A scalable method for finding irreducible state-spaces for stochastic models of biochemical reaction networks

Ankit Gupta and Mustafa Khammash

Department of Biosystems Science and Engineering

ETH Zurich

Mattenstrasse 26

4058 Basel, Switzerland.

May 28, 2019

Abstract

In stochastic models of biochemical reaction networks, the dynamics is usually represented by a Markov process which describes the evolution of the copy-numbers or molecular counts of the constituent species. It is often of biological interest to determine if this Markov process has a unique stationary distribution. This uniqueness will hold if the state-space is *irreducible* in the sense that all the states are reachable from each other in a finite time, with a positive probability. Finding such irreducible state-spaces is quite challenging, because the Markovian dynamics can usually access infinitely many states and the presence of conservation relations among species can constrain the dynamics in complicated ways. The aim of this paper is to develop a computational framework for finding irreducible state-spaces for reaction networks that typically arise in Systems and Synthetic Biology. Our results can help in assessing the long-term behavior of a network and also in explicitly obtaining the stationary distributions in certain cases. Our framework only relies on elementary linear algebra and linear programming, which makes it highly scalable and efficient, even for very large networks. We illustrate the wide applicability of our framework through several examples.

1 Introduction

Many biological processes are described as reaction networks, where certain species interact with each other through a finite number of reaction channels. Such reaction networks can be found in Epidemiology [19], Pharmacology [5], Ecology [4] and most prominently, in Systems and Synthetic Biology [1, 28, 33, 12]. Traditionally, reaction network models have been mathematically studied by expressing the dynamics as a set of ordinary differential equations (ODEs). However it is now well-known that these deterministic formulations become highly inaccurate when the copy-numbers of the reacting species are *small*. This is because the random timing of reactions introduces *noise* into the dynamics, which can significantly change the behavior of the system being modeled [15, 24]. Such situations arise commonly in Systems and Synthetic Biology, since intracellular networks often involve species with low copy-numbers like gene-transcripts, signaling proteins, messenger RNAs, transcription factors etc. [10, 24]. The biochemical noise generated by the intermittency of reactions can be taken into account using stochastic formulations of the dynamics of reaction networks. A common approach is to represent the dynamics as a continuous-time Markov process whose states denote the population sizes or copy-numbers of the constituent species. In recent years, these stochastic models have been extensively used for understanding the role of noise in various biological mechanisms [3, 10].

Suppose that the dynamics of a reaction network is given by a Markov process $(X(t))_{t \geq 0}$ whose states represent the copy-number vectors of the constituent species. This process evolves on a *discrete* state-space $\mathcal{E}$ which must include all the copy-number vectors that can be *accessed* by the random dynamics. This state-space may not be unique, and for many examples of interest it is countably infinite. Furthermore it can

1

depend on the initial conditions or on the quantities that are conserved by the dynamics. To investigate how the noise generated by the random timing of reaction events, affects a biological system in the long-run, one is often interested in knowing if the $\mathcal{E}$-valued Markov process $(X(t))_{t \geq 0}$ has a unique stationary distribution. This uniqueness will hold if the state-space $\mathcal{E}$ is *irreducible*, which means that for any two states $x, y \in \mathcal{E}$, there is a positive probability for the stochastic reaction dynamics to start at x and reach y in finite time. The aim of this paper is to construct a scalable computational framework for finding such irreducible state-spaces. This can help in understanding the long-term behavior of stochastic models of reaction networks as we now explain.

Consider a reaction network whose Markovian dynamics has an irreducible state-space $\mathcal{E}$. Once the existence of a stationary distribution π is verified, we can conclude that the $\mathcal{E}$-valued Markov process representing the reaction network is *ergodic* [17], i.e. for any initial state $X(0) = x_0 \in \mathcal{E}$, the distribution $p_{x_0}(t)$ of the (random) state $X(t)$ at time t, converges to π in an appropriate sense as $t \to \infty$. Ergodicity represents a strong notion of stability for the stochastic model that is similar to having an attracting fixed point for the deterministic model. The ergodicity of the stochastic model of a biological system enables us to draw many desirable conclusions about the long-term behavior of the underlying stochastic dynamics. For example, the long-run proportion of time spent by the stochastic trajectories in any subset of the state-space is equal to the stationary probability of that subset (see (2.9)). Hence observing the whole population of biological systems (e.g. cells) at stationarity is equivalent to observing just one stochastic trajectory of a single system for a sufficiently long time. Furthermore ergodicity implies that statistical quantities associated with the underlying Markov process, like means or variances, converge to their steady-state values with time (see (2.8)). These insights can be used for leveraging different experimental techniques for biological measurements and for designing controllers that steer the mean of some species to specific steady state values [7]. Apart from checking ergodicity, finding the irreducible state-spaces can also confirm the existence of a product-form stationary distribution for certain networks (see [2]).

We now come to the issue of finding irreducible state-spaces for stochastic reaction dynamics. If the reaction network can only access finitely many states, then such state-spaces can be easily identified using matrix methods or state transition diagrams [27, 22]. However many biochemical networks we encounter do not satisfy this property, and as a consequence all the state-spaces must be necessarily infinite. In such a setting, finding the "right" candidate state-spaces and checking their irreducibility is quite challenging. Before we explain these challenges, let us consider a simple birth-death network of the form

$$\emptyset \longrightarrow \mathbf{S} \longrightarrow \emptyset.$$

Here the molecules of a single-species $\mathbf{S}$ are produced and degraded at some positive rates θ_1 and θ_2 respectively. Let $(X(t))_{t \geq 0}$ be the Markov process representing the reaction dynamics in the stochastic setting, where $X(t)$ is the number of molecules at time t. This process evolves on the state-space $\mathbb{N}_0 = \{0, 1, 2, \dots\}$. When the state is $X(t) = n$, then with probability $p_+(n) = \theta_1/(\theta_1 + \theta_2 n)$, the next reaction is a production reaction and the state increases to $(n+1)$ and with probability $p_-(n) = 1 - p_+(n) = (\theta_2 n)/(\theta_1 + \theta_2 n)$, the next reaction is a degradation reaction and the state decreases to $(n-1)$. From these probabilities, it is immediate that for any two distinct states $x, y \in \mathbb{N}_0$, the reaction dynamics starting at x can reach y, either via a sequence of $(y - x)$ production reactions if $x < y$, or via a sequence of $(x - y)$ degradation reactions if $x > y$. In both these cases, the appropriate sequence of reactions has a positive probability of firing in a finite time, and hence the state-space $\mathbb{N}_0$ is irreducible for the reaction network. For a general network with $d \gg 1$ species and a countable state-space $\mathcal{E}$, guaranteeing the existence of such positive-probability sequences of reactions between every two states in $\mathcal{E}$ is very difficult, especially since some reaction channels can *switch-off* at certain states if the requisite number of molecules for each reacting species are not present. These issues make the direct verification of the irreducibility of an infinite state-space $\mathcal{E}$ nearly impossible. The main contribution of this paper is to provide an approach that avoids these issues and still finds irreducible state-spaces for a large class of biochemical reaction networks.

Notice that for the birth-death network mentioned above, we could easily ascertain the irreducibility of the state-space $\mathcal{E} = \mathbb{N}_0$ because the transition structure of the underlying Markov process is very simple. This suggests the possibility that irreducible state-spaces for biochemical networks can perhaps be easily found if we can identify birth-death structures within them. Inspired by this idea we present a computational procedure for finding irreducible state-spaces that is based on the simple observation that many biochemical reaction networks have *cascaded birth and death subnetworks*. This means that there is a set of species that

2

are produced *directly* due to reactions of the form $\emptyset \longrightarrow \mathbf{S}$. These species then produce another set of species which in turn produce another set of species and so on, creating a *birth cascade*. Similarly there is a set of species that are degraded *directly* due to reactions of the from $\mathbf{S} \longrightarrow \emptyset$. These species are then responsible for degrading another set of species which in turn cause degradation of another set of species and so on, forming a *death cascade*. In many networks all the species can be arranged in birth and death cascades. For example, in a simple gene-expression network (see Example 4.1), messenger RNAs (or mRNAs) are produced directly by a gene transcript and these mRNAs then produce proteins, forming a birth cascade involving all the network species (mRNA and protein). On the other hand both mRNA and protein molecules degrade directly, creating a trivial death cascade with all the species. We show that under some mild linear-algebraic conditions, if all the d species in a general reaction network are included in birth and death cascades, then the nonnegative integer orthant $\mathbb{N}_0^d$ is the unique irreducible state-space for the reaction network. In this situation all the species are *free* in the sense that their copy-numbers can take any nonnegative integer value, irrespective of the copy-numbers of other species. This is of course not true if the reaction network consists of *conservation relations*, which are linear relationships among species copy-numbers that are preserved by the dynamics. Presence of such conservation relations can constrain the species copy-numbers and introduce dependencies among them, thereby complicating the task of finding irreducible state-spaces. We shall explore how these conservation relations can be taken into account for constructing a list of likely candidates for irreducible state-spaces and how the irreducibility of each such state-space can be verified using a suitable combination of matrix methods [22] and birth/death cascade construction mentioned earlier.

Conservation relations limit the copy-number ranges of the involved species in two distinct ways. Some conservation relations can cause a species to be *bounded*, which means that its copy-numbers have a bounded range. For example, in a gene-expression network (see Example 4.1) where a single gene can occur in *active* or *inactive* forms, there is a conservation relation which says that the copy-numbers of these two species (active-gene and inactive-gene) must sum to 1, thereby ensuring that these species are *bounded* as their copy-numbers can either be 0 or 1. Conservation relations can also force certain species to be *restricted*, in the sense that their copy-numbers are an *affine* function of the copy-numbers of the *free* species. As an example, consider a simple two-species network where molecules of both the species are produced or degraded together (see Example 4.2), and hence the difference of their copy-numbers is a conservation relation. If this difference is initially 0, then the copy-number of both the species will remain identical throughout the dynamics. Therefore we can declare one of the species as *free* and then the other will become *restricted* because its copy-number is simply equal to the copy-number of the *free* species. This also suggests that there is a certain degree of flexibility in choosing the *free* and the *restricted* species. We later explain how this flexibility can be exploited to facilitate the search for irreducible state-spaces.

In large reaction networks many conservation relations may be present and hence identifying the copy-number range of each species is a complicated task. We will use standard linear-algebraic methods for this purpose and classify each species as one of three types, *free*, *bounded* or *restricted*, depending on the limitations on its copy-numbers or lack thereof. Such a classification allows us to easily express the possible state-spaces and check their irreducibility. The linear-algebraic methods that we employ are basic matrix manipulations, solving systems of linear equations and solving Linear Programming Problems (LPPs) [6]. Since these methods can be efficiently applied in very high dimensions, our framework is highly scalable and can easily handle large reaction networks with several species and reactions.

The main ideas behind this paper are generalized from our recent conference paper [18]. However the method we develop in this paper is far more versatile and scalable. In particular it can efficiently handle a large number of conservation relations and it exploits the dynamical independence between *bounded* species to reduce the computational complexity of searching for irreducible state-spaces. The theoretical underpinnings of both [18] and this paper are provided by some recent results on irreducibility of Discrete Reaction Networks given in [29]. However, unlike our paper, the results in [29] do not apply when the network has conservation relations.

This paper is organized as follows. In Section 2 we present the relevant mathematical background for this paper. We formally describe the stochastic model of a reaction network, explain the notion of state-space irreducibility and discuss its importance. We also introduce some preliminary concepts and definitions that will be used throughout the paper. Section 3 contains the main results of this paper and provides an algorithmic procedure to find irreducible state-spaces for general reaction networks. In Section 4 we illustrate the applicability of our framework by considering examples from Systems and Synthetic Biology.

Notation

We now introduce some notation that will be used throughout the paper. Let $\mathbb{R}$, $\mathbb{R}_+$, $\mathbb{Z}$, $\mathbb{N}$ and $\mathbb{N}_0$ denote the sets of all reals, nonnegative reals, integers, positive integers and nonnegative integers respectively. For any set A, we denote its cardinality by $|A|$. The vectors of all zeros and all ones in $\mathbb{R}^n$ are written as $\overline{0}_n$ and $\overline{1}_n$ respectively. Moreover $e_i^{(n)}$ denotes a vector in $\mathbb{R}^n$ whose i-th entry is 1 while the rest are 0. The set of vectors $\{e_1^{(n)}, \ldots, e_n^{(n)}\}$ forms the *standard basis* for $\mathbb{R}^n$. For any vector $v = (v_1, \ldots, v_n) \in \mathbb{R}^n$ we define its support as the set $\operatorname{supp}(v) = \{i = 1, \ldots, n : v_i \neq 0\}$. For any two vectors $v, w \in \mathbb{R}^n$ we say $v > w$ or $v \geq w$ if the corresponding inequality holds component-wise. A vector $v \in \mathbb{R}^n$ is called nonzero, nonnegative and strictly positive if $v \neq \overline{0}_n$, $v \geq \overline{0}_n$ and $v > \overline{0}_n$ respectively. A vector which is both nonnegative and nonzero is simply called positive. We denote the standard inner product in $\mathbb{R}^n$ by $\langle \cdot, \cdot \rangle$.

While multiplying a matrix with a vector we always regard the vector as a column vector. Let $\mathbb{M}(m, n)$ refer to the set of all $m \times n$ matrices with real entries. For any matrix $M \in \mathbb{M}(m, n)$, we denote its rank by $\operatorname{Rank}(M)$, its transpose by M^T and its left nullspace by $\mathcal{L}(M) = \{\gamma \in \mathbb{R}^m : \gamma^T M = \overline{0}_n\}$. Given n vectors $v_1, \ldots, v_n \in \mathbb{R}^m$ the matrix in $\mathbb{M}(m, n)$ with these vectors as its columns is denoted by $\operatorname{Col}(v_1, \ldots, v_n)$. Moreover for any $M \in \mathbb{M}(m, n)$ and any $k \leq l \leq m$, the projection matrix $\operatorname{Proj}(M, k, l) \in \mathbb{M}(l - k + 1, n)$ is the submatrix formed by rows $(k + 1), (k + 2), \ldots, l$ of matrix M. For any positive integer n, I_n represents the $n \times n$ identity matrix. The dimension of any vector space V is denoted by $\dim(V)$ and this vector space is called trivial or nontrivial depending on whether $\dim(V) = 0$ or $\dim(V) > 0$. If $v_1, \ldots, v_n$ are the columns of M then for any $A \subset \mathbb{R}$, the set $\operatorname{Colspan}_A(M)$ stands for

$$\left\{ x \in \mathbb{R}^m : x = \sum_{i=1}^{n} a_i v_i \text{ for some } a_1, \ldots, a_n \in A \right\}.$$

2 Mathematical Background

We start this section with the description of the stochastic model of a reaction network (Section 2.1). We then discuss the notion of ergodicity (Section 2.2) and the related notion of state-space irreducibility (Section 2.3). Finally we conclude this section by introducing and explaining some preliminary concepts that will be used later in the paper (Section 2.4).

2.1 Stochastic model of a reaction network

Consider a reaction network with d species $\mathbf{S}_1, \ldots, \mathbf{S}_d$. These species interact through K reaction channels of the form

$$\sum_{i=1}^{d} \nu_{ik} \mathbf{S}_i \longrightarrow \sum_{i=1}^{d} \rho_{ik} \mathbf{S}_i, \qquad k = 1, \ldots, K \tag{2.1}$$

where ν_{ik} and ρ_{ik} denote the number of molecules of $\mathbf{S}_i$ that are consumed and produced by reaction k. Define vectors ν_k and ρ_k in $\mathbb{N}_0^d$ as $\nu_k = (\nu_{1k}, \ldots, \nu_{dk})$ and $\rho_k = (\rho_{1k}, \ldots, \rho_{dk})$ respectively. Under the classical well-stirred assumption [13], the state of the network at any time is given by a vector in $\mathbb{N}_0^d$, whose i-th component is the number of molecules (or copy-number) of $\mathbf{S}_i$. When the state is x, the k-th reaction fires with rate $\lambda_k(x)$ and it displaces the state to $y = x - \nu_k + \rho_k$. The functions $\lambda_1, \ldots, \lambda_K$ are called *propensity* functions for the reaction network. We assume that these functions satisfy the following property: for any $k = 1, \ldots, K$ and $x \in \mathbb{N}_0^d$ if $\lambda_k(x) > 0$ then $x \geq (\nu_k - \rho_k)$. This property ensures that the reaction dynamics never leaves the nonnegative integer orthant $\mathbb{N}_0^d$. For our purpose, we need stricter conditions on the propensity functions that we introduce later (see Section 2.4).

Define a couple of $d \times K$ matrices by $\mathcal{V} = \operatorname{Col}(\nu_1, \ldots, \nu_K)$ and $\mathcal{O} = \operatorname{Col}(\rho_1, \ldots, \rho_K)$. Let $\Lambda : \mathbb{N}_0^d \to \mathbb{R}_+^K$ be the *propensity* map given by

$$\Lambda(x) = (\lambda_1(x), \ldots, \lambda_K(x)). \tag{2.2}$$

The above reaction network with d species and K reactions can be fully described by the *triplet*

$$\mathcal{N} = (\mathcal{V}, \mathcal{O}, \Lambda). \tag{2.3}$$

In the stochastic model for network $\mathcal{N}$ we represent the reaction dynamics as a $\mathbb{N}_0^d$-valued Markov process with generator[1] given by

$$\mathbb{A}_{\mathcal{N}} f(x) = \sum_{k=1}^{K} \lambda_k(x) \left(f(x - \nu_k + \rho_k) - f(x) \right),$$

where f is any bounded real-valued function on $\mathbb{N}_0^d$. Under mild conditions on the propensity functions, one can show that for any initial state $x_0 \in \mathbb{N}_0^d$, there exists a well-defined Markov process $(X(t))_{t \geq 0}$ with generator $\mathbb{A}_{\mathcal{N}}$ and $X(0) = x_0$. A state-space for this stochastic model of network $\mathcal{N}$ is any subset of the nonnegative integer orthant which is *closed* under the reaction dynamics.

Definition 2.1 *Consider the reaction network $\mathcal{N} = (\mathcal{V}, \mathcal{O}, \Lambda)$ with d species and K reactions. A non-empty set $\mathcal{E} \subset \mathbb{N}_0^d$ is called a state-space for this network if the following is satisfied for each $x \in \mathcal{E}$ and $k = 1, \ldots, K$:*

$$\text{if} \quad \lambda_k(x) > 0 \quad \text{then} \quad (x + \rho_k - \nu_k) \in \mathcal{E}.$$

Observe that according to this definition, $\mathbb{N}_0^d$ is always a state-space for any network with d species, because we have assumed that if $\lambda_k(x) > 0$ then $x \geq (\nu_k - \rho_k)$ and hence $(x + \rho_k - \nu_k) \in \mathbb{N}_0^d$. However depending upon the network structure and reaction stoichiometries, there may exist smaller state-spaces for a given reaction network. Once a state-space $\mathcal{E}$ has been selected for a reaction network $\mathcal{N} = (\mathcal{V}, \mathcal{O}, \Lambda)$, any Markov process $(X(t))_{t \geq 0}$ with generator $\mathbb{A}_{\mathcal{N}}$ and initial state $X(0) \in \mathcal{E}$, will satisfy $X(t) \in \mathcal{E}$ for all $t \geq 0$. Therefore $\mathcal{E}$ serves as a generic state-space for all Markov processes representing the reaction dynamics and starting with an initial state in $\mathcal{E}$.

