

Minimal Output Unstable Configurations in Chemical Reaction Networks and Deciders

Robert Brijder

the date of receipt and acceptance should be inserted later

Abstract We study the set of output stable configurations of chemical reaction deciders (CRDs). It turns out that CRDs with only bimolecular reactions (which are almost equivalent to population protocols) have a special structure that allows for an algorithm to efficiently compute their finite set of minimal output unstable configurations. As a consequence, a relatively large set of configurations may be efficiently checked for output stability.

We also provide a number of observations regarding the semilinearity result of Angluin et al. [Distrib. Comput., 2007] from the context of population protocols (which is a central result for output stable CRDs). In particular, we observe that the computation-friendly class of totally stable CRDs has equal expressive power as the larger class of output stable CRDs.

Keywords Chemical Reaction Network · Population Protocol · Vector Addition System · Output Stability · Chemical Reaction Decider

1 Introduction

In scenarios where the number of molecules in a chemical reaction network (CRN) is small, traditional continuous models for CRNs based on mass action kinetics are not suitable and one may need to consider discrete CRNs. In discrete CRNs, the number of molecules of each species is represented by a nonnegative integer and probabilities are assigned to each reaction. The computational power of discrete CRNs has been formally studied in [17] (see also [8]), where it is shown that Turing-universal computation is possible with arbitrary small (but nonzero) error probability. The implementability of arbitrary CRNs has been studied using strand displacement reactions as a primitive [18]. As observed in [17], discrete CRNs are similar to population protocols [1,4] and results carry over from one domain to the other. Recent work related to CRNs include the calculus of chemical systems [16], the study of timing issues in CRNs [10], and the study of rate-independent

continuous CRNs [7]. From now on we consider only discrete CRNs, and so we omit the adjective “discrete”.

We continue in this paper the study of CRNs that has for each given input a deterministic output [6]. Thus, we are concerned here with error-free computation and so probabilities are irrelevant and only reachability is important. A given input is accepted by such a “deterministic” CRN, or more precisely *output stable chemical reaction decider* (CRD) [6], if at the end of the “useful” computation we obtain an accept configuration c , which is a configuration where at least one yes voter is present and none of the no voters (each species is marked by the CRD as either a yes or a no voter). Otherwise, the input is rejected and c is a reject configuration, which is a configuration where at least one no voter is present and none of the yes voters. The configuration c may still change, but it stays an accept configuration when c is an accept configuration (and similar for reject). In this case c is called *output stable*.

In Section 3, we provide a number of observations regarding the semilinearity result for population protocols of [1,2]. First we mention that this result has a small gap in its proof which is easily fixable, except for the corner case where the semilinear set contains the zero vector. Next, we define a stricter variant of the notion of output stable, called *totally stable*. In contrast to output stable CRDs, totally stable CRDs eventually (completely) halt for every input. For totally stable CRDs it is computationally easy to determine when the computation has ended. We mention that the semilinearity result of [1,2] works also for totally stable CRDs, and consequently the class of totally stable CRDs has equal expressive power as the larger class of output stable CRDs.

CRNs are similar to Petri nets [15] and vector addition systems (VASs) [12], see [17]. However, Petri nets and VASs operate as “generators” where the computation starts in the given fixed starting configuration (called the initial marking) and one is (generally) interested in the reachable configurations. In contrast, a CRD is a decider where one is (generally) interested in determining the set of inputs that is accepted by the CRD. Despite these differences, various results concerning Petri nets and VASs can be carried over to CRDs.

In Section 4, we take a closer look at the notion of output stable. First, using some well-known results for VASs, we show that determining whether or not a configuration is output stable for an output stable CRD is decidable. Next, we turn to bimolecular CRNs, i.e., CRNs where each reaction has two reactants and two products. It turns out that bimolecular CRDs provide a special structure on the set of output stable configurations. More precisely, it turns out that the set of minimal elements M of the upward-closed set of output-unstable configurations may be efficiently determined for bimolecular CRDs, cf. Theorem 6 — this is the main result of the paper. By efficiently determine we mean here that the fraction of time complexity divided by size of the output is relatively small (note that the size of the output is a naive lower bound on the time complexity). Given M , it is then computationally easy to determine if a given configuration c is output stable. Consequently, the algorithm to determine M provides an efficient method to test a relatively large number of configurations for output stability (the preprocessing cost to generate M becomes smaller, relatively, when testing more configurations for output stability).

A preliminary conference version of this paper was presented at DNA 20 [5].

2 Chemical Reaction Networks and Deciders and Population Protocols

2.1 Chemical Reaction Networks

The notation and terminology of this subsection and the next are similar as in [11].

Let $\mathbb{N} = \{0, 1, \dots\}$. Let A be a finite set. The set of vectors over \mathbb{N} indexed by A (i.e., the set of functions $\varphi : A \rightarrow \mathbb{N}$) is denoted by \mathbb{N}^A . For $x \in \mathbb{N}^A$, we define the *size* of x , denoted by $\|x\|$, as $\sum_{i \in A} x(i)$. We denote the restriction of x to $\Sigma \subseteq A$ by $x|_{\Sigma}$. For $x, y \in \mathbb{N}^A$ we write $x \leq y$ if and only if $x(i) \leq y(i)$ for all $i \in A$. For notational convenience we now also denote vectors in \mathbb{N}^A , which can be regarded as multisets, by their string representations. Thus we denote $c \in \mathbb{N}^A$ by the string $A_1^{c(A_1)} \dots A_n^{c(A_n)}$ (or any permutation of these letters) where $A = \{A_1, \dots, A_n\}$.

Let A be a finite set. A *reaction* α over A is a tuple (r, p) with $r, p \in \mathbb{N}^A$; r and p are called the *reactants* and *products* of α , respectively. A reaction is commonly written in an additive fashion, where for example $A + 2B \rightarrow C$ denotes the reaction (r, p) where AB^2 and C are string representations for r and p , respectively. We say that α is *mute* if $r = p$. We say that α is *nonincreasing* if $\|r\| \geq \|p\|$ and *bimolecular* if $\|r\| = \|p\| = 2$. A *chemical reaction network* (CRN, for short) is a tuple $\mathcal{R} = (A, R)$ with A a finite set and R a finite set of reactions over A . The elements of A are called the *species* of \mathcal{R} . The elements of \mathbb{N}^A are called the *configurations* of \mathcal{R} . For a configuration c , $\|c\|$ is the number of *molecules* of c .

