

Computational identification of irreducible state-spaces for stochastic reaction networks

Ankit Gupta and Mustafa Khammash
 Department of Biosystems Science and Engineering
 ETH Zurich
 Mattenstrasse 26
 4058 Basel, Switzerland.

Abstract

Stochastic models of reaction networks are becoming increasingly important in Systems Biology. In these models, the dynamics is generally represented by a continuous-time Markov chain whose states denote the copy-numbers of the constituent species. The state-space on which this process resides is a subset of non-negative integer lattice and for many examples of interest, this state-space is countably infinite. This causes numerous problems in analyzing the Markov chain and understanding its long-term behavior. These problems are further confounded by the presence of conservation relations among species which constrain the dynamics in complicated ways. In this paper we provide a linear-algebraic procedure to disentangle these conservation relations and represent the state-space in a special decomposed form, based on the copy-number ranges of various species and dependencies among them. This decomposed form is advantageous for analyzing the stochastic model and for a large class of networks we demonstrate how this form can be used for finding all the closed communication classes for the Markov chain within the infinite state-space. Such communication classes are irreducible state-spaces for the dynamics and they support all the extremal stationary distributions for the Markov chain. Hence our results provide important insights into the long-term behavior and stability properties of stochastic models of reaction networks. We discuss how the knowledge of these irreducible state-spaces can be used in many ways such as speeding-up stochastic simulations of multiscale networks or in identifying the stationary distributions of complex-balanced networks. We illustrate our results with several examples of gene-expression networks from Systems Biology.


Keywords: stochastic reaction networks; state-space analysis; communication classes; state-space irreducibility; stationarity; ergodicity.

Mathematical Subject Classification (2010): 60J22; 60J27; 60H35; 65C05.

1 Introduction

Many biological processes are described as reaction networks, where finitely many species interact with each other through some fixed reaction channels [1, 2, 3, 4, 5, 6]. Traditionally, reaction networks 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 timing of reactions becomes random, introducing *noise* into the dynamics, which can significantly change the behavior of the system being modeled [7]. Such situations arise commonly in Systems Biology, since intracellular networks often involve species with low copy-numbers like gene-transcripts, signaling proteins, messenger RNAs, transcription factors etc. [8, 7]. The biochemical noise generated by the intermittency of reactions can be taken into account using stochastic models of reaction networks. A common approach is to represent the dynamics as a continuous-time Markov chain (CTMC) whose states denote the copy-numbers of the constituent species. For recent surveys on these stochastic models and the methods available for analyzing them see [9, 10, 11]. Over the past few years, such models have been extensively used for understanding the role of noise in various biological mechanisms [12, 8].

We now formally describe such a stochastic model for a reaction network. Throughout this paper \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. Consider a network with d species $\mathbf{S}_1, \dots, \mathbf{S}_d$ which interact through K reaction channels of the form



where ν_{ik} and ρ_{ik} belong to \mathbb{N}_0 and they denote the number of molecules of \mathbf{S}_i that are consumed and produced by reaction k . Under the classical well-mixed assumption [13], the network's state at any time is described by a vector $x = (x_1, \dots, x_d) \in \mathbb{N}_0^d$ of copy-number counts of all the species, i.e. x_i denotes the copy-number of species \mathbf{S}_i . The net change in the state due to reaction k is simply given by the *stoichiometry vector* $\zeta_k \in \mathbb{Z}^d$ whose i -th component is $\zeta_{ik} = (\rho_{ik} - \nu_{ik}) \in \mathbb{Z}$. The usual CTMC model for the reaction dynamics [10] stipulates that the transition rate from state x to state $(x + \zeta_k)$ is $\lambda_k(x)$ for each $k = 1, \dots, K$, where $\lambda_1, \dots, \lambda_K : \mathbb{N}_0^d \rightarrow \mathbb{R}_+$ are called the *propensity* functions for the network and they are assumed to satisfy:

$$\text{for any } x \in \mathbb{N}_0^d \text{ and } k = 1, \dots, K \text{ if } \lambda_k(x) > 0 \text{ then } (x + \zeta_k) \in \mathbb{N}_0^d. \quad (1.2)$$

This condition ensures that the set \mathbb{N}_0^d is a valid *state-space* for the CTMC, because it is *closed* under the reaction dynamics i.e. if the starting state is in \mathbb{N}_0^d then the dynamics remains in \mathbb{N}_0^d throughout its trajectory.

Under mild conditions on the propensity functions (see Chapter 6 in [14]), one can ensure that the CTMC

$$(X(t))_{t \geq 0} = (X_1(t), \dots, X_d(t))_{t \geq 0}$$

with the above transition structure is well-defined for any initial state $x_0 \in \mathbb{N}_0^d$. Define the probability that the reaction dynamics is at state $y \in \mathbb{N}_0^d$ at time t by

$$p_{x_0}(t, y) = \mathbb{P}(X(t) = y). \quad (1.3)$$

Then the dynamics of $p_{x_0}(t, \cdot)$ is given by the Chemical Master Equation (CME) [15] which has the following form:

$$\frac{dp_{x_0}(t, y)}{dt} = \sum_{k=1}^K (p_{x_0}(t, y - \zeta_k) \lambda_k(y - \zeta_k) - p_{x_0}(t, y) \lambda_k(y)), \quad (1.4)$$

for each $y \in \mathbb{N}_0^d$ which can be accessed by the CTMC. It can be shown that CME has a unique solution provided that the CTMC is non-explosive (see [10]). For most examples in Systems Biology, the number of accessible states is either infinite or very large, and hence solving the CME to obtain the probabilities $p_{x_0}(t, y)$ is nearly impossible. One generally estimates these values by simulating the CTMC $(X(t))_{t \geq 0}$ using Monte Carlo methods such as Gillespie's *Stochastic Simulation Algorithm* (SSA) [13]. Another popular approach for obtaining approximate solutions is the *Finite State Projection* (FSP) method which efficiently truncates the state-space to a small finite set and then solves the CME over this finite set [16]. Both SSA and FSP based approaches work well for smaller networks and over finite time-intervals, and hence they do not help in satisfactorily assessing the long-term behavior and stability properties of the stochastic model. For Markov chains over a finite state-space, one way to assess this long-term behavior is by computing the disjoint *closed communication classes* in the state-space (see Section 2.1), using matrix or graphical methods [17]. However there do not exist methods for systematically finding all the closed communication classes for stochastic reaction network models with infinite state-spaces. The main goal of this paper is to develop such a method that can *provably* find all such classes under biologically reasonable assumptions on the network. This has important implications regarding the long-term behavior of the stochastic model because these closed communication classes support the stationary distributions (see Section 2.2), which are like attracting fixed-points for the CME (1.4) in the space of probability distributions. Moreover each closed communication class serves as an *irreducible state-space* for the underlying CTMC (see Section 2.1). For this reason, we refer to a closed communication class as an irreducible state-space in this paper.

For many networks, the state-space \mathbb{N}_0^d is simply *too large*, in the sense that it contains several extraneous states that are never visited by the dynamics. This is mainly due to the *conservation relations* present in the network which impose constraints on the copy-number ranges of the involved species. Hence our first task is to develop a systematic procedure to *weed-out* these extraneous states and obtain a *smaller* non-empty set \mathcal{E}_0 which is also a valid state-space (i.e. (1.2) holds with \mathbb{N}_0^d replaced with \mathcal{E}_0). Furthermore we would like the representation of \mathcal{E}_0 to be explicit enough to ensure that the copy-number ranges of various species can be easily identified. Apart from enabling the search for irreducible state-spaces, which is the primary goal of this paper, this explicit state-space representation allows us to gain a better understanding of the network dynamics. We now explore these issues in a greater detail.

Let us return to the CTMC $(X(t))_{t \geq 0}$ starting at some initial state $x_0 \in \mathbb{N}_0^d$. The ideal or the smallest state-space \mathcal{E} for this process would simply be the set of all states in \mathbb{N}_0^d that the reaction dynamics has a positive probability of reaching in a finite time¹, i.e.

$$\mathcal{E} = \{y \in \mathbb{N}_0^d : p_{x_0}(t, y) > 0\}, \quad (1.5)$$

for some $t > 0$. The set \mathcal{E} is non-empty (as $x_0 \in \mathcal{E}$) and using Chapman-Kolmogorov inequalities (see [18]) one can easily check that \mathcal{E} is a valid state-space and so the CTMC $(X(t))_{t \geq 0}$ resides in this set throughout its trajectory. However the set \mathcal{E} is difficult to characterize because the probabilities $p_{x_0}(t, \cdot)$ are usually unknown. Hence we look for a bigger set \mathcal{E}_0 which contains \mathcal{E} and is also a valid state-space. The standard choice is the *stoichiometry compatibility class* (see [10]) defined by

$$\mathcal{E}_0 = (x_0 + \text{Range}(S)) \cap \mathbb{N}_0^d, \quad (1.6)$$

where $S = \text{Col}(\zeta_1, \dots, \zeta_K)$ is the $d \times K$ *stoichiometry matrix* whose columns are the reaction stoichiometry vectors and $\text{Range}(S)$ is the range or column space of S . To see the containment $\mathcal{E} \subset \mathcal{E}_0$ note that if $y \in \mathcal{E}$ then y must be reachable from x_0 in finitely many transition steps in each of the directions ζ_1, \dots, ζ_K . Let r_k be the number of steps needed in direction k , then setting $r = (r_1, \dots, r_K) \in \mathbb{N}_0^K$ we must have $y = x_0 + Sr$ which ensures that $(y - x_0) \in \text{Range}(S)$ and hence $y \in \mathcal{E}_0$. We next write \mathcal{E}_0 in terms of conservation relations for the network, which are the nonzero vectors in the *left nullspace* of the stoichiometry matrix S

$$\mathcal{L}(S) = \{\gamma \in \mathbb{R}^d : \gamma^T S = \mathbf{0}^T\}, \quad (1.7)$$

where $\mathbf{0}$ is the vector of all zeros in \mathbb{R}^K . For any nonzero $\gamma = (\gamma_1, \dots, \gamma_d) \in \mathcal{L}(S)$, the relation $\gamma^T S = \mathbf{0}^T$ implies that for each reaction $k = 1, \dots, K$, the stoichiometry vector ζ_k is orthogonal to γ , i.e. $\langle \gamma, \zeta_k \rangle = 0$, where $\langle \cdot, \cdot \rangle$ denotes the standard inner product in \mathbb{R}^d . Therefore the stochastic reaction dynamics $(X(t))_{t \geq 0}$ will satisfy

$$\langle \gamma, X(t) \rangle = \langle \gamma, X(0) \rangle = \langle \gamma, x_0 \rangle \quad \text{for all } t \geq 0, \quad (1.8)$$

which means that the copy-numbers of species included in the support set

$$\text{supp}(\gamma) := \{i = 1, \dots, d : \gamma_i \neq 0\} \quad (1.9)$$

conserve the linear constraint specified by vector γ . Suppose that for some integer $n > 0$ we have $\text{Rank}(S) = (d - n)$ and so the dimension of the subspace $\mathcal{L}(S)$ is n . Let $\{\gamma_1, \dots, \gamma_n\}$ be a *basis* for $\mathcal{L}(S)$ and define a $d \times n$ matrix Γ and a $n \times 1$ vector c by

$$\Gamma = \text{Col}(\gamma_1, \dots, \gamma_n) \quad \text{and} \quad c = \Gamma^T x_0. \quad (1.10)$$

From (1.8) and the orthogonality of vector spaces $\mathcal{L}(S)$ and $\text{Range}(S)$, we can equivalently express the set \mathcal{E}_0 as


$$\mathcal{E}_0 = \{x \in \mathbb{N}_0^d : \Gamma^T x = c\}. \quad (1.11)$$

Taking a cue from the terminology used in [19], we call Γ as the *conservation matrix*, c as the *conservation vector* and the pair (Γ, c) as the *conservation data* for the network. If the subspace $\mathcal{L}(S)$ is trivial (i.e.

¹The choice of this finite time is not important because for a CTMC if $p_{x_0}(t, y) > 0$ for some $t > 0$ then the same holds for all $t > 0$ (see Theorem 3.2.1 in [18])

$n = 0$) then $\mathcal{E}_0 = \mathbb{N}_0^d$, but in the other situation when this subspace is nontrivial (i.e. $n > 0$), the states in \mathbb{N}_0^d outside \mathcal{E}_0 are not reachable by the dynamics because \mathcal{E}_0 is a valid state-space for the CTMC $(X(t))_{t \geq 0}$ and it contains the initial state $X(0) = x_0$.

The representations (1.6) or (1.11) for \mathcal{E}_0 are quite abstract, making it difficult to get a sense of the dynamics and particularly the copy-number ranges of all the species. We remedy this problem by developing a method that systematically screens the space of conservation relations $\mathcal{L}(S)$ and expresses \mathcal{E}_0 in a more explicit form. We briefly describe the steps needed for this purpose. We call a nonzero vector $\gamma \in \mathcal{L}(S)$ a *semi-positive* conservation relation if all its nonzero entries are positive. Up to independence, the only other type of conservation relations are the *mixed-sign* ones which have at least two nonzero entries with opposite signs. Semi-positive conservation relations ensure that the species in their support-sets (see (1.9)) are *bounded* in the sense that their copy-numbers have a bounded range. This is evident from (1.8) because if γ is a semi-positive conservation relation then for any $i \in \text{supp}(\gamma)$, we have $\gamma_i > 0$ and since $\langle \gamma, X(t) \rangle = \sum_{j \in \text{supp}(\gamma)} \gamma_j X_j(t)$, we must also have $X_i(t) \in [0, \langle \gamma, x_0 \rangle / \gamma_i]$. We scan the space of semi-positive conservation relations and identify all the *bounded* species along with their suitable copy-number ranges. It is possible that these semi-positive relations do not span the whole space of conservation relations $\mathcal{L}(S)$, and so we then look for any mixed-sign conservation relations among the remaining *unbounded* species. We show that such conservation relations force a certain subset of *unbounded* species (called *restricted* species) to mimic the dynamics of the remaining *unbounded* species (called *free* species) according to an appropriately constructed affine map. For example, the following ATP-hydrolysis reaction occurs in many living cells:



Generally **ATP** and **H₂O** molecules are present in high concentrations in the cytosol, and so the dynamics of low copy-number species $\mathbf{S}_1 = \mathbf{ADP}$ and $\mathbf{S}_2 = \mathbf{P}$ can be well approximated by the simple network



The stoichiometry matrix for this network is

$$S = \begin{bmatrix} 1 & -1 \\ 1 & -1 \end{bmatrix}$$

and the space $\mathcal{L}(S)$ of conservation relations has dimension 1. One can see that there are no semi-positive conservation relations and the only independent conservation relation is $\gamma = (1, -1)$ which is mixed-sign. Moreover (1.8) forces $X_1(t) = \phi(X_2(t))$ for all $t \geq 0$, where ϕ is the affine map given by $\phi(x) = x + \langle \gamma, x_0 \rangle$.

In large networks there are several conservation relations and each species can participate in many such relations. To account for all the possibilities, we will employ standard linear-algebraic methods, such as basic matrix manipulations, solving linear-algebraic systems and Linear Programs (LPs) [20], to classify each species as one of three types: *free*, *bounded* or *restricted*. Under fairly general conditions satisfied by most biological networks, we prove that by relabeling the species, we can express the state-space \mathcal{E}_0 in a special decomposed form

$$\mathcal{E}_0 = \mathcal{E}_b \times \Phi, \quad (1.13)$$

where \mathcal{E}_b is a finite set in $\mathbb{N}_0^{d_b}$ and Φ is the *graph* of an *affine* function $\phi : \mathbb{R}^{d_f} \rightarrow \mathbb{R}^{d_r}$ restricted to the nonnegative integer orthant, i.e.

$$\Phi = \left\{ (x, \phi(x)) : x \in \mathbb{N}_0^{d_f} \text{ and } \phi(x) \in \mathbb{N}_0^{d_r} \right\}. \quad (1.14)$$

Here d_f, d_b and d_r are nonnegative integers denoting the number of *free*, *bounded* and *restricted* species respectively. The finite set $\mathcal{E}_b \subset \mathbb{N}_0^{d_b}$ contains 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* and *restricted* species, with the latter being “locked” in a fixed affine relationship (given by function ϕ) with the former. Notice that in comparison to both (1.6) or (1.11), the form (1.13) for state-space \mathcal{E}_0 is *more explicit* as it clearly expresses the copy-number ranges for each species as well as the relationships among them. This enables a better understanding of the

dynamics which can be leveraged to improve the efficiency of existing analytical methods. For instance, one can use the form (1.13) to design optimal state-space truncations for the *Finite State Projection* (FSP) method for solving CMEs [16]. One can also use this form as a guide for automatically separating the high-copy-number and the low-copy-number species for the *method of conditional moments* (MCM) approach [21] or for deriving suitable hybrid approximations of the dynamics [22]. We do not explore these ideas in this paper but instead focus on using this form for the analysis of the long-term behavior of the underlying Markov process.

Suppose there exists a conservation relation in $\mathcal{L}(S)$ such that all its d components are strictly positive. Such a situation generally occurs when the network satisfies some form of global *mass conservation relation* and in this case all the species are *bounded* (i.e. $d_b = d$, $d_f = d_r = 0$), and so the state-space $\mathcal{E}_0 = \mathcal{E}_b$ is finite. Using elementary matrix or graph-theoretic methods [17] we can easily find all the closed communication classes or the irreducible state-spaces (see Section 2.1) $\mathcal{E}_1, \dots, \mathcal{E}_Q$ within \mathcal{E}_0 , where the dynamics eventually lies starting from any initial state x_0 . These classes are mutually disjoint and each \mathcal{E}_q supports a unique stationary distribution π_q for the stochastic reaction dynamics. Moreover these distributions π_1, \dots, π_Q are the *extremal points* of the simplex formed by all the stationary distributions of the network

$$\Sigma = \left\{ \pi = \sum_{q=1}^Q \alpha_q \pi_q : \text{ each } \alpha_1, \dots, \alpha_Q \geq 0 \text{ and } \sum_{q=1}^Q \alpha_q = 1 \right\}. \quad (1.15)$$

Typically for reaction networks in Systems Biology, a conservation relation with all strictly positive entries does not exist. This is because such networks are abstract representations of the actual processes where many details, such as the dynamics of abundant species (like ATP molecules, enzymes etc.), are intentionally omitted to make the analysis more tractable and pertinent to a given problem. Consequently, unlike classical chemical kinetics, global mass conservation fails because mass is allowed to be created and destroyed. Hence it is important to consider this general case, where *free* and *restricted* species are present, and the state-space \mathcal{E}_0 must be necessarily infinite provided certain natural assumptions are satisfied (see Assumption 2.1). In such a scenario the stationary distributions do not always exist, but if the existence can be guaranteed on each irreducible state-space (see Section 2.1), then again the extremal points π_1, π_2, \dots of the simplex of stationary distributions Σ are the unique stationary distributions supported on all the irreducible state-spaces $\mathcal{E}_1, \mathcal{E}_2, \dots$, within the infinite state-space \mathcal{E}_0 . The primary goal of this paper is to develop a method that explicitly identifies all these irreducible state-spaces using the decomposed state-space form (1.13). For this purpose we adapt and generalize the ideas contained in [23], with the main observation being that for most biological networks, the *free* species can be organized in the form of *birth and death cascades*, depending on the minimum number of reactions that the species requires to be created from nothing (denoted by \emptyset) or get reduced to it. As our examples suggest, for many Systems Biology networks these cascades have a natural correspondence to gene expression or signaling stages in the dynamics (see Section 6). Combining this cascade construction along with matrix methods that are used in the finite state-space case, our method provably determines all the irreducible state-spaces for a large class of networks satisfying some biologically reasonable criteria. The knowledge of irreducible state-spaces has several important applications which we discuss in Section 2.2.

In comparison to the method in [23], the method we develop in this paper is far more versatile and can analyze a much larger class of networks arising in Systems and Synthetic Biology. The method in [23] has several shortcomings that severely restrict its range of applicability. Most prominently, it only applies to networks where there are no mixed-sign conservation relations (i.e. no *restricted* species) and all the *free* species can be arranged in birth-cascades. These requirements are often violated by important networks (see Section 6) and this motivated us to devise the method presented in this paper, which can handle all types of conservation relations and intricate interactions among the dynamics of *bounded*, *free* and *restricted* species. The main novelty of the approach we develop in this paper is in the elaborate construction of *tree-like* structures to define reachability relationships between birth and death cascades (see Section 4.4). It is shown that the *leaves* of these cascade trees contain all the information about irreducible state-spaces.

This paper is organized as follows. In Section 2 we introduce some preliminary concepts that will be used throughout the paper. The method for computing the decomposed form of state-space is explained in Section 3 and the process for finding all the irreducible state-spaces is described in Section 4. The algorithms for implementing these procedures are provided in Section 5. In Section 6 we illustrate the applicability and

usefulness of our methods for state-space analysis by considering several examples from Systems Biology. In particular we discuss how this state-space analysis can contribute towards our understanding of the underlying biological process. In Section 7 we conclude and the detailed proof of our main result, Theorem 4.4 is given in the Appendix.

Notation

Most of the notation used in this paper has already been introduced in Section 1. However some additional notation and clarifications are needed which we mention now. For any set A , we denote its cardinality by $|A|$. The vector of all zeros in any dimension is denoted by $\mathbf{0}$. Similarly in any dimension d , the i -th standard basis is denoted by e_i and it is the vector whose i -th entry is 1 while the rest are zeros. The identity matrix is denoted by \mathbf{I} . For any $m \times n$ matrix M and any $k \leq l \leq m$, the $(l - k + 1) \times n$ matrix formed by rows $(k + 1), (k + 2), \dots, l$ of matrix M is denoted by $\text{Proj}(M, k, l)$. If v_1, \dots, v_n are the columns of M then for any $A \subset \mathbb{R}$, the set $\text{Colspan}_A(M)$ stands for

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

The dimension of any vector space V is denoted by $\dim(V)$ and this vector space is called trivial if $\dim(V) = 0$. While multiplying a matrix with a vector we always regard the vector as a column vector. All inequalities involving vectors or matrices must be interpreted componentwise.

2 Preliminaries

Consider the reaction network \mathcal{N} described in Section 1 with d species and K reactions of the form (1.1). Let $\Lambda : \mathbb{N}_0^d \rightarrow \mathbb{R}_+^K$ be the *propensity* map given by

$$\Lambda(x) = (\lambda_1(x), \dots, \lambda_K(x)). \quad (2.16)$$

We represent this network by the triplet $\mathcal{N} = (\mathcal{V}, \mathcal{O}, \Lambda)$, where \mathcal{V} and \mathcal{O} are two $d \times K$ matrices whose entries at row i and column k are given by ν_{ik} and ρ_{ik} respectively. We call these matrices as the *reactant* matrix and the *product* matrix respectively, because they tabulate the number of molecules of each species that are created and removed by each of the reactions. Note that the stoichiometry matrix S for this network is simply $S = (\mathcal{O} - \mathcal{V})$. The left nullspace $\mathcal{L}(S)$ of this matrix is the space of all conservation relations for the network and as explained in Section 1, the constraints imposed by these relations can be described by some conservation data (Γ, c) , and with this data at hand the designated state-space \mathcal{E}_0 is given by (1.11). In the rest of this section we present some concepts and assumptions associated with the reactions networks we consider in this paper.

2.1 The reachability relation and irreducible state-spaces

For any $x, y \in \mathcal{E}_0$ let $p_x(t, y)$ (see (1.3)) be 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

$$x \xrightarrow{\mathcal{N}} y. \quad (2.17)$$

This relation does not depend on the time-value t (see Theorem 3.2.1 in [18]) and it 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 [24]). We say that two states $x, y \in \mathcal{E}_0$ *communicate* if we have both $x \xrightarrow{\mathcal{N}} y$ and $y \xrightarrow{\mathcal{N}} x$, and this relation is denoted by

$$x \xleftrightarrow{\mathcal{N}} y.$$

It is known that $\xrightarrow{\mathcal{N}}$ is an *equivalence* relation on \mathcal{E}_0 and it partitions the set \mathcal{E}_0 into disjoint equivalence classes which are referred to as *communication classes* in the Markov chain literature (see [18]). The communication classes can be further classified as *closed* or *open*. A communication class $\mathcal{E} \subset \mathcal{E}_0$ is called *closed* if and only if for any $x \in \mathcal{E}$ and $y \in \mathcal{E}_0$, if $x \xrightarrow{\mathcal{N}} y$ then $y \in \mathcal{E}$. By definition a closed communication class \mathcal{E} needs to be closed under the reaction dynamics (i.e. (1.2) holds with \mathbb{N}_0^d replaced with \mathcal{E}), and so it is a valid state-space for the underlying CTMC. Moreover as all the states in \mathcal{E} are reachable from each other, \mathcal{E} is an *irreducible* state-space for the CTMC. Therefore we refer to such a set \mathcal{E} as an irreducible state-space in the paper.

When the state-space \mathcal{E}_0 is finite, then all the irreducible state-spaces correspond exactly to all the positive-recurrent classes for the underlying Markov chain (see [18]), and starting from any initial state in \mathcal{E}_0 , the CTMC will eventually get *trapped* in one of these disjoint irreducible state-spaces. The same holds true for infinite state-spaces if the existence of stationary distributions on all the irreducible state-spaces can be guaranteed. This can be done using methods developed by Meyn and Tweedie [25, 26]. In particular, Theorem 4.5 in [26] shows that stationary distributions will exist if one can find a *Foster-Lyapunov* function $V : \mathcal{E}_0 \rightarrow \mathbb{R}_+$ satisfying $V(x) \rightarrow \infty$ as $\|x\| \rightarrow \infty$, such that the \mathbb{R}_+ -valued process $(V(X(t)))_{t \geq 0}$ experiences a negative drift outside some compact (finite) set $C \subset \mathcal{E}_0$. In a recent paper [27] we develop a computational framework for constructing such Foster-Lyapunov functions for a large class of biochemical reaction networks, which includes several well-known examples from Systems and Synthetic Biology.

The problem of finding closed communication classes for stochastic reaction networks is quite challenging, as establishing the reachability (2.17) between any two states $x, y \in \mathcal{E}_0$ is tantamount to showing that there exists a sequence of n reactions $k_1, \dots, k_n \in \{1, \dots, K\}$ such that $y = x + \sum_{i=1}^n \zeta_{k_i}$ and $\lambda_{k_j}(z_j) > 0$ for each $j = 1, \dots, n$, where $z_j = x + \sum_{i=1}^{j-1} \zeta_{k_i}$. These conditions ensure that starting from state x , firing of reactions k_1, \dots, k_n in this order, takes the state to y , and this firing of reactions is a positive-probability event because at all the *intermediate* states (z_j -s), the propensity $\lambda_{k_j}(z_j)$ for the next reaction in this sequence is positive. When the state-space \mathcal{E}_0 is finite, matrix methods (see [17]) can be used for computing the closed communication classes, without having to explicitly find any positive-probability reaction sequences between states. In this paper we provide an extension of this approach that can handle a large class of biological networks with infinite state-spaces (see Section 4). One of the main difficulty that we have to deal with is that some reaction channels may *switch-off* at certain states, due to their propensities being zero at those states, and hence the set of possible transition directions is not the same for all the states in \mathcal{E}_0 . To account for this *switching-off* of reactions we need to impose some conditions on the propensity functions, as we mention in Section 2.3. However, first we discuss the applications of our computational method for identifying all the closed communication classes or irreducible state-spaces for stochastic reaction networks with infinite state-spaces.

