

Computational Complexity of Atomic Chemical Reaction Networks

David Doty
doty@ucdavis.edu

Shaopeng Zhu
spzhu@ucdavis.edu

February 21, 2017

Abstract

Informally, a chemical reaction network is “atomic” if each reaction may be interpreted as the rearrangement of indivisible units of matter. There are several reasonable definitions formalizing this idea. We investigate the computational complexity of deciding whether a given network is atomic according to each of these definitions.

Our first definition, *primitive atomic*, which requires each reaction to preserve the total number of atoms, is shown to be equivalent to mass conservation. Since it is known that it can be decided in polynomial time whether a given chemical reaction network is mass-conserving [30], the equivalence gives an efficient algorithm to decide primitive atomicity.

Another definition, *subset atomic*, further requires that all atoms are species. We show that deciding whether a given network is subset atomic is in NP, and the problem “is a network subset atomic with respect to a given atom set” is strongly NP-complete.

A third definition, *reachably atomic*, studied by Adleman, Gopalkrishnan et al. [1,22], further requires that each species has a sequence of reactions splitting it into its constituent atoms. We show that there is a polynomial-time algorithm to decide whether a given network is reachably atomic, improving upon the result of Adleman et al. that the problem is decidable. We show that the reachability problem for reachably atomic networks is PSPACE-complete.

Finally, we demonstrate equivalence relationships between our definitions and some special cases of another existing definition of atomicity due to Gnacadja [21].

1 Introduction

A *chemical reaction network* is a set of reactions such as $A + B \rightleftharpoons C$ and $X \rightarrow 2Y$, intended to model molecular species that interact, possibly combining or splitting in the process. The model is syntactically equivalent to Petri nets: molecules correspond to “tokens”, species correspond to “places”, reactions correspond to “transitions”, and configurations correspond to “markings”. Indeed, the study of chemical reaction networks has profited from this connection [3,5,17]. However, due to their different modeling goals (concurrent systems and well-mixed chemistry, respectively), sometimes different questions are germane in each model.

For 150 years [23], the model has been a popular language for describing natural chemicals that react in a well-mixed solution. Several recent wet-lab experiments demonstrate the systematic engineering of *custom-designed* chemical reactions [10,31,34,37,38,43,47], and it is known that in theory *any* set of reactions can be implemented by synthetic DNA complexes [46]. Thus chemical reaction networks are now equally appropriate as a *programming* language that can be compiled into real chemicals. With advances in synthetic biology heralding a new era of sophisticated biomolecular

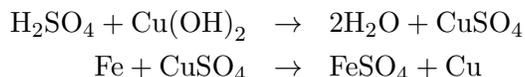
engineering, chemical reaction networks will gain prominence as a natural high-level language for designing molecular control circuitry.

There has been a flurry of recent progress in understanding the ability of chemical reaction networks to carry out computation: computing functions [2, 5, 7–9, 14–17, 39, 41, 45], as well as other computational tasks such as space- and energy-efficient search [49], signal processing [25, 40], linear I/O systems [33], machine learning [32], and even identifying function computation in existing biological chemical reaction networks [6]. These studies generally assume that any set of reactions is permissible, but not all are physically realistic. Consider, for example, the reaction $X \rightarrow 2X$, which appears to violate the law of conservation of mass. Typically such a reaction is a shorthand for a more realistic reaction such as $F + X \rightarrow 2X$, where F is an anonymous and plentiful source of “fuel” providing the necessary matter for the reaction to occur. The behavior of the two is approximately equal only when the number of executions of $X \rightarrow 2X$ is far below the supplied amount of F , and if F runs out then the two reactions behave completely differently. Thus, although $X \rightarrow 2X$ may be implemented approximately, to truly understand the long-term behavior of the system requires studying its more realistic implementation $F + X \rightarrow 2X$. A straightforward generalization of this “realism” constraint is that each chemical species S may be assigned a *mass* $m(S) \in \mathbb{R}^+$, where in each reaction the total mass of the reactants equals that of the products. Indeed, *conservative* Petri nets formalize this very idea [28, 30], and it is straightforward to decide algorithmically if a given network is conservative by reducing to a question of linear algebra.

The focus of this paper is a more stringent condition: that the network should be *atomic*, i.e., each reaction rearranges discrete, indivisible units (atoms), which may be of different noninterchangeable types.¹ (In contrast, mass conservation requires each reaction to rearrange a conserved quantity of continuous, generic “mass”.) We emphasize that this is not intended as a study of the atoms appearing in the periodic table of the elements. Instead, the goal is to model chemical systems whose reactions rearrange certain units, but never split, create, or destroy those units. For example, DNA strand displacement systems [46, 50] have individual DNA strands as indivisible components, and each reaction merely rearranges the secondary structure among the strands (i.e., which bases on the strands are hybridized to others).

In contrast to the idea of mass conservation, there is not a single “obviously correct” definition of what it means for a chemical reaction network to be atomic, as we will discuss. Furthermore, at least two inequivalent definitions exist in the literature [1, 21]. It is not the goal of this paper to identify a single correct definition. Instead, our goal is to evaluate the choices that must be made in formalizing a definition, to place existing and new definitions in this context to see how they relate to each other, and to study the computational complexity of deciding whether a given network is atomic. This is a step towards a more broad study of the computational abilities of “physically realistic” chemical reaction networks.

Consider the following system with 2 (real world) reactions for example:



If we ignore the change of oxidation number in the second reaction (which Chemical Reaction Networks will not be modelling anyway), we may as well just let the “atom” set– set of invariant

¹This usage of the term “atomic” is different from its usage in traditional areas like operating system or syntactic analysis, where an “atomic” execution is an uninterruptable unit of operation [44].

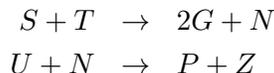
component in this reaction system – be

$$\Delta_1 = \{X := \text{H}, Y := \text{OH}, Z := \text{Cu}, U := \text{Fe}, W := \text{SO}_4\}$$

where we define all *indecomposable components* in the *specific* reaction system as atoms; and the let the molecule² set be:

$$M = \{S := \text{H}_2\text{SO}_4, T := \text{Cu}(\text{OH})_2, G := \text{H}_2\text{O}, N := \text{CuSO}_4, P := \text{FeSO}_4\}$$

and rewrite the reaction system as



and note that the collection of compositions (i.e., atomic makeups):

$$\begin{aligned} \{S &\mapsto \{2X, 1W\}, T \mapsto \{1Z, 2Y\}, G \mapsto \{1X, 1Y\}, \\ N &\mapsto \{1Z, 1W\}, P \mapsto \{1U, 1W\}, A \mapsto \{1A\} (\forall A \in \Delta_1)\} \end{aligned}$$

will maintain all the composition information in the two-reaction system. Instead of using the entire periodic table, *we now shrink the size of atom set into 5*. Note, also, that we may alternatively define Δ in a more “canonical” way—the subset of elements in the periodic table that are actually used in the system, which in this case will be $\Delta_2 = \{\text{H}, \text{O}, \text{S}, \text{Cu}, \text{Fe}\}$, and the corresponding set of compositions shall be:

$$\begin{aligned} \{S &\mapsto \{2\text{H}, 4\text{O}, 1\text{S}\}, T \mapsto \{2\text{H}, 2\text{O}, 1\text{Cu}\}, G \mapsto \{2\text{H}, 1\text{O}\}, \\ N &\mapsto \{4\text{O}, 1\text{S}, 1\text{Cu}\}, P \mapsto \{4\text{O}, 1\text{S}, 1\text{Fe}\}, A \mapsto \{1A\} (\forall A \in \Delta_2)\} \end{aligned}$$

which shows that *we may have multiple choices of “atoms” set* and corresponding set of compositions when modelling the chemical reactions.

1.1 Summary of Results and Connection with Existing Work

The most directly related previous work is that of Adleman, Gopalkrishnan, Huang, Moisset, and Reishus [1] and of Gnacadja [21], which we now discuss in conjunction with our results.

We identify two fundamental questions to be made in formalizing a definition of an “atomic” chemical reaction network:

1. Are atoms required to be species? (for example, if the only reaction is $2\text{H}_2 + \text{O}_2 \rightleftharpoons 2\text{H}_2\text{O}$; then H and O are atoms but not species that appear in a reaction)
2. Is each species required to be separable into its constituent atoms via reactions?

A negative answer to (1) implies a negative answer to (2). (If some atom is not a species, then it cannot be the product of a reaction.) Thus there are three non-degenerate answers to the above two questions: no/no, yes/no, and yes/yes. We respectively call these *primitive atomic*, *subset*

²In this abstract model when we call H_2SO_4 and CuSO_4 “molecules”, we also ignore the caveat that substance like H_2SO_4 and CuSO_4 does not exist as an actual molecule (in the chemical sense) in water solution.

atomic, and *reachably atomic*, defined formally in Section 3. Intuitively, a network is *primitive atomic* if each species may be interpreted as composed of one or more atoms, which themselves are not considered species (a species can be composed of just a single atom, but they will have different “names”). More formally, if Λ is the set of species, there is a set Δ of atoms, such that each species $S \in \Lambda$ has an *atomic composition* $\mathbf{d}_S \in \mathbb{N}^\Delta \setminus \{\mathbf{0}\}$ describing the atoms that constitute S , such that each reaction preserves the atoms. A network is *subset atomic* if it is primitive atomic and the atoms are themselves considered species; i.e., if $\Delta \subseteq \Lambda$. A network is *reachably atomic* if it is subset atomic, and furthermore, for each species $S \in \Lambda$, there is a sequence of reactions, starting with a single copy of S , resulting in a configuration consisting only of atoms. (If each reaction conserves the atomic count, then this configuration must be unique and equal to the atomic composition of S .)

A long-standing open problem in the theory of chemical reaction networks is the global attractor conjecture [13,24], of which even the following special case remains open: is every network satisfying detailed balance *persistent*, i.e., if started with all positive concentrations, do concentrations stay bounded away from 0? Adleman, Gopalkrishnan, Huang, Moisset, and Reishus [1] defined reachably atomic chemical reaction networks and proved the global attractor conjecture holds for such networks. Gnacadja [21], citing similar goals, defined a notion of atomicity called “species composition” and showed a similar result. We establish links between our definitions and those of both [21] and [1]. We discuss related complexity issues in Sections 6 and 8. In particular, Adleman et al. [1] showed that it is decidable whether a given network is reachably atomic. This is not obvious since the condition of a species being separable into its constituent atoms via reactions appears to require an unbounded search. We improve this result, showing it is decidable in polynomial time.

As mentioned, Petri nets are syntactically equivalent to chemical reaction networks, and we exploit this connection earnestly. Mayr and Weihmann [30] proved that configuration reachability graphs for mass conserving chemical reaction networks (i.e., conservative Petri nets) are at most exponentially large in the size of the binary representation of the network, implying via Savitch’s theorem [42] a polynomial-space algorithm for deciding reachability in mass-conserving networks. We use these results in analyzing the complexity of reachability problems in reachably atomic chemical reaction networks in Section 6.

It is clear that any reasonable definition of atomicity should imply mass conservation: simply assign all atoms to have mass 1, noting that any reaction preserving the atoms necessarily preserves their total count. Perhaps surprisingly, the conditions of primitive atomic and mass-conserving are in fact equivalent, so it is decidable in polynomial time whether a network is primitive atomic and what is an atomic composition for each species. A key technical tool is Chubanov’s algorithm [11] for finding exact rational solutions to systems of linear equations with a strict positivity constraint.

Lastly, we note that there have been other models addressing different aspects of atomicity (not necessarily using the term “atomic”). They focus on features of chemical reaction networks not modeled in this paper. For example, Johnson, Dong and Winfree [26] study, using concept called *bisimulation*, the complexity of deciding whether one network \mathcal{N}_1 “implements” another \mathcal{N}_2 . The basic idea involves assigning species S in \mathcal{N}_1 to represent sets of species $\{A_1, \dots, A_k\}$ in \mathcal{N}_2 . Thus S may be intuitively thought of as being “composed” of A_1, \dots, A_k ; however, an allowed composition is \emptyset , to account for the fact that some species in \mathcal{N}_1 have the goal of mediating interactions between species in \mathcal{N}_2 without “representing” any of them. Molecules, on the other hand, are always composed of at least one atom.

3

³One may tend to think at first glance that, for a given chemical reaction network $\mathcal{C} = (\Lambda, R)$ atomic with respect

Mann, Nahar, Schnorr, Backofen, Stadler, and Flamm [29] studied a notion of atomicity based on the idea that shared electrons between atoms must be conserved. This model is more chemically detailed than ours, but also more limited, being unable to model higher-level notions of “shared subcomponents” such as DNA strands whose secondary structure, but not their primary structure, is altered by strand displacement reactions [46,50]. Tapia and Faed [48] developed software for finding atomic compositions in a similarly detailed setting motivated by specific biochemical experiments.

1.2 Structure

With the expectation to fully study the computational power as well as other computationally interesting properties of atomic chemical reaction networks in future, we decided to focus this paper on *different definitions of atomicity* and the *computational complexity for deciding atomicity* under each definition.

Section 2 provides some definitions and notations as background for subsequent sections.

Section 3 introduces the core concepts of this paper: *primitive atomicity*, *subset atomicity*, *reachable atomicity* and *reversibly-reachable atomicity*. Intuitively and informally, *primitive atomicity* requires that

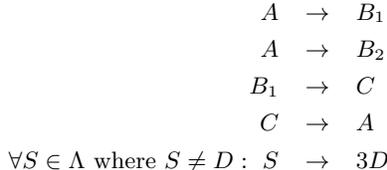
1. There exists a composition for each species into indecomposable components such that each reaction preserves the total count of each type of indecomposable components.
2. Each *composite species* can be decomposed into the linear combination of indecomposable components (and thus its composition should contain at least 2 units of indecomposable components), while each *elementary species* — which are indecomposable components that are themselves species, should be its own “composition”. And,
3. Each elementary species appears in the composition of at least one composite species; thus no elementary species is “redundant”.

to Δ via composition matrix $\mathbf{D} = (\mathbf{d}_S)_{S \in \Lambda}$, the mapping

$$\begin{aligned} \mathbf{d} : \Lambda &\rightarrow \mathbb{N}^\Lambda \\ S &\mapsto \mathbf{d}_S \end{aligned}$$

, as defined in our future chapters, gives an interpretation between $\mathcal{C} = (\Lambda, R)$ and $\mathcal{C}' = (\Delta, R')$ where $R' \subseteq \mathbb{N}^\Delta$ is the natural linear extension of $R \subseteq \mathbb{N}^\Lambda$ — that is, for all $(\mathbf{r}, \mathbf{p}) \in R$, construct $(\mathbf{r}', \mathbf{p}') \in R'$ s.t. $\mathbf{r}' = \sum_{S \in [\mathbf{r}]} \mathbf{d}_S$, $\mathbf{p}' = \sum_{S \in [\mathbf{p}]} \mathbf{d}_S$. For example, when $\mathcal{C} = (\{S_1, S_2\}, \{S_1 \rightarrow 2S_2\})$, $\mathcal{C}' = (\{S_2\}, \{2S_2 \rightarrow 2S_2\})$, and the interpretation would be $m = \{(S_1, 2S_2), (S_2, S_2)\}$. (To avoid confusion one may also relabel the species in \mathcal{C}' : for example, in the previous case, $\mathcal{C}' = (\{S'_2\}, \{2S'_2 \rightarrow 2S'_2\})$ and $m = \{(S_1, 2S'_2), (S_2, S'_2)\}$).

But this reduction does not necessarily work, as the correct interpretation defined in [26] requires “equivalence of trajectories”. Consider a slightly modified example from [26]: let $\mathcal{C} = \{\Lambda, R\}$ where $\Lambda = \{A, B_1, B_2, C, D\}$ and $R =$



By our conjecture above, this would give an interpretation from \mathcal{C} to \mathcal{C}' , where $\mathcal{C}' = (\{D'\}, \{3D' \rightarrow 3D'\})$ and $m = \{(A, 3D'), (B_1, 3D'), (B_2, 3D'), (C, 3D'), (D, D')\}$. Yet this is NOT a correct interpretation, for, as [26] pointed out, \mathcal{C} is subject to deadlock yet \mathcal{C}' is not.

It is called “primitive” because it captures our primitive sense of what properties a chemical reaction network with indecomposable components (“atoms”) should have. *Subset atomicity*, on the other hand, requires all the above as well as an additional condition: all indecomposable components are themselves elementary species. As mentioned in previous sections, *reachable-atomicity* further requires that each composite species can be actually decomposed into elementary species via sequences of reactions. Lastly, *reversibly-reachable atomicity* requires in addition to all that are true for *reachably atomic* networks, plus that each composite species can be constructed from its atomic compositions via sequences of reactions. It will be shown in subsequent sections that the problems of deciding different atomicity type belong to different computation complexities.

Section 4 proves the equivalence between the concept of mass conservation and primitive atomicity. Intuitively, when the set of atoms might be defined independently from the set of species, one may translate between “a mass assignment to each species that preserves the total mass” and “an assignment of number of *same type of atoms* that preserves the total count of atoms” using some algebraic manipulation of the given assignment.

Section 5 attends to the complexity of deciding whether a chemical reaction network is subset atomic. We define two languages, SUBSET-FIXED-ATOMIC and SUBSET-ATOMIC, the first describing chemical reaction networks paired with a *given* atom set with respect to which it is indeed subset atomic, and the second describing subset atomic chemical reaction networks without a given atom set. We proved both languages in NP, with the first one NP-hard. We conjecture that the second is also NP-hard.

Proof that both languages are in NP is based on a reduction to the Integer-Programming problem, which is itself in NP. The reduction is necessary because no obvious short witness for SUBSET-FIXED-ATOMIC and SUBSET-ATOMIC exists due to potentially large size of composition. Proof of the strong NP-hardness of SUBSET-FIXED-ATOMIC is based on a reduction from the “Monotone 1-in-3” problem.

Section 6 looks into the complexity of deciding reachable-atomicity as well as the configuration reachability problem under the reachably atomic constraint. By exhibiting a polynomial time bottom-up algorithm based on the special structure of configuration reachability graphs of reachably atomic chemical reaction networks, we solved the complexity of deciding reachable-atomicity; By first applying Savitch’s Theorem and then modifying [30]’s result into a simulation of polynomial space Turing machines with reachably atomic chemical reaction networks, we proved that reachability problem for reachably atomic chemical reaction networks are PSPACE-complete.

Section 8, as mentioned above, mainly discusses about the relationships between our model and some models established in [21]. By defining the corresponding mappings and then treading through the definitions, we proved that a chemical reaction network is subset atomic if and only if it admits a core composition, disallows isomerism of elementary species and contains no “redundant” elementary species. We then gave two proofs, one by defining an auxiliary subspace of \mathbb{R}^n (n being the number of types of elementary species) and arguing its set-containment relationship with other relevant subspaces, and another one by directly applying a Theorem in [21], that reachable-atomicity implies constructivity of chemical reaction networks. Lastly, we gave a combinatorics-based proof that reversibly-reachable atomicity is equivalent to explicitly-reversibly constructivity plus the restriction that no isomerism is allowed among elementary species.

Section 7 mentioned some open problems left for future works.

2 Preliminaries

Let $\mathbb{Z}, \mathbb{N}, \mathbb{R}$ respectively denote the set of integers, nonnegative integers, and reals. Let Λ be a finite set. We write \mathbb{N}^Λ to denote $\{f : \Lambda \rightarrow \mathbb{N}\}$. Equivalently, by assuming a “canonical” ordering on Λ , an element $\mathbf{c} \in \mathbb{N}^\Lambda$ can also be viewed as a $|\Lambda|$ -dimensional vector of natural numbers, with each coordinate labeled by $S \in \Lambda$ interpreted as the count of S . $\mathbf{c} \in \mathbb{N}^\Lambda$ interpreted this way is called a *configuration*. We sometimes use multiset notation, e.g., $\{3A, 2B\}$ to denote the configuration with 3 copies of A , 2 of B , and 0 of all other species. $\mathbb{Z}^\Lambda, \mathbb{R}^\Lambda, \mathbb{N}^{\Lambda \times \Delta}$ (where Δ is also a finite set) are defined analogously.