For a configuration $c \in \mathbb{N}^A$ and a reaction α over A , we say that $\alpha = (r, p)$ is *applicable* to c if $r \leq c$. If α is applicable to c , then the *result* of applying α to c , denoted by $\alpha(c)$, is $c' = c - r + p$. Note that $\alpha(c) \in \mathbb{N}^A$. In this case, we also write $c \rightarrow_{\alpha} c'$. Moreover, we write $c \rightarrow_{\mathcal{R}} c'$ if $c \rightarrow_{\alpha} c'$ for some reaction α of \mathcal{R} . The transitive and reflexive closure of $\rightarrow_{\mathcal{R}}$ is denoted by $\rightarrow_{\mathcal{R}}^*$. We say that c' is *reachable* from c in \mathcal{R} if $c \rightarrow_{\mathcal{R}}^* c'$. If \mathcal{R} is clear from the context, then we simply write \rightarrow and \rightarrow^* for $\rightarrow_{\mathcal{R}}$ and $\rightarrow_{\mathcal{R}}^*$, respectively.

We remark that a CRN is similar to a Petri net N [15] without the initial marking M : the set A corresponds to the set of places of N and the set of reactions R corresponds to the set of transitions of N . While in a Petri net distinct transitions in N may correspond to a single reaction in R (i.e., there may be “copies” of each transition), this is irrelevant for our purposes.

A CRN is also similar to a vector addition system (VAS) [12]. A *VAS* V is a tuple (A, S) with A a finite set and S a finite subset of \mathbb{Z}^A . Again, the elements of \mathbb{N}^A are the *configurations* of V . One is interested in the binary relation \rightarrow over \mathbb{N}^A , where $c \rightarrow c'$ if and only if $c' = c + x$ for some $x \in S$. Reachability problems concerning CRNs can be straightforwardly translated to VASs (or Petri nets) and vice versa, see [17, Appendix A.6].

2.2 Chemical Reaction Deciders

A (*leaderless*) *chemical reaction decider* (CRD, for short) is a tuple $\mathcal{D} = (A, R, \Sigma, \Upsilon)$, where (A, R) is a CRN, $\Sigma \subseteq A$, $\Upsilon : A \rightarrow \{0, 1\}$. The elements of Σ , $\Upsilon^{-1}(0)$, and $\Upsilon^{-1}(1)$ are called the *input species*, *no voters*, and *yes voters* of \mathcal{D} , respectively. Notation and terminology concerning CRNs carry over to CRDs. For example, we may speak of a configuration of \mathcal{D} . An *initial configuration* of \mathcal{D} is a nonzero

configuration c of \mathcal{D} where $c|_{A \setminus \Sigma} = 0$ (by abuse of notation we denote the zero vector over suitable alphabet by 0). A CRD is called *nonincreasing* (*bimolecular*, resp.) if all reactions of R are nonincreasing (bimolecular, resp.).

We define the following function $\Phi_{\mathcal{D}} : \mathbb{N}^A \rightarrow \{0, 1, \text{und}\}$. For $x \in \mathbb{N}^A$, let $I_x = \{S \in A \mid x(S) > 0\}$. Then, for $i \in \{0, 1\}$, we have $\Phi_{\mathcal{D}}(x) = i$ if and only if both $I_x \cap \mathcal{Y}^{-1}(i) \neq \emptyset$ and $I_x \cap \mathcal{Y}^{-1}(1-i) = \emptyset$ (as usual, \mathcal{Y}^{-1} denotes the preimage of \mathcal{Y}). If x is zero or $I_x \cap \mathcal{Y}^{-1}(0) \neq \emptyset \neq I_x \cap \mathcal{Y}^{-1}(1)$, then $\Phi_{\mathcal{D}}(x) = \text{und}$. Here, the value und is regarded as “undefined”.

A configuration c is called *totally stable* (*t-stable* for short) in \mathcal{D} if both $\Phi_{\mathcal{D}}(c) \in \{0, 1\}$ and, for all c' with $c \rightarrow^* c'$, we have $c' = c$. Note that if c is t-stable in \mathcal{D} , then for all c' with $c \rightarrow c'$, we have $c' = c$. A configuration c is called *output stable* (*o-stable* for short) in \mathcal{D} if both $\Phi_{\mathcal{D}}(c) \in \{0, 1\}$ and, for all c' with $c \rightarrow^* c'$, $\Phi_{\mathcal{D}}(c') = \Phi_{\mathcal{D}}(c)$. Note that every t-stable configuration is o-stable. A configuration that is not o-stable (t-stable, resp.) and nonzero is called *o-unstable* (*t-unstable*, resp.).

We say that \mathcal{D} *o-stably decides* (*t-stably decides*, resp.) the function $\varphi : \mathbb{N}^{\Sigma} \setminus \{0\} \rightarrow \{0, 1\}$ if for each initial configuration c of \mathcal{D} and each configuration c' with $c \rightarrow^* c'$, we have $c' \rightarrow^* c''$ where c'' is o-stable (t-stable, resp.) in \mathcal{D} and $\varphi(c|_{\Sigma}) = \Phi_{\mathcal{D}}(c'')$. In this case, we also say that \mathcal{D} *o-stably decides* (*t-stably decides*, resp.) the set $\varphi^{-1}(1)$ and that \mathcal{D} is *o-stable* (*t-stable*, resp.). Note that $\varphi^{-1}(1)$ along with the set Σ , uniquely determine φ . In [1] (and [11]), only o-stable CRDs are considered, and as a result the prefix *output* is omitted there.

Remark 1 We adopt here the definition of o-stably deciding a function/set from [2, Section 2]. In the original definition of o-stably decides from [1], an initial configuration may be the zero vector and the domain of φ contains the zero vector. Since the zero vector corresponds to an input without any molecules and the number of molecules in a bimolecular CRD stays fixed, no molecule can be introduced and, in particular, none of the yes or no voters can be introduced. As a result, there exist no o-stable bimolecular CRDs when (strictly) using the definition of [1]. Finally, we remark that there are (leaderless) CRDs that are o-stable using the definition of [1], since we may then have reactions (r, p) with r the zero vector. Since p may then be produced at any point in time, an o-stable CRD o-stably decides either \mathbb{N}^{Σ} (when p contains only yes voters) or the empty set (when p contains only no voters). Note that the CRD cannot be o-stable when p contains both yes and no voters. Thus this notion is also not interesting for the (larger) class of CRDs.

2.3 Population Protocols

The notion of population protocol [1, 4] is almost equivalent to the notion of bimolecular CRD. The only difference is that, in a population protocol, the set of reactions R is replaced by a *transition function* $\delta : A^2 \rightarrow A^2$. In this setting, $\delta(A, B) = (C, D)$ corresponds to the reaction (r, p) with $r = AB$ and $p = CD$ (recall that we may denote vectors by strings). Note that the tuples (A, B) and (C, D) are ordered. Note also that, for given $A, B \in A$, there are at most two non-mute reactions with A and B as reactants (since we have a transition for (A, B) and for (B, A)), while for bimolecular CRDs there can be arbitrary many such reactions.

Reactions, molecules, and species are called *transitions*, *agents*, and *states*, respectively, in the context of population protocols.

An important property of bimolecular CRDs is that the number of molecules stays fixed, i.e., if $c \rightarrow^* c'$, then $\|c\| = \|c'\|$.