2.2 Applications

It is interesting to note that for many networks, simply knowing the irreducible state-spaces $\mathcal{E}_1, \mathcal{E}_2, \dots$ allows one to compute the corresponding stationary distributions π_1, π_2, \dots , either analytically (see [28] and Example 6.5) or numerically [29]. Knowledge of these irreducible state-spaces has many other important consequences. For example, if the initial state $X(0)$ of the stochastic reaction dynamics $(X(t))_{t \geq 0}$ lies in the irreducible state-space \mathcal{E}_q , then the ideal state-space \mathcal{E} (see (1.5)) consisting of *only* the reachable states, coincides exactly with \mathcal{E}_q . Furthermore due to Theorem 1.10.2 in [18], for any bounded real-valued function f on \mathcal{E}_q we have

$$\lim_{t \rightarrow \infty} \mathbb{E}(f(X(t))) = \sum_{y \in \mathcal{E}_q} f(y) \pi_q(y) \quad (2.18)$$

and the following limit holds with probability 1

$$\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t f(X(s)) ds = \sum_{y \in \mathcal{E}_q} f(y) \pi_q(y). \quad (2.19)$$

Using (2.18) one can show that various statistical quantities (means, variances, covariances etc.) associated with the stochastic reaction dynamics converge to their *steady state* values as $t \rightarrow \infty$ (see [27]), which

is helpful in the design of controllers that can robustly steer the mean copy-number of some species to specific reference values [30]. Relation (2.19) 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 techniques such as *Flow-Cytometry* and *Time-Lapse Microscopy* in the study of isogenic cell populations (see [27]). One can also use (2.19) to *speed-up* the estimation of the stationary distribution π_q using Monte Carlo simulations. Moreover the knowledge about the exact support of π_q (viz. \mathcal{E}_q) can be used to sample from this stationary distribution more efficiently. Commonly we find that for networks arising in Systems Biology there is only one irreducible state-space \mathcal{E}_1 within \mathcal{E}_0 and hence the simplex Σ collapses to a unique stationary distribution π_1 (see Section 6). In such a scenario the underlying CTMC is *ergodic*, and relations (2.18) and (2.19) hold for any bounded real-valued function f on the full state-space \mathcal{E}_0 and for any initial condition $X(0) \in \mathcal{E}_0$. Checking ergodicity of networks is important for many applications. For example, under ergodicity one can apply tools from *Transition Path Theory* [31] to study the topology of networks by analyzing the statistical properties of trajectories that *flow* between two subsets of the irreducible state-space. Another important application area where checking ergodicity is crucial, is for *speeding-up* the stochastic simulations of multiscale networks, which have reactions firing at multiple timescales [32]. Such networks are common in Systems Biology and it is known that their exact stochastic simulation, using Gillespie’s SSA for example, is highly cumbersome, because most of the simulation time is spent in generating the *fast* reaction events. To circumvent these problems approximate simulation approaches have been developed [33, 34, 35] that apply the *quasi-stationary assumption* (QSA) on *fast* ergodic subnetworks, by supposing that their stochastic dynamics relaxes to stationarity between subsequent reactions at the *slower* timescales. These *fast* subnetworks can change their structure, along with their ergodic properties, depending on the states of the *slow* variables which may determine the set of available *fast* reactions. In this context, our results provide a way for automated discovery of *fast* ergodic subnetworks, during the simulation run, and aid the correct application of QSA. We illustrate this point through an example in Section 6.4. This example also shows how *restricted* species can arise naturally when a *fast* subnetwork within the bigger network, is considered in isolation.

2.3 Conditions on the propensity functions

To facilitate the search for irreducible state-spaces, we shall assume that the network $\mathcal{N} = (\mathcal{V}, \mathcal{O}, \Lambda)$ satisfies the following:

Assumption 2.1 *For each reaction $k = 1, \dots, K$ and each $x = (x_1, \dots, x_d) \in \mathcal{E}_0$, we have $\lambda_k(x) > 0$ if and only if $x_i \geq \nu_{ik}$ for every $i = 1, \dots, d$.*

In other words, at state x , reaction k has a positive probability of firing if and 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}). The “only if” part of this condition is almost 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 the propensity function for a reaction is zero even though all the required molecules for having the reaction are present. However such situations do not typically arise for networks in Systems Biology as we now explain.

Observe that Assumption 2.1 is certainly satisfied if we have *mass-action kinetics* [9] where each propensity function $\lambda_k : \mathbb{N}_0^d \rightarrow \mathbb{R}_+$ has the form

$$\lambda_k(x) = \theta_k \prod_{i=1}^d \frac{x_i(x_i - 1) \dots (x_i - \nu_{ik} + 1)}{\nu_{ik}!}, \quad (2.20)$$

for some rate constant $\theta_k > 0$. Apart from mass-action kinetics, networks in Systems Biology generally have propensity functions describing either Michaelis–Menten kinetics for enzyme-substrate interactions or Hill kinetics for ligand-protein binding dynamics [36]. In both these cases, the propensity functions have a rational form $\lambda_k(x) = p_k(x)/q_k(x)$, where the denominator $q_k(x)$ is always positive and the numerator $p_k(x)$ satisfies the criterion in Assumption 2.1. Consequently the network satisfies Assumption 2.1 even though the propensity functions are not of mass-action form (2.20). Furthermore even if a network does not satisfy Assumption 2.1, it can often be modified in such a way that this assumption is satisfied and its dynamics remains the same (see Section 6.1). We end this section with a simple proposition.

Proposition 2.2 Consider a network \mathcal{N} satisfying Assumption 2.1 with conservation data (Γ, c) and state-space \mathcal{E}_0 given by (1.11). Then the reachability relation $\xrightarrow{\mathcal{N}}$ is **positively additive** on \mathcal{E}_0 , which is to say that for any $x, y \in \mathcal{E}_0$ and $z \in \mathbb{N}_0^d$, if $x \xrightarrow{\mathcal{N}} y$ and $(x + z) \in \mathcal{E}_0$ then we also have $(y + z) \in \mathcal{E}_0$ and $(x + z) \xrightarrow{\mathcal{N}} (y + z)$.

Proof. This is a simple consequence of Assumption 2.1 because for any reaction k and state $u \in \mathcal{E}_0$, if $\lambda_k(u) > 0$ then $\lambda_k(u + z) > 0$ for any $z \in \mathbb{N}_0^d$. Therefore the positive-probability reaction sequence that leads the state from x to y , also serves as a positive-probability reaction sequence that takes the state from $(x + z)$ to $(y + z)$ (recall the discussion in Section 2.1). \square

2.4 Inverse of a reaction network

We now define the inverse $\mathcal{N}_{\text{inv}} = (\mathcal{V}_{\text{inv}}, \mathcal{O}_{\text{inv}}, \Lambda_{\text{inv}})$ of the reaction network $\mathcal{N} = (\mathcal{V}, \mathcal{O}, \Lambda)$, which is obtained by *flipping* the arrows in (1.1). In other words, the K reactions in \mathcal{N}_{inv} are given by

$$\sum_{i=1}^d \rho_{ik} \mathbf{S}_i \longrightarrow \sum_{i=1}^d \nu_{ik} \mathbf{S}_i, \quad k = 1, \dots, K. \quad (2.21)$$

Hence $\mathcal{V}_{\text{inv}} = \mathcal{O}$, $\mathcal{O}_{\text{inv}} = \mathcal{V}$ and we define the propensity map

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

by letting each $\lambda_{k,\text{inv}}$ to have the mass-action form (see (2.20)) with the rate constant $\theta_k = 1$. Observe that the stoichiometry matrix S_{inv} for \mathcal{N}_{inv} is simply the negative of the stoichiometry matrix for \mathcal{N} , and so the space of conservation relations (1.7) as well as the state-space \mathcal{E}_0 (1.11) remain the same for both networks.

By construction, \mathcal{N}_{inv} always satisfies Assumption 2.1 and if the original network \mathcal{N} also satisfies this assumption then we have the following correspondence between the reachability relations induced by the two networks

$$x \xrightarrow{\mathcal{N}} y \quad \text{if and only if} \quad y \xrightarrow{\mathcal{N}_{\text{inv}}} x, \quad \text{for any} \quad x, y \in \mathcal{E}_0. \quad (2.22)$$

To see this correspondence suppose that under network \mathcal{N} , the dynamics can reach y from x in one reaction step. In such a scenario for some reaction k we have $\lambda_k(x) > 0$ and $y = x + \zeta_k$. Since $x_i \geq \nu_{ik}$ and $\zeta_{ik} = (\rho_{ik} - \nu_{ik})$, we have $y_i \geq \rho_{ik}$ for each $i = 1, \dots, d$, which ensures that $\lambda_{k,\text{inv}}(y) > 0$ and so under the inverse network \mathcal{N}_{inv} , the 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 (see Section 2.1) one can conclude that (2.22) holds.

2.5 Reaction network under a permutation

Our description of the stochastic model for network $\mathcal{N} = (\mathcal{V}, \mathcal{O}, \Lambda)$, stipulates that if the state is $x = (x_1, \dots, x_d)$ then x_i denotes the copy-number of species i (viz. \mathbf{S}_i). However in order to simplify the representation for state-space we would often need to redefine the correspondence between the species and the location of their copy-numbers in the state vector. This can be conveniently done by *permuting* the original network to obtain a dynamically equivalent network, as we describe below.

We denote the set of all species labels by $\mathcal{D} = \{1, \dots, d\}$. Let $\sigma : \mathcal{D} \rightarrow \mathcal{D}$ be any permutation (one-to-one and onto) map. Let P_σ be the $d \times d$ permutation matrix given by

$$P_\sigma = \text{Col}(e_{\sigma^{-1}(1)}, \dots, e_{\sigma^{-1}(d)}), \quad (2.23)$$

where e_1, \dots, e_d are the standard basis vectors in \mathbb{R}^d and let σ^{-1} denote the inverse of map σ . Let $\mathcal{V}^\sigma = P_\sigma \mathcal{V}$, $\mathcal{O}^\sigma = P_\sigma \mathcal{O}$ and the *propensity* map $\Lambda^\sigma : \mathbb{N}_0^d \rightarrow \mathbb{R}_+^K$ be as in (2.16) with each λ_k replaced by λ_k^σ given by

$$\lambda_k^\sigma(x) = \lambda_k(P_\sigma^T x) \quad \text{for each} \quad k = 1, \dots, K, \quad (2.24)$$

where $P_\sigma^T = P_{\sigma^{-1}}$ denotes the transpose of matrix P_σ . We define the permuted network as $\mathcal{N}^\sigma = (\mathcal{V}^\sigma, \mathcal{O}^\sigma, \Lambda^\sigma)$. One can see that \mathcal{N}^σ is dynamically equivalent to \mathcal{N} in the following sense. If $(X(t))_{t \geq 0}$ represents the stochastic reaction dynamics for network \mathcal{N} then $(X^\sigma(t))_{t \geq 0}$ defined by

$$X^\sigma(t) = P_\sigma X(t) \quad \text{for all } t \geq 0, \quad (2.25)$$

represents the dynamics for the permuted network \mathcal{N}^σ . In other words, if $X^\sigma(t) = (x_1, \dots, x_d)$ then for each $i = 1, \dots, d$, x_i denotes the copy-number of species $\sigma(i)$ at time t . Therefore the appropriate conservation data and state-space for network \mathcal{N}^σ are (Γ_σ, c) and \mathcal{E}_0^σ respectively, where

$$\Gamma_\sigma = P_\sigma \Gamma \quad \text{and} \quad \mathcal{E}_0^\sigma = P_\sigma \mathcal{E}_0 := \{P_\sigma x : x \in \mathcal{E}_0\}. \quad (2.26)$$

One can also see that if network \mathcal{N} satisfies Assumption 2.1, then the same holds for the permuted network \mathcal{N}^σ . Due to the dynamical equivalence (2.25) we have the following proposition.

Proposition 2.3 *A set $\mathcal{E}^\sigma \subset \mathcal{E}_0^\sigma$ is an irreducible state-space for network \mathcal{N}^σ if and only if the set $\mathcal{E} = P_\sigma^T \mathcal{E}^\sigma \subset \mathcal{E}_0$ is an irreducible state-space for network \mathcal{N} .*

3 Computing the decomposed form of state-space

The aim of this section is to provide a procedure to obtain the decomposed form (1.13) for state-space \mathcal{E}_0 for a network $\mathcal{N} = (\mathcal{V}, \mathcal{O}, \Lambda)$ with conservation data (Γ, c) . For this purpose, we may need to permute this network according to some permutation σ and work with the equivalent network \mathcal{N}^σ with state-space \mathcal{E}_0^σ (see Section 2.5).

The decomposed form can be constructed in two simple steps. Firstly by scanning the space of all semi-positive conservation relations, the *bounded* species are identified and their appropriate finite state-space \mathcal{E}_b is found (Section 3.1). Secondly the rest of the species are classified as *free* or *restricted* depending on any mixed-sign conservation relations between them. Moreover, the affine function ϕ (see (1.14)), which gives the static relationship between these two sets of species, is determined (Section 3.2). The detailed algorithm for performing state-space decomposition is presented in Section 5.

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

Note that the space of all conservation relations can be expressed as $\mathcal{L}(S) = \{\Gamma\alpha : \alpha \in \mathbb{R}^n\}$. Suppose that $\gamma = \Gamma\alpha$ is a semi-positive conservation relation for some $\alpha \in \mathbb{R}^n$. From (1.8) and (1.10), one can see that for any $i \in \text{supp}(\gamma)$ we have

$$0 \leq X_i(t) \leq \frac{\langle c, \alpha \rangle}{\gamma_i} \quad \text{for all } t \geq 0, \quad (3.27)$$

where $X_i(t)$ denotes the copy-number of species i at time t . This shows that species i is *bounded*, but since it may be involved in several semi-positive conservation relations, the upper-bound $\langle c, \alpha \rangle / \gamma_i$ for its copy-numbers may not be *sharp*. To systematically account for all these relations and obtain a sharper upper-bound b_i we solve the following Linear Program (LP):

$$\begin{aligned} b_i &= \min_{\alpha \in \mathbb{R}^n} \langle c, \alpha \rangle \\ \text{subject to} \quad & \Gamma\alpha \geq \mathbf{0} \quad \text{and} \quad \langle e_i, \Gamma\alpha \rangle = 1, \end{aligned} \quad (3.28)$$

where e_i is the i -th standard basis vector in \mathbb{R}^d . If the feasible region of this LP is empty, then $b_i = \infty$ and species i is not involved in any semi-positive conservation relation. Henceforth we partition the set of *all* species $\mathcal{D} = \{1, \dots, d\}$ into the set of *bounded* species $\mathcal{D}_b = \{i \in \mathcal{D} : b_i < \infty\}$ and the set of *unbounded* species $\mathcal{D}_u = \{i \in \mathcal{D} : b_i = \infty\}$. Let $d_b = |\mathcal{D}_b|$ and $d_u = |\mathcal{D}_u| = d - d_b$ be the cardinalities of these two sets.

Choose a permutation map $\sigma_1 : \mathcal{D} \rightarrow \mathcal{D}$ satisfying

$$\sigma_1(l) \in \begin{cases} \mathcal{D}_b & \text{for } l = 1, \dots, d_b \\ \mathcal{D}_u & \text{for } l = (d_b + 1), \dots, d \end{cases} \quad (3.29)$$

and consider the reaction dynamics of the network \mathcal{N}^σ under permutation $\sigma = \sigma_1$ (see Section 2.5). Now the entries in rows $1, \dots, d_b$ of the state vectors will contain the copy-numbers of *bounded* species in \mathcal{D}_b . These copy-numbers, arranged as vectors in $\mathbb{N}_0^{d_b}$, will always lie in the finite rectangular set

$$\mathcal{R}_b^\sigma = \{(x_1, \dots, x_{d_b}) \in \mathbb{N}_0^{d_b} : 0 \leq x_l \leq b_{\sigma(l)} \text{ for each } l = 1, \dots, d_b\}, \quad (3.30)$$

but all the elements in this set may not be reachable from each other due to conservation relations among *bounded* species (see Section 6). We deal with these conservations relations now.

Let S_σ be the stoichiometry matrix for network \mathcal{N}^σ and let (Γ_σ, c) be its conservation data (see (2.26)). Setting $S_\sigma^b := \text{Proj}(S_\sigma, 1, d_b)$, the conservation relations among the *bounded* species are given by nonzero vectors in its left nullspace $\mathcal{L}(S_\sigma^b)$. Suppose $n_b = \dim(\mathcal{L}(S_\sigma^b)) \geq 1$ and let $\{\hat{\gamma}_1, \dots, \hat{\gamma}_{n_b}\}$ denote a *basis* for $\mathcal{L}(S_\sigma^b)$. For each $j = 1, \dots, n_b$, let $\bar{\gamma}_j = (\hat{\gamma}_j, \mathbf{0}) \in \mathbb{R}^d$ and set $\hat{c}_j = \langle \alpha_j, c \rangle$ where $\alpha_j \in \mathbb{R}^n$ is the unique solution of the linear-system $\Gamma_\sigma \alpha_j = \bar{\gamma}_j$. This solution α_j is unique because the columns of Γ_σ are independent and they span $\mathcal{L}(S_\sigma)$ which contains $\bar{\gamma}_j$. For the permuted network \mathcal{N}^σ , the state vectors for all the *bounded* species in \mathcal{D}_b will always lie in the finite set

$$\mathcal{E}_b^\sigma = \{x \in \mathcal{R}_b^\sigma : \langle \hat{\gamma}_j, x \rangle = \hat{c}_j \text{ for each } j = 1, \dots, n_b\}. \quad (3.31)$$

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. Letting $S_\sigma^u = \text{Proj}(S_\sigma, d_b + 1, d)$, we define the number of *free* species d_f and the number of *restricted* species d_r by

$$d_f := \text{Rank}(S_\sigma^u) \quad \text{and} \quad d_r := d_u - d_f = d - d_b - d_f. \quad (3.32)$$

Note 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). Observe that d_r is the dimension of the left nullspace $\mathcal{L}(S_\sigma^u)$ of matrix S_σ^u , which corresponds to the space of all conservation relations among the *unbounded* species. Any nonzero vector $\gamma \in \mathcal{L}(S_\sigma^u)$ must necessarily be a mixed-sign conservation relation, because otherwise the species in its support set (1.9) would be *bounded* which is not the case. Recall that $n = \dim(\mathcal{L}(S_\sigma))$ is the total number of conservation relations while n_b are the number of conservation relations among the *bounded* species. Hence we must have $n \geq n_b + d_r$. As it turns out, usually for biological networks (see Section 6) this inequality is strict (i.e. $n = n_b + d_r$) which is same as saying that the following assumption holds.

Assumption 3.1 *There are no conservation relations involving both the bounded and the free/restricted species i.e.*

$$\text{Rank}(S_\sigma) = \text{Rank}(S_\sigma^b) + \text{Rank}(S_\sigma^u),$$

where $S_\sigma^b = \text{Proj}(S_\sigma, 1, d_b)$ and $S_\sigma^u = \text{Proj}(S_\sigma, d_b + 1, d)$.

This assumption will be required to hold from now on. Suppose that *restricted* species exist (i.e. $d_r \geq 1$). Hence, the space $\mathcal{L}(S_\sigma^u)$ is nontrivial. Let $\{\delta'_1, \dots, \delta'_{d_r}\}$ be a basis for $\mathcal{L}(S_\sigma^u)$. For any subset $I = \{i_1, \dots, i_{d_f}\} \subset \{1, \dots, d_u\}$ with $|I| = d_f$ elements, let A_I be the $d_u \times d_u$ matrix given by

$$A_I = \text{Col}(e_{i_1}, \dots, e_{i_{d_f}}, \delta'_1, \dots, \delta'_{d_r}), \quad (3.33)$$

where e_i -s are the standard basis vectors in \mathbb{R}^{d_u} . Define another set

$$\mathcal{I}_f = \{I \subset \{1, \dots, d_u\} : |I| = d_f \text{ and } \text{Rank}(A_I) = d_u\}. \quad (3.34)$$

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 $\{1, \dots, d_u\}$. We define the set \mathcal{D}_f of *free* species and the set \mathcal{D}_r of *restricted* species as $\mathcal{D}_f = \{\sigma(d_b + i) : i \in I\}$ and $\mathcal{D}_r = \{\sigma(d_b + i) : i \in I^c\}$. These two sets partition the set \mathcal{D}_u . We now choose another permutation map $\sigma_2 : \mathcal{D} \rightarrow \mathcal{D}$ satisfying

$$\sigma_2(l) = \sigma_1(l) \quad \text{for } l = 1, \dots, d_b \quad \text{and} \quad \sigma_2(l) \in \begin{cases} \mathcal{D}_f & \text{for } l = (d_b + 1), \dots, (d_b + d_f) \\ \mathcal{D}_r & \text{for } l = (d_b + d_f + 1), \dots, d. \end{cases} \quad (3.35)$$

Let \mathcal{N}^σ be the network under permutation $\sigma = \sigma_2$ (see Section 2.5), and let S_σ and (Γ_σ, c) be its stoichiometry matrix and conservation data respectively. For each $i = 1, \dots, d_r$, the vector $\bar{\delta}_i = (\mathbf{0}, \delta'_i) \in \mathbb{R}^d$ belongs to $\mathcal{L}(S_\sigma)$ and hence the vector $\hat{\delta}_i = P_{\sigma_2} P_{\sigma_1}^T \bar{\delta}_i$ belongs to $\mathcal{L}(S_\sigma)$. Since the permutations σ_1 and σ_2 are identical on $\{1, \dots, d_b\}$, each $\hat{\delta}_i$ must have the form $\hat{\delta}_i = (\mathbf{0}, \delta_i^{(1)}, \delta_i^{(2)})$ for some vectors $\delta_i^{(1)} \in \mathbb{R}^{d_f}$ and $\delta_i^{(2)} \in \mathbb{R}^{d_r}$. Define a $d_f \times d_r$ matrix $\Delta_1 = \text{Col}(\delta_1^{(1)}, \dots, \delta_{d_r}^{(1)})$ and a $d_r \times d_r$ matrix $\Delta_2 = \text{Col}(\delta_1^{(2)}, \dots, \delta_{d_r}^{(2)})$. Observe that if A_I is the matrix given by (3.33), then there exists a $d_u \times d_u$ permutation matrix Q such matrix QA_I has the form

$$QA_I = \begin{bmatrix} \mathbf{I} & \Delta_1 \\ \mathbf{0} & \Delta_2 \end{bmatrix}, \quad (3.36)$$

where \mathbf{I} is the $d_f \times d_f$ identity matrix and $\mathbf{0}$ 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 $S_\sigma^b = \text{Proj}(S_\sigma, 1, d_b)$, $S_\sigma^u = \text{Proj}(S_\sigma, d_b + 1, d)$, $S_\sigma^f = \text{Proj}(S_\sigma^u, 1, d_f)$ and $S_\sigma^r = \text{Proj}(S_\sigma^u, d_f + 1, d_f + d_r)$. For each $i = 1, \dots, d_r$, the vector $\hat{\delta}_i = (\mathbf{0}, \delta_i^{(1)}, \delta_i^{(2)})$ belongs to $\mathcal{L}(S_\sigma)$, where $\mathbf{0}$ denotes the vector of zeros in \mathbb{R}^{d_b} . Hence we must have $\Delta_1^T S_\sigma^f = -\Delta_2^T S_\sigma^r$ which allows us to write

$$S_\sigma^r = -(\Delta_2^T)^{-1} \Delta_1^T S_\sigma^f = -(\Delta_2^{-1})^T \Delta_1^T S_\sigma^f. \quad (3.37)$$

This also shows that $\text{Rank}(S_\sigma^f) = d_f$ because $d_f = \text{Rank}(S_\sigma^u)$ (see (3.32)) and S_σ^u has the form

$$S_\sigma^u = \begin{bmatrix} S_\sigma^f \\ S_\sigma^r \end{bmatrix} = \begin{bmatrix} \mathbf{I} \\ -(\Delta_2^{-1})^T \Delta_1^T \end{bmatrix} S_\sigma^f.$$

The state-space for network \mathcal{N}^σ is $\mathcal{E}_0^\sigma = P_\sigma \mathcal{E}_0 = \{x \in \mathbb{N}_0^d : \Gamma_\sigma x = c\}$ (see (2.26)). Any element in this set can be expressed as $x = (x_b, x_f, x_r)$ where $x_b \in \mathcal{E}_b^\sigma$, $x_f \in \mathbb{N}_0^{d_f}$ and $x_r \in \mathbb{N}_0^{d_r}$. Since the columns of Γ_σ span $\mathcal{L}(S_\sigma)$, for each $i = 1, \dots, d_r$, there exists a unique $n \times d_r$ matrix M such that

$$\Gamma_\sigma M = \text{Col}(\hat{\delta}_1, \dots, \hat{\delta}_{d_r}).$$

Since Assumption 3.1 holds, the last $(d_f + d_r)$ rows of the condition $\Gamma_\sigma x = c$ yield $\Delta_1^T x_f + \Delta_2^T x_r = \hat{c}$ for $\hat{c} = M^T c$. This shows that $x_r = \phi(x_f)$, where the affine map $\phi : \mathbb{R}^{d_f} \rightarrow \mathbb{R}^{d_r}$ is defined by

$$\phi(x) = (\Delta_2^T)^{-1} \hat{c} - (\Delta_2^T)^{-1} \Delta_1^T x. \quad (3.38)$$

This analysis proves the following proposition.

Proposition 3.2 *Suppose Assumption 3.1 holds. Then state-space \mathcal{E}_0^σ for network \mathcal{N}^σ can be expressed as*

$$\hat{\mathcal{E}}_0^\sigma = \mathcal{E}_b^\sigma \times \Phi, \quad (3.39)$$

where Φ is the graph (1.14) of function ϕ defined by (3.38).

Note that this state-space decomposition result only depends on the reaction stoichiometries but not on their propensities. Indeed the propensity functions can be completely general as long as they satisfy the basic assumption (1.2) which ensures that the dynamics is contained in the positive orthant. We end this section with an important remark.

Remark 3.3 *Note that the classification of unbounded species into free and restricted species depends on the set I which can be chosen to be any element in the set \mathcal{I}_f given by (3.34). This flexibility will be useful in the next section.*

4 Identifying the irreducible state-spaces

In this section we shall assume that network $\mathcal{N} = (\mathcal{V}, \mathcal{O}, \Lambda)$ satisfies Assumption 2.1. As a consequence, the permuted network $\mathcal{N}^\sigma = (\mathcal{V}^\sigma, \mathcal{O}^\sigma, \Lambda^\sigma)$ also satisfies this assumption, which is a property that will

play a crucial role in our search for irreducible state-spaces within the infinite state-space \mathcal{E}_0^σ . Recall that these state-spaces are the closed communication classes for relation $\xleftrightarrow{\mathcal{N}^\sigma}$ on \mathcal{E}_0^σ (see Section 2.1). From the discussion in Section 3.2 it is immediate that *restricted* species have no independent dynamics of their own and they essentially *mimic* the *free* species according to the affine map 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. We now describe this step formally.

