

RESEARCH ARTICLE

Beyond the chemical master equation: Stochastic chemical kinetics coupled with auxiliary processes

Davin Lunz ^{1,2,3,4*}, Gregory Batt ^{3,4}, Jakob Ruess ^{3,4}[✉], J. Frédéric Bonnans ^{1,2}[✉]

1 Inria Saclay – Île de France, Palaiseau, France, **2** École Polytechnique, CMAP, Palaiseau, France, **3** Inria Paris, Paris, France, **4** Institut Pasteur, Paris, France

 These authors contributed equally to this work.

* davin.lunz@inria.fr



Abstract

The chemical master equation and its continuum approximations are indispensable tools in the modeling of chemical reaction networks. These are routinely used to capture complex nonlinear phenomena such as multimodality as well as transient events such as first-passage times, that accurately characterise a plethora of biological and chemical processes. However, some mechanisms, such as heterogeneous cellular growth or phenotypic selection at the population level, cannot be represented by the master equation and thus have been tackled separately. In this work, we propose a unifying framework that augments the chemical master equation to capture such auxiliary dynamics, and we develop and analyse a numerical solver that accurately simulates the system dynamics. We showcase these contributions by casting a diverse array of examples from the literature within this framework and applying the solver to both match and extend previous studies. Analytical calculations performed for each example validate our numerical results and benchmark the solver implementation.

OPEN ACCESS

Citation: Lunz D, Batt G, Ruess J, Bonnans JF (2021) Beyond the chemical master equation: Stochastic chemical kinetics coupled with auxiliary processes. *PLoS Comput Biol* 17(7): e1009214. <https://doi.org/10.1371/journal.pcbi.1009214>

Editor: Vassily Hatzimanikatis, Ecole Polytechnique Fédérale de Lausanne, SWITZERLAND

Received: November 9, 2020

Accepted: June 25, 2021

Published: July 28, 2021

Copyright: © 2021 Lunz et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: The code described in the text and used to produce all the results is freely available in the gitlab repository accessible at: <https://gitlab.inria.fr/dlunz/flips>.

Funding: This work was partially supported by the Inria Project Lab CoSy [GB, JR, JFB] and by the ANR grants ANR-18-CE91-0002 (CyberCircuits) [JR] and ANR-16-CE33-0018 (MEMIP) [GB]. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Author summary

Populations of genetically identical cells tend to exhibit remarkable variability. This seemingly counter-intuitive observation has broad and fascinating implications, and has thus been a focal point of biological modeling. Many important processes act on this cellular heterogeneity at the population level, leading to an intricate coupling between the single-cell and the population-level dynamics. For example, selection pressures or growth rates may depend crucially on the expression of a particular gene (or gene family). Classical single-cell modeling approaches, such as the chemical master equation, can accurately describe the mechanisms driving cellular noise, however, they cannot encapsulate how the aforementioned auxiliary processes affect the population composition. In this work, we propose a unifying framework that extends the classical chemical master equation to faithfully capture the single-cell variability alongside the population-level evolution. We develop, analyse, and showcase an open-source numerical tool to simulate the dynamics

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

of such populations in time. The tool is designed for straightforward use by a non-technical audience: a high-level description of the underlying chemical and population-level processes suffices to simulate complex system dynamics. Simultaneously, we retain high customisability of the underlying mathematical representation for the more advanced user. Ultimately, the unifying framework and the associated computational tool open new horizons in the study of how fundamental microscopic dynamics give rise to complex macroscopic phenomena.

This is a *PLOS Computational Biology Methods* paper.

Introduction

The chemical master equation (CME) [1] governs the evolution of the probability distribution of the configuration, or state, of a reaction network. Typically, the configuration describes the number of molecules of various species, such as chemical reactants undergoing reactions. The CME and its continuum approximations, in particular the Fokker–Planck approximation [2], are fundamental modeling tools used for describing chemical reaction networks. These reaction networks find broad application across the quantitative sciences, providing accurate descriptions of a wide variety of chemical, biological and social phenomena [3]. Crucially, this family of models accounts for stochasticity inherent in reaction processes that play functional roles in biochemical contexts [4, 5]. In a population of individuals each governed by identical reaction kinetics, such as gene expression in a population of cells, this stochasticity manifests in cell-to-cell variability. When a state-dependent selection pressure is exerted upon the population, for example, phenotypic growth or cell-fate decisions, individuals are affected heterogeneously due to this cell-to-cell variability [6, 7]. The dynamics of such phenomena cannot be captured simply by a stochastic reaction network [8]. Models that incorporate population-level selection may break detailed balance and introduce nonlinearity into the master equation [9], rendering explicit solutions unknown in all but the simplest of circumstances [10]. We depict the example of state-dependent growth in a population of cells in Fig 1.

Recently this problem was addressed by Duso and Zechner [11] by tracking the number of cells with each possible internal state (as opposed to tracking only the state of a single cell). Biochemical reactions driving a single cell to transition from one state to another are encapsulated by an associated change to the population count: one fewer in the old state, one more in the new state. While the total population is preserved under such transitions, changes to the population composition, such as the introduction or removal of cells, may be included by laws governing non-conservative counter increments or decrements. This formalism is equivalent to a Markov chain on the state space of mappings from internal configurations to the natural numbers. While conceptually appealing, such a state space is an exponential inflation of the single-cell state space. For example, the dynamics of fewer than 100 molecules in each cell (of just a single species) in fewer than 50 cells, corresponds to a state space of size 10^{100} , far more than the number of known atoms in the universe. This renders calculations of the full distribution infeasible, and motivates the moment closure approximations studied by Duso and Zechner [11].




Fig 1. Schematic depiction of an auxiliary process modifying the population-level density. The measurable concentration of a nominal “blue” chemical species within a cell that is the downstream result of some chemical reaction network. The probability density of concentration in a single cell (left) is identical to the number density in a large population of identical cells (centre). However, when a state-dependent pressure (such as amplified growth for higher blue concentration) is exerted on the population (right), the distribution is affected in a way that may not be encapsulated by a classical CME. In this paper, we study a framework in which chemical reaction networks are coupled to such auxiliary processes, and we introduce a software package to simulate the associated dynamics.

<https://doi.org/10.1371/journal.pcbi.1009214.g001>

In this study, we develop a modeling framework that reaches beyond the classical CME to incorporate such population-level processes while effectively preserving the state space of the single-cell Markov chain. Keeping in mind that when the typical number of a species is large the CME is challenging to solve, we allow a continuum approximation in the form of the Fokker–Planck equation. We also allow non-local transitions to incorporate phenomena discontinuous in state such as cell division and fragmentation [12, ch. 4], as well as phenomena where production machinery operates on timescales significantly shorter than the timescale of interest [13], such as in protein production in “bursts”. This class of models finds application beyond biochemical physics, such as in communications protocols [14].

Crucially, not all species need be cast in the continuum, and thus we retain a hybrid structure: species of typically large number are in the continuum, while the remaining species retain a discrete description. This discrete structure allows us not only to capture discrete molecule numbers, but also to model abstract system states. For example, we may consider transcription factors in gene expression in a bound or unbound state, or distinguish cells based on phases of the cell cycle, as well as cell-fate decisions such as differentiation and recombination. This demonstrates the importance of retaining a discrete structure in the framework, which allows us to capture both discrete molecule numbers, as well as discrete abstract states. In this study, we leverage the discrete structure to capture first-passage times [15], providing insights that lend a new and extended perspective to existing results.

Some population-level processes may be resolved by leveraging this same discrete structure in non-physical ways. In this case, the discrete states are not simply abstract states, as they may no longer correspond to physical configurations at all. For example, we may need to consider negative “reaction” rates, breaching the physical description of a classical chemical reaction network. We highlight the distinction by calling these phantom states. By constructing systems incorporating phantom states, we show that we are able to couple CME dynamics to non-standard processes by casting them within our general framework.

Finally, we show how other classes of auxiliary processes fall within this framework, such as growth–fragmentation processes used to describe cell growth and division, as well as several other physical phenomena.

Our approach is conceptually akin to that pursued by Thomas [7], whereby averaging stochastic fluctuations due to small population sizes allows us to track the population-density evolution on the original state space. This approach is very useful for large populations where such fluctuations may be neglected. We generalise the class of problems studied in Thomas [7] by including a continuum formulation for internal states, allowing state-dependent rates, and resolving transient distributions. For small population sizes where stochastic fluctuations must be accounted for, the framework developed by Duso and Zechner [11] is more appropriate.

With the goal of open and collaborative progress, we make our code publicly available [16] in the hope that this modeling tool can be of broader use for the community. The code is generic and has been written for arbitrary reaction networks, growth–fragmentation and similar models. The interface is intended to be straightforward for the first-time user, while also allowing more advanced control of the numerical scheme, including discretisation and time-stepping details.

