

A Scalable Computational Framework for Establishing Long-Term Behavior of Stochastic Reaction Networks

Ankit Gupta¹, Corentin Briat¹, Mustafa Khammash*

Department of Biosystems Science and Engineering (D-BSSE), Swiss Federal Institute of Technology–Zürich (ETH-Z), Basel, Switzerland



Abstract

Reaction networks are systems in which the populations of a finite number of species evolve through predefined interactions. Such networks are found as modeling tools in many biological disciplines such as biochemistry, ecology, epidemiology, immunology, systems biology and synthetic biology. It is now well-established that, for small population sizes, stochastic models for biochemical reaction networks are necessary to capture randomness in the interactions. The tools for analyzing such models, however, still lag far behind their deterministic counterparts. In this paper, we bridge this gap by developing a constructive framework for examining the long-term behavior and stability properties of the reaction dynamics in a stochastic setting. In particular, we address the problems of determining ergodicity of the reaction dynamics, which is analogous to having a globally attracting fixed point for deterministic dynamics. We also examine when the statistical moments of the underlying process remain bounded with time and when they converge to their steady state values. The framework we develop relies on a blend of ideas from probability theory, linear algebra and optimization theory. We demonstrate that the stability properties of a wide class of biological networks can be assessed from our sufficient theoretical conditions that can be recast as efficient and scalable linear programs, well-known for their tractability. It is notably shown that the computational complexity is often linear in the number of species. We illustrate the validity, the efficiency and the wide applicability of our results on several reaction networks arising in biochemistry, systems biology, epidemiology and ecology. The biological implications of the results as well as an example of a non-ergodic biological network are also discussed.

Citation: Gupta A, Briat C, Khammash M (2014) A Scalable Computational Framework for Establishing Long-Term Behavior of Stochastic Reaction Networks. *PLoS Comput Biol* 10(6): e1003669. doi:10.1371/journal.pcbi.1003669

Editor: Ruth Nussinov, National Cancer Institute, United States of America and Tel Aviv University, Israel

Received: October 1, 2013; **Accepted:** April 30, 2014; **Published:** June 26, 2014

Copyright: © 2014 Gupta et al. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Funding: This work was supported by the Swiss Federal Institute of Technology ETH. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The authors have declared that no competing interests exist.

* Email: mustafa.khammash@bsse.ethz.ch

† These authors contributed equally to this work.

This is a *PLOS Computational Biology Methods* article.

Introduction

Reaction networks are used as modeling tools in many areas of science. Examples include chemical reaction networks [1], cell signalling networks [2], gene expression networks [3], metabolic networks [4], pharmacological networks [5], epidemiological networks [6] and ecological networks [7]. Traditionally, reaction networks are mathematically analyzed by representing the dynamics as a set of ordinary differential equations. Such a deterministic model is reasonably accurate when the number of network participants is *large*. However, when this is not the case, the discreteness in the interactions becomes important and the dynamics inherently *noisy*. This *random* component of the dynamics cannot be ignored as it can strongly influence the system's behavior [8–10]. To understand the effects of this randomness, stochastic models are needed, and the most common approach is to model the reaction dynamics as a continuous-time Markov process, whose states denote the current population size. Many recent works have employed such stochastic models to study the impact of noise [11–14].

In stochastic models, the underlying Markov process $(X(t))_{t \geq 0}$ is a *pure-jump* process whose state space \mathcal{S} contains all the population size vectors that are *reachable* by the random dynamics. The probability distribution of $(X(t))_{t \geq 0}$ evolves according to a system of linear ordinary differential equations (ODEs), known as the *Chemical Master Equation* (CME) or Forward Kolmogorov Equation [15]. The dimension of the system of ODEs is equal to the number of elements in the state space \mathcal{S} , with each element representing a possible combination of reacting species abundances. When \mathcal{S} is finite and small in size, the CME can be solved analytically since it is simply a small and finite system of linear differential equations. However, for infinite state-spaces an exact solution to the CME is difficult to obtain except in some special cases [16,17]. Beyond these special cases, current methods often rely on truncating the infinite state-space to obtain finite approximations of the CME [18], and then resorting to efficient numerical methods for their solutions. Such methods include Expokit [19], which is based on Krylov Subspace Identification, or the backward Euler method proposed in [20], among others. Such an approach works well only for relatively small systems, as the curse-of-dimensionality renders the numerical solution of the truncated master equation of larger systems prohibitive. Nevertheless, recent methods based on Tensor Train (TT) and

Author Summary

In many biological disciplines, computational modeling of interaction networks is the key for understanding biological phenomena. Such networks are traditionally studied using deterministic models. However, it has been recently recognized that when the populations are small in size, the inherent random effects become significant and to incorporate them, a stochastic modeling paradigm is necessary. Hence, stochastic models of reaction networks have been broadly adopted and extensively used. Such models, for instance, form a cornerstone for studying heterogeneity in clonal cell populations. In biological applications, one is often interested in knowing the long-term behavior and stability properties of reaction networks even with incomplete knowledge of the model parameters. However for stochastic models, no analytical tools are known for this purpose, forcing many researchers to use a simulation-based approach, which is highly unsatisfactory. To address this issue, we develop a theoretical and computational framework for determining the long-term behavior and stability properties for stochastic reaction networks. Our approach is based on a mixture of ideas from probability theory, linear algebra and optimization theory. We illustrate the broad applicability of our results by considering examples from various biological areas. The biological implications of our results are discussed as well.

Quantized Tensor Train (QTT) representations [21,22] show that for CME problems that admit bounded TT ranks, storage costs and computational complexity that grow linearly in the number of species may be achieved. These and other methods for the numerical solutions of the CME remain active topics of research.

When \mathcal{S} is infinite or very large in size, the most common approach for approximating the solutions of a CME is by simulating a large number of trajectories of the underlying Markov process $(X(t))_{t \geq 0}$, and using the sample values of $X(t)$ to estimate the distribution at time t . Such simulations are performed using Monte Carlo procedures such as Gillespie's *stochastic simulation algorithm* (SSA) or its variants [23–25]. Since the simulation time of SSA depends linearly on the number of reactions that occur during the simulation time period, these procedures can be cumbersome for large networks. It is well-known that the stochastic effects caused by the random timing of reactions become less important when the population size is large. The dynamical law of large numbers proved by Kurtz [26] shows that under an appropriate scaling relationship between the population size, reaction rates and the system size, the stochastic model of a reaction network converges to the deterministic model, as the system size goes to infinity. Under this scaling relationship, one can also approximate the stochastic dynamics with certain stochastic differential equations (SDEs) that are easier to simulate and analyze [27,28]. However, these SDE approximations can only work when the population sizes of all the species in the reaction network are large, which is often not the case. For a detailed survey on the topic of estimating the solution of a CME, we refer the readers to the paper [29] which contains an exhaustive list of methods for this purpose.

In many biological applications, one is interested in analyzing the long-term behavior or stability properties of a reaction network. This is fairly straightforward for deterministic models because many tools from the theory of ordinary differential equations can be used for this analysis [30]. However, the stability properties of stochastic models for reaction networks are difficult to verify for the following reasons. Let us consider a stochastic

reaction network whose dynamics is represented by the Markov process $(X(t))_{t \geq 0}$ with state space \mathcal{S} . The evolution of the distributions of this Markov process is given by $(p(t))_{t \geq 0}$ which is the solution of the CME corresponding to the reaction network. Heuristically, we regard the stochastic dynamics to be stable when the family of distributions $(p(t))_{t \geq 0}$ is “well-behaved” with time. In this paper, we consider several notions of “well-behaved” dynamics. The strongest of these notions is the concept of *ergodicity* [31] which means that there exists a unique stationary distribution π for the Markovian dynamics, such that $p(t) \rightarrow \pi$ as $t \rightarrow \infty$, irrespective of the initial distribution $p(0)$. This is analogous to having a globally attracting fixed point in the deterministic setting. If \mathcal{S} is finite, the process is ergodic if and only if it is *irreducible*, in the sense that all the states in \mathcal{S} are reachable from each other. It is hence enough to check irreducibility of the process using, for example, matrix methods [32,33]. Contrary to this situation, our main interest in this paper is in analyzing the stability properties of stochastic reaction networks with an infinite state space \mathcal{S} . Note that in such cases, irreducibility no longer implies ergodicity, since the trajectories of the Markov process may blow up with time (see the carcinogenesis example in the discussion section). In this regard, ergodicity cannot be considered as a *generic property* of reaction networks with infinite state-spaces since both ergodic and non-ergodic processes can be found in nature. Assuming ergodicity without verifying it beforehand seems to be therefore unreasonable from both theoretical and practical perspectives. The direct verification of stability properties like ergodicity is generally not possible as the CME cannot be explicitly solved, except in some restrictive cases [16,17]. The common approach of using Monte Carlo simulations for estimating the solutions of a CME is inadequate for assessing the long-term behavior and stability properties of a stochastic reaction network, because one can only simulate finitely many trajectories and those too for a finite amount of time. Some methods for analyzing stability properties without the need for simulations exist, but they either work for specific networks [17,34], very special classes of networks such as zero-deficiency networks [35], or assume system size approximations where the stochastic dynamics is represented by an SDE [36,37]. Such system size approximations do not hold when some species are present in low copy numbers, and even if they hold, the approximation error generally blows up with time [28]. Hence the stochastic dynamics and the corresponding SDE may have completely different long-term behaviors. Our aim, in this paper, is to develop a theoretical and computational framework for analyzing the long-term behavior and stability properties of stochastic models for reaction networks that do not rely on computationally expensive Monte Carlo simulations or on system size approximations of the stochastic dynamics. A similar goal is also achieved in the works [38,39] where results on stability and moments bounds are also obtained. The approach proposed in [40] is built upon a Foster-Lyapunov criterion [31] and a quadratic Foster-Lyapunov function in order to estimate the location of the stationary distribution. In the same, yet different, spirit, the proposed approach also relies on a Foster-Lyapunov condition but using a linear Foster-Lyapunov function that allows us to establish ergodicity, moment bounds, moment convergence and the existence of attractive sets for moments. While the approach in [40] is fully computational, the one we propose is also theoretical and allows us to conclude on structural properties of classes of networks such as *structural ergodicity*, *structural boundedness of moments* and *structural convergence of moments*. Our approach relies on a mixture of simple ideas from stochastic analysis, linear algebra, polynomial analysis and optimization. Even though our conditions are only sufficient, we demonstrate their broad applicability by

successfully establishing stability properties of several reaction networks taken from the literature.