Fix a state-space $\mathcal{E}$ for network $\mathcal{N} = (\mathcal{V}, \mathcal{O}, \Lambda)$. Let $\mathcal{P}(\mathcal{E})$ be the space of all probability distributions on $\mathcal{E}$. We endow this space with the Prohorov metric which generates the weak topology [11]. Pick a $\mu \in \mathcal{P}(\mathcal{E})$ and let $(X(t))_{t \geq 0}$ be the Markov process with generator $\mathbb{A}_{\mathcal{N}}$ and initial distribution μ. Hence $(X(t))_{t \geq 0}$ represents the reaction dynamics and the initial state $X(0)$ is distributed according to μ. For any $y \in \mathcal{E}$, the probability that the reaction dynamics is at state y at time t is

$$p_\mu(t, y) = \mathbb{P}\left(X(t) = y \right). \tag{2.4}$$

The dynamics of $p_\mu(t)$ is given by the Chemical Master Equation (CME) [14] which has the following form. For each $y \in \mathcal{E}$

$$\frac{dp_\mu(t, y)}{dt} = \sum_{k=1}^{K} \left(p_\mu(t, y - \rho_k + \nu_k) \lambda_k(y - \rho_k + \nu_k) - p_\mu(t, y) \lambda_k(y) \right) \tag{2.5}$$

where $p_\mu(0, y) = \mu(\{y\})$. Observe that this system consists of as many equations as the number of elements in $\mathcal{E}$, which is typically infinite or very large, and hence solving the CME is nearly impossible for most examples. Commonly solutions of CME are estimated by simulating the process $(X(t))_{t \geq 0}$ using Gillespie's *Stochastic Simulation Algorithm* (SSA) [13] or its variants. Even though this simulation-based approach is useful for analyzing the stochastic model over a finite time-interval, it does not help in assessing its long-term behavior satisfactorily. We need new tools to analyze the long-term behavior and the contribution of our paper in this regard is discussed next.

2.2 Ergodicity as a notion of stability

Define $p_\mu(t, A) = \sum_{y \in A} p_\mu(t, y)$ for any $A \subset \mathcal{E}$. We can view $p_\mu(t)$ as a probability distribution over $\mathcal{E}$ and hence as an element in $\mathcal{P}(\mathcal{E})$. A distribution $\pi \in \mathcal{P}(\mathcal{E})$ is called a *stationary distribution* for the reaction network $\mathcal{N} = (\mathcal{V}, \mathcal{O}, \Lambda)$ if $p_\pi(t, y) = \pi(y)$ for all $t \geq 0$ and $y \in \mathcal{E}$, where $\pi(y) = \pi(\{y\})$. This means that if the initial distribution is π, then the distribution of the network state at any time t is also π. The form of the CME (2.5) implies π is a stationary distribution if and only if for each $y \in \mathcal{E}$ we have

$$\sum_{k=1}^{K} \pi(y - \rho_k + \nu_k) \lambda_k(y - \rho_k + \nu_k) = \pi(y) \sum_{k=1}^{K} \lambda_k(y). \tag{2.6}$$

---

[1]The generator of a stochastic process is an operator which captures the rate of change of the distribution of the process. See Chapter 4 in [11] for more details.

From this relation it is clear that a stationary distribution π can be viewed as a fixed point for the CME (2.5) which describes the time-evolution of the probability distributions $(p_\mu(t))_{t \geq 0}$ in the space $\mathcal{P}(\mathcal{E})$. If there exists a unique such fixed point π which is *globally attracting* in $\mathcal{P}(\mathcal{E})$, then we say that the stochastic reaction dynamics is *ergodic*. This is same as saying that for any $\mu \in \mathcal{P}(\mathcal{E})$, the $\mathcal{E}$-valued Markov process $(X(t))_{t \geq 0}$ with generator $\mathbb{A}_\mathcal{N}$ and initial distribution μ is *ergodic* [25], in the sense that the probability distribution $p_\mu(t)$, defined by (2.4), converges to the stationary distribution π in the *total-variation* norm on $\mathcal{P}(\mathcal{E})$:

$$\lim_{t \to \infty} \sup_{A \subset \mathcal{E}} |p_\mu(t, A) - \pi(A)| = 0, \tag{2.7}$$

for any $\mu \in \mathcal{P}(\mathcal{E})$. This also means that for any $A \subset \mathcal{E}$, the probability of the event $\{X(t) \in A\}$ converges to $\pi(A)$ as $t \to \infty$, irrespective of the initial distribution μ.

Ergodicity represents a strong notion of *stability* for stochastic reaction dynamics, which can have many practical applications as we now discuss. Ergodicity implies that for any bounded real-valued function f on $\mathcal{E}$ we have

$$\lim_{t \to \infty} \mathbb{E}(f(X(t))) = \sum_{y \in \mathcal{S}} f(y)\pi(y) \tag{2.8}$$

and the following limit holds with probability 1

$$\text{and } \lim_{t \to \infty} \frac{1}{t} \int_0^t f(X(s))ds = \sum_{y \in \mathcal{S}} f(y)\pi(y). \tag{2.9}$$

For a proof of these relations see Theorem 1.10.2 in [27]. Even though these relations generally hold only for the class of bounded functions, for many reaction networks it is possible to extend this class to include all polynomially growing functions [17]. This extension allows us to use (2.8) to show that moments (means, variances, covariances etc.) of the stochastic reaction dynamics converge to their *steady state* values as $t \to \infty$ (see [17]). Such results can be used to design synthetic controllers that robustly steer the moments of certain species to specific steady-state values [7]. Relation (2.9) shows that the stationary distribution of the population can be computed by evaluating the proportion of time spent in various states by a single stochastic trajectory $(X(t))_{t \geq 0}$ over a long period of time. Such an insight can help in leveraging experimental or computational techniques for estimating the long-term behavior of a population of identical cells, where each cell includes the same reaction network with ergodic Markovian dynamics. For example, the steady-state behavior of a cell population (r.h.s. of (2.9)), as observed by *Flow-Cytometry*, will closely resemble the time-averaged behavior of a single cell (l.h.s. of (2.9)), as observed by *Time-Lapse Microscopy*. Hence the "right" experimental technique can be chosen for a given situation, based on convenience, and the long-term behavior of the stochastic system can be studied. Similarly one can use (2.9) to *speed-up* the estimation of the stationary distribution π using computer simulations.

We now come to the problem of checking ergodicity. Clearly a direct approach is fraught with many complications, because the CME (2.5) cannot be solved in most cases, and even if we can solve it or estimate its solutions using simulations, it is impossible to verify that for some stationary distribution π, (2.7) holds for any initial distribution $\mu \in \mathcal{P}(\mathcal{E})$. Due to these issues we need an indirect approach for establishing ergodicity. Fortunately such an indirect approach exists for continuous-time Markov chain models and the key step is to show that the underlying state-space $\mathcal{E}$ is *irreducible*. In the next section, we define this property and discuss the challenges of finding irreducible state-spaces for stochastic models of reaction networks.

For now assume that an irreducible state-space $\mathcal{E}$ has been found for a reaction network. This ensures the uniqueness of the stationary distribution π in $\mathcal{P}(\mathcal{E})$ (see Theorem 8.18 in [20]), and to establish ergodicity over this state-space we only need to show that such a stationary distribution π exists. If the state-space $\mathcal{E}$ is compact (or finite) such a stationary distribution certainly exists (see Section 1.7 in [27]). However its existence cannot be guaranteed in the more common scenario of $\mathcal{E}$ being countably infinite. In this situation, the existence of a stationary distribution can be checked using the results by Meyn and Tweedie [25, 26]. In particular, Theorem 4.5 in [26] shows that a stationary distribution π will exist if one can find a function $V : \mathcal{E} \to \mathbb{R}_+$ satisfying $V(x) \to \infty$ as $\|x\| \to \infty$, another function $f : \mathcal{E} \to [1, \infty)$, a compact set $C \subset \mathcal{E}$ and some constants $c, d > 0$ such that

$$\mathbb{A}_\mathcal{N} V(x) \leq -cf(x) + d\mathbb{1}_C(x) \qquad \text{for all } x \in \mathcal{E}. \tag{2.10}$$

This condition is called *Foster-Lyapunov* criterion in the literature [26] and it essentially says that outside some compact set $C \subset \mathcal{E}$, the Markovian dynamics experiences a *negative drift* given by function f. In a recent paper [17] we develop a computational framework for constructing such norm-like functions V satisfying the above Foster-Lyapunov criterion for a large class of biochemical reaction networks which includes several well-known examples from Systems and Synthetic Biology. Interestingly these functions have a simple *linear* form given by $V(x) = \langle v, x \rangle$, where $v \in \mathbb{R}^d_+$ is a positive vector which is suitably chosen using optimization techniques, such as Linear or Semidefinite Programming [6].

2.3 Irreducibility of a state-space

Consider the reaction network $\mathcal{N} = (\mathcal{V}, \mathcal{O}, \Lambda)$ with some state-space $\mathcal{E}$ (recall Definition 2.1). For any $x \in \mathcal{E}$, let δ_x be the Dirac distribution concentrated at x. Define $p_x(t)$ to be $p_\mu(t)$ for $\mu = \delta_x$. For any $y \in \mathcal{E}$, $p_x(t, y)$ is the probability that the stochastic dynamics starts at x and reaches y at time t. If $p_x(t, y) > 0$ for some $t \geq 0$, then we say that state y is *reachable* from state x, and we denote this relation as

$$x \xrightarrow{\mathcal{N}} y. \tag{2.11}$$

Since we are in the continuous-time setting, $p_x(t, y) > 0$ for some $t \geq 0$ if and only if $p_x(t, y) > 0$ for all $t \geq 0$ (see Theorem 3.2.1 in [27]). Hence the condition for reachability $p_x(t, y) > 0$ does not depend on the particular time-value t. Moreover the reachability relation $\xrightarrow{\mathcal{N}}$ is **transitive**: i.e. for any $x, y, z \in \mathcal{E}$, if $x \xrightarrow{\mathcal{N}} y$ and $y \xrightarrow{\mathcal{N}} z$ then $x \xrightarrow{\mathcal{N}} z$ (see Chapter 6 in [16]).

We say that the state-space $\mathcal{E}$ for a reaction network $\mathcal{N} = (\mathcal{V}, \mathcal{O}, \Lambda)$ is *irreducible* if all the states in $\mathcal{E}$ are reachable from each other, i.e. relation (2.11) holds between any two states $x, y \in \mathcal{E}$. The problem of checking irreducibility can be quite challenging when $\mathcal{E}$ is countably infinite. The main difficulty arises due to the fact that some reaction channels may *switch-off* at certain states, and hence the set of possible transition directions is not the same for all states in $\mathcal{E}$. When the state is x, the reaction channel k will switch-off if its propensity $\lambda_k(x)$ is 0. This switching-off of reactions must be taken into account while checking the irreducibility of $\mathcal{E}$.

Observe that (2.11) certainly holds if $x = y$. In order to prove this relation between two distinct states $x, y \in \mathcal{E}$, we need to show that there is a sequence of n reactions $k_1, \ldots, k_n \in \{1, \ldots, K\}$ such that:

$$y = x + \sum_{i=1}^{n} (\rho_{k_i} - \nu_{k_i}) \quad \text{and} \quad \lambda_{k_j}(z_j) > 0, \tag{2.12}$$

for each $j = 1, \ldots, n$, where $z_j = x + \sum_{i=1}^{j-1}(\rho_{k_i} - \nu_{k_i})$. These conditions imply that if the initial state is x and reactions $k_1, \ldots, k_n$ fire in this order, then the final state of the network will be y. Of course these reactions can only fire in this order, if at all the *intermediate* states (z_j-s), the propensity $\lambda_{k_j}(z_j)$ for the next reaction in this sequence is positive. Proving the existence of a sequence of reactions that satisfies these conditions is technically difficult. Hence it is hard to directly verify the irreducibility of an infinite state-space $\mathcal{E}$ by showing that (2.11) holds for every pair of states x and y in $\mathcal{E}$.

To avoid the difficulties mentioned above, we adopt an indirect approach for checking state-space irreducibility for reaction networks. Our approach is inspired by the structure of networks commonly found in Systems and Synthetic Biology. We present this approach in detail in Section 3 and it does not involve finding sequences of reactions connecting states in the state-space. Instead it only relies on elementary linear algebra and linear programming, which makes it highly scalable for large reaction networks.

Apart from establishing ergodicity, finding irreducible state-spaces can also enable us to compute the exact stationary distribution for certain networks whose propensity functions satisfy mass-action kinetics (see (2.15)) and whose deterministic models admit *complex balanced* equilibria. Indeed the results in [2] prove that the stochastic models for such networks have the following property: for each irreducible state-space $\mathcal{E}$, the unique stationary distribution $\pi \in \mathcal{P}(\mathcal{E})$ is given by the product-form

$$\pi(x) = M_\mathcal{E} \prod_{i=1}^{d} \frac{c_i^{x_i}}{x_i!} e^{-c_i} \quad \text{for} \quad x = (x_1, \ldots, x_d) \in \mathcal{E}, \tag{2.13}$$

where d is the number of species, $c = (c_1, \ldots, c_d)$ is a positive vector in $\mathbb{R}^d$ and $M_{\mathcal{E}}$ is the normalizing constant given by

$$M_{\mathcal{E}} = \left(\sum_{x \in \mathcal{E}} \prod_{i=1}^d \frac{c_i^{x_i}}{x_i!} e^{-c_i} \right)^{-1}. \tag{2.14}$$

Our method not only finds the irreducible state-spaces $\mathcal{E}$ on which this result can be applied, but in cases where $\mathcal{E}$ is countably infinite, it expresses $\mathcal{E}$ in such a way that the infinite sum in (2.14) can be replaced by a finite sum. This allows the normalizing constant $M_{\mathcal{E}}$ to be easily calculated without any truncation errors (see Example 4.5), yielding the exact stationary distribution. Moreover for many networks our method can provably find *all the irreducible state-spaces* $\mathcal{E}_1, \ldots, \mathcal{E}_Q$. Assuming we can exactly compute the stationary distribution $\pi_q \in \mathcal{P}(\mathcal{E}_q)$ of the form (2.13), for each $q = 1, \ldots, Q$, these distributions $\pi_1, \ldots, \pi_Q$ will form the *extremal points* of the simplex formed by all the stationary distributions of the network (see [2]). We end this section with an important remark.

Remark 2.2 *Suppose that $\mathcal{E}_1$ and $\mathcal{E}_2$ are two irreducible state-spaces for a network $\mathcal{N}$. Then these two state-spaces must be necessarily disjoint $\mathcal{E}_1 \cap \mathcal{E}_2 = \emptyset$. To see this, note that if $\mathcal{E}_1 \cap \mathcal{E}_2$ is non-empty, then it is also a valid state-space (see Definition 2.1). Hence the states in $\mathcal{E}_1 \cap \mathcal{E}_2^c$ or $\mathcal{E}_1^c \cap \mathcal{E}_2$ are not reachable from the states in $\mathcal{E}_1 \cap \mathcal{E}_2$, violating the irreducibility of $\mathcal{E}_1$ and $\mathcal{E}_2$ respectively.*

2.4 Preliminaries

In this section we present some preliminary concepts that will be used throughout the paper. We begin by imposing some restrictions on the form of the propensity functions and defining the *inverse* of a reaction network. Often it will be difficult to find irreducible state-spaces for the original network directly, and so we will construct an *equivalent* reaction network by rearranging the species. We describe this construction formally and then discuss how the presence of *conservation relations* in the network complicates the search for irreducible state-spaces. Finally we define the notion of *conservation data* for a network, and explain how it must be taken into account while finding state-spaces for a network.

2.4.1 Restrictions on the propensity functions

Recall the definition of a reaction network $\mathcal{N} = (\mathcal{V}, \mathcal{O}, \Lambda)$ with d species and K reactions of the form (2.1) (see Section 2.1). Here $\Lambda(x) = (\lambda_1(x), \ldots, \lambda_K(x))$ is the positive vector denoting the propensities of the K reactions at state x and $\mathcal{V} = \mathrm{Col}(\nu_1, \ldots, \nu_K)$ is the matrix whose columns contain the vectors of reactant molecular counts that are *consumed* by the K reactions. Throughout the paper we only consider networks that satisfy the following:

Assumption 2.3 *For each reaction $k = 1, \ldots, K$ and each $x \in \mathbb{N}_0^d$, we have $\lambda_k(x) > 0$ if and only if $x \geq \nu_k$.*

This assumption essentially means that when the state is $x = (x_1, \ldots, x_d)$, reaction k has a positive probability of firing if any only if for each species $\mathbf{S}_i$, the number of available molecules (x_i) exceeds the number of molecules consumed by the reaction (ν_{ik}). Observe that the "only if" part of this condition is nearly always satisfied, because a reaction cannot fire unless for each species, the required number of molecules are present for consumption, but the "if" part of this condition may get violated if there is a reaction k and a state x satisfying $x \geq \nu_k$ and $\lambda_k(x) = 0$. However such situations do not typically arise for biochemical reaction networks as we now explain.

Note that Assumption 2.3 is certainly satisfied if we have *mass-action kinetics* [15] where each propensity function $\lambda_k : \mathbb{N}_0^d \to \mathbb{R}_+$ has the form

$$\lambda_k(x) = \theta_k \prod_{i=1}^d \frac{x_i(x_i - 1) \ldots (x_i - \nu_{ik} + 1)}{\nu_{ik}!}, \tag{2.15}$$

for some rate constant $\theta_k > 0$. Apart from mass-action kinetics, networks in Systems and Synthetic Biology generally have propensity functions describing either Michaelis–Menten or Hill-type kinetics [21]. Michaelis–Menten kinetics usually appears when the network involves enzyme-substrate interactions [30] while Hill

type kinetics usually arises in gene-expression networks with feedback regulation [31]. In both these cases, the propensity functions have a rational form given by

$$\lambda_k(x) = \frac{p_k(x)}{q_k(x)},$$

where the denominator $q_k(x)$ is always positive and the numerator $p_k(x)$ satisfies the criterion in Assumption 2.3. As a consequence, the network satisfies Assumption 2.3 even though its propensity functions are not of *mass-action* type. Finally we point out even if a network does not satisfy Assumption 2.3, it can often be modified in such a way that its dynamics remains the same but the modified network satisfies Assumption 2.3 (see Example 4.1). We end this section with a simple proposition.