We write $\mathbf{c} \leq \mathbf{c}'$ to denote that $(\forall X \in \Lambda) \mathbf{c}(X) \leq \mathbf{c}'(X)$, and $\mathbf{c} < \mathbf{c}'$ if $\mathbf{c} \leq \mathbf{c}'$ and $\mathbf{c} \neq \mathbf{c}'$. We say \mathbf{c} and \mathbf{c}' are *incomparable* if $\mathbf{c} \not\leq \mathbf{c}'$ and $\mathbf{c}' \not\leq \mathbf{c}$.

Definition 2.1. *Given a finite set of chemical species Λ , a reaction over Λ is a pair $\alpha = (\mathbf{r}, \mathbf{p}) \in \mathbb{N}^\Lambda \times \mathbb{N}^\Lambda$, specifying the stoichiometry of the reactants and products respectively.⁴*

A chemical reaction network is a pair $\mathcal{C} = (\Lambda, R)$, where Λ is a finite set of chemical species, and R is a finite set of reactions over Λ .

A chemical reaction network is *reversible* if $(\forall (\mathbf{r}, \mathbf{p}) \in R) (\mathbf{p}, \mathbf{r}) \in R$.

For configurations $\mathbf{c}_1, \mathbf{c}_2 \in \mathbb{N}^\Lambda$, we write $\mathbf{c}_1 \Rightarrow_{\mathcal{C}}^* \mathbf{c}_2$ (read “ \mathcal{C} reaches \mathbf{c}_2 from \mathbf{c}_1 ”) if there exists a finite reaction sequence (including the empty sequence) that starts with \mathbf{c}_1 and ends with \mathbf{c}_2 . For the sake of simplicity, we write $\mathbf{c}_1 \Rightarrow^* \mathbf{c}_2$ (read “ \mathbf{c}_2 is reachable from \mathbf{c}_1 ”) when \mathcal{C} is clear from context.

Remark 2.2. Note that $\mathbf{r} \Rightarrow^* \mathbf{r}$ is vacuously true for any $(\mathbf{r}, \mathbf{p}) \in R$, by setting the reaction sequence to be empty (length-0). There may also be positive-length reaction sequences for certain networks and certain $(\mathbf{r}, \mathbf{p}) \in R$ to reach \mathbf{r} from \mathbf{r} : for example, $\mathcal{C} = (\{S_1, S_2, S_3, S_4\}, \{S_1 \rightarrow S_2 + S_3, S_2 + S_3 \rightarrow S_4, S_4 \rightarrow S_1\})$. Consider the reaction $(\{1S_1\}, \{1S_2, 1S_3\}) \in R$, we have $\{1S_1\} \Rightarrow^* \{1S_1\}$ by executing the three reactions in order.

Note also that this does not contradict our previous assumption that $(\forall (\mathbf{r}, \mathbf{p}) \in R) \mathbf{r} \neq \mathbf{p}$, for $\mathbf{r} \Rightarrow^* \mathbf{r}$ does not require that $(\mathbf{r}, \mathbf{r}) \in R$.

Definition 2.3. *Given $\mathbf{c} \in \mathbb{N}^\Lambda$ (or $\mathbb{Z}^\Lambda, \mathbb{R}^\Lambda$ etc. analogously), the support of \mathbf{c} , written as $[\mathbf{c}]$, is the set $\{S \in \Lambda \mid \mathbf{c}(S) \neq 0\}$.*

A few more notation conventions are listed here: We write $\mathbf{e}_A \in \mathbb{N}^\Lambda$ as the unit vector that has count 1 on $A \in \Lambda$ and 0 on everything else. Given a vector $\mathbf{x} \in \mathbb{N}^\Lambda$, write $\|\mathbf{x}\| = \|\mathbf{x}\|_1 = \sum_{S \in \Lambda} \mathbf{x}(S)$. When \cdot is any data, write $\langle \cdot \rangle$ for its binary representation as a string, so $|\langle \cdot \rangle|$ is the length of the binary representation of \cdot . Given $f : A \rightarrow B$ and $C \subseteq A$, $f \upharpoonright C$ is the function $C \rightarrow B$, $c \mapsto f(c)$ ($\forall c \in C$). Lastly, when \mathbf{M} is a matrix, write \mathbf{M}^T as its transposition.

3 Definitions of “Atomic”

This section addresses definitions of several classes of networks, some computational complexity result of which will be exhibited later.

⁴There is typically a positive real-valued *rate constant* associated to each reaction, but we ignore reaction rates in this paper and consequently simplify the definition.

3.1 Primitive Atomic

We now arrive at our first definition of “atomic”, which we call *primitive atomic*. All later definitions of atomic will imply this definition (i.e, the set of primitive atomic networks is a superset of the set defined by other definitions).

Intuitively, \mathcal{C} is primitive atomic if all species can be decomposed into the combination of some atoms. Atoms are not required to be a subset of the set of species. Each reaction conserves the total number of each type of atom in the species involved (i.e., the reaction can only rearrange atoms but not create or destroy them).

Definition 3.1 (primitive atomic). *Let Δ be a nonempty finite set and $\mathcal{C} = (\Lambda, R)$ a chemical reaction network. \mathcal{C} is primitive atomic with respect to Δ if for all $S \in \Lambda$, there is $\mathbf{d}_S \in \mathbb{N}^\Delta \setminus \{\mathbf{0}\}$ such that*

1. $(\forall (\mathbf{r}, \mathbf{p}) \in R)(\forall A \in \Delta) \sum_{S \in \Lambda} \mathbf{r}(S) \cdot \mathbf{d}_S(A) = \sum_{S \in \Lambda} \mathbf{p}(S) \cdot \mathbf{d}_S(A)$. (*reactions preserve atoms*)
2. $(\forall A \in \Delta)(\exists S \in \Lambda) \mathbf{d}_S(A) \neq 0$. (*each atom appears in the composition of some species*)

For $S \in \Lambda$, call \mathbf{d}_S in Condition (1) the (atomic) composition of S . We say \mathcal{C} is primitive atomic if there is a nonempty finite set Δ such that \mathcal{C} is primitive atomic with respect to Δ .

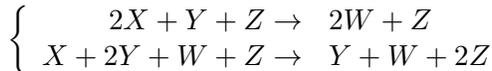
Condition (1) embodies the intuition above. Condition (2) prescribes that each atom appears in the composition of at least one species.

Next, we introduce the definitions of stoichiometric matrix and composition matrix, in order to facilitate a more terse and transparent way of expressing conservation laws. In particular, \mathbf{A} encodes the net change of species that applying a reaction causes, and \mathbf{D} compiles all composition vectors into one data structure.

Definition 3.2 (Stoichiometric Matrix). *The stoichiometric matrix \mathbf{A} for a chemical reaction network $\mathcal{C} = (R, \Lambda)$ is the $|R| \times |\Lambda|$ matrix where the entry $\mathbf{A}_{(\mathbf{r}, \mathbf{p}), S} = \mathbf{p}(S) - \mathbf{r}(S)$ for each $(\mathbf{r}, \mathbf{p}) \in R$ and $S \in \Lambda$.*

Notationwise, $\mathbf{A}_{(\mathbf{r}, \mathbf{p}), S}$ is the entry whose row is labeled by the reaction (\mathbf{r}, \mathbf{p}) and column by the species S . Each row of the stoichiometric matrix represents the change of count of each species via execution of 1 unit of (\mathbf{r}, \mathbf{p}) .

Example 3.3. Consider the network $\mathcal{C} = (\{X, Y, W, Z\}, \{((2, 1, 0, 1)^T, (0, 0, 2, 1)^T), ((1, 2, 1, 1)^T, (0, 1, 1, 2)^T)\})$. The set of reactions can be described as:



For \mathcal{C} , the stoichiometric matrix $\mathbf{A} = \begin{bmatrix} -2 & -1 & 2 & 0 \\ -1 & -1 & 0 & 1 \end{bmatrix}$.

Definition 3.4 (Composition Matrix). *Let $\mathcal{C} = (\Lambda, R)$ be primitive atomic with respect to Δ . The composition matrix \mathbf{D}_Δ for \mathcal{C} with respect to Δ is the $|\Lambda| \times |\Delta|$ matrix whose row vectors are $(\mathbf{d}_S \upharpoonright \Delta)^T$ ($S \in \Lambda$).*

The composition matrix composes the information of atomic composition for each species in the form of a matrix.

Remark 3.5. We note that a composition matrix has the following properties:

1. Let $\mathbf{c} \in \mathbb{N}^\Lambda$ be a configuration for primitive atomic $\mathcal{C} = (\Lambda, R)$, then the vector $\mathbf{D}_\Delta^T \cdot \mathbf{c} \in \mathbb{N}^\Delta$ illustrates the current count of each atom in the system. That is, $(\mathbf{D}_\Delta^T \cdot \mathbf{c})(A)$, the entry of $\mathbf{D}_\Delta^T \cdot \mathbf{c}$ labeled by $A \in \Delta$, represents the current count of atom A in the system. In fact, for each $A \in \Delta$,

$$\begin{aligned} \text{current count of } A \text{ in system} &= \sum_{S \in \Lambda} \mathbf{d}_S(A) \cdot \mathbf{c}(S) \\ &= (\mathbf{D}_\Delta^T)_{(A,:)} \cdot \mathbf{c}, \end{aligned}$$

where $(\mathbf{D}_\Delta^T)_{(A,:)}$ stands for the row vector of \mathbf{D}_Δ^T labeled by A .

2. If $\mathcal{C} = (\Lambda, R)$ is primitive atomic with respect to Δ , then $\mathbf{A} \cdot \mathbf{D}_\Delta = \mathbf{0}$, where \mathbf{A} is the stoichiometric matrix in Definition 3.2 above, and $\mathbf{0}$ is the $|R| \times |\Delta|$ zero matrix. Intuitively, this illustrates the fact that the number of each type of atom is preserved throughout all reactions. Indeed, for each $(\mathbf{r}, \mathbf{p}) \in R$, $A \in \Delta$, the entry

$$\begin{aligned} (\mathbf{A} \cdot \mathbf{D}_\Delta)_{((\mathbf{r}, \mathbf{p}), A)} &= \sum_{S \in \Lambda} (\mathbf{p}(S) - \mathbf{r}(S)) \cdot \mathbf{d}_S(A) \\ &= \sum_{S \in \Lambda} 0 \quad (\text{by -Atomicity}) \\ &= 0 \end{aligned}$$

Conversely, let Δ be a set and \mathbf{D}_Δ be a $|\Lambda| \times |\Delta|$ matrix with rows labeled by each $S \in \Lambda$ and columns labeled by each $A \in \Delta$. If $\mathbf{A} \cdot \mathbf{D}_\Delta = \mathbf{0}$, then by the same arithmetics above and by Definition 3.1, \mathcal{C} is primitive atomic with respect to Δ . It follows that one may rewrite Condition (1) of Definition 3.1 as:

$$(a) (\exists \Delta \mid 0 < |\Delta| < \infty)(\exists \mathbf{D}_\Delta \in \mathbb{N}^{\Lambda \times \Delta} \mid \forall S \in \Lambda, (\mathbf{D}_\Delta)_{(S,:)} \neq (0^\Delta)^T) \text{ such that } \mathbf{A} \cdot \mathbf{D}_\Delta = \mathbf{0}.$$

We use these two expressions for condition (1) interchangeably in future parts of this paper. Furthermore, when \mathcal{C}, Δ is clear from the context, we use \mathbf{D} as shorthand for \mathbf{D}_Δ , and we say \mathcal{C} is primitive atomic with respect to Δ via \mathbf{D} .

Example 3.6. For the network defined in Example 3.3:

$$\mathcal{C} = (\{X, Y, W, Z\}, \{((2, 1, 0, 1)^T, (0, 0, 2, 1)^T), ((1, 2, 1, 1)^T, (0, 1, 1, 2)^T)\})$$

and the set of atoms $\Delta = \{H, O\}$, a valid composition matrix could be:

$$\mathbf{D} = \begin{bmatrix} 2 & 0 \\ 0 & 2 \\ 2 & 1 \\ 2 & 2 \end{bmatrix}$$

It can also be verified that

$$\begin{aligned} \mathbf{A} \cdot \mathbf{D} &= \begin{bmatrix} -2 & -1 & 2 & 0 \\ -1 & -1 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} 2 & 0 \\ 0 & 2 \\ 2 & 1 \\ 2 & 2 \end{bmatrix} \\ &= \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} \end{aligned}$$

So \mathcal{C} is primitive atomic with respect to $\Delta = \{H, O\}$ via \mathbf{D} .

3.2 Subset Atomic

The next definition requires all atoms to be species themselves.

Definition 3.7 (subset atomicity). *Let $\mathcal{C} = (\Lambda, R)$ be a chemical reaction network and let $\Delta \subseteq \Lambda$ be nonempty. We say that \mathcal{C} is subset- Δ -atomic if \mathcal{C} is primitive atomic with respect to Δ and, for each $S \in \Lambda$, $S \in \Delta \implies \mathbf{d}_S = \{S\}$ and $S \notin \Delta \implies \|\mathbf{d}_S\| \geq 2$. We say \mathcal{C} is subset atomic if there exists a nonempty $\Delta \subseteq \Lambda$ such that \mathcal{C} is subset- Δ -atomic.*

By definition 3.7, no two atoms can have the same atomic composition, but it allows for two distinct molecular species to have the same composition. In this case we say the two species are *isomers* (reminiscent of isomers in nature that are composed of the same atoms in different geometrical arrangements). As for the requirement that each non-atom species decompose to a vector of size at least 2, that is to illustrate the idea that a molecule should be composed of at least 2 atoms.

3.3 Reachably Atomic

The next definition imposes further requirement that *each composition can be “realized” via a sequence of reactions*. As discussed in Subsection 1.1, this definition was originally developed in [1] to help their approach to the Global Attractor Conjecture in the field of mass action kinetics. Considering the convention for most networks, we relax their requirement of reversibility for each reaction.

Definition 3.8 (reachable atomicity). *A chemical reaction network $\mathcal{C} = (\Lambda, R)$ is reachably atomic if*

1. \mathcal{C} is subset atomic with respect to some $\Delta \subseteq \Lambda$, and
2. for each $S \in \Lambda \setminus \Delta$, $\{1S\} \Rightarrow^* \mathbf{d}_S$.

Here and wherever necessary, with slight abuse of notation, \mathbf{d}_S , which represents the atomic composition of S , simultaneously represents a configuration in \mathbb{N}^Δ reachable from $\{1S\}$.

Condition 2 is a strong restriction. For example, the atom set of a reachably atomic network is unique:

Lemma 3.9. *If (Λ, R) is reachably atomic, then the choice of Δ with respect to which (Λ, R) is reachably atomic is unique. Moreover, for each $S \in \Lambda$, \mathbf{d}_S is unique, i.e., if $\{1S\} \Rightarrow^* \mathbf{c} \in \mathbb{N}^\Delta$, then $\mathbf{c} = \mathbf{d}_S$.*

Proof. The intuition of the proof by contradiction is to show that should there exist an element A in the set difference between two atom sets, the composition of A with respect to one atom set violates the preservation of atoms in another. Assume for the sake of contradiction that for some reachably atomic network \mathcal{C} , there exist $\Delta_1 \neq \Delta_2$ with respect to both of which \mathcal{C} is reachably atomic, respectively via composition matrices \mathbf{D}_1 and \mathbf{D}_2 . Note that $(\Delta_1 \setminus \Delta_2) \cup (\Delta_2 \setminus \Delta_1) \neq \emptyset$. Take $A \in (\Delta_1 \setminus \Delta_2) \cup (\Delta_2 \setminus \Delta_1)$:

1. If $A \in \Delta_1 \setminus \Delta_2$, then $\{1A\}$ is decomposable into some $\mathbf{c} \in \mathbb{N}^\Lambda \mid [\mathbf{c}] \subseteq \Delta_2$ via a sequence of reactions, with $\|\mathbf{c}\|_1 \geq 2$. There is no way for this sequence of reactions to preserve atoms with respect to Δ_1 , for the initial configuration has a single atom $A \in \Delta_1$ while the final configuration has no atom A .
2. Similarly, if $A \in \Delta_2 \setminus \Delta_1$, there will be a sequence of reactions violating preservation of atoms with respect to Δ_2 .

We next prove the uniqueness of composition vectors \mathbf{d}_S for all $S \in \Lambda$, or equivalently, the uniqueness of composition matrix \mathbf{D} . Suppose not, then there exists $S \in \Lambda \setminus \Delta$ s.t. $\{1S\} \Rightarrow^* \mathbf{y}_1, \mathbf{y}_2 \in \mathbb{N}^\Lambda$, $\mathbf{y}_1 \neq \mathbf{y}_2$ and $[\mathbf{y}_1], [\mathbf{y}_2] \subseteq \Delta$, via reaction sequences $\mathbf{s}_1, \mathbf{s}_2$ respectively. Assume without loss of generality that \mathbf{s}_1 preserves the number of atoms in each reaction, which means $\mathbf{y}_1 = \mathbf{d}_S$. Then there must be one or more actions in \mathbf{s}_2 that does (do) not preserve the number of atoms, for $\mathbf{s}_1, \mathbf{s}_2$ share the same initial configuration $\{1S_i\}$ yet reach different final count of atoms, while no atoms are allowed to be isomeric to each other. \square

Conservation laws in “-atomic” networks reminds us of a more familiar type of conservation law, which is mass conservation. The next section exhibits some observations on the relationship between these two types of conservation laws.

4 Mass-Conservation and primitive atomicity

This section shows that primitive atomicity and mass conservation are equivalent. We first formalize the latter concept, that “mass can neither be created nor destroyed”:

Definition 4.1 (Mass Conservation). *A chemical reaction network $\mathcal{C} = (\Lambda, R)$ is mass conserving if*

$$(\exists \mathbf{m} \in \mathbb{R}_{>0}^\Lambda)(\forall (\mathbf{r}, \mathbf{p}) \in R) \sum_{S \in \Lambda} \mathbf{r}(S) \cdot \mathbf{m}(S) = \sum_{S \in \Lambda} \mathbf{p}(S) \cdot \mathbf{m}(S)$$

Equivalently, if \mathbf{A} is the stoichiometric matrix in Definition 3.2, then \mathcal{C} is mass conserving if $(\exists \mathbf{m} \in \mathbb{R}_{>0}^\Lambda) \mathbf{A} \cdot \mathbf{m} = \mathbf{0}$. We call \mathbf{m} a mass distribution vector.

“Mass Conserving” captures the feature that for every reaction in \mathcal{C} , the total mass of reactants are equal to the total mass of products. Difference between the definitions of Mass Conserving and Primitive Atomicity (as well as all “-atomicity” definitions descended therefrom) become clear if we compare the matrix form of their respective conservation laws: mass conservation requires a single conservation relation ($\mathbf{A} \cdot \mathbf{m} = \mathbf{0}^{|R|}$), while primitive atomicity requires $|\Delta|$ of them ($\mathbf{A} \cdot \mathbf{D} = \mathbf{0}$ where \mathbf{D} is $|\Lambda| \times |\Delta|$).