We mentioned before that the stochastic and the deterministic models of a reaction network are connected through the dynamical law of large numbers [26]. It might be tempting to think that the stability properties of a stochastic model can be assessed by studying the stability properties of the corresponding deterministic model. However in general, the stochastic and deterministic models can have very different stability properties. This is because a deterministic model cannot capture noise induced effects which may have a significant impact on the long-term behavior of a system. For example, in the synthetic Toggle Switch by Gardner [41], the deterministic model exhibits bistability and hence starting from different initial values, the system can converge to two different steady states. On the other hand, the corresponding stochastic model is ergodic (see network (35)) and hence the solution of the CME converges to the same stationary distribution irrespective of the initial distribution. A similar phenomenon occurs with the repressilator (see [42] and network (36)), where the stochastic model is ergodic while the deterministic model exhibits oscillations. On the other hand, it is also possible to find networks for which the deterministic model has a locally asymptotically stable equilibrium point, implying that whenever the initial condition is contained within its region of attraction, the trajectories converge to it. If the initial condition lies outside this region of attraction, then the trajectories of such a network become unbounded with time. In the stochastic setting, the randomness causes each trajectory to leave the region of attraction in finite time, and then become unbounded suggesting that there is no stationary distribution for the dynamics (see network (22) and Figure 1). This lack of stationary distribution is because the stochastic dynamics can jump potential wells from one macroscopic fixed point which is stable to another fixed point which is unstable [43]. A more striking example of divergent deterministic and stochastic behaviors is given by network (26) (see also Figure 2). While the deterministic model has a unique globally stable fixed point, the stochastic model is non-ergodic and all the moments grow unboundedly with time. In this example it is impossible to predict the stochastic behavior from

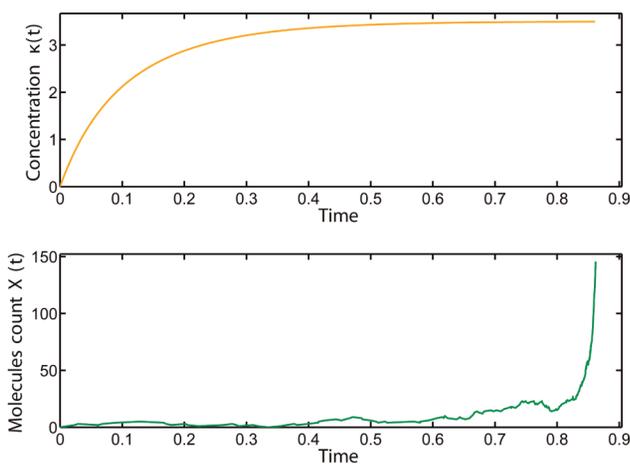


Figure 1. Trajectory of the state of the deterministic system (23) with initial condition $\kappa_0=0$ (top); Sample path of the Markov process describing the network (22) with initial condition $x_0=0$ (bottom). Whereas the trajectory of the state of the deterministic model converges to a stationary value, the trajectory of the state of the stochastic model goes unbounded. doi:10.1371/journal.pcbi.1003669.g001

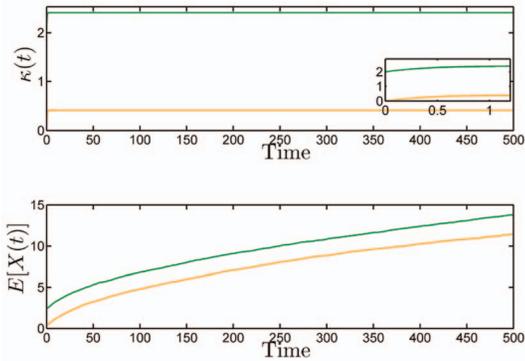


Figure 2. Comparison of the trajectories of the deterministic and stochastic (first-order moments) models of the reaction network (22) with initial condition $\kappa_1(0)=0$, $\kappa_2(0)=2$, $X_1(0)=0$ and $X_2(0)=2$ for the deterministic (top) and stochastic dynamics (bottom), respectively. We can see that while the deterministic trajectories converge to their equilibrium point, the first-order moments grow without bound. doi:10.1371/journal.pcbi.1003669.g002

the deterministic model. The above examples illustrate that the stability properties of the stochastic dynamics can, in general, not be assessed from the stability properties of the deterministic dynamics.

Our results can help in understanding the stability properties of the moments of a Markov process $(X(t))_{t \geq 0}$ representing a reaction network. In particular, we present a method to check if these moments remain bounded with time and if they converge to their steady state values as time goes to infinity. Such results can help in verifying the suitability of a model for a given system and in designing biological controllers that drive the moments to specific steady state values. We provide easily computable bounds for the moments that hold uniformly in time. We also determine bounds for the steady state moment values, which can help in understanding the properties of the steady state distribution, even if this distribution is not explicitly computable. In many biological applications, it is of great interest to explicitly compute the first few moments of the process $(X(t))_{t \geq 0}$ without solving the corresponding CME. One can easily express the dynamics of these moments as a system of ordinary differential equations, but generally this system is not closed when the network has nonlinear interactions. Many moment closure methods that suggest schemes to close these equations to obtain approximations for the moments have been proposed (see e.g. [44,45] and references therein). The results obtained in this paper can be used to ascertain the correctness of a given moment closure method for a specific network (see the example based on the network (29)). Furthermore, several moment closure methods are developed under an implicit assumption that the *moment-generating* function corresponding to the solution of the CME exists for all times. One of our results provides a way to easily check that this assumption is indeed valid.

Reaction networks

Let us now formally describe reaction networks. Motivated by the literature on chemical kinetics, we refer to the network participants as *molecules* which may belong to one of d species S_1, \dots, S_d . There are K reactions in the network and for any $k=1, \dots, K$, the *stoichiometric* vector $\zeta_k = (\zeta_{k,1}, \dots, \zeta_{k,d})$ denotes the change in the number of molecules in each of the species due to the k -th reaction.

Deterministic models

Consider the deterministic model for the reaction network described above. In this setting, the state of the system is described by a vector of *concentrations* of the d species which we denote by $\kappa \in \mathbb{R}_{\geq 0}^d$. The concentration of a species is simply its molecular count divided by the system volume. Let $\tilde{\lambda}_k(\kappa)$ be the *flux* associated with the k -th reaction (see [8]). To ensure positivity of the system, we require that $\tilde{\lambda}_k(\kappa) = 0$ whenever $\kappa_i = 0$ and $\zeta_{k,i} < 0$. If the initial state is κ_0 , then the evolution of concentrations is given by $(\phi_{\kappa_0}(t))_{t \geq 0}$ which satisfies the Reaction Rate Equations (RRE) of the form

$$\frac{d\phi_{\kappa_0}(t)}{dt} = \sum_{k=1}^K \tilde{\lambda}_k(\phi_{\kappa_0}(t)) \zeta_k \text{ with } \phi_{\kappa_0}(0) = \kappa_0. \quad (1)$$

We are interested in the long-term behavior and stability of our reaction dynamics. More precisely, we would like to check if the following conditions are satisfied.

DC1 For any κ_0 , there is a compact set $\mathcal{K}(\kappa_0)$ such that $\phi_{\kappa_0}(t) \in \mathcal{K}(\kappa_0)$ for all $t \geq 0$.

DC2 There exists a compact set \mathcal{K}_0 such that for any κ_0 , we have $\phi_{\kappa_0}(t) \in \mathcal{K}_0$ for large values of t .

DC3 There is a κ_{eq} such that for any κ_0 we have $\phi_{\kappa_0}(t) \rightarrow \kappa_{eq}$ as $t \rightarrow \infty$.

The first condition, **DC1**, says that for any κ_0 , the entire trajectory $(\phi_{\kappa_0}(t))_{t \geq 0}$ stays within some compact set. We would expect this to be true for most *realistic* systems. Hence a violation of this property may suggest a *flaw* in the deterministic model. The second condition, **DC2**, says that there is an *attractor* set for the dynamics, where all the trajectories eventually lie, irrespective of their starting point. The last condition, **DC3**, says that there is a globally attracting fixed point for the deterministic model. Using techniques from the theory of dynamical systems [30,46], one can verify these conditions, without the need of simulating the deterministic model. There is also a general theory to check condition **DC3** for reaction networks satisfying mass-action kinetics (see [47–50]). Broadly speaking, these three conditions present different ways of saying that the reaction dynamics is “well-behaved”. Our goal in this paper is to develop a theoretical and computational framework for verifying conditions similar to **DC1**, **DC2** and **DC3** for stochastic models of reaction networks.

Stochastic models

Consider the stochastic model corresponding to the reaction network described above. In this setting, the firing of reactions are discrete events and the state of the system refers to the vector of molecular counts of the d species. When the state is x , the k -th reaction fires after a random time which is exponentially distributed with rate $\lambda_k(x)$. The functions $\lambda_1, \dots, \lambda_K$ are known as the *propensity* functions in the literature. To ensure positivity of the system, we require that if $x + \zeta_k \notin \mathbb{N}_0^d$, then $\lambda_k(x) = 0$, where \mathbb{N}_0 is the set of non-negative integers. The dynamics can be represented by the Markov process $(X_{x_0}(t))_{t \geq 0}$ where x_0 is the initial state. Note that if $X_{x_0}(t) = (X_1(t), \dots, X_d(t))$, then $X_i(t)$ is the *number* of molecules of \mathbf{S}_i at time t .

It is important to select a suitable state space \mathcal{S} for the Markov process representing the reaction dynamics. We choose \mathcal{S} to be a non-empty subset of \mathbb{N}_0^d satisfying the following properties:

- (A) If $x \in \mathcal{S}$ and $\lambda_k(x) > 0$ for some $k = 1, \dots, K$, then $x + \zeta_k \in \mathcal{S}$.
- (B) There is no proper subset $\mathcal{S}_1 \subset \mathcal{S}$ satisfying part (A).

Observe that part (A) ensures that if $x_0 \in \mathcal{S}$ then $X_{x_0}(t) \in \mathcal{S}$ for all $t \geq 0$ and hence \mathcal{S} can be taken to be the state space of all the

Markov processes describing the stochastic reaction network with an initial state x_0 in \mathcal{S} . Part (B) implies that the reaction dynamics cannot be contained in a proper subset of \mathcal{S} . The role of this assumption will become clear in the next section, when we discuss the issue of state space *irreducibility*. Note that in certain cases, such as the *pure-birth* network $\emptyset \rightarrow \mathbf{S}_1$, a suitable state space satisfying the above criteria cannot be found. There also exist cases where the above criteria restricts the choice of state space. For example, for the *pure-death* network $\mathbf{S}_1 \rightarrow \emptyset$, the only possible choice for state space is $\mathcal{S} = \{0\}$. Finally we remark that if the reactions in a network satisfy a conservation relation then the state space must be chosen with an initial condition in mind. For example, for the network $\mathbf{S}_1 \rightleftharpoons \mathbf{S}_2$, the sum of molecular counts of \mathbf{S}_1 and \mathbf{S}_2 is preserved by the reactions. Hence if we wish to study the stochastic dynamics with the initial sum as n , then the correct choice for state space is $\mathcal{S} = \{(x_1, x_2) \in \mathbb{N}_0^2 : x_1 + x_2 = n\}$.

Let $\mathcal{P}(\mathcal{S})$ denote the space of probability distributions over \mathcal{S} , endowed with the weak topology which is metrized by the Prohorov metric (see [51]). For any $x, y \in \mathcal{S}$ let $p_x(t, y)$ denote the following probability

$$p_x(t, y) = \mathbb{P}(X_x(t) = y). \quad (2)$$

Defining $p_x(t)(A) = \sum_{y \in A} p_x(t, y)$, for any $A \subset \mathcal{S}$, we can view $p_x(t)$ as an element in $\mathcal{P}(\mathcal{S})$. In fact, $p_x(t)$ is the distribution at time t of the Markov process $(X_x(t))_{t \geq 0}$. The dynamics of $p_x(t)$ is given by the Chemical Master Equation (CME) which has the following form:

$$\frac{dp_x(t, y)}{dt} = \sum_{k=1}^K (p_x(t, y - \zeta_k) \lambda_k(y - \zeta_k) - p_x(t, y) \lambda_k(y)), \quad (3)$$

where $p_x(0, y) = 1$ if $x = y$ and $p_x(0, y) = 0$ for all $y \neq x$. Theoretically, one can find $p_x(t, y)$ for any $t \geq 0$ and $y \in \mathcal{S}$, by solving this system. However this system consists of as many equations as the number of elements in \mathcal{S} . Hence an explicit solution is only possible when \mathcal{S} is finite, which only happens in very restrictive cases where all the reactions preserve some conservation relation. Typically, \mathcal{S} is infinite and solving this system analytically or even numerically is nearly impossible, except in some restrictive cases (see [16]). From now on, we assume that \mathcal{S} is infinite.

The above discussion shows that at the level of distributions, we can view the stochastic dynamics $(X_{x_0}(t))_{t \geq 0}$ as the deterministic dynamics $(p_{x_0}(t))_{t \geq 0}$, which satisfies the CME. However, the major difficulty in analyzing this deterministic dynamics is that it occurs over an infinite dimensional space $\mathcal{P}(\mathcal{S})$. Nevertheless we can recast the conditions **DC1**, **DC2** and **DC3** in the stochastic setting as below.

SC1 For any x_0 , there is a compact set $\mathcal{K}(x_0) \subset \mathcal{P}(\mathcal{S})$ such that $p_{x_0}(t) \in \mathcal{K}(x_0)$ for all $t \geq 0$.

SC2 There exists a compact set $\mathcal{K}_0 \subset \mathcal{P}(\mathcal{S})$ such that for any $x_0 \in \mathcal{S}$ we have $p_{x_0}(t) \in \mathcal{K}_0$ for large values of t .