Proposition 2.4 *Suppose $\mathcal{N}$ is a reaction network satisfying Assumption 2.3 and let $\mathcal{E} \subset \mathbb{N}_0^d$ be a state-space for this network. Then relation $\xrightarrow{\mathcal{N}}$ is* **additive***: i.e. for any $x, y \in \mathcal{E}$ and $z \in \mathbb{N}_0^d$, if $x \xrightarrow{\mathcal{N}} y$, $(x + z) \in \mathcal{E}$ and $(y + z) \in \mathcal{E}$ then $(x + z) \xrightarrow{\mathcal{N}} (y + z)$.*

Proof. Due to Assumption 2.3, for any reaction k and state $u \in \mathcal{E}$, if $\lambda_k(u) > 0$ then $\lambda_k(u + z) > 0$ for any $z \in \mathbb{N}_0^d$. Therefore if $k_1, \ldots, k_n$ is the sequence of reactions satisfying (2.12), then the same sequence of reactions will also satisfy (2.12) with x and y replaced by $(x + z)$ and $(y + z)$ respectively. This proves $(x + z) \xrightarrow{\mathcal{N}} (y + z)$ and completes the proof of this proposition. $\qquad\qquad\square$

2.4.2 Inverse of a reaction network

Consider a reaction network $\mathcal{N} = (\mathcal{V}, \mathcal{O}, \Lambda)$ with d species and K reactions of the form (2.1). We now define another reaction network $\mathcal{N}_{\text{inv}}$ which can be viewed as the *inverse* of $\mathcal{N}$. This new network has the same number of species and reactions, but its reactions are obtained by *flipping* the arrows in (2.1). In other words, the K reactions in $\mathcal{N}_{\text{inv}}$ are given by

$$\sum_{i=1}^d \rho_{ik} \mathbf{S}_i \longrightarrow \sum_{i=1}^d \nu_{ik} \mathbf{S}_i, \qquad k = 1, \ldots, K. \tag{2.16}$$

To each reaction k we assign the propensity function $\lambda_{k,\text{inv}}$ by the mass-action form (2.15) with $\theta_k = 1$ and ν_{ik} replaced by ρ_{ik} for each i. Let the propensity map $\Lambda_{\text{inv}} : \mathbb{N}_0^d \to \mathbb{R}_+^K$ be given by

$$\Lambda_{\text{inv}}(x) = (\lambda_{1,\text{inv}}(x), \ldots, \lambda_{K,\text{inv}}(x)).$$

Setting $\mathcal{V}_{\text{inv}} = \mathcal{O}$ and $\mathcal{O}_{\text{inv}} = \mathcal{V}$, the inverse reaction network $\mathcal{N}_{\text{inv}}$ can be formally described by the *triplet* $\mathcal{N}_{\text{inv}} = (\mathcal{V}_{\text{inv}}, \mathcal{O}_{\text{inv}}, \Lambda_{\text{inv}})$. Note that we have chosen the propensity map Λ_{inv} in such a way, that network $\mathcal{N}_{\text{inv}}$ will satisfy Assumption 2.3 even if network $\mathcal{N}$ does not.

Suppose $x, y \in \mathbb{N}_0^d$ are two states such that the stochastic dynamics under the original reaction network $\mathcal{N}$ can reach state y from state x by a single firing of reaction k. In such a scenario we have $\lambda_k(x) > 0$ and $y = x - \nu_k + \rho_k$. Assuming that network $\mathcal{N}$ satisfies Assumption 2.3, we must have $x \geq \nu_k$ which implies that $y \geq \rho_k$ and hence $\lambda_{k,\text{inv}}(y) > 0$. Since $x = y - \rho_k + \nu_k$ and $\lambda_{k,\text{inv}}(y) > 0$, under the inverse network $\mathcal{N}_{\text{inv}}$, the stochastic dynamics can reach state x from state y by a single firing of reaction k. Extending this idea to incorporate a sequence of intermediate states and reactions, one can easily see that

$$x \xrightarrow{\mathcal{N}} y \qquad \text{if and only if} \qquad y \xrightarrow{\mathcal{N}_{\text{inv}}} x, \tag{2.17}$$

where the relation $\xrightarrow{\mathcal{N}_{\text{inv}}}$ is defined in the same way as relation $\xrightarrow{\mathcal{N}}$ (see Section 2.3). A simple consequence of this relationship is given as a remark below.

Remark 2.5 *Assume that $\mathcal{E} \subset \mathbb{N}_0^d$ serves as a state-space for both networks $\mathcal{N}$ and $\mathcal{N}_{\text{inv}}$ (see Definition 2.1). Then $\mathcal{E}$ is irreducible for network $\mathcal{N}$ if and only if it is irreducible for network $\mathcal{N}_{\text{inv}}$.*

2.4.3 Reaction network under a permutation

Consider a network $\mathcal{N} = (\mathcal{V}, \mathcal{O}, \Lambda)$ with K reactions and d species $\mathbf{S}_1, \ldots, \mathbf{S}_d$. In the rest of the paper we refer to $\mathbf{S}_i$ as "species i" and denote the set of all species by $\mathcal{D} = \{1, \ldots, d\}$. Let $\sigma : \mathcal{D} \to \mathcal{D}$ be any permutation (one-to-one and onto) map. When we say that $x = (x_1, \ldots, x_d)$ is the state vector under permutation σ, we imply that x_i denotes the copy-number of species $\sigma(i)$ for each $i = 1, \ldots, m$. Essentially the map σ defines the correspondence between the species and the location of their copy-numbers in the state vector. In our original set-up (see Section 2.1) we described the Markovian state-dynamics $(X(t))_{t \geq 0}$ under the identity permutation σ_{id}[2]. However in order to facilitate the search for irreducible state-spaces for $\mathcal{N}$, we will often need to work with an equivalent network $\mathcal{N}^\sigma$ which describes the dynamics under a suitably constructed permutation map σ.

To define $\mathcal{N}^\sigma$ formally, let P_σ be the following $d \times d$ permutation matrix

$$P_\sigma = \mathrm{Col}\left(e^{(d)}_{\sigma^{-1}(1)}, \ldots, e^{(d)}_{\sigma^{-1}(d)}\right), \tag{2.18}$$

where $e^{(d)}_1, \ldots, e^{(d)}_d$ are the standard basis vectors in $\mathbb{R}^d$ and let σ^{-1} denote the inverse of map σ. Note that P_σ is an orthogonal matrix and so its inverse satisfies $P_\sigma^{-1} = P_\sigma^T = P_{\sigma^{-1}}$. From now on, for any $A \subset \mathbb{R}^d$, the set $P_\sigma A \subset \mathbb{R}^d$ is defined as $P_\sigma A = \{P_\sigma x : x \in A\}$. For each reaction k let vectors $\nu_k^\sigma, \rho_k^\sigma \in \mathbb{N}_0^d$ and the propensity function $\lambda_k^\sigma : \mathbb{N}_0^d \to \mathbb{R}_+$ be given by

$$\nu_k^\sigma = P_\sigma \nu_k, \quad \rho_k^\sigma = P_\sigma \rho_k \quad \text{and} \quad \lambda_k^\sigma(x) = \lambda_k(P_\sigma^T x). \tag{2.19}$$

Moreover let the *propensity* map $\Lambda^\sigma : \mathbb{N}_0^d \to \mathbb{R}_+^K$ be as in (2.2) with each λ_k replaced by λ_k^σ. Setting the two $d \times K$ matrices $\mathcal{V}^\sigma$ and $\mathcal{O}^\sigma$ as $\mathcal{V}^\sigma = P_\sigma \mathcal{V}$ and $\mathcal{O}^\sigma = P_\sigma \mathcal{O}$, we define the new *permuted* reaction network $\mathcal{N}^\sigma$ by the triplet

$$\mathcal{N}^\sigma = (\mathcal{V}^\sigma, \mathcal{O}^\sigma, \Lambda^\sigma). \tag{2.20}$$

The networks $\mathcal{N}$ and $\mathcal{N}^\sigma$ are dynamically equivalent, because if $(X(t))_{t \geq 0}$ represents the stochastic reaction dynamics under network $\mathcal{N}$ then $(X^\sigma(t))_{t \geq 0}$ represents the stochastic reaction dynamics under the permuted network $\mathcal{N}^\sigma$ where

$$X^\sigma(t) = P_\sigma X(t) \qquad \text{for all } t \geq 0. \tag{2.21}$$

Due to this relation we have the following proposition.

Proposition 2.6 *A state-space $\mathcal{E}^\sigma$ is irreducible for reaction network $\mathcal{N}^\sigma$ if and only if the state-space $\mathcal{E} = P_\sigma^T \mathcal{E}^\sigma$ is irreducible for reaction network $\mathcal{N}$.*

In Section 3 we will construct a permutation map σ such that the irreducible state-spaces for the permuted network $\mathcal{N}^\sigma$ can be easily found. Then this proposition will help us in recovering the corresponding irreducible state-spaces for the original network $\mathcal{N}$. We end this section with a simple observation.

Remark 2.7 *Suppose that network $\mathcal{N} = (\mathcal{V}, \mathcal{O}, \Lambda)$ satisfies Assumption 2.3. Then the permuted network $\mathcal{N}^\sigma$ will also satisfy this assumption since for each reaction k and state $x \in \mathbb{N}_0^d$ we have*

$$\lambda_k^\sigma(x) > 0 \quad \Leftrightarrow \quad \lambda_k(P_\sigma^T x) > 0 \quad \Leftrightarrow \quad P_\sigma^T x \geq \nu_k \quad \Leftrightarrow \quad P_\sigma P_\sigma^T x \geq P_\sigma \nu_k \quad \Leftrightarrow \quad x \geq \nu_k^\sigma,$$

where '$\Leftrightarrow$' denotes "if and only if".

2.4.4 Incorporating conservation relations into the network dynamics

In Section 1, we mentioned that the presence of *conservation relations* in a network can introduce complex dependencies among the copy-numbers of various species. We now discuss how these dependencies must be taken into account while choosing a suitable state-space for the network. Fix a reaction network $\mathcal{N} = (\mathcal{V}, \mathcal{O}, \Lambda)$ with d species and K reactions. We define its $d \times K$ *stoichiometry matrix* by $S = \mathrm{Col}(\rho_1 -$

---

[2]defined by $\sigma_{\mathrm{id}}(i) = i$ for each $i \in \mathcal{D}$

$\nu_1, \ldots, \rho_K - \nu_K) = \mathcal{O} - \mathcal{V}$. The k-th column of S contains the displacement $(\rho_k - \nu_k)$ caused by reaction k to the state of the network. A conservation relation for the network $\mathcal{N}$ is any nonzero vector in the *left nullspace* of S:

$$\mathcal{L}(S) = \{\gamma \in \mathbb{R}^d : \gamma^T S = \overline{0}_K^T\}. \tag{2.22}$$

To see why any nonzero $\gamma \in \mathcal{L}(S)$ is a conservation relation, observe that the displacement vector $(\rho_k - \nu_k)$ is orthogonal to γ: i.e. $\langle \gamma, \rho_k - \nu_k \rangle = 0$ for each $k = 1, \ldots, K$. Hence any Markov process $(X(t))_{t \geq 0}$ representing the stochastic reaction dynamics will satisfy

$$\langle \gamma, X(t) \rangle = \langle \gamma, X(0) \rangle \qquad \text{for all } t \geq 0, \tag{2.23}$$

which is a conservation relation among the copy-numbers of species included in the set $\text{supp}(\gamma)$. This set is nonempty because γ is nonzero.

If $d = \text{Rank}(S)$ then the left nullspace $\mathcal{L}(S)$ will be trivial and hence there are no conservation relations. Let us consider the other situation where $\mathcal{L}(S)$ is non-trivial and has dimension $n = \dim(\mathcal{L}(S)) > 0$. We choose a basis $\{\gamma_1, \ldots, \gamma_n\}$ for $\mathcal{L}(S)$ ans define a $d \times n$ matrix by $\Gamma = \text{Col}(\gamma_1, \ldots, \gamma_n)$. We know from (2.23) that for each γ_i, the quantity $\langle \gamma_i, X(t) \rangle$ will have a constant value c_i throughout the dynamics. These constant values $c_1, \ldots, c_n$ must be chosen at the outset according to the initial conditions or other system specific parameters (see Examples 4.1 and 4.4). We refer to $c = (c_1, \ldots, c_n) \in \mathbb{R}^n$ as the *conservation vector*, $\Gamma \in \mathbb{M}(d, n)$ as the *conservation matrix* and the pair (Γ, c) as the *conservation data*. A network $\mathcal{N}$ along with its conservation data (Γ, c) is called a *conservation network* and denoted by $(\mathcal{N}, \Gamma, c)$. If the process $(X(t))_{t \geq 0}$ represents the reaction dynamics for such a network, then using (2.23) we obtain

$$\Gamma^T X(t) = c \qquad \text{for all } t \geq 0. \tag{2.24}$$

This implies that any state-space for the conservation network $(\mathcal{N}, \Gamma, c)$ must be contained in the following set which is also valid state-space according to Definition 2.1

$$\mathcal{E}_0 = \{x \in \mathbb{N}_0^d : \Gamma^T x = c\}. \tag{2.25}$$

It is difficult to directly check the irreducibility of $\mathcal{E}_0$ or find its irreducible subsets, because this definition does not clearly express the constraints on the range of copy-numbers of each species. To remedy this problem, we will construct a permutation map σ such that for the equivalent permuted network $\mathcal{N}^\sigma$ (see Section 2.4.3), the set $\mathcal{E}_0^\sigma$ corresponding to $\mathcal{E}_0$ (see Proposition 2.6) has a simpler form. In particular, the copy-number ranges appear explicitly in $\mathcal{E}_0^\sigma$, which facilitates the search for irreducible subsets (see Section 3). We end this section with a remark about the *conservation data* for the permuted network $\mathcal{N}^\sigma$.

Remark 2.8 *Consider a conservation network $(\mathcal{N}, \Gamma, c)$. For a permutation map $\sigma : \mathcal{D} \to \mathcal{D}$, let $\mathcal{N}^\sigma = (\mathcal{V}^\sigma, \mathcal{O}^\sigma, \Lambda^\sigma)$ be the equivalent reaction network constructed as in Section 2.4.3 and let $S_\sigma = (\mathcal{O}^\sigma - \mathcal{V}^\sigma)$ be its stoichiometry matrix. Then conservation relations for network $\mathcal{N}^\sigma$ are nonzero vectors in the left nullspace $\mathcal{L}(S_\sigma) = P_\sigma \mathcal{L}(S)$. Moreover relations (2.24) and (2.21) imply that the conservation data for network $\mathcal{N}^\sigma$ must be (Γ_σ, c) where $\Gamma_\sigma = P_\sigma \Gamma$. We refer to $(\mathcal{N}^\sigma, \Gamma_\sigma, c)$ as the permuted conservation network.*

3 Procedure to find irreducible state-spaces

Let us fix a network $\mathcal{N} = (\mathcal{V}, \mathcal{O}, \Lambda)$ with d species and K reactions, along with its *conservation data* (Γ, c) (see Section 2.4.4). Our aim is to find irreducible state-spaces for this conservation network inside the set $\mathcal{E}_0$ defined by (2.25). Note that if there are no conservation relations then there are no restrictions on the copy-number ranges and $\mathcal{E}_0$ is simply the nonnegative integer orthant $\mathbb{N}_0^d$. However the presence of conservation relations creates complex dependencies among the species copy-numbers and also constrains their possible values. In what follows, we will use linear-algebraic techniques to *unravel* these dependencies and classify each species as *free*, *bounded* or *restricted* (see Section 1) according to its copy-number range. We then choose a permutation map σ so that for the equivalent permuted conservation network $(\mathcal{N}^\sigma, \Gamma_\sigma, c)$ (see Section 2.4.3), the candidates for the irreducible state-spaces have the following form:

$$\mathcal{E}^\sigma = \mathcal{F}_b \times \Phi, \tag{3.26}$$

where $\mathcal{F}_b$ is a finite set in $\mathbb{N}_0^{d_b}$ and Φ is the *graph* of an *affine* function $\phi : \mathbb{N}_0^{d_f} \to \mathbb{N}_0^{d_r}$ defined by

$$\Phi = \left\{ (x, \phi(x)) \in \mathbb{N}_0^{d_f + d_r} : x \in \mathbb{N}_0^{d_f} \right\}. \tag{3.27}$$

Here d_f, d_b and d_r are nonnegative integers that denote the number of *free*, *bounded* and *restricted* species respectively. The finite set $\mathcal{F}_b \subset \mathbb{N}_0^{d_b}$ serves as a state-space for the dynamics of *bounded* species, while the infinite set $\Phi \subset \mathbb{N}_0^{d_f + d_r}$ serves as a state-space for the dynamics of both *free* species and *restricted* species. As the definition of Φ indicates, these *restricted* species are "locked" in a fixed affine relationship (given by function ϕ) with *free* species. Notice that in comparison to $\mathcal{E}_0$, the state-space $\mathcal{E}^\sigma$ has a simpler form which clearly expresses the copy-number ranges for each species as well as the relationships among them. This simple form is useful for checking irreducibility and explicitly computing the stationary distribution in certain cases (see Example 4.5). Observe that if $d_r = d_f = 0$ then all the species are *bounded* and the irreducibility of the finite set $\mathcal{E} = \mathcal{F}_b$ can be checked using matrix methods (see [22]), while on the other hand if $d_b = d_r = 0$ then all the species are *free* and the irreducibility of $\mathcal{E} = \mathbb{N}_0^{d_f}$ can be checked using birth-death cascade construction mentioned in Section 1 (see also [18]). In this paper we efficiently combine these two approaches to handle large reaction networks with many conservation relations. Note that if $\mathcal{E}^\sigma$ is an irreducible state-space for the permuted conservation network $(\mathcal{N}^\sigma, \Gamma_\sigma, c)$ then $\mathcal{E} = P_\sigma^T \mathcal{E}^\sigma$ is an irreducible state-space for the original conservation network $(\mathcal{N}, \Gamma, c)$ (see Proposition 2.6). Throughout this section we assume that network $\mathcal{N}$ satisfies Assumption 2.3.

Our approach for finding irreducible state-spaces can be described as follows. We first identify the *bounded* species, partition them into distinct *interaction* classes and find their optimal finite state-space $\mathcal{E}_b$ (Section 3.1). We then classify the rest of the species as *free* or *restricted*, and also determine the affine function ϕ (see (3.27)) which gives the static relationship between these two sets of species (Section 3.2). Since the *restricted* species have no dynamics of their own, we "remove" them in a certain way to obtain a simpler network whose dynamics is essentially equivalent to the original network (Section 3.3). Lastly we find the irreducible state-spaces for this "reduced" network by suitably combining the two strategies of arranging the *free* species into birth-death cascades and applying matrix methods to find the irreducible subsets within the finite state-space $\mathcal{E}_b$ for *bounded* species (Section 3.4). We also explain how the pre-partitioning of *bounded* species into interaction classes helps in gaining computational efficiency. The complete procedure for finding irreducible state-spaces is presented in Section 3.5.