However, it is very intuitive that these two types of conservation laws are closely related. In fact, the freedom of defining Δ independent of Λ provides us a choice for making Δ a singleton, which enables us to prove the following equivalence:

Proposition 4.2. *For any network \mathcal{C} , \mathcal{C} is primitive atomic $\Leftrightarrow \mathcal{C}$ is mass conserving. Further, there exists an $O(|\langle \mathbf{A} \rangle|^5)$ algorithm to decide if \mathcal{C} is primitive atomic, with \mathbf{A} the stoichiometric matrix of \mathcal{C} .*

Proof. Intuitively, converting a mass distribution vector into integral and interpreting it as composition vectors proves the “mass conservation \Rightarrow primitive atomicity” direction, while assigning mass 1 to each atom for a primitive atomic chemical reaction network gives the mass distribution vector.

1. primitive atomic \Leftrightarrow mass conserving:

(a) mass Conserving \Rightarrow primitive Atomic:

If a chemical reaction network \mathcal{C} is mass conserving, then there exists a mass distribution vector $\mathbf{m} \in \mathbb{R}_{>0}^\Lambda$ s.t. $\mathbf{A} \cdot \mathbf{m} = 0^{|R|}$, with \mathbf{A} the stoichiometric matrix of \mathcal{C} . We shall exhibit a constructive way of finding a *rational* solution to the linear system $\mathbf{A} \cdot \mathbf{m} = 0^{|R|}$, thereby enabling further manipulation to finally yield an integral solution that could be used to construct a desired composition matrix \mathbf{D} .

- i. Knowing the existence of the distribution vector $\mathbf{m} \in \mathbb{R}_{>0}^\Lambda$ s.t. $\mathbf{A} \cdot \mathbf{m} = 0^{|R|}$, we run Sergei’s “strictly positive solution finder for linear system” algorithm [11, 12] to find $\mathbf{m}' \in \mathbb{Q}_{>0}^\Lambda \subseteq \mathbb{R}_{>0}^\Lambda$, s.t. $\mathbf{A} \cdot \mathbf{m}' = 0^{|R|}$. The algorithm has complexity $O(\max\{|\Lambda|^4, |R|^4\} \cdot |\langle \mathbf{A} \rangle|) = O(|\langle \mathbf{A} \rangle|^5)$ [11, 12].
- ii. Now that we have obtained a solution vector $\mathbf{m}' \in \mathbb{Q}_{>0}^\Lambda$, for each $S \in \Lambda$ we may write $\mathbf{m}'(S) = \frac{a_S}{b_S} \mid a_S \in \mathbb{Z}, b_S \in \mathbb{Z}_{\neq 0}$. Compute $c = \text{lcm}\{b_S\}_{S \in \Lambda}$, and let $\mathbf{m}'' = \mathbf{m}' \cdot c$. Since $\mathbf{A} \cdot \mathbf{m}'' = c \cdot (\mathbf{A} \cdot \mathbf{m}') = 0^{|R|}$, by definition $\mathbf{m}'' \in \mathbb{N}_{>0}^{|\Lambda|}$ is also a valid mass distribution vector.

Complexitywise, one may apply the binary Euclidean Algorithm (Let us denote the algorithm as $\text{gcd}(a, b)$) for $|\Lambda| - 1$ times, finding $\text{gcd}(b_{S_1}, b_{S_2}), \text{gcd}(b_{S_1}, b_{S_2}, b_{S_3}) = \text{gcd}(\text{gcd}(b_{S_1}, b_{S_2}), b_{S_3}), \dots, \text{gcd}(b_{S_1}, b_{S_2}, \dots, b_{S_{|\Lambda|}})$ dynamically, and then compute

$c = \text{lcm}\{b_S\}_{S \in \Lambda} = \frac{\prod_{S \in \Lambda} b_S}{\text{gcd}\{b_S\}_{S \in \Lambda}}$. Since the complexity of Euclidean Algorithm $\text{gcd}(a, b)$ is $O((\max\{|\langle a \rangle|, |\langle b \rangle|\})^3)$ [27], complexity for computing $\text{gcd}(b_1, \dots, b_{S_{|\Lambda|}})$ as well as $c = \text{lcm}(b_1, \dots, b_{S_{|\Lambda|}})$ is $O(|\Lambda| \cdot (\max\{|\langle \mathbf{m}'(S) \rangle|\})^3)$.

Lastly, it is shown in the footnote that computing $\mathbf{m}'' = \mathbf{m}' \cdot c$ expands the binary size of \mathbf{m}' by at most a factor of $|\Lambda|$,⁵ so altogether, the complexity for this Step (1(a)ii) is $O(|\Lambda|^2 \cdot (\max\{|\langle \mathbf{m}'(S) \rangle|\})^3)$. Since \mathbf{m}' is obtained from Step (1(a)i) with size-bound $|\langle \mathbf{A} \rangle|^5$, we have

$$O(|\Lambda|^2 \cdot (\max\{|\langle \mathbf{m}'(S) \rangle|\})^3) = O(|\langle \mathbf{A} \rangle|^{17}) \quad (4.1)$$

⁵To make the notations a little simpler, let the i -th entry of \mathbf{m}' be written as $\frac{c_i}{d_i}$, so the the numerator and denominator are $\log(|c_i|)$ and $\log(|d_i|)$ bits long, respectively ($\forall i \in [1, |\Lambda|]$). Define

$$t := \underbrace{\sum_{j=1}^{|\Lambda|} (\log(c_j) + \log(d_j))}_{\text{binary size of } \mathbf{m}'}$$

In the worst case where $d_i, d_j (\forall i, j)$ are pairwise co-prime, multiplying each entry with $\text{gcd}(d_1, d_2, \dots, d_{|\Lambda|})$ is equivalent to first setting all d_i ’s to 1, then expending each c_i by $(\sum_{j=1}^{|\Lambda|} \log(d_j)) - \log(d_i)$ bits. The net effect is expanding the

iii. Output $\Delta = \{A_1\}$ where $A_1 \notin \Lambda$, and $\mathbf{D} = 2\mathbf{m}''$.

Note that \mathbf{D} is a length- $|\Lambda|$ column vector when $|\Delta| = 1$, and so is $2\mathbf{m}''$, hence their dimensions match.

Intuitively, the last step is assigning mass 1 to the single atom A_1 , and then decomposing each $S \in \Lambda$ into $2\mathbf{m}''(S)$ of A_1 's. We formally verify the correctness of this output:

- i. $\mathbf{A} \cdot \mathbf{D} = 2(\mathbf{A} \cdot \mathbf{m}'') = 0^{|R|}$, so Condition (1) for primitive atomicity is satisfied;
- ii. A_1 as the single atom in Δ is used in the composition of each molecular species, verifying Condition (2) as well.

Complexitywise, Step (1(a)ii) dominates, so (4.1) is the complexity of the whole algorithm.

(b) primitive atomic \Rightarrow mass conserving:

We shall prove that given any chemical reaction network \mathcal{C} primitive atomic with respect to Δ via composition matrix \mathbf{D} , there exists a polynomial time algorithm to compute a valid mass distribution vector \mathbf{m} . For each $S \in \Lambda$, let $\mathbf{m}(S) = \sum_{A \in \Delta} \mathbf{d}_S(A) = \|\mathbf{D}_{(S, \cdot)}\|_1$;

that is, we assign mass 1 to all atoms, making the mass of a molecular species equal to the total count of atoms in its atomic composition. Then for each $(\mathbf{r}, \mathbf{p}) \in R$, the entry

$$\begin{aligned}
 (\mathbf{A} \cdot \mathbf{m})_{(\mathbf{r}, \mathbf{p})} &= \sum_{S \in \Lambda} (\mathbf{A}_{((\mathbf{r}, \mathbf{p}), S)}) \cdot \sum_{A \in \Delta} \mathbf{d}_S(A) \\
 &= \sum_{S \in \Lambda} \sum_{A \in \Delta} (\mathbf{A}_{((\mathbf{r}, \mathbf{p}), S)}) \cdot \mathbf{D}_{(S, A)} \\
 &= \sum_{A \in \Delta} \sum_{S \in \Lambda} (\mathbf{A}_{((\mathbf{r}, \mathbf{p}), S)}) \cdot \mathbf{D}_{(S, A)} \\
 &= \|(\mathbf{A} \cdot \mathbf{D})_{((\mathbf{r}, \mathbf{p}), \cdot)}\|_1 \\
 &= 0,
 \end{aligned}$$

as desired. Note that $\|\langle \mathbf{m} \rangle\|_1 = O(\|\langle \mathbf{D} \rangle\|)$, which gives the upper bound of complexity.

2. There exists an $O(|\langle A \rangle|^5)$ algorithm to decide if \mathcal{C} is primitive atomic.

By the equivalence relationship shown right above, deciding primitive atomicity is equivalent to deciding mass conservation, which is in turn equivalent to deciding if there exists a strictly positive solution to the linear system $\mathbf{A} \cdot \mathbf{m} = 0^{|R|}$ with \mathbf{A} the stoichiometric matrix of \mathcal{C} . The latter problem is decidable by an $O(|\langle \mathbf{A} \rangle|^5)$ algorithm, as mentioned in Step 1(a)i [11, 12]. \square

size of the input by

$$\begin{aligned}
 \sum_{i=1}^{|\Lambda|} (((\sum_{j=1}^{|\Lambda|} \log(d_j)) - \log(d_i)) - (\log d_i - 1)) &= \sum_{i=1}^{|\Lambda|} ((\sum_{j=1}^{|\Lambda|} \log(d_j)) - 2\log(d_i) + 1) \\
 &= (|\Lambda| - 2)(\sum_{j=1}^{|\Lambda|} \log(d_j)) + |\Lambda| \\
 &< |\Lambda| \cdot t + |\Lambda| \\
 &= O(|\Lambda| \cdot t)
 \end{aligned}$$

bits, as desired.

Remark 4.3. The equivalence with Mass Conserving is **not true** for *subset atomic* networks: Consider $\mathcal{C} = (\{X, Y\}, \{(1, 0), (0, 1)\})$. This network involves a single reaction: $X \rightarrow Y$. There are only three choices for the set of atoms Δ , which is now required to satisfy $|\Delta| > 0, \Delta \subseteq \Lambda$: $\Delta = \{X\}$, $\Delta = \{Y\}$ or $\Delta = \{X, Y\}$, respectively.

\mathcal{C} is Mass Conserving, the mass distribution vector being $\mathbf{m}(X) = \mathbf{m}(Y) = 1$; Nonetheless, no choice of Δ makes \mathcal{C} subset atomic. Take $\Delta = \{X, Y\}$ for example, and consider the atom X : the reactants have a single atom of X but the products have no X (since Y is itself an atom it cannot be composed of X). If $\Delta = \{X\}$, then Y is not an atom, but the definition of atomic composition requires that non-atomic species have a composition with at least two atoms, which means that since X is the only atom, \mathbf{d}_Y contains at least two X 's, while the reactants have only one X . A similar argument applies if $\Delta = \{Y\}$.

Recall that subset atomicity imposes the restriction that $\Delta \subseteq \Lambda$. As we'll show in the following section, this single restriction increases the computational complexity in deciding subset atomicity, compared to primitive atomicity.

5 Complexity of subset atomic

We shall find in this chapter the computational complexity for deciding the subset atomicity of networks. First, we define languages of encodings of chemical reaction networks with subset atomicity:

Definition 5.1. *We define the following languages:*

$$\begin{aligned} \text{SUBSET-ATOMIC} &= \{ \langle \Lambda, R \rangle \mid (\exists \Delta \subseteq \Lambda) (\langle \Lambda, R \rangle \text{ is subset atomic with respect to } \Delta) \} \\ \text{SUBSET-FIXED-ATOMIC} &= \{ \langle \Lambda, R, \Delta \rangle \mid \langle \Lambda, R \rangle \text{ is subset atomic with respect to } \Delta \} \end{aligned}$$

By definition, SUBSET-ATOMIC is the language whose elements are the encoding of a *subset atomic* chemical reaction network. SUBSET-FIXED-ATOMIC, on the other hand, is the language consisting of the encoding of a (network, atom set) pair where the network is *subset atomic* with respect to the given atom set. In this section we determine the complexity classes of these languages.

5.1 Subset-Fixed-Atomic and Subset-Atomic are in NP

It is not obviously true that there exists a short witness for either language (which if true would have proved that both languages are in NP immediately). We therefore reduce SUBSET-FIXED-ATOMIC to INTEGER-PROGRAMMING, which is in NP [35].

Lemma 5.2. *SUBSET-FIXED-ATOMIC is polynomial-time many-one reducible to INTEGER-PROGRAMMING (hereinafter, "IP").*

Proof. The proof is done by translating the axioms/constraints in the definition of subset atomicity into linear equations or inequalities. We transform the input into the encoding of the following equivalent linear system: In the following, $b_1, b_2, \dots, b_{m-n}, c_1, c_2, \dots, c_n \in \mathbb{N}$ are slack variables so

that we may express an inequality as an equality together with a nonnegativity constraint on b_s, c_a ($1 \leq s \leq m - n, 1 \leq a \leq n$).

$$\begin{aligned}
& \text{for } a \in [1, n] : \\
& \quad \text{for } r \in [1, k] : \\
& \qquad \sum_{s=1}^m \mathbf{v}_r(S_s) \cdot x_{sa} = 0 \\
& \text{for } s \in [1, m - n] : \\
& \qquad \left(\sum_{a=1}^n x_{s,a} \right) - b_s = 2 \\
& \qquad \qquad b_s \geq 0 \\
& \text{for } a \in [1, n] : \\
& \qquad \left(\sum_{s=1}^{m-n} x_{s,a} \right) - c_a = 1 \\
& \qquad \qquad c_a \geq 0 \\
& \text{for } a \in [1, n] : \\
& \quad \text{for } a' \in [1, n] : \\
& \qquad \quad \text{if } a = a' \\
& \qquad \qquad \qquad x_{a'+m-n,a} = 1 \\
& \qquad \quad \text{else} \\
& \qquad \qquad \qquad x_{a'+m-n,a} = 0 \\
& \text{for } s \in [1, m] : \\
& \quad \text{for } a \in [1, n] : \\
& \qquad \qquad x_{sa} \geq 0
\end{aligned} \tag{5.1}$$

The equivalence between the linear system and the subset atomicity of (Λ, R, Δ) follows from Definitions 5.1 and 3.7. In fact, the first equation of (5.1) asks the number of each atom to be preserved across each reaction; the second equation and third inequality prescribes that for all non-atomic species $S \in \Lambda \setminus \Delta$, $\|\mathbf{d}_S\|_1 \geq 2$; the fourth equation and fifth inequality prescribes that for all atomic species $A \in \Delta$, $\sum_{S \in \Lambda \setminus \Delta} \mathbf{d}_S(A) \geq 1$, which is equivalent to saying that each A appears in at least one atomic composition for some molecules; the sixth and seventh equation translates to $\forall A \in \Delta, \mathbf{d}_A = \mathbf{e}_A$ (recall \mathbf{e}_A is the unit vector that is 1 on A and 0 on everything else); and the last inequality restricts atomic counts in species to be nonnegative integers.

Note, also, that $\mathbf{d}_S \neq 0^\Lambda$ has been ensured by the respective restrictions on composition for non-atom species and atoms.

Let us analyse the complexity of the reduction. To construct a constraint system

$$\mathbf{Ax} = \mathbf{b}, \mathbf{x} \in \mathbb{N}^n$$

we first observe that

$$|\mathbf{x}| = |\{x_{s,a} \mid s \in [1, m], a \in [1, n]\}| + |\{d_s \mid s \in [1, m - n]\}| + |c_a \mid a \in [1, n]| = mn + m$$

and that

$$\begin{aligned}
|\mathbf{b}| &= \text{number of equations in (5.1)} \\
&= nk + m + n^2
\end{aligned}$$

which means the matrix \mathbf{A} is $(nk + m + n^2) \times (mn + m)$, a polynomial in m, n, k . Further, each entry of \mathbf{b} is an integer in $[0, 2]$, while the absolute value of each integral entry of \mathbf{A} is bounded by $\max\{\max_{r \in R, s \in [1, m]} \{\mathbf{v}_r(S_s)\}, 1\}$. This shows that the linear system is of size polynomial in m, n, k and the binary size of $\langle \Lambda, R, \Delta \rangle$, and the process reducing $\langle \Lambda, R, \Delta \rangle$ to the linear system is also of time polynomial to the same parameters.

Since $\langle \Lambda, R, \Delta \rangle \in \text{SUBSET-FIXED-ATOMIC}$ if and only if $\Delta \subseteq \Lambda$ and the linear system (5.1) has a solution, and since constructing the encoding of the linear system (5.1) takes polynomial time, we conclude that $\text{SUBSET-FIXED-ATOMIC}$ is polynomial-time many-one reducible to Integer Programming. \square

We immediately obtain the following corollaries:

Corollary 5.3. $\text{SUBSET-FIXED-ATOMIC} \in \text{NP}$.

Proof. Papadimitriou [35] proved that for the $m \times n$ integer matrix \mathbf{A} and m -vector \mathbf{b} , the problem of deciding whether there exists $\mathbf{x} \in \mathbb{N}^n$ such that

$$\mathbf{A}\mathbf{x} = \mathbf{b}, \mathbf{x} \in \mathbb{N}^n \tag{5.2}$$

is contained in NP . Since $\text{SUBSET-FIXED-ATOMIC} \leq_m^p \text{IP}$, the conclusion follows. \square

Corollary 5.4. $\text{SUBSET-ATOMIC} \in \text{NP}$.

Proof. The proof is given by an polynomial time verification algorithm using the polynomial-time verifier of $\text{SUBSET-FIXED-ATOMIC}$ as an oracle. From Corollary 5.3 we know that there exists a polynomial time verifier V' for the language $\text{SUBSET-FIXED-ATOMIC}$, who takes an instance $\langle \Lambda, R, \Delta \rangle$ and a witness $\langle \mathbf{D} \rangle$, the latter being encoding of a composition matrix \mathbf{D} , and ACCEPTS (resp. REJECTS) if and only if $\langle \Lambda, R, \Delta \rangle \in \text{SUBSET-FIXED-ATOMIC}$ (resp. $\langle \Lambda, R, \Delta \rangle \notin \text{SUBSET-FIXED-ATOMIC}$) via \mathbf{D} .

We exhibit a polynomial time verifier V for SUBSET-ATOMIC . On instance $c = \langle \Lambda, R \rangle$ and witness $w = \langle \Delta, \mathbf{D} \rangle$, the algorithm V :

- 1 *Parses* $\langle c, w \rangle$ into $\langle c', w' \rangle$ where $c' = \langle \Lambda, R, \Delta \rangle$ and $w' = \mathbf{D}$;
- 2 *Runs* V' on $\langle c', w' \rangle$ and echos its output;

Algorithm 1: Verifier V for SUBSET-ATOMIC

A valid witness $w = \langle \Delta, \mathbf{D} \rangle$ ⁶ has a size polynomial of the input size since a valid $\Delta \subseteq \Lambda$, while Corollary 5.3 ensures that a valid \mathbf{D} has size polynomial of the input as well. Step 1 therefore takes linear time and by Corollary 5.3 again, step 2 takes polynomial time too. \square

5.2 Subset-Fixed-Atomic is NP-hard

Our proof shall be based on reduction from $\text{MONOTONE-1-IN-3-SAT}$. Recall that a monotone 3-CNF C is a conjunctive normal form with no negations, and a 1-in-3 satisfying assignment for C is an assignment of Boolean values to all variables such that for each clause in C , exactly one variable is assigned true.

⁶In fact the witness could even be a single $\langle \mathbf{D} \rangle$, as one may read each row of \mathbf{D} and decide if the species represented by that row is a molecule (sum of entries in the row is at least 2) or an atom (the row would be a unit vector) immediately, thereby determining Δ .