SC3 There is a $\pi \in \mathcal{P}(\mathcal{S})$ such that for any x_0 we have $p_{x_0}(t) \rightarrow \pi$ as $t \rightarrow \infty$.

Each of the above conditions give an important insight about the long-term behavior and stability of the stochastic dynamics. The first condition, **SC1**, says that for every $\epsilon \in (0, 1)$ we can find a finite set $A_\epsilon \subset \mathcal{S}$ such that each $p_{x_0}(t)$ puts at least $(1 - \epsilon)$ of its mass in A_ϵ . In other words, the probability that the state of the underlying Markov process at any time t is inside A_ϵ is greater than $(1 - \epsilon)$. We would expect this to be true for most realistic models. If condition **SC2** holds then the evolution of distributions have a compact attractor set in $\mathcal{P}(\mathcal{S})$, where all the trajectories eventually lie irrespective of their starting point. This suggests that

in the long run, the family of processes $\{(X_{x_0}(t))_{t \geq 0} : x_0 \in \mathcal{S}\}$, spend most of their time on the same set of states. The last condition **SC3** says that the evolution of distributions have a globally attracting fixed point π . If this holds, then the Markov process representing the reaction dynamics is *ergodic* with π as the unique stationary distribution. For understanding the long-term behavior of a stochastic process, ergodicity is a desirable property to have. In the long-run, the proportion of time spent by any trajectory of an ergodic process, in any subset of the state space is equal to the stationary probability of that subset (see (12)). In other words, information about the stationary distribution can be obtained by observing just one trajectory for a sufficiently long time. Such a result can have important applications. For example, consider a culture with a large number of identical cells with each cell having the same reaction network. If we can show that this intracellular network is ergodic, then by observing the long-term reaction dynamics in a single cell, using for example, time-lapse microscopy, we can obtain statistical information about all the cells at stationarity. Conversely, ergodicity allows us to obtain the stationary distribution of a single-cell by observing the distribution over the population, using for example flow cytometry.

In this paper we develop a general framework for checking conditions **SC1**, **SC2** and **SC3**. However, the scope of our paper is broader than that. As mentioned in the introduction, we obtain easily computable bounds for the statistical moments of the underlying Markov process and investigate when these moments converge with time. We also present conditions for the distribution of the process to be *light-tailed*.

Results

Preliminaries

In this section we discuss the main results of our paper. In particular, we explain how conditions **SC1**, **SC2** and **SC3** can be verified without having to simulate the trajectories of the Markov process representing the reaction dynamics. Intuitively, these conditions can only hold if the Markov process has a low probability of hitting states that have a very large *size*. In our case, the states are vectors in \mathbb{R}^d and so we can measure their *size* by using any norm on \mathbb{R}^d . The central theme of this paper is to demonstrate that for many networks, long-term behavior can be easily analyzed by choosing the *right* norm for measuring the state sizes. This *right* norm has the form

$$\|x\|_v = \sum_{i=1}^d v_i |x_i|, \tag{4}$$

where v is a positive vector in \mathbb{R}^d satisfying the following condition.

Condition 1 (Drift-Diffusivity Condition). For a positive vector $v \in \mathbb{R}^d$, there exist positive constants c_1, c_2, c_3, c_4 and a nonnegative constant c_5 such that for all $x \in \mathcal{S}$

$$\sum_{k=1}^K \lambda_k(x) \langle v, \zeta_k \rangle \leq c_1 - c_2 \langle v, x \rangle \text{ and} \tag{5a}$$

$$\sum_{k=1}^K \lambda_k(x) \langle v, \zeta_k \rangle^2 \leq c_3 + c_4 \langle v, x \rangle + c_5 \langle v, x \rangle^2. \tag{5b}$$

Here $\langle \cdot, \cdot \rangle$ denotes the standard inner product on \mathbb{R}^d . If we consider the process $(\|X_{x_0}(t)\|_v)_{t \geq 0}$, then its dynamics can be seen

to have two components *drift* and *diffusion* which have the form $\sum_{k=1}^K \lambda_k(x) \langle v, \zeta_k \rangle$ and $\sum_{k=1}^K \lambda_k(x) \langle v, \zeta_k \rangle^2$ respectively when $X_{x_0}(t) = x$ (see page 2 in the Supplementary Material **S1**). Condition 1 gives upper-bounds for the magnitude of these two components and hence we call it the *drift-diffusivity condition* (abbreviated to Condition **DD** from now on; the abbreviations **DD1** and **DD2** stand for the first and second inequality, respectively). Observe that when the process $(\|X_{x_0}(t)\|_v)_{t \geq 0}$ goes above c_1/c_2 then it experiences a negative drift, suggesting that it will move downwards. This fact will be crucial for our analysis.

For now, we assume that a vector v satisfying Condition **DD** has been found. In later sections we demonstrate how v can be determined for a large class of networks by solving suitably constructed optimization problems.

For any positive integer r , let $m_{x_0}^r(t)$ denote the r -th moment of $\|X_{x_0}(t)\|_v$ defined by

$$m_{x_0}^r(t) = \mathbb{E} \left(\|X_{x_0}(t)\|_v^r \right) = \sum_{y \in \mathcal{S}} \|y\|_v^r p_{x_0}(t, y). \tag{6}$$

Similarly let $\Psi^r(x_0, t)$ denote the r -th moment of $X_{x_0}(t)$ at time t . Then $\Psi^r(x_0, t)$ is a tensor of rank r whose entry at index $(i_1, \dots, i_r) \in \{1, 2, \dots, d\}^r$ is given by

$$\Psi_{i_1 \dots i_r}^r(x_0, t) = \sum_{y \in \mathcal{S}} y_{i_1} \dots y_{i_r} p_{x_0}(t, y), \tag{7}$$

where $y = (y_1, \dots, y_d)$ and $p_{x_0}(t)$ is the distribution of $X_{x_0}(t)$.

Suppose that for some positive constants r and $C_r(x_0)$ we have

$$\sup_{t \geq 0} m_{x_0}^r(t) \leq C_r(x_0). \tag{8}$$

For any $M > 0$, let K_M be the compact (finite) set defined by $K_M = \{x \in \mathcal{S} : \|x\|_v \leq M\}$ and let K_M^c denote its complement. Markov's inequality (see [52]) implies that for any $\epsilon > 0$ we can choose M large enough to satisfy

$$\begin{aligned} \sup_{t \geq 0} p_{x_0}(t, K_M^c) &= \sup_{t \geq 0} \mathbb{P} \left(\|X_{x_0}(t)\|_v^r > M^r \right) \leq \\ \sup_{t \geq 0} \mathbb{E} \left(\|X_{x_0}(t)\|_v^r \right) &\leq \frac{C_r(x_0)}{M^r} < \epsilon. \end{aligned}$$

Hence Prohorov's theorem (see Chapter 3 in [51]) ensures that condition **SC1** holds. Similarly we can prove that condition **SC2** will hold if for some $r > 0$ there exists a constant \hat{C}_r such that

$$\limsup_{t \rightarrow \infty} m_{x_0}^r(t) \leq \hat{C}_r \text{ for all } x_0 \in \mathcal{S}. \tag{9}$$

Relations (8) and (9) give uniform and asymptotic upper-bounds for $m_{x_0}^r(t)$. Using these relations we can also obtain uniform and asymptotic upper-bounds for the entries of $\Psi^r(x_0, t)$. Such moment bound results have applications in queuing theory and control theory (see [53]). In Theorem 2 we show that under certain conditions, (8) and (9) hold and the upper-bounds can be easily computed.

Instead of the r -th moment of the process $(\|X_{x_0}(t)\|_v)_{t \geq 0}$, one can ask if the exponential moment of this process is uniformly bounded from above. This will happen if for some $\gamma > 0$ we have

$$\sup_{t \geq 0} \mathbb{E} \left(e^{\gamma \|X_{x_0}(t)\|_v} \right) = \sup_{t \geq 0} \sum_{y \in \mathcal{S}} e^{\gamma \|y\|_v} p_{x_0}(t, y) < \infty. \quad (10)$$

If (10) holds, then the distribution $p_{x_0}(t)$ is *light-tailed* (a distribution is called light-tailed if its tails are majorized by an exponential decay) uniformly in t . This shows that all the cumulants of the distribution $p_{x_0}(t)$ exist, which is an important result for the following reason. There is a considerable body of research dedicated to estimating the moments of the process $(X_{x_0}(t))_{t \geq 0}$ directly without computing the distribution functions $p_{x_0}(t)$. For any integer $r > 0$, one can easily write the differential equations for the dynamics of the first r moments. However when the reaction network has nonlinear interactions, this system of equations is not closed for any r . Various moment closure methods (see [54,55]) exist that specify ways to close these equations artificially and estimate the moments approximately. A popular moment closure method is the cumulant-neglect method which ignores the higher order cumulants of the distribution $p_{x_0}(t)$ for all $t \geq 0$. Of course this method is only valid when the higher order cumulants exist. This is guaranteed if (10) holds. In Theorem 3 we give conditions for verifying (10).

We now come to the question of checking condition **SC3** which says that the process $(X_{x_0}(t))_{t \geq 0}$ is ergodic. This can only happen if the state space \mathcal{S} is *irreducible*, which means that all the states are accessible from each other. Recall the definition of $p_x(t, y)$ from (2). Mathematically, we say that \mathcal{S} is irreducible if for all $x, y \in \mathcal{S}$, we have $p_x(t_1, y) > 0$ and $p_y(t_2, x) > 0$ for some $t_1, t_2 > 0$. In order to check the irreducibility of \mathcal{S} , one has to verify that there is no proper subset $\mathcal{S}_1 \subset \mathcal{S}$, such that once the process reaches a state in \mathcal{S}_1 , it stays in \mathcal{S}_1 forever. For reaction networks with mass-action kinetics, methods for checking irreducibility have recently been reported in [56] and [57]. These methods can be extended to situations where the propensity functions are positive in the positive orthant. When the propensity functions vanish inside the positive orthant, the problem of checking irreducibility can become much more complicated, and to the best of our knowledge no methods exist in the literature for this purpose.

We mentioned before that the vector v is chosen so that the process $(\|X_{x_0}(t)\|_v)_{t \geq 0}$ has a negative drift at large values. Assuming irreducibility, this is sufficient to verify ergodicity of $(X_{x_0}(t))_{t \geq 0}$ (see Proposition 4).

Suppose that condition **SC3** is satisfied and the process $(X_{x_0}(t))_{t \geq 0}$ is ergodic with stationary distribution π . For any positive integer r , let Π^r denote the r -th moment of the stationary distribution π . Then Π^r is a tensor of rank r defined in the same way as $\Psi^r(x_0, t)$ (see (7)), with $p_{x_0}(t, y)$ replaced by $\pi(y)$. Using Theorem 2 we can determine the values of r for which Π^r is finite (componentwise) and $\Psi^r(x_0, t) \rightarrow \Pi^r$ as $t \rightarrow \infty$ (see Theorem 5). We can also identify functions $f : \mathcal{S} \rightarrow \mathbb{R}$ for which

$$\lim_{t \rightarrow \infty} \mathbb{E}(f(X_{x_0}(t))) = \sum_{y \in \mathcal{S}} f(y)\pi(y) < \infty \quad (11)$$

holds for any $x_0 \in \mathcal{S}$. If f is such a function, then the ergodic theorem for Markov processes (see [58]) says that

$$\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t f(X_{x_0}(s)) ds = \sum_{y \in \mathcal{S}} f(y)\pi(y) \text{ almost surely,} \quad (12)$$

for any $x_0 \in \mathcal{S}$. Lastly, we also obtain conditions to check if the stationary distribution π is *light-tailed* (see Theorem 6).

General results

In this section, we formally present the main results of our paper. Their proofs are given in the Supplementary Material **S1**.

Moment bounds. Our first result establishes that for certain values of r , we can obtain uniform and asymptotic moment bounds for the r -th moment of the process $(\|X_{x_0}(t)\|_v)_{t \geq 0}$.