3.1 Identifying the *bounded* species and their state-space

Recall that (Γ, c) is the *conservation data* for our network $\mathcal{N}$. Here Γ is a $d \times n$ matrix whose columns span $\mathcal{L}(S)$ (see Section 2.4.4). Therefore we can write $\mathcal{L}(S) = \{\Gamma\alpha : \alpha \in \mathbb{R}^n\}$. Suppose for some $\alpha \in \mathbb{R}^n$, $\gamma = \Gamma\alpha \in \mathcal{L}(S)$ is a conservation relation such that all its nonzero entries have the same sign. Replacing α with $-\alpha$ if necessary, we can assume that this sign is positive. If $(X(t))_{t \geq 0}$ is the Markov process representing the reaction dynamics under network $\mathcal{N}$, then relation (2.24) implies that

$$\langle \gamma, X(t) \rangle = \sum_{i=1}^{d} \gamma_i X_i(t) = \langle c, \alpha \rangle \qquad \text{for all } t \geq 0. \tag{3.28}$$

Hence for any $i \in \text{supp}(\gamma)$ (i.e. $\gamma_i > 0$), species i is *bounded* because throughout the dynamics its copy-number $X_i(t)$ is bounded above by $\langle c, \alpha \rangle / \gamma_i$.

Of course each *bounded* species may be involved in several positive conservation relations and this must be taken into account for computing a *sharp* upper-bound for its copy-numbers. This sharpness is necessary for finding irreducible state-spaces, because otherwise certain copy-numbers would be unreachable. For each species $i \in \mathcal{D}$, we find the optimal upper-bound b_i for its copy-number by solving the following Linear Programming Problem (LPP):

$$b_i = \min_{\alpha \in \mathbb{R}^n} \langle c, \alpha \rangle \tag{3.29}$$

$$\text{subject to} \qquad \Gamma\alpha \geq \overline{0}_d \quad \text{and} \quad \langle e_i^{(d)}, \Gamma\alpha \rangle = 1.$$

Note that the optimal value b_i is set to ∞ if the feasible region of this LPP is empty, and in this case species i is not bounded because i does not lie in the support of any positive conservation relation. From now on, we denote the set of *bounded* species by $\mathcal{D}_b = \{i \in \mathcal{D} : b_i < \infty\}$ and the set of *unbounded* species by its complement $\mathcal{D}_u = \mathcal{D}_b^c = \{i \in \mathcal{D} : b_i = \infty\}$. Let $d_b = |\mathcal{D}_b|$ and $d_u = |\mathcal{D}_u| = d - d_b$ be the cardinalities of these two sets.

In order to handle the dynamics of *bounded* species more efficiently we partition them according to their interaction relationships. Define an *interaction* relation $I_b \subset \mathcal{D}_b \times \mathcal{D}_b$ by

$$I_b = \{(i,j) \in \mathcal{D}_b \times \mathcal{D}_b : \{i,j\} \subset \mathrm{supp}(\nu_k) \cup \mathrm{supp}(\rho_k) \text{ for some } k = 1, \ldots, K\}. \tag{3.30}$$

Two *bounded* species $i, j \in \mathcal{D}_b$ are said to *interact* if $(i,j) \in I_b$ which occurs when both these species are *involved* in some reaction k, either as products or as reactants. Note that relation I_b is *reflexive* and *symmetric*. Hence its *transitive closure* I_b^* is an equivalence relation which partitions $\mathcal{D}_b$ into M disjoint equivalence classes denoted by $\mathcal{D}_{b,1}, \ldots, \mathcal{D}_{b,M}$. From now on we refer to $\mathcal{D}_{b,m}$ as the m-th *interaction class*.

For each $m = 1, \ldots, M$ let $d_{b,m} = |\mathcal{D}_{b,m}|$ and let $s_m = \sum_{j=1}^{m-1} d_{b,j}$. Choose a permutation map $\sigma_1 : \mathcal{D} \to \mathcal{D}$ satisfying

$$\sigma_1(l) \in \left\{ \begin{array}{ll} \mathcal{D}_{b,m} & \text{for each } l = s_m + 1, \ldots, s_{m+1} \text{ and } m = 1, \ldots, M \\ \mathcal{D}_u & \text{for } l = (d_b + 1), \ldots, d. \end{array} \right. \tag{3.31}$$

Consider the reaction dynamics of the permuted network $\mathcal{N}^{\sigma_1}$ under permutation σ_1 (see Section 2.4.3). Now the entries in rows $(s_m + 1), \ldots, s_{m+1}$ of the state vectors will contain the copy-numbers of *bounded* species in $\mathcal{D}_{b,m}$. These copy-numbers, arranged as vectors in $\mathbb{N}_0^{d_{b,m}}$, will always lie in the finite rectangular set

$$\mathcal{R}_{b,m}^{\sigma_1} = \{(x_1, \ldots, x_{d_{b,m}}) \in \mathbb{N}_0^{d_{b,m}} : x_l \leq b_{\sigma_1(s_m + l)} \quad \text{for each } l = 1, \ldots, d_{b,m}\}, \tag{3.32}$$

but all the elements in this set may not be reachable from each other due to conservation relations among *bounded* species (see Example 4.1 and 4.4). We deal with these conservations relations now.

Let S_{σ_1} be the stoichiometry matrix for network $\mathcal{N}^{\sigma_1}$ and let (Γ_{σ_1}, c) be its conservation data (see Remark 2.8). If $S_{\sigma_1}^{b,m} = \mathrm{Proj}(S_{\sigma_1}, s_m + 1, s_{m+1})$ is the $d_{b,m} \times K$ matrix consisting of rows $(s_m + 1), \ldots, s_{m+1}$ of S_{σ_1}, then conservation relations among the *bounded* species in the m-th interaction class $\mathcal{D}_{b,m}$ are given by nonzero vectors in its left nullspace $\mathcal{L}(S_{\sigma_1}^{b,m})$. Suppose $n_{b,m} = \dim\left(\mathcal{L}(S_{\sigma_1}^{b,m})\right) \geq 1$ and let $\{\widehat{\gamma}_1^{(m)}, \ldots, \widehat{\gamma}_{n_{b,m}}^{(m)}\}$ denote a *basis* for $\mathcal{L}(S_{\sigma_1}^{b,m})$. For each $j = 1, \ldots, n_{b,m}$, let $\overline{\gamma}_j^{(m)} = (\overline{0}_{s_m}, \widehat{\gamma}_j^{(m)}, \overline{0}_{d - s_{m+1}}) \in \mathbb{R}^d$ and set $\widehat{c}_j^{(m)} = \langle \alpha_j^{(m)}, c \rangle$ where $\alpha_j^{(m)} \in \mathbb{R}^n$ is the unique solution of the linear-system $\Gamma_{\sigma_1} \alpha_j^{(m)} = \overline{\gamma}_j^{(m)}$. Relation (2.24) implies that under permutation σ_1 the state vectors for the *bounded* species in $\mathcal{D}_{b,m}$ will always lie in the finite set

$$\mathcal{E}_{b,m}^{\sigma_1} = \{x \in \mathcal{R}_{b,m}^{\sigma_1} : \langle \widehat{\gamma}_j^{(m)}, x \rangle = \widehat{c}_j^{(m)} \quad \text{for each } j = 1, \ldots, n_{b,m}\}. \tag{3.33}$$

Therefore the state-vectors for all the *bounded* species will lie in the finite set

$$\mathcal{E}_b^{\sigma_1} = \mathcal{E}_{b,1}^{\sigma_1} \times \cdots \times \mathcal{E}_{b,M}^{\sigma_1}. \tag{3.34}$$

Such a *fragmented* structure of $\mathcal{E}_b^{\sigma_1}$ allows us to exploit the *sparsity* in the interaction relationships among the *bounded* species and obtain significant reductions in the computational effort required for finding irreducible state-spaces, as we shall see in Section 3.4.

3.2 Identifying the *free* and the *restricted* species

We now partition the set $\mathcal{D}_u$ of *unbounded* species, into a set $\mathcal{D}_f$ of *free* species and a set $\mathcal{D}_r$ of *restricted* species. Let $S_{\sigma_1}^u = \mathrm{Proj}(S_{\sigma_1}, d_b + 1, d) \in \mathbb{M}(d_u, K)$ and define the number of *free* species d_f and the number of *restricted* species d_r by

$$d_f = \mathrm{Rank}(S_{\sigma_1}^u) \quad \text{and} \quad d_r = \dim(\mathcal{L}(S_{\sigma_1}^u)). \tag{3.35}$$

13

From the fundamental theorem of linear algebra we have $d_f + d_r = d - d_b = d_u$, which also shows that the total number of species (d) is equal to the sum of the number of *free* (d_f), *bounded* (d_b) and *restricted* species (d_r).

Suppose $d_r \geq 1$ and so the left nullspace $\mathcal{L}(S^u_{\sigma_1})$ in nontrivial. The nonzero vectors in $\mathcal{L}(S^u_{\sigma_1})$ correspond to conservation relations among the *unbounded* species, and these relations will create *restricted* species as we shall soon see. Unlike the *bounded* species (see Section 3.1), these *restricted* species arise due to conservation relations whose components are *not all of the same sign*. Let $\{\delta'_1, \ldots, \delta'_{d_r}\}$ be a basis for $\mathcal{L}(S^u_{\sigma_1})$. For any subset $I = \{i_1, \ldots, i_{d_f}\} \subset \{1, \ldots, d_u\}$ with $|I| = d_f$ elements, let $A_I \in \mathbb{M}(d_u, d_u)$ be the matrix given by

$$A_I = \mathrm{Col}\left(e^{(d_u)}_{i_1}, \ldots, e^{(d_u)}_{i_{d_f}}, \delta'_1, \ldots, \delta'_{d_r}\right). \tag{3.36}$$

Define another set

$$\mathcal{I}_f = \{I \subset \{1, \ldots, d_u\} : |I| = d_f \quad \text{and} \quad \mathrm{Rank}(A_I) = d_u\}. \tag{3.37}$$

Note that this set is nonempty and its cardinality is bounded above by $\binom{d_u}{d_f}$.

Fix a $I \in \mathcal{I}_f$ and let I^c denote its complement in the set $\{1, \ldots, d_u\}$. We define the set $\mathcal{D}_f$ of *free* species and the set $\mathcal{D}_r$ of *restricted* species as $\mathcal{D}_f = \{\sigma_1(d_b + i) : i \in I\}$ and $\mathcal{D}_r = \{\sigma_1(d_b + i) : i \in I^c\}$. These two sets are certainly disjoint and $\mathcal{D}_u = \mathcal{D}_f \cup \mathcal{D}_r$ (see (3.31)). We now choose another permutation map $\sigma_2 : \mathcal{D} \to \mathcal{D}$ satisfying

$$\sigma_2(l) = \sigma_1(l) \quad \text{for } l = 1, \ldots, d_b \quad \text{and} \quad \sigma_2(l) \in \left\{ \begin{array}{ll} \mathcal{D}_f & \text{for } l = (d_b + 1), \ldots, (d_b + d_f) \\ \mathcal{D}_r & \text{for } l = (d_b + d_f + 1), \ldots, d. \end{array} \right. \tag{3.38}$$

Let $\mathcal{N}^{\sigma_2}$ be the network under permutation σ_2 (see Section 2.4.3), and let S_{σ_2} and (Γ_{σ_2}, c) be its stoichiometry matrix and conservation data respectively (see Remark 2.8). For each $i = 1, \ldots, d_r$, the vector $\overline{\delta}_i = (\overline{0}_{d_b}, \delta'_i) \in \mathbb{R}^d$ belongs to $\mathcal{L}(S_{\sigma_1})$ and hence the vector $\widehat{\delta}_i = P_{\sigma_2} P^T_{\sigma_1} \overline{\delta}_i$ belongs to $\mathcal{L}(S_{\sigma_2})$. Since the permutations σ_1 and σ_2 are identical on $\{1, \ldots, d_b\}$, each $\widehat{\delta}_i$ must have the form $\widehat{\delta}_i = (\overline{0}_{d_b}, \delta^{(1)}_i, \delta^{(2)}_i)$ for some vectors $\delta^{(1)}_i \in \mathbb{R}^{d_f}$ and $\delta^{(2)}_i \in \mathbb{R}^{d_r}$. Define matrices $\Delta_1 = \mathrm{Col}(\delta^{(1)}_1, \ldots, \delta^{(1)}_{d_r}) \in \mathbb{M}(d_f, d_r)$ and $\Delta_2 = \mathrm{Col}(\delta^{(2)}_1, \ldots, \delta^{(2)}_{d_r}) \in \mathbb{M}(d_r, d_r)$. Observe that if A_I is the matrix given by (3.36), then there exists a permutation matrix $P \in \mathbb{M}(d_u, d_u)$ such matrix PA_I has the form

$$PA_I = \left[\begin{array}{cc} I_{d_f} & \Delta_1 \\ 0 & \Delta_2 \end{array} \right], \tag{3.39}$$

where I_{d_f} is the $d_f \times d_f$ identity matrix and 0 is the $d_r \times d_f$ matrix of all zeroes. Matrix A_I is invertible because $I \in \mathcal{I}_f$, and hence matrix Δ_2 is also invertible.

From now on let $S^b_{\sigma_2} = \mathrm{Proj}(S_{\sigma_2}, 1, d_b)$, $S^u_{\sigma_2} = \mathrm{Proj}(S_{\sigma_2}, d_b + 1, d)$, $S^f_{\sigma_2} = \mathrm{Proj}(S^u_{\sigma_2}, 1, d_f)$ and $S^r_{\sigma_2} = \mathrm{Proj}(S^u_{\sigma_2}, d_f + 1, d_f + d_r)$. Since for each $i = 1, \ldots, d_r$, the vector $\widehat{\delta}_i = (\overline{0}_{d_b}, \delta^{(1)}_i, \delta^{(2)}_i)$ belongs to $\mathcal{L}(S_{\sigma_2})$ we must have $\Delta^T_1 S^f_{\sigma_2} = -\Delta^T_2 S^r_{\sigma_2} = 0$ which allows us to write

$$S^r_{\sigma_2} = -(\Delta^T_2)^{-1}\Delta^T_1 S^f_{\sigma_2} = -(\Delta_2^{-1})^T \Delta^T_1 S^f_{\sigma_2}. \tag{3.40}$$

This also shows that $\mathrm{Rank}(S^f_{\sigma_2}) = d_f$ because $d_f = \mathrm{Rank}(S^u_{\sigma_1}) = \mathrm{Rank}(S^u_{\sigma_2})$ (see (3.35)) and $S^u_{\sigma_2}$ has the form

$$S^u_{\sigma_2} = \left[\begin{array}{c} S^f_{\sigma_2} \\ S^r_{\sigma_2} \end{array} \right] = \left[\begin{array}{c} I_{d_f} \\ -(\Delta_2^{-1})^T\Delta^T_1 \end{array} \right] S^f_{\sigma_2}.$$

Let $(X^{\sigma_2}(t))_{t \geq 0}$ denote the stochastic reaction dynamics representing the conservation network $(\mathcal{N}^{\sigma_2}, \Gamma_{\sigma_2}, c)$. Then at any time t we can write

$$X^{\sigma_2}(t) = (X^{\sigma_2}_b(t), X^{\sigma_2}_f(t), X^{\sigma_2}_r(t)), \tag{3.41}$$

where $X^{\sigma_2}_b(t)$, $X^{\sigma_2}_f(t)$ and $X^{\sigma_2}_r(t)$ denote the state vectors for species in $\mathcal{D}_b$, $\mathcal{D}_f$ and $\mathcal{D}_r$ under permutation σ_2. Since the columns of Γ_{σ_2} span $\mathcal{L}(S_{\sigma_2})$, for each $i = 1, \ldots, d_r$, there exists a unique $\alpha_i \in \mathbb{R}^n$ such that

$\Gamma_{\sigma_2}\alpha_i = \widehat{\delta}_i$. Let $\widehat{c} = (\langle\alpha_1, c\rangle, \ldots, \langle\alpha_{d_r}, c\rangle) \in \mathbb{R}^{d_r}$. Relation (2.24) implies that for all $t \geq 0$ we must have $\Delta_1^T X_f^{\sigma_2}(t) + \Delta_2^T X_r^{\sigma_2}(t) = \Delta_1^T X_f^{\sigma_2}(0) + \Delta_2^T X_r^{\sigma_2}(0) = \widehat{c}$ which shows that

$$X_r^{\sigma_2}(t) = \phi(X_f^{\sigma_2}(t)), \tag{3.42}$$

where the *affine* map $\phi : \mathbb{R}^{d_f} \to \mathbb{R}^{d_r}$ is defined by

$$\phi(x) = (\Delta_2^T)^{-1}\widehat{c} - (\Delta_2^T)^{-1}\Delta_1^T x. \tag{3.43}$$

We already know that $X_b^{\sigma_2}(t)$ will lie in the set $\mathcal{E}_b^{\sigma_2} := \mathcal{E}_b^{\sigma_1}$ (given by (3.34)) for all $t \geq 0$. Therefore

$$\mathcal{E}_0^{\sigma_2} = \{(x_b, x_f, x_r) \in \mathbb{N}_0^d : x_b \in \mathcal{E}_b^{\sigma_2}, \quad x_f \in \mathbb{N}_0^{d_f} \quad \text{and} \quad x_r = \phi(x_f)\}, \tag{3.44}$$

serves as the *maximal* state-space for the conservation network $(\mathcal{N}^{\sigma_2}, \Gamma_{\sigma_2}, c)$. In other words, it includes every other state-space for this conservation network. We end this section with an important remark.

Remark 3.1 *Note that the classification of* unbounded *species into* free *and* restricted *species depends on the set I which can chosen to be any element in the set $\mathcal{I}_f$ given by* (3.37). *This flexibility in the choice of I can be quite useful for our purpose as we shall explain in the next section.*

3.3 Network reduction by elimination of *restricted* species

From the discussion in the previous section it is immediate that the *restricted* species have no independent dynamics of their own and they essentially *mimic* the *free* species according to some mapping ϕ. This suggests that for finding irreducible state-spaces we can simply remove the *restricted* species and concentrate on the dynamics of the *bounded* and the *free* species. In other words we can construct a "reduced" network $\widetilde{\mathcal{N}}^{\sigma_2}$ with $(d_b + d_f)$ species in the set $\mathcal{D}_b \cup \mathcal{D}_f$ where $\mathcal{D}_b = \{\sigma_2(1), \ldots, \sigma_2(d_b)\}$ and $\mathcal{D}_f = \{\sigma_2(d_b + 1), \ldots, \sigma_2(d_b + d_f)\}$.

We now describe the network $\widetilde{\mathcal{N}}^{\sigma_2}$ more formally. For each reaction $k = 1, \ldots, K$ define a function $\widetilde{\lambda}_k^{\sigma_2} : \mathbb{N}_0^{d_b} \times \mathbb{N}_0^{d_f} \to \mathbb{R}_+$ by

$$\widetilde{\lambda}_k^{\sigma_2}(x_b, x_f) = \lambda_k^{\sigma_2}(x_b, x_f, \phi(x_f)) \tag{3.45}$$

where $x_b \in \mathbb{N}_0^{d_b}, x_f \in \mathbb{N}_0^{d_f}$ and the function $\lambda_k^{\sigma_2}$ is defined by (2.19). Let the *propensity* map $\widetilde{\Lambda}^{\sigma_2} : \mathbb{N}_0^{d_b+d_f} \to \mathbb{R}_+^K$ be as in (2.2) with each λ_k replaced by $\widetilde{\lambda}_k^{\sigma_2}$. Define $\widetilde{\mathcal{V}}^{\sigma_2} = \text{Proj}(\mathcal{V}^{\sigma_2}, 1, d_b + d_f)$ and $\widetilde{\mathcal{O}}^{\sigma_2} = \text{Proj}(\mathcal{O}^{\sigma_2}, 1, d_b + d_f)$ to be the matrices containing the first $(d_b + d_f)$ rows of $\mathcal{V}^{\sigma_2}$ and $\mathcal{O}^{\sigma_2}$ respectively (see Section 2.4.3). We define the reduced network as $\widetilde{\mathcal{N}}^{\sigma_2} = (\widetilde{\mathcal{V}}^{\sigma_2}, \widetilde{\mathcal{O}}^{\sigma_2}, \widetilde{\Lambda}^{\sigma_2})$.