Remark 2 In [1], $\delta(A, B) = (C, D)$ is interpreted as follows: a molecule of type A is transformed into a molecule of type C and simultaneously a molecule of type B is transformed into a molecule of type D . As a consequence, applying the “reaction” $\delta(A, B) = (B, A)$ would result in a different configuration. However, in [2] this interpretation is abandoned and $\delta(A, B) = (B, A)$ is considered a mute reaction. We adopt the convention of [2].

3 Semilinearity

In this section we state a number of modest, but useful, observations we made when studying the proof of the semilinearity result of [1].

Let A be a finite set. A set $S \subseteq \mathbb{N}^A$ is called *linear* (over A) if there are $v_0, \dots, v_n \in \mathbb{N}^A$ such that $S = \{v_0 + \sum_{i=1}^n k_i v_i \mid k_i \in \mathbb{N}, i \in \{1, \dots, n\}\}$. A set $S \subseteq \mathbb{N}^A$ is called *semilinear* (over A) if S is the union of a finite number of linear sets over A .

It is stated in [1] that every semilinear set S is o-stably decidable by a population protocol (i.e., a bimolecular CRD). While this result is often cited in the literature, it is straightforward to verify that the result fails if S contains the zero vector. Indeed, by definition of semilinear sets may contain the zero vector, while the domain of φ in the above definition of stably deciding a set is restricted to nonzero vectors (recall from Remark 1 that we have to use the definition of [2] instead of [1]). This small counterexample led us to revisit the proof of [1]. It turns out that Lemma 5 of [1] implicitly assumes that there are at least 2 agents (i.e., molecules), which translate into an initial configuration of size at least 2. Fortunately, this proof can be straightforwardly modified to allow for initial configurations of size 1, by letting, in [1, Lemma 5], I map σ_i to $(1, b, a_i)$ with $b = 1$ if and only if $a_i < c$ for case 1, and with $b = 1$ if and only if $a_i = c \pmod m$ for case 2 (instead of to $(1, 0, a_i)$) — note that these terms, such as I and σ_i , are taken from [1, Lemma 5]. In [2] (see also [3]), it is shown that if $S \subseteq \mathbb{N}^A$ is o-stably decidable by a population protocol, then S is semilinear. Thus we have the following (attributed, of course, to [1, 2]).

Theorem 1 ([1, 2]) *For every $S \subseteq \mathbb{N}^\Sigma$, S is o-stably decidable by a population protocol (i.e., a bimolecular CRD) if and only if S is both semilinear and does not contain the zero vector.*

As recalled in [6], the result from [2] that the sets o-stably decidable by population protocols are semilinear holds not only for population protocols, but for any reflexive and transitive relation \rightarrow^* that respects addition (i.e., for $c, c', x \in \mathbb{N}^\Sigma$, $c \rightarrow^* c'$ implies $c + x \rightarrow^* c' + x$). Hence, Theorem 1 holds also for the (broader) family of all CRDs.

Another observation one can make when studying [1] is that the proof concerning o-stable CRDs holds unchanged for the smaller class of t-stable CRDs. By expressive power of a family \mathcal{F} of CRDs we mean the family of sets decidable by \mathcal{F} . As the result follows from the proof of [1], we attribute it to [1].

Theorem 2 ([1]) *The family of t -stable bimolecular CRDs have equal expressive power as the family of o -stable CRDs. Equivalently, the sets that are t -stably decidable by bimolecular CRDs are precisely the semilinear sets without the zero vector.*

Proof First recall, by the comment below Theorem 1, that the expressive powers of the families of o -stable CRDs and o -stable bimolecular CRDs are equal. Now, the family of t -stable bimolecular CRDs is a subset of the family of o -stable bimolecular CRDs. Thus it suffices to show that the if-direction of Theorem 1 holds for t -stable bimolecular CRDs.

The essential part of the if-direction of the proof of Theorem 1 above is Lemma 3 and Lemma 5 from [1]. In the proof of Lemma 5 in [1] a population protocol P is described that eventually reaches a configuration c which is called “stable” in [1], and which, in fact, is easily seen to be t -stable (by checking the three conditions of “stable” in [1]). The proof of Lemma 3 in [1] trivially holds for t -stable bimolecular CRDs. \square

Since the bimolecular CRDs form a subset of the CRDs, Theorem 2 holds also when omitting the word “bimolecular”.

The family of t -stable CRDs form an interesting subclass of CRDs. Indeed, it is easy to verify, during a run of a t -stable CRD, whether or not a configuration is t -stable: one simply needs to verify whether or not there is an applicable (non-mute) reaction. In other words, it is easily verified whether or not the computation has ended. In the larger class of o -stable CRDs, it is not clear whether or not it is computationally easy to verify if a given configuration is o -stable or not. We revisit this latter problem in Section 4.

The concept of *CRDs with leaders* was introduced in [6] (it is simply called a CRD in [6]). The difference with (leaderless) CRDs is that for CRDs with leaders an additional vector $\sigma \in \mathbb{N}^{A \setminus \Sigma}$ is given and that the initial configurations c have the condition that $c|_{A \setminus \Sigma}$ is equal to σ (instead of equal to 0). Moreover, in the definition of o/t -stably deciding a function φ , the domain of φ is \mathbb{N}^Σ instead of $\mathbb{N}^\Sigma \setminus \{0\}$. Using Theorem 1, we now straightforwardly observe that CRDs with leaders decide all semilinear sets. The difference between Theorem 1 and Theorem 3 is because the zero vector for (leaderless) CRDs cannot be represented by an initial configuration of positive size, while the zero vector for CRDs with leaders can be represented by an initial configuration of positive size.

Theorem 3 ([6]) *For every $S \subseteq \mathbb{N}^\Sigma$, S is o -stably decidable by a CRD with leaders if and only if S is semilinear.*

Proof Again, by [2], every set o -stably decidable by a CRD with leaders is semilinear.

Conversely, let $S \subseteq \mathbb{N}^\Sigma$ be semilinear. Consider $\Sigma' = \{t\} \cup \Sigma$, where t is an element outside Σ . Let $S' = \{x \in \mathbb{N}^{\Sigma'} \mid x(t) = 1, x|_\Sigma \in S\}$. It is easy to verify that S' is semilinear. Indeed, let v_0, \dots, v_n be the vectors (cf. the definition of linear set) for one of the linear sets that together make up S . Then by adding an entry for t with value 1 for v_0 and value 0 for the other vectors, we see that the obtained vectors define a corresponding linear set for S' . Consequently, S' is semilinear. Note that S' does not contain the zero vector. By Theorem 1, there is a CRD $\mathcal{D} = (A, R, \Sigma', \mathcal{I})$ that o -stably decides S' . Consider now the CRD $\mathcal{D}' = (A, R, \Sigma, \mathcal{I}, \sigma)$ with leaders where $\sigma \in \mathbb{N}^{A \setminus \Sigma}$ is such that $\sigma(t) = 1$ and $\sigma(i) = 0$ if $i \in A \setminus \Sigma'$.