4.1 Network reduction by elimination of *restricted* species

We construct a “reduced” network $\tilde{\mathcal{N}}^\sigma$ with $(d_b + d_f)$ species in the set $\mathcal{D}_b \cup \mathcal{D}_f$ where $\mathcal{D}_b = \{\sigma(1), \dots, \sigma(d_b)\}$ and $\mathcal{D}_f = \{\sigma(d_b + 1), \dots, \sigma(d_b + d_f)\}$. For each reaction $k = 1, \dots, K$, let its propensity function $\tilde{\lambda}_k^\sigma : \mathbb{N}_0^{d_b} \times \mathbb{N}_0^{d_f} \rightarrow \mathbb{R}_+$ be given by

$$\tilde{\lambda}_k^\sigma(x_b, x_f) = \lambda_k^\sigma(x_b, x_f, \phi(x_f)) \quad (4.40)$$

where $x_b \in \mathbb{N}_0^{d_b}$, $x_f \in \mathbb{N}_0^{d_f}$ and the function λ_k^σ is defined by (2.24). Let the *propensity* map $\tilde{\Lambda}^\sigma : \mathbb{N}_0^{d_b + d_f} \rightarrow \mathbb{R}_+^K$ be as in (2.16) with each λ_k replaced by $\tilde{\lambda}_k^\sigma$. Setting $\tilde{\mathcal{V}}^\sigma = \text{Proj}(\mathcal{V}^\sigma, 1, d_b + d_f)$ and $\tilde{\mathcal{O}}^\sigma = \text{Proj}(\mathcal{O}^\sigma, 1, d_b + d_f)$, we define the reduced network as $\tilde{\mathcal{N}}^\sigma = (\tilde{\mathcal{V}}^\sigma, \tilde{\mathcal{O}}^\sigma, \tilde{\Lambda}^\sigma)$.

Unfortunately the reduced network $\tilde{\mathcal{N}}^\sigma$ may not satisfy Assumption 2.1 even though the original network does (see Example 6.4). We need to avoid this problem because our approach requires this property. Fortunately for most biological examples this can be done by exploiting the flexibility in the choice of set I (see Remark 3.3), which classifies 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. Hence the irreducible state-spaces for the original network can be found with any I chosen as per our convenience. We sequentially examine each element in the finite set \mathcal{I}_f (see (3.34)) until we either find a I for which the affine function ϕ satisfies the following assumption or we exhaust all the possibilities (see Algorithm 1).

Assumption 4.1 Consider an affine map $f : \mathbb{R}^{d_f} \rightarrow \mathbb{R}^{d_r}$

$$f(x) = F_0 + F_1 x,$$

where F_0 is a vector in \mathbb{R}^{d_r} and F_1 is a $d_r \times d_f$ matrix. We say that this map is compatible with network $\mathcal{N}^\sigma = (\mathcal{V}^\sigma, \mathcal{O}^\sigma, \Lambda^\sigma)$ if F_0 and F_1 have all the entries in \mathbb{N}_0 and the matrix inequality

$$F_0 \mathbf{1}^T + F_1 \mathcal{V}_f^\sigma \geq \mathcal{V}_r^\sigma \quad (4.41)$$

holds, where $\mathcal{V}_f^\sigma = \text{Proj}(\mathcal{V}^\sigma, d_b + 1, d_b + d_f)$, $\mathcal{V}_r^\sigma = \text{Proj}(\mathcal{V}^\sigma, d_b + d_f + 1, d)$ and $\mathbf{1}$ is the vector of all ones in \mathbb{R}^K .

If the affine map ϕ satisfies this compatibility condition, then the network reduction will automatically satisfy Assumption 2.1. To see this note that (4.41) implies that for any reaction k , if $x_f \geq \nu_{k,f}$ then $\phi(x_f) \geq \nu_{r,f}$, where $\nu_{k,f}$ ($\nu_{r,f}$) denotes the k -th column of matrix \mathcal{V}_f^σ (\mathcal{V}_r^σ). Therefore the reduced network $\tilde{\mathcal{N}}^\sigma$ satisfies Assumption 2.1 because

$$\begin{aligned} \tilde{\lambda}_k^\sigma(x_b, x_f) > 0 &\iff \lambda_k^\sigma(x_b, x_f, \phi(x_f)) > 0 \\ &\iff (x_b, x_f, \phi(x_f)) \geq (\nu_{k,b}, \nu_{k,f}, \nu_{k,r}) \\ &\iff (x_b, x_f) \geq (\nu_{k,b}, \nu_{k,f}) \quad \text{and} \quad \phi(x_f) \geq \nu_{r,f} \\ &\iff (x_b, x_f) \geq (\nu_{k,b}, \nu_{k,f}), \end{aligned}$$

where $\nu_{k,b}$ is the k -th column of matrix $\text{Proj}(\mathcal{V}^\sigma, 1, d_b)$, \iff denotes “if and only if” and the last relation holds because $\phi(x_f) \geq \nu_{r,f}$ follows from (4.41).

It is easy to check that the reduced network $\tilde{\mathcal{N}}^\sigma$ preserves all the semi-positive conservation relations among the *bounded* species. Therefore if ϕ satisfies Assumption 4.1 the state-space for this reduced network is

$$\tilde{\mathcal{E}}_0^\sigma = \mathcal{E}_b^\sigma \times \mathbb{N}_0^{d_f}.$$

In what follows, we shall identify all the irreducible state-spaces for the communication relation $\xleftrightarrow{\tilde{\mathcal{N}}^\sigma}$ induced by the network $\tilde{\mathcal{N}}^\sigma$ on $\tilde{\mathcal{E}}_0^\sigma$. The next proposition then allows us to recover all the irreducible state-spaces for the original network \mathcal{N}^σ .

Proposition 4.2 *Assume that the affine map ϕ , given by (3.38). Then for any $F \subset \mathcal{E}_b^\sigma$ and $G \subset \mathbb{N}_0^{d_f}$, the set $F \times G$ is an irreducible state-space for relation $\xleftrightarrow{\tilde{\mathcal{N}}^\sigma}$ if and only if the set $F \times \Phi_G$ is an irreducible state-space for relation $\xleftrightarrow{\mathcal{N}^\sigma}$. Here Φ_G is the graph of function ϕ (see (1.14)), restricted to the domain G .*

Proof. The proof follows simply from the construction of the reduced network $\tilde{\mathcal{N}}^\sigma$ and the fact that the dynamics of the *restricted* species is “tied” to the dynamics of the *free* species according to map ϕ . \square

4.2 Networks with only *bounded* species

We first consider the case when there are no *free* species (i.e. $d_f = 0$), and hence the state-space $\tilde{\mathcal{E}}_0^\sigma = \mathcal{E}_b^\sigma$ for network $\tilde{\mathcal{N}}^\sigma$ is finite². In such a situation all the irreducible state-spaces can be found using simple matrix manipulations. We briefly describe this approach in the context of reaction networks and introduce the relevant concepts that will be useful later in the paper.

Let $\mathcal{E}_b^\sigma = \{y_1, \dots, y_{N_b}\}$ be the finite state-space with $N_b = |\mathcal{E}_b^\sigma|$ states. For each reaction k , let $\bar{v}_k^\sigma \in \mathbb{N}_0^{d_b}$ and $\bar{\zeta}_k^\sigma \in \mathbb{Z}^{d_b}$ be the vectors denoting the k -th column of the reactant and the stoichiometry matrices of the network. As network $\tilde{\mathcal{N}}^\sigma$ satisfies Assumption 2.1, only the following reactions have a positive propensity of firing

$$\mathcal{K}_r(y) = \{k = 1, \dots, K : y \geq \bar{v}_k^\sigma\}. \quad (4.42)$$

Define a $N_b \times N_b$ matrix Z by

$$Z_{ij} = \begin{cases} 1 & \text{if } y_j = y_i + \bar{\zeta}_k^\sigma \text{ for some } k \in \mathcal{K}_r(y_i) \\ 0 & \text{otherwise.} \end{cases} \quad (4.43)$$

We can view Z as the *zero-pattern* matrix³ of a finite Markov chain [17] and use it to study reachability and communication relations corresponding to network $\tilde{\mathcal{N}}^\sigma$ (see Section 2.1). Define the $N_b \times N_b$ *reachability* matrix Ω by

$$\Omega = (\mathbf{I} + Z)^{N_b - 1}, \quad (4.44)$$

where \mathbf{I} is the $N_b \times N_b$ identity matrix. For any two states $y_i, y_j \in \mathcal{E}_b^\sigma$, state y_i is reachable from state y_j if and only if the ij -th entry of this reachability matrix is positive (i.e. $\Omega_{ij} > 0$). Therefore based on this matrix we can define the communication relation Θ on \mathcal{E}_b^σ by

$$\Theta = \{(y_i, y_j) \in \mathcal{E}_b^\sigma \times \mathcal{E}_b^\sigma : \Omega_{ij} > 0 \text{ and } \Omega_{ji} > 0\}. \quad (4.45)$$

As we mentioned in Section 2.1, Θ is an equivalence relation on \mathcal{E}_b^σ which partitions this set into distinct equivalence or communication classes. Let $\bar{\Theta} = [\bar{\Theta}_{ij}]$ be the $N_b \times N_b$ matrix representing this relation, i.e. $\bar{\Theta}_{ij} = 1$ if $(y_i, y_j) \in \Theta$ and $\bar{\Theta}_{ij} = 0$ otherwise. One can verify that states y_i and y_j will belong to the same communication class if and only if rows i and j are identical in matrix $\bar{\Theta}$. Let n_c be the number of unique rows of matrix $\bar{\Theta}$ and let U be the $n_c \times N_b$ matrix formed by these rows. Each row i of U corresponds to a distinct communication class made up of those states y_j for which $U_{ij} = 1$. To study the interaction among communication classes we define another $n_c \times n_c$ matrix R by

$$R = UZU^T - \mathbf{I}. \quad (4.46)$$

²Observe that as we have no *free/restricted* species, we do not need to permute the network (i.e. σ can be chosen as the identity permutation) or reduce it to eliminate the *restricted* species (i.e. $\tilde{\mathcal{N}}^\sigma = \mathcal{N}$). However we employ this complicated notation to be consistent with the general case where both *bounded* and *free/restricted* species are present.

³The *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

This matrix captures the reachability relations among the communication classes (see Chapter 8 in [37]). In other words, if $R_{ij} > 0$, then there exists a reaction that takes a state in the i -th communication class to a state in the j -th communication class. If the i -th row of matrix R only consists of zeros then the i -th communication class is *closed* and otherwise it is *open*. Starting from any initial state in \mathcal{E}_b^σ the dynamics of the *bounded* species will eventually get trapped in one of the closed communication classes and these are the **only** irreducible state-spaces for network $\tilde{\mathcal{N}}^\sigma$ without any *free* species.

4.3 Networks with only *unbounded* species

In this section we assume that there are no *bounded* species (i.e. $d_b = 0$) and all the irreducible state-spaces must lie inside the nonnegative integer orthant $\tilde{\mathcal{E}}_0^\sigma = \mathbb{N}_0^{d_f}$, after the *restricted* species have been removed (see Section 4.1). We will find the irreducible state-spaces by adopting a simple scheme that attempts to arrange the *free* species into *birth and death cascades*. We begin by formalizing the notion of birth-cascades for the network $\tilde{\mathcal{N}}^\sigma = (\tilde{\mathcal{V}}^\sigma, \tilde{\mathcal{O}}^\sigma, \tilde{\Lambda}^\sigma)$.

For each reaction k , let $\tilde{\nu}_k^\sigma$ and $\tilde{\rho}_k^\sigma$ be the vectors in $\mathbb{N}_0^{d_f}$, denoting the k -th column of matrices $\tilde{\mathcal{V}}^\sigma$ and $\tilde{\mathcal{O}}^\sigma$ respectively. As $d_b = 0$ the set of all *free* species is $\mathcal{D}_f = \{\sigma(1), \dots, \sigma(d_f)\}$. From now on let $\mathcal{F} = \{1, \dots, d_f\}$ denote the set of addresses of all the *free* species under the map $\sigma(\cdot)$. For any $A \subset \mathcal{F}$, let $\mathcal{D}_f^\sigma(A)$ be the subset of *free* species given by

$$\mathcal{D}_f^\sigma(A) := \{\sigma(i) \in \mathcal{D}_f : i \in A\} \quad (4.47)$$

and let $\mathbb{B}(A) \subset \mathcal{F}$ be defined by

$$\mathbb{B}(A) = \{i \in \mathcal{F} : i \notin A, \text{ supp}(\tilde{\nu}_k^\sigma) \subset A \text{ and } i \in \text{supp}(\tilde{\rho}_k^\sigma) \text{ for some reaction } k = 1, \dots, K\}.$$

The set $\mathcal{D}_f^\sigma(\mathbb{B}(A))$ represents the *free* species that do not belong to the set $\mathcal{D}_f^\sigma(A)$ and are produced by a reaction that only consumes the *free* species in $\mathcal{D}_f^\sigma(A)$.

Using this mapping \mathbb{B} , we define a sequence of subsets of \mathcal{F} , indexed by nonnegative integer *levels* $l = 0, 1, \dots$, as follows: let $G_0 = \emptyset$ and for each $l \geq 1$ let

$$G_l = G_{l-1} \cup \mathbb{B}(G_{l-1}).$$

This sequence of sets $\{G_l : l = 0, 1, 2, \dots\}$ represent the *birth-cascades* for network $\tilde{\mathcal{N}}^\sigma$. At any level l , the set $\mathcal{D}_f^\sigma(G_l)$ consists of all those *free* species that either belong to the previous cascade $\mathcal{D}_f^\sigma(G_{l-1})$ or it is produced by a reaction that only consumes the *free* species in this previous cascade $\mathcal{D}_f^\sigma(G_{l-1})$. In particular for level $l = 1$ the set $\mathcal{D}_f^\sigma(G_1)$ consists of all those *free* species that can be produced *from nothing*, due to reactions of the form $\emptyset \longrightarrow \mathbf{S}$. From this cascade construction one can expect that all the *free* species in each of these birth-cascades, can have arbitrarily high copy-numbers due to the reaction dynamics. This is precisely what we show later in the paper (see Lemma 7.1).

As the number of *free* species is finite, the following number is well-defined

$$l_b = \max\{l \in \mathbb{N} : G_l \neq G_{l-1}\} = \min\{l \in \mathbb{N} : \mathbb{B}(G_{l-1}) = \emptyset\}.$$

The monotonically increasing sequence of sets G_0, G_1, \dots stops growing beyond level l_b and so we have $G_l = G_{l_b}$ for all $l \geq l_b$. We define the *Birth-Cascade Path* (BCP) as the directed graph

$$\emptyset \Rightarrow G_1 \Rightarrow G_2 \Rightarrow \dots \Rightarrow G_{l_b} \quad (4.48)$$

and refer to G_{l_b} as its terminal node. The set of all the *free* species that can be arranged into birth-cascades is

$$\mathcal{B} = \mathcal{D}_f^\sigma(G_{l_b}).$$

We define two reaction-sets as

$$\mathcal{K}_r = \{k = 1, \dots, K : \text{supp}(\tilde{\nu}_k^\sigma) \subset G_{l_b}\} \quad \text{and} \quad \mathcal{K}_p = \{k = 1, \dots, K : \text{supp}(\tilde{\rho}_k^\sigma) \subset G_{l_b}\}. \quad (4.49)$$

The set \mathcal{K}_r (\mathcal{K}_p) consists of all those reactions that do **not** consume (produce) the *free* species that lie outside the set \mathcal{B} . Since G_{l_b} is the terminal node of the BCP we must have that

$$\mathcal{K}_r \subset \mathcal{K}_p, \quad (4.50)$$

which is to say that any reaction that only consumes the *free* species in \mathcal{B} must not produce any *free* species outside this set.

Analogous to birth-cascades we now construct the death-cascades for network $\tilde{\mathcal{N}}^\sigma$, but we restrict our attention to only the species in \mathcal{B} and only the reactions in \mathcal{K}_r . Let $\tilde{\mathcal{N}}^\sigma(\mathcal{B})$ be the network restricted to only these species and reactions and let $\tilde{\mathcal{N}}_{\text{inv}}^\sigma(\mathcal{B})$ be its inverse (see Section 2.4). We construct the birth-cascades for this inverse network $\tilde{\mathcal{N}}_{\text{inv}}^\sigma(\mathcal{B})$ as before, and define the l -th death-cascade \hat{G}_l for network $\tilde{\mathcal{N}}^\sigma$ as the l -th birth-cascade for network $\tilde{\mathcal{N}}_{\text{inv}}^\sigma(\mathcal{B})$. Correspondingly the *Death-cascade Path* (DCP) for network $\tilde{\mathcal{N}}^\sigma$ is the BCP for network $\tilde{\mathcal{N}}_{\text{inv}}^\sigma(\mathcal{B})$ and we can represent it as

$$\emptyset \Rightarrow \hat{G}_1 \Rightarrow \hat{G}_2 \Rightarrow \cdots \Rightarrow \hat{G}_{l_d}, \quad (4.51)$$

where \hat{G}_{l_d} is the terminal node of DCP and the set

$$\mathcal{X} = \mathcal{D}_f^\sigma(\hat{G}_{l_d})$$

consists of all the *free* species that can be arranged into death-cascades. Note that since the set of species is restricted to \mathcal{B} we must have $\mathcal{X} \subset \mathcal{B}$. The dynamical relation (2.22) satisfied by a network and its inverse, provides us with the following interpretation of death-cascades: the *free* species in $\mathcal{D}_f^\sigma(\hat{G}_1)$ for level $l = 1$ degrade spontaneously due to reactions in \mathcal{K}_r of the form $\mathbf{S} \rightarrow \emptyset$, while for higher levels $l > 1$ each *free* species \mathbf{S}_i in $\mathcal{D}_f^\sigma(\hat{G}_l)$ is converted by a reaction in \mathcal{K}_r to some *free* species \mathbf{S}_j at a lower level, i.e. species \mathbf{S}_j belongs to the set $\mathcal{D}_f^\sigma(\hat{G}_{l-1})$. For most biological networks, spontaneous degradation is very common and so typically the first death-cascade $\mathcal{D}_f^\sigma(\hat{G}_1)$ is heavily populated (see Section 6). The species in \mathcal{B} that do not belong to $\mathcal{D}_f^\sigma(\hat{G}_1)$ usually belong to one of the higher death-cascades because generally they undergo a sequence of conversions to eventually produce a spontaneously degradable species (see Section 6). These remarks suggest that for most biological networks we can expect to have

$$\mathcal{B} = \mathcal{X}, \quad (4.52)$$

which is to say that all the *free* species in \mathcal{B} can be arranged into death-cascades by the procedure we just described.

For our main result we need to ensure that each and every molecule of the *free* species in \mathcal{B} can be *flushed-out* from the system. Due to this reason we need to impose certain stoichiometric restrictions on the network as we describe now. Let \tilde{S}_σ be the $d_f \times K$ stoichiometry matrix for network $\tilde{\mathcal{N}}^\sigma$ and let $\tilde{S}_\sigma[\mathcal{K}_p]$ be the matrix formed by restricting \tilde{S}_σ to only those columns that correspond to reactions in \mathcal{K}_p (see (4.49)). For any $i \in \mathcal{F}$ the *free* species $\sigma(i) \in \mathcal{D}_f$ is said to be *singularly-degradable* if

$$-e_i \in \text{Colspan}_{\mathbb{N}_0} \left(\tilde{S}_\sigma[\mathcal{K}_p] \right), \quad (4.53)$$

where e_i denotes the i -th standard basis vector in \mathbb{R}^{d_f} . Note that the column span in this condition is over nonnegative integers \mathbb{N}_0 . Observe that if a *free* species can degrade spontaneously because of a reaction like $\mathbf{S} \rightarrow \emptyset$ then it is certainly singularly-degradable. Hence condition (4.53) primarily pertains to those *free* species that require a sequence of conversion reactions to produce a spontaneously degradable species. We can easily identify all the singularly-degradable *free* species by computing a modified Hermite normal form⁴ of the transpose of the matrix $-\tilde{S}_\sigma[\mathcal{K}_p]$. For obtaining this normal form, the only admissible operation is the addition of one row to another, possibly after multiplication by a positive integer. Once this Hermite normal form has been computed, the *free* species $\sigma(i)$ is singularly-degradable if and only if there exists a row with the leading entry of 1 in column i of the Hermite normal form.

⁴The Hermite normal form is an analogue of the row-reduced echelon form for integer matrices. For more details see [38].

We now state the main result of this section Proposition 4.3, which determines an irreducible state-space for network $\tilde{\mathcal{N}}^\sigma$ and also provides a way to easily check if it is the *only irreducible state-space* for the network in the infinite state-space $\tilde{\mathcal{E}}_0^\sigma = \mathbb{N}_0^{d_f}$. To state our main result, we shall consider the network dynamics under the permutation $\sigma = \sigma_3 : \mathcal{D} \rightarrow \mathcal{D}$ defined by

$$\sigma_3(i) = \begin{cases} j_i & \text{for } i = 1, \dots, d_{fb} \\ m_{i-d_{fb}} & \text{for } i = (d_{fb} + 1), \dots, d_f \\ \sigma_2(i) & \text{for } i = (d_f + 1), \dots, (d_f + d_r), \end{cases} \quad (4.54)$$

where $d_{fb} = |\mathcal{B}|$ be the total number of *free* species that can be arranged into birth-cascades, the set \mathcal{B} is $\mathcal{B} = \{j_1, \dots, j_{d_{fb}}\}$ and the set of remaining *free* species is $\mathcal{D}_f \cap \mathcal{B}^c = \{m_1, \dots, m_{d_f-d_{fb}}\}$. Observe that under this permutation the copy-numbers of all the *free* species in \mathcal{B} occupy the first d_{fb} components of the state vector.

Proposition 4.3 *Assume that network $\tilde{\mathcal{N}}^\sigma$ satisfies Assumption 2.1 and it has no bounded species (i.e. $d_b = 0$). Suppose that permutation σ is defined by (4.54) and (4.52) holds. Then we have the following:*

- (A) *If all the free species in \mathcal{B} are singularly-degradable then $\mathbb{N}_0^{d_{fb}} \times \{\mathbf{0}\}$ is an irreducible state-space for network $\tilde{\mathcal{N}}^\sigma$, where $\mathbf{0}$ denotes the $(d_f - d_{fb})$ -dimensional vector of zeros.*
- (B) *Additionally if all the free species in $\mathcal{D}_f \cap \mathcal{B}^c$ are also singularly-degradable then $\mathbb{N}_0^{d_{fb}} \times \{\mathbf{0}\}$ is the **only** irreducible state-space for network $\tilde{\mathcal{N}}^\sigma$.*

Proof. This proposition is a special case of Theorem 4.4 that is proved in the Appendix. \square

We end this section with a simple example that emphasizes the importance of imposing the condition of being singularly-degradable on the *free* species. Consider a network in which a single species \mathbf{S} participates in the following two reactions



We assume that the propensity functions follow mass-action kinetics. In this network the molecules of \mathbf{S} are produced and degraded in *pairs*. One can easily check that for this network $\mathcal{B} = \mathcal{X} = \{\mathbf{S}\}$, but \mathbb{N}_0 is **not** an irreducible state-space because species \mathbf{S} is not singularly-degradable. One can check that this network has two irreducible state-spaces: the set of all odd integers $\mathbb{O} = \{1, 3, 5, \dots\}$ and the set of all even integers $\mathbb{E} = \{0, 2, 4, \dots\}$.

4.4 Networks with both *unbounded* and *bounded* species

In this section we consider the general case where network $\tilde{\mathcal{N}}^\sigma = (\tilde{\mathcal{V}}^\sigma, \tilde{\mathcal{O}}^\sigma, \tilde{\Lambda}^\sigma)$ has both *unbounded* and *bounded* species and so all the irreducible state-spaces will lie in the infinite set $\tilde{\mathcal{E}}_0^\sigma = \mathcal{E}_b^\sigma \times \mathbb{N}_0^{d_f}$ after the *restricted* species have been removed (see Section 4.1). We now develop a procedure to identify all these state-spaces by *intertwining* the matrix-based approach in Section 4.2 with the birth-death cascade construction described in Section 4.3. However instead of linear *path-like* cascades we need to construct *tree-like* structures as we explain below.

We first fix some notation. For each reaction $k = 1, \dots, K$ let \bar{v}_k^σ , \tilde{v}_k^σ , $\bar{\rho}_k^\sigma$ and $\tilde{\rho}_k^\sigma$ denote the k -th column of matrices $\text{Proj}(\tilde{\mathcal{V}}^\sigma, 1, d_b)$, $\text{Proj}(\tilde{\mathcal{V}}^\sigma, d_b + 1, d_b + d_f)$, $\text{Proj}(\tilde{\mathcal{O}}^\sigma, 1, d_b)$ and $\text{Proj}(\tilde{\mathcal{O}}^\sigma, d_b + 1, d_b + d_f)$ respectively. Recall that the set of all *free* species is $\mathcal{D}_f = \{\sigma(d_b + 1), \dots, \sigma(d_b + d_f)\}$ and let $\mathcal{F} = \{1, \dots, d_f\}$ denote the set of addresses of all the *free* species under the map $\sigma(d_b + \cdot)$. For any $A \subset \mathcal{F}$, let $\mathcal{D}_f^\sigma(A) \subset \mathcal{D}_f$ be given by (4.47) with $\sigma(\cdot)$ replaced by $\sigma(d_b + \cdot)$. Corresponding to such a set A we define a set of reactions as

$$\mathcal{K}_r(A) = \{k = 1, \dots, K : \text{supp}(\tilde{v}_k^\sigma) \subset A\},$$

which includes only those reactions that do **not** consume any *free* species that lie outside the set $\mathcal{D}_f^\sigma(A)$. For any $y \in \mathcal{E}_b^\sigma$ define

$$\mathcal{K}_r(y, A) = \mathcal{K}_r(y) \cap \mathcal{K}_r(A),$$

where $\mathcal{K}_r(y)$ is given by (4.42). Let $Z(A)$ be the zero-pattern matrix defined by (4.43) with $\mathcal{K}_r(y)$ replaced by $\mathcal{K}_r(y, A)$. Using the procedure described in Section 4.2 we can compute the associated communication relation $\Theta(A)$ (see (4.45)) along with the corresponding set of equivalence or communication classes. We can also compute the reachability relations $R(A)$ (see (4.46)) among these classes and determine which classes are closed and open. From now on let $\mathcal{C}(A)$ denote the set of closed communication classes under relation $\Theta(A)$.

Suppose A_1, A_2 are two subsets of \mathcal{F} such that $A_1 \subset A_2$. Then certainly $\mathcal{K}_r(A_1) \subset \mathcal{K}_r(A_2)$ and hence $Z_{ij}(A_1) \leq Z_{ij}(A_2)$ for each i, j . Consequently $\Theta(A_2)$ has fewer equivalence classes than $\Theta(A_1)$, and each class of $\Theta(A_1)$ is contained in a unique class of $\Theta(A_2)$. Let C_1 and C_2 be two closed communication classes in the sets $\mathcal{C}(A_1)$ and $\mathcal{C}(A_2)$ respectively. We say that C_1 *reaches* C_2 if either $C_1 \subset C_2$, or C_1 is a subset of an open class O for $\Theta(A_2)$ and the closed class C_2 is reachable from O under relation $R(A_2)$. We define a map from $\mathcal{C}(A_1)$ to a subset of $\mathcal{C}(A_2)$ by

$$\Psi_{(A_1, A_2)}(C_1) = \{C_2 \in \mathcal{C}(A_2) : C_1 \text{ reaches } C_2\}.$$

Observe that this map is transitive in the sense that for any $A_1 \subset A_2 \subset A_3 \subset \mathcal{F}$ if $C_2 \in \Psi_{(A_1, A_2)}(C_1)$ and $C_3 \in \Psi_{(A_2, A_3)}(C_2)$ then $C_3 \in \Psi_{(A_1, A_3)}(C_1)$. This transitivity holds because $\mathcal{K}_r(A_1) \subset \mathcal{K}_r(A_2) \subset \mathcal{K}_r(A_3)$.

