Brief Announcement: Reachability in Deletion-Only Chemical Reaction Networks

Bin Fu
University of Texas Rio Grande Valley, Edinburg, TX, USA

Timothy Gomez
Massachusetts Institute of Technology, Cambridge, MA, USA

Ryan Knobel
University of Texas Rio Grande Valley, Edinburg, TX, USA

Austin Luchsinger
University of Texas Rio Grande Valley, Edinburg, TX, USA

Aiden Massie
University of Texas Rio Grande Valley, Edinburg, TX, USA

Marco Rodriguez
Massachusetts Institute of Technology, Cambridge, MA, USA

Adrian Salinas
University of Texas Rio Grande Valley, Edinburg, TX, USA

Robert Schweller
University of Texas Rio Grande Valley, Edinburg, TX, USA

Tim Wylie
University of Texas Rio Grande Valley, Edinburg, TX, USA

Abstract

For general discrete Chemical Reaction Networks (CRNs), the fundamental problem of reachability – the question of whether a target configuration can be produced from a given initial configuration – was recently shown to be Ackermann-complete. However, many open questions remain about which features of the CRN model drive this complexity. We study a restricted class of CRNs with *void rules*, reactions that only decrease species counts. We further examine this regime in the motivated model of step CRNs, which allow additional species to be introduced in discrete stages. With and without steps, we characterize the complexity of the reachability problem for CRNs with void rules. We show that, without steps, reachability remains polynomial-time solvable for bimolecular systems but becomes NP-complete for larger reactions. Conversely, with just a single step, reachability becomes NP-complete even for bimolecular systems. Beyond what is contained in this brief announcement, we also investigate optimization variants of reachability, provide approximation results for maximizing species deletion, establish ETH-based lower bounds for NP-complete cases, and prove hardness for counting reaction sequences.

2012 ACM Subject Classification Theory of computation → Models of computation; Theory of computation → Problems, reductions and completeness

Keywords and phrases CRN, Chemical Reaction Network, Reachability, Void Reactions

Digital Object Identifier 10.4230/LIPIcs.SAND.2025.23

Funding This research was supported in part by National Science Foundation Grant CCF-2329918.

1 Introduction

We study reachability [8, 9, 10] in a restricted class of Chemical Reaction Networks [4, 5], that considers deletion-only systems [1, 2] (called *void rules*) where reactions only ever consume species and reduce the size of the system. Despite the closure of this problem for general

© Bin Fu, Timothy Gomez, Ryan Knobel, Austin Luchsinger, Aiden Massie, Marco Rodriguez, Adrian Salinas, Robert Schweller, and Tim Wylie;
licensed under Creative Commons License CC-BY 4.0

Reachability Complexity	
Void Rules	**Basic CRNs**
$(2,0)$	$O(\lvert\Lambda\rvert^2\lvert\Gamma\rvert\log(\lvert\Lambda\rvert))$ [1]
$(2,1)$	NL [Thm. 5]
$(2,0)$, $(2,1)$	$O(\lvert\Lambda\rvert^2\lvert\Gamma\rvert\log(\lvert\Lambda\rvert))$ [Thm. 7]
$(k,k-1)^+$	$O(\lvert\Lambda\rvert^2\lvert\Gamma\rvert)$ [Thm. 9]
$(k\geq 3, g\leq k-2)$	NPC [Thm. 11]

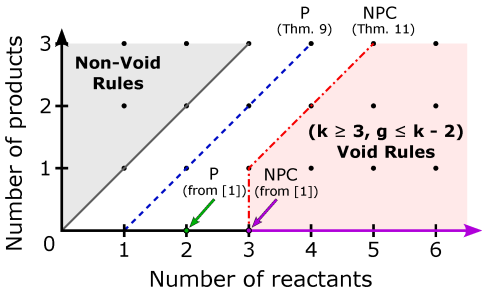

Figure 1 (Left) A table of reachability results for various void rule sizes. (Right) A plot depicting the complete characterization of reachability complexity for basic CRNs with uniform-size void rules.

CRNs [6, 7], and in fact due to how complicated these systems can be, there is a natural motivation to explore reachability for these more restricted systems. This limited class of systems permits tractable and intractable problems, placing it along an interesting complexity boundary. We also consider *step* CRNs, an experimentally motivated model introduced in [2, 3], where species are added in discrete stages to reflect the progressive addition of molecules in laboratory protocols [11, 12]. In this paper we show that adding even a single step drastically changes the computational complexity of the reachability problem, moving from P (and even NL) to NP-complete in many cases.