The rest of the paper is structured as follows. First, we introduce the encompassing framework by defining the class of problems under consideration. We then outline the numerical software developed to solve problems cast in this framework. The results are presented as a series of case studies of auxiliary processes not classically captured by the CME: a first-passage time problem in self-regulated gene expression, a birth–death process under population-wide growth pressure, and a growth–fragmentation model. Finally, we discuss the results and their implications. For the sake of completeness, in Section 1 in [S1 Appendix](#) we include an extensive description of the numerical scheme implemented in the software.

Methodology

Augmented CME framework

We consider a stochastic reaction network of d species, which we write in vector form $\mathbf{X} := (X_1, X_2, \dots, X_d)^T$, undergoing reactions labeled by $i \in \mathcal{I}$, which we write as



for reactant species and quantities determined by $\mathbf{c}_i \in \mathbb{Z}^d$ and product species and quantities determined by $\mathbf{d}_i \in \mathbb{Z}^d$. The reaction rate R_i may depend on the system state and time. The system state is a continuous-time Markov chain and may be described by its law, $P(\mathbf{X}, t)$, the probability of being in state \mathbf{X} at time t . The chemical master equation (CME), governing the evolution of the law [1], takes the form

$$\frac{\partial P(\mathbf{X}, t)}{\partial t} = \sum_{i \in \mathcal{I}} R_i(\mathbf{X} + \mathbf{c}_i - \mathbf{d}_i, t) P(\mathbf{X} + \mathbf{c}_i - \mathbf{d}_i, t) - R_i(\mathbf{X}, t) P(\mathbf{X}, t). \tag{2}$$

We highlight a gentle abuse of notation: \mathbf{X} denotes species labels in the reaction network description (1), but is an independent variable in \mathbb{Z}^d in the law in (2). Only the latter will be used henceforth as we focus on the law of the stochastic process.

We assume that the species typically occur in large numbers, and seek to approximate the discrete description on the continuum. Defining $\mathbf{e}_i := \mathbf{d}_i - \mathbf{c}_i$, the Fokker–Planck

approximation [17, Ch. 5] is given by

$$\frac{\partial}{\partial t} p(\mathbf{x}, t) = \sum_{i \in \mathcal{I}} -\mathbf{e}_i^T \nabla [r_i(\mathbf{x}, t) p(\mathbf{x}, t)] + \frac{1}{2\Omega} \mathbf{e}_i^T \vec{\mathbf{H}}(r_i(\mathbf{x}, t) p(\mathbf{x}, t)) \mathbf{e}_i, \tag{3}$$

where $\Omega \gg 1$ is the system size, $\mathbf{x} = \mathbf{X}/\Omega$, $r_i = R_i/\Omega$ and $\vec{\mathbf{H}}$ denotes the Hessian. The probability mass $P(\mathbf{X}, t)$ is approximated by $p(\mathbf{x}, t)/\Omega^d$. Eq (3) is a particular form of the more general Fokker–Planck equation [18].

We may extend the class of reactions by modeling non-local state changes to describe, for example, species production in “bursts”. We emphasise that we use the terms “non-local” and “local” with respect to the state; spatial location is not accounted for in this study. Such a process, indexed by $j \in \mathcal{J}$, may be associated with a kernel $\mathcal{B}_j(y, t)$, typically describing the probability density of a state change of size $y \in \mathbb{R}$ at time t , and the rate $f_j(\mathbf{x}, t)$, describing the rate of such changes occurring from state \mathbf{x} at time t . The non-local model then takes the form

$$\begin{aligned} \frac{\partial}{\partial t} p(\mathbf{x}, t) = & \sum_{i \in \mathcal{I}} -\mathbf{e}_i^T \nabla [r_i(\mathbf{x}, t) p(\mathbf{x}, t)] + \frac{1}{2\Omega} \mathbf{e}_i^T \vec{\mathbf{H}}(r_i(\mathbf{x}, t) p(\mathbf{x}, t)) \mathbf{e}_i \\ & + \sum_{j \in \mathcal{J}} -f_j(\mathbf{x}, t) p(\mathbf{x}, t) + \|\mathbf{e}_j\| \int_{\mathbf{x}-z\mathbf{e}_j \in \mathbb{R}_+^d} f_j(\mathbf{x}-z\mathbf{e}_j, t) p(\mathbf{x}-z\mathbf{e}_j, t) \mathcal{B}_j(z \|\mathbf{e}_j\|, t) dz. \end{aligned} \tag{4}$$

This is also closely related to droplet breakup and coalescence, where the equation is called the Smoluchowski coagulation equation, and has been used extensively in models from aerosols to cosmological structure formation [19]. While our focus is on the breakup process rather than coalescence (or aggregation), this framework may be extended to include such terms.

Lastly, we introduce discrete states indexed by $k \in \mathcal{K}$, whereby the probability is distributed among the family of densities $\{p_k\}_{k \in \mathcal{K}}$. There is a flux of probability from state $k \in \mathcal{K}$ to state $\ell \in \mathcal{K}$ at a rate of $g_{k\ell}(\mathbf{x}, t)$, that may depend on the continuum state and time. Importantly, all other rates (reaction and non-local) may depend on the discrete state k , and are subscripted to reflect this. In this way, we avoid having to consider reactions and non-local processes, indexed by \mathcal{I} and \mathcal{J} , respectively, that depend on the discrete state $k \in \mathcal{K}$. Instead, we may consider \mathcal{I} and \mathcal{J} to include the union of all local and non-local interactions, and set the rates to zero for the appropriate $k \in \mathcal{K}$. The final form of the system is

$$\begin{aligned} \left\{ \frac{\partial}{\partial t} p_k(\mathbf{x}, t) = \sum_{i \in \mathcal{I}} -\mathbf{e}_i^T \nabla [r_{ik}(\mathbf{x}, t) p_k(\mathbf{x}, t)] + \frac{1}{2\Omega} \mathbf{e}_i^T \vec{\mathbf{H}}(r_{ik}(\mathbf{x}, t) p_k(\mathbf{x}, t)) \mathbf{e}_i \right. \\ + \sum_{j \in \mathcal{J}} -f_{jk}(\mathbf{x}, t) p_k(\mathbf{x}, t) + \|\mathbf{e}_j\| \int_{\mathbf{x}-z\mathbf{e}_j \in \mathbb{R}_+^d} f_{jk}(\mathbf{x}-z\mathbf{e}_j, t) p_k(\mathbf{x}-z\mathbf{e}_j, t) \mathcal{B}_{jk}(z \|\mathbf{e}_j\|) dz \\ \left. + \sum_{\ell \in \mathcal{K}} -g_{k\ell}(\mathbf{x}, t) p_k(\mathbf{x}, t) + g_{\ell k}(\mathbf{x}, t) p_\ell(\mathbf{x}, t) \right\}_{k \in \mathcal{K}}. \end{aligned} \tag{5}$$

We henceforth call system (5) the Augmented Chemical Master Equation (ACME). The ACME framework describes a hybrid model that generalises the continuum Fokker–Planck equation (the first line of (5)) and the discrete CME (the last line of (5)) by encapsulating both descriptions simultaneously alongside non-local transitions of the continuum state (the second line of (5)). Conceptually, the continuum components were motivated by the need for efficient approximations of the CME for large state spaces, and they typically preserve the underlying

master equation structure. Therefore, the simplest conception of the ACME framework remains fundamentally classical; describing the evolution of the probability of a single cell to have a particular internal state. This begs the question: in what sense, beyond merely piecing together the hybrid formulation, is the ACME augmenting the CME? We answer this question below by discussing two technical assumptions made in this section that may be relaxed. The relaxed assumptions violate the classical framework, but in so doing allow the ACME framework to reach beyond the CME in physically meaningful ways, which we demonstrate with concrete example applications.

The augmentation and its interpretation

It is instructive to first contemplate what the discrete states model. Retaining a discrete description is crucial for molecules not expected to be present in large numbers. In this case, we consider d to denote the number of continuum species. The discrete states may also be used to describe abstract system states in which different behaviour governs the reaction network. For example, light-induced recombination [20] allows cells to alter their genetic structure, thereby changing the reaction network behaviour. This transition to a different governing dynamic may be encapsulated by transition between discrete states. A third use-case arises when discrete transition rates are negative, whereby some discrete states no longer correspond to a physical system configuration: the density may become negative and thus no longer corresponds to a probability measure. We will show that these phantom states enable a broader class of auxiliary population-level dynamics to be encapsulated as the model retains a physically meaningful interpretation.

The non-local contributions introduced above are inspired by long-range changes of cell state, such as production in large bursts. We will demonstrate that particular jump kernels that are not probability distributions violate classical conservation properties. Nevertheless, these can simultaneously allow us to capture certain auxiliary population-level processes and thus the ACME framework remains a useful population description.