Consequently, the difference between \mathcal{D} and \mathcal{D}' is that index t is not part of the input species. Hence, \mathcal{D}' o-stably decides S . \square

Of course, (the proof of) Theorem 3 also holds by replacing o-stable by t-stable and/or replacing CRDs by bimolecular CRDs.

4 Determining the output stable configurations

In this section we consider the problem of determining whether or not the “useful” computation of an o-stable CRD has ended. More precisely, we consider the problem of determining whether or not a given configuration of a o-stable CRD is output stable. Recall from the previous section that it *is* straightforward to determine whether or not a given configuration c is t-stable: one simply needs to check whether or not a non-mute reaction is applicable to c (and check that $\Phi_{\mathcal{D}}(c) \in \{0, 1\}$). In Subsection 4.1 we consider the o-stable CRDs \mathcal{D} in general and in Subsection 4.2 we consider the case where \mathcal{D} is bimolecular as this subclass turns out to enjoy special properties regarding this problem.

4.1 The general case

Similar as done in [17, Theorem 4.2], we formulate now [12, Corollary 4.1] (defined in the context of VASs) in terms of CRNs.

Proposition 1 ([12]) *For given CRN \mathcal{R} and configurations x, y of \mathcal{R} , it is decidable whether or not $x \rightarrow^* y'$ for some configuration $y' \geq y$.*

A much more involved result is known as the decidability of the reachability problem for vector addition systems, shown in [14] (see [13] for a simplified proof).

Proposition 2 ([14]) *For given CRN \mathcal{R} and configurations x, y of \mathcal{R} , it is decidable whether or not $x \rightarrow^* y$.*

The precise complexity of the reachability problem of Proposition 2 is famously unknown (see, e.g., [13]).

By Propositions 1 and 2 we straightforwardly obtain the following result.

Theorem 4 *For a given o-stable CRD \mathcal{D} and configuration c of \mathcal{D} , it is decidable whether or not c is o-stable in \mathcal{D} .*

Proof Testing whether or not $\Phi_{\mathcal{D}}(c) \in \{0, 1\}$ is clearly decidable. Let $\Phi_{\mathcal{D}}(c) = j$. Let, for $X \in \Lambda$, y_X be the configuration with $\|y_X\| = 1$ and $y_X(X) = 1$. By Proposition 1 it is decidable, for each $X \in \mathcal{Y}^{-1}(1 - j)$, whether or not there exists a c' such that $c \rightarrow^* c'$ and $c' \geq y_X$, i.e., $c'(X) > 0$. Hence if c contains only yes voters, then we can decide if there is a reachable configuration with no voters (and analogously if c contains only no voters). The only case left to decide is whether or not $c \rightarrow^* 0$ (again, 0 denotes the zero vector over Λ). By Proposition 2 it is decidable if the zero vector is reachable. Consequently, it is decidable if c is o-stable in \mathcal{D} . \square

We now look at properties of o-stable configurations.

Let \mathcal{D} be an o-stable CRD. We now consider the set $U_{\mathcal{D}}$ of all output unstable configurations of \mathcal{D} . If \mathcal{D} is clear from the context, then we simply write U for $U_{\mathcal{D}}$. We now recall a useful result from [2, Lemma 10]. For convenience, we also recall its short proof.

Proposition 3 ([2]) *Let \mathcal{D} be an o-stable CRD. Then U is closed upward under \leq . In other words, for all $c, c' \in \mathbb{N}^A$ with $c \leq c'$, if $c \in U$, then $c' \in U$.*

Proof Let $c \in U$ and $c \leq c'$. If $\Phi_{\mathcal{D}}(c) = \text{und}$, then c contains both yes and no voters (since $c \in U$, c is nonzero). Thus c' also contains both yes and no voters and we have $c' \in U$. Assume that $\Phi_{\mathcal{D}}(c) \in \{0, 1\}$. If $\Phi_{\mathcal{D}}(c') = \text{und}$, then there is nothing to prove. Thus assume that $\Phi_{\mathcal{D}}(c) = \Phi_{\mathcal{D}}(c')$. Since $c \in U$, there is a c'' with $c \rightarrow^* c''$ with $\Phi_{\mathcal{D}}(c'') \neq \Phi_{\mathcal{D}}(c)$. Let $x := c' - c \in \mathbb{N}^A$. Then $c' = c + x \rightarrow^* c'' + x$ with $\Phi_{\mathcal{D}}(c'' + x) \neq \Phi_{\mathcal{D}}(c) = \Phi_{\mathcal{D}}(c')$ and $c' \in U$. \square

Remark 3 In some papers, such as [6], not all species in CRDs need to be voters. In other words, in the definition of CRD we have $\Upsilon : E \rightarrow \{0, 1\}$ for some $E \subseteq A$ (instead of $E = A$). We remark that Proposition 3 fails in this more general setting. Indeed, if nonzero c contains no voters, then $c \in U$, but by extending c with, say, a yes voter may result in an output stable configuration.

By Proposition 3, the set U is characterized by the set $\min(U)$ of minimal elements of U under \leq . By Dickson's lemma, recalled below, $\min(U)$ is a finite set.

Proposition 4 (Dickson's lemma [9]) *Let A be a finite set. Then for every $S \subseteq \mathbb{N}^A$, $\min(S)$ is finite.*

Given an o-stable CRD \mathcal{D} and the set $\min(U)$, it is straightforward to verify if a given configuration c is o-stable in \mathcal{D} . Indeed, c is o-stable in \mathcal{D} if and only if $u \not\leq c$ for all $u \in \min(U)$. Thus, to check whether or not c is o-stable in \mathcal{D} takes $O(|\min(U)|)$ configuration comparisons, i.e., $O(|\min(U)| \cdot |A|)$ comparisons of molecule counts. Note that this complexity bound depends only on \mathcal{D} , i.e., it is independent of c .

This complexity bound can be improved if the vectors of $\min(U)$ are stored in the k -fold tree $T_b(k)$ described in [20]. In general, given a set S of n vectors of dimension k and vectors a and b of dimension k , the topic of orthogonal range querying is concerned with determining the set $X = \{x \in S \mid a \leq x \leq b\}$. A related question is to determine only $|X|$. Using the data structure called k -fold trees one can efficiently determine X or $|X|$. In particular, using the k -fold tree $T_b(k)$ described in [20] it takes $O(\log^{k-\frac{1}{2}}(n))$ vector comparisons to determine whether or not $X = \emptyset$. Now, assume $\min(U)$ is stored in $T_b(k)$. Then a given vector c is o-stable in \mathcal{D} if and only if $X = \{u \in \min(U) \mid u \leq c\}$ is equal to the empty set. Consequently, taking a equal to the zero vector and b equal to c , we have the following.

