O})$ are also consumed within it.*

A *semi-organisation* in a MCRN is closed and semi-self-maintaining according to definition 15. Since this definition of semi-self-maintenance is more restrictive, the al-

Function computeOrgs:
Input: MCRN $\langle \mathcal{M}, \mathcal{R} \rangle$
Output: set of all organisations in result
result $\leftarrow \emptyset$;
SOsToCheck $\leftarrow \{\text{closure}(\{\})\}$;
while SOsToCheck $\neq \emptyset$ **do**
 current $\leftarrow \text{getSmallestSO}(\text{SOsToCheck})$;
 SOsToCheck $\leftarrow \text{SOsToCheck} \cup$
 CSOsDirectlyAbove(current);
 SOsToCheck $\leftarrow \text{SOsToCheck} \setminus \{\text{current}\}$;
 result $\leftarrow \text{result} \cup \{\text{current}\}$;
end
End Function

Figure 1: Returns all organisations within a reaction network

gorithm can ignore sets of species that only satisfy the original definition of semi-self-maintenance. Note that semi-self-maintenance according to definition 15 does not imply self-maintenance. For example, each species is both produced and consumed within the reaction network shown in equations 1 and 2, so the network is semi-self-maintaining, but reaction $b \rightarrow c$ induces an acyclic edge, so the network cannot be self-maintaining according to theorem 7.

Function computeOrgs (Fig. 1) represents the main loop of the algorithm. It is identical to the computeSemiOrganisations function (Centler et al., 2008) except it operates on a MCRN. The function finds all semi-organisations within the reaction network. The set SOsToCheck is initialised with the smallest semi-organisation which is the closure of the empty set ($\text{closure}(\{\})$). In a MCRN there are no inflow reactions so $\text{closure}(\{\})$ is the empty set. $\text{closure}(\text{set})$ computes the smallest closed set containing the set of input species set. This is done by adding all species to set that can be produced from all reactions $(A \rightarrow B) \in \mathcal{R}$ with $A \subseteq \text{set}$ and then repeating this process until no novel species can be added. The function shown in Fig. 3 also makes use of the closure function.

The SOsDirectlyAbove function (Centler et al., 2008) is called by the computeSemiOrganisations function and returns at least all semi-organisations *directly above* input semi-organisation so. A semi-organisation is directly above so if it contains so and does not contain any other semi-organisations that contain so. However, a combinatorial explosion can lead to a very large number of organisations. For example, a reaction network containing n species but no reactions contains 2^n organisations because every combination of species is closed and self-maintaining. Hence, it can be computationally costly to compute all organisations. Also, enumerating sets of non-interacting species is unlikely to be of interest to researchers investigating the chemical origins of life. The ConnectedSOsDirectlyAbove

Function CSOsDirectlyAbove:
Input: semi-organisation so, MCRN $\langle \mathcal{M}, \mathcal{R} \rangle$
Output: set of all semi-organisations directly above so in result
result $\leftarrow \emptyset$;
usable $\leftarrow \emptyset$;
if so = $\{\}$ **then**
 usable $\leftarrow \cup_{s \in \mathcal{M}} \{\{s\}\}$;
else
 foreach $u \in \mathcal{R}$ with reactants $(u) \not\subseteq \text{so}$ **do**
 if $\exists s \in \text{so}$ with $s \in \text{reactants}(u) \cup$
 products (u) **then**
 usable $\leftarrow \text{usable} \cup$
 $\{\text{reactants}(u) \setminus \text{so}\}$;
 end
 end
end
foreach set $\in \text{usable}$ **do**
 result $\leftarrow \text{result} \cup$
 SOsDirectlyAboveContaining(so,
 set);
end
End Function

Figure 2: Returns all semi-organisations directly above the input one

function (Centler et al., 2008) is the same as the original SOsDirectlyAbove function, except it only returns reactive connected semi-organisations, thereby ensuring the algorithm only returns organisations containing interacting species. The CSOsDirectlyAbove function (Fig. 2) is identical to the original ConnectedSOsDirectlyAbove function (Centler et al., 2008) except it operates on a MCRN and is called by computeOrgs. We choose to use CSOsDirectlyAbove instead of SOsDirectlyAbove because of its practical advantages. However, it could be swapped with a function identical to SOsDirectlyAbove that operates on a MCRN.

Function computeOrgs takes the smallest semi-organisation current from SOsToCheck, adds the semi-organisations returned by CSOsDirectlyAbove(current) to SOsToCheck, removes current from SOsToCheck and then repeats this entire process until SOsToCheck is empty. In practice, SOsToCheck would be some kind of hash structure that ensures the same semi-organisation cannot be added to SOsToCheck more than once.