The ACME framework augments the classical framework in violating these two classical assumptions—that the discrete transition rates $g_{k\ell}$ are non-negative, and that the jump kernels \mathcal{B}_{jk} are probability densities. The result is that the density described by the ACME framework may no longer describe the probability distribution of an underlying stochastic process of a single cell, since the density may have negative components or be non-conservative (and thus not sum to unity). Instead, it captures the evolution of the prevalence of different states within a population. In other words, the interpretation of a probability density is replaced by that of a population density. Such populations are governed not only by the changes of internal state of single cells, but additionally by auxiliary processes, those that act to change the composition of the population. In this way, the positive components may still describe a physically valid population density and they need not be conservative. This paradigm shift is illustrated in Fig 1: the vertical axes on the left panel (classical CME conception) describe probability densities while on the right (ACME framework) they describe number densities.

It is worth classifying which auxiliary processes are incorporated in this formulation. Unary auxiliary processes—those originating from a single individual—are incorporated. This includes the duplication or removal of individuals, e.g. due to growth or death of a single cell, as well as the redistribution of a single cell, e.g. cell division. Nullary processes—those not originating from the existing population, such as external introduction of new individuals into the population—may also be incorporated via abstract states. Binary and higher-order processes, such as multiple cells fusing, are not presently included.

Equipped with the ACME framework of system (5), we proceed to outline the implementation of a numerical solver capable of solving system (5) in the general case, which we call the Flips solver.

The Flips library

In this section we briefly survey numerical solvers in the chemical kinetics literature to establish the context of the numerical scheme we employ. We stress that all of these solutions remain grounded within the classical CME framework, and thus fall short from the population-level perspective.

Numerical simulation of stochastic reaction networks has an extensive and richly developed literature. Rather than an exhaustive survey, we provide an outline of the approaches and their motivations and relative merits. Reaction networks are commonly modeled as a continuous-time Markov chain (CTMC) whose state distribution is governed by the CME. When the possible number of states becomes large, solving the master equation becomes prohibitively challenging. An alternative to solving the CME is to generate exact sample trajectories of the stochastic process via the Stochastic Simulation Algorithm (SSA) [21, 22]. Stochastic simulation of the underlying CTMC is equivalent to the master equation in the sense that each simulated trajectory is a sample from the distribution given by the solution of the master equation. In other words, the underlying stochastic process is identical, the CME captures the full distribution while the SSA merely samples from it. For accurate resolution of the state distribution, large number of sample paths are often required, which has led to a rich set of refinements [23] and approximations [24]. Hybrid approaches that attempt to efficiently capture dynamics at different state [25] or time [26] scales remain a topic of active research. Ultimately, producing a sufficient number of trajectories to accurately resolve the distribution continues to be a challenge.

In an attempt to avoid stochastic sampling but accelerate solution of the full CME, several model reduction approaches have been considered, such as the Finite State Projection method [27], leveraging timescale separation [28], lumping (or aggregating) states [29, and references therein], as well as various spectral approximations [30, 31]. Hybrid combinations of these techniques have also been composed [32, and references therein].

One prevalent model-reduction technique is the class of continuum approximations of the CME, in particular, the Fokker–Planck equation [17, 18]. The underlying assumption is that, when only a small number of molecules need to be simulated the computational complexity is manageable, but when species are expected to have large copy numbers, a continuum approximation is appropriate. Just as the Fokker–Planck equation is the continuum analogue of the CME, the Langevin equation is the continuum analogue of the continuous-time Markov chain. That is, the Langevin equation describes the underlying stochastic process whose distribution is captured by the Fokker–Planck equation.

Various numerical schemes for solving the Fokker–Planck approximation of the CME appear in the literature [33, 34]. Just as with the trajectory sampling, many studies adopt hybrid approaches. In Haseltine and Rawlings [35] and Salis and Kaznessis [36] reactions are partitioned into fast and slow, and sample paths are generated by the Langevin equation or variants of the SSA discussed above. However, the Monte Carlo sampling requires a large number of samples just as for the aforementioned discrete sampling. In Safta et al. [37], the dynamics for each species are partitioned depending on molecule count. A discrete description is implemented for small molecule count (where the continuum assumption breaks down), while the continuum approximation is adopted for larger numbers of molecules. In Sjöberg [38], species are separated into those that are expected to have only small copy numbers for which the discrete description is retained, and those expected to exist in large number where

the continuum description is accurate, although non-local contributions in the continuum species are not incorporated.

As we will discuss, in the extended setting of the ACME framework, the quantity p governed by (5) does not always represent a probability distribution. Therefore, the direct link to the underlying stochastic process is severed, rendering stochastic sampling techniques not directly applicable. We instead develop a solver based on the system of differential equations (5), which we call the Flips solver (loosely acronymous of the Fokker–Planck system it is designed to solve). Species are separated into those whose description remains discrete, and those whose description are approximated on the continuum, in which we allow non-local dynamics. Most substantially, the solver is built to capture the augmentation of the classical chemical reaction structure, capable of encapsulating auxiliary population-level processes and thus appreciably expanding the model scope, as we demonstrate in the sequel.

To solve system (5) numerically, we begin by discretising the state space. Despite having d continuum-state dimensions, each continuum reaction $i \in \mathcal{I}$ in the model (5) acts to advect and diffuse probability density only in the \mathbf{e}_i direction. Since $\mathbf{e}_i \in \mathbb{Z}^d$, there is a natural discretisation for an arbitrary network on uniform grids that preserves this one-dimensional structure. Similarly, each burst production changes the state by $\mathbf{e}_j \in \mathbb{Z}^d$ for $j \in \mathcal{J}$, and thus the non-local integral term describes an exchange of probability in the \mathbf{e}_j direction, which is similarly preserved on the uniform grid. This observation means that the differential and integral operators act in one-dimension even in a high-dimensional system with multiple species thereby inducing no spurious diffusion in orthogonal directions.

Our approach is to adopt a uniform grid in state space, where we truncate the space at some sufficiently large boundary where we expect only negligible probability density to accumulate. For such a numerical scheme to yield computational advantage over solving the underlying discrete system (2), the state step Δx must satisfy two competing constraints. On the one hand, $\Delta x \ll 1$ is to be sufficiently small so that the scheme is accurate, while on the other hand, it must be nominally larger than a single discrete molecule, that is, $\Delta x > 1/\Omega$, where Ω is a typical number of molecules (see the scalings below (3)). In practice, the molecule copy number is typically $\Omega \gtrsim 10^4$, while we will consider Δx on the order of $\mathcal{O}(10^{-2})$, thus both of these constraints are well met.

The conservative finite-volume numerical scheme, based on Kurganov and Tadmor [39], is presented and analysed in Section 1 in S1 Appendix. The Flips solver is made open source [16] to encourage its use and development. We emphasise that the Flips solver has been written with two distinct types of user in mind.

First, the practitioner principally interested in simulation results with less focus on the numerical implementation details. With this in mind, the code is distributed as a package in the open-source and widely distributed Python programming language. The interface allows for a direct encoding of the reaction network structure without reference to the underlying mathematical abstraction, as depicted in Fig 2.

Second, for the user interested in more control of the numerical implementation, the interface provides options to calibrate the state discretisation and tweak the flux limiter, as well as adjust the time stepping method and its order. Furthermore, the code has been written in a modular way that allows the more advanced user to retain the framework structure and implement a completely custom discretisation with minimal effort. For example, an implementation of a first-order finite difference scheme is included in the code, and achieved in little more than 20 lines of code. Building upon this, second- and third-order extensions are achieved in little more than 10 lines of code each.