This reduced network may not satisfy Assumption 2.3 even if the original network does (see Example 4.2), which is a problem because our approach requires this assumption. Fortunately, one can deal with this problem by exploiting the *flexibility* in the choice of set I (see Remark 3.1) which determines the reduced network by classifying each *unbounded* species as *free* or *restricted*. Note that different choices of I will yield different reduced networks but they correspond to the same dynamics for the original network. Therefore irreducible state-spaces for the original network can be found with any choice of I and one can sequentially examine each element in the finite set $\mathcal{I}_f$ (see (3.37)) until one finds a I for which the reduced network satisfies Assumption 2.3. We illustrate this in Example 4.2.

Let $\widetilde{S}_{\sigma_2} = \widetilde{\mathcal{O}}^{\sigma_2} - \widetilde{\mathcal{V}}^{\sigma_2} \in \mathbb{M}(d_b + d_f, K)$ be the stoichiometry matrix for network $\widetilde{\mathcal{N}}^{\sigma_2}$ and let $\widetilde{S}_{\sigma_2}^b = \text{Proj}(\widetilde{S}_{\sigma_2}, 1, d_b)$ and $\widetilde{S}_{\sigma_2}^f = \text{Proj}(\widetilde{S}_{\sigma_2}, d_b + 1, d_b + d_f)$. Observe that $\widetilde{S}_{\sigma_2}^b = S_{\sigma_2}^b$ and $\widetilde{S}_{\sigma_2}^f = S_{\sigma_2}^f$, where the matrices $S_{\sigma_2}^b$ and $S_{\sigma_2}^f$ were defined in Section 3.2. Hence $\text{Rank}(\widetilde{S}_{\sigma_2}^f) = d_f$ and so there are no conservation relations among *free* species. This shows that $\mathcal{D}_b$ and $\mathcal{D}_f$ are exactly the sets of *bounded* and *free* species for network $\widetilde{\mathcal{N}}^{\sigma_2}$, and this network does not contain any *restricted* species.

Recall that permutations σ_1 and σ_2 agree on the set $\{1, \ldots, d_b\}$ (see (3.38)) and hence $\mathcal{L}(S_{\sigma_2}^b) = \mathcal{L}(S_{\sigma_1}^b)$ and this left nullspace is spanned by the set of vectors $\{\widehat{\gamma}_1, \ldots, \widehat{\gamma}_{n_b}\}$ (see Section 3.1). For each $j = 1, \ldots, n_b$ define $\widetilde{\gamma}_j = (\widehat{\gamma}_j, \overline{0}_{d_f})$ and $\widetilde{c}_j = \widehat{c}_j$ (see Section 3.1). This discussion shows that the conservation data for network $\widetilde{\mathcal{N}}^{\sigma_2}$ is simply $(\widetilde{\Gamma}_{\sigma_2}, \widetilde{c})$ where $\widetilde{c} = (\widetilde{c}_1, \ldots, \widetilde{c}_{n_b})$ and $\widetilde{\Gamma}_{\sigma_2} = \text{Col}(\widetilde{\gamma}_1, \ldots, \widetilde{\gamma}_{n_b})$. These conservation relations are automatically satisfied if the dynamics of the *bounded* species lies in the set $\mathcal{E}_b^{\sigma_2} := \mathcal{E}_b^{\sigma_1}$ defined

15

by (3.33). Therefore the set $\widetilde{\mathcal{E}}_0^{\sigma_2} = \mathcal{E}_b^{\sigma_2} \times \mathbb{N}_0^{d_f}$ is a *maximal* state-space for network $\widetilde{\mathcal{N}}^{\sigma_2}$. We now come to the main result of this section which allows us to use the reduced network $\widetilde{\mathcal{N}}^{\sigma_2}$ to find the irreducible state-spaces for the actual network.

Proposition 3.2 *For any $A_1 \subset \mathcal{E}_b^{\sigma_2}$ and $A_2 \subset \mathbb{N}_0^{d_f}$, the set $\widetilde{\mathcal{E}}_1^{\sigma_2} = A_1 \times A_2$ is an irreducible state-space for conservation network $(\widetilde{\mathcal{N}}^{\sigma_2}, \widetilde{\Gamma}_{\sigma_2}, \widetilde{c})$ if and only if the set*

$$\mathcal{E}_1^{\sigma_2} = \{(x_b, x_f, x_r) \in \mathcal{E}_0^{\sigma_2} : x_b \in A_1, \quad x_f \in A_2 \quad \text{and} \quad x_r = \phi(x_f)\}$$

is an irreducible state-space for conservation network $(\mathcal{N}^{\sigma_2}, \Gamma_{\sigma_2}, c)$.

Proof. The proof follows simply from the construction of the reduced network $\widetilde{\mathcal{N}}^{\sigma_2}$ and the fact that the dynamics of the *restricted* species is "tied" to the dynamics of the *free* species according to map ϕ (see (3.42)). $\qquad\square$

3.4 Networks with only *free* and *bounded* species

We now assume that we have a "reduced" network $\widetilde{\mathcal{N}}^{\sigma_2} = (\widetilde{\mathcal{V}}^{\sigma_2}, \widetilde{\mathcal{O}}^{\sigma_2}, \widetilde{\Lambda}^{\sigma_2})$ without any *restricted* species which satisfies Assumption 2.3 with $(d_b + d_f)$ species that are in the set $\mathcal{D}_b \cup \mathcal{D}_f$. The conservation data for this network is $(\widetilde{\Gamma}_{\sigma_2}, \widetilde{c})$ and the associated conservation relations are automatically satisfied if the dynamics lies in the set $\widetilde{\mathcal{E}}_0^{\sigma_2} = \mathcal{E}_b^{\sigma_2} \times \mathbb{N}_0^{d_f}$ which is a *maximal* state-space for the network. Therefore all its irreducible state-spaces must necessarily belong to this set $\widetilde{\mathcal{E}}_0^{\sigma_2}$. In the rest of this section, network $\widetilde{\mathcal{N}}^{\sigma_2}$ always refers to the conservation network $(\widetilde{\mathcal{N}}^{\sigma_2}, \widetilde{\Gamma}_{\sigma_2}, \widetilde{c})$.

Our first task is identify the candidate sets $\mathcal{F}_b^{\sigma_2} \subset \mathcal{E}_b^{\sigma_2}$ for which the set $\mathcal{F}_b^{\sigma_2} \times \mathbb{N}_0^{d_f}$ has the possibility of being an irreducible state-space for network $\widetilde{\mathcal{N}}^{\sigma_2}$. Recall the partitioning of *bounded* species into M distinct interaction classes $\mathcal{D}_{b,1}, \ldots, \mathcal{D}_{b,m}$ and the form of the finite set $\mathcal{E}_b^{\sigma_2} = \mathcal{E}_{b,1}^{\sigma_1} \times \cdots \times \mathcal{E}_{b,M}^{\sigma_1}$ from Section 3.1. Let $\mathcal{E}_{b,m}^{\sigma_2} = \mathcal{E}_{b,m}^{\sigma_1}$ for each $m = 1, \ldots, M$. Since the *bounded* species in different interaction classes do not interact, for $\mathcal{F}_b^{\sigma_2} \times \mathbb{N}_0^{d_f}$ to be an irreducible state-space for network $\widetilde{\mathcal{N}}^{\sigma_2}$, the set $\mathcal{F}_b^{\sigma_2}$ must have the form

$$\mathcal{F}_b^{\sigma_2} = \mathcal{F}_{b,1}^{\sigma_2} \times \mathcal{F}_{b,2}^{\sigma_2} \times \cdots \times \mathcal{F}_{b,M}^{\sigma_2}, \tag{3.46}$$

where $\mathcal{F}_{b,m}^{\sigma_2} \subset \mathcal{E}_{b,m}^{\sigma_2}$ for each m. We now find the "right" $\mathcal{F}_{b,1}^{\sigma_2}, \ldots, \mathcal{F}_{b,M}^{\sigma_2}$ for which the irreducibility of $\mathcal{F}_b^{\sigma_2} \times \mathbb{N}_0^{d_f}$ is plausible.

For each m, let $d_{b,m}$ and s_m be as in Section 3.1. For each $k = 1, \ldots, K$, let $\overline{\nu}_{k,m}^{\sigma_2}, \overline{\rho}_{k,m}^{\sigma_2}, \widehat{\nu}_k^{\sigma_2}$ and $\widehat{\rho}_k^{\sigma_2}$ denote the k-th column of matrices $\mathrm{Proj}(\widetilde{\mathcal{V}}^{\sigma_2}, s_m+1, s_m+d_{b,m})$, $\mathrm{Proj}(\widetilde{\mathcal{O}}^{\sigma_2}, s_m+1, s_m+d_{b,m})$, $\mathrm{Proj}(\widetilde{\mathcal{V}}^{\sigma_2}, d_b+1, d_b+d_f)$ and $\mathrm{Proj}(\widetilde{\mathcal{O}}^{\sigma_2}, d_b+1, d_b+d_f)$ respectively. Note that vectors $\overline{\nu}_{k,m}^{\sigma_2}$ and $\overline{\rho}_{k,m}^{\sigma_2}$ pertain to the *bounded* species in the m-th interaction class $\mathcal{D}_{b,m}$ while the vectors $\widehat{\nu}_k^{\sigma_2}$ and $\widehat{\rho}_k^{\sigma_2}$ pertain to the *free* species. Let $F = \{1, \ldots, d_f\}$ and for any $A \subset F$ define a set of *free* species as $\mathcal{D}_f^{\sigma_2}(A) = \{\sigma_2(d_b + i) : i \in A\} \subset \mathcal{D}_f$. From now on, when we say that the *free* species in $\mathcal{D}_f^{\sigma_2}(A)$ are *abundantly available* we mean that the copy-number of each species $\sigma_2(d_b + i) \in \mathcal{D}_f^{\sigma_2}(A)$ is greater than $\max\{\widehat{\nu}_{ik}^{\sigma_2} : k = 1, \ldots, K\}$, which is the maximum number of molecules of this species that are consumed by any reaction. For any $A \subset F$ and any $y \in \mathcal{E}_{b,m}^{\sigma_2}$ let

$$\mathcal{K}_m(y, A) = \{k = 1, \ldots, K : \mathrm{supp}(\widehat{\nu}_k^{\sigma_2}) \subset A \text{ and } y \geq \overline{\nu}_{k,m}^{\sigma_2}\}, \tag{3.47}$$

be the set of reactions that have a positive probability of firing when the *free* species in $\mathcal{D}_f^{\sigma_2}(A)$ are abundantly available and when the dynamics of *bounded* species in the m-th interaction class $\mathcal{D}_{b,m}$ is at state y. Borrowing ideas from the theory of finite Markov chains [22], we can study reachability relationships in $\mathcal{E}_{b,m}^{\sigma_2}$ under these reactions, by forming a *zero-pattern* matrix[3] and computing the communication classes of the induced equivalence relation. Suppose $N_{b,m} = |\mathcal{E}_{b,m}^{\sigma_2}|$ and $\mathcal{E}_{b,m}^{\sigma_2} = \{y_{1,m}, \ldots, y_{N_b,m}\}$. Define the zero-pattern matrix $Z_m(A) \in \mathbb{M}(N_{b,m}, N_{b,m})$ by

$$Z_{ij,m}(A) = \begin{cases} 1 & \text{if } y_{j,m} = y_{i,m} + \overline{\rho}_{k,m}^{\sigma_2} - \overline{\nu}_{k,m}^{\sigma_2} \text{ for some } k \in \mathcal{K}_m(y_{i,m}, A) \\ 0 & \text{otherwise.} \end{cases} \tag{3.48}$$

---

[3]The *zero-pattern* matrix corresponding to a finite Markov chain is obtained by setting all the positive entries in its probability transition matrix to 1 and all the rest to 0

For any $y_{i,m}, y_{j,m} \in \mathcal{E}_{b,m}^{\sigma_2}$, state $y_{i,m}$ is reachable from state $y_{j,m}$ if and only if $\Omega_{ij,m}(x_f) > 0$, where $\Omega_{ij,m}(x_f)$ is the ij-th entry of the *reachability* matrix $\Omega_m(x_f) \in \mathbb{M}(N_{b,m}, N_{b,m})$ defined by

$$\Omega_m(A) = \left(I_{N_{b,m}} + Z_m(A) \right)^{N_{b,m}-1}. \tag{3.49}$$

Based on this matrix we can define an *equivalence* relation $\Theta_m(A)$ on $\mathcal{E}_{b,m}^{\sigma_2}$ as

$$\Theta_m(A) = \{ (y_{i,m}, y_{j,m}) \in \mathcal{E}_{b,m}^{\sigma_2} \times \mathcal{E}_{b,m}^{\sigma_2} : \Omega_{ij,m}(A) > 0 \quad \text{and} \quad \Omega_{ji,m}(A) > 0 \}. \tag{3.50}$$

The equivalence classes for this relation are known as *communication classes* in the Markov chain literature. A communication class C is called *closed* if for any $y_{i,m} \in C$ and $y_{j,m} \in \mathcal{E}_{b,m}^{\sigma_2}$, $Z_{ij,m}(A) = 1$ implies that $y_{j,m} \in C$. Let $\mathcal{C}_m(A)$ be the set of all closed communication classes for relation $\Theta_m(A)$. This set can be easily computed by performing simple operations on the matrix $\Omega_m(A)$ (see [22]). When the *free* species in $\mathcal{D}_f^{\sigma_2}(A)$ are abundantly available, the dynamics of the *bounded* species in the m-th interaction class $\mathcal{D}_{b,m}$ will eventually get *trapped* in one of the closed equivalence classes in $\mathcal{C}_m(A)$.

Observe that for $A_1 \subset A_2$ we have $\mathcal{K}_m(y, A_1) \subset \mathcal{K}_m(y, A_2)$ for any $y \in \mathcal{E}_{b,m}^{\sigma_2}$. Hence as more *free* species become abundantly available, more transition channels open up for dynamics of *bounded* species. This shows that if $A_1 \subset A_2$ then $|\mathcal{C}_m(A_2)| \leq |\mathcal{C}_m(A_1)|$ and the number of closed equivalence classes ($|\mathcal{C}_m(A)|$) is least when all the *free* species are abundantly available ($A = F$). Choose a set $\mathcal{F}_{b,m}^{\sigma_2} \in \mathcal{C}_m(F)$ for each $m = 1, \ldots, M$ and let $\mathcal{F}_b^{\sigma_2}$ be given by (3.46). Then $\mathcal{F}_b^{\sigma_2}$ is certainly a state-space for network $\widetilde{\mathcal{N}}^{\sigma_2}$ (recall Definition 2.1), because for any $y = (y_1, \ldots, y_M) \in \mathcal{F}_b^{\sigma_2}$, with $y_m \in \mathcal{F}_{b,m}^{\sigma_2}$ for each m, the set of possible transitions is always a subset of $\cap_{m=1}^M \mathcal{K}(y_m, F)$ and the finite set $\mathcal{F}_b^{\sigma_2}$ is "closed" w.r.t. these transitions. In fact only such state-spaces can be irreducible if all the *free* species are abundant. This motivates our definition of candidate state-spaces whose irreducibility will be subsequently checked.

Definition 3.3 *Let $\mathcal{C}(F)$ be the set given by*

$$\mathcal{C}(F) = \{ C_1 \times C_2 \times \cdots \times C_M : C_m \in \mathcal{C}_m(F) \text{ for each } m = 1, \ldots, M \}.$$

For each $\mathcal{F}_b^{\sigma_2} \in \mathcal{C}(F)$, the set $\mathcal{F}_b^{\sigma_2} \times \mathbb{N}_0^{d_f}$ is called a Candidate Irreducible State-Space **(CISS)** *for network $\widetilde{\mathcal{N}}^{\sigma_2}$.*

Note that each CISS is certainly an irreducible state-space when $d_f = 0$. In the rest of the section, we consider the situation when both *bounded* and *free* species are present, and develop a procedure to check the irreducibility of each CISS. Fix a CISS $\mathcal{F}_b^{\sigma_2} \times \mathbb{N}_0^{d_f}$ for some $\mathcal{F}_b^{\sigma_2} = \mathcal{F}_{b,1}^{\sigma_2} \times \mathcal{F}_{b,2}^{\sigma_2} \times \cdots \times \mathcal{F}_{b,M}^{\sigma_2}$. We now discuss how the irreducibility of this CISS can be verified by arranging the *free* species into *birth and death cascades*. Let A be any subset of $F = \{1, \ldots, d_f\}$. Since $\mathcal{F}_b^{\sigma_2} \in \mathcal{C}(F)$, for any $C \in \mathcal{C}_m(A)$ we either have $C \cap \mathcal{F}_{b,m}^{\sigma_2} = \emptyset$ or $C \subset \mathcal{F}_{b,m}^{\sigma_2}$. Let $\mathcal{C}_m(A, \mathcal{F}_{b,m}^{\sigma_2})$ be the set of those classes in $\mathcal{C}_m(A)$ which are contained in $\mathcal{F}_{b,m}^{\sigma_2}$. This set of closed communication classes can be directly computed from the equivalence relation $\Theta_m(A)$ constructed as above by replacing $\mathcal{E}_{b,m}^{\sigma_2}$ with $\mathcal{F}_{b,m}^{\sigma_2}$. For any $C \in \mathcal{C}_m(A, \mathcal{F}_{b,m}^{\sigma_2})$ let

$$\mathcal{K}_m(C, A) = \bigcup_{y \in C} \mathcal{K}_m(y, A),$$

be the set of reactions that have the possibility of firing when the *free* species in $\mathcal{D}_f^{\sigma_2}(A)$ are abundantly available and when the dynamics of *bounded* species in the m-th interaction class $\mathcal{D}_{b,m}$ is inside C. For each $l = 1, 2, \ldots$ and $m = 1, \ldots, M$ we define a subset of F by

$$G_l(m) = \Big\{ i \in F : i \notin H_l \text{ and for each } C \in \mathcal{C}_m(H_l, \mathcal{F}_{b,m}^{\sigma_2}) \tag{3.51}$$

$$\text{there exists a } k \in \mathcal{K}_m(C, H_l) \text{ such that } i \in \mathrm{supp}(\widehat{\rho}_k^{\sigma_2}) \Big\},$$

where

$$H_l = \bigcup_{j=1}^{l-1} G_j \quad \text{and} \quad G_l = \bigcup_{m=1}^M G_l(m). \tag{3.52}$$

Let $B_l(m) = \mathcal{D}_f^{\sigma_2}(G_l(m)) = \{\sigma_2(d_b + i) : i \in G_l(m)\}$. We call the set $B_l = \mathcal{D}_f^{\sigma_2}(G_l) = \cup_m B_l(m)$ of *free* species as the l-th birth cascade for network $\widetilde{\mathcal{N}}^{\sigma_2}$. The set B_l contains all those *free* species that do not belong to any of the previous birth-cascades $B_1, \ldots, B_{l-1}$ and for some $m = 1, \ldots, M$, get produced by a reaction in $\mathcal{K}_m(C, H_l)$ for each closed communication class $C \in \mathcal{C}_m(H_l, \mathcal{F}_{b,m}^{\sigma_2})$. Note that any such reaction can only *consume* the *free* species in the previous birth-cascades $B_1, \ldots, B_{l-1}$.