In this paper we present a number of results for the various types of deletion-only interaction rules that, when taken together, yield (1) a complete characterization of reachability complexity for uniform-size void rules in standard CRNs, (2) a nearly complete characterization of reachability for arbitrary combinations of void rule sizes in standard CRNs, and (3) a complete characterization of reachability for void rules in step CRNs. These results are outlined in Figure 1. In cases where reachability is NP-complete, we derive lower bounds under the Exponential Time Hypothesis and establish nearly matching upper bounds, though these are not included in this brief announcement. Also deferred to the full version are results on optimization and counting variants of reachability, including approximation algorithms for species deletion and hardness results for counting reaction sequences.

2 Preliminaries

A discrete Chemical Reaction Network is a pair (Λ,Γ) that consists of a finite set of species Λ and a finite set of reactions Γ that dictate how species interact with one another. A configuration is a vector C over Λ specifying the count of each species. A reaction $R\in\Gamma$ is an ordered pair of configurations (R_r,R_p) representing reactants and products. A reaction is *applicable* to a configuration C if C has sufficient species to satisfy R_r, and applying R results in a new configuration $C'=C-R_r+R_p$. A configuration is called *terminal* with respect to a CRN (Λ,Γ) if no rule R can be applied to it. We denote the set of reachable configurations (from initial configuration I) of a CRN as $REACH_{I,\Lambda,\Gamma}$. We define the subset of reachable configurations that are terminal as $TERM_{I,\Lambda,\Gamma}$.

A *void* rule is any rule that does not create any new copies of any species types (only deletes). Formally, $R=(R_r,R_p)$ is a *void* rule if $R_p[i]\leq R_r[i]$ for all species i, with at least one strict inequality. Thus, a void rule either has no products or has products that are a subset of its reactants (in which case these products are termed *catalysts*). We classify void rules by their size: a rule has size (i,j) if the total number of reactants is i and the total number of products is j. A reaction is trimolecular if $i=3$, bimolecular if $i=2$, and unimolecular if $i=1$.

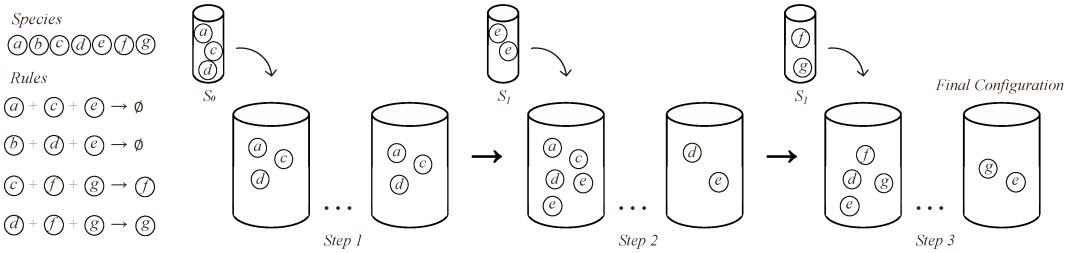

Figure 2 An example step CRN system. The test tubes show the species added at each step and the system with those elements added. The CRN species and void rule-set are shown on the left.

A step CRN is a CRN augmented with a finite sequence of *steps*, where after reaching a terminal configuration at each step, additional species are added to the system. Formally, a step CRN of k steps is an ordered pair $((\Lambda, \Gamma), (s_0, s_1, s_2, \ldots s_{k-1}))$, where the first element of the pair is a normal CRN (Λ, Γ), and the second is a sequence of configuration vectors denoting how many copies of each species type to add after each step. Figure 2 illustrates a simple step CRN system. Given a step CRN, REACH_1 is the set of configurations reachable from the initial configuration s_0, and $\text{TERM}_1 \subseteq \text{REACH}_1$ is the subset of terminal configurations. For each subsequent step i, REACH_i is defined as the union of configurations reachable from s_0 by adding s_{i-1} to each terminal configuration in TERM_{i-1}. Similarly, TERM_i is the subset of terminal configurations within REACH_i. The set of reachable configurations for a k-step CRN is the set REACH_k, and the set of terminal configurations is TERM_k. A classical CRN can be represented as a step CRN with $k = 1$ steps and an initial configuration of $I = s_0$.