As a well-established result, the following language is NP-complete [19].

$$\text{MONOTONE-1-IN-3-SAT} = \{ \langle V, C \rangle \mid C \text{ is a monotone 3-CNF over } V = \{v_1, v_2, \dots, v_n\}, \\ \text{and there exists a 1-in-3 satisfying assignment for } C \}$$

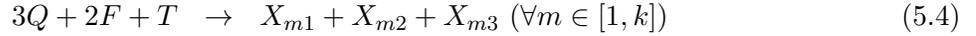
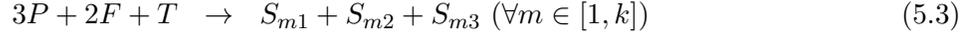
We shall prove that MONOTONE-1-IN-3-SAT is polynomial time many-one reducible to SUBSET-FIXED-ATOMIC.

Lemma 5.5. MONOTONE-1-IN-3-SAT \leq_m^p SUBSET-FIXED-ATOMIC.

Proof. We design a chemical reaction network whose molecular species consist of 2 atoms while atoms in Δ represent “True” and “False”, and whose reactions force the network to be subset- Δ -atomic if and only if the original input instance is in MONOTONE-1-IN-3-SAT. For each instance $\langle V, C \rangle$ of MONOTONE-1-IN-3-SAT,

$$\text{let } \Delta = \{T, F, P, Q\}, \Lambda = \{S_1, S_2, \dots, S_n, X_1, X_2, \dots, X_n\} \cup \Delta.$$

To construct R , we denote $C = c_1 \wedge c_2 \wedge \dots \wedge c_k$. For the i -th clause $c_i \in C$, let v_{ij} denote the j -th literal of c . Same indexing convention applies for $\{S_i\}_{i=1}^n$ and $\{X_i\}_{i=1}^n$, hence each S_{ij} (resp. X_{ij}) denotes an element in $\{S_i\}_{i=1}^n$ (resp. $\{X_i\}_{i=1}^n$).⁷ The set R contains the following reactions⁸:



Note that we only need to construct $4 + 2n$ species and $2k + n$ reactions whose coefficients are constant, so this transformation is polynomial in time and space. We argue that $\langle V, C \rangle \in \text{MONOTONE-1-IN-3-SAT}$ if and only if $\langle \Lambda, R, \Delta \rangle \in \text{SUBSET-FIXED-ATOMIC}$.

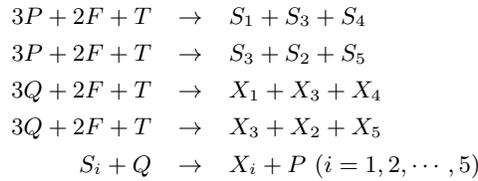
\Rightarrow : If $\langle V, C \rangle \in \text{MONOTONE-1-IN-3-SAT}$, there exists a $\phi : V \rightarrow \{0, 1\}$ under which $\exists (n_1, n_2, \dots, n_q) \subseteq (1, 2, \dots, n)$ s.t. $\phi(v_{n_i}) = 1$ ($\forall i \in [1, q]$), $\phi(v_j) = 0$ ($j \in ([1, n] \setminus (n_1, n_2, \dots, n_q))$), and for each $c_i \in C$, exactly one in three of its literals evaluates to 1. Let:

$$\begin{aligned} \Lambda_{TP} &= \{S_{n_j} \mid j \in [1, q]\} \\ \Lambda_{FP} &= \Lambda \setminus (\Delta \cup \Lambda_{TP} \cup \{X_i \mid i \in [1, n]\}) \\ \Lambda_{TQ} &= \{X_{n_j} \mid j \in [1, q]\} \\ \Lambda_{FQ} &= \Lambda \setminus (\Delta \cup \Lambda_{TQ} \cup \Lambda_{TP} \cup \Lambda_{FP}) \end{aligned}$$

Then $\langle \Lambda, R, \Delta \rangle \in \text{SUBSET-FIXED-ATOMIC}$ because with all atoms listed in the order: $\{T, F, P, Q\}$,

⁷For example, for $V = \{v_1, v_2, \dots, v_5\}$, $C = (v_1 \vee v_3 \vee v_4) \wedge (v_3 \vee v_2 \vee v_5)$, $v_{11} = v_1, v_{12} = v_3, v_{13} = v_4, \dots, v_{23} = v_5$. Correspondingly, $S_{11} = S_1, X_{11} = X_1, \dots, S_{23} = S_5, X_{23} = X_5$.

⁸To continue the example in the previous footnote, the set of reactions shall be:



one may make the following composition:

$$\begin{aligned}
\mathbf{d}_U \upharpoonright \Delta &= (1, 0, 1, 0)^T, \forall U \in \Lambda_{TP} \\
\mathbf{d}_V \upharpoonright \Delta &= (0, 1, 1, 0)^T, \forall V \in \Lambda_{FP} \\
\mathbf{d}_W \upharpoonright \Delta &= (1, 0, 0, 1)^T, \forall W \in \Lambda_{TQ} \\
\mathbf{d}_Z \upharpoonright \Delta &= (0, 1, 0, 1)^T, \forall Z \in \Lambda_{FQ}
\end{aligned}$$

Because of the way $\{n_j\}_{j=1}^q$ was taken, for each reaction in (5.3), exactly one of the product species decompose to $1T$ and $1P$, with the other two decomposing to $1F$ and $1P$. Similar argument applies for reactions in (5.4), substituting X_i for S_i and Q for P . Arithmetics show that all three reactions (5.3) through (5.5) preserve the number of atoms, each atom appears in the composition of at least one molecular species, and the number of atoms in the composition of each species complies with the Definition 5.1. Therefore $\langle \Lambda, R, \Delta \rangle \in \text{SUBSET-FIXED-ATOMIC}$.

\Leftarrow : If $\langle \Lambda, R, \Delta \rangle \in \text{SUBSET-FIXED-ATOMIC}$, (5.3) ensures that each molecular species contains exactly two atoms. That is because for each $i \in [1, n]$,

$$\mathbf{d}_{S_i}(T) + \mathbf{d}_{S_i}(F) + \mathbf{d}_{S_i}(P) + \mathbf{d}_{S_i}(Q) \geq 2 \quad (5.6)$$

by Definition 5.1, so for each $m \in [1, k]$,

$$\sum_{j=1}^3 (\mathbf{d}_{S_{m_j}}(T) + \mathbf{d}_{S_{m_j}}(F) + \mathbf{d}_{S_{m_j}}(P) + \mathbf{d}_{S_{m_j}}(Q)) \geq 3 \times 2 = 6 \quad (5.7)$$

However the total number of atoms on the left hand side of 5.3 is exactly 6, meaning the equal sign has to be taken everywhere in (5.7) for any $m \in [1, k]$, forcing (5.6) to take equal sign as well. Similarly, (5.4) ensures $\|\mathbf{d}_{X_i}\|_1 = 2$ for each $i \in [1, n]$.

The reaction series (5.5) implies that each S_i has at least one P and each X_i has at least one Q . Furthermore,

1. if any S_i decomposes to $2P$, its corresponding X_i shall be composed of PQ , contradicting (5.4) which says that no X_i contains any P ;
2. if any S_i decomposes to PQ , it contradicts with (5.3) which says that no S_i contains any Q .

Therefore all S_i shall decompose to either $\{1F, 1P\}$ $((0, 1, 1, 0)^T)$ or $\{1T, 1P\}$ $((1, 0, 1, 0)^T)$. Correspondingly, X_i decompose to either $(0, 1, 0, 1)^T$ or $(1, 0, 0, 1)^T$.

Construct ϕ such that $\phi(v_j) = 1$ for all $v_j \in \{v_j \mid \mathbf{d}_{S_j} \upharpoonright \Delta = (1, 0, 1, 0)^T\}$, and $\phi(v_m) = 0$ for all $v_m \in V \setminus \{v_j \mid \mathbf{d}_{S_j} \upharpoonright \Delta = (1, 0, 1, 0)^T\}$. Exactly one in three of the products in the right hand side of (5.3) decomposes to $(1, 0, 1, 0)$, so exactly one in three of the variables (literals) in each clause of C evaluates to 1. It follows that $\langle V, C \rangle \in \text{MONOTONE-1-IN-3-SAT}$. \square

We notice that the coefficients of all species in all the reactions (5.3) \sim (5.5) are constants, so the numerical parameters – entries of each $(\mathbf{r}, \mathbf{p}) \in R$ – of the instances $\langle \Lambda, R, \Delta \rangle$ constructed above are bounded by the constant 3, which is again bounded by a polynomial of the length of $\langle V, C \rangle$, presuming the encoding scheme is “reasonable and concise” [18]. We therefore conclude that:

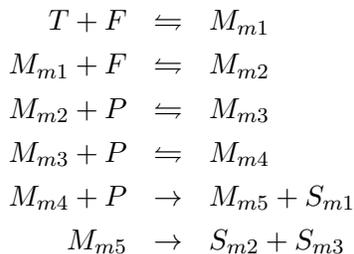
Corollary 5.6. *SUBSET-FIXED-ATOMIC is strongly NP-hard.*

Combining Corollary 5.6 and Corollary 5.3, we immediately obtain:

Theorem 5.7. SUBSET-FIXED-ATOMIC is strongly NP-complete. □

Remark 5.8. Another interesting thing to note is that SUBSET-FIXED-ATOMIC remains NP-complete even restricted to instances where R contains only unimolecular and bimolecular reactions.

In fact, one may convert, for each $m \in [1, k]$ (recall that k is the number of clauses in C), any reaction in the series (5.3) ($3P + 2F + T \rightarrow S_{m1} + S_{m2} + S_{m3}$ ($\forall m \in [1, k]$)) into the following series:



And apply similar methods to the X_i species. Such conversion creates $2 \times 10k = 20k$ extra reactions and $2 \times 5k = 10k$ intermediate species, which is polynomial in both time and space.

However, the case of SUBSET-ATOMIC is more complicated. We first note that in many cases, for some $\langle \Lambda, R \rangle \in \text{SUBSET-ATOMIC}$, there exist multiple possible Δ 's with respect to whom (Λ, R) is subset atomic. It remains open as to what complexity class SUBSET-ATOMIC belongs to, but we have the following conjecture:

Conjecture 5.9. SUBSET-ATOMIC is NP-complete.

6 Complexity of reachably atomic

Without regurgitating the intuition of the definition of reachably atomicity which has been explained in subsections 1.1 and 3.3, we proceed with the corresponding definition of languages for deciding reachably atomicity and the reachability problem in reachably atomic networks.

Definition 6.1. We define the following languages:

$$\begin{aligned}
 \text{REACHABLY-ATOMIC} &= \{ \langle \Lambda, R \rangle \mid (\exists \Delta \subseteq \Lambda) ((\Lambda, R) \text{ is reachably atomic with respect to } \Delta) \} \\
 \text{REACHABLY-FIXED-ATOMIC} &= \{ \langle \Lambda, R, \Delta \rangle \mid (\Lambda, R) \text{ is reachably atomic with respect to } \Delta \}
 \end{aligned}$$

Distinction between REACHABLY-FIXED-ATOMIC and REACHABLY-ATOMIC is analogous to that between SUBSET-FIXED-ATOMIC and SUBSET-ATOMIC. However, given the uniqueness of choice of Δ as well as the composition matrix \mathbf{D} (Lemma 3.9), there is no reason to distinguish between “REACHABLY-FIXED-ATOMIC” and REACHABLY-ATOMIC. Hence, for subsequent discussions we shall only consider REACHABLY-ATOMIC.

6.1 Reachably-Atomic is in P

As mentioned before, the requirement that $\{1S\} \Rightarrow \mathbf{d}_S \ (\forall S \in \Lambda)$ seems strong. The complexity results in this subsection confirm this.

Lemma 6.2. *If a network $\mathcal{C} = (\Lambda, R)$ is reachably atomic with respect to Δ via decomposition matrix \mathbf{D} (or equivalently, via the set of composition vectors $\{\mathbf{d}_S\}_{S \in \Lambda}$), then $\exists S \in \Lambda \setminus \Delta$ and $(\mathbf{r}, \mathbf{p}) \in R$ s.t. $\mathbf{r} = \{1S\}$ and $\mathbf{p} = \mathbf{d}_S$.*

Proof. The claim is saying that if a network is reachably atomic, then there exists a molecular species that can be decomposed into its atomic composition in a single reaction. Proof is done by supposing otherwise and chasing the composition sequence to find an infinite descending chain, contradicting to the finiteness of species set. Suppose not, then for all reactions with $\mathbf{r} = \{1S\}$ for some $S \in \Lambda \setminus \Delta$, either $[\mathbf{p}] \cap (\Lambda \setminus \Delta) \neq \emptyset$, or $[\mathbf{p}] \subseteq \Delta$ but $\mathbf{p} \neq \mathbf{d}_S$. The last case cannot happen, due to the uniqueness of atomic composition for reachably atomic networks (Recall Lemma 3.9). Hence for all (\mathbf{r}, \mathbf{p}) with $\mathbf{r} = \{1S\}$ for some $S \in \Lambda \setminus \Delta$, $[\mathbf{p}] \cap (\Lambda \setminus \Delta) \neq \emptyset$ [*].

[*], together with the reachable-atomicity, implies that for each $S \in \Lambda \setminus \Delta$ one may find a S' s.t. $\|\mathbf{d}_S\|_1 > \|\mathbf{d}_{S'}\|_1$ [**]. To see this, consider an arbitrary $S_i \in \Lambda \setminus \Delta$: any (\mathbf{r}, \mathbf{p}) with $\mathbf{r} = \{1S_i\}$ has either $\|\mathbf{p}\|_1 = 1$, or $\|\mathbf{p}\|_1 \geq 2$. In the second case we are done, for any $S_j \in [\mathbf{p}] \cap (\Lambda \setminus \Delta)$ satisfies $\|\mathbf{d}_{S_j}\|_1 < \|\mathbf{d}_{S_i}\|_1$; in the first case, we have found some S_{i+1} s.t. $\mathbf{d}_{S_{i+1}} = \mathbf{d}_{S_i}$ (and we call such (\mathbf{r}, \mathbf{p}) an *isomerization reaction*), so we recursively inspect into all $(\mathbf{r}_1, \mathbf{p}_1)$ with $\mathbf{r}_1 = \{1S_{i+1}\}$. Such a recursion must finally terminate with some S_{i+m} that satisfies $(\exists(\mathbf{r}_m, \mathbf{p}_m) \mid \mathbf{r}_m = \{1S_{i+m}\}) \|\mathbf{p}_m\|_1 \geq 2$, for otherwise \mathbf{d}_{S_i} would not be reachably decomposable into \mathbf{d}_{S_i} via any reaction sequence. It follows that any $S_{i+m+1} \in [\mathbf{p}_m] \cap (\Lambda \cap \Delta)$ satisfies $\|\mathbf{d}_{S_i}\|_1 > \|\mathbf{d}_{S_{i+m+1}}\|_1$.

We have argued that our initial assumption (for the sake of contradiction) implies [**]. But [**] would imply that there exists no molecular species with minimal size, contradicting the finiteness of Λ . \square

Theorem 6.3. REACHABLY-ATOMIC \in P.

Proof. Intuitively, we would need to exhibit a polynomial time algorithm to identify if there exists a separation of Λ into two non-empty, disjoint sets M and Δ , with elements in M decomposable via sequences of reactions into combination of elements in Δ . To achieve this goal, we set $M = \{S \in \Lambda \mid (\exists(\mathbf{r}, \mathbf{p}) \in R) \mathbf{r} = \{1S\}\}$, which would later be justified as the very set of molecular species; upon successfully obtaining such a separation (should it exist), we first check if there exists a molecular species decomposable into its atomic composition *via one single reaction*, which should exist if the network is reachably atomic, by Lemma 6.2; then recursively, we check if there exist elements in M that can be decomposed into combination of atoms and previously identified decomposable molecules *via one single reaction*. If this process terminates, finding (candidate) atomic composition for all molecules, we proceed to check if the compositions preserve molecules and if each atom appears in the composition of at least one molecule. We only accept when the network passes all those tests. describe this polynomial time algorithm in more details, argue for its correctness, and then exhibit the pseudo-code.

Our algorithm will compose of the following steps:

1. Enumerate all reactions in R : for each reaction starting with $\{1S\}$ for some $S \in \Lambda$, put S in the set M of molecules. After the enumeration, define $\Delta = \Lambda \setminus M$. If $\Delta = \emptyset$ or $M = \emptyset$, reject.

2. Find if there exists some molecular species $S \in M$ such that S is decomposable into some \mathbf{d}_S with $[\mathbf{d}_S] \subseteq \Delta$ by execution of a single reaction. If any of such \mathbf{d}_S is of size 0 or 1, reject. Keep track of each composition vector \mathbf{d}_S ;

Make a subset M' of M s.t. M' contains all molecular species which CANNOT be decomposed into \mathbf{d}_S with $[\mathbf{d}_S] \subseteq \Delta$ by execution of one reaction. If $M' = M$, then reject.

3. while M' is not empty, iterate and try to find an $S \in M'$ that satisfies this condition: $\exists(\mathbf{r}, \mathbf{p}) \in R$ s.t. $\mathbf{r} = \{1S\}$ and $[\mathbf{p}] \subseteq (M \setminus M') \cup \Delta$. Note that all elements S'' in $M \setminus M'$ satisfies $\{1S''\} \Rightarrow^* \mathbf{d}_{S''}$, hence if S can be decomposed into complexes consisting solely of elements in $(M \setminus M') \cup \Delta$ via execution of one reaction, then S itself satisfies $\{1S\} \Rightarrow^* \mathbf{d}_S$ as well. Keep track of \mathbf{d}_S and exclude such S from M' .

If in some iteration we cannot find any such $S \in M'$, then reject; else, the iteration will finally halt excluding all such S 's and making M' empty.

4. We have obtained \mathbf{d}_S for each $S \in \Lambda$ (note that $\mathbf{d}_A = \mathbf{e}_A$ for all $A \in \Delta$) s.t. $[\mathbf{d}_S] \subseteq \Delta$ and $\{1S\} \Rightarrow^* \mathbf{d}_S$. By Lemma 3.9, such set of composition is unique. Now, check if this composition conserves atoms by composing stoichiometric matrix \mathbf{A} and composition matrix \mathbf{D} and examine if $\mathbf{A} \cdot \mathbf{D} = \mathbf{0}$, and finally check if each atom A appears at least once in some \mathbf{d}_S for some $S \in M$.

We first prove that if reachably atomic then the algorithm will halt in ACCEPT:

if \mathcal{C} is reachably atomic, then we claim that the set $\{S \in \Lambda \mid \exists(\mathbf{r}, \mathbf{p}) \in R \text{ s.t. } \{1S\} = \mathbf{r}\}$ is exactly the set of molecules M , with its complement $\Delta = \Lambda \setminus M$. To see this, recall that we prescribed there being no " $\mathbf{r} \rightarrow \mathbf{r}$ " reactions in R , so all reactions $(\mathbf{r}, \mathbf{p}) \in R \mid \|\mathbf{r}\| = 1$ is either an isomerization reaction ($\|\mathbf{p}\|_1 = 1, \mathbf{p} \neq \mathbf{r}$) or disassociation reaction ($\|\mathbf{p}\|_1 \geq 2$). Both types of reactions can only happen when $S \in [\mathbf{r}]$ is a molecule; it follows that $\{S \in \Lambda \mid \exists(\mathbf{r}, \mathbf{p}) \in R \text{ s.t. } \{1S\} = \mathbf{r}\} \subseteq M$. Conversely, when $S \in M$, reachably atomicity gives $S \in \{S \in \Lambda \mid \exists(\mathbf{r}, \mathbf{p}) \in R \text{ s.t. } \{1S\} = \mathbf{r}\}$.⁹

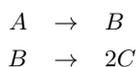
Neither M nor Δ would be empty, for $(\Delta = \emptyset) \Rightarrow (\mathcal{C} \text{ is not reachably atomic})$ and $(M = \emptyset) \Rightarrow (R = \emptyset)$. Hence the algorithm passes Step 1, correctly identifying the partition (M, Δ) of Λ .