Theorem 2 Assume that Condition **DD** holds. Let r_{\max} be given by

$$r_{\max} = \begin{cases} 1 + \frac{2c_2}{c_5} & \text{if } c_5 > 0 \\ \infty & \text{if } c_5 = 0. \end{cases} \quad (13)$$

For any positive integer r , if $r < r_{\max}$ then there exist positive constants $C_r(x_0)$ and \tilde{C}_r such that (8) and (9) hold.

The values of the constants $C_r(x_0)$ and \tilde{C}_r can be explicitly computed using a recursive relationship (see the Supplementary Material **S1**). Note that if $v = (v_1, \dots, v_d)$, then for any $y = (y_1, \dots, y_d) \in \mathcal{S}$ we have $y_i \leq \|y\|_v / v_i$ for any i . Hence for any $i_1, \dots, i_r \in \{1, 2, \dots, d\}$ we have $\Psi_{i_1 \dots i_r}^r(x_0, t) \leq m_{x_0}^r(t) / \prod_{j=1}^r v_{i_j}$. Therefore using Theorem 2, we can obtain uniform and asymptotic moment bounds for the reaction dynamics $(X_{x_0}(t))_{t \geq 0}$ (see the Supplementary Material **S1**).

Observe that if $c_5 = 0$ then $r_{\max} = \infty$. In this case, Theorem 2 says that for each positive integer r and $x_0 \in \mathcal{S}$ there exists a constant $C_r(x_0)$ such that (8) holds. By showing that we have a $C > 0$ such that $C_r(x_0) \leq r! C^r$ for all positive integers r , we obtain our next result, which gives sufficient conditions to check (10).

Theorem 3 (Uniform Light-Tailedness) Suppose that Condition **DD** holds with $c_5 = 0$. Given an initial state $x_0 \in \mathcal{S}$ there exists a $\gamma > 0$ such that

$$\sup_{t \geq 0} \mathbb{E} \left(e^{\gamma \|X_{x_0}(t)\|_v} \right) = \sup_{t \geq 0} \sum_{y \in \mathcal{S}} e^{\gamma \|y\|_v} p_{x_0}(t, y) < \infty.$$

Ergodicity and Moment Convergence. The next result verifies the ergodicity of a reaction network satisfying Condition **DD**. It follows from Theorem 7.1 in Meyn and Tweedie [59].

Proposition 4 (Ergodicity) Assume that the state space \mathcal{S} of the Markov process $(X_{x_0}(t))_{t \geq 0}$ is irreducible and Condition **DD1** holds. Then this process is exponentially ergodic in the sense that there exists a unique distribution $\pi \in \mathcal{P}(\mathcal{S})$ along with constants $B, c > 0$ such that for any $x_0 \in \mathcal{S}$

$$\sup_{A \subset \mathcal{S}} |p_{x_0}(t, A) - \pi(A)| \leq B e^{-ct} \text{ for all } t \geq 0.$$

This result says that as $t \rightarrow \infty$, the distribution $p_{x_0}(t)$ converges to π exponentially fast. Henceforth we assume that the process $(X_{x_0}(t))_{t \geq 0}$ is ergodic with stationary distribution π .

Let $f : \mathcal{S} \rightarrow \mathbb{R}$ be a function such that for some positive integer $r < (r_{\max} - 1)$, there exists a $C > 0$ satisfying $|f(x)| \leq C(1 + \|x\|_v^r)$ for all $x \in \mathcal{S}$. Using Theorem 2 we can prove that for such a f , the relations (11) and (12) hold. As a consequence we obtain the following result about the convergence of moments with time.

Theorem 5 (Moment Convergence) Assume that Condition **DD** holds. Let r be any positive integer satisfying $r < (r_{\max} - 1)$. Then Π^r is finite (componentwise) and $\Psi^r(x_0, t) \rightarrow \Pi^r$ as $t \rightarrow \infty$.

If $f(x) = \|x\|_v^r$ then Theorem 2 and (11) imply that for any positive integer $r < (r_{\max} - 1)$ there exists a positive constant \tilde{C}_r such that

$$\sum_{y \in \mathcal{S}} \|y\|_v^r \pi(y) \leq \hat{C}_r. \tag{14}$$

In particular, if $c_5 = 0$ then $r_{\max} = \infty$ and (14) holds for each r . By proving the existence of a constant $C > 0$ such that $\hat{C}_r \leq r! C^r$ for all positive integers r we get our last result which shows that the stationary distribution is *light-tailed*.

Theorem 6 (Light-Tailedness at stationarity) *Suppose that Condition DD holds with $c_5 = 0$. Then there exists a $\gamma > 0$ such that*

$$\sum_{y \in \mathcal{S}} e^{\gamma \|y\|_v} \pi(y) < \infty.$$

The framework described above is very general and can be applied to any network that satisfies Condition DD. In what follows, we specialize the results for two wide classes of networks with mass-action kinetics, namely reaction networks with monomolecular and bimolecular reactions. It will be, however, pointed out in the examples that the scope of our approach is much broader since more general propensities, such as those involving Hill functions or more general mass-action kinetics, can be considered.

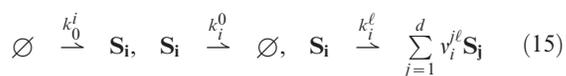
Methods

Using the analytical tools developed in the previous sections, several general results can be stated for the class of unimolecular reaction networks and bimolecular reaction networks. In what follows, when we say that a moment is bounded, we mean that it is bounded uniformly in time (as in (8)). This can be established using Theorem 2 once Condition DD is verified. Furthermore, when we say that a moment is globally converging, we mean that it converges to its equilibrium value as time tends to infinity, irrespective of the initial state x_0 . Once, Condition DD is verified, this can be established using Theorem 5.

The main aim of the section is to develop a theoretical and computational framework for checking Condition DD.

Results for stochastic unimolecular reaction networks

Let us then consider a unimolecular reaction network which involves d species that interact through K reaction channels of the form:



where $i = 1, \dots, d$, $\ell \in \{1, \dots, N_i\}$, $N_i > 0$ and $v_i^{j\ell} \in \mathbb{N}_0$. The reaction rates k_0^i , k_i^0 and k_i^ℓ are positive real numbers. In accordance with (3), the reactions are indexed from $n = 1$ to K , and corresponding propensities and stoichiometries are denoted by $\lambda_n(x)$ and ζ_n , respectively.

Motivations. The unimolecular case may seem quite restrictive at first sight and not of particular practical interest. We demonstrate below that, on the contrary, the proposed results on unimolecular reaction networks complete existing ones and are, therefore, of practical and theoretical interests. Although some explicit solutions for the CME are indeed known for some particular unimolecular reactions [16], it is still unknown whether the CME admits an closed-form solution for all possible type of unimolecular reactions. Note that no simplification nor assumption is made on the

problem in our work. Hence, we are dealing with the very general unimolecular case.

The results developed of this section are useful in several ways. First of all, all types of unimolecular reactions can be handled with the proposed approach, making it more general than existing ones in this regard. Moreover, given a specific reaction network, the method allows one to establish whether a unique stationary distribution exists *without solving the CME*. This is particularly important since unimolecular networks *may not be ergodic*. In this case, the network can exhibit unstable behaviour which may suggest a flaw in the model if the considered real-world system exhibits stable trajectories. Moreover, in certain design applications such as those in synthetic biology, it seems natural to design networks that have well-behaved dynamics. Checking ergodicity provides a convenient way to determine if the network dynamics is well-behaved. Note, furthermore, that it is, in general, difficult to infer ergodicity directly from the solution of the CME (when it is known) since proving the existence of a unique globally attractive stationary distribution amounts to checking convergence of the solution to the CME to the same distribution for all possible initial distributions, which are in infinite number in our setup. This fact is even more true when large networks are considered since the explicit form of the solution to the CME is, in this case, very intricate [16]. The proposed results allow one to circumvent this difficulty and demonstrate that ergodicity can be assessed by very simple means, i.e. using basic notions of linear algebra. The results can be furthermore used to assess the *structural ergodicity* of a reaction network, that is, the ergodicity of a network for any combination of the rate parameters, by very simple means. This very strong and practically relevant notion is extremely difficult, again, to check from the solution of the CME since it would require to check the convergence of the solution of the CME to the same stationary distribution for all initial conditions and all positive values of the rate parameters, a very cumbersome task, even for small networks. Finally, the results pertaining to unimolecular networks will also turn out to play an important role in the ergodicity analysis of bimolecular reaction networks.

Theoretical results. Let us start with several theoretical results that characterize the long-term behavior of unimolecular networks of the form (15).

Proposition 7 (Ergodicity of unimolecular networks)

Let us consider the general unimolecular reaction network (15) and assume that the state-space of the underlying Markov process is irreducible. Let the matrices $A \in \mathbb{R}^{d \times d}$ and $b \in \mathbb{R}_{\geq 0}^d$, $\|b\| \neq 0$, be further defined as

$$\sum_{n=1}^K \lambda_n(x) \langle v, \zeta_n \rangle = x^T A v + b^T v. \tag{16}$$

Then, the following statements are equivalent:

1. *The matrix A is Hurwitz-stable, i.e. all its eigenvalues lie in the open left half-plane.*
2. *There exists a vector $v \in \mathbb{R}_{> 0}^d$ such that $Av < 0$.*

Moreover, when one of the above statements holds, the Markov process describing the reaction network is exponentially ergodic and all the moments are bounded and globally converging.

The above result shows that, for unimolecular networks, ergodicity and the existence of moment bounds can be directly inferred from the properties of the matrix A defined in (16). The second statement, which characterizes the Hurwitz-stability of A in an implicit way, will turn out to play a key role in the analysis of

unimolecular and bimolecular reaction networks since checking whether $Av < 0$ for some $v > 0$ is a linear programming problem.

It is important to stress that, in the result above, if we simply demand that the moments be bounded and converging, then A may be allowed to have zero eigenvalues in certain cases. Note, however, that the moments will converge to values that may depend on the initial conditions.

In the case that the structure of the network (the reactions and the stoichiometries) is exactly known, but that the reaction rates are subject to uncertainties, the above theorem can be *robustified* to account for these uncertainties. To this aim, suppose that the matrix A depends on a vector $\delta \in [-1, 1]^\eta$ where $\eta \in \mathbb{N}$ is the number of distinct uncertain parameters. We write this matrix as $A(\delta)$ and assume that there exists a matrix $A_+ \in \mathbb{R}^{d \times d}$ satisfying the following properties:

1. $A(\delta) \leq A_+$ (in the componentwise sense) for all $\delta \in [-1, 1]^\eta$
2. There exists a $\delta^* \in [-1, 1]^\eta$ such that $A_+ = A(\delta^*)$.

Note that such a matrix A_+ may not exist, especially when some entries are not independent. However, when A_+ exists we have the following result.

Proposition 8 (Robust ergodicity) *Let us consider the general unimolecular reaction network (15) described by some uncertain matrices $A(\delta)$ and $b(\delta)$, $\|b(\delta)\| \neq 0$. Assume further the matrix $A(\delta)$ admits the upper-bound A_+ defined above and that the state-space of the underlying Markov process is irreducible for all uncertain parameter values $\delta \in [-1, 1]^\eta$. Then, the following statements are equivalent:*

1. The matrix $A(\delta)$ is Hurwitz-stable for all $\delta \in [-1, 1]^\eta$.
2. The matrix A_+ is Hurwitz-stable.
3. There exists a positive vector $v \in \mathbb{R}^d$ such that $A_+v < 0$.

Moreover, when one of the above statements holds, the Markov process describing the reaction network is robustly exponentially ergodic and all the moments are bounded and globally converging.

Observe that checking the Hurwitz-stability property of each $A(\delta)$ is equivalent to checking it for only A_+ . Hence we can conclude that, in this case, *checking ergodicity of a family of networks is not more complicated than checking ergodicity of a single network*. The case when the matrix A_+ is not defined is more complicated and is discussed in the supplementary material **S1**.