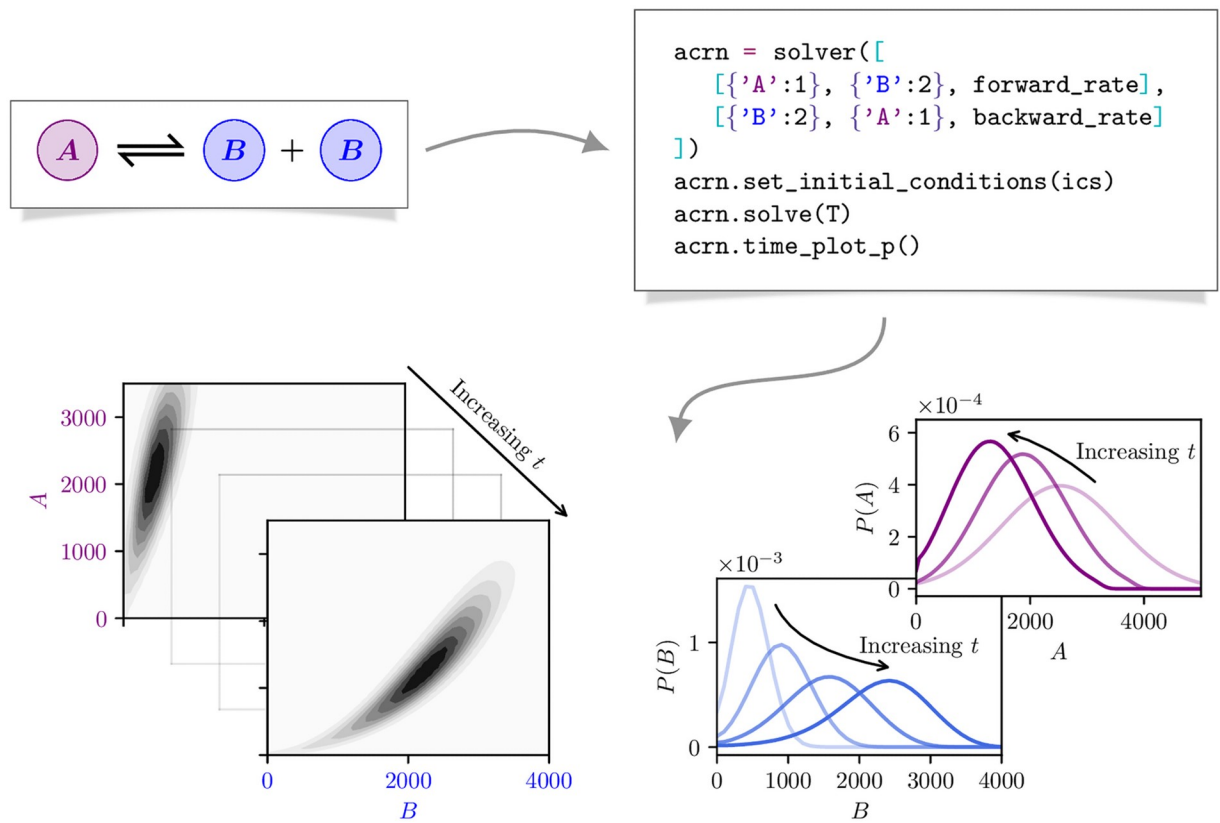


Fig 2. Schematic depiction of the Flips solver pipeline. Beginning with a (possibly augmented) chemical reaction network, we encode the reactions, set the initial distribution, and instruct the solver to determine the state distribution up to a terminal time $t = T$ or terminal event. We may then plot the evolution of the state distribution, and its marginals, over time. For the sake of clarity, at this point we illustrate a classical chemical reaction network to highlight the simple one-to-one encoding of the reaction network, and the flavour of results the solver generates. The abstract and non-physical augmentation that distinguishes the ACME framework is demonstrated on more substantial networks in the Applications section.

<https://doi.org/10.1371/journal.pcbi.1009214.g002>


We proceed to describe case studies in which we cast problems of interest from the literature in the ACME framework and solve them using the Flips software.

Applications

Our aim is to combine a representation of highly non-trivial systems involving chemical reactions coupled to auxiliary processes to demonstrate how the hybrid structure of the ACME formulation of system (5) can be exploited to tackle an extensive class of problems. We benchmark the Flips solver by comparing simulations of each example problem with analytical results.

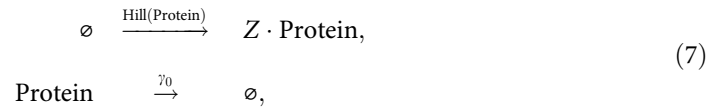
First-passage times in self-regulated gene expression

Our first example is a model of regulated protein production as described in Friedman et al. [13] and Lin and Doering [40]. The reaction network may be written as



The system comprises the transcription of mRNA molecules, the translation of protein molecules, and the degradation of both species. The dynamics are self-regulated because the mRNA transcription is regulated by the quantity of protein present, inducing a feedback loop. The notation $\text{Hill}(\cdot)$ represents that the feedback is modulated by a Hill function.

As described in Friedman et al. [13] and Lin and Doering [40], the characteristic timescale of the mRNA dynamics $1/\gamma$ is often exceeded by the protein lifetime timescale $1/\gamma_0$. In the limit as $\gamma/\gamma_0 \rightarrow \infty$, the mRNA acts instantaneously to produce a burst of proteins, and the reaction network reduces to



where production in bursts is represented by the random variable Z , such that bursts occur at a rate $\text{Hill}(\text{Protein})$ with size distributed by $Z \sim \text{Geo}(1/B)$.

Thus we have two models from the literature describing a single genetic network. The finite mRNA lifetime model (6) is more general but less easy to analyse, while the infinitely fast mRNA model (7) is less general but simpler. Our aim is to use these two models as a first demonstration of how both descriptions (6) and (7) may be encapsulated and studied within the ACME framework of system (5). We then turn our attention to an insightful first-passage time problem studied for the second system (7). We describe how this problem too may be simulated and studied using the Flips solver, and significantly extend results previously reported.

Since the protein copy number is typically large, the proteins are taken on the continuum: represented by x where $d = 1$. Both the studies of Friedman et al. [13] and Lin and Doering [40] neglect the stochastic noise expressed by the diffusive term of order $1/\Omega$, and thus, for the sake of comparison, we too consider the noiseless limiting case as $\Omega \rightarrow \infty$.

In system (6), the mRNA will remain discrete, and the system may be written as

$$\begin{aligned} \frac{\partial}{\partial t} p_k(x, t) &= -\frac{\partial}{\partial x} [(k\gamma B - \gamma_0 x) p_k(x, t)] - (H(x) + \gamma k) p_k(x, t) \\ &+ \gamma(k + 1) p_{k+1}(x, t) + \mathbb{1}_{k>0} H(x) p_{k-1}(x, t), \end{aligned} \tag{8}$$

where $\mathbb{1}_{k>0}$ denotes the indicator function, and H denotes the Hill function given by

$$H(x) = r_0 + r_1 \frac{x^n}{1 + x^n}. \tag{9}$$

Each discrete state k represents the number of mRNA molecules present. In practice, we truncate the discrete mRNA state at some finite maximum $0 \leq k \leq K$.

System (7) may be written as

$$\frac{\partial}{\partial t} p(x, t) = \frac{\partial}{\partial x} [\gamma_0 x p(x, t)] - H(x) p(x) + \int_0^x H(x - z) p(x - z, t) \frac{e^{-z/b}}{b} dz, \tag{10}$$

for the mean continuum-scaled burst size $b = B/\Omega$. The exponential burst kernel is a standard form for the burst-production jump density [13], and may be derived as the continuum scaling limit of the (discrete) geometrically distributed burst sizes [40].

Both reaction networks (6) and (7) fall within the ACME framework (5), and may then be solved numerically by application to the forms (8) and (10), respectively. To demonstrate the simplicity of using the Flips software, in Section 2 in S1 Appendix we present the code required to set up and simulate both reaction networks. Most of the code is simply a one-to-one translation of the discrete reaction network (each reaction is described by its reactants, products and

rate). Intuitive instructions tell the solver to ignore diffusion, fix uniform initial conditions, solve over a specified time interval, and plot the distribution. This makes the software accessible to the end-user without the need to deal with the technicalities of the continuum approximation or its analysis. At the same time, the software allows fine-grained control of the discretisation of the state space and the differential operators, as well as the time stepping method and associated properties.

For the sake of comparing to results in Lin and Doering [40], we adopt the same parameter values, namely

$$r_0 = 2, \quad r_1 = 10, \quad n = 4, \quad B = 40, \quad \gamma_0 = 1, \quad \Omega = 200. \quad (11)$$

First, we explore the stationary distributions of these systems by comparing them to benchmark solutions. In the case of finite mRNA lifetime, the CME (with discrete state-space) associated with (6) is solved using the Flips solver to provide the stationary distribution of the discrete model. For infinitely fast-lived mRNA, the stationary distribution of the continuum (10), which we denote $p_\infty(x)$, is given in Lin and Doering [40] by

$$p_\infty(x) = \frac{c}{\gamma_0} e^{-x/b} x^{r_0/\gamma_0 - 1} (1 + x^n)^{r_1/(m_0)}, \quad (12)$$

where r_0 , r_1 and n are the Hill function parameters, and c is a normalisation constant. These two distributions (black curves) are a useful comparison for the Flips solver (coloured curves) as illustrated in Fig 3.

In Fig 3A we plot the stationary distributions of (8) as given by the Flips solver (solid curves) and the full CME (dotted curves) for finite, increasing values of the mRNA production/degradation rate γ . The limiting case of $\gamma = \infty$ is plotted by simulation of (10) and compared to the exact continuum solution (dashed curve) given in (12). We see good agreement even though the Flips solver has lumped every ten protein molecules into a single discrete




Fig 3. Comparison of self-regulated gene expression models. The stationary distributions of models (6) and (7) are compared while varying the typical mRNA production/degradation rate γ . (A) For finite γ , the solid curves show steady-state solutions of the continuum ACME form (8) compared to dotted curves which are the steady-state solution of the associated discrete CME system, retaining $K = 30$ discrete mRNA states in both cases. For the limiting case $\gamma = \infty$, the solid curve depicts the steady-state solution of the continuum ACME form (10) compared to the dashed curve which is the exact continuum solution p_∞ in (12). The distributions are the marginal distributions of protein molecules, summing over the mRNA distribution. All continuum simulations used $\Delta x = 0.05$, corresponding to lumping 10 discrete states together. (B) Convergence of the steady-state solution of the continuum ACME form (8) to the exact continuum solution p_∞ in (12) in different L^n norms. The triangle shows the $1/\gamma$ convergence rate. We employed a discretisation of $\Delta x = 0.025$, corresponding to lumping 5 discrete states together.