By Lemma 6.2, reachable-atomicity implies that the algorithm will find at least one molecular species S that directly decomposes to its atomic composition \mathbf{d}_S and grantedly, $\|\mathbf{d}_S\| \geq 2$, so Step (2) will be passed.

Further, applying the same argument in Lemma 6.2 on the set M' , the while loop must shrink the cardinality of M' by at least 1 per iteration, and finally exit by making M' empty,¹⁰ passing Step (3);

Finally, the composition must preserve atoms for all reactions, and all atoms must appear in the composition of at least one molecule, which make both tests in Step (4) passed.

⁹ We point out that the set of atoms $M \neq \{S \in \Lambda \mid \exists(\mathbf{r}, \mathbf{p}) \in R \text{ s.t. } (\{1S\} = \mathbf{r}) \wedge (\|\mathbf{p}\| \geq 2)\}$, so we have to test the $\|\mathbf{d}_S\| \geq 2$ condition in later steps. This is because it might be the case that the only reaction (\mathbf{r}, \mathbf{p}) with $\mathbf{r} = \{1S\}$ turns out to be an isomerization reaction. A counter example would be:



By our definition $M = \{S \in \Lambda \mid \exists(\mathbf{r}, \mathbf{p}) \in R \text{ s.t. } \{1S\} = \mathbf{r}\}$, we shall correctly identify $M = \{A, B\}$, yet the added condition $\|\mathbf{p}\| \geq 2$ would make $M = \{B\}$, a mis-identification.

¹⁰That is, if $(\forall S \in M')(\forall(\mathbf{r}, \mathbf{p}) \in R) (\mathbf{r} = \{1S\} \Rightarrow [\mathbf{p}] \cap M' \neq \emptyset)$, then for each species S in M' there will be $S' \in M'$ s.t. $\|S\|_1 > \|S'\|_1$, contradicting the finiteness of M' .

It remains to show that if \mathcal{C} is not reachably atomic then our algorithm will halt in REJECT. We claim that: if \mathcal{C} is not reachably atomic, then exactly one of the following will be true:

1. There is no valid separation of Λ into M and Δ . That is, either $\{S \in \Lambda \mid \exists(\mathbf{r}, \mathbf{p}) \in R \text{ s.t. } \{1S\} = \mathbf{r}\} = \emptyset$ (no species is the single reactant of an isomerization or disassociation reaction, so no species S is decomposable from $\{1S\}$), or $\{S \in \Lambda \mid \exists(\mathbf{r}, \mathbf{p}) \in R \text{ s.t. } \{1S\} = \mathbf{r}\} = \Lambda$ (every species is the single reactant of some isomerization or disassociation reaction, which contradicts the definition of subset atomicity for atoms should be neither isomerizable nor decomposable). An example where $\{S \in \Lambda \mid \exists(\mathbf{r}, \mathbf{p}) \in R \text{ s.t. } \{1S\} = \mathbf{r}\} = \emptyset$ would be $(\Lambda = \{A, B, C\}, R = \{2A + 3B \rightarrow 4C\})$, while $(\Lambda' = \{A, B\}, R' = \{A \rightarrow B, B \rightarrow A\})$ would be an instance where $\{S \in \Lambda \mid \exists(\mathbf{r}, \mathbf{p}) \in R \text{ s.t. } \{1S\} = \mathbf{r}\} = \Lambda$.

Observe such a valid separation (M, Δ) of Λ , if existing, is unique for a certain $\mathcal{C} = (\Lambda, R)$, since $S \in M$ if and only if $\exists(\mathbf{r}, \mathbf{p}) \in R \text{ s.t. } \mathbf{r} = \{1S\}$, and this property is uniquely decided by R .

2. There exists the unique valid separation (M, Δ) of Λ , but there exists no molecular species directly decomposable into its atomic composition via execution of one single reaction. That is, $(\forall S \in M)(\forall(\mathbf{r}, \mathbf{p}) \in R)(\mathbf{r} = \{1S\} \Rightarrow [\mathbf{p}] \cap M \neq \emptyset)$. An example of this is $(\Lambda = \{A, B, C\}, R = \{A \rightarrow B + 5C, B \rightarrow A + 5C\})$. We would successfully identify $M = \{A, B\}, \Delta = \{C\}$, but for all reactions we $(\mathbf{r}, \mathbf{p}) \in R$ have $[\mathbf{p}] \cap M \neq \emptyset$.

3. There exists the unique valid separation (M, Δ) of Λ and $\{S \in M \mid (\exists(\mathbf{r}_S, \mathbf{p}_S) \in R)((\mathbf{r}_S = \{1S\}) \wedge ([\mathbf{p}_S] \subseteq \Delta))\} \neq \emptyset$, but for some $S \in \{S \in M \mid (\exists(\mathbf{r}_S, \mathbf{p}_S) \in R)((\mathbf{r}_S = \{1S\}) \wedge ([\mathbf{p}_S] \subseteq \Delta))\}$, $\|\mathbf{p}\| \leq 1$. That is, we have some reaction $S_1 \rightarrow A_1$ with $S_1 \in M$ and $A_1 \in \Delta$, or $S_1 \rightarrow \emptyset$. In this case, either a molecule decomposes to a single atom, or it vanishes.

Typical Examples are: $\mathcal{C}_1 = (\{A, B, C\}, \{A \rightarrow B, B \rightarrow C\})$, $\mathcal{C}_2 = (\{A, B, C\}, \{A \rightarrow 2C, B \rightarrow \emptyset\})$. In both cases one would identify $M_{\mathcal{C}_1} = M_{\mathcal{C}_2} = \{A, B\}$; for both networks, $\{S \in M \mid (\exists(\mathbf{r}_S, \mathbf{p}_S) \in R)((\mathbf{r}_S = \{1S\}) \wedge ([\mathbf{p}_S] \subseteq \Delta))\} = \{A, B\}$. But In \mathcal{C}_1 , B decomposes to a single atom C ; in \mathcal{C}_2 , B vanishes. This disqualifies both sets from being reachably atomic by placing them in the third case.

4. There exists the unique valid separation (M, Δ) of Λ , and $\{S \in M \mid (\exists(\mathbf{r}_S, \mathbf{p}_S) \in R)((\mathbf{r}_S = \{1S\}) \wedge ([\mathbf{p}_S] \subseteq \Delta))\} \neq \emptyset$; further, each $S \in \{S \in M \mid (\exists(\mathbf{r}_S, \mathbf{p}_S) \in R)((\mathbf{r}_S = \{1S\}) \wedge ([\mathbf{p}_S] \subseteq \Delta))\}$ satisfies $\|\mathbf{p}\|_1 \geq 2$. However, there exists some indecomposable molecular species. That is, $\exists S' \in M \text{ s.t. } \forall \mathbf{c} \in \mathbb{N}^\Lambda \text{ where } \{1S'\} \Rightarrow^* \mathbf{c}, [\mathbf{c}] \cap M \neq \emptyset$.

An example for this case is $\mathcal{C} = (\{A, B, C, D, E\}, \{A \rightarrow B, B \rightarrow C, D \rightarrow 3E\})$. One may identify $M = \{A, B, D\}$ and find $\{D\} = \{S \in M \mid (\exists(\mathbf{r}, \mathbf{p}) \in R)((\mathbf{r} = \{1S\}) \wedge ([\mathbf{p}] \subseteq \Delta))\}$. Further, the reaction $D \rightarrow 3E$ where $\mathbf{r} = \{1D\}$ satisfies $\|\mathbf{p}\| = 3$. This network does not belong to any of the first few cases but it does belong to Case 4, for $\forall \mathbf{c}$ where $\{1A\} \Rightarrow^* \mathbf{c}, [\mathbf{c}] \subseteq \{B, C\} \subseteq M$.

5. There is a unique valid composition (M, Δ) of Λ and $(\forall S \in M) (\exists \mathbf{c}_S \text{ with } [\mathbf{c}_S] \subseteq \Delta) (\{1S\} \Rightarrow^* \mathbf{c}_S) \wedge (\|\mathbf{c}_S\|_1 \geq 2)$, but the composition does not preserve atoms for some reaction. That is, with \mathbf{A} the stoichiometric matrix and \mathbf{D} the composition matrix (row vectors being the \mathbf{c}_S 's restricted to Δ), $\mathbf{A} \cdot \mathbf{D} \neq \mathbf{0}$. Note that for reachably atomic networks, atomic composition vectors (or equivalently, matrix) should be unique.

One example of this would be $(\{A, B, C, D\}, \{A \rightarrow B, B \rightarrow 3C, A + B \rightarrow 5C + D\})$. Here we have $M = \{A, B\}$ and $\{3C\} \underset{B \rightarrow 3C}{=} \mathbf{d}_B \underset{A \rightarrow B}{=} \mathbf{d}_A \underset{A+B \Rightarrow 5C+D, B \rightarrow 3C}{=} \{2C + D\}$, contradicting the preservation of atoms. Note that this happens to be another example where a network is mass conserving (Just set $\mathbf{m}(A) = \mathbf{m}(B) = 3\mathbf{m}(C) = 3\mathbf{m}(D) = 3$) but not subset atomic (and hence not reachably atomic).

6. There is a unique valid composition (M, Δ) of Λ and $(\forall S \in M) (\exists \mathbf{c}_S$ with $[\mathbf{c}_S] \subseteq \Delta)$ $(\{1S\} \Rightarrow^* \mathbf{c}_S) \wedge (\|\mathbf{c}_S\|_1 \geq 2)$, and the composition preserves atoms ($\mathbf{A} \cdot \mathbf{D} = 0$), but some atoms are "redundant": $\exists A \in \Delta$ s.t. $\forall S \in M, A \notin [\mathbf{d}_S] = [\mathbf{c}_S]$.

One may "fix" the last example into this case: $(\{A, B, C, D\}, \{A \rightarrow B, B \rightarrow 3C, A + B \rightarrow 6C\})$. Here we have $M = \{A, B\}$ and $\{3C\} \underset{B \rightarrow 3C}{=} \mathbf{d}_B \underset{A \rightarrow B}{=} \mathbf{d}_A \underset{A+B \Rightarrow 6C, B \rightarrow 3C}{=} \{3C\}$, so every condition for reachably atomic is satisfied, except that $(\forall S) D \notin [\mathbf{d}_S]$.

All six cases contradict each other, so *at most* one case could hold; on the other hand, the negation of the disjunction of all six cases says that there exists a non-empty proper subset of Λ and a composition matrix \mathbf{D} , such that all three conditions of primitive atomicity holds with respect to Λ via \mathbf{D} , and $\{1S\} \Rightarrow^* \mathbf{d}_S$ for all S . This is exactly the definition of reachably atomicity. So taking contraposition, non-reachable-atomicity implies *at least* one of the six cases hold.

Instances satisfying Case (1) will be rejected in Step (1), while Cases (2) and (3) will get rejected in Step (2). In case (4), the loop for finding composition vectors must terminate before M' is emptied, so it will get rejected by Step (3); Cases (5) and (6) triggers rejection in Step (4).

We exhibit the following pseudocode for the decider:

```

1 Initialize global set  $M, M', \Delta, D = \emptyset$  //  $D$ : {composition vectors}.
2 for  $(\mathbf{r}, \mathbf{p}) \in R$  do
3   | if  $(\exists S \in \Lambda) \mathbf{r} = \{1S\}$  then
4   |   |  $M \leftarrow M \cup \{S\}$ 
5   |   end
6   end
7  $\Delta \leftarrow \Lambda \setminus M$ 
8  $M' \leftarrow M$ 
9 if  $M = \emptyset$  or  $\Delta = \emptyset$  then
10  | REJECT
11 end
12  $D \leftarrow D \cup \{\mathbf{e}_A\}_{A \in \Delta}$ 
13 for  $(\mathbf{r}, \mathbf{p}) \in R$  where  $(\exists S \in M) \mathbf{r} = \{1S\}$  do
14  | if  $[\mathbf{p}] \subseteq \Delta$  then
15  |   | if  $\|\mathbf{p}\|_1 \leq 1$  then
16  |   |   | REJECT
17  |   |   end
18  |   end
19  |    $D \leftarrow D \cup \{\langle \mathbf{d}_S = \mathbf{p} \rangle\}$ 
20  |    $M' \leftarrow M' \setminus \{S\}$ 
21 end
22 if  $M' = M$  then
23  | REJECT
24 end
25 while  $M' \neq \emptyset$  do
26  | if  $(\forall S \in M') (\forall (\mathbf{r}, \mathbf{p}) \in R \mid \mathbf{r} = \{1S\}) ([\mathbf{p}] \cap M' \neq \emptyset)$  then
27  |   | REJECT
28  |   end
29  |   else
30  |     for  $(S \in M' \mid (\exists (\mathbf{r}, \mathbf{p}) \in R \mid \mathbf{r} = \{1S\} \text{ and } [\mathbf{p}] \cap M' = \emptyset))$  do
31  |       |  $D \leftarrow D \cup \{\langle \mathbf{d}_S = \sum_{S' \in [\mathbf{p}]} \mathbf{d}_{S'} \rangle\}$ 
32  |       |  $M' \leftarrow M' \setminus \{S\}$ 
33  |       end
34  |     end
35 end
36 Compose  $\mathbf{A}$  (stoichiometric matrix) and  $\mathbf{D}$  (composition matrix)
37 if  $\mathbf{A} \cdot \mathbf{D} \neq \mathbf{0}$  then
38  | REJECT
39 end
40 if  $(\exists A \in \Delta) (\forall S \in M) A \notin [\mathbf{d}_S]$  then
41  | REJECT
42 end
43 ACCEPT

```

Algorithm 2: Decider for REACHABLY-ATOMIC

Let us briefly mention that this is a polynomial time algorithm. The first `for`-loop takes $O(|R|)$ time; the second `for`-loop takes at most $O(|R|)$ iterations, and each iteration takes $O(|\Lambda|^3)$ time; as for the `while` loop, note that it either shrinks the size of M' by 1 per iteration, or `REJECTS`. Hence the `while` loop takes at most $O(|\Lambda|)$ to exit. The `if`-statement inside the `while`-loop takes $O(|\Lambda| \cdot |R| \cdot |\Lambda|^2)$ to evaluate. Lastly, composing and multiplying $\mathbf{A} \cdot \mathbf{D}$ takes $O(|R||\Lambda| \cdot |\Lambda|^2)$ time, and verifying each $A \in \Delta$ is “used” by the composition of some molecule is $O(|\Lambda| \cdot |\Lambda| \cdot |\Lambda|)$. The times complexity is therefore dominated by the `while`, which is $O(|R||\Lambda|^4)$. No input, output or intermediate encoding takes more than polynomial space to record, so `REACHABLY-ATOMIC` $\in \mathbf{P}$, as desired. \square

6.2 Reachable-Reach is PSPACE-complete

We begin this section by introducing the definition of configuration reachability graphs, followed by a result proved in relevant literature [30] (see also Subsection 1.1), the idea of which we shall be using to prove the PSPACE-completeness of `REACHABLE-REACH`.

Definition 6.4 (Configuration Reachability Graph). *An \mathbf{i} -initiated Configuration Reachability Graph $G_{C,\mathbf{i}}$ of the chemical reaction network $C = (\Lambda, R)$ is a directed graph (V, E) , where:*

1. each $v_{\mathbf{c}} \in V$ ($\mathbf{c} \in \mathbb{N}^\Lambda$) is labeled by a reachable configuration \mathbf{c} of C ;
2. $v_{\mathbf{i}} \in V$ ($\mathbf{i} \in \mathbb{N}^\Lambda$) is the vertice labeled by the initializing configuration \mathbf{i} ;
3. the ordered pair $(v_{\mathbf{c}_1}, v_{\mathbf{c}_2}) \in E$ if and only if $\mathbf{c}_1 \Rightarrow^1 \mathbf{c}_2$.

Remark 6.5. For the sake of simplicity, we use $G_{C,\mathbf{i}}$ as shorthand for $G_{C,v_{\mathbf{i}}}$.

For the same C , Configuration Reachability Graphs can be far from isomorphic due to parameterization by different initialization vectors.

Example 6.6. Consider $C = (\Lambda, R)$ where $\Lambda = \{S_1, S_2, S_3, A_1\}$ (we use the order exhibited above for \mathbb{N}^Λ), and R , in its explicit form, consists of



Now consider two initialization vectors, $\mathbf{i} = (0, 2, 1, 0)^T$ (that is, $\{2S_2, 1S_3\}$) and $\mathbf{i}' = (2, 1, 3, 0)^T$ (i.e., $\{2S_1, 1S_2, 3S_3\}$). For \mathbf{i} , the only possible reaction is (6.2) and the only possible reaction path is two consecutive executions of (6.2). Hence $G_{C,\mathbf{i}} = (V, E)$ where $V = \{v_{\mathbf{i}}, v_{(0,1,1,9)^T}, v_{(0,0,1,18)^T}\}$ and $E = \{(v_{\mathbf{i}}, v_{(0,1,1,9)^T}), (v_{(0,1,1,9)^T}, v_{(0,0,1,18)^T})\}$.

The case for \mathbf{i}' is more complicated. Potential execution paths include: (6.3); (6.1) \rightarrow (6.1) \rightarrow (6.2); (6.1) \rightarrow (6.2) \rightarrow (6.1); (6.2) \rightarrow (6.1) \rightarrow (6.1). We may construct the Configuration Reachability Graph $G_{C,\mathbf{i}'}$ following these paths. Figures of $G_{C,\mathbf{i}}$ and $G_{C,\mathbf{i}'}$ are shown below:

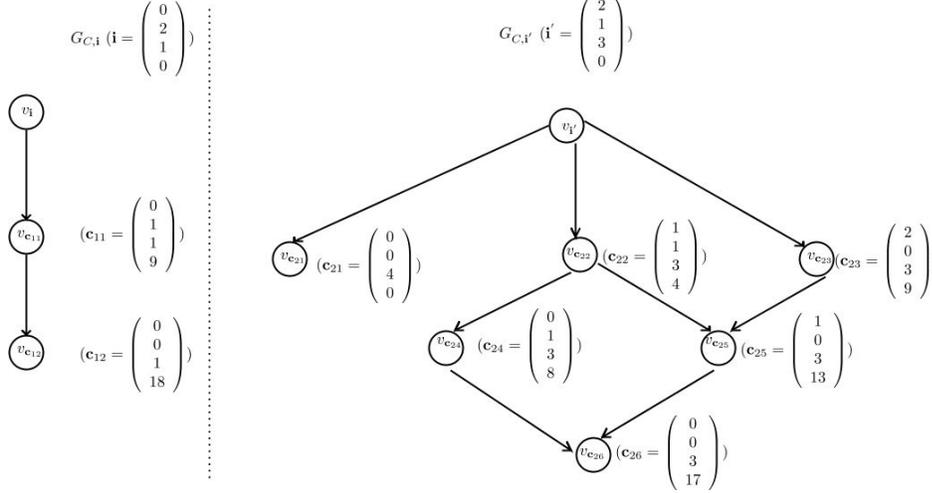


Figure 1: Configuration Reachability Graphs $G_{C,i}, G_{C,i'}$

We are about ready to prove the conclusion on the complexity of the reachability problem for reachably atomic networks. But first, we point out that the following Theorem is a direct result of a finding in [30], giving the complexity class of reachability problems for mass-conserving chemical reaction networks.