Function CSOsDirectlyAbove(so) iterates through sets of species that are directly connected to species in so and for each set finds all organisations directly above so that contain the set. To do this, CSOsDirectlyAbove calls SOsDirectlyAboveContaining(so, species)

Function `SODirectlyAboveContaining`:
Input: semi-organisation `so`, species set `species` to be contained in new semi-orgs., MCRN $\langle \mathcal{M}, \mathcal{R} \rangle$
Output: set of all semi-organisations directly above `so` that contain `species` in result
`result` $\leftarrow \emptyset$;
`closure` $\leftarrow \text{closure}(so \cup \text{species})$;
if `closure` is semi-self-maintaining **then**
 `result` $\leftarrow \{\text{closure}\}$;
else
 `cSets` \leftarrow
 `producerAndConsumerSets(closure)`;
 foreach `set` \in `cSets` **do**
 `result` $\leftarrow \text{result} \cup$
 `SODirectlyAboveContaining(so,`
 `species \cup set)`;
 end
end

End Function

Figure 3: Takes an input semi-organisation and returns all semi-organisations directly above it containing additionally specified species

to find a semi organisation containing both `so` and species set `species`. The `SODirectlyAboveContaining` function (Fig. 3) is almost identical to the original `SODirectlyAboveContaining` function (Centler et al., 2008) except it functions slightly differently because it operates on a MCRN. First, the closure of the union of `so` and `species` is computed and stored in `closure` which is then tested to see if it is semi-self-maintaining. If it is semi-self-maintaining then a semi-organisation has been found and the function returns.

If `closure` is not semi-self-maintaining then function `producerAndConsumerSets` (Fig. 4) is used to retrieve all possible species combinations that produce or consume the species in `closure` that are only consumed or only produced in `closure`'s implied reaction network respectively. `producerAndConsumerSets` is similar to the original `producerSets` function (Centler et al., 2008). The difference is that `producerSets` does not necessarily operate on a MCRN and therefore only returns all possible species combinations that produce the species in `closure` that are only consumed in `closure`'s implied reaction network. For each species combination returned by `producerAndConsumerSets`, `SODirectlyAboveContaining(so, species)` is called recursively but with the species combination added to `species`.

Each semi-organisation returned by `computeOrgs` that is self-maintaining is an organisation. Checking if each semi-organisation is an organization is a linear programming

Function `producerAndConsumerSets`:
Input: semi-organisation `so` and matter-conserving reaction network $\langle \mathcal{M}, \mathcal{R} \rangle$
Output: set of all species combinations in result that can produce (consume) all species that are only consumed (produced) in `so`
`result` $\leftarrow \emptyset$;
foreach `m` \in `so` where `m` is consumed and not produced within $R(\text{so})$ **do**
 `cSetsm` $\leftarrow \emptyset$;
 foreach `u` \in \mathcal{R} with $s_{m,u} > 0$ **do**
 `cSetsm` $\leftarrow \text{cSets}_m \cup \{\text{reactants}(u) \setminus \text{so}\}$;
 end
end
foreach `m` \in `so` where `m` is produced and not consumed within $R(\text{so})$ **do**
 `cSetsm` $\leftarrow \emptyset$;
 foreach `u` \in \mathcal{R} with $s_{m,u} < 0$ **do**
 `cSetsm` $\leftarrow \text{cSets}_m \cup \{\text{reactants}(u) \setminus \text{so}\}$;
 end
end
repeat
 `current` $\leftarrow \emptyset$;
 foreach `m` \in `so` for which $\exists \text{cSets}_m$ **do**
 select a set `cSetsM` from `cSetsm`;
 `current` $\leftarrow \text{current} \cup \text{cSetsM}$;
 end
 `result` $\leftarrow \text{result} \cup \{\text{current}\}$;
until all possible set combinations have been considered;

End Function

Figure 4: Returns all combinations of species facilitating reactions consuming (producing) all species that are only produced (consumed) by the input semi-organisation

problem in which a flux vector r must be found that satisfies the conditions of self-maintenance (Centler et al., 2008). Finding an optimal flux vector is not necessary since meeting the condition of self-maintenance only requires $(Sr)_i$ to be non-negative (Centler et al., 2008). After all organisations have been identified, the algorithm completes.