Computational results. We now present several computational results that accompany the theoretical results of the previous section. It is possible to extract many computational results from our general framework, but for simplicity we only address the problems of checking ergodicity and computing the first-order moment bounds. The asymptotic first-order moment bound, defined in Theorem 2, is given by $\hat{C}_1 = c_1/c_2$. So the question arises: what is the smallest value for such a ratio? Or, in other words, what is the smallest attractive compact set for the first-order moment of $\langle v, X(t) \rangle$? Several numerical methods, solving exactly or approximately this problem, are discussed in the supplementary material **S1**. One of them is the following optimization problem which is *fully equivalent to Proposition 7*:

Optimization problem 9 *Let us consider the general unimolecular reaction network (15) and assume that the state-space of the underlying Markov process is irreducible. Assume further that the optimization problem*

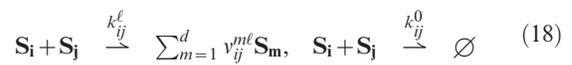
$$\max_{z,v} \quad s.t. \quad \begin{aligned} z > 0, v > \varepsilon \\ (zI + A)v \leq 0 \end{aligned} \quad (17)$$

is feasible with (z^, v^*) as minimizer. Then, we have $\hat{C}_1^* \leq b^\top v^* / z^*$ and Proposition 7 holds.*

A striking feature about the above optimization program is that the numbers of variables and constraints are given by $d + 1$ and $2d + 1$, respectively. This means that the optimization problem scales *linearly* with respect to the number of species (d) in the network, and is independent of the number of reactions K . Therefore, from the point of view of this optimization problem, the size of a unimolecular network can be identified with the number of species, and not the number of reactions. The above optimization problem can be efficiently solved using a bisection algorithm over z that is globally and geometrically converging to z^* . Each iteration consists of solving a linear program, a class of optimization problems known to be very tractable, and for which numerous advanced solvers exist [60]. These properties, altogether, make the overall approach highly scalable, which is necessary for dealing with very large networks.

Results for stochastic bimolecular reaction networks

Similar results are now presented for stochastic *bimolecular reaction networks* which, in addition to the unimolecular reactions (15), also involve bimolecular reactions of the form:



where $i, j = 1, \dots, d$, $\ell \in \{1, \dots, N_{ij}\}$, $N_{ij} > 0$, and $v_{ij}^{m\ell} \in \mathbb{N}_0$. The reaction rates k_{ij}^ℓ and k_{ij}^0 are positive real numbers.

Theoretical results for bimolecular networks. When *bimolecular reaction networks* of the form (15)–(18) are considered, the left-hand side of condition (5a) can be expressed as

$$\sum_{i=1}^K \lambda_k(x) \langle v, \zeta_k \rangle = x^\top M(v)x + x^\top Av + b^\top v \quad (19)$$

where $M(v) \in \mathbb{R}^{d \times d}$ is symmetric, $A \in \mathbb{R}^{d \times d}$ and $b \in \mathbb{R}_{\geq 0}^d$. Let $S := [\zeta_1 \ \dots \ \zeta_K]$ be the stoichiometry matrix of the *bimolecular reaction network* (15)–(18), and let S_q be the restriction of S to *bimolecular reactions*, only. Further define a set

$$\mathcal{N}_q := \{v \in \mathbb{R}^d : v > 0, v^\top S_q = 0\}.$$

When $v \in \mathcal{N}_q$, the quadratic term $x^\top M(v)x$ in (19) vanishes, and equation (19) reduces to

$$\sum_{i=1}^K \lambda_k(x) \langle v, \zeta_k \rangle = x^\top Av + b^\top v$$

which is exactly the same expression as in the case of unimolecular networks. This means that, with the additional constraint that $v \in \mathcal{N}_q$, all the results derived for unimolecular networks directly apply to *bimolecular networks* as well. This allows us to obtain the following result.

Proposition 10 (Ergodicity of networks) *Let us consider the bimolecular reaction network of the form (15)–(18) such that $\|b\| \neq 0$ in (19) and assume that the state-space of the underlying Markov process is irreducible. Assume further that the network admits a non-empty \mathcal{N}_q .*

If there exists a vector $v \in \mathcal{N}_q$ such that the inequality $Av < 0$ holds, then the stochastic bimolecular reaction network (15)–(18) is ergodic and all the moments are bounded and globally converging.

It is important to mention that the existence of a non-empty set \mathcal{N}_q is a prerequisite for utilizing the above result. Non-emptiness of \mathcal{N}_q is equivalent to the existence of a conservation

relation for all the bimolecular reactions, i.e. the value of (at least) a positive linear combination of the species populations remains unchanged when any of the bimolecular reactions fires. Note that this definition extends to more general mass-action kinetics as well. A necessary condition for the non-emptiness of \mathcal{N}_q is that S_q is not full-row rank. This non-emptiness condition may seem restrictive at first sight, but it will be shown that several important reaction networks from the literature satisfy this condition.

Whenever \mathcal{N}_q is empty or there is no $v \in \mathcal{N}_q$ such that $Av < 0$ holds, the next result can be used.

Proposition 11 (Ergodicity of bimolecular networks)

Let us consider the bimolecular reaction network of the form (15)–(18) such that $\|b\| \neq 0$ in (19) and assume that the state-space of the underlying Markov process is irreducible. Assume further that one of the following statements holds:

1. There exists $v \in \mathbb{R}_{>0}^d$ such that $Av < 0$ and $M(v) \leq 0$ hold.
2. There exists $v \in \mathbb{R}_{>0}^d$ such that $M(v)$ is negative definite.

Then, the stochastic bimolecular reaction network (15)–(18) is ergodic and all the moments up to order $(\lfloor 1 + 2c_2/c_5 \rfloor - 2)$ are bounded and globally converging.

In the above result, the first statement can be checked using a linear program since the inequalities are componentwise. Checking the second statement, however, requires a semidefinite program, which is a more general convex program, that can be solved using solvers such as SeDuMi [61] and SDPT3 [62]. More details on the above result can be found in the supplementary material **S1**.

Computational results for bimolecular networks. It is shown here that, once again, the theoretical results can be easily turned into linear programs that can be checked in a very efficient way. The following result is the numerical translation of Proposition 10.

Optimization problem 12 Let us consider a bimolecular reaction network (15)–(18) and assume that the state-space of the underlying Markov process is irreducible. Assume further that $\mathcal{N}_q \neq \emptyset$ and that the optimization problem

$$\begin{aligned} \max_{z,v} \quad & z \text{ s.t.} \quad z > 0, v > \epsilon \\ & (zI + A)v \leq 0 \\ & v^\top S_q = 0. \end{aligned} \tag{20}$$

is feasible with (z^*, v^*) as minimizer. Then, we have $\hat{C}_1^* \leq b^\top v^* / z^*$ and Proposition 10 holds.

The computational complexity of this optimization problem scales linearly with the number of species and can therefore be solved for large networks. The following optimization problem is the computational counterpart of the first statement of Proposition 11.

Optimization problem 13 Let us consider a bimolecular reaction network of the form (15)–(18) and assume that the state-space of the underlying Markov process is irreducible. Assume further that the nonlinear optimization problem

$$\begin{aligned} \max_{z,v} \quad & z \text{ s.t.} \quad z > 0, v > \epsilon \\ & (zI + A)v \leq 0 \\ & M(v) \leq 0. \end{aligned} \tag{21}$$

is feasible with (z^*, v^*) as minimizer. Then, we have $\hat{C}_1^* \leq b^\top v^* / z^*$ and Proposition 11 holds.

The above optimization problem does not scale as nicely as (20) since, in the worst case, the number of constraints related to $M(v)$ is quadratic in the number of species. The problem, however, remains tractable due to the linear programming structure.

Qualitative differences between deterministic and stochastic dynamics

In this section we illustrate that stochastic and deterministic models of the same reaction network may exhibit very different qualitative behaviors. Therefore assessing ergodicity or the convergence of moments of a stochastic model from the stability properties of the corresponding deterministic model is, in general, incorrect. To support this claim, we consider two reaction networks.

Jumping potential wells. Our first example shows that stochastic dynamics can jump potential wells and leave the stability regions of the deterministic dynamics, resulting in an unstable behavior. Consider the following reaction network:



where $0 < \alpha < \beta$. The deterministic dynamics for this network is given by

$$\dot{\kappa} = f(\kappa) := \kappa^2 - (\alpha + \beta)\kappa + \alpha\beta \tag{23}$$

where $\kappa \in \mathbb{R}_{\geq 0}$ denotes the concentration of **S**. The fixed points for the dynamics are $\kappa_- = \alpha$ and $\kappa_+ = \beta$, respectively. From the graph $\{(\kappa, f(\kappa)) \in \mathbb{R}_{\geq 0} \times \mathbb{R} : \kappa \in \mathbb{R}_{\geq 0}\}$, it is immediate that the fixed point $\kappa_- = \alpha$ is locally asymptotically stable with the region of attraction as $[0, \beta)$ while the other fixed point $\kappa_+ = \beta$ is unstable.

We now consider the stochastic version of this network and let \mathbb{A} be the generator of the corresponding Markov process. For the identity function $f(x) = x$ we have

$$\mathbb{A}f(x) = \frac{1}{2}x^2 - \left(\alpha + \beta + \frac{1}{2}\right)x + \alpha\beta. \tag{24}$$

The polynomial on the right-hand side has two positive roots that are

$$x_{\pm} = \alpha + \beta + \frac{1}{2} \pm \sqrt{\left(\alpha + \beta + \frac{1}{2}\right)^2 - 2\alpha\beta}. \tag{25}$$

This means that for all $x \in \mathbb{N}_0$ satisfying $x \geq 1 + x_+$, we have $\mathbb{A}f(x) \geq \epsilon$, for some $\epsilon > 0$, implying that the drift is positive. So if the state of the network reaches a value that is greater than $1 + x_+$, then there is a possibility that the trajectories become unbounded with time.

To demonstrate this, we pick $\alpha = 7/2$ and $\beta = 21/2$. In such a case, the largest root of the polynomial on the right-hand side of (24) is $x_+ = (29 + \sqrt{547})/2 \approx 26.194 > \beta$. We can see that the region where the drift $\mathbb{A}f(x)$ is negative is actually larger than the region of attraction of the locally asymptotically stable fixed point for the deterministic dynamics. This is due to the fact that the

propensity function of the bimolecular reaction differs from whether we are in the deterministic or in the stochastic setting.

Let us now set the initial condition $\kappa_0=0$ for the deterministic model and $x_0=0$ for the stochastic one. Note that they both lie within the region of attraction of the fixed point of the deterministic dynamics and in the region of negative drift for the stochastic dynamics. We then perform 1000 SSA runs over 100 seconds and stop the simulation when the propensity function $x(x-1)/2$ of the bimolecular reaction exceeds the value corresponding to 15000 molecules (approx. 1.12×10^8). At this rate value, the bimolecular reaction fires, on average, every 10^{-8} seconds, leading to an explosion of the state of the system and to unbounded trajectories. Out of 1000 SSA runs, all were stopped before the end of the simulation time-period (100 seconds). This behavior strongly indicates that the system is *not ergodic* despite the fact that the deterministic model has a locally asymptotically stable fixed point. Figure 1 illustrates the above discussion.

Globally stable deterministic dynamics does not imply moments stability. In the previous example, the stochastic and deterministic behaviors were different, but one can still understand stochastic instability through the deterministic model. The deterministic dynamics possesses a region in which the solutions explode and the randomness in the stochastic dynamics allows it to enter this region in finite time and grow unbounded thereafter. We now present an example which is more striking in the sense that the deterministic model cannot be used in any way to infer the instability of the stochastic model. In this example, the deterministic dynamics has a unique fixed point which is exponentially stable, while the stochastic dynamics is not ergodic with all its moments growing unboundedly with time.

Consider the reaction network given by



Let $\kappa \in \mathbb{R}_{\geq 0}^2$ be the vector of concentrations. The state of the deterministic model evolves according to

$$\begin{aligned} \dot{\kappa}_1(t) &= 1 - \kappa_1(t)\kappa_2(t) \\ \dot{\kappa}_2(t) &= 1 - \kappa_1(t)\kappa_2(t). \end{aligned} \tag{27}$$

Assume that the initial conditions satisfy $\kappa_2(0) - \kappa_1(0) = \alpha$, for some $\alpha \in \mathbb{R}$. Then we have the following result.