<https://doi.org/10.1371/journal.pcbi.1009214.g003>

volume. Importantly, results in Lin and Doering [40] were obtained only in the limiting case of $\gamma/\gamma_0 \rightarrow \infty$, and thus intermediate values of γ/γ_0 were beyond reach (see [40, Appendix F]). By retaining the finite lifetime of the mRNA molecules in formulation (8), the Flips software is able to match the CME over the entire range of relevant γ/γ_0 values, and we observe the convergence of the finite mRNA lifetime scheme towards the infinitely fast-lived scheme. In Fig 3B we plot the norm of the discrepancy between the stationary distributions of these two schemes for increasing mRNA production/degradation rates γ . We observe a linear convergence of order $1/\gamma$, which provides an error estimate not previously reported.

Having demonstrated the ACME framework’s utility in studying model accuracy by comparing stationary distributions, we proceed to study an important transient phenomenon. One striking feature of the regulated protein production is the bimodality of the stationary distribution: the nonlinear bursting supports both a low mode and a high mode where protein production and degradation are in balance. This stationary distribution is a dynamic equilibrium: in any single cell, the stochastic expression will be continually changing the number of protein molecules. The switching between these two modes can shed light on biological event timing. Choosing an arbitrary boundary of $x_c = 0.825$ proteins separating the high and low protein modes, we may ask: given any initial number of proteins x , how long will it take to cross the boundary $x = x_c$, which we interpret to be a transition into the other mode. The mean switching time is given explicitly in Lin and Doering [40] by

$$T_{\text{low} \rightarrow \text{high}}(x) = \frac{1 - \gamma_0 x_c V(x_c) e^{-M(x_c)}}{H(x_c)} + \int_{x_c}^x e^{-M(y)} V(y) dy, \tag{13a}$$

$$T_{\text{high} \rightarrow \text{low}}(x) = \int_0^x e^{-M(y)} [V(y) - V(\infty)] dy, \tag{13b}$$

where

$$M(x) = \log \frac{x}{H(x)} - \frac{x}{b} + \int_0^x \frac{H(y)}{\gamma_0 y} dy, \quad V(x) = - \int_0^x \left(\frac{1}{b\gamma_0 y} + \frac{1}{\gamma_0 y H(y)} \frac{dH}{dy} \right) e^{M(y)} dy. \tag{13c}$$

We would like to go beyond the mean switching time to study the full switching time distribution. Strategically leveraging the hybrid structure of the ACME framework, we may augment system (10) with an additional abstract state, so that the system is described by the densities p_0 and p_1 . The original dynamics (10) are simulated on p_0 , in addition to a transition to the absorbing state p_1 at a rate α , conditioned on whether the state has reached the opposite region. In other words, for an initial condition in the high-mode region $x > x_c$, the absorption rate is $\alpha \mathbb{1}_{x < x_c}$, while for an initial condition in the low-mode region $x < x_c$ the absorption rate is $\alpha \mathbb{1}_{x > x_c}$. In the limit as $\alpha \rightarrow \infty$, this models the opposite region as an immediately absorbing state. In practice, a sufficiently large value of α provides a good approximation. The switching time density is simply the probability mass that crosses the critical state x_c for the first time. With the solution of this augmented system, the switching time density is approximated by the instantaneous absorption rate.

In Fig 4 we plot the cumulative switching time density, along with the mean, median and mode. The exact solution (13) is plotted as a black dashed curve, with which the numerical solution closely agrees. Strikingly, the switching time distribution is highly skewed: for an initial number of proteins $x \in (0.35, 2.25)$, the mean is at least an order of magnitude larger than the mode and is uniformly larger than the median. It is important for both the experimentalist and theoretician to keep this heavy tail in mind: with a small number of samples, switching times near the mode (and not the mean) are likely going to dominate the sample.

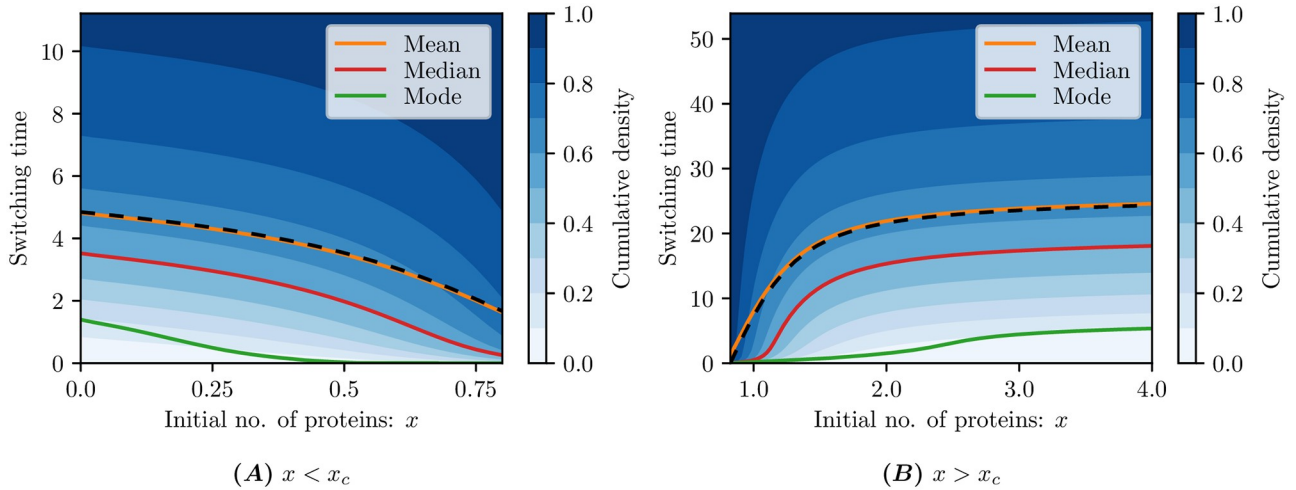


Fig 4. Switching time distributions. For different initial conditions x on either side of the critical $x_c = 0.825$ with an absorption rate $\alpha = 1000$ and spatial discretisation of $\Delta x = 0.025$ (i.e. lumping 5 states together) up to time $t = 150$. The exact mean switching time (13), is plotted as a black dashed curve.

<https://doi.org/10.1371/journal.pcbi.1009214.g004>

We plot a selection of the switching time densities in Fig 5, for initial conditions in both the low- and high-mode regions, to demonstrate just how far the mean is from the mode near the critical number of proteins x_c . The heavy tails of these distributions are illustrated in Fig 5B, where we distinguish different large-time asymptotic behaviour for initial conditions below/above the critical boundary x_c . The exponential decay in the switching time density suggests that the profile p converges to a separable solution of the form $p(x, t) \sim T(t)Y(x)$ for $T(t) = e^{-\lambda t}$ motivating further analytical study.

In the models of self-regulated gene expression studied in this paper, the Flips solver is able to quantify model accuracy and convergence rates. By augmenting an abstract discrete state to capture a first passage, the Flips solver is able to produce the full first-passage time distribution, revealing a remarkably rich structure of heavy-tailed distributions, distinguished by the critical




Fig 5. Extended switching time distributions. From the same data used to produce Fig 4 with the initial conditions $x \in \{0.7, 0.75, 0.8, 0.85, 0.9, 0.95\}$. (A) Zooming in to a small time interval resolves the modes, while (A inset) zooming out far enough to see the mean switching times shows the distribution skew. (B) The heavy tails are shown on a logarithmic scale for large times.

<https://doi.org/10.1371/journal.pcbi.1009214.g005>

boundary between the high and low modes of protein expression. We proceed to study a population-level auxiliary process that cannot be captured by a master equation, showing that careful manipulation of physical constraints make the problem tractable within the ACME framework.

Phenotypic selection

Many phenomena—such as changes to the ambient environment or the presence of media that affect metabolic or cell-cycle processes—can exert selection pressures. Importantly, such selection pressure may affect cells in a manner that depends on cell phenotype. For instance, recent studies provide an increasingly sharp understanding of how growth rates depend on gene expression [41]. This dependence motivates the development of models where an auxiliary process (capturing a selection pressure such as growth) is coupled to the single-cell gene expression kinetics. The ACME framework is capable of describing such population-level evolution, as we demonstrate in this second example.