Theorem 6.7. *For all mass conserving chemical reaction networks \mathcal{C} and initial configuration \mathbf{i} of \mathcal{C} , $|\langle G_{C,i} \rangle| \in O(2^{\text{poly}(|\langle \mathcal{C}, \mathbf{i} \rangle|)})$. That is, the binary size of the encoding of the configuration reachability graph $G_{C,i}$ is at most exponential to the binary size of the encoding of the pair $(\mathcal{C}, \mathbf{i})$.*

Furthermore, reachability problem for mass conserving networks is PSPACE-complete. That is, it is PSPACE-complete to decide if an instance is in the following language:

$$\{ \langle \Lambda, R, \mathbf{c}_1, \mathbf{c}_2 \mid (\Lambda, R) \text{ is mass conserving; } \mathbf{c}_1, \mathbf{c}_2 \in \mathbb{N}^\Lambda; \mathbf{c}_1 \Rightarrow^* \mathbf{c}_2 \rangle \}$$

Built on Theorem 6.7, we now exhibit the proof that the decision problem “Given a Reachably Atomic network, is \mathbf{c}_2 reachable from \mathbf{c}_1 ” is PSPACE-Complete. More formally,

Definition 6.8. *We define the language*

$$\text{REACHABLE-REACH} = \{ \langle \Lambda, R, \mathbf{c}_1, \mathbf{c}_2 \rangle \mid (\Lambda, R) \text{ is reachably atomic; } \mathbf{c}_1, \mathbf{c}_2 \in \mathbb{N}^\Lambda; \mathbf{c}_1 \Rightarrow^* \mathbf{c}_2 \}$$

Referencing the proof ideas and conclusions in Theorem 6.7, we shall show the complexity class of REACHABLE-REACH below, using the exponential size bound on configuration reachability graph of reachably atomic networks and a simulation of polynomial space Turing Machine by reachably atomic network.

Theorem 6.9. *REACHABLE-REACH is PSPACE-complete.*

Proof. We prove the containment in PSPACE by the PSPACE-completeness of reachability problem in a mass conserving chemical reaction network (Theorem 6.7). Hardness is shown by simulating polynomial space Turing Machines. Let $\langle \Lambda, R, \mathbf{c}_1, \mathbf{c}_2 \rangle$ be an instance, and let $Z := |\langle \Lambda, R, \mathbf{c}_1, \mathbf{c}_2 \rangle|$. In Theorem 6.3 we proved that REACHABLY-ATOMIC $\in P$ so we may run the polynomial decider on $\langle \Lambda, R \rangle$ and REJECT if $\langle \Lambda, R \rangle \notin \text{REACHABLY-ATOMIC}$.

If the REACHABLY-ATOMIC decider halts in ACCEPT, we would obtain $\Delta \subseteq \Lambda$ with respect to which (Λ, R) is reachably atomic, as well as the set $\{\mathbf{d}_S\}_{S \in \Lambda}$ of composition vectors. Further, we would have confirmed that $\mathcal{C} = (\Lambda, R)$ is mass-conserving, for this is implied by reachably atomicity. Recall from [30] that the number of vertices in a configuration reachability graph $G_{\mathcal{C}, \mathbf{c}_1}$ for mass-conserving network \mathcal{C} is at most exponential to the binary size of the input. Now, let $n = |V_{G_{\mathcal{C}, \mathbf{c}_1}}|$, then by Savitch's Theorem [36],

$$\text{REACHABLE-REACH} \in \mathbf{SPACE}((\log n)^2) = \mathbf{SPACE}((\log(O(2^{\text{poly}(Z)})))^2) = \mathbf{SPACE}(O(\text{poly}(L))^2)$$

It follows that REACHABLE-REACH \in PSPACE.

As for the PSPACE hardness, we shall prove by simulating a polynomial-space Turing Machine. That is, consider the language

$$L := \{\langle M, x, 0^{|x|^c} \rangle \mid M \text{ is an } O(|x|^c) \text{ - space, clocked Turing Machine, } x \in \{0, 1\}^* : M(x) \rightarrow 1\}$$

Just to clarify the notation, “ $M(x) \rightarrow 1$ ” means M on the input x runs for $O(|x|^c)$ time and ACCEPTS. We shall construct an REACHABLE-REACH instance $\langle \Lambda, R, \mathbf{c}_1, \mathbf{c}_2 \rangle$ by a polynomial time reduction from an instance $\langle M, x, 0^{|x|^c} \rangle$, and show that $\langle M, x, 0^{|x|^c} \rangle \in L$

if and only if $\langle \Lambda, R, \mathbf{c}_1, \mathbf{c}_2 \rangle \in \text{REACHABLE-REACH}$.

Without loss of generality, assume the initial configuration of M is $q_1 \in Q_M = \{q_1, q_2, \dots, q_{t-2}, q_A, q_R\}$ (where $t := |Q_M|$, Q_A is the accept state and Q_R is the reject state), and that the TM blank the tape cells and return the tape head to the leftmost position before halting. Let p denote the maximum number of tape cells that M may use on input x (Note that $p \in O(|x|^c)$). Define the following set of species:

$$\Lambda = \{A, \underbrace{Q_1, \dots, Q_{t-2}, Q_A, Q_R}_{\text{machine states}}, \underbrace{P_1, \dots, P_p}_{\text{head positions}}, \underbrace{T_1^0, T_1^1, \dots, T_p^0, T_p^1}_{\text{tape contents}}\}$$

and configurations:

$$\begin{aligned} \mathbf{c}_1 &= \{P_1, Q_1, T_1^{x[1]}, \dots, T_{|x|}^{x[|x|]}, T_{|x|+1}^\sqcup, \dots, T_p^\sqcup\} \\ \mathbf{c}_2 &= \{1Q_A, T_1^\sqcup, \dots, T_p^\sqcup, P_1\} \end{aligned}$$

Recalling that the transition function $\delta_M : Q_M \setminus \{Q_A, Q_R\} \times \Gamma \rightarrow Q_M \times \Gamma \times \{-1, +1\}$, we construct the set R of reactions in the following way:

```

1 for ( $\forall q_i \in Q_M$ ) ( $\forall b \in \{0, 1, \sqcup\}$ ) ( $\forall k \in \{1, 2, \dots, p\}$ ) do
2   | if  $\delta(q_i, b) \rightarrow (q_j, b', m)$  then
3   |   | Add Reaction  $Q_i + T_k^b + P_k \rightarrow Q_j + T_k^{b'} + P_{k+m} // m \in \{\pm 1\}$  : tape head moving
4   |   | direction.
5   | end
6 for  $S \in \Lambda \setminus \{A\}$  do
7   | Add Reaction  $S \rightarrow 2A$ 
8 end

```

Algorithm 3: Construction of R

Observe that (Λ, R) is a reachably atomic network with respect to $\Delta = \{A\} \subseteq \Lambda$, for any molecular species can be decomposed to $\{2A\}$ explicitly via Lines 11-12, A appears in the composition of all molecular species, and all reactions preserve the number of atoms.

Further,

$$\begin{aligned}
\langle M, X, 0^{|x|^c} \rangle \in L &\Leftrightarrow M(x) \rightarrow 1 \\
&\Leftrightarrow \exists \text{computation path } (q_1, (x[1], x[2], \dots, x[|x|], \underbrace{\sqcup, \dots, \sqcup}_{p-|x|})) \Rightarrow^* (q_A, (\underbrace{\sqcup, \sqcup, \dots, \sqcup}_p)) \\
&\Leftrightarrow \mathbf{c}_1 \Rightarrow^* \mathbf{c}_2 \\
&\Leftrightarrow \langle \Lambda, R, \mathbf{c}_1, \mathbf{c}_2 \rangle \in \text{REACHABLE-REACH}
\end{aligned}$$

Finally, $|\Lambda| = 1 + t + 3p$; $|R| \in O(3pq + |\Lambda|)$, $\|\mathbf{c}\|_1 = \|\mathbf{c}\|_2 = p + 2$. All coefficients of reactions are constant Hence this reduction is polynomial in Z both timewise and spacewise. \square

Remark 6.10. The fact that all coefficients of reactions are constant also implies that REACHABLE-REACH is PSPACE-hard (and hence complete) in the strong sense.

Remark 6.11. The argument above fails if we require all reactions to be reversible. In fact, if that were the case, one may run the disassociation reactions first and run the reversed disassociation reactions to create any “input species” (species representing the tape cell) as wanted, breaking the equivalence.

As mentioned in Subsection 1.1, there are some interesting relationship between some property of network defined in [21] and ours, which we shall look into in this section.

We begin by introducing some definitions in [21]. We disclaim that all the following definitions and notations (but not remarks) before Subsection (6.3) are from [21], possibly with slight modification of notations and/or interpretations:

Definition 6.12. A species composition map, or simply a composition of chemical reaction network \mathcal{C} is a map $\mathcal{E} : \Lambda \rightarrow \mathbb{N}^n \setminus \{0^n\}$, where $n \in \mathbb{N}_{>0}$.

Remark 6.13. We donot confuse \mathcal{E} with \mathbf{d} because they are completely different mappings. In particular, \mathcal{E} could map different species to the same \mathbf{e}_i , which means $\mathcal{E} \upharpoonright \mathcal{E}^{-1}(\mathbf{e}_i)$ could be non-injective, while we donot allow this for atomic composition \mathbf{d} . This point is further illustrated in the following Lemmas, such as Lemma 6.23.

Definition 6.14.

1. A species $S \in \Lambda$ is \mathcal{E} -elementary if $\|\mathcal{E}(S)\|_1 = 1$;
2. A species $S \in \Lambda$ is \mathcal{E} -composite if $\|\mathcal{E}(S)\|_1 \geq 2$;
3. $S, S' \in \Lambda$ are \mathcal{E} -isomeric if $\mathcal{E}(S) = \mathcal{E}(S')$; equivalently, we say S, S' belong to the same \mathcal{E} -isomeric class.

Definition 6.15. The extended composition $\tilde{\mathcal{E}}$ of \mathcal{E} is defined as in Equation (6.5).

We denote $\Theta = \text{span}(\{(\mathbf{p} - \mathbf{r})\}_{(\mathbf{r}, \mathbf{p}) \in R}) \subseteq \mathbb{R}^\Lambda$. With slight abuse of notation, we sometimes also write $\Theta = \text{span}(R)$ with R the set of reaction vectors.

Definition 6.16. a composition \mathcal{E} is near-core if:

1. $\ker(\tilde{\mathcal{E}}) \supseteq \Theta$, which is equivalent to saying \mathcal{C} is \mathcal{E} -conservative; and

2. $\mathbf{e}_1, \dots, \mathbf{e}_n \in \text{range}(\mathcal{E})$.

Definition 6.17. A composition \mathcal{E} is core if \mathcal{E} is near-core and further, $\ker(\tilde{\mathcal{E}}) = \Theta$.

Remark 6.18. Intuitively, a core composition \mathcal{E} is the composition whose linear extension $\tilde{\mathcal{E}}$ has the **smallest** kernel containing the reaction vector space $\Theta = \text{span}(R)$ as subspace. That ensures the uniqueness (up to isomorphism) of $\tilde{\mathcal{E}}$, and avoids including vectors that are not reachable by reactions into the kernel. Theorem 2.11, 2.12 of [21] has a detailed and more formal discussion on this matter.

Definition 6.19. A reaction network \mathcal{C} is constructive if it admits a core composition.

Definition 6.20.

1. A species Y is explicitly constructible (resp. explicitly destructible) if there are isomerization reactions $Y_0 \rightarrow \dots \rightarrow Y_l$ (resp. $Y_l \rightarrow \dots \rightarrow Y_0$), where $l \in \mathbb{N}$, such that Y_0 is the target of a binding reaction (resp. the source of a dissociation reaction) and $Y_l = Y$.

Binding reactions are $(\mathbf{r}, \mathbf{p}) \in R$ s.t. $\|\mathbf{r}\|_1 \geq 2$ and $\|\mathbf{p}\|_1 = 1$, and dissasociation reactions have similar definition with \mathbf{r} and \mathbf{p} swapped. Isomerizations are $(\mathbf{r}, \mathbf{p}) \in R$ s.t. $|\mathbf{r}|_1 = |\mathbf{p}|_1 = 1$: note that the reactant and product of an isomerization reaction are isomers.

2. A species X is explicitly constructive (resp. explicitly destructive) if there is a binding reaction $Q \rightarrow Y$ (resp. a dissociation reaction $Y \rightarrow Q$ such that $X \in [Q]$).

Remark 6.21. Intuitively, a species is explicit constructible if it is “eventually” a product of a binding reaction (up to having some isomerization reactions in between), while explicit constructivity means a species directly participate in a binding reaction as reactant.

And finally,

Definition 6.22. A chemical reaction network \mathcal{C} is explicitly-reversibly constructive if:

1. \mathcal{C} is constructive;
2. Each composite species is both explicitly constructible and explicitly destructible; and
3. Each elementary species is both explicitly constructive and explicitly destructible.

6.3 Atomic chemical reaction networks with core or near core compositions

We would like to begin this section by showing an equivalence relationship between our definition of subset atomicity and [21]’s definition of networks admitting near-core compositions with certain restrictions.

Lemma 6.23. A chemical reaction network $\mathcal{C} = (\Lambda, R)$ is subset atomic if and only if $\exists n \in \mathbb{N}_{>0}, \mathcal{E} : \Lambda \rightarrow \mathbb{N}^n \setminus \{0\}^n$, s.t.

1. \mathcal{E} is a near-core composition of \mathcal{C} ;
2. $\mathcal{E} \upharpoonright \mathcal{E}^{-1}(\{\mathbf{e}_i\}_{i=1}^n)$ is one to one, and

3. $(\forall i \in [1, n]) (\exists S \in \Lambda \setminus \bigcup_{i=1}^n \mathcal{E}^{-1}(\mathbf{e}_i)) (\mathcal{E}(S))_i > 0$.

Remark 6.24. Adopting the definition that $\mathcal{X}_i := \{S \in \Lambda \mid \mathcal{E}(S) = \mathbf{e}_i\}$, then condition 2 is saying that $p_i := |\mathcal{X}_i| = 1$ for each i . It intuitively translates to “no isomerization is allowed for \mathcal{E} -elementary species”. Note also the similarity between Condition 3 above and Condition 2 of Definition 3.1. They will translate to each other by construction in the proof below.

Proof. We note that atomic composition in a subset atomic network \mathcal{C} describes a similar phenomenon of a near-core composition of \mathcal{C} , and we prove the equivalence by mathematically translating the constraints in one definition into the other, and vice versa.

1. \Rightarrow :

Suppose \mathcal{C} is subset atomic with respect to Δ via composition matrix \mathbf{D} . Let $n := |\Delta|$. Consider

$$\begin{aligned} \mathcal{E} : \Lambda &\rightarrow \mathbb{N}^n \setminus \{0^n\} : \\ S &\mapsto \mathbf{d}_S, \forall S \in \Lambda \setminus A \\ A &\mapsto \mathbf{d}_A = \mathbf{e}_A, \forall A \in \Lambda \end{aligned} \tag{6.4}$$

We shall argue that \mathcal{E} has the desired property. By construction, $\mathbf{e}_1 = \mathbf{e}_{A_1}, \dots, \mathbf{e}_n = \mathbf{e}_{A_n} \in \mathcal{E}(\Lambda)$, and \mathcal{E} restricted to the preimage of $\{\mathbf{e}_i\}_{i=1}^n$ is one-to-one; subset atomicity inherits Condition (2) of Definition (3.1) (Primitive Atomic), which implies that

$$(\forall i \in [1, n]) (\exists S \in \Lambda \setminus \bigcup_{i=1}^n \mathcal{E}^{-1}(\mathbf{e}_i) = \Lambda \setminus \Delta) (\mathcal{E}(\mathbf{S}))_i = \mathbf{d}_S(A_i) > 0$$

It remains to show that \mathcal{C} is \mathcal{E} -conservative, which, by [21], is equivalent to $\ker(\tilde{\mathcal{E}}) \supseteq \Theta$ where Θ is the span of the reaction vectors and $\tilde{\mathcal{E}}$ is the linear extension $\tilde{\mathcal{E}}$ of \mathcal{E} :

$$\begin{aligned} \tilde{\mathcal{E}} : \mathbb{R}^\Lambda &\rightarrow \mathbb{R}^n : \\ \mathbf{c} &\mapsto \left(\sum_{S \in \Lambda} \mathbf{c}(S) \cdot (\mathcal{E}(S))_1, \dots, \sum_{S \in \Lambda} \mathbf{c}(S) \cdot (\mathcal{E}(S))_n \right)^T, \mathbf{c} \in \mathbb{R}^\Lambda \end{aligned} \tag{6.5}$$

Observe that $\tilde{\mathcal{E}}(\cdot)$ operates on $\mathbf{c} \in \mathbb{R}^\Lambda$ as left multiplication by \mathbf{D}^T , the transpose of the decomposition matrix. By definition of subset atomicity, any reaction preserves the count/concentration of each atom, so $\mathbf{D}^T \cdot (\mathbf{p} - \mathbf{r}) = 0^n$ ($\forall (\mathbf{r}, \mathbf{p}) \in R$). Hence $a_1(\mathbf{p}_1 - \mathbf{r}_1) + \dots + a_k(\mathbf{p}_k - \mathbf{r}_k) \in \ker \tilde{\mathcal{E}}$ ($k = |R|$) for any $a_1(\mathbf{p}_1 - \mathbf{r}_1) + \dots + a_k(\mathbf{p}_k - \mathbf{r}_k) \in \Theta$, as desired.

2. \Leftarrow :

Suppose we have a function $\mathcal{E} : \Lambda \rightarrow \mathbb{N}^n \setminus \{0\}^n$ satisfying the three described properties. Then define

$$\Delta = \{\mathcal{E}^{-1}(\mathbf{e}_i)\}_{i=1}^n \tag{6.6}$$

We argue that \mathcal{C} is subset atomic with respect to Δ . Indeed, define the composition vectors:

$$\begin{aligned} \mathbf{d} : \Lambda &\rightarrow \mathbb{N}^n \setminus \{0\}^n, \\ S &\mapsto \mathcal{E}(S), \forall S \in \Lambda \end{aligned} \quad (6.7)$$

Since \mathbf{d} coincides with \mathcal{E} everywhere and since \mathcal{C} is \mathcal{E} -conservative, for each $S \in \Lambda$, $(\mathbf{r}, \mathbf{p}) \in R$ and $A_i \in \Delta$, we have

$$\begin{aligned} \sum_{S \in \Lambda} (\mathbf{p}(S) - \mathbf{r}(S)) \cdot \mathbf{d}_S(A_i) &= \sum_{S \in \Lambda} (\mathbf{p}(S) - \mathbf{r}(S)) \cdot (\mathcal{E}(S))_i \\ &= (\tilde{\mathcal{E}}(\mathbf{p} - \mathbf{r}))_i \\ &\stackrel{\text{underbrace}}{=} 0, \end{aligned} \quad (6.8)$$

$\mathcal{E} \text{ conservative} \Rightarrow \mathbf{p} - \mathbf{r} \in \ker \tilde{\mathcal{E}}$

which gives the atom-preservation condition of subset atomic.

By construction, $(\forall i) A_i \in \Delta \subseteq \Lambda$, $\mathbf{d}_{A_i} = \mathbf{e}_i$; To see that $\|\mathbf{d}(S)\|_1 \geq 2$ for all $S \in \Lambda \setminus \Delta$, recall that $\mathcal{E} \upharpoonright \mathcal{E}^{-1}(\{\mathbf{e}_i\}_{i=1}^n)$ is one to one, which means $\forall S \in \Lambda \setminus \Delta$, $\mathcal{E}(S) \notin \{\mathbf{e}_i\}_{i=1}^n$. Given that $\text{range}(\mathcal{E}) = \mathbb{N}^n \setminus \{0\}^n$, this means $\mathbf{d}(S)$ is some non-trivial linear combination of \mathbf{e}_i 's, which gives $\|\mathbf{d}(S)\|_1 \geq 2$ as desired;

Lastly, $(\forall i \in [1, n]) (\exists S \in \Lambda \setminus \bigcup_{i=1}^n \mathcal{E}^{-1}(\mathbf{e}_i)) (\mathcal{E}(S))_i > 0$ translates to $(\forall A \in \Delta) (\exists S \in \Lambda \setminus \Delta) A \in [\mathbf{d}_S]$ by definition of \mathbf{d} . \square

In order to further describe the relationship between atomic networks and core compositions, we make the following definitions first.