Experimental Results

We conducted a short experiment to verify that the altered algorithm presented in this paper works and is faster than the original version of the algorithm. We randomly generate 100 Spiky-RBN (Krastev et al., 2016) chemistries, each with 12 instances of 12 unique atomic particle species. We will briefly summarise the Spiky-RBN AChem. Each atomic particle species is a Random Boolean Network (RBN) with 12 nodes and 2 inputs per node. A bonding site, or 'interaction list', is a grouping of RBN nodes that can interlink with the nodes of another interaction list. Each of these in-




Figure 5: Execution time of both algorithms and number of organisations found per reaction network

teraction lists has a ‘spike’ which is determined by the dynamics of the RBN. We attempt 10,000 reactions for each chemistry in an aspatial, matter-conserving reactor (Krashev et al., 2017). A single reaction attempt consists of randomly picking two atomic particles, attempting to randomly choose an unlinked interaction list on both particles, and interlinking them if their spikes sum to zero. We infer that a new link cannot form between two particles that are already linked. Link formation changes the underlying RBN and can therefore change the spikes. So, directly after link formation every linked pair of interaction lists in the new composite particle is checked to see if the spikes still sum to zero. Any link that does not meet this criterion is removed by reversing the interlinking process. A reaction network is generated for all reactions that take place. A reaction is ignored if the multiset of products is equal to the multiset of reactants. Two particles or composite particles are considered the same species if their underlying RBNs are identical. We find the complete set of reactive connected organisations within each chemistry’s reaction network using Centler et al.’s original algorithm for finding reactive connected organisations (Centler et al., 2008) and the altered version of this algorithm proposed in this paper. We will refer to Centler et al.’s algorithm as the *original* algorithm, and our altered version as the *altered* algorithm hereafter. For each reaction network, both algorithms found the exact same set of organisations.

Fig. 5 shows the execution time of the original algorithm (vertical axis) and the altered algorithm (horizontal axis) for every reaction network, each of which is represented by a single data point. A line shows where the execution times of the algorithms would be equal. The shade of each data point represents how many organisations containing at least two species were found (crosses indicate zero found). Logarithmic scales were used except for organisation counts between zero and one which were made linear (due to zero counts). The plot clearly shows that in the majority of cases the execution time of the altered algorithm is much faster than the execution time of the original algorithm. This is

not true for 16 out of the 100 reaction networks. In these cases the cost of the altered algorithm’s iteration over combinations of consuming sets may outweigh the cost of the original algorithm’s iteration over semi-organisation candidates ignored by the altered algorithm. The altered algorithm was 8.92 times faster than the original on average. This average was obtained by dividing the execution time of the original algorithm by the execution time of the altered one for each reaction network and calculating the mean of these values. The number of organisations found appears to increase with execution time. This implies that there are more organisations to be found in reaction networks which are more difficult for the algorithm to process. The mean number of organisations is 143.38, although no organisations were found in 31 out of 100 reaction networks. The difference in execution time between algorithms appears to increase with the number of organisations found, which implies that the benefit of using the altered algorithm increases as the reaction network becomes more difficult to process.

Conclusions and Future Work

We have shown that it is possible to compute the complete set of organisations in a reaction network more efficiently if we know the network has been generated by a matter-conserving chemistry. Adding conservation-of-matter to an AChem has been shown to produce interesting behaviour (Hickinbotham and Stepney, 2015; Krashev et al., 2017) and the theorems and algorithm we provide could facilitate further research into these behaviours. For example, researchers may want to understand if adding conservation-of-matter to an AChem changes the diversity and evolvability of organisations. Such research could contribute knowledge about what factors may prohibit or permit the evolution of self-maintaining sets of molecular species on the way to the emergence of minimal artificial lifeforms. We are currently using the altered algorithm to compare the diversity and evolvability of organisations between different variants of the Spiky-RBN AChem (Krashev et al., 2016). The tools provided here can also be applied to reaction networks for natural systems and may therefore be useful in other fields of research. For example, COT has been used to study atmospheric photochemistry models of Mars (Centler and Dittrich, 2007) and models of sugar metabolism in *E. coli* (Centler et al., 2007). Models of natural reaction networks may exist which are matter-conserving and could therefore be analysed using our more efficient algorithm. In future research, theorem 7 could be further exploited to find the complete set of organisations even faster.

Acknowledgements

We thank the anonymous reviewers for helpful comments. Jonathan is funded by the EPSRC and AHRC Centre for Doctoral Training in Media and Arts Technology.