Lemma 1 *Given an o-stable CRD \mathcal{D} and assume that the set $\min(U)$ is stored in a k -fold tree $T_b(k)$ with $k = |A|$. Checking whether or not a configuration c is o-stable in \mathcal{D} takes $O(\log^{k-\frac{1}{2}}(n))$ configuration comparisons, with $n = |\min(U)|$.*

Notice the improvement in complexity of Lemma 1 compared to the straightforward $O(n)$ method mentioned above. Of course, there is a penalty associated with storing $\min(U)$ in a k -fold tree: inserting a vector in a k -fold tree takes $O(\log^{k-\frac{1}{2}}(n))$ vector comparisons and thus it takes $O(n \log^{k-\frac{1}{2}}(n))$ vector comparisons to set up this data structure. Consequently, the k -fold tree method is to be used for testing a (relatively large) set of configurations for output stability (instead of just a single configuration c).

4.2 The bimolecular case

We now show that $\min(U)$ can be efficiently determined when \mathcal{D} is bimolecular. By efficiently determine we mean here that the fraction of time complexity divided by size $|\min(U)|$ of the output is relatively small — note that $|\min(U)|$ is a naive lower bound on the time complexity. This is particularly useful when one wants to test for o-stability for some large (finite) set of configurations (instead of just a single configuration).

Let, for $k \geq 0$, $\mathcal{C}_{\leq k}$ ($\mathcal{C}_{=k}$, resp.) be the set of configurations $c \in \mathbb{N}^A$ with $\|c\| \leq k$ ($\|c\| = k$, resp.).

We remark that the naive approach to determine whether or not a particular configuration c is o-stable in a o-stable bimolecular CRD \mathcal{D} , would compute the set R_c of all configurations reachable from c and then verify that $\Phi_{\mathcal{D}}(c') = \Phi_{\mathcal{D}}(c)$ for all $c' \in R_c$. Note that $R_c \subseteq \mathcal{C}_{=k}$ with $k = \|c\|$ since \mathcal{D} is bimolecular. Thus, in the worst case, one needs to compute in the order of $|\mathcal{C}_{=k}|$ configurations. The value of $|\mathcal{C}_{=k}|$ is equal to the number of multisets of cardinality k over A . This number (called figurate number, simplex number, or multiset coefficient), sometimes denoted by $\binom{|A|}{k}$, is equal to the binomial coefficient $\binom{|A|+k-1}{k}$, see, e.g., [19, Section 1.2].

We start with introducing a new binary relation \leftrightarrow . For $c, c' \in \min(U)$, denote $c \leftrightarrow_{\alpha} c'$ if $c \rightarrow_{\alpha} c' + b$ where $\alpha = (r, p)$ is a reaction and b is some configuration with $b \leq p$ and $b \neq p$. We write $c \leftrightarrow c'$ if $c \leftrightarrow_{\alpha} c'$ for some reaction α . It is important to realize that \leftrightarrow is a binary relation on $\min(U)$. Again, the transitive and reflexive closure of \leftrightarrow is denoted by \leftrightarrow^* .

We now provide some intuition regarding the notion of \leftrightarrow . Intuitively, one may also view \leftrightarrow as a graph G where $\min(U)$ is the set of vertices and $c \leftrightarrow c'$ denotes an arrow from c to c' . Lemma 2 below proves a number of properties of \leftrightarrow . Although Lemma 2 does not require that the CRD \mathcal{D} is bimolecular, assume for didactical purposes that \mathcal{D} is bimolecular. Now, Statement 5 of Lemma 2 says that for every vertex of G there is a path in G to a vertex in $M_1 \cup M_2$. The vertices of $M_1 \cup M_2 = M_1 \cup T$ are all of size 2 and can be readily computed by the first three statements of Lemma 2 (because M_1 and T can be readily computed). So, to compute all of $\min(U)$ we simply start from $M_1 \cup T$ and obtain all vertices by moving against the arrows. This is accomplished by Algorithm 1. Of course, G is not part of the input and so the arrows of G have to be found dynamically and the candidate vertices have to be checked for membership of $\min(U)$. Statement 6 describes the possible differences in size of any two adjacent vertices, which significantly restricts the search space for arrows.

For the next result, recall again that we may denote vectors by strings. Also, note that we do not require in Lemma 2 that the CRD \mathcal{D} is bimolecular.

Lemma 2 *Let $\mathcal{D} = (A, R, \Sigma, \Upsilon)$ be an o -stable CRD. Let $M_1 = \{c \in \min(U) \mid \Phi(c) = \text{und}\}$, $M_2 = \{c \in \min(U) \mid \Phi(c) \in \{0, 1\}, c \rightarrow c' \text{ for some } c' \text{ with } \Phi(c') \neq \Phi(c)\}$, and $T = \{r \in \min(U) \mid (r, p) \in R, \text{ and either } \|p\| = 0 \text{ or } \Upsilon(A) \neq \Upsilon(B) \text{ for some } A, B \in \Lambda \text{ with } r(A) \neq 0 \neq p(B)\}$. We have the following.*

1. $M_1 = \{AB \mid A, B \in \Lambda, \Upsilon(A) \neq \Upsilon(B)\}$.
2. $M_2 \subseteq T$.
3. $T \subseteq M_1 \cup M_2$.
4. If $c \rightarrow_\alpha c'$ for some $\alpha \in R$, $c \in \min(U)$, and $c' \in U$, then there is a $c'' \in \min(U)$ with $c \hookrightarrow_\alpha c''$.
5. If $\|r\| \geq \|p\|$ for all $(r, p) \in R$, then, for all $c \in \min(U)$, $c \hookrightarrow^* c'$ for some $c' \in M_1 \cup M_2$.
6. If $c \hookrightarrow c'$ and $2 = \|r\| \geq \|p\|$ for all $(r, p) \in R$, then $\|c'\| = \|c\|$ or $\|c'\| = \|c\| - 1$.

Proof The nonzero configurations where $\Phi(c) = \text{und}$ are those where there are $A, B \in \Lambda$ such that both $c(A) > 0$ and $c(B) > 0$, and $\Upsilon(A) \neq \Upsilon(B)$. The minimal such configurations are such that $c(A) = c(B) = 1$ and $c(X) = 0$ for all other species X , and so we obtain the first statement.

We now turn to the second statement. Let $c \in M_2$. Thus $c \in \min(U)$ with $\Phi(c) \in \{0, 1\}$ and $c \rightarrow c'$ for some c' with $\Phi(c') \neq \Phi(c)$. Without loss of generality, assume that $\Phi(c) = 0$, i.e., c contains only no voters. Let $\alpha = (r, p)$ be the reaction of \mathcal{D} such that $c \rightarrow_\alpha c'$. Since $\Phi(c') \neq \Phi(c)$, either $\|p\| = 0$ or a yes voter has been introduced by α . As $c \in \min(U)$, we have $c = r$. Also, if a yes voter has been introduced by α , then we have $\Upsilon(A) = 0 \neq 1 = \Upsilon(B)$ for some $A, B \in \Lambda$ with $r(A) \neq 0 \neq p(B)$.