Definition 6.25 (Associated Composition). *Given a subset atomic chemical reaction network $\mathcal{C} = (\Lambda, R)$ with respect to Δ via composition matrix \mathbf{D} , the associated composition of \mathbf{D} is the function \mathcal{E} constructed in (6.4). \mathcal{E} 's unique linear extension $\tilde{\mathcal{E}}$, constructed in (6.5), is defined as the extended associated composition of \mathbf{D} .*

Remark 6.26. Note that $\tilde{\mathcal{E}}$, the linear extension of \mathcal{E} , is defined in (6.5) independent of atomic decompositions.

Next we prove that reachably atomic networks admits a core composition. But first, we exhibit some auxiliary definitions.

Definition 6.27. *Given a subset atomic chemical reaction network (Λ, R) with respect to Δ via \mathbf{D} , a single-molecule composition vector $\mathbf{d}'_{S_i} \in \mathbb{N}^\Lambda \setminus \{0^n\}$ is defined as $\mathbf{d}_{S_i} - \mathbf{e}_{S_i}$. The set of single-molecule composition vector is denoted as $U := \{\mathbf{d}'_{S_i}\}_{S_i \in \Lambda \setminus \Delta}$.*

Recall that \mathbf{d}_{S_i} is the composition vector of S_i whose first $|\Lambda| - |\Delta \cap \Lambda|$ coordinates are 0 and last $|\Delta|$ coordinates correspond to the count of each atom in the molecule. for subset atomic networks, $\mathbf{d}_{S_i} \in \mathbb{N}^\Lambda \setminus \{0^n\}$, so \mathbf{d}'_{S_i} is well-defined by replacing 0 with -1 on the i -th molecular position.

We explore the relationship between atomicity and core compositions by inspecting into the relationship between $\ker(\tilde{\mathcal{E}})$ and Θ , the span of reaction vectors. This is in turn done by inspecting the relation between $\ker \tilde{\mathcal{E}}$ and the space spanned by U .

Next, let Υ denote the vector space spanned by U , as a subspace of \mathbb{R}^n . We now give the first approach to the implication “reachably atomicity \Rightarrow Core-Composition Admission”. To do this, we will prove that for subset atomic networks with $\tilde{\mathcal{E}}$ defined as previously defined in (6.5,6.4), $\ker \tilde{\mathcal{E}} = \Upsilon$; for reachably atomic networks (which are by definition also subset atomic), $\Upsilon \subseteq \Theta$. The two relations combined would give $\ker \tilde{\mathcal{E}} \subseteq \Theta$, which is exactly the missing bit from near-core to core compositions. The proofs will be carried out from Lemma 6.28 through Lemma 6.31.

Lemma 6.28. *Vectors in U are linearly independent. Since they also span Υ , this means U is a basis for Υ .*

Proof. Observe that the -1 on the i -th position ($\forall 1 \leq i \leq |\Lambda| - |\Delta|$) cannot be obtained by linear combination of other vectors in U , the i -th position of which are all 0’s.

This also shows that $\dim(\Upsilon) = |U| = |\Lambda - \Delta| = |\Lambda| - |\Delta|$. □

Lemma 6.29 (Kernel-Span Equivalence). *For subset atomic networks with $\tilde{\mathcal{E}}$ defined as in (6.5,6.4), $\ker(\tilde{\mathcal{E}}) = \Upsilon$.*

Proof. By verification of definitions.

1. $\Upsilon \subseteq \ker(\tilde{\mathcal{E}})$:

Take $\mathbf{u} = a_1 \mathbf{d}'_{S_1} + \cdots + a_{|\Lambda|-|\Delta|} \mathbf{d}'_{S_{|\Lambda|-|\Delta|}} \in \Upsilon$. Then $\forall i \in [1, n]$, we have

$$\begin{aligned} (\tilde{\mathcal{E}}(\mathbf{u}))_i &= \sum_{S_j \in \Lambda} \mathbf{u}(S_j) \cdot (\mathcal{E}(S_j))_i \\ &\stackrel{(6.4)}{=} \sum_{S_j \in \Lambda \setminus \Delta} \mathbf{u}(S_j) \cdot \mathbf{d}_{S_j}(A_i) + \sum_{A_k \in \Delta} \mathbf{u}(A_k) \cdot \mathbf{d}_{A_k}(A_i) \\ &= \sum_{S_j \in \Lambda \setminus \Delta} (-a_j) \cdot \mathbf{d}_{S_j}(A_i) + \mathbf{u}(A_i) \cdot 1 \\ &= \sum_{S_j \in \Lambda \setminus \Delta} (-a_j) \cdot \mathbf{d}_{S_j}(A_i) + \sum_{S_j \in \Lambda \setminus \Delta} a_j \cdot \mathbf{d}_{S_j}(A_i) \\ &= 0 \end{aligned}$$

2. $\Upsilon \supseteq \ker(\tilde{\mathcal{E}})$:

Take $\mathbf{v} \in \ker(\tilde{\mathcal{E}})$, and let $a_1, \dots, a_{|\Lambda|-|\Delta|}$ denote the first $|\Lambda| - |\Delta|$ coordinates of \mathbf{v} . We claim that

$$\mathbf{v} = - \sum_{i=1}^{|\Lambda|-|\Delta|} a_i \mathbf{d}'_{S_i} \tag{6.9}$$

Indeed, for each $S_j \in \Lambda \setminus \Delta$,

$$\mathbf{v}(S_j) = -a_j \cdot (-1) = \sum_{i \neq j} (-a_i) \cdot \underbrace{\mathbf{d}'_{S_i}(S_j)}_0 + (-a_j) \cdot \underbrace{\mathbf{d}'_{S_j}(S_j)}_{-1} = - \sum_{i=1}^{|\Lambda|-|\Delta|} a_i \mathbf{d}'_{S_i}(S_j)$$

It remains to verify that (6.9) holds the last $|\Delta| = n$ positions. Let b_1, \dots, b_n denote the last n positions of \mathbf{v} , c_1, \dots, c_n denote the last n positions of $-\sum_{i=1}^{|\Lambda|-|\Delta|} a_i \mathbf{d}'_{S_i}$. Because $\mathbf{v} \in \ker(\tilde{\mathcal{E}})$, we know that for each $i \in [1, n]$,

$$\begin{aligned}
(\tilde{\mathcal{E}}(\mathbf{v}))_i &= \sum_{S_j \in \Lambda} \mathbf{v}(S_j) \cdot (\mathcal{E}(S_j))_i \\
&\stackrel{(6.4)}{=} \underbrace{\sum_{S_j \in \Lambda \setminus \Delta} \mathbf{v}(S_j) \cdot \mathbf{d}_{S_j}(A_i)} + \sum_{A_k \in \Delta} \mathbf{v}(A_k) \cdot \mathbf{d}_{A_k}(A_i) \\
&= \sum_{S_j \in \Lambda \setminus \Delta} (a_j) \cdot \mathbf{d}_{S_j}(A_i) + \mathbf{v}(A_i) \cdot 1 \\
&= \sum_{S_j \in \Lambda \setminus \Delta} (a_j) \cdot \mathbf{d}_{S_j}(A_i) + b_i \cdot 1 \\
&= 0,
\end{aligned}$$

which gives that

$$\begin{aligned}
b_i &= - \sum_{j=1}^{|\Lambda|-|\Delta|} a_j \cdot \mathbf{d}_{S_j}(A_i) \\
&= - \sum_{j=1}^{|\Lambda|-|\Delta|} a_j \cdot \mathbf{d}'_{S_j}(A_i) \\
&= c_i
\end{aligned}$$

as desired, completing the proof. \square

Remark 6.30. The lemmas above directly imply that $\dim(\ker(\tilde{\mathcal{E}})) = \dim(\Upsilon) = |\Lambda| - |\Delta|$.

The Remark above resonates Gilles' Theorem 3.3, which gives a equation for general cases where $\mathcal{E} \upharpoonright \mathcal{E}^{-1}(\{\mathbf{e}_i\}_{i=1}^n)$ is not necessarily one-to-one. Since the general cases are not directly related to our discussion on network atomicity,¹¹ we refrain from further discussion thereon.

Lemma 6.31. *If a subset atomic chemical reaction network \mathcal{C} is reachably atomic, then \mathcal{C} admits a core composition. That is, \mathcal{C} is constructive.*

Proof. By Lemma (6.23), since \mathcal{C} is subset atomic, it admits a near-core composition. By Lemma (6.29), $\ker(\tilde{\mathcal{E}}) = \Upsilon$, so it suffices to prove $\Upsilon \subseteq \Theta$ when \mathcal{C} is reachably atomic.

To prove $\Upsilon \subseteq \Theta$, recall that by Lemma (6.28), U is a basis for Υ . Hence we only need to argue that each basis vector $\mathbf{d}'_{S_i} \in U$ is a linear combination of reaction vectors. Indeed, $\forall S_i \in \Lambda \setminus \Delta$, $\mathbf{e}_{S_i} \Rightarrow^* \mathbf{d}_{S_i}$, so there exists $\mathbf{p}_1 - \mathbf{r}_1, \dots, \mathbf{p}_k - \mathbf{r}_k \in R$ for some k with $\mathbf{r}_1 = \mathbf{e}_{S_i}$ s.t. $\mathbf{e}_{S_i} + \sum_k (\mathbf{p}_k - \mathbf{r}_k) = \mathbf{d}_{S_i}$. But this means $\mathbf{d}'_{S_i} = \mathbf{d}_{S_i} - \mathbf{e}_{S_i} = \sum_k (\mathbf{p}_k - \mathbf{r}_k) \in \Theta$.

Having proved $\ker \tilde{\mathcal{E}} = \Gamma = \Theta$, we conclude that \mathcal{C} admits a core composition, as desired. \square

¹¹For the fact that any species cannot have two atomic compositions in a single composition matrix. Note that molecular species may admit different atomic compositions, but they belong to different composition matrices (that is, atomic compositions of other species have to change accordingly).

Remark 6.32. We exhibit an alternative approach for Lemma 6.31 by directly applying Theorem 4.2 of [21].

Theorem 4.2 of [21] states that if $\mathcal{E} : \Lambda \rightarrow \mathbb{N}^n \setminus \{0^n\}$ satisfies the following, then \mathcal{E} is a core composition:

1. \mathcal{E} is near-core;
2. all \mathcal{E} -elementary species of the same \mathcal{E} -isomeric class are stoichiometrically-isomeric ($X, Y \in \mathcal{X}_i \Rightarrow Y - X \in \Theta = \text{span}(R)$); here, \mathcal{E} -isomeric classes are defined as species having the same $\mathcal{E}(\cdot)$ -value. Restricted to \mathcal{E} -elementary species, the \mathcal{E} -isomeric classes are $\mathcal{X}_i := \{S \in \Lambda \mid \mathcal{E}(S) = \mathbf{e}_i\}$ ($i \in [1, n]$).
3. For every \mathcal{E} -composite Y with $\mathcal{E}(Y) = \alpha \in \mathbb{Z}_{\geq 0}^n \setminus \{0^n\}$, there exist one “representative” elementary species from each elementary isomeric class such that Y and the α -linear combination of these elementary species are stoichiometrically compatible. That is, $\exists W_1 \in \mathcal{X}_1, \dots, W_n \in \mathcal{X}_n$, s.t. $Y - \sum_{i=1}^n \alpha_i W_i \in \Theta$.

Note that for a reachably atomic network \mathcal{C} and its associated composition \mathcal{E} , all elementary \mathcal{E} -isomeric classes are singleton, so $(\forall i)(\forall A_i \in \mathcal{X}_i)(A_i - A_i = 0^n \in \Theta)$ ¹²; Also, for each $Y \in \Lambda \setminus \Delta$, $Y - \sum_i \underbrace{\mathbf{d}_Y(A_i)}_{\alpha_i} A_i = (-1) \cdot \mathbf{d}'_Y \in \Theta$, as $\mathbf{d}'_Y = \mathbf{d}_Y - \mathbf{e}_Y = \sum_j (\mathbf{p}_j - \mathbf{r}_j) \in \Theta$ is guaranteed by the reachably atomicity. The reachably atomicity implies subset atomicity, which again implies \mathcal{E} is near-core. Therefore \mathcal{C} adopts a core composition, as desired.

Although in the alternative proof exhibited above we talked about the concept of “span” as linear combinations with real coefficients, in fact our model of reachably atomic networks is not necessarily equipped with the property that each reaction is reversible. This fact does not break the proof, but it justifies our decision to keep both approaches. Further, the above observation indicates the possibility that networks with the reversible property may itself guarantee some more interesting structures. In fact, let us study the relationship between reversibly-reachably atomic networks and explicitly-reversibly constructive networks in the following section.

6.4 Reversibly-Reachably Atomicity and Explicitly-Reversibly Constructiveness

In this section we show an equivalence between our definition of *reversibly-reachably atomic* and [21]’s definition of *explicitly-reversibly constructive* with an additional restriction.

We first note that the following subclass of reachably atomic networks has such a property that each molecular species can be explicitly constructed from its atomic makeup via reactions.

Definition 6.33 (Reversibly-Reachably Atomic). *A chemical reaction network $\mathcal{C} = (\Lambda, R)$ is reversibly-reachably saturated atomic if:*

1. *It is reachably atomic with respect to some $\Delta \subseteq \Lambda$ via the composition matrix \mathbf{D} ;*
2. $\forall S_j \in \Lambda \setminus \Delta, \mathbf{d}_{S_j} \Rightarrow^* \{1S_j\}$.

¹²Note that this doesn’t contradict the assumption that $(\mathbf{r}, \mathbf{r}) \notin R$, since $0^n = 1 \cdot (\mathbf{p} - \mathbf{r}) + (-1) \cdot (\mathbf{p} - \mathbf{r})$ can be obtained by linear combination of reactions $(\mathbf{r}, \mathbf{p}) \in R$ where $\mathbf{r} \neq \mathbf{p}$.

By definition, reversibly-reachably atomic networks \subsetneq reachably atomic networks \subsetneq subset atomic networks \subsetneq primitive atomic networks; or equivalently, for a given \mathcal{C} , primitive atomicity \Rightarrow subset atomicity \Rightarrow reachably atomicity \Rightarrow reversibly-reachably atomicity, while the reversed arrows do not necessarily hold.

Correspondingly, we define the following language of encodings of networks with the reversibly-reachably atomic property:

Definition 6.34.

$$\text{REVERSIBLY-REACHABLY-ATOMIC} = \{ \langle \Lambda, R \rangle \mid (\exists \Delta \subseteq \Lambda) (\langle \Lambda, R \rangle \text{ is reversibly-reachably-atomic with respect to } \Delta) \}$$

Corollary 6.35. $\text{REVERSIBLY-REACHABLY ATOMIC} \in P$.

Proof. The proof is almost mirroring the previous proof of 6.3. It takes polynomial time to decide if an instance $\langle \Lambda, R \rangle \in \text{REACHABLY-ATOMIC}$, as shown above. we extend the REACHABLY-ATOMIC decider to decide whether \mathcal{C} , if confirmed to be reachably atomic, further satisfies $\mathbf{d}_S \Rightarrow^* \{1S\}$ for each $S \in \Lambda \setminus \Delta = M$:

Construct $M'' = M$. While M'' is not empty, iterate and try to find an $S'' \in M''$ that satisfies this condition: $\exists(\mathbf{r}, \mathbf{p}) \in R$ s.t. $\mathbf{p} = \{1S''\}$ and $[\mathbf{r}] \subseteq (M \setminus M') \cup \Delta$. Note that all elements S'' in $M \setminus M''$ satisfies $\{1S''\} \Leftarrow^* \mathbf{d}_{S''}$, hence if $\{S''\}$ is the product of a reaction whose reactants consist solely of elements in $(M \setminus M'') \cup \Delta$, then S'' itself satisfies $\{1S''\} \Leftarrow^* \mathbf{d}_{S''}$ as well. Keep track of $\mathbf{d}_{S''}$ and exclude such S'' from M'' .

If in some iteration we cannot find such $S'' \in M''$, then REJECT; else, the iteration will finally halt excluding all such S'' 's and making M'' empty, in which case we ACCEPT.

Proof of correctness works analogously as the proof above, with the \Rightarrow^* reversed to \Leftarrow^* and disassociation reactions changed into association reactions for consideration.

We omit the pseudocode for the algorithm described above, as it highly resembles Lines 25-35 of Algorithm 2 (with \mathbf{r} and \mathbf{p} reversed) and can be reconstructed from the verbal description above. Lastly, the complexity is dominated by the reachably atomic deciding process. □

Theorem 6.36. *For a chemical reaction network \mathcal{C} , the following are equivalent:*

1. \mathcal{C} is reversibly-reachably atomic;
2. \mathcal{C} is explicitly-reversibly constructive, with $p_i = 1$ ($\forall i \in [1, n]$) where $p_i := |\mathcal{X}_i| = |\mathcal{E}^{-1}(e_i)|$.

Note that condition (2) translates to “no \mathcal{E} -elementary species is \mathcal{E} -isomeric”.

Proof. Proof is done by similar techniques in 6.2: that is, assuming otherwise, then there will be an infinite descending chain of species ordered by number of atoms in their respective composition. This contradicts the fact that the set of species is finite.

1. $1 \Rightarrow 2$:

By Definition 6.33, \mathcal{C} is reachably atomic with respect to some $\Delta \subseteq \Lambda$. Let $|\Delta| = n$, then \mathcal{C} is constructive by lemma (6.31). By Lemma (6.23), since a reachably atomic network with $|\Delta| = n$ is subset atomic, $p_i = 1$ for each $i \in [1, n]$.

the weakly reversibility of $\{1S_i\} \Rightarrow^* \mathbf{d}_{S_i}$ ensures that each composite species is both explicitly constructible and explicitly destructible, because it is ensured that each molecular species directly or indirectly participates in at least one disassociation (resp. binding) reaction as reactant (resp. product);

We also claim that each elementary species is both explicitly constructive and explicitly destructive:

suppose for the sake of contradiction $A_i \in \Delta$ is not explicitly destructive. Then for any reactions where A_i participates as product, the reactant has to contain at least 2 species. In particular, since $A_i \in [\mathbf{d}_{S_i}]$ for some $S_i \in \Lambda \setminus \Delta$, if we consider the sequence $\{1S_i\} \Rightarrow^* \mathbf{d}_{S_i}$, there must be a reaction in this sequence written as



with the multisets $\{a_j Y_j\}_{j=1}^k \neq \{b_j Q_j\}_{j=1}^s$ and $A_i \in \{Q_j\}_{j=1}^s$, s.t.

$$\exists S_j \in (\Lambda \setminus \Delta) \cap \{Y_j\}_{j=1}^k \text{ with } A_i \in [\mathbf{d}_{S_j}]. \quad (6.11)$$

This is because A_i cannot be directly obtained from $\{1S_i\}$, and hence must be obtained from some intermediate molecular species. Apparently $\|\mathbf{d}_{S_j}\|_1 < \|\mathbf{d}_{S_i}\|_1$ by conservativity of atoms. But then consider the decomposition series $\{1S_i\} \Rightarrow^* \mathbf{d}_{S_j}$ and apply the same argument, we obtain $S'_j \in \Lambda \setminus \Delta$ s.t. $\|\mathbf{d}_{S'_j}\|_1 < \|\mathbf{d}_{S_j}\|_1 < \|\mathbf{d}_{S_i}\|_1$. By the finiteness of $\Lambda \setminus \Delta$, this repeated process terminates with a “smallest” molecule containing A_i ; that is, there exists some $S_m \in \Lambda \setminus \Delta$ where $A_i \in [\mathbf{d}_{S_m}]$ and $\forall n \neq m$, if $A_i \in [\mathbf{d}_{S_n}]$, then $\|\mathbf{d}_{S_m}\|_1 < \|\mathbf{d}_{S_n}\|_1$.