The network $\widetilde{\mathcal{N}}^{\sigma_2}$ is called *birth-exhaustive* for CISS $\mathcal{F}_b^{\sigma_2} \times \mathbb{N}_0^{d_f}$ if for some positive integer l_0 we have

$$B_{l_0} = \mathcal{D}_f, \tag{3.53}$$

which means that all the *free* species can be arranged into birth-cascades constructed according to the scheme described above. As Lemma 4.6 shows, this birth-exhaustivity implies that the reaction dynamics can lead to arbitrarily high copy-numbers for all the *free* species. Let $\widetilde{\mathcal{N}}_{\text{inv}}^{\sigma_2}$ be the inverse of network $\widetilde{\mathcal{N}}^{\sigma_2}$ (see Section 2.4.2). Note that the stoichiometry matrix for $\widetilde{\mathcal{N}}_{\text{inv}}^{\sigma_2}$ is just the negative of the stoichiometry matrix for network $\widetilde{\mathcal{N}}^{\sigma_2}$. Using this one can show that $\mathcal{F}_b^{\sigma_2} \times \mathbb{N}_0^{d_f}$ is also a CISS for this inverse network $\widetilde{\mathcal{N}}_{\text{inv}}^{\sigma_2}$. We defne the l-th *death-cascade* D_l for network $\widetilde{\mathcal{N}}^{\sigma_2}$ as the l-th birth-cascade for the inverse network $\widetilde{\mathcal{N}}_{\text{inv}}^{\sigma_2}$. We say that network $\widetilde{\mathcal{N}}^{\sigma_2}$ is *death-exhaustive* for CISS $\mathcal{F}_b^{\sigma_2} \times \mathbb{N}_0^{d_f}$ if and only if the inverse network $\widetilde{\mathcal{N}}_{\text{inv}}^{\sigma_2}$ is *birth-exhaustive* for CISS $\mathcal{F}_b^{\sigma_2} \times \mathbb{N}_0^{d_f}$.

A necessary condition for $\mathcal{F}_b^{\sigma_2} \times \mathbb{N}_0^{d_f}$ to be irreducible is that the dynamics of the *free* species can reach any state in $\mathbb{N}_0^{d_f}$ from any other state in $\mathbb{N}_0^{d_f}$ through a combination of displacements in directions $(\widehat{\rho}_k^{\sigma_2} - \widehat{\nu}_k^{\sigma_2})$ for $k = 1, \ldots, K$. This condition can be formulated as

$$\text{Colspan}_{\mathbb{N}_0}(\widetilde{S}_{\sigma_2}^f) = \mathbb{Z}^{d_f}, \tag{3.54}$$

where $\widetilde{S}_{\sigma_2}^f \in \mathbb{M}(d_f, K)$ is the matrix defined in Section 3.3. Note that for (3.54) to hold, it is necessary that $\text{Rank}(\widetilde{S}_{\sigma_2}^f) = d_f$ which is indeed the case (see Section 3.3). When the network is birth-exhaustive, (3.54) follows from another condition which can be more easily verified for biochemical networks. This condition can be stated as

$$-\mathbb{N}_0^{d_f} \subset \text{Colspan}_{\mathbb{N}_0}(\widehat{S}_{\sigma_2}), \tag{3.55}$$

where $\widehat{S}_{\sigma_2}$ is the matrix consisting of only those columns of matrix $\widetilde{S}_{\sigma_2}^f$ which correspond to reactions that *exclusively involve* the *free* species. The set of such reactions is given by

$$\mathcal{K}_f^0 = \{k = 1, \ldots, K : \quad \text{supp}(\overline{\nu}_{k,m}^{\sigma_2}) = \emptyset \quad \text{and} \quad \text{supp}(\overline{\rho}_{k,m}^{\sigma_2}) = \emptyset \quad \text{for each} \quad m = 1, \ldots, M\}.$$

Condition (3.55) is equivalent to having $-e_i^{(d_f)} \subset \text{Colspan}_{\mathbb{N}_0}(\widehat{S}_{\sigma_2})$ for each $i = 1, \ldots, d_f$. This can be easily checked by computing a modified Hermite normal form[4] (see [8]) of the integer matrix $-\widehat{S}_{\sigma_2}^T$. For obtaining this normal form, the only admissible operation is the addition of one row to another, possibly after multiplication by a positive integer. If the Hermite normal form is such that for each $i = 1, \ldots, d_f$ there exists a row with the leading entry of 1 at column i, then we can conclude that (3.55) holds. Generally for biochemical networks this condition is trivially satisfied because all the *free* species degrade naturally due to reactions of the form $\mathbf{S}_i \longrightarrow \emptyset$. Moreover even for networks where this does not happen, this condition mostly holds because each *free* species usually undergoes a sequence of conversions, eventually resulting in a set of naturally degrading species. We now come to the main result of this paper.

Theorem 3.4 *Suppose that* (3.55) *holds and network $\widetilde{\mathcal{N}}^{\sigma_2}$ is both birth and death exhaustive for CISS $\mathcal{F}_b^{\sigma_2} \times \mathbb{N}_0^{d_f}$. Then $\mathcal{F}_b^{\sigma_2} \times \mathbb{N}_0^{d_f}$ is an irreducible state-space for network $\widetilde{\mathcal{N}}^{\sigma_2}$.*

Proof. Throughout this proof we denote the relation $\xrightarrow{\widetilde{\mathcal{N}}^{\sigma_2}}$ by $\longrightarrow$ and $\mathcal{F}_b^{\sigma_2}$ by $\mathcal{F}$. Without loss of generality we can assume that there is only one interaction class ($M = 1$). To prove the theorem it suffices to show that for any $z_1, z_2 \in \mathcal{F}$ and $x \in \mathbb{N}_0^{d_f}$ we have

$$(z_1, \overline{0}_{d_f}) \longrightarrow (z_2, x) \longrightarrow (z_1, \overline{0}_{d_f}). \tag{3.56}$$

---

[4] The Hermite normal form is an analogue of the row-reduced echelon form for integer matrices.

If this holds then $\mathcal{F} \times \mathbb{N}_0^{d_f}$ is irreducible because for any $(z_1, x_1), (z_2, x_2) \in \mathcal{F} \times \mathbb{N}_0^{d_f}$ we have $(z_1, x_1) \longrightarrow (z_2, x_2)$. This is due to the following chain of accessibility relations $(z_1, x_1) \longrightarrow (z_1, \overline{0}_{d_f}) \longrightarrow (z_2, x_2)$ and the fact that relation $\longrightarrow$ is transitive.

We now prove (3.56) for any fixed $z_1, z_2 \in \mathcal{F}$ and $x \in \mathbb{N}_0^{d_f}$. Since network $\widetilde{\mathcal{N}}^{\sigma_2}$ is death exhaustive, Lemma 4.6 and relation (2.17) imply that for any $z_1, z_2 \in \mathcal{F}$ and $r_0 \in \mathbb{N}^{d_f}$ there exists a vector $x' \in \mathbb{N}^{d_f}$ satisfying $x' \geq r_0$ and $(z_2, x') \longrightarrow (z_2, \overline{0}_{d_f})$. We assume that r_0 is as in Lemma 4.7. Relation $(z_2, x') \longrightarrow (z_2, \overline{0}_{d_f})$ also implies $(z_2, x + x') \longrightarrow (z_2, x)$ due to Proposition 2.4. From Lemma 4.6 we can find a vector $y \geq (x + x')$ satisfying $(z_1, \overline{0}_{d_f}) \longrightarrow (z_2, y)$. Since $(x + x') \geq r_0$, using Lemma 4.7 we obtain $(z_2, y) \longrightarrow (z_2, x + x')$. We now have an accessibility chain $(z_1, \overline{0}_{d_f}) \longrightarrow (z_2, y) \longrightarrow (z_2, x + x') \longrightarrow (z_2, x)$, and since $\longrightarrow$ is transitive we have the first part of (3.56). We now show the second part $(z_2, x) \longrightarrow (z_1, \overline{0}_{d_f})$. Using Lemma 4.6 and death-exhaustivity of network $\widetilde{\mathcal{N}}^{\sigma_2}$, we can find vectors $x_1, x_2 \geq r_0$ such that $(z_2, x_2) \longrightarrow (z_1, \overline{0}_{d_f})$, $x_1 \geq x_2$ and $(z_2, \overline{0}_{d_f}) \longrightarrow (z_2, x_1)$. The last relation also implies $(z_2, x) \longrightarrow (z_2, x_1 + x)$ due to Proposition 2.4. Since $(x_1 + x) \geq x_2 \geq r_0$ we have $(z_2, x_1 + x) \longrightarrow (z_2, x_2)$ due to Lemma 4.7. This gives us the following chain of accessibility relations $(z_2, x) \longrightarrow (z_2, x_1 + x) \longrightarrow (z_2, x_2) \longrightarrow (z_1, \overline{0}_{d_f})$ which shows the second part of (3.56) and completes the proof of this theorem. $\qquad\square$

Remark 3.5 *Suppose using Theorem 3.4 we can show that $\mathcal{F}_b^{\sigma_2} \times \mathbb{N}_0^{d_f}$ is an irreducible state-space for network $\widetilde{\mathcal{N}}^{\sigma_2}$. Then due to Proposition 3.2, the corresponding irreducible state-space for the conservation network $(\mathcal{N}^{\sigma_2}, \Gamma_{\sigma_2}, c)$ is $\mathcal{F}_b^{\sigma_2} \times \Phi$, where Φ is the graph (see (3.27)) of the affine mapping ϕ used in reducing the network (see (3.43)). Furthermore, the corresponding irreducible state-space for the original conservation network $(\mathcal{N}, \Gamma, c)$ can be found using Proposition 2.6.*

Using Theorem 3.4 we can determine if a particular CISS $\mathcal{F}_b^{\sigma_2} \times \mathbb{N}_0^{d_f}$ is an irreducible state-space for the reduced network $\widetilde{\mathcal{N}}^{\sigma_2}$. If $\eta = |\mathcal{C}(F)|$ then there are η possible choices for CISS (see Definition 3.3). If all these CISSs are irreducible state-spaces for network $\widetilde{\mathcal{N}}^{\sigma_2}$ then it is possible to find *all the irreducible state-spaces* for the original conservation network $(\mathcal{N}, \Gamma, c)$ as the next remark indicates. Recall that $\mathcal{E}_b^{\sigma_2} \times \mathbb{N}_0^{d_f}$ is the *maximal* state-space for network $\widetilde{\mathcal{N}}^{\sigma_2}$.

Remark 3.6 *Let $\mathcal{C}(F) = \{\mathcal{F}_{b,1}^{\sigma_2}, \mathcal{F}_{b,2}^{\sigma_2}, \ldots, \mathcal{F}_{b,\eta}^{\sigma_2}\}$. Then there are η possible choices for CISS given by $\mathcal{F}_{b,1}^{\sigma_2} \times \mathbb{N}_0^{d_f}, \mathcal{F}_{b,2}^{\sigma_2} \times \mathbb{N}_0^{d_f}, \ldots, \mathcal{F}_{b,\eta}^{\sigma_2} \times \mathbb{N}_0^{d_f}$. Suppose using Theorem 3.4 we can prove that all these CISSs are irreducible state-spaces for network $\widetilde{\mathcal{N}}^{\sigma_2}$. If*

$$\mathcal{E}_b^{\sigma_2} = \bigcup_{i=1}^{\eta} \mathcal{F}_{b,i}^{\sigma_2} \tag{3.57}$$

holds then these CISSs are the only irreducible state-spaces for the reduced network $\widetilde{\mathcal{N}}^{\sigma_2}$. This is because any other irreducible state-space must be disjoint from all these CISSs (see Remark 2.2) which cannot happen since $\mathcal{E}_b^{\sigma_2} \times \mathbb{N}_0^{d_f}$ is the maximal state-space for this network. Similar conclusions can be drawn regarding the irreducible state-spaces for conservation networks $(\mathcal{N}^{\sigma_2}, \Gamma_{\sigma_2}, c)$ and $(\mathcal{N}, \Gamma, c)$ (see Remark 3.5).

We now explain how the pre-partitioning of *bounded* species into interaction classes helps us in saving computational effort. This partitioning is not needed for computing the CISSs and the birth/death cascades, but without it, one would have to consider the dynamics of *bounded* species in a finite state-space of size $N_b = |\mathcal{E}_b^{\sigma_2}| = \prod_{m=1}^{M} N_{b,m}$. Thus for finding the closed communication classes, we would need to compute the $N_b \times N_b$ reachability matrix (see (3.49)) which requires roughly $N_b^3 \log_2 N_b$ computational steps, assuming that matrix powers are computed by repeated squaring [9]. However with the partitioning, the same can be achieved in $\sum_{m=1}^{M} N_{b,m}^3 \log_2(N_{b,m})$ steps which is always smaller than $N_b^3 \log_2 N_b$, and the difference can be large when several independently interacting *bounded* species (like gene-transcripts) are present. Since the computation of closed communication classes is needed for first determining the CISSs and then at each step of the cascade construction, the overall computational savings realized by initially partitioning the *bounded* species can be enormous. Note that this extra partitioning step only requires roughly $d_b K$ computational steps, which is insignificant in comparison to the whole procedure.

19

3.5 Main Algorithm

In this section we describe our procedure to find irreducible state-spaces in detail. Consider a network $\mathcal{N} = (\mathcal{V}, \mathcal{O}, \Lambda)$ with K reactions and d species in the set $\mathcal{D} = \{1, \ldots, d\}$. Assume that the conservation data for this network is (Γ, c). Given the conservation network $(\mathcal{N}, \Gamma, c)$ as input, the method FINDIRREDUCIBLESTATESPACES$(\cdot)$ (see Algorithm 1) uses our approach to find as many irreducible state-spaces as possible. This method starts by identifying the *bounded* species, classifying them according to their interaction relationships and finding their optimal state-space (see Section 3.1). It then computes the numbers of *free* (d_f) and *restricted* (d_r) species, and if $d_r > 0$, it reduces the network by eliminating the *restricted* species (see Sections 3.2 and 3.3). This reduction step is performed by calling the method FINDREDUCEDNETWORK$(\cdot)$ (see Algorithm 2). Finally all the CISSs for the reduced network are identified and for each CISS $\mathcal{F}_b^{\sigma_2} \times \mathbb{N}_0^{d_f}$, birth and death exhaustivity is checked (see Section 3.4) using the method CHECKBIRTHEXHAUSTIVITY$(\cdot)$ (see Algorithm 3). If both these exhaustivity conditions are met, then we can apply Theorem 3.4 to conclude that $\mathcal{F}_b^{\sigma_2} \times \mathbb{N}_0^{d_f}$ is an irreducible state-space for the reduced network. Hence the corresponding irreducible state-space for the original conservation network can be easily identified (see Remark 3.5).

Algorithm 1 Finds irreducible state-spaces for a network $\mathcal{N} = (\mathcal{V}, \mathcal{O}, \Lambda)$ with conservation data (Γ, c)

Require: Network $\mathcal{N}$ satisfies Assumption 2.3

1: **function** FINDIRREDUCIBLESTATESPACES$(\mathcal{N}, \Gamma, c)$
2:      For each species $i \in \mathcal{D}$ solve the LPP (3.29) to compute b_i.
3:      Set $\mathcal{D}_b = \{i \in \mathcal{D} : b_i < \infty\}$ and $\mathcal{D}_u = \{i \in \mathcal{D} : b_i = \infty\}$. Also set $d_b = |\mathcal{D}_b|$ and $d_u = |\mathcal{D}_u|$.
4:      Compute the interaction relation I_b given by (3.30) along with its transitive closure I_b^*.
5:      Partition $\mathcal{D}_b$ according to the equivalence relation I_b^*. Let this partition be $\mathcal{D}_{b,1}, \ldots, \mathcal{D}_{b,M}$.
6:      Select the permutation map $\sigma_1 : \mathcal{D} \to \mathcal{D}$ according to (3.31). Construct the permuted network $\mathcal{N}^{\sigma_1}$ (see Section 2.4.3) along with its conservation data (Γ_{σ_1}, c) (see Remark 2.8).
7:      For each $m = 1, \ldots, M$ define the finite set $\mathcal{E}_{b,m}^{\sigma_1}$ by (3.33).
8:      Evaluate d_f and d_r according to (3.35).
9:      **if** $d_r > 0$ **then** set $(\sigma_2, \phi, \widetilde{\mathcal{N}}^{\sigma_2}, \mathcal{D}_f) = $ FINDREDUCEDNETWORK$(\mathcal{N}^{\sigma_1}, \Gamma_{\sigma_1}, c)$.
10:      **else** set $\sigma_2 = \sigma_1$, $\mathcal{D}_f = \mathcal{D}_u$ and $\widetilde{\mathcal{N}}^{\sigma_2} = \mathcal{N}^{\sigma_1}$.
11:      **end if**
12:      Set $F = \{1, \ldots, d_f\}$. For each $m = 1, \ldots, M$ compute the set of closed communication classes $\mathcal{C}_m(F)$ for relation $\Theta_m(F)$ given by (3.49). Let $\mathcal{C}(F)$ be as in Definition 3.3.
13:      **for all** $\mathcal{F}_b^{\sigma_2} \in \mathcal{C}(F)$ **do**
14:          **if** CHECKBIRTHEXHAUSTIVITY$(\widetilde{\mathcal{N}}^{\sigma_2}, \mathcal{F}_b^{\sigma_2} \times \mathbb{N}_0^{d_f})$ **And** Network $\widetilde{\mathcal{N}}^{\sigma_2}$ satisfies (3.55) **then**
15:              Construct the inverse network $\widetilde{\mathcal{N}}_{\text{inv}}^{\sigma_2}$ (see Section 2.4.2).
16:              **if** CHECKBIRTHEXHAUSTIVITY$(\widetilde{\mathcal{N}}_{\text{inv}}^{\sigma_2}, \mathcal{F}_b^{\sigma_2} \times \mathbb{N}_0^{d_f})$ **then**
17:                  Let Φ be the graph of map ϕ (see (3.27)).
18:                  Define the permutation matrix P_{σ_2} by (2.18) and let

$$\mathcal{E} = \{P_{\sigma_2}^T x \in \mathbb{N}_0^d : x \in \mathcal{F}_b^{\sigma_2} \times \Phi\}.$$

19:                  **Output:** The state-space $\mathcal{E}$ is irreducible for the conservation network $(\mathcal{N}, \Gamma, c)$.
20:              **end if**
21:          **end if**
22:      **end for**
23: **end function**

The network-reduction method FINDREDUCEDNETWORK$(\cdot)$ proceeds as following. It first computes the set $\mathcal{I}_f$ (see (3.37)), whose elements determine all plausible partitions of the set of *unbounded* species $(\mathcal{D}_u)$ into the sets of *free* $(\mathcal{D}_f)$ and *restricted* $(\mathcal{D}_r)$ species. For any $I \in \mathcal{I}_f$, the reduced network is constructed as described in Sections 3.2 and 3.3. If this reduced network satisfies Assumption 2.3 then the procedure is stopped after returning this reduced network (along with its parameters) to the main method

FINDIRREDUCIBLESTATESPACES($\cdot$). Otherwise the procedure moves to a different choice of $I \in \mathcal{I}_f$. If all the choices have been exhausted and no reduced network satisfying Assumption 2.3 is found, then this method returns **Quit** which means that our approach will not work for the considered conservation network $(\mathcal{N}, \Gamma, c)$.