We turn to the third statement. Let $\alpha = (r, p)$ be a reaction of \mathcal{D} such that $r \in \min(U)$ and either $\|p\| = 0$ or $\Upsilon(A) \neq \Upsilon(B)$ for some $A, B \in \Lambda$ with $r(A) \neq 0 \neq p(B)$. Assume $r \notin M_1$, i.e., $\Phi(r) \in \{0, 1\}$. Then $r \rightarrow_\alpha p$ with $\Phi(p) \neq \Phi(r)$ since either $\|p\| = 0$ or $\Upsilon(A) \neq \Upsilon(B)$ for some $A, B \in \Lambda$ with $r(A) \neq 0 \neq p(B)$. Consequently, $r \in M_2$.

We now turn to the fourth statement. Let $\alpha = (r, p)$. Since $c \in \min(U)$, we have that $c - r = c' - p \notin U$. Since $c' \in U$, we have $c'' = c' - b \in \min(U)$ for some configuration $b \leq p$ and $b \neq p$. Therefore, $c \hookrightarrow_\alpha c''$.

We now turn to the fifth statement. If $c \in M_1$, then we are done. For all $c \in \min(U) \setminus M_1$, $c \hookrightarrow^* x \rightarrow y$ for some configurations x and y with $\Phi(x) \neq \Phi(y)$. For all such c , we assign the value (k, l) where $k = \|c\|$ and l is minimal such that $c \rightarrow^l x \rightarrow y$ for some configurations x and y with $\Phi(x) \neq \Phi(y)$ (by \rightarrow^l we mean the l -th power of the relation \rightarrow). We show the result by induction on (k, l) . If $l = 0$, then $c \in M_2$ and we are done. Assume $l > 0$. Then, by the fourth statement, $c \hookrightarrow c''$ and $c = c'' + b + r - p$. As $\|r\| \geq \|p\|$, we have $\|c''\| \leq \|c\|$. If $\|c''\| < \|c\|$, then, by the induction hypothesis, $c'' \hookrightarrow^* c'$ with $c' \in M_1 \cup M_2$ and so $c \hookrightarrow^* c'$. If $\|c''\| = \|c\|$, then $c'' \rightarrow^{l-1} x \rightarrow y$. This also leads, by the induction hypothesis, to $c'' \hookrightarrow^* c'$ with $c' \in M_1 \cup M_2$ and so $c \hookrightarrow^* c'$.

For the sixth statement, by the definition of the relation \hookrightarrow , we have $c' + b = c - r + p$, $b \leq p$ and $b \neq p$. Thus $c' = c - (r - p + b)$. If $2 = \|r\| \geq \|p\|$, then $1 \leq \|p - b\| \leq \|p\| \leq 2$ and so $0 \leq \|r - p + b\| \leq 1$. \square

Lemma 2 above is key for Theorem 5 below. The strategy in the proof of Theorem 5 is to discover all elements of $\min(U)$ ordered by size: first all elements of $\min(U)$ of size k are computed, before any of the elements of $\min(U)$ of size

$k + 1$ are computed. This ensures that the generated candidates c can be tested for minimality in U , i.e., it can be tested whether or not $c \in \min(U)$. Otherwise, the number of generated candidates could potentially grow unbounded.

Theorem 5 *Let $\mathcal{D} = (A, R, \Sigma, \Upsilon)$ be an o -stable bimolecular CRD. Given \mathcal{D} , Algorithm 1 computes $\min(U)$.*

Proof First, we initialize $M := M_1 \cup M_2 = M_1 \cup T$ with M_1 , M_2 , and T from Lemma 2, see Lines 2-4. Note that the requirement $r \in \min(U)$ in the definition of T is mute as \mathcal{D} is bimolecular: $r \in U$ always holds and since \mathcal{D} is bimolecular, for all $c \in \min(U)$ we have $\|c\| \geq 2 = \|r\|$, and thus $r \in \min(U)$. Also, since \mathcal{D} is bimolecular, the case $\|p\| = 0$ in the definition of T is mute. The second (and final) phase is to iteratively augment M with the elements from $\min(U) \setminus (M_1 \cup M_2)$ as prescribed by Statements 5 and 6 of Lemma 2.

We show by induction that at Line 15, we have $M_{\text{it}} = \min(U) \cap \mathcal{C}_{=k}$ and $M = \min(U) \cap \mathcal{C}_{\leq k}$.

We first consider the basis case $k = 2$. Note that, by Lemma 2, $\min(U) \cap \mathcal{C}_{=2} = \min(U) \cap \mathcal{C}_{\leq 2}$ is obtained from $M_1 \cup M_2$ by adding all c' such that $c' \rightarrow^* c$ and $c \in M_1 \cup M_2$. Note that each such c' is minimal in U as $\|c'\| = 2$. This is accomplished in Lines 6-14.

We now consider the induction step. Let $k \geq 2$. Consider the set $X = \{c' \mid c' \rightarrow_\alpha c + B, \text{ for some } \alpha \in R, c \in \min(U) \cap \mathcal{C}_{=k}, B \in \Lambda, c'' \not\preceq c' \text{ for all } c'' \in \min(U) \cap \mathcal{C}_{\leq k}\}$, where we identify here $B \in \Lambda$ by the configuration b with $\|b\| = 1$ and $b(B) = 1$. Note that $X \subseteq U$. Since for all $c' \in X$, $\|c'\| = k + 1$ and $c' \in U$, we have that $c'' \not\preceq c'$ for all $c'' \in \min(U) \cap \mathcal{C}_{\leq k}$ if and only if $c'' \not\preceq c'$ for all $c'' \in \min(U)$. Hence $X \subseteq \min(U) \cap \mathcal{C}_{=k+1}$. The set X is computed in Lines 16-21. Now, by Statements 5 and 6 of Lemma 2, $\min(U) \cap \mathcal{C}_{=k+1}$ is obtained from X by adding the configurations c' such that $c' \rightarrow^* c$ with $c \in X$ and $c'' \not\preceq c'$ for all $c'' \in \min(U)$. Again, since $\|c'\| = k + 1$ and $c' \in U$, we have that $c'' \not\preceq c'$ for all $c'' \in \min(U)$ if and only if $c'' \not\preceq c'$ for all $c'' \in \min(U) \cap \mathcal{C}_{\leq k}$. These additional configurations c' are (again) computed in Lines 6-14.

The algorithm halts as by Dickson's Lemma (Proposition 4), $\min(U)$ is finite. \square

We now consider the time complexity of Algorithm 1.