We now develop the notion of a *Birth-Cascade Tree* (BCT) by generalizing the ideas in Section 4.3. This tree is developed in levels or generations indexed by nonnegative integers $l = 0, 1, 2, \dots$, and it is a *directed graph* with nodes in the set

$$\mathbb{T} = \{(C, A) : A \subset \mathcal{F} \text{ and } C \in \mathcal{C}(A)\} \quad (4.55)$$

and edges in the set $\mathbb{T} \times \mathbb{T}$. For generation $l = 0$ the BCT is initialized by simply adding the nodes (C, \emptyset) for each $C \in \mathcal{C}(\emptyset)$. For any generation $l > 1$, the BCT is extended as follows: for each node (C, A) that was added to the BCT in the previous generation $(l - 1)$ we compute the set

$$\mathbb{B}(C, A) = \{i \in \mathcal{F} : i \notin A \text{ and there exists a } k \in \mathcal{K}_r(C, A) \text{ such that } i \in \text{supp}(\tilde{\rho}_k^\sigma)\}, \quad (4.56)$$

where

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

is the set of those reactions that only consume *free* species in the set $\mathcal{D}_f^\sigma(A)$ as reactants and have a positive probability of firing when the dynamics of *bounded* species is in the closed communication class C . If $\mathbb{B}(C, A) \neq \emptyset$ then we set $A' = A \cup \mathbb{B}(C, A)$ and for each $C' \in \Psi_{(A, A')}(C)$ we add the node (C', A') to the BCT along with the directed edge $(C, A) \Rightarrow (C', A')$.

The interpretation of the set $\mathbb{B}(C, A)$ is similar to the interpretation of $\mathbb{B}(A)$ in Section 4.3 with the only difference being that now the states of *bounded* species can move around in the closed communication class C . As the number of *free* species is finite, the BCT will stop growing beyond some generation l_b and this point the construction of BCT is complete and any node (C, A) from which there are no outgoing edges is called a *leaf* of the BCT. For such a leaf node we have

$$\mathbb{B}(C, A) = \emptyset, \quad (4.57)$$

and there exists a set of $(l + 1)$ BCT nodes, $\{(C_i, G_i) : i = 0, \dots, l\}$ such that $G_0 = \emptyset$, $C_l = C$, $G_l = A$ and the BCT has the following directed path

$$(C_0, G_0) \Rightarrow (C_1, G_1) \Rightarrow \dots \Rightarrow (C_l, G_l). \quad (4.58)$$

The transitivity of the map Ψ mentioned above implies that for any $i \in \{0, \dots, (l - 1)\}$ we have $C_l \in \Psi_{(G_i, G_l)}(C_i)$. Moreover since the set of reactions $\mathcal{K}_r(G_1), \mathcal{K}_r(G_2), \dots$ is monotonically increasing, if $C_0 \subset C_l$ then $C_i \subset C_l$ for each $i = 1, \dots, (l - 1)$.

Let \mathbb{L} be the set of all leaf nodes in the BCT. A leaf node $(C, A) \in \mathbb{L}$ is called *minimal* if there is no other leaf node that is strictly contained in (C, A) i.e. if $(C', A') \in \mathbb{L}$, $C' \subset C$ and $A' \subset A$ then $C = C'$ and $A = A'$. Henceforth we denote the set of all minimal leaf nodes by \mathbb{L}_{\min} . Our main result of the paper,

Theorem 4.4, will show that under certain conditions these minimal leaf nodes exactly characterize all the irreducible state-spaces for the network \tilde{N}^σ in the infinite state-space $\tilde{\mathcal{E}}_0^\sigma = \mathcal{E}_b^\sigma \times \mathbb{N}_0^{d_f}$. However before we state this result we need an appropriate notion of death-cascades and we need to impose some stoichiometric restrictions on the network as in Section 4.3.

Pick a minimal leaf node $(C, A) \in \mathbb{L}_{\min}$ and let $\tilde{N}^\sigma(C, A)$ be the network formed by removing all the *free* species outside $\mathcal{D}_f^\sigma(A)$ and discarding all the reactions outside $\mathcal{K}_r(C, A)$. Moreover we also replace the finite state-space \mathcal{E}_b^σ of the *bounded* species with its subset C . This replacement is feasible because (4.57) holds ensuring that reactions in $\mathcal{K}_r(C, A)$ cannot produce any *free* species outside $\mathcal{D}_f^\sigma(A)$, and C is a closed communication class for the states of *bounded* species under these reactions. Let $\tilde{N}_{\text{inv}}^\sigma(C, A)$ be the inverse of network $\tilde{N}^\sigma(C, A)$ (see Section 2.4). We compute the BCT for network $\tilde{N}_{\text{inv}}^\sigma(C, A)$ and refer to it as the *Death-cascade Tree* (DCT) for network $\tilde{N}^\sigma(C, A)$. Let $\mathbb{L}_d(C, A)$ denote the set of all leaf nodes of this DCT. Note that for any leaf node $(C', A') \in \mathbb{L}_d(C, A)$ we certainly have $C' \subset C$ and $A' \subset A$. We say that (C, A) is *death-exhaustive* if there exists a leaf node $(C', A') \in \mathbb{L}_d(C, A)$ such that $A' = A$. If such a leaf node exists then it is the only leaf node in $\mathbb{L}_d(C, A)$ since $C' = C$. This is shown in the proof of Lemma 7.2.

Consider a minimal leaf node $(C, A) \in \mathbb{L}_{\min}$ and let $\tilde{\mathcal{K}}_p(A)$ be the set of all those reactions that do not involve any *bounded* species and do not produce any *free* species outside the set $\mathcal{D}_f^\sigma(A)$, i.e.

$$\tilde{\mathcal{K}}_p(A) = \{k = 1, \dots, K : \text{supp}(\bar{\nu}_k^\sigma) = \text{supp}(\bar{\rho}_k^\sigma) = \emptyset \text{ and } \text{supp}(\bar{\rho}_k^\sigma) \subset A\}. \quad (4.59)$$

Let $\hat{S}_\sigma = \text{Proj}(\tilde{S}_\sigma, d_b + 1, d_b + d_f)$ be the $d_f \times K$ matrix formed by the last d_f rows of the $(d_b + d_f) \times K$ stoichiometry matrix \tilde{S}_σ for network \tilde{N}^σ and let $\hat{S}_\sigma[\tilde{\mathcal{K}}_p(A)]$ be the matrix formed by restricting \hat{S}_σ to only those columns that correspond to reactions in $\tilde{\mathcal{K}}_p(A)$. Recall that the set of all *free* species is $\mathcal{D}_f = \{\sigma(d_b + i) : i \in \mathcal{F}\}$. For any $i \in \mathcal{F}$ we say that the *free* species $\sigma(d_b + i) \in \mathcal{D}_f$ is *singularly-degradable* w.r.t. A if (4.53) is satisfied with matrix $\hat{S}_\sigma[\tilde{\mathcal{K}}_p(A)]$. To obtain the irreducible state-space corresponding to the leaf node $(C, A) \in \mathbb{L}_{\min}$, we shall permute the network according to the permutation $\sigma_3(A) : \mathcal{D} \rightarrow \mathcal{D}$ which ensures that the copy-numbers of all the *free* species in $\mathcal{D}_f^\sigma(A)$ occupy the components $(d_b + 1), \dots, (d_b + |A|)$ of the state vector. Such a permutation can be defined by setting it to be the same as σ_2 on the set $\{1, \dots, d_b\}$ and letting its image on the set $\{(d_b + 1), \dots, (d_b + d_f)\}$ be determined by (4.54) with $\mathcal{B} = \mathcal{D}_f^\sigma(A)$. We now state the main result of our paper.

Theorem 4.4 *Suppose Assumption 3.1 holds, network \tilde{N}^σ satisfies Assumption 2.1 and its state-space is $\tilde{\mathcal{E}}_0^\sigma = \mathcal{E}_b^\sigma \times \mathbb{N}_0^{d_f}$. Let \mathbb{L} and \mathbb{L}_{\min} be the sets of all leaf nodes and minimal leaf nodes for the BCT of the network. Then we have the following:*

- (A) *Pick any minimal leaf node $(C, A) \in \mathbb{L}_{\min}$. Suppose that this leaf node is death-exhaustive and all the free species in $\mathcal{D}_f^\sigma(A)$ are singularly-degradable w.r.t. A . Then the set $C \times \mathbb{N}_0^{|A|} \times \{\mathbf{0}\}$ is an irreducible state-space for network \tilde{N}^σ permuted according to permutation $\sigma = \sigma_3(A)$, where $\mathbf{0}$ denotes the $(d_f - |A|)$ -dimensional vector of zeros. Additionally if all the free species in the set $\mathcal{D}_f \cap (\mathcal{D}_f^\sigma(A))^c$ are also singularly-degradable w.r.t. A then $C \times \mathbb{N}_0^{|A|} \times \{\mathbf{0}\}$ is **the only** irreducible state-space for network \tilde{N}^σ that can have a nonempty intersection with the set $C \times \mathbb{N}_0^{d_f}$.*
- (B) *Suppose that all leaf nodes are minimal (i.e. $\mathbb{L} = \mathbb{L}_{\min}$), each leaf node $(C, A) \in \mathbb{L}$ is death exhaustive and all the free species in \mathcal{D}_f are singularly-degradable w.r.t. A . Then up to the relabeling of free species in \mathcal{D}_f , all the irreducible state-spaces for network \tilde{N}^σ are given by*

$$C \times \mathbb{N}_0^{|A|} \times \{\mathbf{0}\} \quad \text{for all} \quad (C, A) \in \mathbb{L}.$$

Proof. This theorem is proved in the Appendix and it generalizes Proposition 4.3 by taking the dynamics of *bounded* species into account. \square

We end this section with a couple of useful remarks.

Remark 4.5 Commonly in Systems Biology networks (see Section 6) there is only one leaf node (C, A) and there is only one class C_0 in $\mathcal{C}(\emptyset)$ such that $C_0 \subset C$. Hence all leaf nodes are minimal (because $\mathbb{L} = \mathbb{L}_{\min} = \{(C, A)\}$) and one can see that the BCT reduces to a single path

$$(C_0, \emptyset) \Rightarrow (C_1, G_1) \Rightarrow \cdots \Rightarrow (C_{l_b}, G_{l_b}) = (C, A).$$

In such a scenario we can talk about the birth-cascades for the network in the same way as in Section 4.3. In particular $\mathcal{D}_f^\sigma(G_l)$ is the set of free species that belong to the l -th birth cascade. If the leaf node (C, A) is death-exhaustive as well, then the Death-Cascade Tree (DCT) for network $\tilde{\mathcal{N}}^\sigma(C, A)$ is also a single path and so we can also talk about the death-cascades for the network in the same way as in Section 7.2.

Remark 4.6 Suppose from part (A) of Theorem 4.4 we obtain an irreducible state-space for network $\tilde{\mathcal{N}}^\sigma$. Then the corresponding irreducible state-space for the original network \mathcal{N} can be easily identified using Propositions 4.2 and 2.3. Similarly if **all** the irreducible state-spaces for network $\tilde{\mathcal{N}}^\sigma$ can be found using part (B) of Theorem 4.4, then we can identify all the irreducible state-spaces for the original network \mathcal{N} using Propositions 4.2 and 2.3.

5 Algorithms

The aim of this section is to provide detailed algorithmic descriptions of various procedures that can be used to apply the results in this paper. We start with network $\mathcal{N} = (\mathcal{V}, \mathcal{O}, \Lambda)$ with d species in the set $\mathcal{D} = \{1, \dots, d\}$ and K reactions of the form (see Section 2). We assume that this network satisfies Assumption 2.1 and its conservation data is (Γ, c) . Our first goal is to find a decomposed state-space \mathcal{E}_0^σ of the form (1.13) under some suitably constructed permutation $\sigma : \mathcal{D} \rightarrow \mathcal{D}$. This is accomplished in the method FINDDECOMPOSEDSTATESPACE(\cdot) (see Algorithm 1). This method starts by identifying the *bounded* species 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 $d_r > 0$, then it tries to classify the *unbounded* species into *free* and *restricted* species in such a way that the associated affine map ϕ (see (3.38)) is compatible with the reaction network (see Assumption 4.1). Such a compatible map is found by iterating through the set \mathcal{I}_f (see (3.34)), whose elements determine all plausible partitions of the set of *unbounded* species (\mathcal{D}_u) into the sets of *free* (\mathcal{D}_f). The outputs returned by FINDDECOMPOSEDSTATESPACE(\cdot) are the permutation map σ and the decomposed state-space \mathcal{E}_0^σ .

If the network has any *restricted* species (i.e. $d_r > 0$), then we construct the reduced network $\tilde{\mathcal{N}}^\sigma = (\tilde{\mathcal{V}}, \tilde{\mathcal{O}}, \tilde{\Lambda})$ along with the associated permutation $\sigma = \sigma_2$ (see (3.35)), by systematically removing the *restricted* species as described in Section 4.1. Assuming that a compatible affine map ϕ was discovered by FINDDECOMPOSEDSTATESPACE(\cdot), the reduced network $\tilde{\mathcal{N}}^\sigma$ will also satisfy Assumption 2.1, which is necessary for the process of identifying all the irreducible state-spaces. This process is accomplished by our next method FINDIRREDUCIBLESTATESPACES(\cdot) (see Algorithm 2), which works by first constructing the Birth-Cascade Tree (BCT) for network $\tilde{\mathcal{N}}^\sigma$ and then examining all the minimal leaf nodes of this BCT (see Section 4.4). The construction of BCT is performed by method CONSTRUCTBCT(\cdot) (see Algorithm 3) which returns the set of leaves \mathbb{L} of the BCT.

For each minimal leaf node $(C, A) \in \mathbb{L}_{\min}$ the method FINDIRREDUCIBLECLASSES(\cdot) checks if the set $C \times \mathbb{N}_0^{|A|} \times \{\mathbf{0}\}$ is an irreducible state-space for network $\tilde{\mathcal{N}}^\sigma$ under the permutation $\sigma = \sigma_3(A)$ (see Section 4.4) along with whether all the *free* species are singularly-degradable w.r.t. A . If both of these are true, then part (A) of Theorem 4.4 tells us that $C \times \mathbb{N}_0^{|A|} \times \{\mathbf{0}\}$ is the only irreducible state-space which can have common elements with the set $C \times \mathbb{N}_0^{d_f}$. If this is the case for all minimal leaf nodes and $\mathbb{L} = \mathbb{L}_{\min}$, then using FINDIRREDUCIBLECLASSES(\cdot) we recover all the irreducible state-spaces for our network due to part (B) of Theorem 4.4. Remark 4.6 tells us how these irreducible state-spaces for network $\tilde{\mathcal{N}}^\sigma$ can be mapped to the corresponding irreducible state-space for the original network \mathcal{N} .

Algorithm 1 Finds a decomposed state-space \mathcal{E}_0^σ of the form (1.13) for a network $\mathcal{N} = (\mathcal{V}, \mathcal{O}, \Lambda)$ with conservation data (Γ, c)

Require: Network \mathcal{N} satisfies Assumption 2.1

```

1: function FINDDECOMPOSEDSTATESPACE( $\mathcal{N}, \Gamma, c$ )
2:   For each species  $i \in \mathcal{D}$  solve the LP (3.28) to compute  $b_i$ .
3:   Set  $\mathcal{D}_b = \{i \in \mathcal{D} : b_i < \infty\}$  and  $\mathcal{D}_u = \{i \in \mathcal{D} : b_i = \infty\}$  to be the sets of bounded and unbounded
   species respectively. Also set  $d_b = |\mathcal{D}_b|$  and  $d_u = |\mathcal{D}_u|$ .
4:   Select the permutation map  $\sigma_1 : \mathcal{D} \rightarrow \mathcal{D}$  according to (3.29).
5:   Set  $\sigma = \sigma_1$  and construct the permuted network  $\mathcal{N}^\sigma$  (see Section 2.5) along with its conservation
   data  $(\Gamma_\sigma, c)$ .
6:   Compute the finite set  $\mathcal{E}_b^\sigma$  according to (3.31). This is the state-space for bounded species in  $\mathcal{D}_b$ .
7:   Evaluate  $d_f$  and  $d_r$  according to (3.32).
8:   if  $d_r = 0$  then
9:     Output: The decomposed state-space is  $\mathcal{E}_0^\sigma = \mathcal{E}_b^\sigma \times \mathbb{N}_0^{d_f}$  under permutation  $\sigma$ .
10:  else
11:    Verify Assumption 3.1 and if it does not hold then return QUIT
12:    Compute the set  $\mathcal{I}_f$  given by (3.34).
13:    for all  $I \in \mathcal{I}_f$  do
14:      Set  $\mathcal{D}_f = \{\sigma(d_b + i) : i \in I\}$  and  $\mathcal{D}_r = \{\sigma(d_b + i) : i \in I^c\}$  to be the sets of free and restricted
      species respectively.
15:      Select the permutation map  $\sigma_2 : \mathcal{D} \rightarrow \mathcal{D}$  according to (3.35).
16:      Set  $\sigma = \sigma_2$  and construct the permuted network  $\mathcal{N}^\sigma$  along with its conservation data  $(\Gamma_\sigma, c)$ 
17:      Define the affine map  $\phi : \mathbb{R}^{d_f} \rightarrow \mathbb{R}^{d_r}$  according to (3.38).
18:      if  $\phi$  satisfies Assumption 4.1 then
19:        Output: Map  $\phi$  is compatible with network  $\mathcal{N}^\sigma$ .
20:        Exit the for-loop and go to step 23
21:      end if
22:    end for
23:    If  $\phi$  does not satisfy Assumption 4.1 then return QUIT. Otherwise let  $\Phi$  be the graph of  $\phi$  given
    by (1.14) .
24:    Output: The decomposed state-space is  $\mathcal{E}_0^\sigma = \mathcal{E}_b^\sigma \times \Phi$  under permutation  $\sigma$ 
25:  end if
26: end function

```

Algorithm 2 Identifies the irreducible state-spaces for network $\tilde{\mathcal{N}}^\sigma = (\tilde{\mathcal{V}}, \tilde{\mathcal{O}}, \tilde{\Lambda})$ in the state-space $\mathcal{E}_b^\sigma \times \mathbb{N}_0^{d_f}$

Require: Network $\tilde{\mathcal{N}}^\sigma$ only has *bounded* and *free* species in the sets $\mathcal{D}_b = \{\sigma(i) : i = 1, \dots, d_b\}$ and $\mathcal{D}_f = \{\sigma(d_b + i) : i = 1, \dots, d_f\}$ respectively.

```

1: function FINDIRREDUCIBLESTATESPACES( $\tilde{\mathcal{N}}^\sigma, \mathcal{E}_b^\sigma \times \mathbb{N}_0^{d_f}$ )
2:   Let  $\mathbb{L} = \text{CONSTRUCTBCT}(\tilde{\mathcal{N}}^\sigma, \mathcal{E}_b^\sigma \times \mathbb{N}_0^{d_f})$ 
3:   Identify the set of minimal leaf nodes  $\mathbb{L}_{\min} \subset \mathbb{L}$ 
4:   Initialize  $\mathbb{L}' = \emptyset$ .
5:   for all  $(C, A) \in \mathbb{L}_{\min}$  do
6:     Let  $\tilde{\mathcal{N}}^\sigma(C, A)$  be the network formed by removing all the free species outside  $\mathcal{D}_f^\sigma(A)$  and discarding
       all the reactions outside  $\mathcal{K}_r(C, A)$  (see Section 4.4).
7:     Let  $\tilde{\mathcal{N}}_{\text{inv}}^\sigma(C, A)$  be the inverse of network  $\tilde{\mathcal{N}}^\sigma(C, A)$  (see Section 2.4).
8:     Set  $\mathbb{L}_d(C, A) = \text{CONSTRUCTBCT}(\tilde{\mathcal{N}}_{\text{inv}}^\sigma(C, A), C \times \mathbb{N}_0^{d_f})$ 
9:     Let  $\tilde{\mathcal{K}}_p(A)$  be the set of reactions given by (4.59).
10:    Construct the matrix  $\hat{S}_\sigma[\tilde{\mathcal{K}}_p(A)]$  by restricting the stoichiometry matrix for network  $\tilde{\mathcal{N}}^\sigma$  to only
       its last  $d_f$  rows and only the columns that correspond to reactions in  $\tilde{\mathcal{K}}_p(A)$ .
11:    Compute the set of free species  $\mathcal{D}_{sd}(A) \subset \mathcal{D}_f$  that are singularly-degradable w.r.t.  $A$  i.e.
       
$$\mathcal{D}_{sd}(A) = \{\sigma(d_b + i) : i = 1, \dots, d_f \text{ and (4.53) is satisfied with matrix } \hat{S}_\sigma[\tilde{\mathcal{K}}_p(A)]\}.$$

12:    if  $\mathbb{L}_d(C, A) = \{(C, A)\}$  AND  $\mathcal{D}_{sd}(A) \subset \mathcal{D}_f^\sigma := \{\sigma(d_b + i) : i \in A\}$  then
13:      Set  $\sigma = \sigma_3(A)$  (see Section 4.4) and construct the permuted network  $\tilde{\mathcal{N}}^\sigma$  (see Section 2.5).
14:      Output: The set  $C \times \mathbb{N}_0^{|A|} \times \{\mathbf{0}\}$  is an irreducible state-space for network  $\tilde{\mathcal{N}}^\sigma$  under permutation
        $\sigma$ .
15:      if  $\mathcal{D}_{sd}(A) = \mathcal{D}_f$  then
16:        Update  $\mathbb{L}' \leftarrow \mathbb{L}' \cup \{(C, A)\}$ .
17:      end if
18:    end if
19:  end for
20:  if  $\mathbb{L} = \mathbb{L}_{\min} = \mathbb{L}'$  then
21:    Output: The irreducible state-spaces found in step 14 are the only irreducible state-spaces for
       network  $\tilde{\mathcal{N}}^\sigma$  in the state-space  $\mathcal{E}_b^\sigma \times \mathbb{N}_0^{d_f}$  up to the relabeling of free species in  $\mathcal{D}_f$ .
22:  end if
23: end function

```

Algorithm 3 Constructs the Birth-Cascade Tree (BCT) for network $\tilde{\mathcal{N}}^\sigma = (\tilde{\mathcal{V}}, \tilde{\mathcal{O}}, \tilde{\Lambda})$ with state-space $\mathcal{E}_b^\sigma \times \mathbb{N}_0^{d_f}$. This BCT is a directed graph with nodes in the set \mathbb{T} given by (4.55).

Require: Network $\tilde{\mathcal{N}}^\sigma$ only has *bounded* and *free* species in the sets $\mathcal{D}_b = \{\sigma(i) : i = 1, \dots, d_b\}$ and $\mathcal{D}_f = \{\sigma(d_b + i) : i = 1, \dots, d_f\}$ respectively.

```

1: function CONSTRUCTBCT( $\tilde{\mathcal{N}}^\sigma, \mathcal{E}_b^\sigma \times \mathbb{N}_0^{d_f}$ )
2:   Initialize  $l = 0$ ,  $\mathbb{L} = \emptyset$  and  $G_0 = \{(C, \emptyset) : \text{for each } C \in \mathcal{C}(\emptyset)\}$ .
3:   Add each node in  $G_0$  to the BCT.
4:   repeat
5:     Set  $G_{l+1} = \emptyset$ 
6:     for all  $(C, A) \in G_l$  do
7:       Compute the set  $B = \mathbb{B}(C, A)$  where the operator  $\mathbb{B}$  is defined by (4.56).
8:       if  $B \neq \emptyset$  then
9:         Set  $A' = A \cup B$ .
10:        for all  $C' \in \Psi_{(A, A')}(C)$  do
11:          Update  $G_{l+1} \leftarrow G_{l+1} \cup \{(C', A')\}$ .
12:          Add the node  $(C', A')$  to the BCT along with the directed edge  $(C, A) \Rightarrow (C', A')$ .
13:        end for
14:      else
15:        Update  $\mathbb{L} \leftarrow \mathbb{L} \cup \{(C, A)\}$ 
16:      end if
17:    end for
18:    Update  $l \leftarrow l + 1$ 
19:  until  $G_l = \emptyset$ 
20:  Output: The set of leaf nodes of the BCT is  $\mathbb{L}$ .
21: end function

```

6 Examples

In this section we illustrate our results using several networks from Systems Biology. We start by considering a family of simple Gene-Expression networks which illustrate various theoretical ideas developed in this paper (see Section 6.1). Next we consider a couple of Circadian Clock models (Section 6.2) and a Bacterial Heat-Shock response model (Section 6.3). The networks underlying these models have many species and reactions, and we discuss how the state-space analysis presented in this paper can help in understanding network design as well as the long-term behavior of the associated stochastic models. In Section 6.4 we provide a simple example of a Toxin-Antitoxin network to demonstrate how our results can aid the automatic, real-time application of the *quasi-stationary* approximation to speed up stochastic simulations of the multiscale network. Finally in Section 6.5, we present a class of networks, where our analysis along with certain existing results, provide the exact stationary distribution for the stochastic model.

In all the examples, the reactions propensity functions (λ_k -s) are assumed to have the *mass-action* form (2.20) unless otherwise stated. For correct interpretation of our results we provide a “Species Chart” that encodes the names of network species into the notation used in our paper i.e. $\mathbf{S}_1, \mathbf{S}_2, \dots$. Throughout this section the copy-number of species \mathbf{S}_i is denoted by x_i . Often we would need to permute the network (see Section 2.5) for our analysis. The permutation σ for which the final results are shown is given as a vector $\sigma = (\sigma(1), \dots, \sigma(d))$, and under this permutation the i -th component of the state-vector corresponds to the copy-number of species $\mathbf{S}_{\sigma(i)}$.

6.1 Family of Gene-Expression networks

We now consider several variants of the simple Gene-Expression network given in [39]. In these networks there is a Gene (G), which is responsible for the transcription of *massenger* RNA or mRNA (M) molecules, that later translate into the Protein (P) of interest. Both mRNA and Protein molecules degrade spontaneously and as in [40] we allow for the Gene to switch between an *active* (G_{on}) and an *inactive* (G_{off}) state. For now

we assume that the transcription of mRNA is only possible in the active gene state G_{on} .

The first model (see Network 0 in Figure 1A) we examine consists of 4 species (see the Species Chart in Table 1) and 6 reactions displayed in Figure 1B. For this network there is only one (independent) conservation relation $\gamma = (1, 1, 0, 0)$ which is semi-positive, and so there are two *bounded* species \mathbf{S}_1 and \mathbf{S}_2 , whose copy-numbers x_1 and x_2 are “locked” in the relation $x_1 + x_2 = c = 1$, which simply says that the total number of Gene copies c is conserved by the dynamics. The other two species \mathbf{S}_3 and \mathbf{S}_4 are not involved in this conservation relation, and are hence *free* species.