Algorithm 2 Reduces the conservation network $(\mathcal{N}^{\sigma_1}, \Gamma_{\sigma_1}, c)$ by eliminating the *restricted* species.

Require: The sets of *bounded* and *unbounded* species for network $\mathcal{N}^{\sigma_1}$ are given by $\mathcal{D}_b = \{\sigma_1(i) : i = 1, \ldots, d_b\}$ and $\mathcal{D}_u = \{\sigma_1(d_b + i) : i = 1, \ldots, (d_f + d_r)\}$ respectively.

 1: **function** FINDREDUCEDNETWORK($\mathcal{N}^{\sigma_1}, \Gamma_{\sigma_1}, c$)
 2:     Compute the set $\mathcal{I}_f$ given by (3.37).
 3:     **for all** $I \in \mathcal{I}_f$ **do**
 4:         Set $\mathcal{D}_f = \{\sigma_1(d_b + i) : i \in I\}$ and $\mathcal{D}_r = \{\sigma_1(d_b + i) : i \in I^c\}$.
 5:         Select the permutation map $\sigma_2 : \mathcal{D} \to \mathcal{D}$ according to (3.38).
 6:         Construct the permuted network $\mathcal{N}^{\sigma_2} = (\mathcal{V}^{\sigma_2}, \mathcal{O}^{\sigma_2}, \Lambda^{\sigma_2})$ (see Section 2.4.3) along with its conservation data (Γ_{σ_2}, c) (see Remark 2.8).
 7:         Define the affine map $\phi : \mathbb{R}^{d_f} \to \mathbb{R}^{d_r}$ according to (3.43).
 8:         Let $\widetilde{\mathcal{V}}^{\sigma_2} = \mathrm{Proj}(\mathcal{V}^{\sigma_2}, 1, d_b + d_f)$ and $\widetilde{\mathcal{O}}^{\sigma_2} = \mathrm{Proj}(\mathcal{O}^{\sigma_2}, 1, d_b + d_f)$.
 9:         Let $\widetilde{\Lambda}^{\sigma_2} = (\widetilde{\lambda}_1^{\sigma_2}, \ldots, \widetilde{\lambda}_K^{\sigma_2})$ where each $\widetilde{\lambda}_k^{\sigma_2}$ is given by (3.45).
10:         Define the reduced network as $\widetilde{\mathcal{N}}^{\sigma_2} = (\widetilde{\mathcal{V}}^{\sigma_2}, \widetilde{\mathcal{O}}^{\sigma_2}, \widetilde{\Lambda}^{\sigma_2})$.
11:         **if** Network $\widetilde{\mathcal{N}}^{\sigma_2}$ satisfies Assumption 2.3 **then return** $(\sigma_2, \phi, \widetilde{\mathcal{N}}^{\sigma_2}, \mathcal{D}_f)$
12:         **end if**
13:     **end for**
14:     **return QUIT**
15: **end function**

The method CHECKBIRTHEXHAUSTIVITY($\cdot$) arranges the *free* species of a network into birth-cascades as described in Section 3.4. The cascade-construction process is terminated when an *empty* cascade is reached because all the subsequent cascades will also be empty. Upon termination, this method returns **True** or **False** depending of whether all the *free* species belong to one of the birth-cascades or not.

Algorithm 3 Checks birth-exhaustivity of a network $\widetilde{\mathcal{N}}^{\sigma_2} = (\widetilde{\mathcal{V}}^{\sigma_2}, \widetilde{\mathcal{O}}^{\sigma_2}, \widetilde{\Lambda}^{\sigma_2})$ with state-space $\mathcal{F}_b^{\sigma_2} \times \mathbb{N}_0^{d_f}$


Require: Network $\widetilde{\mathcal{N}}^{\sigma_2}$ has $(d_b + d_f)$ species in the sets $\mathcal{D}_b$ and $\mathcal{D}_f$ given by $\mathcal{D}_b = \{\sigma_2(i) : i = 1, \ldots, d_b\}$ and $\mathcal{D}_f = \{\sigma_2(d_b + i) : i = 1, \ldots, d_f\}$ respectively.

 1: **function** CHECKBIRTHEXHAUSTIVITY($\widetilde{\mathcal{N}}^{\sigma_2}, \mathcal{F}_b^{\sigma_2} \times \mathbb{N}_0^{d_f}$)
 2:     Set $l = 0$ and $H_l = \emptyset$
 3:     **repeat**
 4:         $l \leftarrow l + 1$
 5:         For each $m = 1, \ldots, M$ compute the set by $G_l(m)$ by (3.51).
 6:         Set $G_l = \bigcup_{m=1}^{M} G_l(m)$ and update $H_l \leftarrow H_{l-1} \cup G_l$.
 7:     **until** $G_l = \emptyset$
 8:     **if** $H_l = \{1, \ldots, d_f\}$ **then return True**
 9:     **else return False**
10:     **end if**
11: **end function**

---

4 Biological Examples

In this section we will illustrate our procedure for finding irreducible state-spaces using examples from Systems and Synthetic Biology. In all the examples, the reactions are numbered in the order of their appearance and the propensity functions (λ_k-s) have *mass-action* form (2.15) unless otherwise stated.

Example 4.1 (Gene-expression networks) As our first example we consider the gene-expression network [32] along with its many variants. The simplest such network has three species: Gene (G), mRNA (M) and protein (P), and the following four reactions:

$$G \longrightarrow G + M, \qquad M \longrightarrow M + P, \qquad M \longrightarrow \emptyset \qquad \text{and} \qquad P \longrightarrow \emptyset. \qquad (4.58)$$

These reactions represent the transcription of mRNA by the gene, the translation of mRNA into protein, and the degradation of mRNA and protein respectively. Since the copy-numbers of G are not changed by these reactions, we can disregard its dynamics and replace the first reaction with $\emptyset \longrightarrow M$. Hence this network has effectively only two species $\mathbf{S}_1 = M$ and $\mathbf{S}_2 = P$ which are both *free*. The consumption vectors $(\nu_k$-s) of the four reactions are: $\nu_1 = (0,0), \nu_2 = (1,0), \nu_3 = (1,0)$ and $\nu_4 = (0,1)$. This network satisfies Assumption 2.3 and condition (3.55) holds because both species degrade naturally. One can easily check that for this network $\mathcal{N}$ is birth-exhaustive because the first two birth-cascades are: $B_1 = \{M\}$ and $B_2 = \{P\}$. It is also death-exhaustive as its first death-cascade, or the first birth-cascade of the inverse network $\mathcal{N}_{\mathrm{inv}}$ is $D_1 = \{M, P\}$. Therefore by Theorem 3.4 and Remark 3.6 we can conclude that $\mathbb{N}_0^2$ is the unique irreducible state-space for network $\mathcal{N}$.

Let us change the above network to incorporate transcriptional feedback by the protein molecules [23]. The propensity of the mRNA transcription reaction is now given by Hill-type kinetics of the form:

$$\lambda_1(x_1, x_2) = \theta_1 + \theta_{\mathrm{fb}} \frac{x_2^n}{c + x_2^n}, \qquad (4.59)$$

where x_2 is the number of protein molecules, and $\theta_1, \theta_{\mathrm{fb}}, c$ and n are strictly positive parameters. This modified network, denoted by $\mathcal{N}'$, still satisfies Assumption 2.3, and we can check the irreducibility of state-space $\mathbb{N}_0^2$ for network $\mathcal{N}'$ in exactly the same way as above. However if $\theta_1 = 0$, then network $\mathcal{N}'$ violates Assumption 2.3, because for $(x_1, x_2) = (1, 0)$ we have $x \geq \nu_1 = (0, 0)$ but $\lambda_1(x_1, x_2) = 0$. In this case, we can ensure that network $\mathcal{N}'$ satisfies Assumption 2.3 by changing the mRNA transcription reaction to $P \longrightarrow P + M$. Note that the dynamics remains unaffected by this change and Assumption 2.3 holds for this network because ν_1 is now $(0, 1)$, and hence we can apply our procedure to find irreducible state-spaces. However this new network is not birth-exhaustive which suggests that $\mathbb{N}_0^2$ may not be irreducible, which is indeed the case because $(0, 0)$ is an *absorbing* state for the network where all propensities are zero. This simple example shows that even though our procedure only provides sufficient criteria for determining irreducibility, it can still be useful for detecting its failure.

We now consider the above gene expression network (4.58) in the situation where the Gene (G) can spontaneously switch between an active (G^*) an an inactive (G^0) form [23], and the transcription of mRNA (M) is only possible in the active form. This network, denoted by $\widehat{\mathcal{N}}$ can be described by four species: active Gene $(\mathbf{S}_1 = G^*)$, inactive Gene $(\mathbf{S}_2 = G^0)$, mRNA $(\mathbf{S}_3 = M)$ and protein $(\mathbf{S}_4 = P)$, and six reactions given by

$$G^* \longrightarrow G^0, \quad G^0 \longrightarrow G^*, \quad G^* \longrightarrow G^* + M, \quad M \longrightarrow M + P, \quad M \longrightarrow \emptyset \quad \text{and} \quad P \longrightarrow \emptyset. \qquad (4.60)$$

One can check that $\gamma = (1, 1, 0, 0)$ is the only (independent) conservation relation for network $\widehat{\mathcal{N}}$. Let Γ be the 4×1 matrix whose single column is γ and let $c = 1$. We define (Γ, c) to be the conservation data for this network, and it implies that throughout the dynamics the copy-number $X_1(t)$ of species $\mathbf{S}_1$ and the copy-number $X_2(t)$ of species $\mathbf{S}_2$ satisfy $X_1(t) + X_2(t) = 1$. In other words, only one copy of the Gene (G) is present which can either exist in active or inactive form. Note that γ is a positive conservation relation and it causes the species G^* and G^0 to be *bounded*. The other two species are *free* and there are no *restricted* species in network $\widehat{\mathcal{N}}$. The dynamics of the *bounded* species evolves on the set $\mathcal{E} = \{(1, 0), (0, 1)\}$ and one can verify that $\mathcal{E} \times \mathbb{N}_0^2$ is the only CISS for the network. Moreover this network is birth-exhaustive for CISS $\mathcal{E} \times \mathbb{N}_0^2$ as the first two birth-cascades are $B_1 = \{M\}$ and $B_2 = \{P\}$ as before. Similarly this network is also *death-exhaustive* for CISS $\mathcal{E} \times \mathbb{N}_0^2$ because the first death-cascade is $D_1 = \{M, P\}$. Condition 3.55 is satisfied and hence using Theorem 3.4 along with Remark 3.6, we can conclude that $\mathcal{E} \times \mathbb{N}_0^2$ is the unique irreducible state-space for network $\widehat{\mathcal{N}}$.

Example 4.2 The aim of this example is to illustrate how *restricted* species arise and how they can be eliminated (see Sections 3.2 and 3.3). Consider a reaction network with two species $\mathbf{S}_1$ and $\mathbf{S}_2$, and two

reactions given by

$$\emptyset \longrightarrow \mathbf{S}_1 + \mathbf{S}_2 \qquad \text{and} \qquad \mathbf{S}_1 + \mathbf{S}_2 \longrightarrow \emptyset.$$

Since both the species are produced and degraded together, $\gamma = (1, -1)$ is the only (independent) conservation relation for this network. We define (Γ, c) to be the conservation data, where Γ is the 2×1 matrix with column γ and c is any integer. This conservation data says that throughout the dynamics the copy-number $X_1(t)$ of species $\mathbf{S}_1$ and copy-number $X_2(t)$ of species $\mathbf{S}_2$ satisfy $X_1(t) - X_2(t) = c$, which shows that each species "mimics" the dynamics of the other one. We have the flexibility of declaring any of the two species as *free* and then the other one will be *restricted*. This flexibility manifests in the choice of I from the set $\mathcal{I}_f$ (3.37) which has two elements $\{1\}$ and $\{2\}$ in this case. If we select I to be $\{1\}$ then the reduced network consists of one *free* species $\mathbf{S}_1$ and two reactions given by $\emptyset \longrightarrow \mathbf{S}_1$ and $\mathbf{S}_1 \longrightarrow \emptyset$ with propensity functions $\widetilde{\lambda}_1(x_1) = \theta_1$ and $\widetilde{\lambda}_2(x_1) = \theta_2 x_1(x_1 - c)$. Note that if $c \leq 0$, then this reduced network satisfies Assumption 2.3 and using Theorem 3.4 and Remark 3.6 we can conclude that $\mathbb{N}_0$ is the only irreducible state-space for the reduced network and hence by Proposition 3.2

$$\mathcal{E}_1 = \{(x_1, x_2) \in \mathbb{N}_0^2 : x_2 = x_1 - c\} \tag{4.61}$$

is the only irreducible state-space for the original network. On the other hand if $c > 0$ then this reduced network does not satisfy Assumption 2.3 because $\widetilde{\lambda}_2(x_1) = 0$ for $x_1 = c > 0$. However in this case we can choose I to be $\{2\}$ to obtain another reduced network with one *free* species $\mathbf{S}_2$ and two reactions $\emptyset \longrightarrow \mathbf{S}_2$ and $\mathbf{S}_2 \longrightarrow \emptyset$ with propensity functions $\widetilde{\lambda}_1(x_2) = \theta_1$ and $\widetilde{\lambda}_2(x_2) = \theta_2 x_2(x_2 + c)$ respectively. Observe that this new reduced network satisfies Assumption 2.3 and we can use Theorem 3.4 and Remark 3.6 as before to arrive at the same conclusion that $\mathcal{E}_1$ given by (4.61) is the unique irreducible state-space for the original network.

Example 4.3 (Genetic toggle switch) With this well-known example from Synthetic Biology we demonstrate that non-mass-action kinetics can be easily incorporated in our setup. Consider the network of a synthetic toggle switch [12] with two species $\mathbf{S}_1 = \mathcal{U}$ and $\mathbf{S}_2 = \mathcal{V}$ and the following four reactions

$$\emptyset \longrightarrow \mathcal{U}, \qquad \mathcal{U} \longrightarrow \emptyset, \qquad \emptyset \longrightarrow \mathcal{V} \qquad \text{and} \qquad \mathcal{V} \longrightarrow \emptyset.$$

The propensity functions of reactions 2 and 4 have the mass-action form, while the propensity functions for the other two reactions are given by

$$\lambda_1(x_1, x_2) = \frac{\alpha_1}{1 + x_2^\beta} \qquad \text{and} \qquad \lambda_3(x_1, x_2) = \frac{\alpha_2}{1 + x_1^\gamma},$$

where α_1, α_2, β and γ are certain positive parameter values. One can easily check that this network satisfies Assumption 2.3 even though λ_1 and λ_3 do not have the mass-action form. This network has no conservation relations and so both the species are *free*. Moreover this network is both birth and death exhaustive because the first birth and death cascades are $B_1 = D_1 = \{\mathcal{U}, \mathcal{V}\}$. Condition (3.55) is also satisfied and hence by Theorem 3.4 and Remark 3.6, we can conclude that $\mathbb{N}_0^2$ is the unique irreducible state-space for this network.

Example 4.4 (Circadian clock network) In this example from Systems Biology we consider the circadian clock oscillator moedel in [33], which has 9 species $\mathbf{S}_1, \ldots, \mathbf{S}_9$ and 16 reactions given in the table below.

No.	Reaction	No.	Reaction
1	$\mathbf{S_6} + \mathbf{S_2} \longrightarrow \mathbf{S_7}$	9	$\mathbf{S_2} \longrightarrow \emptyset$
2	$\mathbf{S_7} \longrightarrow \mathbf{S_6} + \mathbf{S_2}$	10	$\mathbf{S_9} \longrightarrow \mathbf{S_9} + \mathbf{S_3}$
3	$\mathbf{S_8} + \mathbf{S_2} \longrightarrow \mathbf{S_9}$	11	$\mathbf{S_8} \longrightarrow \mathbf{S_8} + \mathbf{S_3}$
4	$\mathbf{S_9} \longrightarrow \mathbf{S_8} + \mathbf{S_2}$	12	$\mathbf{S_3} \longrightarrow \emptyset$
5	$\mathbf{S_7} \longrightarrow \mathbf{S_7} + \mathbf{S_1}$	13	$\mathbf{S_3} \longrightarrow \mathbf{S_3} + \mathbf{S_4}$
6	$\mathbf{S_6} \longrightarrow \mathbf{S_6} + \mathbf{S_1}$	14	$\mathbf{S_4} \longrightarrow \emptyset$
7	$\mathbf{S_1} \longrightarrow \emptyset$	15	$\mathbf{S_2} + \mathbf{S_4} \longrightarrow \mathbf{S_5}$
8	$\mathbf{S_1} \longrightarrow \mathbf{S_1} + \mathbf{S_2}$	16	$\mathbf{S_5} \longrightarrow \mathbf{S_4}$

23

Computing the left nullspace of the associated stoichiometry matrix we can see that this network, denoted by $\mathcal{N}$, has two independent conservation relations given by $\gamma_1 = e_6 + e_7$ and $\gamma_2 = e_8 + e_9$, where e_i is the i-th standard basis vector in $\mathbb{R}^9$. Let $\Gamma = \mathrm{Col}(\gamma_1, \gamma_2)$ and $c = (1, 1)$. We define (Γ, c) to be the conservation data for this network, thereby indicating that throughout the dynamics, the sum of the copy-numbers of species $\mathbf{S}_6$ ($\mathbf{S}_8$) and $\mathbf{S}_7$ ($\mathbf{S}_9$) is 1. The species $\mathbf{S}_6$ ($\mathbf{S}_8$) and $\mathbf{S}_7$ ($\mathbf{S}_9$) correspond to the *bound* and the *unbound* form of an activator (promoter) gene. Therefore our conservation data (Γ, c) implies that a single copy of both these genes is present which can either exist in *bound* or *unbound* forms.