We begin by considering the concentration of a particular protein that is heterogeneously distributed within a population due to stochastic gene expression (as in the previous case study). This is modeled as a birth–death process of a chemical species X within a single cell, governed by



The continuum approximation of process (14) takes the form [42]

$$\frac{\partial}{\partial t} p(x, t) = -\frac{\partial}{\partial x} [(\lambda(x) - \mu(x))p(x, t)] + \frac{1}{2\Omega} \frac{\partial^2}{\partial x^2} [(\lambda(x) + \mu(x))p(x, t)]. \tag{15}$$

We now seek to describe a growth process of a population of many such cells. We consider the protein concentration in daughter cells to be inherited from the mother cell. The growth G is to depend monotonically on protein concentration x (for example, when growth positively regulates cell-cycle stages and ultimately the rate of cell division), such that the population-level effect is preferential proliferation of cells with higher protein concentrations. This is modeled by adding a classical exponential growth term $G(x)p(x, t)$ on the right-hand side of (15). Unfortunately, the resulting equation no longer conserves the total mass of p , therefore, p is no longer a probability distribution. This may be salvaged by normalisation [9, 10] where the total mass inflation due to growth is uniformly removed, yielding the governing equation

$$\begin{aligned} \frac{\partial}{\partial t} p(x, t) &= -\frac{\partial}{\partial x} [(\lambda(x) - \mu(x))p(x, t)] + \frac{1}{2\Omega} \frac{\partial^2}{\partial x^2} [(\lambda(x) + \mu(x))p(x, t)] \\ &+ G(x)p(x, t) - \left(\int_0^\infty G(z)p(z, t) dz \right) p(x, t). \end{aligned} \tag{16}$$

By integrating over the state space x , the reader may verify that the total probability $\int_0^\infty p(x, t) dx$ is preserved in formulation (16).

It seems that only in the simplest of cases is equation (16) explicitly tractable. More worryingly, it seems that equation (16) is nonlinear and not of the general ACME form (5). We proceed to demonstrate how the nonlinear equation (16) is solved by a quantity satisfying an appropriately cast ACME equation of the generic form (5), thereby allowing us to use the Flips solver to tackle this population-level problem.

First, consider q , the solution of (16) without the nonlinear normalisation, namely

$$\frac{\partial}{\partial t} q(x, t) = -\frac{\partial}{\partial x} [(\lambda(x) - \mu(x))q(x, t)] + \frac{1}{2\Omega} \frac{\partial^2}{\partial x^2} [(\lambda(x) + \mu(x))q(x, t)] + G(x)q(x, t). \quad (17)$$

The normalised quantity $p = q / \int_0^\infty q(x, t) dx$ then satisfies the original Eq (16) and thus we have reduced the nonlinear problem (16) to the linear Eq (17).

Second, note that the linear Eq (17) may be cast in the general ACME form (5) by considering two discrete states ($k = 0, 1$): the first ($k = 0$) representing the density of the chemical species, and the second ($k = 1$) a phantom state to which a transition occurs from $k = 0$ with a propensity that is the negative of the growth rate, namely,

$$\begin{aligned} \frac{\partial}{\partial t} p_0(x, t) &= -\frac{\partial}{\partial x} [(\lambda(x) - \mu(x))p_0(x, t)] + \frac{1}{2\Omega} \frac{\partial^2}{\partial x^2} [(\lambda(x) + \mu(x))p_0(x, t)] \\ &\quad - [-G(x)]p_0(x, t), \end{aligned} \quad (18)$$

$$\frac{\partial}{\partial t} p_1(x, t) = [-G(x)]p_0(x, t). \quad (19)$$

Ultimately, the ACME form of (18) and (19) allows us to use the Flips solver to tackle the nonlinear population-level Eq (16).

The quantity p_0 is equivalent to q , except posed in a closed ACME form in (18) and (19), and thus is to be understood as a population density (as opposed to a probability density). The initial probability mass is to be distributed in p_0 , while $p_1(x, 0) = 0$. The new state described by p_1 , is a phantom state, in that it does not represent a physical system configuration. To see this, consider, for example, the case of a positive growth rate $G(x) > 0$ then $p_1(x, t) \leq 0$ for all time t , and therefore p_1 has no interpretation as a probability or population density. However, this non-physicality is of no concern: p_0 solves (17), and we may safely ignore the value of the phantom state p_1 .

There are other ways of exploiting the observation that we may split the solution of (16) into a non-normalised form followed by a normalisation. For example, we could implement operator splitting by solving the non-normalised linear Eq (17) and adding a normalisation at each time step. The advantage of this approach would be that no additional discrete state need be introduced, and the numerical quantities remain $\mathcal{O}(1)$ without risk of numerical overflow at large times. When a smaller state space, or large-time simulations are crucial, this option should be considered. However, the principal advantage of our approach is that it does not require bespoke changes to the ACME framework or solver, and thus we pursue this approach for simplicity. Moreover, that the quantity p_0 is a population density allows a direct quantification of the overall growth rate, as we now demonstrate.

For the purpose of numerical simulations, we take the underlying birth–death process studied in Lunz [42], namely

$$\lambda(x) = \Lambda x(1 - x), \quad \mu(x) = x. \quad (20)$$

We highlight that, for $\Lambda > 1$, there is a critical point x_c for which $\lambda(x_c) = \mu(x_c)$. On $x \in (0, x_c)$, birth dominates death $\lambda(x) > \mu(x)$, and the state increases towards x_c , while for $x \in (x_c, 1)$, the opposite is true, and $\lambda(x) < \mu(x)$ and deaths dominate births driving the state down to x_c . In the vicinity of the critical point stochastic effects become important [42].

In Fig 6A we show the solution of (18) with an initial Gaussian distribution centred at $x = 0.1$. The density $p_0(x, t)$ drifts toward higher values of x and from $t \approx 5$ appears to have a stationary centre point after which time the growth is evident. In the inset of Fig 6A we plot




Fig 6. Chemical kinetics coupled to state-dependent growth. Simulations of (18) for $G(x) = gx$ and the birth and death functions given in (20). Initial conditions are $p_0(x, 0) = \exp(-10^3(x - 0.1)^2) \sqrt{10^3/\pi}$ and $p_1(x, 0) = 0$. **(A)** Density p_0 for growth factor $g = 0.5$ and birth factor $\Lambda = 3$ at times $t \in \{0, 0.5, 1, 1.5, 2, 3, 4, 5\}$; **(A inset)** normalised density $p = p_0/P_0$ with P_0 defined in (23) at times $t \in \{0, 0.5, 1, 1.5, 5\}$. The black dashed curve shows the asymptotic approximation of the limiting profile given by (S46) in Section 3 in **S1 Appendix**. **(B)** Growth rate $(dP_0/dt)/P_0$ for growth factor $g = 0.5$ and birth factor $\Lambda \in \{2, 2.5, 3, 3.5, 4\}$; **(B inset)** $\Lambda = 3$ and $g \in \{0.25, 0.5, 1, 2, 4\}$. The black dashed lines show the large time predictions ((S47) and (S50)). The state space discretisation was $\Delta x = 0.007$ (i.e. lumping 7 states together).

<https://doi.org/10.1371/journal.pcbi.1009214.g006>

the normalised density $p = p_0 / \int_0^\infty p_0(x, t) dx$, and find that, for large times, the profile remains largely unchanged.

Motivated by the observation that the normalised density p appears to converge for large times t , we seek a separable solution of Eq (17). The calculations are detailed in Section 3 in **S1 Appendix** where we deduce that the profile tends to

$$p_0(x, t) \sim T(t)Y(x), \tag{21a}$$

where

$$T(t) = e^{r_0 t}, \quad Y(x) = \frac{1}{\sqrt{a\Omega\pi}} e^{-(x-x_c)^2/(a\Omega)}, \tag{21b}$$

The limiting profile (21) is a Gaussian centred at $x = x_c$ of width $\mathcal{O}(1/\sqrt{a\Omega})$ growing at rate r_0 , which is given by

$$r_0 = G(x_c). \tag{22}$$

For $G(x) = gx$, as used in our simulations, this becomes $r_0 = g(1 - 1/\Lambda)$.

These calculations serve as a useful gauge for our numerical simulations as illustrated in **Fig 6**. In the inset of **Fig 6A** we show the limiting normalised density (21), and observe good agreement with the numerical simulation at $t = 5$. We define the average growth rate of the numerical solution by $(dP_0/dt)/P_0$, where P_0 denotes the total density in the $k = 0$ state, that is,

$$P_0 = \int_0^\infty p_0(x, t) dx. \tag{23}$$

In **Fig 6B** we plot the average growth rate, and the predicted limiting growth rate (22) as black dashed lines, to which the numerical simulations converge for a range of system parameters.

To recap, the augmented reaction network considered in this section is not strictly physical, as we have incorporated negative propensities that induce negative “probabilities”. Nevertheless, by carefully crafting such a non-physical network, we were able to couple the single-cell reaction kinetics with an auxiliary growth process while taming the associated nonlinearity. In fact, the numerical simulations motivated the analytical study, which in turn uncovered insightful analytical information regarding the modes and long-time behaviour of the system, and allowed us to validate the simulations. We now turn our attention to an auxiliary fragmentation process that demonstrates the extensive coverage of the ACME framework.

Divide and fragment

It can be crucially important to track the structure of a population, such as via the size, age, etc. of individual members. For example, in cell-cycle modeling, cell division may depend on cell size and cell-cycle stage [7, 43]. This motivates our final example, which does not stem from a classical chemical reaction network but a growth–fragmentation process, describing the evolution of a distribution of cells of different sizes subject to size increase (growth) and cell division (fragmentation). For an accessible introductory treatment, we refer the reader to Perthame [12, ch. 4]. In this class of models, we consider the independent state variable x to represent size, and we emphasise this distinction by denoting the number density $\rho(x, t)$, with the volume density of cells of size x given by $x\rho(x, t)$. It is instructive to have in mind the concrete example of cell growth and division, and we will adopt the corresponding nomenclature, while keeping in mind that the models may be applied significantly more broadly both in biology [11] and beyond [14, 19].