Theorem 14 *The unique equilibrium point of the dynamics (27) given by*

$$\kappa_1^* = \frac{1}{2} \left(-\alpha + \sqrt{\alpha^2 + 4} \right) \text{ and } \kappa_2^* = \frac{1}{2} \left(\alpha + \sqrt{\alpha^2 + 4} \right). \tag{28}$$

is globally exponentially stable.

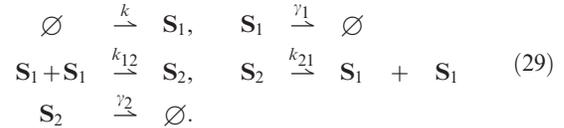
In the stochastic setting, the picture is completely different as the next result indicates.

Theorem 15 *The Markov process corresponding to the stochastic model of network (26) is not ergodic and all its moments grow unboundedly with time. Moreover, if $X_1(0) - X_2(0) = \alpha$ for some $\alpha > 0$, we have that $\mathbb{E}[X_1(t) - X_2(t)] = \alpha$ for all $t \geq 0$.*

To illustrate this result, we simulate the deterministic and the stochastic process (10000 SSA runs) for $\kappa_1(0)=0$, $\kappa_1(0)=\alpha$, $X_1(0)=0$, $X_2(0)=\alpha$ and $\alpha=2$. The results are shown in Figure 2.

Finding an attractive compact set for the first-order moments

The goal of this section is to compute a compact set that is attractive for the first-order moment of $\langle v, X(t) \rangle$ using the optimization problems (17) or (20). Due to the moment closure problem [54], analytical expressions for the steady-state values of the moments of bimolecular reaction networks are not available, and hence this is an important class of networks to analyze. Consider the following bimolecular reaction network



representing a dimerization process, i.e. \mathbf{S}_1 dimerizes to \mathbf{S}_2 . It is easily seen that this network is irreducible since any point in the state-space can be reached from any other point in a finite number of reactions having nonzero propensities. Choosing v in \mathcal{N}_q , e.g. $v^\top = [1 \ 2]$, yields that $c_1^* = k$ and $c_2^* = \min\{\gamma_1, \gamma_2\}$, hence the network is exponentially ergodic, and all the moments are bounded and converging. On solving the optimization problem (20) with numerical values $k=1$, $\gamma_1=\gamma_2=0.2$, $k_{12}=1$ and $k_{21}=0.1$, we get that $\hat{C}_1 = c_1^*/c_2^* = 5$ which coincides with the theoretical value $k/\min\{\gamma_1, \gamma_2\}$. One can regard $\{(x_1, x_2) \in \mathbb{R}_{>0}^2 : v^\top x \leq \hat{C}_1\}$ to be an attractive compact set in which the first-order moments of $\langle v, X(t) \rangle$ eventually lie. To validate this calculation, Monte-Carlo simulations were performed which yield

$$\lim_{t \rightarrow \infty} \mathbb{E}[\langle v, X(t) \rangle] = 5.024 \pm 0.05, \tag{30}$$

showing the correctness of the attractive compact set. To further illustrate this result, several trajectories of $\mathbb{E}[X_1(t)]$ and $\mathbb{E}[X_2(t)]$ for different initial conditions are plotted in Figure 3.

We now discuss how the computation of an attractive compact set for the first-order moments can be used to assess whether a closure method leads to a result that is consistent with the

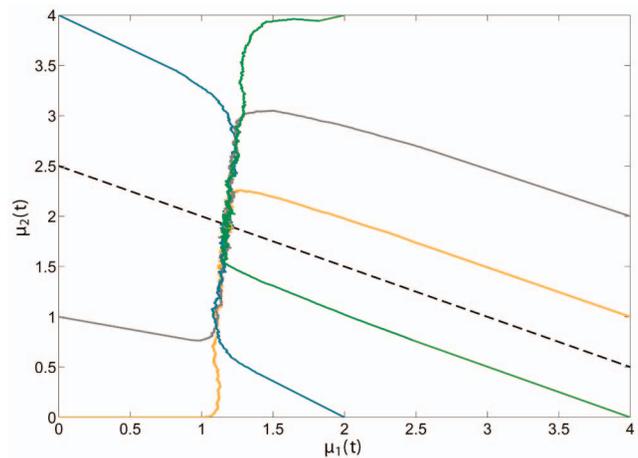


Figure 3. Trajectories of the first order moments $\mu_1(t) = \mathbb{E}[X_1(t)]$ and $\mu_2(t) = \mathbb{E}[X_2(t)]$ of network (29) for different initial conditions (averaging is performed over 5000 cells). The trajectories converge to the unique steady-state value located inside the compact set (the surface below the dashed line), very close to the boundary. doi:10.1371/journal.pcbi.1003669.g003

stochastic dynamics. The idea is to check whether the closed system converges towards a value which lies within the compact set. Let us consider the reaction network (29) and close the first-order moments equations by neglecting the second order cumulant, i.e. neglecting the variance. By doing so, we get the model

$$\begin{aligned} \dot{\tilde{\mu}}_1(t) &= k - \gamma_1 \tilde{\mu}_1(t) - k_{12} \tilde{\mu}_1(t)(\tilde{\mu}_1(t) - 1) + 2k_{21} \tilde{\mu}_2(t) \\ \dot{\tilde{\mu}}_2(t) &= k_{12} \tilde{\mu}_1(t)(\tilde{\mu}_1(t) - 1) - \gamma_2 \tilde{\mu}_2(t) \end{aligned} \quad (31)$$

where $\tilde{\mu}_1$ and $\tilde{\mu}_2$ are the approximate first-order moments of the system. The unique positive equilibrium point for this model is given by

$$\begin{aligned} \tilde{\mu}_1^* &= \frac{1}{2k} \left(-\gamma_1 + \frac{k_{12}\gamma_2}{\gamma_2 + k_{21}} + \sqrt{\Delta} \right) \\ \tilde{\mu}_2^* &= \frac{k_{12}}{2(\gamma_2 + k_{21})} \tilde{\mu}_1^*(\tilde{\mu}_1^* - 1) \end{aligned} \quad (32)$$

where $\Delta = \left(-\gamma_1 + \frac{k_{12}\gamma_2}{\gamma_2 + k_{21}} \right)^2 + \frac{4kk_{12}\gamma_2}{\gamma_2 + k_{21}}$.

With the same parameter values as before, we find that $\tilde{\mu}_1^* = 1.6238$ and $\tilde{\mu}_2^* = 1.6881$ and therefore $v^T \tilde{\mu}^* = 5$ for $v^T = [1 \ 2]$, showing that the state of the closed system converges to the boundary of the compact set. Note that SSA also predicts that the trajectories of the first-order moments converge to the boundary of this set. However the actual equilibrium values for the first-order moments of the stochastic dynamics are $\mu_1^* \simeq 1.1450$ and $\mu_2^* \simeq 1.9350$, which differ from the ones obtained with the closure method. This discrepancy is expected since the variance has been neglected.

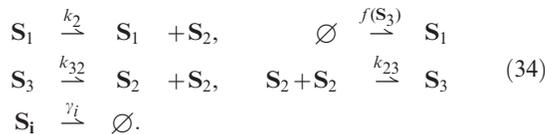
This example shows how attractive compact sets for the moments can be used as a test for the moment-closure methods by checking whether the closed system predicts trajectories that converge inside those compact sets. However, note that in the current state, these compact sets can only be used to obtain a lower bound on the closure-error whenever the trajectories of the closed dynamics converge to a point outside the compact set. In such a case, the lower bound on the closure-error ε is simply given by the distance between the equilibrium point of the closed-system

$$\varepsilon \geq \inf_{\theta \in C} \|\tilde{\mu}^* - \theta\|_2 \quad (33)$$

where C is the attractive (convex) compact set and $\tilde{\mu}^*$ is the equilibrium point of the closed dynamics.

Feedback loop

Let us consider the feedback loop network of Figure 4 represented by the reaction network



where S_1 is mRNA and S_2 is the corresponding protein. The dimer S_3 acts back on the gene expression through an arbitrary bounded nonnegative function $f(\cdot)$.

We have the following result:

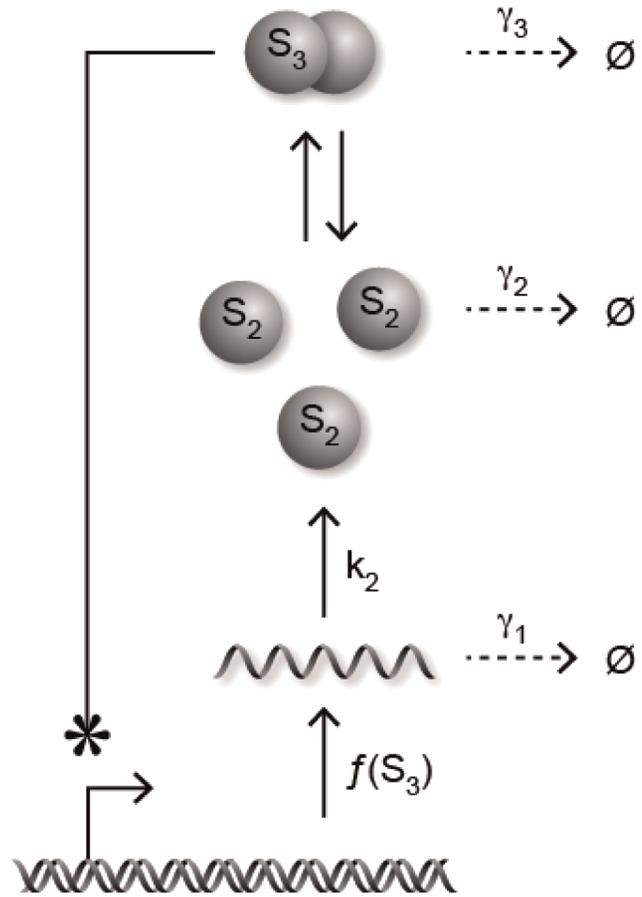
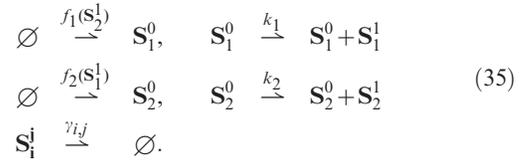


Figure 4. Feedback loop with arbitrary feedback rule.
doi:10.1371/journal.pcbi.1003669.g004

Result 16 For any positive values of the rate parameters and any bounded nonnegative function $f(\cdot)$, the feedback loop with dimerization (34) is ergodic and all the moments are bounded and globally converging.

Stochastic switch

Let us consider the stochastic switch of [63] described by the unimolecular stochastic reaction network



Above S_i^0 and S_i^1 represent mRNAs and proteins of gene i , respectively. The functions $f_1(\cdot)$ and $f_2(\cdot)$ are arbitrary bounded nonnegative functions. We have the following result:

Result 17 For any positive values of the rate parameters and any bounded nonnegative functions $f_1(\cdot)$ and $f_2(\cdot)$, the stochastic switch (35) is ergodic and all the moments are bounded and globally converging.

Repressilator

We consider here the stochastic repressilator of Figure 5 (see also [42]) involving N genes. The reaction network corresponding to this N -gene repressilator is given by

which is an open analogue of the deterministic Lotka-Volterra system of [68]. The first set of reactions represent immigration, the second one reproduction, the third one competition due to overpopulation and the last one deaths/migrations. We obtain then the following result, which is a stochastic analogue of the results in [69] obtained in the deterministic setting:

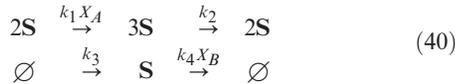
Theorem 22 *Let us define $\Gamma(v)=[v_i\gamma_{ij}]$ and assume that one of the following conditions hold:*

1. there exists $v > 0$ such that the matrix $\Gamma(v) + \Gamma(v)^\top$ is positive definite;
2. there exists $v > 0$ such that the $\Gamma(v) + \Gamma(v)^\top$ is copositive, i.e. $x^\top(\Gamma(v) + \Gamma(v)^\top)x \geq 0$ for all $x \geq 0$, and $\beta_i - \delta_i < 0$ for all $i = 1, \dots, n$.

Then, the stochastic reaction network (39) is ergodic and all the moments up to order $\lfloor 1 + \frac{2c_2}{c_5} \rfloor - 2$ are bounded and globally converging.