Due to the conservation relations, the species in the set $\mathcal{D}_b = \{6, 7, 8, 9\}$ are *bounded* while the rest of the species $\mathcal{D}_f = \{1, 2, 3, 4, 5\}$ are *free*. Note that species $\mathbf{S}_6$ and $\mathbf{S}_7$ do not directly interact with species $\mathbf{S}_8$ and $\mathbf{S}_9$, i.e. there is no reaction involving species from both the sets $\mathcal{D}_{b,1} = \{6, 7\}$ and $\mathcal{D}_{b,2} = \{8, 9\}$. Hence $\mathcal{D}_{b,1}$ and $\mathcal{D}_{b,2}$ are two interaction classes that partition the set $\mathcal{D}_b$ of *bounded* species. As in Section 3.1, we define a permutation map $\sigma : \{1, \ldots, 9\} \to \{1, \ldots, 9\}$ by $\sigma(1) = 6, \sigma(2) = 7, \sigma(3) = 8, \sigma(4) = 9$ and $\sigma(i) = (i-4)$ for $i = 5, 6, 7, 8$ and 9. We will consider network $\mathcal{N}^\sigma$ under permutation σ. Observe that the state-space for the *bounded* species in both interaction classes $\mathcal{D}_{b,1}$ and $\mathcal{D}_{b,2}$ is $\mathcal{E} = \{(0, 1), (1, 0)\}$. Hence the set $\mathcal{E}_0^\sigma = \mathcal{E} \times \mathcal{E} \times \mathbb{N}_0^5$ is a *maximal* state-space for the network $\mathcal{N}^\sigma$. The dynamics of the two *bounded* species in each interaction class $\mathcal{D}_{b,m}$ will only induce one closed equivalence class, which is either $\mathcal{E}$ or $\{(1, 0)\}$ depending on whether species $\mathbf{S}_2$ is *abundantly available* or not. In other words for any $A \subset F = \{1, 2, 3, 4, 5\}$ and each $m = 1, 2$ we have

$$\mathcal{C}_m(A) = \left[\begin{array}{cc} \{\mathcal{E}\} & \text{if } 2 \in A \\ \{(1, 0)\} & \text{otherwise} \end{array} \right]. \tag{4.62}$$

Using (4.62) we now construct the birth cascades for network $\mathcal{N}^\sigma$ as in Section 3.4. One can check that $B_1(1) = \{1\}$ and $B_1(2) = \{3\}$ and hence the first birth cascade is $B_1 = B_1(1) \cup B_1(2) = \{1, 3\}$. Similarly the second and third birth cascades are $B_2 = \{2, 4\}$ and $B_3 = \{5\}$ implying that network $\mathcal{N}^\sigma$ is *birth-exhaustive* because $\mathcal{D}_f = B_1 \cup B_2 \cup B_3$. Performing this cascade construction for the inverse network $\mathcal{N}_{\mathrm{inv}}^\sigma$ we find that the first two death cascades for network $\mathcal{N}^\sigma$ are $D_1 = \{1, 2, 3, 4\}$ and $D_2 = \{5\}$ implying that network $\mathcal{N}^\sigma$ is *death-exhaustive* as well. Since condition (3.55) is satisfied, using Theorem 3.4, Remark 3.5 and Remark 3.6, we can conclude that $\mathcal{E}_0 = \mathbb{N}_0^5 \times \mathcal{E} \times \mathcal{E}$ is the unique irreducible state-space for the original conservation network $(\mathcal{N}, \Gamma, c)$.

Example 4.5 (Networks with product-form stationary distributions) We mentioned in Section 2.3 that for many networks, the unique stationary distribution $\pi \in \mathcal{P}(\mathcal{E})$ corresponding to any irreducible state-space $\mathcal{E}$, has the product form (2.13) (see Theorem 4.3 in [2]), with the normalizing constant $M_\mathcal{E}$ is given by (2.14). Using our procedure we can identify the irreducible state-spaces on which this result can be applied. Moreover in cases where these state-spaces are countably infinite, our procedure represents them in such a form, that the normalizing constant can be easily evaluated by replacing the infinite sum in (2.14) by an exact finite sum, thereby preventing any truncation errors. To see this, consider a conservation network $(\mathcal{N}, \Gamma, c)$ with $d_f > 0$ *free* species and $d_b \geq 0$ *bounded* species. Suppose using Theorem 3.4 we can find an irreducible state-space of the form $\mathcal{E}^\sigma = \mathcal{F}_b^\sigma \times \mathbb{N}_0^{d_f}$ for the permuted conservation network $(\mathcal{N}^\sigma, \Gamma_\sigma, c)$ with respect to some permutation σ. If this network has the product form stationary distribution (2.13), then for computing the normalizing constant $M_{\mathcal{E}^\sigma}$ (2.14) we can simply replace the infinite sum over $\mathcal{E}^\sigma$ by a finite sum over $\mathcal{F}_b^\sigma$:

$$M_{\mathcal{E}^\sigma} = \left[\sum_{x \in \mathcal{E}^\sigma} \prod_{i=1}^{d_b+d_f} \frac{c_i^{x_i}}{x_i!} e^{-c_i} \right]^{-1} = \left[\sum_{y \in \mathcal{F}_b^\sigma} \left(\prod_{i=1}^{d_b} \frac{c_i^{y_i}}{y_i!} e^{-c_i} \right) \right]^{-1}$$

because the sum of a Poisson probability mass function over $\mathbb{N}_0$ is 1 (i.e. $\sum_{j \in \mathbb{N}_0} (\alpha^j / j!) e^{-\alpha} = 1$ for any $\alpha > 0$). This example demonstrates how our method can help in computing the exact stationary distribution for certain networks.

Appendix

In this section we provide a couple of lemmas that are needed for proving the main result Theorem 3.4.

Lemma 4.6 *Suppose that network $\widetilde{\mathcal{N}}^{\sigma_2}$ is birth-exhaustive for CISS $\mathcal{F}_b^{\sigma_2} \times \mathbb{N}_0^{d_f}$. Then for any $z_1, z_2 \in \mathcal{F}_b^{\sigma_2}$ and any $r_0 \in \mathbb{N}^{d_f}$ there exists a vector $x \in \mathbb{N}^{d_f}$ such that $x \geq r_0$ and*

$$(z_1, \overline{0}_{d_f}) \xrightarrow{\widetilde{\mathcal{N}}^{\sigma_2}} (z_2, x).$$

Proof. We prove this lemma by induction. Throughout this proof we denote the relation $\xrightarrow{\widetilde{\mathcal{N}}^{\sigma_2}}$ by $\longrightarrow$ and the finite set $\mathcal{F}_b^{\sigma_2}$ by $\mathcal{F}$. Let sets G_l and H_l be as defined in Section 3.4. We say that level l is *satisfiable* if for any $z_1 \in \mathcal{F}$ and $r \in \mathbb{N}_0^{d_f}$ with $\mathrm{supp}(r) \subset H_l \cup G_l$, there exists a $z_2 \in \mathcal{F}$ and $x \in \mathbb{N}_0^{d_f}$ such that $x \geq r$ and $(z_1, \overline{0}_{d_f}) \longrightarrow (z_2, x)$.

Assume that the *free* species in $\mathcal{D}_f^{\sigma_2}(H_l)$ are abundantly available and pick any $i \in G_l$. For any initial state $y_0 \in \mathcal{F}$, the dynamics of *bounded* species will lie in a closed communication class $C \in \mathcal{C}(H_l, \mathcal{F})$ after finitely many reactions. After some more reactions, the dynamics will be at a state $y_1 \in C$ for which there exists a reaction $k \in \mathcal{K}(y_1, H_l)$ which produces the *free* species $\sigma_2(d_b + i)$. Note that all these reactions will only consume the *free* species in $\mathcal{D}_f^{\sigma_2}(H_l)$ which are abundantly available. Repeating this sequence of reactions multiple times, we can obtain as many molecules of species $\sigma_2(d_b + i)$ as needed. Performing this procedure for each $i \in G_l$ proves that level l is satisfiable. Since level 1 is trivially satisfiable, by induction we can conclude that all levels are satisfiable.

The birth-exhaustivity of the network implies that there exists a $l_0 \in \mathbb{N}$ such that $H_{l_0} = F = \{1, \ldots, d_f\}$. Satisfiability of level l_0 guarantees that for any $z_1 \in \mathcal{F}$ and $r \in \mathbb{N}^{d_f}$ there exists a $z_3 \in \mathcal{F}$ and $x' \in \mathbb{N}^{d_f}$ such that $x' \geq r$ and $(z_1, \overline{0}_{d_f}) \longrightarrow (z_3, x')$. Since $\mathcal{F} \in \mathcal{C}(F)$ is a closed communication class, all the states in $\mathcal{F}$ are accessible from each other when all the *free* species are are abundantly available. Therefore for any $z_2 \in \mathcal{F}$ and $r_0 \in \mathbb{N}^{d_f}$ we can pick the components of r large enough to ensure that $(z_3, x') \longrightarrow (z_2, x)$ for some vector $x \in \mathbb{N}^{d_f}$ satisfying $x \geq r_0$. The transitivity of relation $\longrightarrow$ proves this lemma. $\qquad\square$

Lemma 4.7 *Suppose that (3.55) is satisfied. Then there exists a vector $r_0 \in \mathbb{N}^{d_f}$ such that for any $x_1, x_2 \in \mathbb{N}^{d_f}$ satisfying $x_1 \geq x_2 \geq r_0$ we have $(z, x_1) \xrightarrow{\widetilde{\mathcal{N}}^{\sigma_2}} (z, x_2)$ for any $z \in \mathcal{F}_b^{\sigma_2}$.*

Proof. The proof of this lemma is inspired by the proof of Theorem 3.4 in [29]. Throughout this proof we denote relation $\xrightarrow{\widetilde{\mathcal{N}}^{\sigma_2}}$ by $\longrightarrow$ and $e_i^{(d_f)}$ by e_i, which is i-th standard basis vector in $\mathbb{R}^{d_f}$. We first show that for each $i \in F = \{1, \ldots, d_f\}$ there exists a $r_i \in \mathbb{N}^{d_f}$ such that

$$(z, r_i) \longrightarrow (z, r_i - e_i) \tag{4.63}$$

for any $z \in \mathcal{F}_b^{\sigma_2}$. Fix any $i \in F$. Note that since (3.55) holds we have $-e_i \in \mathrm{Colspan}_{\mathbb{N}_0}(\widehat{S}_{\sigma_2})$. This implies that there is a sequence of reactions $k_1, \ldots, k_n \in \mathcal{K}_f^0$ such that $-e_i = \sum_{j=1}^{n} (\widehat{\rho}_{k_j}^{\sigma_2} - \widehat{\nu}_{k_j}^{\sigma_2})$. Let $y_1 = \overline{0}_{d_f}$ and for each $m = 2, \ldots, n$ let $y_m = \sum_{j=1}^{m-1} (\widehat{\rho}_{k_j}^{\sigma_2} - \widehat{\nu}_{k_j}^{\sigma_2})$. By choosing a $r_i \in \mathbb{N}^{d_f}$ with large enough entries we can ensure that $(r_i + y_m) \geq \widehat{\nu}_{k_m}^{\sigma_2}$ for each $m = 1, \ldots, n$. Since the reactions in $\mathcal{K}_f^0$ do not involve the *bounded* species and Assumption 2.3 is satisfied, such a choice of r_i also ensures that each reaction k_m has a positive probability of firing when the state of the *free* species is $(r_i + y_m)$. Hence (4.63) holds for any $z \in \mathcal{F}_b^{\sigma_2}$. Let $r_0 = \max_{i \in F}\{r_i\}$ and let $x \in \mathbb{N}^{d_f}$ be any vector satisfying $x \geq r_0$. Using Proposition 2.4 we can conclude that $(z, x) \longrightarrow (z, x - e_i)$ for any $z \in \mathcal{F}_b^{\sigma_2}$ and any $i \in F$. Now select any $x_1, x_2 \in \mathbb{N}^{d_f}$ satisfying $x_1 \geq x_2 \geq r_0$ and fix any $z \in \mathcal{F}_b^{\sigma_2}$. Let $\alpha = (x_1 - x_2) \in \mathbb{N}_0^{d_f}$ and express it as $\alpha = \sum_{i=1}^{d_f} \alpha_i e_i$. Exploiting the transitivity of relation $\longrightarrow$ and using $(z, x) \longrightarrow (z, x - e_i)$, α_i times for each i we obtain the accessibility chain $(z, x_2 + \alpha) \longrightarrow (z, x_2 + \alpha - \alpha_1 e_1) \longrightarrow (z, x_2 + \alpha - \alpha_1 e_1 - \alpha_2 e_2) \longrightarrow \ldots \longrightarrow (z, x_2 + \alpha - \sum_{i=1}^{d_f} \alpha_i e_i)$. But $x_2 + \alpha = x_1$ and $(x_2 + \alpha - \sum_{i=1}^{d_f} \alpha_i e_i) = x_2$ and hence the proof of this lemma is complete. $\qquad\square$

References

[1] U. Alon. *An introduction to systems biology : design principles of biological circuits.* Chapman & Hall/CRC mathematical and computational biology series. Chapman & Hall/CRC, Boca Raton (Fla.), 2007. 1

[2] D. F. Anderson, G. Craciun, and T. G. Kurtz. Product-form stationary distributions for deficiency zero chemical reaction networks. *Bulletin of Mathematical Biology*, 72(8):1947–1970, 2010. 1, 2.3, 2.3, 4.5

[3] A. P. Arkin, C. V. Rao, and D. M. Wolf. Control, exploitation and tolerance of intracellular noise. *Nature*, 420:231–237, 2002. 1

[4] J. Bascompte. Structure and dynamics of ecological networks. *Science*, 329(5993):765–766, 2010. 1

[5] S. I. Berger and R. Iyengar. Network analyses in systems pharmacology. *Bioinformatics*, 25(19):2466–2472, 2009. 1

[6] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004. 1, 2.2

[7] C. Briat, A. Gupta, and M. Khammash. Antithetic integral feedback ensures robust perfect adaptation in noisy biomolecular networks. *Cell Systems*, 2(1):15–26, 2016. 1, 2.2

[8] H. Cohen. *A course in computational algebraic number theory*, volume 138 of *Graduate Texts in Mathematics*. Springer-Verlag, Berlin, 1993. 3.4

[9] T. Cormen, C. Leiserson, R. Rivest, and C. Stein. *Introduction to Algorithms*. McGraw-Hill Science/Engineering/Math, New York, second edition, 2003. 3.4

[10] M. B. Elowitz, A. J. Levine, E. D. Siggia, and P. S. Swain. Stochastic gene expression in a single cell. *Science*, 297(5584):1183–1186, 2002. 1

[11] S. N. Ethier and T. G. Kurtz. *Markov processes : Characterization and Convergence*. Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics. John Wiley & Sons Inc., New York, 1986. 2.1, 1

[12] T. S. Gardner, C. R. Cantor, and J. J. Collins. Construction of a genetic toggle switch in escherichia coli. *Nature*, 403(6767):339–342, 2000. 1, 4.3

[13] D. T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *The Journal of Physical Chemistry*, 81(25):2340–2361, 1977. 2.1, 2.1

[14] D. T. Gillespie. A rigorous derivation of the chemical master equation. *Physica A: Statistical Mechanics and its Applications*, 188(13):404 – 425, 1992. 2.1

[15] J. Goutsias. Classical versus stochastic kinetics modeling of biochemical reaction systems. *Biophysical Journal*, 92(7):2350–2365, 2007. 1, 2.4.1

[16] G. R. Grimmett and D. R. Stirzaker. *Probability and random processes*. Oxford University Press, New York, third edition, 2001. 2.3

[17] A. Gupta, C. Briat, and M. Khammash. A scalable computational framework for establishing long-term behavior of stochastic reaction networks. *PLoS Comput Biol*, 10(6):e1003669, 06 2014. 1, 2.2, 2.2

[18] A. Gupta and M. Khammash. Determining the long-term behavior of cell populations: A new procedure for detecting ergodicity in large stochastic reaction networks. In *Proceedings of the 19th IFAC World Congress*, 2014. 1, 3

[19] H. Hethcote. The mathematics of infectious diseases. *SIAM Review*, 42(4):599–653, 2000. 1

[20] O. Kallenberg. *Foundations of modern probability*. Probability and its Applications (New York). Springer-Verlag, New York, second edition, 2002. 2.2

[21] J. P. Keener and J. Sneyd. *Mathematical physiology*, volume Interdisciplinary applied mathematics. Springer, 2009. 2.4.1

[22] J. G. Kemeny and J. L. Snell. *Finite Markov chains*. The University Series in Undergraduate Mathematics. D. Van Nostrand Co., Inc., Princeton, N.J.-Toronto-London-New York, 1960. 1, 3, 3.4, 3.4

[23] T. B. Kepler and T. C. Elston. Stochasticity in transcriptional regulation: Origins, consequences, and mathematical representations. *Biophysical Journal*, 81(6):3116 – 3136, 2001. 4.1, 4.1

[24] H. H. McAdams and A. Arkin. Stochastic mechanisms in gene expression. *Proc. Natl. Acad. Sci., Biochemistry*, 94:814–819, 1997. 1

[25] S. Meyn and R. L. Tweedie. *Markov chains and stochastic stability.* Cambridge University Press, Cambridge, second edition, 2009. With a prologue by Peter W. Glynn. 2.2, 2.2

[26] S. P. Meyn and R. L. Tweedie. Stability of Markovian processes. III. Foster-Lyapunov criteria for continuous-time processes. *Adv. in Appl. Probab.*, 25(3):518–548, 1993. 2.2, 2.2

[27] J. R. Norris. *Markov chains*, volume 2 of *Cambridge Series in Statistical and Probabilistic Mathematics.* Cambridge University Press, Cambridge, 1998. Reprint of 1997 original. 1, 2.2, 2.3

[28] J. A. Papin, T. Hunter, B. O. Palsson, and S. Subramaniam. Reconstruction of cellular signalling networks and analysis of their properties. *Nat Rev Mol Cell Biol*, 6(2):99–111, 02 2005. 1

[29] L. Pauleve, G. Craciun, and H. Koeppl. Dynamical properties of discrete reaction networks. *Journal of Mathematical Biology*, pages 1–18, 2013. 1, 4

[30] S. Schnell and C. Mendoza. Closed form solution for time-dependent enzyme kinetics. *Journal of Theoretical Biology*, 187(2):207 – 212, 1997. 2.4.1

[31] P. Smadbeck and Y. Kaznessis. Stochastic model reduction using a modified hill-type kinetic rate law. *The Journal of Chemical Physics*, 137(23), 12 2012. 2.4.1

[32] M. Thattai and A. van Oudenaarden. Intrinsic noise in gene regulatory networks. *Proceedings of the National Academy of Sciences*, 98(15):8614–8619, 2001. 4.1

[33] J. M. G. Vilar, H. Y. Kueh, N. Barkai, and S. Leibler. Mechanisms of noise-resistance in genetic oscillator. *Proc. Natl. Acad. Sci.*, 99(9):5988–5992, 2002. 1, 4.4