Species Chart	
$\mathbf{S}_1 = G_{\text{off}}$	$\mathbf{S}_3 = M$
$\mathbf{S}_2 = G_{\text{on}}$	$\mathbf{S}_4 = P$

Table 1: Species chart for Gene-expression networks in Section 6.1. G_{on} and G_{off} denote the gene in active and inactive states. M denotes the mRNA and P denotes the corresponding Protein.

Using Algorithms 1 and 2, we can identify the decomposed state-space and the irreducible state-spaces for Network 0 (see Table 2). One can verify that the state-space $\mathcal{E}_b^\sigma = \{(1, 0), (0, 1)\}$ for *bounded* species is a closed communication class in $\mathcal{C}(\emptyset)$ and so the Birth-Cascade Tree (BCT) for this network will only consist of one path whose nodes give us the birth-cascades of *free* species (see Remark 4.5). It is easy to see that the length of this path is 3, the first birth-cascade is $\{\mathbf{S}_3\}$ and the second one is $\{\mathbf{S}_4\}$. Observe that these two cascades correspond naturally to the two stages in the network, transcription and translation, thereby suggesting that our cascade construction approach is a useful tool for understanding the structure of large Systems Biology networks. This point is reinforced by examples in Sections 6.2 and 6.3. We can check that the only leaf node of BCT is death-exhaustive and recover the death-cascades for the network (see Remark 4.5). In Network 0, both *free* species degrade spontaneously and so they are both singularly-degradable w.r.t. any set. Other conditions of Theorem 4.4 also hold and so we can use this result to conclude that the full state-space $\mathcal{E}_b^\sigma \times \mathbb{N}_0^2$ is the unique irreducible state-space for Network 0. This concludes our state-space analysis for Network 0.

We now create Networks 1-4 by adding feedback from protein molecules (see [40]) to various reactions in Network 0 (see Figure 1A). This feedback is added to reaction k by multiplying its original mass-action propensity function $\lambda_k(x)$ by a Hill-type factor of the form

$$\theta_{\text{fb}} \frac{x_4^n}{c + x_4^n}, \quad (6.60)$$

where x_4 is the number of protein molecules, and θ_{fb}, c and n are strictly positive parameters. Note that $\lambda_k(x) = 0$ if $x_4 = 0$ which means that if we just multiply the propensity function for reaction k in Network 0 by this factor (6.60), then the modified network will not satisfy Assumption 2.1 which is required for our analysis. However we can circumvent this problem by simply adding a molecule of species \mathbf{S}_4 to both sides of reaction k , changing it from $A \rightarrow B$ to $A + \mathbf{S}_4 \rightarrow B + \mathbf{S}_4$. This simple trick ensures that the modified network satisfies Assumption 2.1 and its dynamics remains unaffected as the reaction stoichiometry is the same. Incorporating the feedback mechanism this way, we list the reactions for Networks 1-4 in Figure 1B and provide the results from the state-space analysis of these networks in Table 2. Observe that in each of these cases a unique irreducible state-space is guaranteed by Theorem 4.4, but this irreducible state-space varies among networks. Also note that in all the Networks 0-4, the reactions stoichiometries are the same and consequently their decomposed state-space is identical.

We now consider another gene-expression model (see Network 5 in Figure 1) by allowing the transcription of mRNA in the inactive gene state G_{off} , and having feedback from protein molecules to the translation reaction as well as the gene-switching reactions. In this network also the decomposed state-space is the same as before but there are exactly two irreducible state-spaces, $\{(0, 1)\} \times \mathbb{N}_0 \times \{0\}$ and $\{(1, 0)\} \times \mathbb{N}_0 \times \{0\}$ according to Theorem 4.4 (see Table 2).

We end this section by remarking that if all the Hill-type feedback factors have a positive *basal* level, i.e.

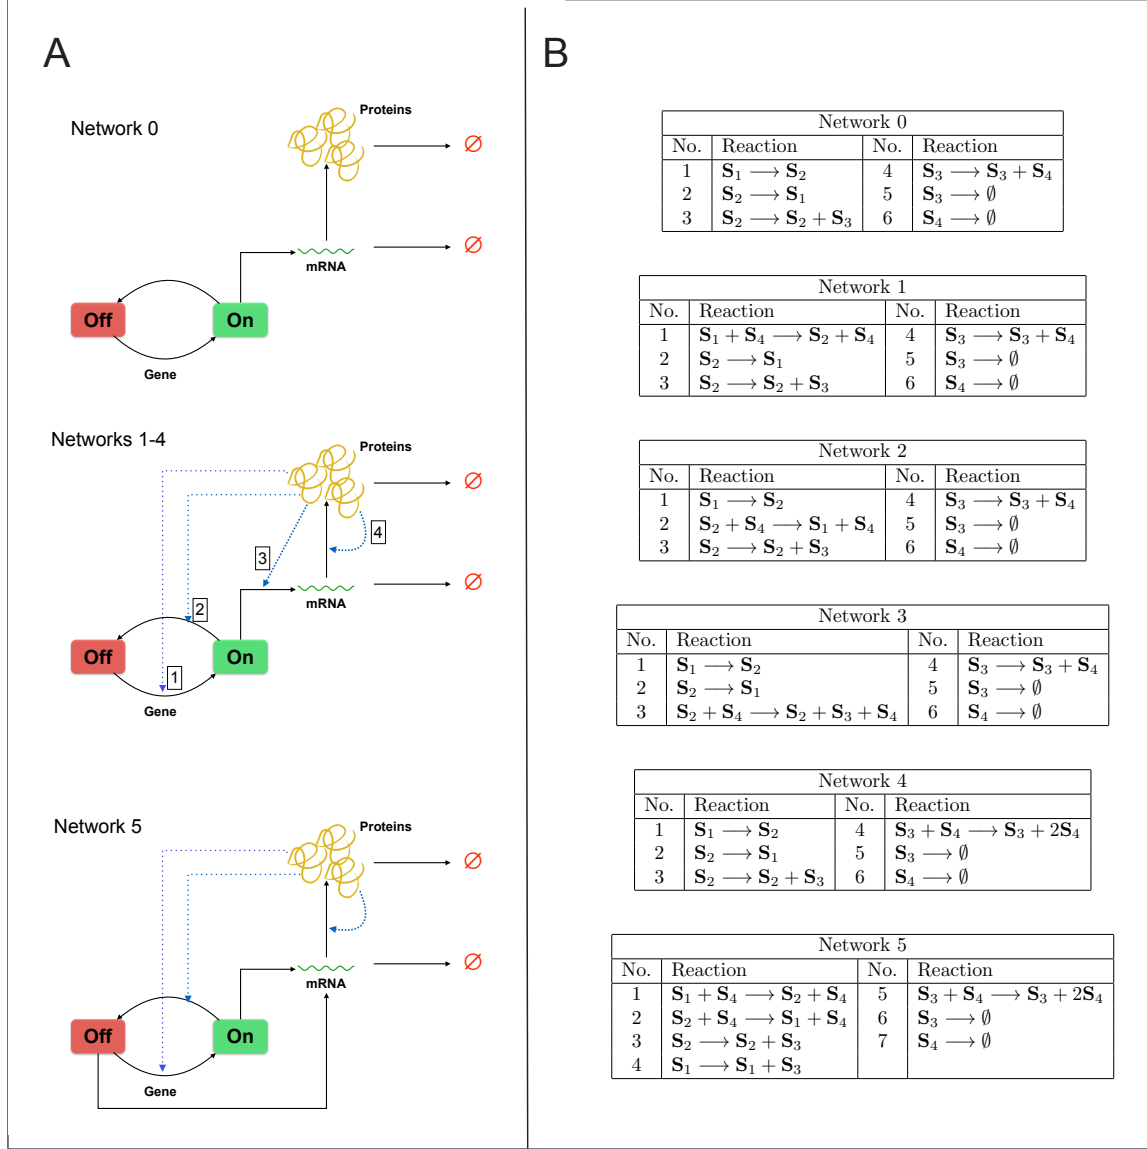


Figure 1: The Gene-expression networks in Section 6.1. The networks are depicted in panel **A** and the corresponding reactions sets are tabulated in panel **B**. The *feedback* interactions are shown with dotted arrows and for each of the Networks 1-4, the corresponding feedback interaction is clearly marked in the middle figure of panel **A**.

instead of (6.60) we have

$$\theta_1 + \theta_{fb} \frac{x_4^n}{c + x_4^n},$$

for some $\theta_1 > 0$, then we do not need to add species S_4 to both sides of feedback reactions to ensure that the modified networks satisfy Assumption 2.1. Indeed now Networks 1-4 will satisfy this assumption with the original set of reactions (same as Network 0) and hence they all will have the full state-space $\mathcal{E}_b^\sigma \times \mathbb{N}_0^2$ as their unique irreducible state-space. The same holds true for Network 5 even though it has an extra reaction. This illustrates the significance of Assumption 2.1 in determining the irreducible state-spaces.

Network	Decomposed State-Space	Irreducible state-spaces
0	$\mathcal{E}_b^\sigma \times \mathbb{N}_0^2$	$\mathcal{E}_b^\sigma \times \mathbb{N}_0^2$
1	$\mathcal{E}_b^\sigma \times \mathbb{N}_0^2$	$\{(1, 0, 0, 0)\}$
2	$\mathcal{E}_b^\sigma \times \mathbb{N}_0^2$	$\mathcal{E}_b^\sigma \times \mathbb{N}_0^2$
3	$\mathcal{E}_b^\sigma \times \mathbb{N}_0^2$	$\mathcal{E}_b^\sigma \times \{(0, 0)\}$
4	$\mathcal{E}_b^\sigma \times \mathbb{N}_0^2$	$\mathcal{E}_b^\sigma \times \mathbb{N}_0 \times \{0\}$
5	$\mathcal{E}_b^\sigma \times \mathbb{N}_0^2$	$\{(0, 1)\} \times \mathbb{N}_0 \times \{0\}$ and $\{(1, 0)\} \times \mathbb{N}_0 \times \{0\}$

Table 2: State-space analysis for Gene-expression networks in Section 6.1. All results reported w.r.t the identity permutation $\sigma = (1, 2, 3, 4)$ and $\mathcal{E}_b^\sigma = \{(1, 0), (0, 1)\}$. For all the five networks *all* irreducible state-spaces were identified by Theorem 4.4.

6.2 Two Circadian Clock models

Circadian clocks are molecular time-keeping devices that coordinate many physiological processes in living organisms [41]. These clocks generate oscillatory rhythms that are usually *entrained* to the periodic cues provided by the day-light cycles [42, 43]. We consider two circadian clock networks in this section and prove using our analysis that there exists a unique irreducible state-space for both these models, thereby indicating that the stationary distributions for the associated stochastic models is unique. The existence of these stationary distribution can be checked using the techniques in [27] and hence the stochastic models for both these networks are ergodic (see Section 1). Therefore under constant inputs the individual stochastic trajectory of a single circadian clock may be oscillatory, but the mean trajectories, corresponding to the bulk or population-level behavior of several uncoupled and identical circadian clocks, **cannot** be oscillatory due to (2.18). This is consistent with both computational [42, 27] and experimental [44] observations in the existing literature. In a recent paper it is argued that this loss of oscillatory activity at the population-level plays an important role in ensuring that the entrainment to periodic cues is robust at the population-level [45].

The first circadian clock model we consider is from Vilar et al. [5] and it is depicted in Figure 2A. It consists of gene-expression modules for an *activator* protein A and a *repressor* protein R which sequesters the activator protein A by forming an inactive complex AR . The activator protein A can enhance the transcription of both the mRNAs (M_A and M_R) by binding to the promoter regions of the activator gene D_A and the repressor gene D_R . When the promoter region is occupied the activator and the repressor genes are denoted by D'_A and D'_R respectively. The overall network consists of 9 species (see the Species Chart in Table 3) and 16 reactions (see Table 4). The results from our state-space analysis on this network are presented in Figure 2B and they show that the network has a unique irreducible state-space which coincides with its decomposed state-space. For this network the situation of Remark 4.5 applies and so we can arrange all the *free* species into birth and death cascades (Figure 2C). These cascades correspond naturally to different stages in the network.

Next we examine the detailed mammalian circadian clock given in [46] (see Figure 3A). It consists of three gene-expression modules corresponding to *Per*, *Cry* and *Bmal1* genes. The proteins created by these genes participate in complex regulatory steps through various mechanisms such as formation of dimeric complexes, translocation in an out of nucleus and undergoing phosphorylation-dephosphorylation cycles. Overall the network consists of 16 species (see the Species Chart in Table 5) and 52 reactions (see Table 6). Interestingly there are no conservation relations for this network and so all the species are *free* and Theorem 4.4 proves that \mathbb{N}_0^{16} is the unique irreducible state-space (see Figure 3). The birth and death cascades, displayed in Figure 3C, signify the various network stages as before.

6.3 Bacterial Heat-Shock response model

We now analyze the bacterial heat-shock response model developed in Kurata et al. [47]. It is known that when bacterial cells are exposed to thermal shocks, certain cellular structures are damaged, causing

Species Chart	
$\mathbf{S}_1 = M_A$	$\mathbf{S}_6 = G_A^b$
$\mathbf{S}_2 = A$	$\mathbf{S}_7 = G_A^u$
$\mathbf{S}_3 = M_R$	$\mathbf{S}_8 = G_R^b$
$\mathbf{S}_4 = R$	$\mathbf{S}_9 = G_R^u$
$\mathbf{S}_5 = AR$	

Table 3: Species chart for the first Circadian clock model (Vilar et al. [5]). A and R denote the activator and repressor proteins. G_A (M_A) and G_R (M_R) denote the gene (mRNA) corresponding to these proteins. The superscript b or u indicates the *bound* or *unbound* from of the gene. AR denotes the dimeric complex between A and R .

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	$\mathbf{S}_5 \longrightarrow \mathbf{S}_4$

Table 4: Reactions for the first Circadian clock model (Vilar et al. [5]). These reactions follow the Species Chart in Table 3. The network is depicted in Figure 2.

Species Chart	
$\mathbf{S}_1 = M_P$	$\mathbf{S}_9 = PC_N$
$\mathbf{S}_2 = M_C$	$\mathbf{S}_{10} = PC_{CP}$
$\mathbf{S}_3 = M_B$	$\mathbf{S}_{11} = PC_{NP}$
$\mathbf{S}_4 = P_C$	$\mathbf{S}_{12} = B_C$
$\mathbf{S}_5 = C_C$	$\mathbf{S}_{13} = B_{CP}$
$\mathbf{S}_6 = P_{CP}$	$\mathbf{S}_{14} = B_N$
$\mathbf{S}_7 = C_{CP}$	$\mathbf{S}_{15} = B_{NP}$
$\mathbf{S}_8 = PC_C$	$\mathbf{S}_{16} = I_N$

Table 5: Species chart for the second Circadian clock model (Leloup and Goldbeter [46]). P , C and B denote the *Per*, *Cry* and *Bmal1* proteins. The dimeric complex between two proteins X and Y is denoted by XY . The subscripts C , N and P stand for “Cytosol”, “Nucleus” and “Phosphorylated” respectively. M_X denotes the mRNA for protein X and I_N is an inactive trimer in the Nucleus.

some proteins to denature or unfold. Accumulation of these denatured proteins within cells can disrupt their normal functioning and hence a regulatory mechanism has evolved to detect and repair unfolded proteins. The model in [47] describes this mechanism in bacterium *E. Coli* and it consists of 5 distinct gene-expression modules responsible for creating various heat-shock specific proteins (*hsps*), such as the heat-shock transcription factor σ^{32} factor, proteases like *HslVu*, *FtsH* etc. and a molecular chaperon *DnaK*, which assists in the refolding of denatured proteins. The sigma factor σ^{32} competes with the dominant sigma factor σ^{70} to bind to the *RNA Polymerase* (RNAP), and when it is able to bind, it can initiate the transcription of the heat-shock gene *ph*, which then produces the other heat-shock proteins. The sigma factor σ^{32} is itself produced by the gene *pg* when RNAP is bound to sigma factor σ^{70} . As explained in

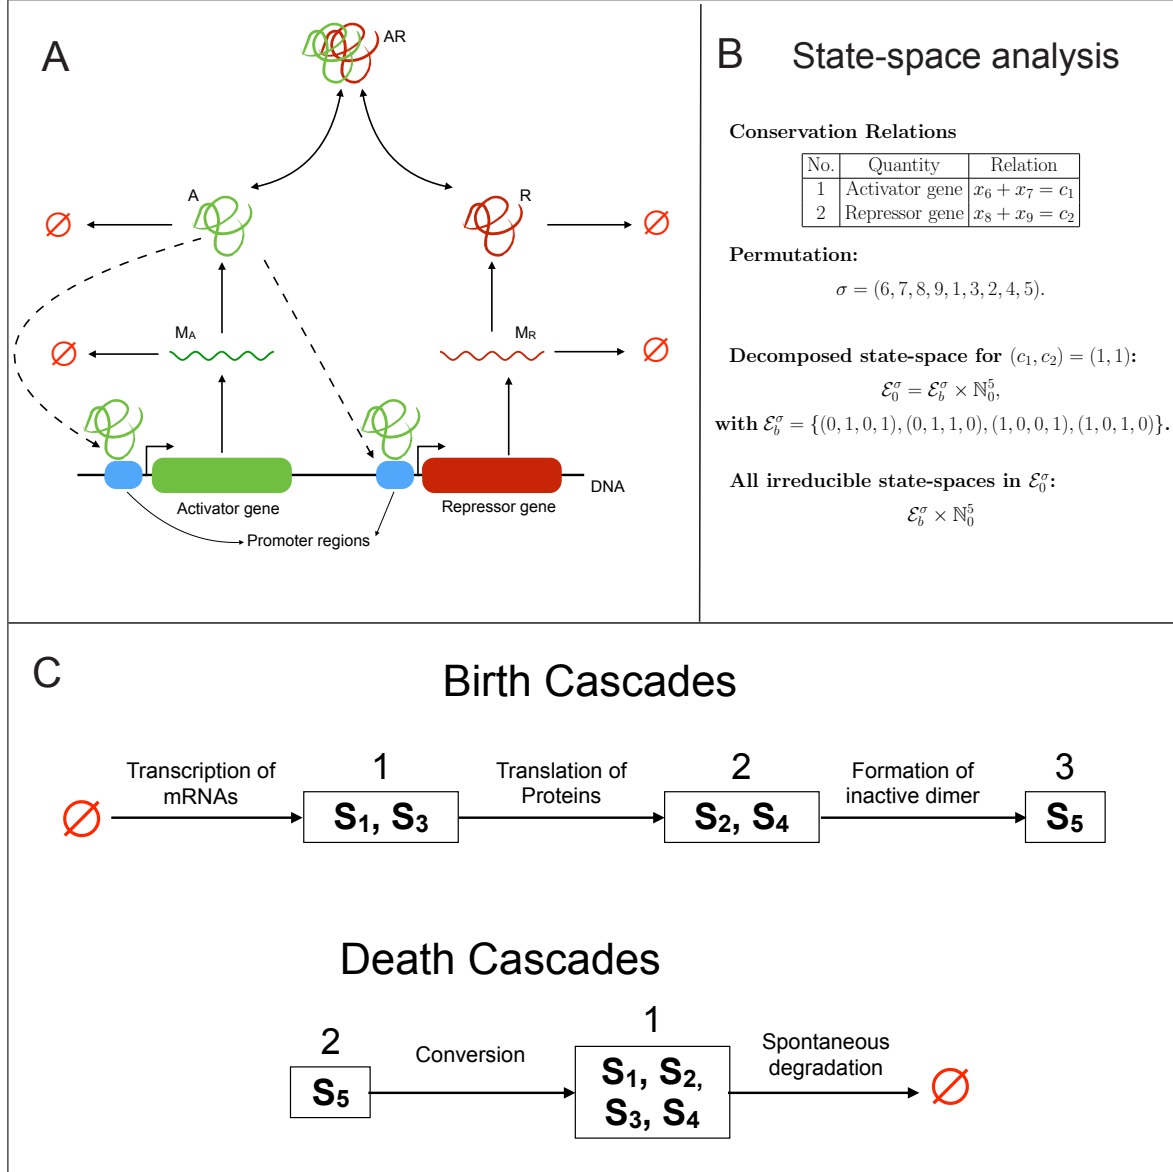


Figure 2: Results for the Circadian clock model (Vilar et al. [5]) depicted in panel **A**. The results of state-space analysis are presented in panel **B** and the birth-death cascades are shown in panel **C**. The birth-cascades signify various stages in the network. Note that in this network there are four *bounded* species ($\sigma(1), \dots, \sigma(4)$) and five *free* species ($\sigma(5), \dots, \sigma(9)$).

[47] the heat-shock proteins interact in complex ways to realize both feedback and feedforward control that confers several performance advantages to the heat-shock response mechanism.

The network underlying the model in [47] is shown in Figure 4**A** and it is fairly large, with 28 species (see the Species Chart in Table 7) and 61 reactions (see Table 8). The network contains 5 semi-positive conservation relations corresponding to various conserved quantities (see Figure 4**B**). Due to these relations 12 species are *bounded* while the rest are *free*. Using our analysis we can find the decomposed state-space for this network and verify the existence of a unique irreducible state-space (see Figure 4**B**). Note that in this example, the irreducible state-space is a strict subset of the full state-space and in particular, the last two components of each state in this irreducible state-space are zeros. These two components being zero implies that the copy-number of unfolded protein molecules is zero. As these unfolded protein molecules

No.	Reaction	No.	Reaction
1	$\emptyset \rightarrow \mathbf{S}_1$	27	$\mathbf{S}_8 \rightarrow \mathbf{S}_9$
2	$\mathbf{S}_1 \rightarrow \emptyset$	28	$\mathbf{S}_8 \rightarrow \emptyset$
3	$\mathbf{S}_1 \rightarrow \emptyset$	29	$\mathbf{S}_9 \rightarrow \mathbf{S}_{11}$
4	$\emptyset \rightarrow \mathbf{S}_2$	30	$\mathbf{S}_{11} \rightarrow \mathbf{S}_9$
5	$\mathbf{S}_2 \rightarrow \emptyset$	31	$\mathbf{S}_9 + \mathbf{S}_{14} \rightarrow \mathbf{S}_{16}$
6	$\mathbf{S}_2 \rightarrow \emptyset$	32	$\mathbf{S}_{16} \rightarrow \mathbf{S}_9 + \mathbf{S}_{14}$
7	$\emptyset \rightarrow \mathbf{S}_3$	33	$\mathbf{S}_9 \rightarrow \emptyset$
8	$\mathbf{S}_3 \rightarrow \emptyset$	34	$\mathbf{S}_{10} \rightarrow \emptyset$
9	$\mathbf{S}_3 \rightarrow \emptyset$	35	$\mathbf{S}_{10} \rightarrow \emptyset$
10	$\mathbf{S}_1 \rightarrow \mathbf{S}_1 + \mathbf{S}_4$	36	$\mathbf{S}_{11} \rightarrow \emptyset$
11	$\mathbf{S}_4 \rightarrow \mathbf{S}_6$	37	$\mathbf{S}_{11} \rightarrow \emptyset$
12	$\mathbf{S}_6 \rightarrow \mathbf{S}_4$	38	$\mathbf{S}_3 \rightarrow \mathbf{S}_{12}$
13	$\mathbf{S}_8 \rightarrow \mathbf{S}_4 + \mathbf{S}_5$	39	$\mathbf{S}_{12} \rightarrow \mathbf{S}_{13}$
14	$\mathbf{S}_4 + \mathbf{S}_5 \rightarrow \mathbf{S}_8$	40	$\mathbf{S}_{13} \rightarrow \mathbf{S}_{12}$
15	$\mathbf{S}_4 \rightarrow \emptyset$	41	$\mathbf{S}_{12} \rightarrow \mathbf{S}_{14}$
16	$\mathbf{S}_2 \rightarrow \mathbf{S}_2 + \mathbf{S}_5$	42	$\mathbf{S}_{14} \rightarrow \mathbf{S}_{12}$
17	$\mathbf{S}_7 \rightarrow \mathbf{S}_5$	43	$\mathbf{S}_{12} \rightarrow \emptyset$
18	$\mathbf{S}_5 \rightarrow \mathbf{S}_7$	44	$\mathbf{S}_{13} \rightarrow \emptyset$
19	$\mathbf{S}_5 \rightarrow \emptyset$	45	$\mathbf{S}_{13} \rightarrow \emptyset$
20	$\mathbf{S}_6 \rightarrow \emptyset$	46	$\mathbf{S}_{14} \rightarrow \mathbf{S}_{15}$
21	$\mathbf{S}_6 \rightarrow \emptyset$	47	$\mathbf{S}_{15} \rightarrow \mathbf{S}_{14}$
22	$\mathbf{S}_7 \rightarrow \emptyset$	48	$\mathbf{S}_{14} \rightarrow \emptyset$
23	$\mathbf{S}_7 \rightarrow \emptyset$	49	$\mathbf{S}_{15} \rightarrow \emptyset$
24	$\mathbf{S}_8 \rightarrow \mathbf{S}_{10}$	50	$\mathbf{S}_{15} \rightarrow \emptyset$
25	$\mathbf{S}_{10} \rightarrow \mathbf{S}_8$	51	$\mathbf{S}_{16} \rightarrow \emptyset$
26	$\mathbf{S}_9 \rightarrow \mathbf{S}_8$	52	$\mathbf{S}_{16} \rightarrow \emptyset$

Table 6: Reactions for the second Circadian clock model (Leloup and Goldbeter [46]). These reactions follow the Species Chart in Table 5. The network is depicted in Figure 3.

can only *leave* the network by converting to a properly folded protein (via reaction 25 in Table 8), we can conclude that the heat-shock response network is such, that despite having noisy, stochastic dynamics, all the unfolded protein molecules will eventually get folded, irrespective of the initial count of these unfolded protein molecules. We base this assertion on the fact that starting from any initial state in the state-space, the Markovian reaction dynamics will eventually get trapped inside the unique irreducible state-space. What makes this assertion powerful is that it holds *just* due to the structure of the network, regardless of the values of the rate constants of the reactions. Also for this network we are in the situation of Remark 4.5. Hence we can arrange all the *free* species into birth and death cascades which correspond to various network stages (see Figure 4C).