References

- Centler, F. and Dittrich, P. (2007). Chemical organizations in atmospheric photochemistries—a new method to analyze chemical reaction networks. *Planetary and Space Science*, 55(4):413–428.
- Centler, F., Kaleta, C., Speroni di Fenizio, P., and Dittrich, P. (2008). Computing chemical organizations in biological networks. *Bioinformatics*, 24(14):1611–1618.
- Centler, F., Speroni di Fenizio, P., Matsumaru, N., and Dittrich, P. (2007). Chemical organizations in the central sugar metabolism of escherichia coli. In *Mathematical Modeling of Biological Systems, Volume I*, pages 105–119. Springer.
- Contreras, D., Pereira, U., Hernandez, V. C., Reynaert, B., and Letelier, J.-C. (2011). A loop conjecture for metabolic closure. In *ECAL*, pages 176–183.
- Dittrich, P. and Speroni di Fenizio, P. (2007). Chemical organisation theory. *Bulletin of mathematical biology*, 69(4):1199–1231.
- Dittrich, P. and Winter, L. (2008). Chemical organizations in a toy model of the political system. *Advances in Complex Systems*, 11(04):609–627.
- Dittrich, P., Ziegler, J. C., and Banzhaf, W. (2001). Artificial chemistries—a review. *Artificial Life*, 7(3):225–275.
- Fontana, W. and Buss, L. W. (1994). ‘The arrival of the fittest’: Toward a theory of biological organization. *Bulletin of Mathematical Biology*, 56(1):1–64.
- Hickinbotham, S., Clark, E., Stepney, S., Clarke, T., Nellis, A., Pay, M., and Young, P. (2010). Specification of the stringmol chemical programming language version 0.2. Technical report, Technical Report YCS-2010-458, Dept Computer Science, University of York.
- Hickinbotham, S. and Stepney, S. (2015). Conservation of matter increases evolutionary activity. In *Artificial Life Conference Proceedings 13*, pages 98–105. MIT Press.
- Hordijk, W., Hein, J., and Steel, M. (2010). Autocatalytic sets and the origin of life. *Entropy*, 12(7):1733–1742.
- Hordijk, W., Steel, M., and Dittrich, P. (2018). Autocatalytic sets and chemical organizations: modeling self-sustaining reaction networks at the origin of life. *New Journal of Physics*, 20(1):015011.
- Hordijk, W., Steel, M., and Kauffman, S. (2012). The structure of autocatalytic sets: Evolvability, enablement, and emergence. *Acta biotheoretica*, 60(4):379–392.
- Kauffman, S. A. (1986). Autocatalytic sets of proteins. *Journal of theoretical biology*, 119(1):1–24.
- Krastev, M., Sebald, A., and Stepney, S. (2016). Emergent bonding properties in the spiky RBN achem. In *Proceedings of the Artificial Life Conference*, pages 600–607.
- Krastev, M. S., Sebald, A. A.-M., and Stepney, S. (2017). Functional grouping analysis of varying reactor types in the spiky-RBN achem. In *European Conference on Artificial Life*, pages 247–254.
- Langton, C. (1989). *Artificial Life: Proceedings Of An Interdisciplinary Workshop On The Synthesis And Simulation Of Living Systems*. Santa Fe Institute Series. Avalon Publishing.
- Matsumaru, N., Centler, F., Speroni di Fenizio, P., and Dittrich, P. (2006a). Chemical organization theory applied to virus dynamics. *IT-Information Technology*, 48(3):154–160.
- Matsumaru, N., Speroni di Fenizio, P., Centler, F., and Dittrich, P. (2006b). On the evolution of chemical organizations. In *Proc. of the 7th German Workshop of Artificial Life*, pages 135–146.
- Milreu, P. V., Acuña, V., Birmelé, E., Crescenzi, P., Marchetti-Spaccamela, A., Sagot, M.-F., Stougie, L., and Lacroix, V. (2010). Enumerating chemical organisations in consistent metabolic networks: Complexity and algorithms. In *International Workshop on Algorithms in Bioinformatics*, pages 226–237. Springer.
- Speroni di Fenizio, P. (2015). The lattice of chemical organisations. In *Artificial Life Conference Proceedings 13*, pages 242–248. MIT Press.
- Veloz, T. (2020). The complexity–stability debate, chemical organization theory, and the identification of non-classical structures in ecology. *Foundations of Science*, 25(1):259–273.
- Veloz, T. and Razeto-Barry, P. (2017). Reaction networks as a language for systemic modeling: On the study of structural changes. *Systems*, 5(2):30.
- Veloz, T., Reynaert, B., Rojas-Camaggi, D., and Dittrich, P. (2011). A decomposition theorem in chemical organizations. In *ECAL*, pages 820–827.