Schlögl model

In order to illustrate that the method can be applied to systems with more general mass-action kinetics, we consider the stochastic version of the well-known Schlögl model [70]:



where **S** is the main molecule in the network. The above model is derived in the supplementary material **S1** where we have assumed that the other molecular populations do not vary over time. Note that in the present form the model has an infinite state-space and involves a single trimolecular reaction. We then have the following result.

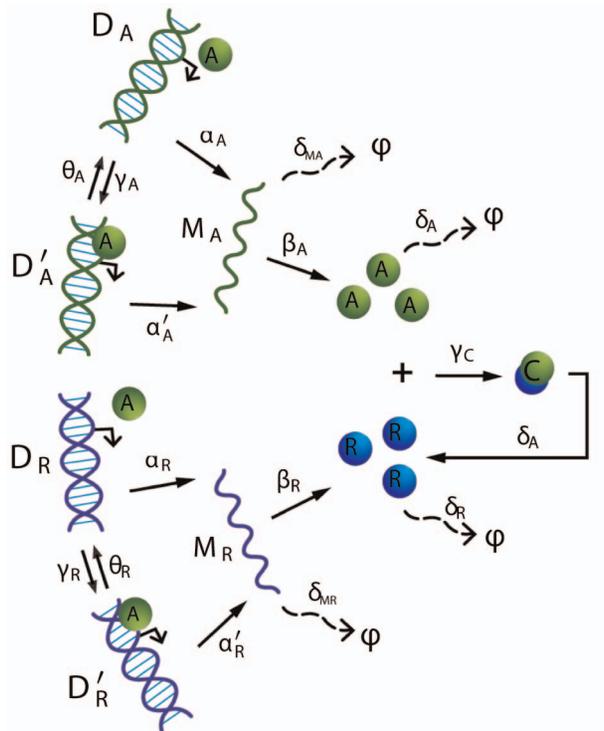


Figure 6. Circadian clock model of [65].
doi:10.1371/journal.pcbi.1003669.g006

Theorem 23 *For any positive values of the rate parameters k_1, k_2, k_3, k_4 and any positive values for X_A and X_B , the Markov process describing the Schlögl model (40) is exponentially ergodic.*

Note, however, that we cannot say anything on the stability of the moments (besides the fact that the first order-moment converges) since the condition **DD2** does not hold here due to the presence of a cubic term. Note that extending the condition **DD2** to handle more general cases, such as this one, might be possible.

Discussion

The central theme of this paper is to verify the ergodicity and moment boundedness of reaction networks in the stochastic setting. Note that even though we mainly consider mass-action kinetics in this paper, the framework also applies to more general kinetics described, for instance, by Hill functions (see the examples on the repressilator and the stochastic switch) and more general mass-action kinetics. These results have several interesting and important biological implications.

For example, the ergodicity of a network shows that population-level information could be obtained by observing a single trajectory for a long time. Such an insight can be used to leverage different experimental techniques for a given application. For example, consider a clonal cell population with each cell having a gene-expression network that is ergodic. Then the stationary distribution (at the population level) of the species involved in this network can be ascertained by *observing a single cell* over time. In other words, to obtain stationary distributions one can either collect samples over time from a single cell (e.g. using time-lapse microscopy) or one can take a snapshot of the entire cell population at some fixed time (e.g. using flow-cytometry). Due to ergodicity, both these approaches will yield the same information. Hence, far from being a technical condition, ergodicity can have far reaching experimental implications.

As a property of a network, ergodicity also sheds important light on the long range behaviors that can be exhibited by that network. One may expect that most endogenous biochemical networks to be ergodic in order to achieve robustness with respect to variability in initial conditions and kinetic parameters, thus ensuring proper biological functions in spite of environmental disturbances. As also mentioned in the introduction, ergodicity is a non-trivial property which needs to be carefully established and cannot be generically

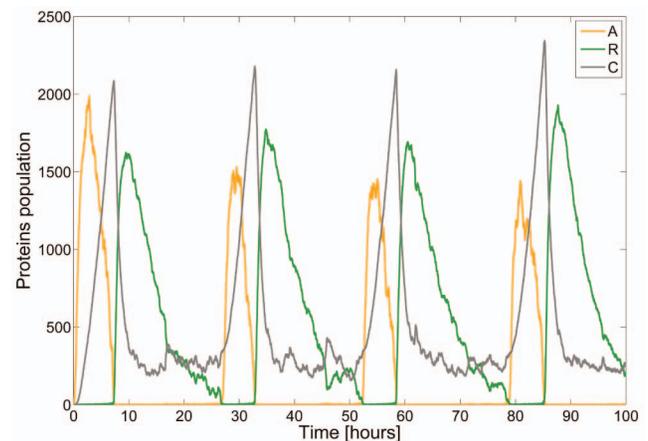


Figure 7. Sample-path of the species of the circadian clock model.
doi:10.1371/journal.pcbi.1003669.g007

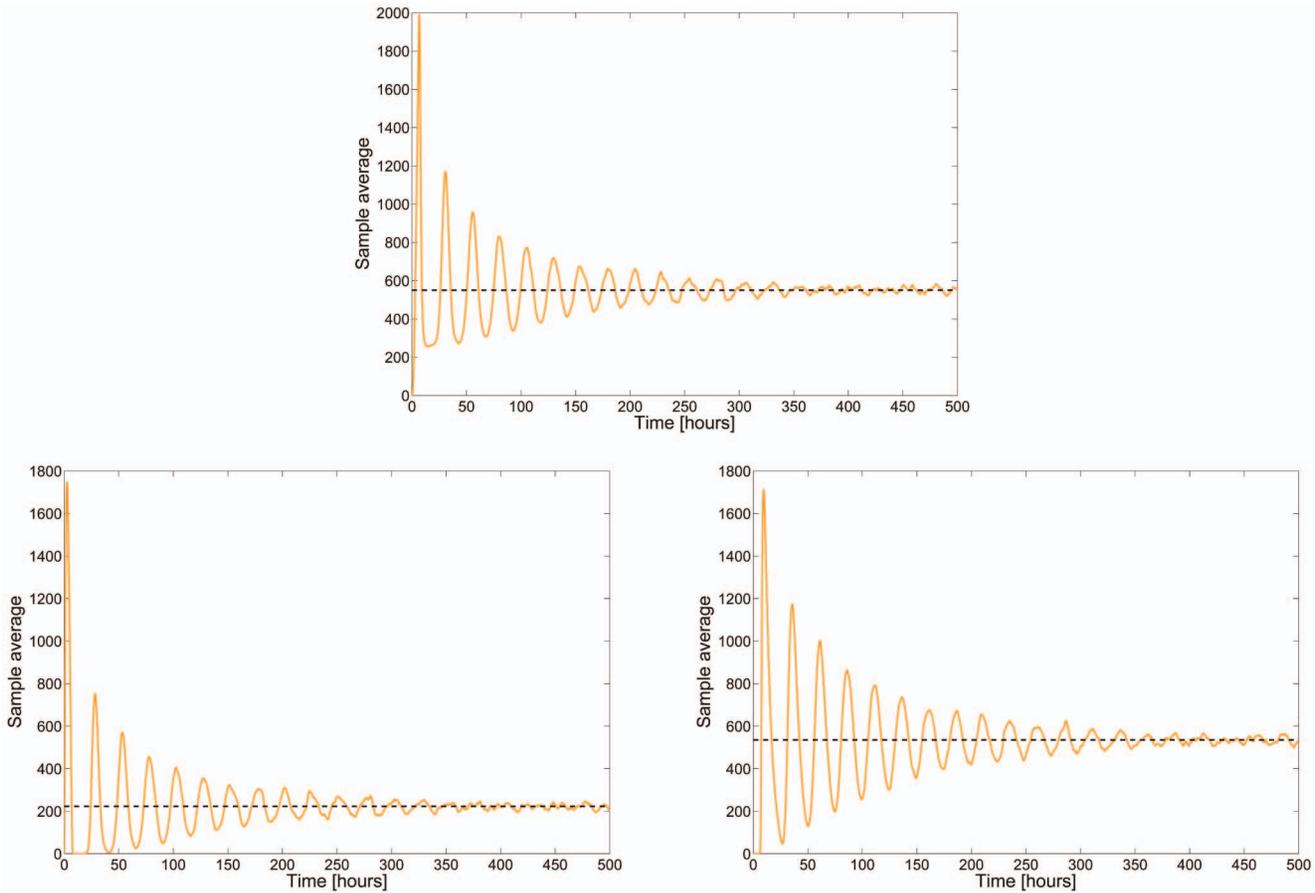
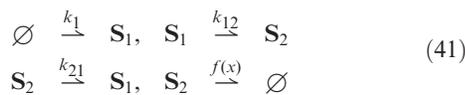


Figure 8. Time evolution of the sample averages of the species A (top), R (left) and C (right) of the circadian clock model (2000 cells averaging). The dashed-lines correspond to the (asymptotic) time-average. doi:10.1371/journal.pcbi.1003669.g008

assumed. To illustrate this, let us consider a simplified version of the model of carcinogenesis considered in [71] which is given by



where $f(x) = \frac{\gamma_2}{\alpha + x_2}$, $\alpha > 0$. When $k_1 > \gamma_2$, the trajectories of the species grow unbounded, as shown in Figure 9, emphasizing then non-ergodicity of the model for this choice of parameters.

The ideas we use for analysis can also be applied for rationally designing circuits in synthetic biology, where it is important that the network be (structurally) ergodic in order to ensure that the dynamics has the desired behavior irrespective of the initial conditions. Such a design is crucial because the initial conditions are usually unknown or difficult to control at certain times, e.g. after cell division or after the transfection of plasmids in the cell.

Our results on boundedness and convergence of statistical moments enable verification of the suitability of a stochastic model and to characterize the properties of its steady-state distributions, even if such a distribution is not explicitly computable. One application of this is to provide justifications and insights for using moment closure techniques which have been extensively used to study stochastic chemical reaction networks. Some of these

techniques [72,73] are based on manipulations of the moment generating function of the underlying stochastic process. The existence of this moment generating function is implicitly assumed in such techniques but it may not always hold, thereby

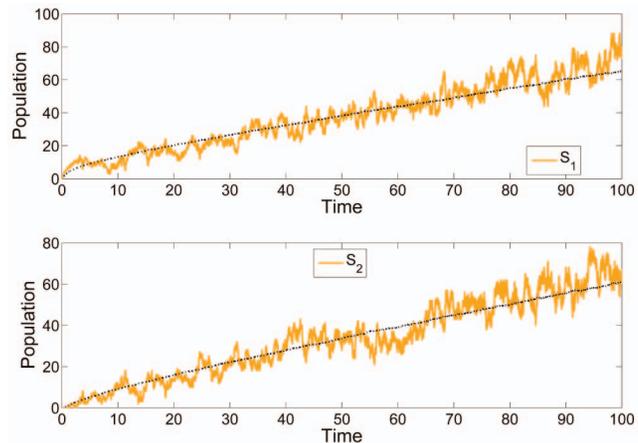


Figure 9. State trajectories of the carcinogenesis model (41) with the parameters $k_1=5$, $k_{12}=1$, $k_{21}=1$, $\gamma_2=4$ and $\alpha=1$. The dashed lines correspond to the average trajectories computed over 1000 cells. doi:10.1371/journal.pcbi.1003669.g009

jeopardizing the validity of the technique. In this article, we show that under certain conditions, the distribution of the stochastic process is uniformly light-tailed, which proves that the moment generating function exists for all time. Certain moment closure techniques (see [74,75]) prescribe ways to approximate higher order moments as a function of lower order moments. Such an approximation is, however, only reasonable if the higher order moments are bounded over time. This can be easily assessed with our approach and one can even quantify the error by explicitly computing the moment bounds as described in this article.

Finally, the techniques developed here will prove invaluable for designing synthetic biological control systems and circuits whose objective is to steer the moments of the network of interest to a specific steady-state value. Until now, no theory has provided guidance for such a design. The specifics are outside the scope of this article and will be pursued elsewhere.