The primary computational problem studied in this paper is *reachability*. Informally, reachability asks if a given initial configuration A can be turned into a target configuration B by applying a sequence of rules from the given CRN. The precise problem statement is given below.

▶ **Definition 1** (Reachability Problem). *Given an initial configuration A, a destination (target) configuration B, and a step CRN $((\Lambda, \Gamma), (s_0 = A, s_1, s_2, \ldots s_{k-1}))$, determine if $B \in REACH_k$, i.e., is configuration B reachable for the given step CRN. In the case of basic CRNs, this simplifies to: given configurations A and B, and basic CRN (Λ, Γ), determine if $B \in REACH_{A,\Lambda,\Gamma}$.*

We initiate our study of deletion-only systems by observing that reachability stays within the class NP with only void rules, even with step-CRNs.

▶ **Theorem 2.** *The reachability problem for void rule step-CRNs is in NP.*

3 Bimolecular Rules

Our first collection of results focus on bimolecular systems with either all size $(2, 0)$ rules (non-catalytic) or $(2, 1)$ rules (catalytic). Recently, $(2, 0)$ rule reachability was proven to be polynomial [1]. For $(2, 1)$, we present a polynomial-time algorithm for reachability. In contrast, we show that in either scenario the problem becomes NP-complete with the addition of a second step. In Section 3.2, we consider the same problems for CRNs that use both $(2, 0)$ and $(2, 1)$ rules together.

3.1 Uniform Bimolecular Rules: $(2, 0)$ or $(2, 1)$

Bimolecular Void Rules Without Catalysts: $(2, 0)$

In [1], the authors proved that reachability in a CRN system with only $(2, 0)$ rules is in P by reducing from the b-matching problem, which is a generalization of matching. We show that the same problem in step CRNs (with just two steps) becomes NP-complete via a reduction from the graph 3-colorability (3-COL) problem.

▶ **Theorem 3.** *Reachability for basic CRNs with binary encoded species with only rules of size (2, 0) is solvable in* $O(|\Lambda|^2 \log(|\Lambda|)(|\Gamma| + |\Lambda| \log(|\Lambda|)))$ *time [1].*

▶ **Theorem 4.** *Reachability for 2-step CRNs with only rules of size (2, 0) is NP-complete, even for unary encoded species counts.*

Bimolecular Void Rules with Catalysts: $(2, 1)$

We go on to show that reachability for size $(2, 1)$ void rules is solvable nondeterministically using only logarithmic space (the class NL) for standard CRNs, yet it immediately becomes NP-complete when a second step is allowed.

▶ **Theorem 5.** *Reachability for basic CRNs with only (2, 1) void rules is in NL.*

▶ **Theorem 6.** *Reachability for 2-step CRNs with only (2, 1) void rules is NP-complete, even for unary encoded species counts.*

3.2 Mixed Bimolecular Rules: $(2, 0)$ and $(2, 1)$

We next consider CRNs that allow a mix of catalytic (2,1) and non-catalytic (2,0) void rules. Despite this generalization, reachability remains tractable in the basic (1-step) model. We prove that reachability for such systems reduces to a generalized form of matching (perfect b-matching) which can be solved in polynomial time.

▶ **Theorem 7.** *Reachability for basic CRNs with both (2,0) and (2,1) void rules is solvable in* $O(|\Lambda|^2 \log(|\Lambda|)(|\Gamma| + |\Lambda| \log(|\Lambda|)))$ *time.*

4 Larger Void Rules

Our next results involve CRNs with reactions that require more than two reactants. If a system's rules have all but one reactant serving as catalysts (i.e., $(k, k-1)$ void rules), then reachability remains polynomial-time solvable. In contrast, reachability for systems with any other form of void rule (with 3 or more reactants) becomes NP-complete.

4.1 Mixed, Mostly-Catalytic Void Rules: $(\mathbf{k}, \mathbf{k-1})^+$

We provide a polynomial-time dynamic programming algorithm to decide reachability for CRNs that use mostly-catalytic void rules of the form $(k, k-1)$. We further argue that reachability remains in P, even for CRNs that use a combination of various size $(k, k-1)$. For simplicity, we refer to void rules of sizes $(k_1, k_1-1), \ldots, (k_b, k_b-1)$, where all $k_i \in \mathbb{N}$, as $(k, k-1)^+$, meaning there is one or more rule of this type.