The number density ρ is governed by the growth–fragmentation equation

$$\frac{\partial}{\partial t}\rho(x, t) = -\frac{\partial}{\partial x}[g(x, t)\rho(x, t)] - B(x, t)\rho(x, t) + \int_x^\infty b(y, x, t)\rho(y, t) dy. \quad (24)$$

The first term on the right-hand side of (24) represents the growth at a rate $g(x, t)$. The final two terms represent the fragmentation. The first of these terms describes the reduction of cells of size x due to fragmentation, where $B(x, t)$ is the rate of fragmentation of cells of size x . The integral term adds all the cells $y > x$ that fragment into cells of size x , where $b(y, x, t)$ is the fragmentation rate (per unit size) of cells of size y that fragment into a piece of size x and another piece of size $y - x$.

A no-flux boundary condition is imposed at $x = 0$, with the flux given by $g(x, t)\rho(x, t)$, and a (sufficiently rapid) vanishing far-field condition as $x \rightarrow \infty$. The model is typically studied in the absence of diffusion.

It is natural to expect the model to possess three physical features. First, the dynamics of the total number of cells is not to depend explicitly on the growth rate g , since growth is exclusively responsible for increasing cell size but not number. Second, the change of volume is not to depend explicitly on fragmentation, via B or b , since fragmentation is exclusively responsible for changing cell numbers but not total volume. Third, the rate of fragmentation of cells of size y into cells of size x is equal to the rate of fragmentation into cells of size $y - x$ (for the case of binary fragmentation), since these events are equivalent. These features are guaranteed by ensuring that B and b satisfy certain constraints. Upon reflection, it is intuitive that B and b cannot be independent, as B quantifies the total fragmentation rate, while b represents the fragmentation rate (per unit size) into particular sizes. For the sake of completeness, in Section 4 in [S1 Appendix](#) we provide a comprehensive derivation of the algebraic conditions necessary to ensure physicality.

The salient point for our purposes is that the fragmentation can be guaranteed to respect the conservation properties described above by suitably relating B and b . In Section 4 in [S1 Appendix](#) we further show how the representation in equation (24) can be transformed to a rate and kernel, equivalent to the burst production process and conforming to the ACME structure of system (5). Thus, despite not being a chemical reaction network, the growth–fragmentation process is solvable with the Flips software.

We now seek a benchmark case with which we can compare numerical solutions. In Section 5 in [S1 Appendix](#) we combine results from Rooney et al. [44] and Cáceres et al. [45] to derive two closed-form solutions for two growth–fragmentation equations. In the first case, we use the time-invariant fragmentation kernel $b(y, x, t) = 2ay^{k-1}$ for positive constants a and k , with no growth, that is, $g \equiv 0$. We may thus derive the form $B(x, t) = ax^k$ (see Section 4 in [S1 Appendix](#)), whereupon the equation takes the form

$$\frac{\partial}{\partial t} \rho(x, t) = -ax^k \rho(x, t) + \int_x^\infty 2ay^{k-1} \rho(y, t) dy, \tag{25}$$

and admits the self-similar solution

$$\rho(x, t) = c(at)^{2/k} e^{-atx^k}, \tag{26}$$

for an arbitrary constant c . A change of variables (see Section 5 in [S1 Appendix](#)) allows us to transform (26) to a solution of the equation

$$\frac{\partial}{\partial t} \rho(x, t) = -\frac{\partial}{\partial x} [x\rho(x, t)] - (akx^k + 1)\rho(x, t) + \int_x^\infty 2aky^{k-1} \rho(y, t) dy, \tag{27}$$

which is given by

$$\rho(x, t) = c e^{-2t} [a(e^{kt} - 1)]^{2/k} e^{-a(e^{kt}-1)(e^{-t}x)^k}. \tag{28}$$

Eq (27) describes growth linear in size, $g(x, t) = x$, and uniform decay.

Physically, the first example (25) describes a pure fragmentation process where the probability of a cell of size y being divided into cells of size x and $y - x$ is uniform for all $0 \leq x \leq y$. Thus, we expect the solution (26) to concentrate an increasing (and diverging) number density in a vicinity of the origin. Indeed, we see from (26) that, for any positive x , $\rho(x, t) \rightarrow 0$ while $\rho(0, t) \rightarrow \infty$ as $t \rightarrow \infty$. The volume density, $x\rho(x, t)$, while eventually vanishing at every point $x > 0$, is conserved in total, as physically expected.

The second example (27) incorporates growth and decay, acting in competition with the fragmentation. How these competing forces balance in asymptotically large time is not trivial. We find from (28) that, in the limit as $t \rightarrow \infty$, the density ρ converges to the profile

$$\rho(x, t) \sim ca^{2/k} e^{-ax^k}. \tag{29}$$

We model Eq (25) by considering a species X undergoing a non-local fragmentation with kernel $b(y, x) = ay^{k-1}$. This corresponds to a fragmentation rate of ax^k and a uniform fragmentation distribution (see Section 4 in [S1 Appendix](#)). Eq (27) requires two extra ingredients: production of X (in the sense of a reaction network, which captures growth in growth–fragmentation processes) and exponential transition to a new discrete state (decay). The species X is produced by a reaction at a rate proportional to its quantity



in the diffusion-free limit $\Omega \rightarrow \infty$. Adding a discrete state allows us to incorporate the decay

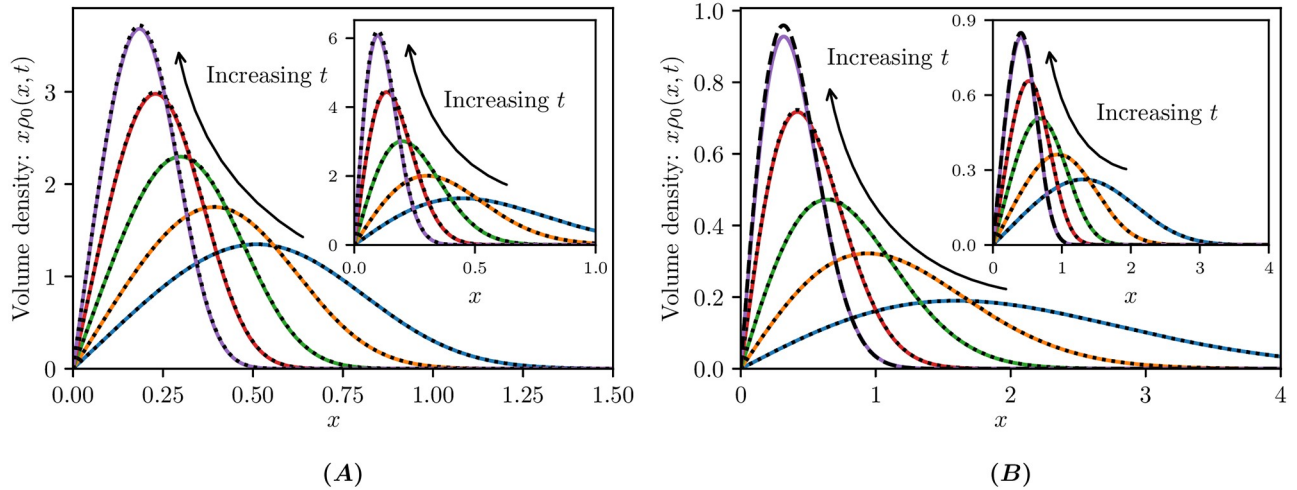


Fig 7. Growth-fragmentation processes. Analytical (black dotted curves) and numerical (coloured curves) solutions of growth-fragmentation equations at various times t . **(A)** Eq (25) with solution (26) at times $t - t_0 \in \{0, 0.6, 2, 5, 10\}$ where $t_0 = 0.5$ and $k = 3$; **(A inset)** $k = 2$; **(B)** Eq (27) with solution (28) at times $t - t_0 \in \{0.04, 0.12, 0.4, 2\}$ where $t_0 = 0.02$ and $k = 2$; **(B inset)** $t - t_0 \in \{0.017, 0.07, 0.2, 2\}$ where $t_0 = 0.01$ and $k = 3$. The black dashed curves in (B) show the limiting profile as $t \rightarrow \infty$ given by (29). Other parameters used are $a = 5$, $\Delta x = 0.005$. Initial conditions are given by the analytical solution at time $t = t_0$.