Theorem 6 *Algorithm 1 takes $O(|R| \cdot |\Lambda| \cdot n \log^{|\Lambda| - \frac{1}{2}}(n))$ configuration comparisons to compute $\min(U)$, where $n = |\min(U)|$.*

Proof There are two inner loops. The first inner loop (at Lines 9-13) checks for every $c \in \min(U)$ and $\alpha \in R$, whether or not a $c' \rightarrow_\alpha c$ exists, and if such a c' exists, whether or not $c'' \not\preceq c'$ for all $c'' \in \min(U) \cap \mathcal{C}_{\leq \|c'\| - 1}$. The second inner loop (at Lines 17-21) checks for every $c \in \min(U)$, $\alpha \in R$, and $B \in \Lambda$, whether or not a $c' \rightarrow_\alpha c + B$ exists, and if such a c' exists, whether or not $c'' \not\preceq c'$ for all $c'' \in \min(U) \cap \mathcal{C}_{\leq \|c'\| - 1}$. Consequently, the second inner loop is dominant and has at most $n \cdot |R| \cdot |\Lambda|$ iterations. Just like as in Lemma 1 we use the k -fold tree $T_b(k)$ from [20] to store the vectors of M . Again, $k = |\Lambda|$ is the dimension of the vectors of M . Recall from Subsection 4.1 that it takes $O(\log^{k - \frac{1}{2}}(N))$ configuration comparisons, where $N = |M|$ is the number of elements in $T_b(k)$, to determine if a vector v is such that $w \not\preceq v$ for all vectors w in $T_b(k)$. Thus, we require

Algorithm 1 Generate the set M of minimal output unstable configurations of an \mathfrak{o} -stable bimolecular CRD $\mathcal{D} = (A, R, \Sigma, \Upsilon)$

```

1: procedure GENMINUNSTABLE( $\mathcal{D}$ )
2:    $T \leftarrow \{r \mid (r, p) \in R, \text{ and either } \|p\| = 0 \text{ or } \Upsilon(A) \neq \Upsilon(B) \text{ for some } A, B \in \Lambda \text{ with } r(A) \neq 0 \neq p(B)\}$ 
3:    $M_{\text{it}} \leftarrow \{AB \mid A, B \in \Lambda, \Upsilon(A) \neq \Upsilon(B)\} \cup T$ 
4:    $M \leftarrow M_{\text{it}}$ 
5:   while  $M_{\text{it}} \neq \emptyset$  do
6:      $M_{\text{new}} \leftarrow M_{\text{it}}$ 
7:     while  $M_{\text{new}} \neq \emptyset$  do
8:        $M_{\text{old}}, M_{\text{new}} \leftarrow M_{\text{new}}, \emptyset$ 
9:       for all  $c \in M_{\text{old}}, \alpha \in R$  do
10:        if  $\exists c'$  with  $c' \rightarrow_{\alpha} c$  and  $c'' \not\preceq c'$  for all  $c'' \in M$  then
11:           $M_{\text{new}}, M_{\text{it}}, M \leftarrow M_{\text{new}} \cup \{c'\}, M_{\text{it}} \cup \{c'\}, M \cup \{c'\}$ 
12:        end if
13:      end for
14:    end while
15:     $\triangleright$  At this point  $M = \min(U) \cap \mathcal{C}_{\leq k}$  and  $M_{\text{it}} = \min(U) \cap \mathcal{C}_{=k}$ .
16:     $M_{\text{itold}}, M_{\text{it}} \leftarrow M_{\text{it}}, \emptyset$ 
17:    for all  $c \in M_{\text{itold}}, \alpha \in R, B \in \Lambda$  do
18:      if  $\exists c'$  with  $c' \rightarrow_{\alpha} c + B$  and  $c'' \not\preceq c'$  for all  $c'' \in M$  then
19:         $M_{\text{it}}, M \leftarrow M_{\text{it}} \cup \{c'\}, M \cup \{c'\}$ 
20:      end if
21:    end for
22:  end while
23:  return  $M$ 
24: end procedure

```

$O(|R| \cdot |\Lambda| \cdot n \log^{|\Lambda| - \frac{1}{2}}(n))$ configuration comparisons. Inserting a vector in $T_b(k)$ takes $O(\log^{k - \frac{1}{2}}(N))$ configuration comparisons and so this step does not dominate. Consequently, we obtain the stated complexity. \square

Note that to obtain the time complexity of Algorithm 1 one multiplies the expression of Theorem 6 by the time complexity of comparing two vectors of dimension $|\Lambda|$.

We now obtain the following corollary to Theorem 6 and Lemma 1.

Corollary 1 *Given an \mathfrak{o} -stable bimolecular CRD \mathcal{D} , checking \mathfrak{o} -stability for a set S of configurations using Algorithm 1 takes $O((|R| \cdot |\Lambda| \cdot n + |S|) \log^{|\Lambda| - \frac{1}{2}}(n))$ configuration comparisons, with $n = |\min(U)|$.*

In view of Theorem 6, it would be interesting to obtain an upper bound on $|\min(U)|$. In fact, it is perhaps reasonable to view $|\min(U)|$ as a measure for the “complexity” of the underlying \mathfrak{o} -stable CRD \mathcal{D} . The set $\min(U)$ is an antichain, as any two elements of $\min(U)$ are incomparable (i.e., if $x, y \in \min(U)$ are distinct, then $x \not\preceq y$ and $y \not\preceq x$). In general, antichains can be arbitrary large for fixed Λ : for example, for every $k \in \mathbb{N}$, $\mathcal{C}_{=k}$ is an antichain with $|\mathcal{C}_{=k}| = \binom{|\Lambda|}{k} > k$ if $|\Lambda| \geq 2$. Note however that, by Lemma 2, if $x \in \min(U)$ with $\|x\| = k$, then for every $l \in \{2, \dots, k-1\}$ there is a $y \in \min(U)$ with $\|y\| = l$. Thus, in particular, $\min(U)$ (for some \mathfrak{o} -stable bimolecular CRD \mathcal{D}) cannot be equal to $\mathcal{C}_{=k}$ for any $k \geq 3$. We expect, but it would be interesting to confirm, that the existence of these “small” configurations in $\min(U)$ significantly restricts the cardinality of the antichain $\min(U)$.

In view of Lemma 2, we notice that Algorithm 1 works unchanged for the slightly larger class of nonincreasing o-stable CRDs \mathcal{D} with $\|r\| = 2$ for all $(r, p) \in R$. At a significant expense of computational efficiency, Algorithm 1 can even be extended to allow for arbitrary nonincreasing o-stable CRDs \mathcal{D} . The main issue in extending Algorithm 1 to this larger class of CRDs is that Statement 6 of Lemma 2 is then violated: if $c \leftrightarrow c'$, then $\|c'\|$ can be more than one smaller than $\|c\|$. Indeed, we have $c' + b = c - r + p$, for some $b \leq p$ with $b \neq p$. Thus $c' = c - (r - p + b)$ and so in the worst case $\|c'\|$ is equal to $\|c\| - (\|r\| - 1)$. As a result, the computational complexity of the dominant second inner loop (at Lines 17-21) is significantly increased. Another, smaller issue in extending Algorithm 1 is that the elements of the set T may not all be of size 2. Therefore, the elements of T need to be added to M at the right time (according to their sizes) within the outer while loop instead of at the beginning of the algorithm.