▶ **Theorem 8.** *Reachability for basic CRNs with only void rules of size* $(k, k-1)$ *is solvable in* $O(|\Lambda|^2|\Gamma|)$.

▶ **Theorem 9.** *Reachability for basic CRNs with only void rules of size* $(k, k-1)^+$ *is solvable in* $O(|\Lambda|^2|\Gamma|)$.

▶ **Corollary 10.** *Reachability for 2-step CRNs with only void rules of size* $(k, k-1)^+$ *is NP-complete, even for unary encoded species counts.*

4.2 Uniform Large Void Rules: $(\mathbf{k \geq 3, g \leq k-2})$

Here we show that reachability for CRNs using any void rules with at least three reactants, and that removes at least two species, becomes NP-complete. To achieve this result, we show that reachability is NP-complete for CRNs using only $(3, 1)$ void rules via a reduction from the Hamiltonian path problem for directed graphs. Since it was previously shown that reachability is NP-complete for CRNs using only $(3, 0)$ rules [1], this implies Corollary 12.

▶ **Theorem 11.** *Reachability for CRNs with only void rules of size* $(3, 1)$ *is NP-complete, even for unary encoded species counts.*

▶ **Corollary 12.** *Reachability for CRNs with only void rules of size* $(k \geq 3, g \leq k-2)$ $(k, g \in \mathbb{N})$ *is NP-complete, even for unary encoded species counts.*

References

1. Robert M. Alaniz, Bin Fu, Timothy Gomez, Elise Grizzell, Andrew Rodriguez, Robert Schweller, and Tim Wylie. Reachability in restricted chemical reaction networks, 2022. doi:10.48550/arXiv.2211.12603.
2. Rachel Anderson, Alberto Avila, Bin Fu, Timothy Gomez, Elise Grizzell, Aiden Massie, Gourab Mukhopadhyay, Adrian Salinas, Robert Schweller, Evan Tomai, and Tim Wylie. Computing threshold circuits with void reactions in step chemical reaction networks. In *10th conference on Machines, Computations and Universality (MCU 2024)*, 2024.
3. Rachel Anderson, Bin Fu, Aiden Massie, Gourab Mukhopadhyay, Adrian Salinas, Robert Schweller, Evan Tomai, and Tim Wylie. Computing threshold circuits with bimolecular void reactions in step chemical reaction networks. In *International Conference on Unconventional Computation and Natural Computation*, pages 253–268. Springer, 2024. doi:10.1007/978-3-031-63742-1_18.
4. Rutherford Aris. Prolegomena to the rational analysis of systems of chemical reactions. *Archive for Rational Mechanics and Analysis*, 19(2):81–99, January 1965. doi:10.1007/BF00282276.
5. Rutherford Aris. Prolegomena to the rational analysis of systems of chemical reactions ii. some addenda. *Archive for Rational Mechanics and Analysis*, 27(5):356–364, January 1968.
6. Wojciech Czerwiński and Łukasz Orlikowski. Reachability in vector addition systems is Ackermann-complete. In *62nd Annual Symposium on Foundations of Computer Science*, FOCS'21, pages 1229–1240. IEEE, 2021.
7. Jérôme Leroux. The reachability problem for Petri nets is not primitive recursive. In *62nd Annual Symposium on Foundations of Computer Science*, FOCS'21. IEEE, 2021.
8. Richard J. Lipton. The reachability problem requires exponential space. Technical Report 62, Yale University, 1976.
9. Ernst W. Mayr. An algorithm for the general Petri net reachability problem. In *Proc. of the* 13^{th} *Annual ACM Symposium on Theory of Computing*, STOC'81, pages 238–246. ACM, 1981. doi:10.1145/800076.802477.
10. George S Sacerdote and Richard L Tenney. The decidability of the reachability problem for vector addition systems (preliminary version). In *Proceedings of the ninth annual ACM symposium on Theory of computing*, pages 61–76, 1977. doi:10.1145/800105.803396.

11 David Soloveichik, Matthew Cook, Erik Winfree, and Jehoshua Bruck. Computation with finite stochastic chemical reaction networks. *natural computing*, 7(4):615–633, 2008. doi:10.1007/S11047-008-9067-Y.

12 Boya Wang, Chris Thachuk, Andrew D Ellington, Erik Winfree, and David Soloveichik. Effective design principles for leakless strand displacement systems. *Proceedings of the National Academy of Sciences*, 115(52):E12182–E12191, 2018.