6.4 A simple Toxin-Antitoxin network

Many intracellular networks have reactions taking place on many different time-scales [32, 35]. It is well known that direct stochastic simulations of such networks, using Gillespie’s SSA [13], is computationally infeasible because most of the resources are spent on simulating the “fast” reactions which are generally less important than the “slow” ones [33]. To remedy this problem, approximate simulation approaches have been developed that *separate the time-scales* by applying the quasi-stationary assumption (QSA) on the fast subcomponents of the network [32]. Under this assumption, the slow reactions are “switched-off” and the dynamics of the slow species involved in these reactions is “frozen in time”, while the stable Markovian dynamics of the fast subnetwork relaxes to a unique stationary distribution (see (2.18)). Using this stationary distribution the propensities of the slow reactions can be estimated and the next slow reaction

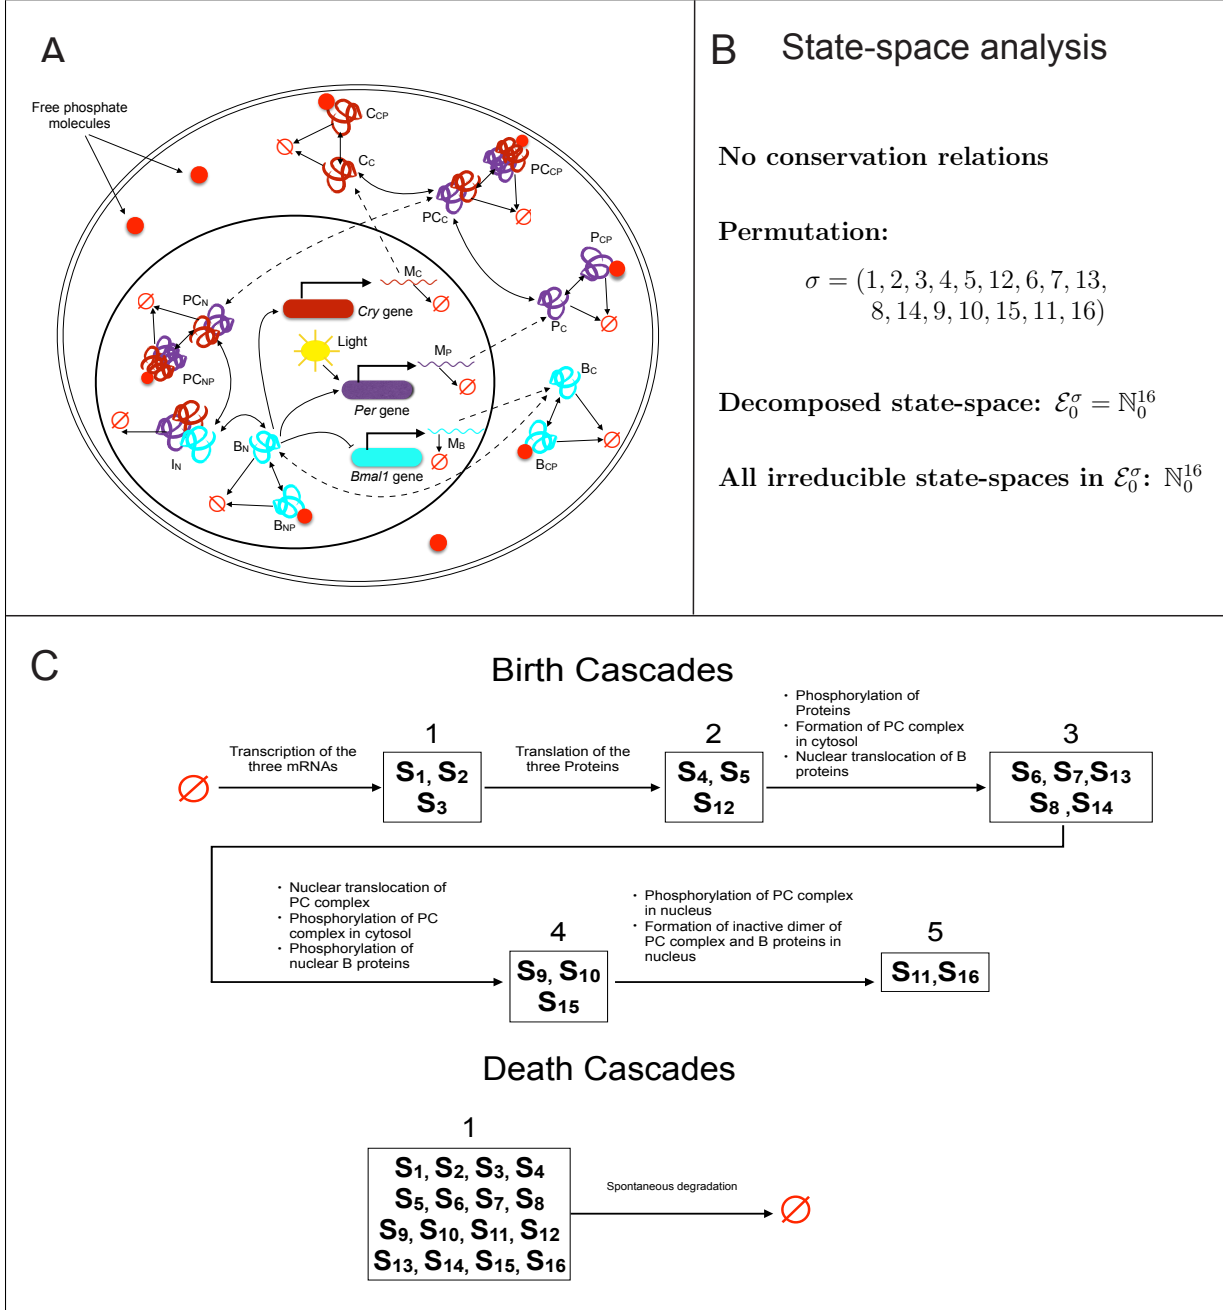


Figure 3: Results for the Circadian clock model (Leloup and Goldbeter [46]) depicted in panel A. The results of state-space analysis are presented in panel B and the birth-death cascades are shown in panel C. The birth-cascades signify various stages in the network. Note that in this network all the species are *free* as there are no conservation relations.

can be sampled. Applying QSA between every consecutive slow reaction events allows one to approximate the original dynamics without having to simulate the fast reactions, thereby reducing the overall computational effort drastically.

For a successful application of QSA it is imperative to ascertain the stability or ergodicity of the fast subnetwork in *real-time* (i.e. during the simulation run), as the set of available fast reactions can depend on the current state of the slow species. Moreover for efficiently estimating the slow reaction propensities it

Species Chart	
$\mathbf{S}_1 = \text{RNAP}$	$\mathbf{S}_{15} = \sigma^{32} : \text{DnaK}$
$\mathbf{S}_2 = \sigma^{70}$	$\mathbf{S}_{16} = \text{P}_{\text{unfolded}}$
$\mathbf{S}_3 = \sigma^{70} : \text{RNAP}$	$\mathbf{S}_{17} = \text{P}_{\text{unfolded}} : \text{DnaK}$
$\mathbf{S}_4 = \sigma^{32}$	$\mathbf{S}_{18} = \text{Protease}$
$\mathbf{S}_5 = \sigma^{32} : \text{RNAP}$	$\mathbf{S}_{19} = \sigma^{32} : \text{DnaK} : \text{Protease}$
$\mathbf{S}_6 = D$	$\mathbf{S}_{20} = \text{HslVu}$
$\mathbf{S}_7 = \text{RNAP} : D$	$\mathbf{S}_{21} = \sigma^{32} : \text{HslVu}$
$\mathbf{S}_8 = \sigma^{70} : \text{RNAP} : D$	$\mathbf{S}_{22} = \text{mRNA}(\text{DnaK})$
$\mathbf{S}_9 = \sigma^{32} : \text{RNAP}$	$\mathbf{S}_{23} = \text{mRNA}(\text{Protease})$
$\mathbf{S}_{10} = pg$	$\mathbf{S}_{24} = \text{mRNA}(\text{HslVu})$
$\mathbf{S}_{11} = ph$	$\mathbf{S}_{25} = \text{mRNA}(\sigma^{32})$
$\mathbf{S}_{12} = \sigma^{70} : \text{RNAP} : pg$	$\mathbf{S}_{26} = \text{mRNA}(\text{FtsH})$
$\mathbf{S}_{13} = \sigma^{32} : \text{RNAP} : ph$	$\mathbf{S}_{27} = \text{FtsH}$
$\mathbf{S}_{14} = \text{DnaK}$	$\mathbf{S}_{28} = \sigma^{32} : \text{DnaK} : \text{FtsH}$

Table 7: Species chart for the Bacterial Heat-Shock response model (Kurata et al. [47]). Here the complex formed by binding two biomolecules A and B is denoted by $A : B$. The main players in this network are the heat-shock proteins *DnaK*, *Protease*, *HslVu* and *FtsH*, and the transcription factor σ^{32} which binds to the RNA Polymerase (RNAP) to initiate the production of the heat-shock proteins via the *ph* gene. This transcription factor σ^{32} is itself created by the gene *pg* when the transcription factor σ^{70} is bound to RNAP. Here D denotes the part of DNA that does not include these genes, and P is the protein of interest that needs to be converted from its denatured form $\text{P}_{\text{unfolded}}$ to its proper form P_{folded} . The mRNA for protein X is denoted by $\text{mRNA}(X)$.

is useful to determine the exact time-varying support-sets of the stationary distributions. These tasks can be *automatically* accomplished by integrating our computational procedures for state-space analysis, with any QSA-based simulation algorithm like the *slow-scale* SSA or ssSSA [33] or the Nested SSA [34, 35]. We illustrate this next using a simple Toxin-Antitoxin network module which is found in many bacterial cells and is believed to lead to the formation of slow-growing *persister* cells that exhibit multi-drug tolerance. This example shows how *restricted* species arise naturally when we restrict our attention to a subnetwork within the larger network. Moreover this example also highlights that the flexibility in the classification of *unbounded* species as *free* or *restricted* (see Remark 3.3) can be really important for applications. In fact for this example, this classification will change randomly with time depending on the states visited by a stochastic trajectory.

The simple Toxin-Antitoxin network we consider is depicted in Figure 5 and it is based on the more detailed model given in [48]. It consists of a single DNA strand containing genes for both Toxin T and Antitoxin A protein molecules. Both these proteins are translated by a common *biscistronic* mRNA M and they both annihilate each other in the sense that they bind to form an inactive complex AT which does not participate in the dynamics. The Antitoxin molecules directly inhibit the transcription of mRNA and the Toxin molecules convert to a protein P that interferes with bacterial metabolism and harms the cells in various ways [49]. Following the Species Chart in Table 9, we can describe our simple Toxin-Antitoxin system as a network with 8 reactions which are listed along with their propensity functions in Table 10. The choice of rate constants in these propensity functions is arbitrary but reflective of the values found in the literature [48]. From these rates it can be inferred that two reactions can be considered fast, namely, the translation of mRNA and the mutual annihilation of Toxin and Antitoxin proteins (see Figure 5A).

Suppose one is interested in simulating the stochastic model for this network in the time period $[0, T]$ for $T = 100$. Performing exact stochastic simulations is cumbersome due to the presence of fast reactions, but an approximate QSA-based algorithm can be used for carrying out these simulations with far lesser computational effort. To apply QSA, we consider the dynamics of the component of the network consisting

No.	Reaction	No.	Reaction
1	$\mathbf{S}_1 + \mathbf{S}_2 \longrightarrow \mathbf{S}_3$	32	$\mathbf{S}_{18} \longrightarrow \emptyset$
2	$\mathbf{S}_3 \longrightarrow \mathbf{S}_1 + \mathbf{S}_2$	33	$\mathbf{S}_{19} \longrightarrow \mathbf{S}_{15}$
3	$\mathbf{S}_4 + \mathbf{S}_1 \longrightarrow \mathbf{S}_5$	34	$\mathbf{S}_{12} \longrightarrow \mathbf{S}_{25} + \mathbf{S}_{12}$
4	$\mathbf{S}_5 \longrightarrow \mathbf{S}_4 + \mathbf{S}_1$	35	$\mathbf{S}_{25} \longrightarrow \emptyset$
5	$\mathbf{S}_1 + \mathbf{S}_6 \longrightarrow \mathbf{S}_7$	36	$\mathbf{S}_{25} \longrightarrow \mathbf{S}_4 + \mathbf{S}_{25}$
6	$\mathbf{S}_7 \longrightarrow \mathbf{S}_1 + \mathbf{S}_6$	37	$\mathbf{S}_4 \longrightarrow \emptyset$
7	$\mathbf{S}_4 + \mathbf{S}_{14} \longrightarrow \mathbf{S}_{15}$	38	$\mathbf{S}_{19} \longrightarrow \mathbf{S}_{18} + \mathbf{S}_{14}$
8	$\mathbf{S}_{15} \longrightarrow \mathbf{S}_4 + \mathbf{S}_{14}$	39	$\mathbf{S}_{21} \longrightarrow \mathbf{S}_{20}$
9	$\mathbf{S}_{16} + \mathbf{S}_{14} \longrightarrow \mathbf{S}_{17}$	40	$\mathbf{S}_{28} \longrightarrow \mathbf{S}_{14} + \mathbf{S}_{27}$
10	$\mathbf{S}_{17} \longrightarrow \mathbf{S}_{16} + \mathbf{S}_{14}$	41	$\mathbf{S}_{13} \longrightarrow \mathbf{S}_{13} + \mathbf{S}_{24}$
11	$\mathbf{S}_3 + \mathbf{S}_6 \longrightarrow \mathbf{S}_8$	42	$\mathbf{S}_{24} \longrightarrow \emptyset$
12	$\mathbf{S}_8 \longrightarrow \mathbf{S}_3 + \mathbf{S}_6$	43	$\mathbf{S}_{24} \longrightarrow \mathbf{S}_{20} + \mathbf{S}_{24}$
13	$\mathbf{S}_5 + \mathbf{S}_6 \longrightarrow \mathbf{S}_9$	44	$\mathbf{S}_{20} \longrightarrow \emptyset$
14	$\mathbf{S}_9 \longrightarrow \mathbf{S}_5 + \mathbf{S}_6$	45	$\mathbf{S}_{21} \longrightarrow \mathbf{S}_4$
15	$\mathbf{S}_3 + \mathbf{S}_{10} \longrightarrow \mathbf{S}_{12}$	46	$\mathbf{S}_4 + \mathbf{S}_{20} \longrightarrow \mathbf{S}_{21}$
16	$\mathbf{S}_{12} \longrightarrow \mathbf{S}_3 + \mathbf{S}_{10}$	47	$\mathbf{S}_{21} \longrightarrow \mathbf{S}_4 + \mathbf{S}_{20}$
17	$\mathbf{S}_5 + \mathbf{S}_{11} \longrightarrow \mathbf{S}_{13}$	48	$\mathbf{S}_{13} \longrightarrow \mathbf{S}_{13} + \mathbf{S}_{26}$
18	$\mathbf{S}_{13} \longrightarrow \mathbf{S}_5 + \mathbf{S}_{11}$	49	$\mathbf{S}_{26} \longrightarrow \emptyset$
19	$\mathbf{S}_{15} + \mathbf{S}_{18} \longrightarrow \mathbf{S}_{19}$	50	$\mathbf{S}_{26} \longrightarrow \mathbf{S}_{26} + \mathbf{S}_{27}$
20	$\mathbf{S}_{19} \longrightarrow \mathbf{S}_{15} + \mathbf{S}_{18}$	51	$\mathbf{S}_{27} \longrightarrow \emptyset$
21	$\mathbf{S}_{13} \longrightarrow \mathbf{S}_{13} + \mathbf{S}_{22}$	52	$\mathbf{S}_{28} \longrightarrow \mathbf{S}_{15}$
22	$\mathbf{S}_{22} \longrightarrow \emptyset$	53	$\mathbf{S}_{15} + \mathbf{S}_{27} \longrightarrow \mathbf{S}_{28}$
23	$\mathbf{S}_{22} \longrightarrow \mathbf{S}_{14} + \mathbf{S}_{22}$	54	$\mathbf{S}_{28} \longrightarrow \mathbf{S}_{15} + \mathbf{S}_{27}$
24	$\mathbf{S}_{14} \longrightarrow \emptyset$	55	$\mathbf{S}_5 \longrightarrow \mathbf{S}_1$
25	$\mathbf{S}_{17} \longrightarrow \mathbf{S}_{14} + P_{\text{folded}}$	56	$\mathbf{S}_{13} \longrightarrow \mathbf{S}_1 + \mathbf{S}_{11}$
26	$\mathbf{S}_{15} \longrightarrow \mathbf{S}_4$	57	$\mathbf{S}_9 \longrightarrow \mathbf{S}_7$
27	$\mathbf{S}_{19} \longrightarrow \mathbf{S}_4 + \mathbf{S}_{18}$	58	$\mathbf{S}_{15} \longrightarrow \mathbf{S}_{14}$
28	$\mathbf{S}_{28} \longrightarrow \mathbf{S}_4 + \mathbf{S}_{27}$	59	$\mathbf{S}_{19} \longrightarrow \mathbf{S}_{14} + \mathbf{S}_{18}$
29	$\mathbf{S}_{13} \longrightarrow \mathbf{S}_{13} + \mathbf{S}_{23}$	60	$\mathbf{S}_{28} \longrightarrow \mathbf{S}_{14} + \mathbf{S}_{27}$
30	$\mathbf{S}_{23} \longrightarrow \emptyset$	61	$\mathbf{S}_{21} \longrightarrow \mathbf{S}_{20}$
31	$\mathbf{S}_{23} \longrightarrow \mathbf{S}_{18} + \mathbf{S}_{23}$		

Table 8: Reactions for the Bacterial Heat-Shock response model (Kurata et al. [47]). These reactions follow the Species Chart in Table 7. The network is depicted in Figure 4A. Reaction 25 represents the refolding of denatured protein P when it is in complex with the chaperon protein $DnaK$. This reaction is the only way in which molecules of protein P exit the network.

Species Chart	
$\mathbf{S}_1 = M$	$\mathbf{S}_3 = A$
$\mathbf{S}_2 = T$	$\mathbf{S}_4 = P$

Table 9: Species chart for the Toxin-Antitoxin model. T and A denote the Toxin and Antitoxin proteins respectively. They are both translated by the *biscistronic* mRNA M which is produced constitutively by a gene. The Toxin protein converts to another protein P which inhibits the metabolism of the cell.

of two fast reactions

$$\mathbf{S}_1 \longrightarrow \mathbf{S}_1 + \mathbf{S}_2 + \mathbf{S}_3 \quad \text{and} \quad \mathbf{S}_2 + \mathbf{S}_3 \longrightarrow \emptyset, \quad (6.61)$$

assuming that all the other *slow* reactions are switched off and the state of the slow variables is fixed.

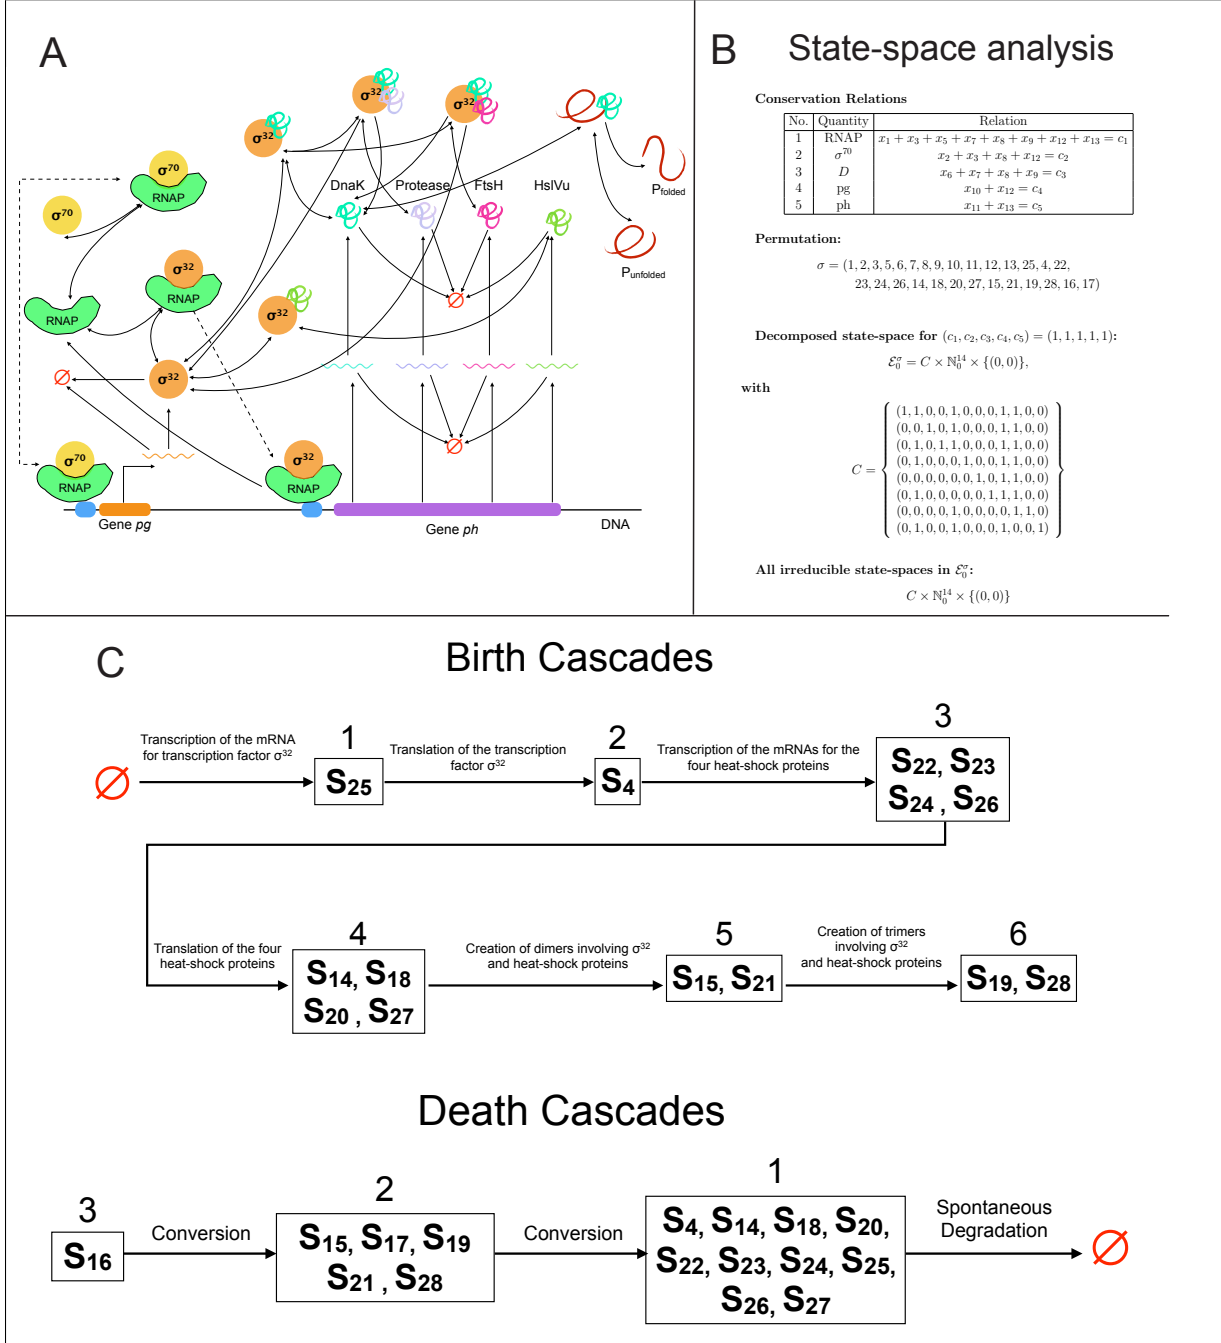


Figure 4: Results for the Bacterial Heat-Shock response model (Kurata et al. [47]) depicted in panel **A**. The mRNAs are not labeled to avoid clutter. The results of state-space analysis are presented in panel **B** and the birth-death cascades are shown in panel **C**. The birth-cascades signify various stages in the network. Note that in this network there are twelve *bounded* species ($\sigma(1), \dots, \sigma(12)$) and fourteen *free* species ($\sigma(13), \dots, \sigma(26)$).

Following [32], the slow variables in this network are x_1 (mRNA copy-number) and $y = x_2 - x_3$ (difference between Toxin and Antotoxin copy-numbers), because these variables are unaffected by the fast reactions in (6.61). Fixing x_1 and y we compute the decomposed state-space for this subnetwork using Algorithm 1. Due to the semi-positive conservation relation $\gamma_1 = (1, 0, 0)$, the species S_1 is *bounded* and its finite state-space

No.	Reaction	Propensity
1	$\emptyset \longrightarrow \mathbf{S}_1$	$\lambda_1(x) = \left(\frac{20}{1+x_3}\right)$
2	$\mathbf{S}_1 \longrightarrow \mathbf{S}_1 + \mathbf{S}_2 + \mathbf{S}_3$	$\lambda_2(x) = \theta_1 x_1$
3	$\mathbf{S}_2 + \mathbf{S}_3 \longrightarrow \emptyset$	$\lambda_3(x) = \theta_2 x_2 x_3$
4	$\mathbf{S}_1 \longrightarrow \emptyset$	$\lambda_4(x) = 2x_1$
5	$\mathbf{S}_2 \longrightarrow \emptyset$	$\lambda_5(x) = 5x_2$
6	$\mathbf{S}_3 \longrightarrow \emptyset$	$\lambda_6(x) = x_3$
7	$\mathbf{S}_2 \longrightarrow \mathbf{S}_4$	$\lambda_7(x) = 0.5x_2$
8	$\mathbf{S}_4 \longrightarrow \emptyset$	$\lambda_8(x) = 0.1x_4$

Table 10: Reactions for the Toxin-Antitoxin model according to the Species Chart in Table 9. The associated propensity functions λ_k -s are also provided. Here x_i denotes the copy-number of species \mathbf{S}_i . We choose $\theta_1 = 100$ and $\theta_2 = 10$, and hence reactions 2 and 3 are much *faster* in comparison to the other reactions. Reaction 3 is the annihilation reaction between Toxin and Antitoxin molecules, which actually represents formation of an inactive complex AT . In reactions 2 and 3, we choose parameters $\theta_1 = 100$ and $\theta = 10$, and so the subnetwork formed by these two reactions can be considered *fast* in the context of the larger network.

is simply the singleton $\{x_1\}$. The only other independent conservation relation is $\gamma_2 = (0, 1, -1)$, which is mixed-sign, and so there exists a *restricted* species. Algorithm 1 will automatically classify one of the remaining species as *free* and the other as *restricted*, depending on the value of y , in such a way that the associated affine map ϕ is compatible with the network (recall Assumption 4.1). If $y \geq 0$, then Algorithm 1 picks species \mathbf{S}_3 as *free*, species \mathbf{S}_2 as *restricted*, the affine map ϕ as $\phi(x_2) = x_3 + y$ and the permutation σ as $\sigma = (1, 3, 2)$. However these choices violate network compatibility when $y < 0$ and so in this situation Algorithm 1 picks species \mathbf{S}_2 as *free*, species \mathbf{S}_3 as *restricted*, the affine map ϕ as $\phi(x_3) = x_2 - y$ and the permutation σ as $\sigma = (1, 2, 3)$. For convenience, let \mathbf{S}_y , ϕ_y and σ_y denote the y -dependent choices of the *free* species, the affine map and the permutation σ respectively. Then according to Algorithm 1, the decomposed state-space for the fast subnetwork permuted with σ_y is simply $\{x_1\} \times \Phi_y$, where Φ_y is the graph (see (1.14)) of map ϕ_y .