References

- Érdi P, Tóth J (1989) *Mathematical models of chemical reactions*. Princeton, NJ: Princeton University Press, xxiv+259 pp.
- Papin JA, Hunter T, Palsson BO, Subramaniam S (2005) Reconstruction of cellular signalling networks and analysis of their properties. *Nat Rev Mol Cell Biol* 6: 99–111.
- Thattai M, van Oudenaarden A (2001) Intrinsic noise in gene regulatory networks. *Proceedings of the National Academy of Sciences* 98: 8614–8619.
- Schuetz R, Zamboni N, Zampieri M, Heinemann M, Sauer U (2012) Multidimensional optimality of microbial metabolism. *Science* 336: 601–604.
- Berger SI, Iyengar R (2009) Network analyses in systems pharmacology. *Bioinformatics* 25: 2466–2472.
- Hethcote H (2000) The mathematics of infectious diseases. *SIAM Review* 42: 599–653.
- Bascompte J (2010) Structure and dynamics of ecological networks. *Science* 329: 765–766.
- Goutsias J (2007) Classical versus stochastic kinetics modeling of biochemical reaction systems. *Biophysical Journal* 92: 2350–2365.
- McAdams HH, Arkin A (1999) It's a noisy business! Genetic regulation at the nanomolar scale. *Trends in genetics : TIG* 15: 65–69.
- Levin MD, Morton-Firth CJ, Abouhamad WN, Bourret RB, Bray D (1998) Origins of individual swimming behavior in bacteria. *Biophysical Journal* 74: 175–181.
- Elowitz MB, Levine AJ, Siggia ED, Swain PS (2002) Stochastic gene expression in a single cell. *Science* 297: 1183–1186.
- Arkin AP, Rao CV, Wolf DM (2002) Control, exploitation and tolerance of intracellular noise. *Nature* 420: 231–237.
- Kierzek AM, Zaim J, Zielenkiewicz P (2001) The effect of transcription and translation initiation frequencies on the stochastic fluctuations in prokaryotic gene expression. *Journal of Biological Chemistry* 276: 8165–8172.
- McAdams HH, Arkin A (1997) Stochastic mechanisms in gene expression. *Proc Natl Acad Sci, Biochemistry* 94: 814–819.
- Gillespie DT (1997) A rigorous derivation of the chemical master equation. *Physica A* 188: 404–425.
- Jahnke T, Huisinga W (2007) Solving the chemical master equation for monomolecular reaction systems analytically. *Journal of Mathematical Biology* 54: 1–26.
- Grima R, Schmidt DR, Newman TJ (2012) Steady-state fluctuations of a genetic feedback loop: An exact solution. *The Journal of Chemical Physics* 137: 035104.
- Munsky B, Khammash M (2006) The finite state projection algorithm for the solution of the chemical master equation. *Journal of Chemical Physics* 124.
- Sidje RB (1998) Expokit: A software package for computing matrix exponentials. *ACM Transactions on Mathematical Software* 24(1): 130–156.
- Jenkinson G, Goutsias J (2012) Numerical integration of the master equations in some models of stochastic epidemiology. *PLOS One* 7(5): e36160.
- Dolgov SV, Khoromskij BN (2012) Tensor-product approach to global time-space-parametric discretization of chemical master equation. Technical report, Max-Planck-Institut für Mathematik in den Naturwissenschaften.
- Kaazev V, Khammash M, Nip M, Schwab C (2014) Direct solution of the chemical master equation using quantized tensor trains. *PLOS Computational Biology* 10(3): e1003359.
- Gillespie DT (1976) A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *Journal of Computational Physics* 22(4): 403–434.
- Gillespie DT (1977) Exact stochastic simulation of coupled chemical reactions. *The Journal of Physical Chemistry* 81(25): 2340–2361.
- Gibson MA, Bruck J (2000) Efficient exact stochastic simulation of chemical systems with many species and many channels. *The Journal of Physical Chemistry A* 104: 1876–1889.
- Kurtz TG (1971) Limit theorems for sequences of jump Markov processes approximating ordinary differential processes. *J Appl Probability* 8: 344–356.
- van Kampen NG (1961) A power series expansion of the master equation. *Canad J Phys* 39: 551–567.
- Kurtz TG (1976) Limit theorems and diffusion approximations for density dependent Markov chains. *Math Programming Stud* : 67–78.
- Goutsias J, Jenkinson G (2013) Markovian dynamics on complex reaction networks. *Physics Reports* 529: 199–264.
- Khalil HK (1992) *Nonlinear systems*. New York: Macmillan Publishing Company, xii+564 pp.
- Meyn S, Tweedie RL (2009) *Markov chains and stochastic stability*. Cambridge: Cambridge University Press, xxviii+594 pp.
- Earnshaw BA, Keener JP (2010) Global asymptotic stability of solutions of nonautonomous master equations. *SIAM J Applied Dynamical Systems* 9: 220–237.
- Schnakenberg J (1976) Network theory of microscopic and macroscopic behavior of master equation systems. *Rev Mod Phys* 48: 571–585.
- Laurenzi IJ (2000) An analytical solution of the stochastic master equation for reversible bimolecular reaction kinetics. *The Journal of Chemical Physics* 113: 3315–3322.
- Anderson DF, Craciun G, Kurtz TG (2010) Product-form stationary distributions for deficiency zero chemical reaction networks. *Bull Math Biol* 72: 1947–1970.
- Lemarchand H (1980) Asymptotic solution of the master equation near a nonequilibrium transition: The stationary solutions. *Physica A: Statistical Mechanics and its Applications* 101: 518–534.
- Malek Mansour M, Van Den Broeck C, Nicolis G, Turner JW (1981) Asymptotic properties of markovian master equations. *Annals of Physics* 131: 283–313.
- Engblom S (2012) On the stability of stochastic jump kinetics. *ArXiv:12023892*.
- Rathinam M (2014) Moment growth bounds on continuous time markov processes on non-negative integer lattices. *Quarterly of Applied Mathematics: [In press]*.
- Dayar T, Hermanns H, Spieler D, Wolf V (2011) Bounding the equilibrium distribution of markov population models. *Numerical Linear Algebra with Applications* 18: 931–946.
- Gardner TS, Cantor CR, Collins JJ (2000) Construction of a genetic toggle switch in *escherichia coli*. *Nature* 403: 339–342.
- Elowitz MB, Leibler S (2000) A synthetic oscillatory network of transcriptional regulators. *Nature* 403: 335–338.
- Vellela M, Qian H (2008) Stochastic dynamics and non-equilibrium thermodynamics of a bistable chemical system: the schlögl model revisited. *Journal of the Royal Society Interface* 6(39): 925–940.
- Hespanha JP (2008) Moment closure for biochemical networks. In: *3rd International Symposium on Communications, Control and Signal Processing*. St. Julian's, Malta, pp. 142–147.
- Grima R (2012) A study of the accuracy of moment-closure approximations for stochastic chemical kinetics. *The Journal of Chemical Physics* 136: 154105.
- Smale S (2004) *Differential Equations, Dynamical Systems & an introduction to Chaos*. Academic Press.
- Feinberg M (1972/73) Complex balancing in general kinetic systems. *Arch Rational Mech Anal* 49: 187–194.
- Feinberg M (1995) The existence and uniqueness of steady states for a class of chemical reaction networks. *Arch Rational Mech Anal* 132: 311–370.
- Horn F, Jackson R (1972) General mass action kinetics. *Arch Rational Mech Anal* 47: 81–116.
- Horn F (1972/73) Necessary and sufficient conditions for complex balancing in chemical kinetics. *Arch Rational Mech Anal* 49: 172–186.

Supporting Information

Text S1 Text supporting information file with full mathematical proofs and more detailed examples. (PDF)

Acknowledgments

The authors are grateful to Stephanie Aoki and Christine Khammash who spent some of their precious time in producing several illustrative pictures.

Author Contributions

Conceived and designed the experiments: AG CB MK. Performed the experiments: CB AG. Analyzed the data: AG CB MK. Wrote the paper: AG CB MK. Developed the mathematical framework: AG. Developed the results for unimolecular and bimolecular reaction networks, and applied them to the examples: CB.

51. Ethier SN, Kurtz TG (1986) Markov processes : Characterization and Convergence. New York: John Wiley & Sons Inc., x+534 pp.
52. Kallenberg O (2002) Foundations of modern probability. New York: Springer-Verlag, xx+638 pp.
53. Pemantle R, Rosenthal JS (1999) Moment conditions for a sequence with negative drift to be uniformly bounded in l_r . Stochastic Processes and their Applications 82: 143–155.
54. Hespanha JP (2008) Moment closure for biochemical networks. In: 3rd International Symposium on Communications, Control and Signal Processing. St. Julian's, Malta, pp. 142–147.
55. Pendar H, Platini T, Kulkarni RV (2013) Exact protein distributions for stochastic models of gene expression using partitioning of poisson processes. Physical Review E 87: 042720.
56. Paulevé L, Craciun G, Koepl H (2013) Dynamical properties of discrete reaction networks. Journal of Mathematical Biology : 1–18.
57. Gupta A, Khammash M (2013) Determining the long-term behavior of cell populations: A new procedure for detecting ergodicity in large stochastic reaction networks. Technical Report arXiv:1312.2879, ETH-Zürich.
58. Norris JR (1998) Markov chains. Cambridge: Cambridge University Press, xvi+237 pp.
59. Meyn SP, Tweedie RL (1993) Stability of Markovian processes. III. Foster-Lyapunov criteria for continuous-time processes. Adv in Appl Probab 25: 518–548.
60. Boyd S, Vandenberghe L (2004) Convex Optimization. Cambridge, MA, USA: Cambridge University Press.
61. Sturm JF (2001) Using SEDUMI 1.02, a Matlab Toolbox for Optimization Over Symmetric Cones. Optimization Methods and Software 11: 625–653.
62. Tütüncü RH, Toh KC, Todd MJ (2003) Solving semidefinite-quadratic-linear programs using SDPT3. Mathematical Programming Ser B 95: 189–217.
63. Tian T, Burrage K (2006) Stochastic models for regulatory networks of the genetic toggle switch. Proc Natl Acad Sci 103(22): 8372–8377.
64. Chen WY, Bokka S (2005) Stochastic modeling of nonlinear epidemiology. J theor Biol 234: 455–470.
65. Vilar JMG, Kueh HY, Barkai N, Leibler S (2002) Mechanisms of noise-resistance in genetic oscillator. Proc Natl Acad Sci 99(9): 5988–5992.
66. Löfberg J (2004) Yalmip : A toolbox for modeling and optimization in MATLAB. <http://control.ee.ethz.ch/~joloef/yalmip.php>
67. Geva-Zatorsky N, Rosenfeld N, Itzkovitz S, Milo R, Sigal A, et al. (2006) Oscillations and variability in the p53 system. Molecular Systems Biology 2: 2006.0033.
68. Gopalsamy K (1984) Global asymptotic stability in volterra's population systems. Journal of Mathematical Biology 19: 157–168.
69. Champagnat N, Jabin P, Raoul G (2010) Convergence to equilibrium in competitive Lotka-Volterra and chemostat systems. Comptes Rendus de l'Académie des Sciences - Mathématique 348(23–24): 1267–1272.
70. Schlögl F (1972) Chemical reaction models for non-equilibrium phase transition. Zeitschrift für Physik 253(2): 147–161.
71. Bois FY, Zeise L, Tozer TN (1990) Precision and sensitivity of pharmacokinetic models for cancer risk assessment: Tetrachloroethylene in mice, rats, and humans. Toxicology and applied pharmacology 102: 300–315.
72. Lee C, Kim KH, Kim P (2009) A moment closure method for stochastic reaction networks. The Journal of Chemical Physics 130(13): 134107.
73. Ale A, Kirk P, Stumpf MP (2013) A general moment expansion method for stochastic kinetic models. The Journal of Chemical Physics 138(17): 174101.
74. Gomez-Urbe CA, Verghese GC (2007) Mass fluctuation kinetics: Capturing stochastic effects in systems of chemical reactions through coupled mean-variance computations. The Journal of Chemical Physics 126(2): 024109.
75. Singh A, Hespanha JP (2011) Approximate moment dynamics for chemically reacting systems. IEEE Transactions on Automatic Control 56(2): 414–418.