We remark that there is no obvious way to extend Algorithm 1 for arbitrary (i.e., also increasing) o-stable CRDs. In particular, it is not clear how to generate the elements of $\min(U)$ in order of their size (as used in the proof of Theorem 5) since increasing o-stable CRDs may generate large minimal configurations from small minimal configurations. In fact, it is not even clear if it is decidable, given an arbitrary o-stable CRD \mathcal{D} and a finite set M of configurations, whether or not $M = \min(U)$.

5 Discussion

Using the semilinearity proof of [1], we found that the class of t-stable CRDs have equal expressive power as the larger class of o-stable CRDs. Also, we shown a subtle difference in expressive power between CRDs and CRDs with leaders. Then, we considered the problem of determining whether or not a given configuration c is output stable. In particular, we have shown that the set $\min(U)$ of minimal output unstable configurations may be efficiently computed provided that we restrict to the class of o-stable bimolecular CRDs. Given $\min(U)$ it is straightforward to verify whether or not a given configuration c is output stable.

Various questions regarding the computational complexity of CRDs are open. For example, is it decidable whether or not a given CRD is o-stable, or whether or not it is t-stable? Also, likely some “bridges” between the domains of CRDs (functioning as acceptors/deciders) and Petri nets (functioning as generators) remain to be discovered. For example, the semilinear sets are precisely the sets of reachable markings of weakly persistent Petri nets [21]. This suggests a possible link between the notions of weak persistence (from the domain of Petri nets) and stable deciders (from the domain of CRDs).

Acknowledgements

We thank Jan Van den Bussche for interesting discussions on CRNs and for useful comments on an earlier version of this paper. We also thank the anonymous reviewers for useful comments on the paper. R.B. is a postdoctoral fellow of the Research Foundation – Flanders (FWO).

References

1. D. Angluin, J. Aspnes, Z. Diamadi, M. J. Fischer, and R. Peralta. Computation in networks of passively mobile finite-state sensors. *Distributed Computing*, 18(4):235–253, 2006.
2. D. Angluin, J. Aspnes, and D. Eisenstat. Stably computable predicates are semilinear. In E. Ruppert and D. Malkhi, editors, *Proceedings of the 25th Annual ACM Symposium on Principles of Distributed Computing (PODC 2006)*, pages 292–299. ACM, 2006.
3. D. Angluin, J. Aspnes, D. Eisenstat, and E. Ruppert. The computational power of population protocols. *Distributed Computing*, 20(4):279–304, 2007.
4. J. Aspnes and E. Ruppert. An introduction to population protocols. *Bulletin of the EATCS*, 93:98–117, 2007.
5. R. Brijder. Output stability and semilinear sets in chemical reaction networks and deciders. In S. Murata and S. Kobayashi, editors, *Proceedings of the 20th International Conference on DNA Computing and Molecular Programming (DNA 20)*, volume 8727 of *Lecture Notes in Computer Science*, pages 100–113. Springer, 2014.
6. H.-L. Chen, D. Doty, and D. Soloveichik. Deterministic function computation with chemical reaction networks. In D. Stefanovic and A. J. Turberfield, editors, *Proceedings of the 18th International Conference on DNA Computing and Molecular Programming (DNA 18)*, volume 7433 of *Lecture Notes in Computer Science*, pages 25–42. Springer, 2012.
7. H.-L. Chen, D. Doty, and D. Soloveichik. Rate-independent computation in continuous chemical reaction networks. In M. Naor, editor, *Innovations in Theoretical Computer Science (ITCS'14)*, pages 313–326. ACM, 2014.
8. M. Cook, D. Soloveichik, E. Winfree, and J. Bruck. Programmability of chemical reaction networks. In A. Condon, D. Harel, J. N. Kok, A. Salomaa, and E. Winfree, editors, *Algorithmic Bioprocesses*, Natural Computing Series, pages 543–584. Springer Berlin Heidelberg, 2009.
9. L. E. Dickson. Finiteness of the odd perfect and primitive abundant numbers with n distinct prime factors. *American Journal of Mathematics*, 35:413–422, 1913.
10. D. Doty. Timing in chemical reaction networks. In C. Chekuri, editor, *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2014)*, pages 772–784. SIAM, 2014.
11. D. Doty and M. Hajiaghayi. Leaderless deterministic chemical reaction networks. In D. Soloveichik and B. Yurke, editors, *Proceedings of the 19th International Conference on DNA Computing and Molecular Programming (DNA 19)*, volume 8141 of *Lecture Notes in Computer Science*, pages 46–60. Springer, 2013.
12. R. M. Karp and R. E. Miller. Parallel program schemata. *Journal of Computer and System Sciences*, 3(2):147–195, 1969.
13. J. Leroux. Vector addition systems reachability problem (a simpler solution). In A. Voronkov, editor, *Proceedings of the Alan Turing Centenary Conference (Turing-100)*, volume 10 of *EPiC Series*, pages 214–228, 2012.
14. E. W. Mayr. An algorithm for the general Petri net reachability problem. In *Proceedings of the 13th Annual ACM Symposium on Theory of Computing (STOC '81)*, pages 238–246, New York, NY, USA, 1981. ACM.
15. J. L. Peterson. Petri nets. *ACM Computing Surveys*, 9(3):223–252, 1977.
16. G. D. Plotkin. A calculus of chemical systems. In V. Tannen, L. Wong, L. Libkin, W. Fan, W.-C. Tan, and M. Fourman, editors, *In Search of Elegance in the Theory and Practice of Computation*, volume 8000 of *Lecture Notes in Computer Science*, pages 445–465. Springer Berlin Heidelberg, 2013.
17. D. Soloveichik, M. Cook, E. Winfree, and J. Bruck. Computation with finite stochastic chemical reaction networks. *Natural Computing*, 7(4):615–633, 2008.
18. D. Soloveichik, G. Seelig, and E. Winfree. DNA as a universal substrate for chemical kinetics. *Proceedings of the National Academy of Sciences*, 107(12):5393–5398, 2010.
19. R. P. Stanley. *Enumerative Combinatorics, Volume 1, 2nd Edition*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2011.
20. D. E. Willard. New data structures for orthogonal range queries. *SIAM Journal on Computing*, 14(1):232–253, 1985.
21. H. Yamasaki. On weak persistency of Petri nets. *Information Processing Letters*, 13(3):94–97, 1981.