As described in Section 4.1 we can reduce the fast subnetwork (6.61) by eliminating the *restricted* species to obtain the following network



where the propensity function for the first reaction is same as before, but for the second reaction it changes to

$$\tilde{\lambda}_y(x_1, z) = \theta_2 z \phi_y(z) = \theta_2 z(z + |y|). \quad (6.63)$$

Here z is the copy-number of the *free* species \mathbf{S}_y . Observe that the compatibility of map ϕ_y ensures that $\tilde{\lambda}_y(x_1, z) = 0$ if and only if $z = 0$, and hence the reduced network satisfies Assumption 2.1. This allows us to apply Algorithm 2 to find all the irreducible state-spaces for this reduced network in the state-space $\{x_1\} \times \mathbb{N}_0$ and this algorithm outputs that there is a single irreducible state-space which is identical to the full state-space $\{x_1\} \times \mathbb{N}_0$. Therefore the stationary distribution for this network must be exactly supported on $\{x_1\} \times \mathbb{N}_0$, and in fact using the results in [28] (see also Section 6.5) we can compute it to be exactly:

$$\tilde{\pi}_y(x_1, z) = \begin{cases} 1 & \text{if } x_1 = 0 \text{ and } z = 0 \\ \frac{1}{L} \left(\frac{\theta_1}{\theta_2}\right)^z x_1^z \left(\frac{|y|!}{z!(z+|y|)!}\right) & \text{if } x_1 > 0 \text{ and } z = 0, 1, 2, \dots \\ 0 & \text{otherwise,} \end{cases} \quad (6.64)$$

where $a!$ denotes the factorial of a and L is the normalization constant given by

$$L = \sum_{z=0}^{\infty} \left(\frac{\theta_1}{\theta_2}\right)^z x_1^z \left(\frac{|y|!}{z!(z+|y|)!}\right).$$

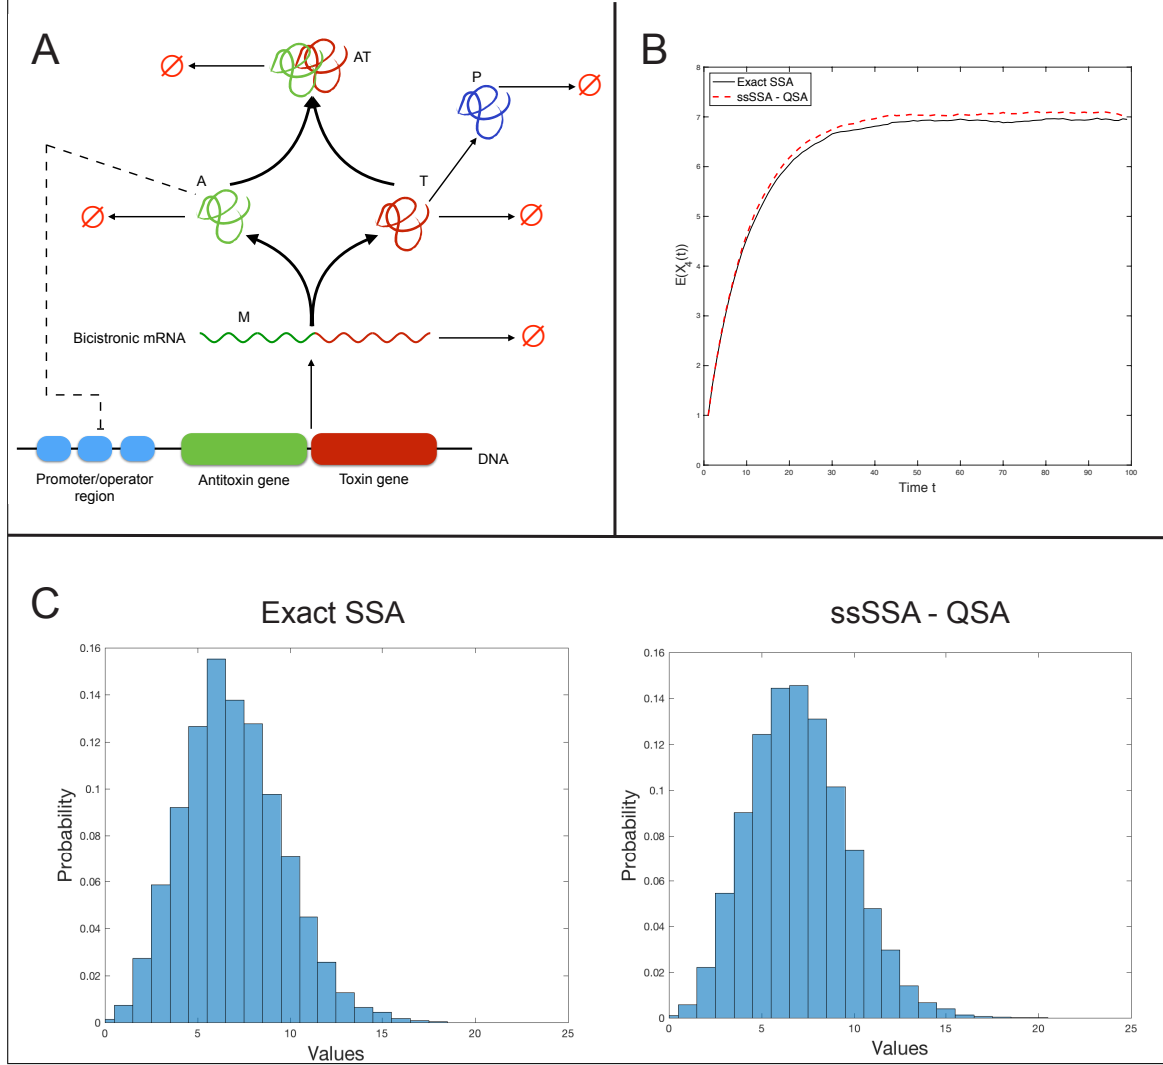


Figure 5: Stochastic simulations for the Toxin-Antitoxin model depicted in panel **A**. The *fast* reactions are indicated with a thicker arrow. In panel **B** the estimated mean-dynamics $\mathbb{E}(X_4(t))$ for copy-numbers of protein P is plotted, and in panel **C**, the estimated probability distributions (histograms) for these copy-numbers at the final time-point $T = 100$ are shown. All these estimations were performed with two simulation schemes - the Exact SSA [13] and the approximate ssSSA [33] that uses the quasi-stationary assumption (QSA). Observe that ssSSA is quite accurate and for this example, the simulations using ssSSA were about 6 times faster than those using the Exact SSA.

It is easy to check that $L < \infty$ and hence $\tilde{\pi}_y$ is a valid stationary distribution over $\{x_1\} \times \mathbb{N}_0$. The uniqueness of this stationary distribution is guaranteed because $\{x_1\} \times \mathbb{N}_0$ is the only irreducible state-space for the reduced network. Due to Proposition 4.2 we can conclude that $\{x_1\} \times \Phi_y$ is the only irreducible state-space for the full subnetwork (6.61) under permutation σ_y and its unique stationary distribution π_y on $\{x_1\} \times \Phi_y$ is just

$$\pi_y(x_1, z, \phi_y(z)) = \tilde{\pi}_y(x_1, z).$$

This also shows that the subnetwork (6.61) is ergodic which is necessary for the application of QSA. For the *slow* reactions $k \in \{1, 4, 5, 6, 7, 8\}$ the propensities needed for the application of QSA can be estimated as

$$\hat{\lambda}_k(x_1, y, x_4) = \sum_{z \in \mathbb{N}_0} \lambda_k(x_1, z, \phi_y(z), x_2) \tilde{\pi}_y(x_1, z).$$

Using these propensities we can simulate the Toxin-Antitoxin network with ssSSA [33] to estimate the probability distribution of copy-numbers of protein P (i.e. $X_4(t)$) as well as the mean dynamics of these copy-numbers. The results are reported in Figure 5, where results from the exact SSA simulations are also shown for comparison. One can see that ssSSA simulations are quite accurate thanks to the correct identification of the ergodic subnetwork and its stationary distribution. Moreover these ssSSA simulations are about 6 times faster than simulations with Exact SSA. This simple example nicely illustrates how our algorithms for state-space analysis can be integrated with a simulation scheme like ssSSA, to aid the application of QSA. Note that in more complicated examples, the exact form of the stationary distribution may not be available, but our analysis can verify its uniqueness and provide a description of its exact support. In such cases this knowledge can be combined with other simulation-based schemes to sample from the stationary distributions and estimate the slow propensities for applying QSA (see [34, 35]).

6.5 Networks with product-form stationary distributions

In this paper we present a method for finding irreducible state-spaces for networks where infinitely many states are accessible. On each of these irreducible state-spaces, the uniqueness of the stationary distribution is guaranteed but its existence needs to be checked by other means (like the analysis in [27]). This is complementary to certain other results in the literature which assume the knowledge of irreducible state-spaces and demonstrate the existence of product-form stationary distributions for a large class of networks [28]. Exploiting this complementarity, we now explore how the combination of our results with the results on product-form stationary distributions, can provide us with a complete characterization of the simplex of stationary distributions (see (1.15)) for several networks.

Consider a reaction network $\mathcal{N} = (\mathcal{V}, \mathcal{O}, \Lambda)$ with d species and K reactions of the form (1.1) (see Section 2). For now we assume that each propensity function λ_k has the mass-action form (2.20) with some rate constant $\theta_k > 0$. In the deterministic setting, the state of the network at time t is a vector of species concentrations $x(t) = (x_1(t), \dots, x_d(t)) \in \mathbb{R}_+^d$, which evolves according to the following ODE

$$\frac{dx}{dt} = \sum_{k=1}^K \theta_k \left(\prod_{i=1}^d x_i^{\nu_{ik}} \right) (\rho_k - \nu_k), \quad (6.65)$$

where $\nu_k = (\nu_{1k}, \dots, \nu_{dk})$ and $\rho_k = (\rho_{1k}, \dots, \rho_{dk})$ are vectors in \mathbb{N}_0^d denoting the k -th column of matrices \mathcal{V} and \mathcal{O} respectively. Observe that the k -th reaction in the network can simply be represented as $\nu_k \longrightarrow \rho_k$. With this association, ν_k and ρ_k are called network *complexes*. Let $\mathbb{C} = \{\nu_k, \rho_k : k = 1, \dots, K\}$ be the set of all network complexes and suppose there exists a strictly positive vector $r = (r_1, \dots, r_d) \in \mathbb{R}_+^d$ such that for each $z \in \mathbb{C}$ we have

$$\sum_{k=1, \nu_k=z}^K \theta_k \left(\prod_{i=1}^d r_i^{\nu_{ik}} \right) = \sum_{k=1, \rho_k=z}^K \theta_k \left(\prod_{i=1}^d r_i^{\rho_{ik}} \right).$$

This relation simply says that when the vector of species concentrations is r , the rate at which a complex z is consumed (l.h.s.) is same as the rate at which this complex is produced (r.h.s.). It can be shown that r is a fixed point for the deterministic state-dynamics (6.65), i.e. r.h.s of (6.65) is $\mathbf{0}$ when $x = r$. Hence the network is called *complex-balanced* and r is called a *complex-balanced* fixed point. This property has several implications regarding the topology of the network as well as the existence, uniqueness and stability of its fixed points (see [50, 51, 52]). The existence of a complex-balanced equilibrium can be verified computationally, but it can also be checked for many networks using the famous *deficiency zero theorem* (see [52]) from Chemical Reaction Network Theory.

In [28], Anderson et al. prove that if a network \mathcal{N} with mass-action propensities is complex-balanced in the sense described above, then there exists a *product-form* stationary distribution π for its stochastic model, on each irreducible state-space \mathcal{E} within the state-space $\mathcal{E}_0 \subset \mathbb{N}_0^d$ of the network. This stationary distribution is given by

$$\pi(x) = \frac{1}{M} \prod_{i=1}^d \frac{r_i^{x_i}}{x_i!} \quad \text{for all } x \in \mathcal{E}, \quad (6.66)$$

where $a!$ denotes the factorial of a , r is the complex-balanced fixed point and M is the normalizing constant given by

$$M = \sum_{x \in \mathcal{E}} \prod_{i=1}^d \frac{r_i^{x_i}}{x_i!}. \quad (6.67)$$

To use this elegant result one needs to know the irreducible state-spaces, which is precisely what is accomplished in this paper. In fact, the examples considered before suggest that for many networks we can provably find all the irreducible state-spaces $\mathcal{E}_1, \dots, \mathcal{E}_Q$ for the network using our main result Theorem 4.4. If this is true, then we can easily compute the corresponding product-form (6.66) stationary distributions supported on these classes and obtain the exact simplex Σ (see (1.15)) of stationary distributions for the network.

Note that if the whole nonnegative integer orthant $\mathcal{E} = \mathbb{N}_0^d$ is an irreducible state-space, then this is the only irreducible state-space, and the simplex Σ consists of just one stationary distribution π given by (6.66) with $\mathcal{E} = \mathbb{N}_0^d$ and $M = \exp(\sum_{i=1}^d r_i)$. In other words, π is just a product of Poisson distributions and in such a scenario, the species copy-numbers are independent at stationarity, with the copy-number distribution of species i being Poisson with mean r_i . As the species are constantly interacting through reactions, having this independence is quite remarkable, and it has been argued that this independence could play an important role in metabolic pathways [53].

We now demonstrate how for certain networks, our results can help in accurately computing the stationary networks by replacing a possibly infinite sum in (6.67) with a finite sum, thereby avoiding any truncation errors associated with the problem of estimating infinite sums. Assume that network \mathcal{N} is complex-balanced and it does not have any *restricted* species. Also suppose that using Theorem 4.4, we obtain the irreducible state-space $\mathcal{E} = \mathcal{E}_b \times \mathbb{N}_0^{d_1} \times \{0\}$, under some permutation σ which we can assume to be identity (i.e. $\sigma(i) = i$ for each i), without any loss of generality. Here \mathcal{E}_b is a finite set in $\mathbb{N}_0^{d_b}$, d_b is the number of *bounded* species and d_1 is some number less than the number of *free* species $d_f = (d - d_b)$. The decomposed form of \mathcal{E} shows that any element $x \in \mathcal{E}$ can be expressed as $x = (x_b, x_f, \mathbf{0})$, where $x_b = (x_1, \dots, x_{d_b}) \in \mathcal{E}_b$ and $x_f = (x_{d_b+1}, \dots, x_{d_b+d_1}) \in \mathbb{N}_0^{d_1}$. Therefore we can write (6.67) as

$$\begin{aligned} M &= \sum_{(x_b, x_f, \mathbf{0}) \in \mathcal{E}} \prod_{i=1}^d \frac{r_i^{x_i}}{x_i!} = \left(\sum_{x_b \in \mathcal{E}_b} \prod_{i=1}^{d_b} \frac{r_i^{x_i}}{x_i!} \right) \left(\sum_{x_f \in \mathbb{N}_0^{d_1}} \prod_{i=d_b+1}^{d_b+d_1} \frac{r_i^{x_i}}{x_i!} \right) \\ &= \exp \left(\sum_{i=d_b+1}^{d_b+d_1} r_i \right) \left(\sum_{x_b \in \mathcal{E}_b} \prod_{i=1}^{d_b} \frac{r_i^{x_i}}{x_i!} \right), \end{aligned}$$

which is a finite sum, that can be easily computed without incurring any truncation errors. In deriving the last relation, we have used that $\sum_{n=0}^{\infty} a^n/n! = \exp(a)$ for any $a \in \mathbb{R}$.

Until now we were assuming that propensity functions of the complex-balanced network \mathcal{N} satisfy mass-action kinetics. Theorem 6.1 in [28] relaxes this assumption of mass-action kinetics and proves the existence of product-form stationary distributions for more general kinetics, where each species i has a “rate of association” function $\kappa_i : \mathbb{N}_0 \rightarrow \mathbb{R}_+$. In this setting, the mass-action formula (2.20) for the propensity function changes to

$$\lambda_k(x) = \theta_k \prod_{i=1}^d \prod_{j=0}^{\nu_{ik}-1} \kappa_i(x_i - j)$$

and the product-form stationary distribution becomes

$$\pi(x) = \frac{1}{M} \prod_{i=1}^d \frac{r_i^{x_i}}{\prod_{j=1}^{x_i} \kappa_i(j)} \quad \text{for all } x \in \mathcal{E},$$

with the normalizing constant M chosen to ensure that $\sum_{x \in \mathcal{E}} \pi(x) = 1$. Observe that this result can be used to compute the stationary distribution (6.64) for network (6.62) in Section 6.4. Recently this result

has been extended even further to consider more general species-specific association rate functions (see [54]). Such results along with our computational framework for identifying irreducible state-spaces, provide a way for characterizing the stationary distributions for complex-balanced networks.

7 Conclusion

The aim of this paper is to provide a new tool for analyzing continuous-time Markov chain (CTMC) models of biomolecular reaction networks. Specifically we are interested in situations where the state-space for the CTMC needs to be countably infinite due to the lack of a global conservation relationship among all the species. Such situations arise frequently in Systems Biology as stochastic models generally describe the activity of a small subnetwork embedded within a larger network. We develop a simple procedure to systematically explore the space of conservation relations among species and represent the state-space of the CTMC in a special decomposed-form based on the copy-number ranges of all the species. This form can help in assessing the reachability relations and the communication structures within the infinite state-space of the underlying CTMC. In this context, the main goal of this paper is to construct a computational method for finding all the *closed communication classes* for the CTMC. Such classes are natural attracting sets for the dynamics and they can also be viewed as *irreducible state-spaces* for the CTMC (see Section 2.1). Under the existence of a suitable Foster-Lyapunov function (see [27, 26] and Section 1), each irreducible state-space supports a unique stationary distribution and these distributions form the vertices of the full simplex (see (1.15)) of stationary distributions of the CTMC.

As we discuss in this paper, finding irreducible state-spaces for networks where infinitely many states are accessible, is a challenging but a very important problem for several reasons. These reasons include understanding the stability and ergodic properties of networks (see Section 1), analyzing network topologies using methods from *Transition Path Theory* [31] and in obtaining the exact stationary distribution for a large class of networks (see Section 6.5). Furthermore our computational procedure for finding irreducible state-spaces can assist in an automated discovery of *fast* ergodic subnetworks for the application of *quasi-stationary assumption* (QSA) in the simulation of multiscale reaction networks (see Section 6.4 and [33, 34, 35]). Using our procedure we demonstrate in Section 6 that generally networks from Systems Biology admit a single irreducible state-space which corresponds to the situation where the CTMC describing the reaction dynamics is ergodic and has a unique stationary distribution. We discuss how this information along with the structure of irreducible state-space can sometimes provide valuable biological insights into the design and functionality of the underlying network (see Sections 6.2 and 6.3).

Our approach works by classifying each species as *free*, *bounded* and *restricted*, based on their admissible copy-number ranges. The *bounded* species have a finite copy-number range and their dynamics evolves in a finite set \mathcal{E}_b . On the other hand the copy-number range of *free* species is the infinite set of all non-negative integers \mathbb{N}_0 , and hence their dynamics evolves in an orthant \mathbb{N}_0^{df} . The *restricted* species imitate the dynamics of *free* species according to some affine function, and they can be removed from the network for the purpose of finding irreducible state-spaces (see Section 4.1). We develop a computational procedure that can provably locate all the irreducible state-spaces for the underlying CTMC within the infinite state-space $\mathcal{E}_b \times \mathbb{N}_0^{df}$. This is accomplished by suitably combining the matrix methods used in the finite state-space case where only *bounded* species are present (see Section 4.2), along with the construction of *birth-death cascades* for the *free* species (see Section 4.4). We demonstrate the versatility of our method through many examples from Systems Biology in Section 6. From these examples one can also observe that the birth-death cascades correspond naturally to various important stages in the network and hence this cascade construction process facilitates a better understanding of the network design.

Finally we would like to mention that since our computational procedure only involves basic linear-algebraic tasks (such as matrix computations, solving linear equations and Linear Programs etc.), it can be efficiently applied in very high dimensions. Hence our method can easily handle large reaction networks with several species and reactions. However computational issues may arise if the size N_b of the finite state-space \mathcal{E}_b for *bounded* species becomes too large, as our method requires several computations with matrices of size $N_b \times N_b$. We hope to resolve these issues in a future work.

Appendix: Proof of Theorem 4.4

Let \mathbb{L}_{\min} be the set of all minimal leaf nodes of the Birth-Cascade Tree (BCT) constructed in Section 4.4. We pick any $(C, A) \in \mathbb{L}_{\min}$ and consider the dynamics of network $\tilde{\mathcal{N}}^\sigma$ under the permutation $\sigma = \sigma_3(A)$ mentioned in the statement of Theorem 4.4. Consider the directed path (4.58) in BCT, starting with the node (C_0, \emptyset) and terminating at the leaf node (C, A) . The next lemma is a simple consequence of the construction of BCT.

Lemma 7.1 *Pick any $z_1 \in C_0$, $z_2 \in C$ and $r_0 \in \mathbb{N}_0^{|A|}$. Then there exists a vector $x \in \mathbb{N}_0^{|A|}$ such that $x \geq r_0$ and*

$$(z_1, \mathbf{0}) \xrightarrow{\tilde{\mathcal{N}}^\sigma} (z_2, x, \mathbf{0}).$$

Moreover this relation also holds for any $z_1 \in C$.

Proof. Throughout this proof we denote the relation $\xrightarrow{\tilde{\mathcal{N}}^\sigma}$ by \longrightarrow . Let C_i -s and G_i -s be as in the path (4.58) consisting of $(l+1)$ nodes. Let $d_i = |G_i|$ for each i and since G_i -s in an increasing family of sets we must have

$$0 = d_0 < d_1 < d_2 < \dots < d_l = |A|.$$

By permuting the network if necessary, we can assume that $G_i = \{1, \dots, d_i\}$ for each i .

Clearly the case $l = 0$ is trivial because in this case $C_0 = C$ and $A = \emptyset$. We now consider the case $l = 1$ and $A = G_1$. Due to reactions in the set $\mathcal{K}_r(\emptyset)$, the dynamics of *bounded* species can reach any state in C_0 from any other state in C_0 . The construction of BCT ensures that for each *free* species \mathbf{S}_i in $\mathcal{D}_f^\sigma(A)$, there exists a state $y \in C_0$ and a reaction $\mathcal{K}_r(y, \emptyset) \subset \mathcal{K}_r(\emptyset)$, which produces this *free* species. By repeated firings of this reaction we can push the molecular-count of species \mathbf{S}_i beyond any positive integer. Note that reactions in $\mathcal{K}_r(\emptyset)$ do not consume any *free* species. Hence we can perform this procedure independently for all the *free* species in $\mathcal{D}_f^\sigma(A)$, and prove that for any $r \in \mathbb{N}_0^{d_1}$ there exists a vector $x' \in \mathbb{N}_0^{d_1}$ along with some state $y \in C_0$ such that $x' \geq r$ and

$$(z_1, \mathbf{0}) \longrightarrow (y, x', \mathbf{0}).$$

However since $C \in \Psi_{(\emptyset, A_1)}(C_0)$, after finitely many reactions in $\mathcal{K}_r(A)$ that only consume the *free* species in $\mathcal{D}_f^\sigma(A)$, the state of the *bounded* species can go from y to $z_2 \in C$, thereby ensuring that

$$(y, x', \mathbf{0}) \longrightarrow (z_2, x, \mathbf{0})$$

for some $x \in \mathbb{N}_0^{d_1}$ with $x \leq x'$. By choosing r with large enough entries we can ensure that $x \geq r_0$. The assertion of this lemma (i.e. $(z_1, \mathbf{0}) \longrightarrow (z_2, x, \mathbf{0})$) then follows for $l = 1$ from the transitivity of relation \longrightarrow (see Section 2.1).

For a general $l > 1$, one can repeat the above arguments at each stage $i = 1, 2, \dots, (l-1)$ to prove the lemma's assertion. At each stage we rely on the fact that C_i is a closed communication class under reactions in $\mathcal{K}_r(G_i)$ and these reactions do not consume any *free* species outside the set $\mathcal{D}_f^\sigma(G_i)$.

Now suppose $z_1 \in C$. After finitely many reactions in the set $\mathcal{K}_r(\emptyset)$, the state of *bounded* species will reach a state z'_1 in some closed communication class $C'_0 \in \mathcal{C}(\emptyset)$. As $z_1 \in C$ and $\mathcal{K}_r(\emptyset) \subset \mathcal{K}_r(A)$ we must have $C'_0 \subset C$. There exists a BCT-path of the form

$$(C'_0, \emptyset) = (C'_0, G'_0) \Rightarrow (C'_1, G'_1) \Rightarrow \dots \Rightarrow (C'_m, G'_m) = (C', A'),$$

culminating in the leaf node $(C', A') \in \mathbb{L}$. Note that as $C'_0 \subset C$ we have $\mathcal{K}_r(C'_0, \emptyset) \subset \mathcal{K}_r(C, A)$ and hence $C'_1 \subset C$ and $G'_1 \subset A$. Repeating this argument $(m-1)$ times we can conclude that $C'_m = C' \subset C$ and $G'_m = A' \subset A$. However as (C, A) is a minimal leaf node we must have $C' = C$ and $A' = A$. The result now follows from the assertion already proven above. \square

For the minimal leaf node $(C, A) \in \mathbb{L}_{\min}$, let $\tilde{\mathcal{N}}^\sigma(C, A)$ be the network formed by restricting the *free* species to the set $\mathcal{D}_f^\sigma(A)$, the reactions to the set $\mathcal{K}_r(C, A)$ and the *bounded* species state-space to the set C . Let $\mathbb{L}_d(C, A)$ denote the set of all leaf nodes for the Death-Cascade Tree (DCT) for network $\tilde{\mathcal{N}}^\sigma(C, A)$. The following lemma pertains to the situation when the leaf node (C, A) is death-exhaustive.

Lemma 7.2 Suppose that the minimal leaf node $(C, A) \in \mathbb{L}_{\min}$ is death-exhaustive, i.e. there exists a node $(C', A') \in \mathbb{L}_d(C, A)$ such that $A = A'$. Then $C' = C$ and for any $z_1, z_2 \in C$ and $r_0 \in \mathbb{N}_0^{|A|}$ there exists a vector $x \in \mathbb{N}_0^{|A|}$ such that $x \geq r_0$ and

$$(z_1, x, \mathbf{0}) \xrightarrow{\tilde{\mathcal{N}}^\sigma} (z_2, \mathbf{0}).$$

Proof. Let $\tilde{\mathcal{N}}_{\text{inv}}^\sigma(C, A)$ be the inverse of network $\tilde{\mathcal{N}}^\sigma(C, A)$. Note that since (C, A) is a leaf node for BCT, relation (4.57) holds and hence the reactions in $\mathcal{K}_r(C, A)$ do not produce any *free* species outside the set $\mathcal{D}_f^\sigma(A)$. This implies that in the inverse network $\tilde{\mathcal{N}}_{\text{inv}}^\sigma(C, A)$ no reaction in $\mathcal{K}_r(C, A)$ can consume any *free* species outside the set $\mathcal{D}_f^\sigma(A)$. Hence all the reactions in $\mathcal{K}_r(C, A)$ can fire in this inverse network as it satisfies Assumption 2.1. Since communication is a symmetric relation and $C' \subset C$ is a closed communication class, $A' = A$ implies that $C' = C$. By definition, the DCT for network $\tilde{\mathcal{N}}^\sigma(C, A)$ is the BCT of network $\tilde{\mathcal{N}}_{\text{inv}}^\sigma(C, A)$. Consider the directed path (4.58) in this BCT, starting with the node (C_0, \emptyset) and terminating at the leaf node (C, A) . Pick any $z_1, z_2 \in C$ and $r_0 \in \mathbb{N}_0^{|A|}$. Applying Lemma 7.1 we can conclude that there exists a vector $x \in \mathbb{N}_0^{|A|}$ such that $x \geq r_0$ and

$$(z_2, \mathbf{0}) \xrightarrow{\tilde{\mathcal{N}}_{\text{inv}}^\sigma(C, A)} (z_1, x, \mathbf{0}),$$

which also implies

$$(z_1, x, \mathbf{0}) \xrightarrow{\tilde{\mathcal{N}}^\sigma} (z_2, \mathbf{0}),$$

due to relation (2.22) and the fact that network $\tilde{\mathcal{N}}^\sigma(C, A)$ is simply a restriction of the network $\tilde{\mathcal{N}}^\sigma$. This completes the proof of this lemma. \square

Lemma 7.3 Let $(C, A) \in \mathbb{L}_{\min}$ be a minimal leaf node such that all the free species in $\mathcal{D}_f^\sigma(A)$ are singularly-degradable w.r.t. A . Then there exists a $r_0 \in \mathbb{N}_0^{|A|}$ such that for any $z \in C$ and $x_1, x_2 \in \mathbb{N}_0^{|A|}$ satisfying $x_1 \geq x_2 \geq r_0$ we have

$$(z, x_1, \mathbf{0}) \xrightarrow{\tilde{\mathcal{N}}^\sigma} (z, x_2, \mathbf{0}).$$

Proof. Throughout this proof we denote the relation $\xrightarrow{\tilde{\mathcal{N}}^\sigma}$ by \longrightarrow . Without loss of generality we can assume that $A = \{1, 2, \dots, |A|\}$. The proof of this lemma is inspired by the proof of Theorem 3.4 in [55]. Fixing a $z \in C$, we first show that for each $i \in A$ there exists a $r_i \in \mathbb{N}_0^{|A|}$ such that

$$(z, r_i, \mathbf{0}) \longrightarrow (z, r_i - e_i, \mathbf{0}), \quad (7.68)$$

where e_i is the i -th standard basis vector in $\mathbb{R}^{|A|}$.