But now we once again apply the argument (involving equations (6.10,6.11) above, getting some $S'_m \in \Lambda \setminus \Delta$ with $A_i \in [\mathbf{d}_{S'_m}]$ s.t. $\|\mathbf{d}_{S'_m}\|_1 < \|\mathbf{d}_{S_m}\|_1$, a contradiction.

Symmetric argument with “products” and “reactants” swapped, given the reachability of $\mathbf{d}_{S_i} \Rightarrow^* \{1S_i\}$ for each $S_i \in \Lambda \setminus \Delta$, proves that all $A_i \in \Delta$ also have to be explicitly constructive.

2. $2 \Rightarrow 1$:

Because \mathcal{C} admits a core composition, in particular it admits a near-core composition. The one-to-one condition allows us to define the set of atoms Δ and composition \mathbf{d} as (6.6) and (6.7) in the proof of Lemma (6.23).

We first argue that condition (2) of Definition 3.1 holds. Suppose not, then $\exists A_i \in \Delta$ s.t. $\forall S_j \in \Lambda \setminus \Delta$, $A_i \notin [\mathbf{d}_{S_j}]$. But then consider the binding reaction $Q \rightarrow Y$ where $|Y| = 1$ and $A_i \in [Q]$; such a reaction has to exist because of the explicit constructivity of A_i . Y has to be a molecule containing A_i in $[\mathbf{d}_Y]$, a contradiction.

By definition 6.7, the above implies that $(\forall i \in [1, n]) (\exists S \in \Lambda \setminus \Delta = \Lambda \setminus \bigcup_{i=1}^n \mathcal{E}^{-1}(\mathbf{e}_i)) \mathbf{d}_S(A_i) = (\mathcal{E}(S))_i > 0$. Together with the condition $p_i = 1$ ($\forall i \in [1, n]$) and the near-core property, we use the equivalence in Lemma (6.23) to conclude that \mathcal{C} is subset atomic.

Next, we argue that $(2) \Rightarrow \forall S_j \in \Lambda \setminus \Delta$, $\mathbf{d}_{S_j} \Rightarrow^* \{1S_j\}$ and $\{1S_j\} \Rightarrow^* \mathbf{d}_{S_j}$. Note that the latter reachability would also imply the reachable-atomicity of \mathcal{C} , given that \mathcal{C} is subset atomic

and that atomic decomposition is naturally unique for reachably atomic networks, by Lemma 3.9.

Consider an arbitrary $S_i \in \Lambda \setminus \Delta$. For the sake of contradiction, assume $\{1S_i\} \not\Rightarrow^* \mathbf{d}_{S_i}$. Since S_i is explicitly destructible, some reaction sequence starting with S_i has to eventually split into a complex Q_i with $|Q_i| \geq 2$. On the other hand, since \mathcal{C} is subset atomic, any reaction sequence starting with $\{1S_i\}$ either reaches \mathbf{d}_{S_i} or reaches some configuration \mathbf{c} with $[\mathbf{c}] \cap (\Lambda \setminus \Delta) \neq \emptyset$. Since we assumed that $1S_i \not\Rightarrow^* \mathbf{d}_{S_i}$, it must be the second case, and in particular, there exists $S_j \in [Q_i] \cap (\Lambda \setminus \Delta)$. Apparently $\|\mathbf{d}_{S_j}\|_1 < \|\mathbf{d}_{S_i}\|_1$ by conservativity of number of atoms.

We apply the same argument to S_j and obtain S'_j s.t. $\|\mathbf{d}_{S'_j}\|_1 < \|\mathbf{d}_{S_j}\|_1 < \|\mathbf{d}_{S_i}\|_1$. Repeat this process, and by finiteness of $\Lambda \setminus \Delta$, we'll find some $S_m \in \Lambda \setminus \Delta$ satisfying $\forall m' \neq m, \|\mathbf{d}_{S_{m'}}\|_1 > \|\mathbf{d}_{S_m}\|_1$. Application of the same argument to S_m yields the contradiction as to the size of composition vector.

The same “infinite descending chain” argument applies to the other direction, with all the arrows reversed and the explicit constructibility property applied. This proves the reachable-atomicity as well as condition (2) in Definition 6.33. \square

Corollary 6.37. *The problem “Given a chemical reaction network, is it explicitly reversibly constructible with no isomeric elementary species” as well as the problem “Given a chemical reaction network \mathcal{C} , is \mathcal{C} reversibly-reachably atomic” are both polynomial time decidable.*

Proof. Immediate from Corollary 6.35 and Theorem 6.36. \square

7 Open Problems

Since we say that the property of mass conservation captures that mass can be neither created nor destroyed, it is natural to conjecture that if a network is *not* mass conserving, then it can either create mass (i.e., there exist configurations $\mathbf{c}_1 < \mathbf{c}_2$ such that $\mathbf{c}_1 \Rightarrow^* \mathbf{c}_2$) or destroy mass (i.e., there exist configurations $\mathbf{c}_1 > \mathbf{c}_2$ such that $\mathbf{c}_1 \Rightarrow^* \mathbf{c}_2$). However, this is not always the case. Consider the reactions:



We first argue that for arbitrary configurations $\mathbf{c}_1, \mathbf{c}_2$ such that $\mathbf{c}_1 \Rightarrow_{\mathcal{C}}^* \mathbf{c}_2$, \mathbf{c}_1 and \mathbf{c}_2 are incomparable. Suppose \mathcal{C} reaches \mathbf{c}_2 from \mathbf{c}_1 via reaction sequence $\{(\mathbf{r}_i, \mathbf{p}_i)\}_{i \in I}$ for some finite index set I . Observe that if $\{(\mathbf{r}_i, \mathbf{p}_i)\}_{i \in I}$ contains reaction (7.1), then execution of (7.1) will reduce the count of X and increase the count of B . The configuration immediately after execution of (7.1) is incomparable with the configuration immediately before execution of (7.1), and the only way to drain the B 's (to produce a strictly smaller configuration) is to execute equal times of (7.2), which will irreversibly increase the count of Z . This means $\mathbf{c}_1 \Rightarrow_{\mathcal{C}}^* \mathbf{c}_2$ implies $((7.1)) \in \{(\mathbf{r}_i, \mathbf{p}_i)\}_{i \in I} \Rightarrow \mathbf{c}_1$ incomparable with \mathbf{c}_2 .

But if $(7.1) \notin \{(\mathbf{r}_i, \mathbf{p}_i)\}_{i \in I}$, then $\{(\mathbf{r}_i, \mathbf{p}_i)\}_{i \in I}$ contains only (7.2) and (7.3), which irreversibly decreases count of Y and increases count of Z . Altogether, $(\mathbf{c}_1 \Rightarrow_{\mathcal{C}}^* \mathbf{c}_2) \Rightarrow \mathbf{c}_1$ is incomparable with \mathbf{c}_2 .

We should also observe that \mathcal{C} is non-mass conserving: for any candidate mass distribution \mathbf{m} , (7.3) requires $\mathbf{m}(Y) > \mathbf{m}(Z)$, which together with (7.2) implies $\mathbf{m}(A) > \mathbf{m}(B)$, yet (7.1) requires $\mathbf{m}(B) > \mathbf{m}(A)$. Hence \mathcal{C} is a non-mass-conserving example where $\forall \mathbf{c}_1, \mathbf{c}_2$ s.t. $\mathbf{c}_1 \Rightarrow_{\mathcal{C}}^* \mathbf{c}_2$, \mathbf{c}_1 and \mathbf{c}_2 are incomparable. Note that \mathcal{C} is not reversible. However, it is open whether a *reversible* network conserves mass if and only if it neither creates nor destroys mass in the reachability sense defined above. Since reversible networks have a symmetric reachability relation, this is equivalent to the following conjecture:

Conjecture 7.1. *Let a chemical reaction network \mathcal{C} be reversible. \mathcal{C} is non-mass conserving if and only if there are $\mathbf{c}_1 < \mathbf{c}_2 \in \mathbb{N}^A$ such that $\mathbf{c}_1 \Rightarrow_{\mathcal{C}}^* \mathbf{c}_2$.*

Next, recall the conjecture 5.9 at the end of Section 3:

Conjecture 7.2. *SUBSET-ATOMIC is NP-complete.*

One may note that there are two sources of indeterminacy in the problem SUBSET-ATOMIC: the choice of Δ and the choice of \mathbf{D} . For example, the network constructed in the proof of NP-hardness of SUBSET-FIXED-ATOMIC would remain subset atomic if we define $\Delta = \{T, F\}$, and let $\mathbf{d}_P = \mathbf{d}_Q = \{kT, sF\}$ for any positive k, s .

There is a formal sense in which chemical reaction networks have been shown to be able to compute functions $f : \mathbb{N}^k \rightarrow \mathbb{N}$ [8] and predicates $\mathbb{N}^k \rightarrow \{0, 1\}$ [4]. A function/predicate can be computed “deterministically” (i.e., regardless of the order in which reactions occur) if and only if it is semilinear (see [20] for a definition).

Problem 7.3. *What is the computational power of atomic chemical reaction networks? What semilinear functions/predicates can atomic chemical reaction networks compute deterministically, and how efficiently? What general functions/predicates can atomic chemical reaction networks compute with high probability, and how efficiently?*

Acknowledgements. The authors are thankful to Manoj Gopalkrishnan, Gilles Gnacadja, Javier Esparza, Sergei Chubanov, and Matthew Cook for their insights and useful discussion.

References

- [1] Leonard Adleman, Manoj Gopalkrishnan, Ming-Deh Huang, Pablo Moisset, and Dustin Reishus. *On the Mathematics of the Law of Mass Action*, pages 3–46. Springer Netherlands, Dordrecht, 2014. URL: http://dx.doi.org/10.1007/978-94-017-9041-3_1, doi: 10.1007/978-94-017-9041-3_1.
- [2] Dan Alistarh, James Aspnes, David Eisenstat, Rati Gelashvili, and Ronald Rivest. Time-space trade-offs in molecular computation. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2560–2579, 2017.
- [3] David Angeli, Patrick De Leenheer, and Eduardo D. Sontag. A Petri net approach to the study of persistence in chemical reaction networks. *Mathematical Biosciences*, 210(2):598–618, 2007.
- [4] Dana Angluin, James Aspnes, Zoë Diamadi, Michael Fischer, and René Peralta. Computation in networks of passively mobile finite-state sensors. *Distributed Computing*, 18:235–253, 2006. Preliminary version appeared in PODC 2004. URL: <http://dx.doi.org/10.1007/s00446-005-0138-3>.
- [5] Robert Brijder, David Doty, and David Soloveichik. Robustness of expressivity in chemical reaction networks. In *DNA 2016: Proceedings of the 22th International Meeting on DNA Computing and Molecular Programming*, 2016.
- [6] Luca Cardelli and Attila Csikász-Nagy. The cell cycle switch computes approximate majority. *Scientific Reports*, 2, 2012.
- [7] Ho-Lin Chen, Rachel Cummings, David Doty, and David Soloveichik. Speed faults in computation by chemical reaction networks. *Distributed Computing*, 2015. to appear. Special issue of invited papers from DISC 2014.
- [8] Ho-Lin Chen, David Doty, and David Soloveichik. Deterministic function computation with chemical reaction networks. *Natural Computing*, 13(4):517–534, 2013. Special issue of invited papers from DNA 2012.
- [9] Ho-Lin Chen, David Doty, and David Soloveichik. Rate-independent computation in continuous chemical reaction networks. In *ITCS 2014: Proceedings of the 5th Conference on Innovations in Theoretical Computer Science*, pages 313–326, 2014.
- [10] Yuan-Jyue Chen, Neil Dalchau, Niranjana Srinivas, Andrew Phillips, Luca Cardelli, David Soloveichik, and Georg Seelig. Programmable chemical controllers made from DNA. *Nature Nanotechnology*, 8(10):755–762, 2013.
- [11] Sergei Chubanov. A polynomial projection algorithm for linear feasibility problems. *Mathematical Programming*, 153(2):687–713, 2015.
- [12] Sergei Chubanov. Private communication. *Email*, 2016.
- [13] Gheorghe Craciun, Alicia Dickenstein, Anne Shiu, and Bernd Sturmfels. Toric dynamical systems. *Journal of Symbolic Computation*, 44(11):1551–1565, 2009.

- [14] Rachel Cummings, David Doty, and David Soloveichik. Probability 1 computation with chemical reaction networks. *Natural Computing*, pages 1–17, 2015. Special issue of invited papers from DNA 2014. URL: <http://dx.doi.org/10.1007/s11047-015-9501-x>, doi:10.1007/s11047-015-9501-x.
- [15] David Doty. Timing in chemical reaction networks. In *SODA 2014: Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 772–784, January 2014.
- [16] David Doty and Monir Hajiaghayi. Leaderless deterministic chemical reaction networks. *Natural Computing*, 14(2):213–223, 2015. Preliminary version appeared in DNA 2013.
- [17] Javier Esparza, Pierre Ganty, Jérôme Leroux, and Rupak Majumdar. Verification of population protocols. *Acta Informatica*, pages 1–25, 2016.
- [18] Michael R Garey and David S Johnson. “strong” np-completeness results: Motivation, examples, and implications. *Journal of the ACM (JACM)*, 25(3):499–508, 1978.
- [19] Michael R. Garey and David S. Johnson. *Computers and Intractability*. W. H. Freeman, New York, 1979.
- [20] Seymour Ginsburg and Edwin H. Spanier. Semigroups, Presburger formulas, and languages. *Pacific Journal of Mathematics*, 16(2):285–296, 1966. URL: <http://projecteuclid.org/euclid.pjm/1102994974>.
- [21] Gilles Gnacadja. Reachability, persistence, and constructive chemical reaction networks (part II): A formalism for species composition in chemical reaction network theory and application to persistence. *Journal of mathematical chemistry*, 49(10):2137, 2011.
- [22] Manoj Gopalkrishnan. Private communication. *Email*, 2016.
- [23] Cato M. Guldberg and Peter Waage. Studies concerning affinity. *Forhandlinger: Videnskabs-Selskabet i Christinia. Norwegian Academy of Science and Letters*, 35, 1864.
- [24] F. J. M. Horn. The dynamics of open reaction systems. In *SIAM-AMS Proceedings VIII*, page 125–137, 1974.
- [25] Hua Jiang, Sayed Ahmad Salehi, Marc D Riedel, and Keshab K Parhi. Discrete-time signal processing with DNA. *ACS Synthetic Biology*, 2(5):245–254, 2013.
- [26] Robert F Johnson, Qing Dong, and Erik Winfree. Verifying chemical reaction network implementations: A bisimulation approach. In *International Conference on DNA-Based Computers*, pages 114–134. Springer, 2016.
- [27] DE Knuth. Seminumerical algorithms, the art of computer programming, vol. 2 addison-wesley. *Reading, MA*, 1997.
- [28] Y Edmund Lien. A note on transition systems. *Information Sciences*, 10(2):347–362, 1976.
- [29] Martin Mann, Feras Nahar, Norah Schnorr, Rolf Backofen, Peter F Stadler, and Christoph Flamm. Atom mapping with constraint programming. *Algorithms for Molecular Biology*, 9(1):23, 2014.

- [30] Ernst W Mayr and Jeremias Weihmann. A framework for classical Petri net problems: Conservative Petri nets as an application. In *International Conference on Applications and Theory of Petri Nets and Concurrency*, pages 314–333. Springer, 2014.
- [31] Kevin Montagne, Raphael Plasson, Yasuyuki Sakai, Teruo Fujii, and Yannick Rondelez. Programming an *in vitro* DNA oscillator using a molecular networking strategy. *Molecular Systems Biology*, 7(1), 2011.
- [32] Nils E Napp and Ryan P Adams. Message passing inference with chemical reaction networks. In *Advances in neural information processing systems*, pages 2247–2255, 2013.
- [33] Kevin Oishi and Eric Klavins. Biomolecular implementation of linear i/o systems. *IET Systems Biology*, 5(4):252–260, 2011.
- [34] Adrien Padirac, Teruo Fujii, and Yannick Rondelez. Nucleic acids for the rational design of reaction circuits. *Current Opinion in Biotechnology*, 24(4):575–580, 2013.
- [35] Christos H Papadimitriou. On the complexity of integer programming. *Journal of the ACM (JACM)*, 28(4):765–768, 1981.
- [36] Christos H Papadimitriou. *Computational complexity*. John Wiley and Sons Ltd., 2003.
- [37] L. Qian, E. Winfree, and J. Bruck. Neural network computation with dna strand displacement cascades. *Nature*, 475(7356):368–372, 2011.
- [38] Lulu Qian and Erik Winfree. Scaling up digital circuit computation with DNA strand displacement cascades. *Science*, 332(6034):1196, 2011.
- [39] Sayed Ahmad Salehi, Keshab K Parhi, and Marc D Riedel. Chemical reaction networks for computing polynomials. *ACS Synthetic Biology*, 2016.
- [40] Sayed Ahmad Salehi, Marc D Riedel, and Keshab K Parhi. Asynchronous discrete-time signal processing with molecular reactions. In *2014 48th Asilomar Conference on Signals, Systems and Computers*, pages 1767–1772. IEEE, 2014.
- [41] Sayed Ahmad Salehi, Marc D Riedel, and Keshab K Parhi. Markov chain computations using molecular reactions. In *2015 IEEE International Conference on Digital Signal Processing (DSP)*, pages 689–693. IEEE, 2015.
- [42] Walter J Savitch. Relationships between nondeterministic and deterministic tape complexities. *Journal of Computer and System Sciences*, 4(2):177–192, 1970.
- [43] Georg Seelig, David Soloveichik, David Yu Zhang, and Erik Winfree. Enzyme-free nucleic acid logic circuits. *Science*, 314(5805):1585–1588, 2006. URL: <http://www.sciencemag.org/cgi/doi/10.1126/science.1132493>, doi:10.1126/science.1132493.
- [44] Abraham Silberschatz, Peter B Galvin, Greg Gagne, and A Silberschatz. *Operating system concepts*, volume 4. Addison-wesley Reading, 2013.
- [45] David Soloveichik, Matthew Cook, Erik Winfree, and Jehoshua Bruck. Computation with finite stochastic chemical reaction networks. *Natural Computing*, 7(4):615–633, 2008. URL: <http://dx.doi.org/10.1007/s11047-008-9067-y>.

- [46] David Soloveichik, Georg Seelig, and Erik Winfree. DNA as a universal substrate for chemical kinetics. *Proceedings of the National Academy of Sciences*, 107(12):5393, 2010. Preliminary version appeared in DNA 2008.
- [47] Niranjan Srinivas. *Programming chemical kinetics: Engineering dynamic reaction networks with DNA strand displacement*. PhD thesis, California Institute of Technology, 2015.
- [48] Jose-Juan Tapia and James R Faeder. The atomizer: extracting implicit molecular structure from reaction network models. In *Proceedings of the International Conference on Bioinformatics, Computational Biology and Biomedical Informatics*, page 726. ACM, 2013.
- [49] Chris Thachuk and Anne Condon. Space and energy efficient computation with DNA strand displacement systems. In *DNA 2012: Proceedings of the 18th International Meeting on DNA Computing and Molecular Programming*, pages 135–149, 2012.
- [50] B. Yurke, A.J. Turberfield, A.P. Mills Jr, F.C. Simmel, and J.L. Neumann. A DNA-fuelled molecular machine made of DNA. *Nature*, 406(6796):605–608, 2000.