Pick any $i \in A$ and let $\tilde{\mathcal{K}}_p(A)$ be the set of reactions given by (4.59). Since each species in $\mathcal{D}_f^\sigma(A)$ is singularly-degradable, there exists a sequence of reactions $k_1, \dots, k_n \in \tilde{\mathcal{K}}_p(A)$ such that

$$(-e_i, \mathbf{0}) = \sum_{j=1}^n (\tilde{\rho}_{k_j}^\sigma - \tilde{\nu}_{k_j}^\sigma), \quad (7.69)$$

where $\tilde{\nu}_k^\sigma$ and $\tilde{\rho}_k^\sigma$ are as in Section 4.4. Note that since for each reaction k_j we have $\text{supp}(\tilde{\rho}_{k_j}^\sigma) \subset A$, we must also have $\text{supp}(\tilde{\nu}_{k_j}^\sigma) \subset A$, or otherwise the last $(d_f - |A|)$ components in the r.h.s. of (7.69) cannot be $\mathbf{0}$. Therefore each of these reactions can only consume the *free* species in $\mathcal{D}_f^\sigma(A)$. For each $m = 1, \dots, n$ let

$$y_m = \sum_{j=1}^{m-1} (\tilde{\rho}_{k_j}^\sigma - \tilde{\nu}_{k_j}^\sigma).$$

By choosing a $r_i \in \mathbb{N}_0^{|A|}$ with large enough entries we can ensure that $(r_i, \mathbf{0}) + y_m \geq \tilde{\nu}_{k_m}^\sigma$ for each m . Since none of these reactions involve the *bounded* species and Assumption 2.1 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 $(r_i, \mathbf{0}) + y_m$, thereby implying that (7.68) is satisfied.

We find such a r_i for each $i \in A$ and compute the maximum $r_0 := \max_{i \in A} \{r_i\}$ of these vectors in the componentwise sense. Using Proposition 2.2 we can conclude that for any $x \in \mathbb{N}_0^{|A|}$ satisfying $x \geq r_0$ we have

$$(z, x, \mathbf{0}) \longrightarrow (z, x - e_i, \mathbf{0}) \quad (7.70)$$

for each $i \in A$. Now select any $x_1, x_2 \in \mathbb{N}_0^{|A|}$ satisfying $x_1 \geq x_2 \geq r_0$ and let $\alpha = (x_1 - x_2) \in \mathbb{N}_0^{|A|}$. We can express α as the sum

$$\alpha = \sum_{i=1}^{|A|} \alpha_i e_i.$$

Exploiting the transitivity of relation \longrightarrow and using (7.70), α_i times for each i we obtain the accessibility chain $(z, x_2 + \alpha, \mathbf{0}) \longrightarrow (z, x_2 + \alpha - \alpha_1 e_1, \mathbf{0}) \longrightarrow (z, x_2 + \alpha - \alpha_1 e_1 - \alpha_2 e_2, \mathbf{0}) \longrightarrow \cdots \longrightarrow (z, x_2 + \alpha - \sum_{i=1}^{|A|} \alpha_i e_i, \mathbf{0})$. But $x_2 + \alpha = x_1$ and $(x_2 + \alpha - \sum_{i=1}^{|A|} \alpha_i e_i) = x_2$ and hence the proof of this lemma is complete. \square

Next we present a simple modification of the above lemma in the case where all the *free* species are singularly-degradable w.r.t. A .

Lemma 7.4 *Let $(C, A) \in \mathbb{L}_{\min}$ be a minimal leaf node such that all the free species in \mathcal{D}_f are singularly-degradable w.r.t. A . Then there exists a $r_0 \in \mathbb{N}_0^{|A|}$ such that for any $z \in C$, $y \in \mathbb{N}_0^{d_f - |A|}$ and any $x \in \mathbb{N}_0^{|A|}$ satisfying $x \geq r_0$ we have*

$$(z, x, y) \xrightarrow{\tilde{N}^\sigma} (z, x, \mathbf{0}).$$

Proof. The proof is similar to the proof of Lemma 7.3. As before we denote the relation $\xrightarrow{\tilde{N}^\sigma}$ by \longrightarrow and assume that $A = \{1, 2, \dots, |A|\}$. Fix a $z \in C$. We first show that for each $i \in \{|A| + 1, \dots, d_f\}$ there exists a $r_i \in \mathbb{N}_0^{|A|}$ such that

$$(z, r_i, e_i) \longrightarrow (z, r_i, \mathbf{0}), \quad (7.71)$$

where e_i is the i -th standard basis vector in $\mathbb{R}^{d_f - |A|}$.

Fix any $i \in \{|A| + 1, \dots, d_f\}$ and since all the *free* species in \mathcal{D}_f are singularly-degradable w.r.t. A , there exists a sequence of reactions $k_1, \dots, k_n \in \tilde{\mathcal{K}}_p(A)$ such that

$$(\mathbf{0}, -e_i) = \sum_{j=1}^n (\tilde{\rho}_{k_j}^\sigma - \tilde{\nu}_{k_j}^\sigma).$$

Note that for each reaction k_j we have $\text{supp}(\tilde{\rho}_{k_j}^\sigma) \subset A$. Hence there must exist a $n_0 \in \{1, \dots, n\}$ such that $\text{supp}(\tilde{\nu}_{k_j}^\sigma) \subset A$ for all $j \neq n_0$ and for reaction k_{n_0} the following must be satisfied:

$$\text{supp}(\tilde{\nu}_{k_{n_0}}^\sigma) \subset A \cup \{i\} \quad \text{and} \quad \tilde{\nu}_{ik_{n_0}}^\sigma = 1.$$

For each $m = 1, \dots, n$ let

$$y_m = \sum_{j=1}^{m-1} (\tilde{\rho}_{k_j}^\sigma - \tilde{\nu}_{k_j}^\sigma).$$

By choosing a $r_i \in \mathbb{N}_0^{|A|}$ with large enough entries we can ensure that $(r_i, e_i) + y_m \geq \tilde{\nu}_{k_m}^\sigma$ for each m . Since none of these reactions involve the *bounded* species and Assumption 2.1 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 $(r_i, e_i) + y_m$, thereby implying that (7.71) is satisfied.

We find such a r_i for each $i \in \{|A| + 1, \dots, d_f\}$ and compute the maximum $r_0 := \max\{r_i\}$ in the componentwise sense. Using Proposition 2.2 we can conclude that for any $x \in \mathbb{N}_0^{|A|}$ satisfying $x \geq r_0$ and any $y = (y_1, \dots, y_{d_f - |A|}) \in \mathbb{N}_0^{d_f - |A|}$ we have

$$(z, x, y) \longrightarrow (z, x, y - e_i) \quad (7.72)$$

as long as $y_i \geq 1$. Exploiting the transitivity of relation \longrightarrow and using (7.72), y_i times for each i we obtain the accessibility chain $(z, x, y) \longrightarrow (z, x, y - y_1 e_1) \longrightarrow (z, x, y - y_1 e_1 - y_2 e_2) \longrightarrow \dots \longrightarrow (z, x, y - y_1 e_1 - y_2 e_2 - \dots - y_{d_f - |A|} e_{d_f - |A|})$. But $y - y_1 e_1 - y_2 e_2 - \dots - y_{d_f - |A|} e_{d_f - |A|} = \mathbf{0}$ and hence the proof of this lemma is complete. \square

We are now ready to prove Theorem 4.4. We shall use the fact that irreducible state-spaces for network $\tilde{\mathcal{N}}^\sigma$ must be necessarily disjoint (see Section 2.1).

Proof.[Proof of Theorem 4.4] Throughout the proof we denote the relation $\xrightarrow{\tilde{\mathcal{N}}^\sigma}$ by \longrightarrow . We pick any minimal leaf node $(C, A) \in \mathbb{L}_{\min}$ and we suppose that this node is death-exhaustive and all the free species in $\mathcal{D}_f(A)$ are singularly-degradable w.r.t. A . We now show that $C \times \mathbb{N}_0^{|A|} \times \{\mathbf{0}\}$ is an irreducible state-space for network $\tilde{\mathcal{N}}^\sigma$ permuted according to the permutation $\sigma = \sigma_3(A)$ defined in Section 4.4. To prove this assertion it suffices to show that for any $z_1, z_2 \in C$ and $x \in \mathbb{N}_0^{|A|}$ we have

$$(z_1, \mathbf{0}) \longrightarrow (z_2, x, \mathbf{0}) \longrightarrow (z_1, \mathbf{0}). \quad (7.73)$$

This is because if (7.73) holds then for any $(z_1, x_1, \mathbf{0}), (z_2, x_2, \mathbf{0}) \in C \times \mathbb{N}_0^{|A|} \times \{\mathbf{0}\}$ we have $(z_1, x_1, \mathbf{0}) \longrightarrow (z_2, x_2, \mathbf{0})$ due to the following chain of accessibility relations $(z_1, x_1, \mathbf{0}) \longrightarrow (z_1, \mathbf{0}) \longrightarrow (z_2, x_2, \mathbf{0})$ and the transitivity of \longrightarrow .

We now prove (7.73) for a fixed $z_1, z_2 \in C$ and $x \in \mathbb{N}_0^{|A|}$. Since the leaf node (C, A) is death-exhaustive, Lemma 7.2 implies that for any $r_0 \in \mathbb{N}_0^{|A|}$ there exists a vector $x' \in \mathbb{N}_0^{|A|}$ satisfying $x' \geq r_0$ and

$$(z_2, x', \mathbf{0}) \longrightarrow (z_2, \mathbf{0}). \quad (7.74)$$

We assume that r_0 is as in Lemma 7.3. Due to Proposition 2.2, relation (7.74) implies $(z_2, x + x', \mathbf{0}) \longrightarrow (z_2, x, \mathbf{0})$. Using Lemma 7.1 we can find a vector $y \in \mathbb{N}_0^{|A|}$ satisfying $y \geq (x + x')$ and $(z_1, \mathbf{0}) \longrightarrow (z_2, y, \mathbf{0})$. Since $(x + x') \geq r_0$, Lemma 7.3 implies that $(z_2, y, \mathbf{0}) \longrightarrow (z_2, x + x', \mathbf{0})$. We now have an accessibility chain $(z_1, \mathbf{0}) \longrightarrow (z_2, y, \mathbf{0}) \longrightarrow (z_2, x + x', \mathbf{0}) \longrightarrow (z_2, x, \mathbf{0})$, and since \longrightarrow is transitive we have the first relation in (7.73).

We now show the converse (i.e. $(z_2, x, \mathbf{0}) \longrightarrow (z_1, \mathbf{0})$). Using Lemma 7.1 and 7.2, we can find vectors $x_1, x_2 \in \mathbb{N}_0^{|A|}$ such that $x_1 \geq x_2 \geq \mathbf{0}$, $(z_2, x_2, \mathbf{0}) \longrightarrow (z_1, \mathbf{0})$ and $(z_2, \mathbf{0}) \longrightarrow (z_2, x_1, \mathbf{0})$. The last relation also implies $(z_2, x, \mathbf{0}) \longrightarrow (z_2, x_1 + x, \mathbf{0})$ due to Proposition 2.2. Since $(x_1 + x) \geq x_2 \geq r_0$ we have $(z_2, x_1 + x, \mathbf{0}) \longrightarrow (z_2, x_2, \mathbf{0})$ due to Lemma 7.3. This gives us the following chain of accessibility relations $(z_2, x, \mathbf{0}) \longrightarrow (z_2, x_1 + x, \mathbf{0}) \longrightarrow (z_2, x_2, \mathbf{0}) \longrightarrow (z_1, \mathbf{0})$ which shows the second relation of (7.73) and proves that $C \times \mathbb{N}_0^{|A|} \times \{\mathbf{0}\}$ is an irreducible state-space for network $\tilde{\mathcal{N}}^\sigma$.

Assuming that all the free species in the set \mathcal{D}_f are singularly-degradable w.r.t. A , we now prove that this is the network's only irreducible state-space that can contain elements in the set $C \times \mathbb{N}_0^{d_f}$ under the permutation $\sigma = \sigma_3(A)$. We prove this claim by contradiction by showing that if $\mathcal{E} \subset \mathcal{E}_b^\sigma \times \mathbb{N}_0^{d_f}$ is any irreducible state-space for network $\tilde{\mathcal{N}}^\sigma$ satisfying $\mathcal{E} \cap (C \times \mathbb{N}_0^{d_f}) \neq \emptyset$ then it cannot be disjoint from $C \times \mathbb{N}_0^{|A|} \times \{\mathbf{0}\}$. Let (z, x) be any state in \mathcal{E} , and as this set is irreducible, any state that can be accessed by (z, x) with relation \longrightarrow , must also be in \mathcal{E} . We can write $x \in \mathbb{N}_0^{d_f}$ as $x = (x_1, x_2)$ where $x_1 \in \mathbb{N}_0^{|A|}$ and $x_2 \in \mathbb{N}_0^{d_f - |A|}$. Let $r_0 \in \mathbb{N}_0^{|A|}$ be as in Lemma 7.4. Using Lemma 7.1 and Proposition 2.2 if necessary, we can assume that $(z, x_1, x_2) \longrightarrow (z, x'_1, x_2)$ for some $x'_1 \geq r_0$. Lemma 7.4 then implies that $(z, x'_1, x_2) \longrightarrow (z, x'_1, \mathbf{0})$, which also means that $(z, x'_1, \mathbf{0}) \in \mathcal{E}$. But $(z, x'_1, \mathbf{0}) \in C \times \mathbb{N}_0^{|A|} \times \{\mathbf{0}\}$ and hence \mathcal{E} cannot be disjoint from $C \times \mathbb{N}_0^{|A|} \times \{\mathbf{0}\}$, giving a contradiction. This completes the proof of part (A) of Theorem 4.4.

We now prove part (B). Let $\mathcal{E} \subset \mathcal{E}_b^\sigma \times \mathbb{N}_0^{d_f}$ be a nonempty irreducible state-space for network $\tilde{\mathcal{N}}^\sigma$. Let (z, x) be any state in \mathcal{E} . After finitely many reactions in the set $\mathcal{K}_r(\emptyset)$, the state of *bounded* species will

reach a state z' in some closed communication class $C'_0 \in \mathcal{C}(\emptyset)$. Hence there exists a $x' \in \mathbb{N}_0^{d_f}$ such that $(z, x) \longrightarrow (z', x') \in \mathcal{E}$. Moreover there exists a BCT-path of the form

$$(C'_0, \emptyset) = (C'_0, G'_0) \Rightarrow (C'_1, G'_1) \Rightarrow \cdots \Rightarrow (C'_m, G'_m) = (C', A'),$$

culminating in the leaf node $(C', A') \in \mathbb{L}$. Since $\mathbb{L} = \mathbb{L}_{\min}$ the leaf node (C', A') is also minimal. We view the dynamics of network $\tilde{\mathcal{N}}^\sigma$ under the permutation $\sigma = \sigma_3(A')$ defined in Section 4.4. Pick any $z_1 \in C'$. Due to Lemma 7.1 and Proposition 2.2 we know that for any $r_0 \in \mathbb{N}_0^{|A'|}$, there exists a $x_1 \in \mathbb{N}_0^{|A'|}$ and $x_2 \in \mathbb{N}_0^{d_f - |A'|}$ such that $(z', x') \longrightarrow (z_1, x_1, x_2) \in \mathcal{E}$. However (z_1, x_1, x_2) is an element in $C' \times \mathbb{N}_0^{|A'|} \times \{\mathbf{0}\}$ and hence, as argued previously, we arrive at a contradiction unless $\mathcal{E} = C' \times \mathbb{N}_0^{|A'|} \times \{\mathbf{0}\}$. This completes the proof of part (B) of Theorem 4.4. \square

References

- [1] H. Hethcote. The mathematics of infectious diseases. *SIAM Review*, 42(4):599–653, 2000.
- [2] Seth I. Berger and Ravi Iyengar. Network analyses in systems pharmacology. *Bioinformatics*, 25(19):2466–2472, 2009.
- [3] Uri 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.
- [4] Jason A. Papin, Tony Hunter, Bernhard O. Palsson, and Shankar Subramaniam. Reconstruction of cellular signalling networks and analysis of their properties. *Nat Rev Mol Cell Biol*, 6(2):99–111, 02 2005.
- [5] 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.
- [6] Timothy S. Gardner, Charles R. Cantor, and James J. Collins. Construction of a genetic toggle switch in escherichia coli. *Nature*, 403(6767):339–342, 2000.
- [7] Harley H. McAdams and Adam Arkin. Stochastic mechanisms in gene expression. *Proc. Natl. Acad. Sci., Biochemistry*, 94:814–819, 1997.
- [8] Michael B. Elowitz, Arnold J. Levine, Eric D. Siggia, and Peter S. Swain. Stochastic gene expression in a single cell. *Science*, 297(5584):1183–1186, 2002.
- [9] John Goutsias. Classical versus stochastic kinetics modeling of biochemical reaction systems. *Biophysical Journal*, 92(7):2350–2365, 2007.
- [10] D.A. Anderson and T.G. Kurtz. Continuous time Markov chain models for chemical reaction networks. In H. Koepl, G. Setti, M. di Bernardo, and D. Densmore, editors, *Design and Analysis of Biomolecular Circuits*. Springer-Verlag, 2011.
- [11] David Schnoerr, Guido Sanguinetti, and Ramon Grima. Approximation and inference methods for stochastic biochemical kineticsa tutorial review. *Journal of Physics A: Mathematical and Theoretical*, 50(9):093001, 2017.
- [12] Adam P. Arkin, Christopher V. Rao, and Denise M. Wolf. Control, exploitation and tolerance of intracellular noise. *Nature*, 420:231–237, 2002.
- [13] Daniel T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *The Journal of Physical Chemistry*, 81(25):2340–2361, 1977.

- [14] 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.
- [15] Daniel T. Gillespie. A rigorous derivation of the chemical master equation. *Physica A: Statistical Mechanics and its Applications*, 188(13):404 – 425, 1992.
- [16] B. Munsky and M. Khammash. The finite state projection algorithm for the solution of the chemical master equation. *Journal of Chemical Physics*, 124(4), 2006.
- [17] John G. Kemeny and J. Laurie Snell. *Finite Markov chains*. The University Series in Undergraduate Mathematics. D. Van Nostrand Co., Inc., Princeton, N.J.-Toronto-London-New York, 1960.
- [18] J. R. Norris. *Markov chains*, volume 2 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 1998. Reprint of 1997 original.
- [19] Stefan Schuster and Thomas Höfer. Determining all extreme semi-positive conservation relations in chemical reaction systems: a test criterion for conservativity. *Journal of the Chemical Society, Faraday Transactions*, 87(16):2561–2566, 1991.
- [20] Stephen Boyd and Lieven Vandenbergh. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004.
- [21] J. Hasenauer, V. Wolf, A. Kazerooni, and F. J. Theis. Method of conditional moments (mcm) for the chemical master equation. *Journal of Mathematical Biology*, 69(3):687–735, 2014.
- [22] Benjamin Hepp, Ankit Gupta, and Mustafa Khammash. Adaptive hybrid simulations for multiscale stochastic reaction networks. *The Journal of chemical physics*, 142(3):034118, 2015.
- [23] 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.
- [24] Geoffrey R. Grimmett and David R. Stirzaker. *Probability and random processes*. Oxford University Press, New York, third edition, 2001.
- [25] Sean Meyn and Richard L. Tweedie. *Markov chains and stochastic stability*. Cambridge University Press, Cambridge, second edition, 2009. With a prologue by Peter W. Glynn.
- [26] Sean 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.
- [27] Ankit Gupta, Corentin Briat, and Mustafa Khammash. A scalable computational framework for establishing long-term behavior of stochastic reaction networks. *PLoS Comput Biol*, 10(6):e1003669, 06 2014.
- [28] David F. Anderson, Gheorghe Craciun, and Thomas G. Kurtz. Product-form stationary distributions for deficiency zero chemical reaction networks. *Bulletin of Mathematical Biology*, 72(8):1947–1970, 2010.
- [29] Ankit Gupta, Jan Mikelson, and Mustafa Khammash. A finite state projection algorithm for the stationary solution of the chemical master equation. *The Journal of Chemical Physics*, 147(15):154101, 2017.
- [30] Corentin Briat, Ankit Gupta, and Mustafa Khammash. Antithetic integral feedback ensures robust perfect adaptation in noisy biomolecular networks. *Cell Systems*, 2(1):15–26, 2016.
- [31] Philipp Metzner, Christof Schütte, and Eric Vanden-Eijnden. Transition path theory for markov jump processes. *Multiscale Modeling & Simulation*, 7(3):1192–1219, 2009.
- [32] Hye-Won Kang and Thomas G. Kurtz. Separation of time-scales and model reduction for stochastic reaction networks. *Ann. Appl. Probab.*, 23(2):529–583, 2013.

- [33] Y. Cao, D.T. Gillespie, and L.R. Petzold. The slow-scale stochastic simulation algorithm. *Journal of Chemical Physics*, 122(1):1–18, 2005.
- [34] E. Weinan, D. Liu, and E. Vanden-Eijnden. Nested stochastic simulation algorithm for chemical kinetic systems with disparate rates. *Journal of Chemical Physics*, 123(19):1–8, 2005.
- [35] Weinan E, Di Liu, and Eric Vanden-Eijnden. Nested stochastic simulation algorithms for chemical kinetic systems with multiple time scales. *J. Comput. Phys.*, 221(1):158–180, January 2007.
- [36] James P. Keener and James Sneyd. *Mathematical physiology*, volume Interdisciplinary applied mathematics. Springer, 2009.
- [37] John G Kemeny, James Laurie Snell, and Gerald Luther Thompson. *Introduction to finite mathematics*. Prentice-Hall Englewood Cliffs, NJ, 1957.
- [38] Henri Cohen. *A course in computational algebraic number theory*, volume 138 of *Graduate Texts in Mathematics*. Springer-Verlag, Berlin, 1993.
- [39] Mukund Thattai and Alexander van Oudenaarden. Intrinsic noise in gene regulatory networks. *Proceedings of the National Academy of Sciences*, 98(15):8614–8619, 2001.
- [40] Thomas B. Kepler and Timothy C. Elston. Stochasticity in transcriptional regulation: Origins, consequences, and mathematical representations. *Biophysical Journal*, 81(6):3116 – 3136, 2001.
- [41] Steven M. Reppert and David R. Weaver. Coordination of circadian timing in mammals. *Nature*, 418(6901):935–941, 2002.
- [42] Neda Bagheri, Stephanie R Taylor, Kirsten Meeker, Linda R Petzold, and Francis J Doyle. Synchrony and entrainment properties of robust circadian oscillators. *Journal of The Royal Society Interface*, 5(Suppl 1):S17–S28, 2008.
- [43] Erik D. Herzog, Sara J. Aton, Rika Numano, Yoshiyuki Sakaki, and Hajime Tei. Temporal precision in the mammalian circadian system: A reliable clock from less reliable neurons. *Journal of Biological Rhythms*, 19(1):35–46, 2004.
- [44] David K. Welsh, Seung-Hee Yoo, Andrew C. Liu, Joseph S. Takahashi, and Steve A. Kay. Bioluminescence imaging of individual fibroblasts reveals persistent, independently phased circadian rhythms of clock gene expression. *Current Biology*, 14(24):2289 – 2295, 2004.
- [45] Ankit Gupta, Benjamin Hepp, and Mustafa Khammash. Noise induces the population-level entrainment of incoherent, uncoupled intracellular oscillators. *Cell Systems*, 3(6):521–531, 2016.
- [46] Jean-Christophe Leloup and Albert Goldbeter. Toward a detailed computational model for the mammalian circadian clock. *Proceedings of the National Academy of Sciences*, 100(12):7051–7056, 2003.
- [47] Hiroyuki Kurata, Hana El-Samad, T-M Yi, Mustafa Khammash, and J Doyle. Feedback regulation of the heat shock response in e. coli. In *Decision and Control, 2001. Proceedings of the 40th IEEE Conference on*, volume 1, pages 837–842. IEEE, 2001.
- [48] Lendert Gelens, Lydia Hill, Alexandra Vandervelde, Jan Danckaert, and Remy Loris. A general model for toxin-antitoxin module dynamics can explain persister cell formation in e. coli. *PLoS Comput Biol*, 9(8):e1003190, 2013.
- [49] Rick A Fasani and Michael A Savageau. Molecular mechanisms of multiple toxin–antitoxin systems are coordinated to govern the persister phenotype. *Proceedings of the National Academy of Sciences*, 110(27):E2528–E2537, 2013.
- [50] F. Horn and R. Jackson. General mass action kinetics. *Arch. Rational Mech. Anal.*, 47:81–116, 1972.
- [51] Martin Feinberg. Complex balancing in general kinetic systems. *Arch. Rational Mech. Anal.*, 49:187–194, 1972/73.

- [52] Martin Feinberg. The existence and uniqueness of steady states for a class of chemical reaction networks. *Arch. Rational Mech. Anal.*, 132(4):311–370, 1995.
- [53] Erel Levine and Terence Hwa. Stochastic fluctuations in metabolic pathways. *Proceedings of the National Academy of Sciences*, 104(22):9224–9229, 2007.
- [54] David F Anderson and Simon L Cotter. Product-form stationary distributions for deficiency zero networks with non-mass action kinetics. *Bulletin of Mathematical Biology*, 78(12):2390–2407, 2016.
- [55] Loc Pauleve, Gheorghe Craciun, and Heinz Koeppl. Dynamical properties of discrete reaction networks. *Journal of Mathematical Biology*, pages 1–18, 2013.