<https://doi.org/10.1371/journal.pcbi.1009214.g007>

term, ultimately arriving at the system

$$\frac{\partial}{\partial t} \rho_0(x, t) = -\frac{\partial}{\partial x} [x\rho_0(x, t)] + \int_x^\infty 2ax^k \rho_0(y, t) dy - ax^k \rho_0(x, t) - \rho_0(x, t), \quad (31)$$

$$\frac{\partial}{\partial t} \rho_1(x, t) = \rho_0(x, t). \quad (32)$$

For both examples, we use the initial condition given by the analytical solution at a nominal time $t = t_0$, that is, $\rho_0(x, 0) = \rho(x, t_0)$.

In Fig 7 we illustrate the analytical (black dotted curves) and numerical (coloured curves) solutions of Eqs (25) and (27) for $a = 5$ and $k = 2, 3$ at different times t . In Fig 7A, we simulate Eq (25) describing only fragmentation, which manifests as the volume density becoming concentrated around smaller sizes (near the origin). In contrast, Fig 7B shows simulations of Eq (27), where fragmentation acts in combination with decay to diminish the total volume density, and in competition with linear growth. The numerical solutions exhibit good agreement with the analytical solution. In Eq (27), the growth, decay, and fragmentation balance at asymptotically large time and the process converges toward the stationary distribution profile (29) as a black dashed curve in Fig 7B. We see that the long-time asymptotic behaviour is well approximated by the numerical solution already by time $t = 2$.

We have pointed out that both the dependent and independent quantities in the growth-fragmentation model (24) differ conceptually from their counterparts in the general reaction network model (5). Nonetheless, the model is of the same general form, and thus the solver is no less applicable. This is a powerful observation as it allows us to couple these two classes of models. As an example application, cell division has been reported to depend on gene expression pathways [46–48]. The ACME framework in tandem with the Flips software provides a generic framework in which such coupled models can be studied, alongside population-level effects and with view to transient dynamics (such as first-passage time problems).

Discussion

The aim of this work is to establish a unifying framework for auxiliary processes coupled to arbitrary internal reaction kinetics that extend beyond classical chemical master equation descriptions. The Augmented Chemical Master Equation (ACME) framework that we introduce, is grounded in the classical CME. Considering the prohibitive challenge of solving the CME for large copy numbers, we take the Fokker–Planck continuum approximation for species that are typically large, while retaining discrete states for the remaining species. For the continuum states, we capture non-local effects, such as production in bursts.

We may leverage the discrete states to describe abstract system configurations that are not simply cardinal quantities, for example, binding/unbinding events, cell-fate decisions, and so forth. We demonstrate this use in modeling transient dynamics of self-regulated gene expression to solve the first-passage time problem. The method provides the complete first-passage time distribution of the bimodal system, extending previous results in the literature.

Another use for the discrete states is to describe non-physical phantom system configurations. The power of this approach is that it allows the underlying equations to capture dynamics beyond the reach of the classical chemical master equation. We demonstrate this by simulating population-level phenotypic selection. Typically, a nonlinear problem, we prove that the nonlinear component may be resolved separately by normalising the solution of a master equation with an exponential growth term, thereby reducing the problem to a linear one. The growth is captured by transition to a phantom state at a (non-physical) negative rate. The phantom state has no physical interpretation (with negative “probabilities”), but its presence makes the extended population-level dynamics tractable within the unified ACME framework, and thus directly solvable with the Flips solver.

Finally, the general problem tackled within the ACME framework is applicable to equations not typically associated with the chemical master equation, such as growth–fragmentation models. We demonstrate that the Flips solver provides an accurate tool for this class of models, and point to physical phenomena described by growth–fragmentation which are coupled to processes classically modeled with the master equation. This hints at further benefits of coupling the master equation to such auxiliary processes, motivating future work in this direction.

Our results were obtained by casting the diverse collection of examples in the ACME framework and solving these using the Flips solver. The run-time of each simulation depends strongly on the time and space discretisation. The time-step is typically dependent on both the space step as well as the specific problem structure and parameters (see (S29) in Section 1 in [S1 Appendix](#)). The examples presented here typically took between a couple of seconds and a couple of minutes on a single core of a standard laptop computer.

The first example of self-regulated gene expression was classical in the sense that the ACME framework describes a probability density of the state of a single cell, and there were no population-level effects. In the latter two examples of phenotypic selection and growth–fragmentation we observe that the density is not always positive nor is its integral conserved, whereby the ACME framework corresponds to the evolution of the expectation of an underlying stochastic process [7]. It is worth crystallising precisely when this does and does not occur. Crucially, system (5) is conservative if all kernels \mathcal{B}_{jk} are probability densities. This can be seen by integrating over the state space (using the zero-flux boundary conditions) to find that the total mass $\sum_{k \in \mathcal{K}} \int_{\mathbb{R}_+^d} p_k(\mathbf{x}, t) \, d\mathbf{x}$ is time invariant. We thus deduce that, when this kernel condition is satisfied, the total mass is conserved. Then, assuming an initial unit mass, it might seem that a probability distribution is retained. However, as is demonstrated in the example of phenotypic

selection, if the rates are negative, we can end up with negative values of p_k , which cannot be probability densities. The underlying cause is the negative discrete transition rates (even if $p_k > 0$ for all k and all time). The classical Markov jump process is defined via the probability of a transition during an infinitesimal interval dt (given by the rates times dt). This is where the probabilistic interpretation breaks down: a negative rate corresponds to a negative transition probability, thus no direct stochastic interpretation is possible.

While bursty production satisfies the aforementioned kernel condition, the growth–fragmentation equation does not (since the kernel \mathcal{B} is two times a probability density, as explained in Section 4 in [S1 Appendix](#), stemming from the fact that the division of a single cell leads to two cells). This results in number densities that are positive but not conserved. We conclude that both the non-local contributions and the extended use of negative rates are augmentations of the classical CME and Fokker–Planck framework, able to capture changes in the population composition due to auxiliary processes.

The current formulation admits some limitations. The ACME framework does not describe higher-order auxiliary processes such as aggregation, which may be of biological relevance. The reasons for this are manifold. While it is certainly possible to extend the solver to handle higher-order dynamics, a significant computational cost would accompany this high-order structure. From a theoretical perspective, the population density in the present formulation is so interpreted as it arises from an expectation of an underlying random variable (see, e.g., [Supplementary Information 1 of Thomas \[7\]](#)). The extension to nonlinear dynamics is nontrivial as the nonlinearity under expectation introduces covariances that are unaccounted for in the current theory.

Even though the continuum setting allows for significant dimensionality reduction by coarse graining, the framework (and hence the Flips solver) is subject to the curse of dimensionality. Machine-learning-based approaches have enormous potential to make high dimensional systems tractable [\[49\]](#) and escape this limitation.

The examples illustrated in this work are intended to demonstrate some of the basic building blocks of the framework and solver. In this spirit, tractable examples were chosen to compare with previous work and analytical calculations. Nevertheless, we stress that more elaborate and sophisticated models may be tackled with this framework. We confined ourselves to no more than one auxiliary process in each example, however, this is not a fundamental limitation, and several such processes may be combined.

The hybrid network structure combining the continuum and discrete descriptions, alongside non-local dynamics, applied in versatile, often non-physical, ways, such as the use of negative reaction rates, significantly extends the scope of the classical master equation formulation. Our hope is that this augmented approach, and the accompanying software, puts a richer class of models and complex coupled processes firmly within reach.

Supporting information

S1 Appendix. Supplementary numerical, analytical, and computational descriptions. Section 1 includes details of the numerical scheme and its analysis. Section 2 provides code examples for the self-regulated genetic expression models described in this paper. In Section 3, an asymptotic analysis is performed for the selection model presented earlier. Section 4 introduces properties of the growth–fragmentation models studied here, including discretisation details. Section 5 develops two explicit solutions of growth–fragmentation equations to serve as benchmarks.

(PDF)

Author Contributions

Conceptualization: Davin Lunz, Gregory Batt, Jakob Ruess, J. Frédéric Bonnans.

Formal analysis: Davin Lunz, Jakob Ruess, J. Frédéric Bonnans.

Funding acquisition: Gregory Batt, Jakob Ruess, J. Frédéric Bonnans.

Investigation: Davin Lunz.

Methodology: Davin Lunz, Gregory Batt, Jakob Ruess, J. Frédéric Bonnans.

Project administration: Davin Lunz, Gregory Batt, Jakob Ruess, J. Frédéric Bonnans.

Software: Davin Lunz.

Supervision: Gregory Batt, Jakob Ruess, J. Frédéric Bonnans.

Validation: Davin Lunz.

Visualization: Davin Lunz.

Writing – original draft: Davin Lunz, Gregory Batt, Jakob Ruess, J. Frédéric Bonnans.

Writing – review & editing: Davin Lunz, Gregory Batt, Jakob Ruess, J. Frédéric Bonnans.

References

1. Gillespie DT. A rigorous derivation of the chemical master equation. *Physica A*. 1992; 188(1):404–425. [https://doi.org/10.1016/0378-4371\(92\)90283-V](https://doi.org/10.1016/0378-4371(92)90283-V)
2. Gillespie DT. The chemical Langevin equation. *J Chem Phys*. 2000; 113(1):297–306. <https://doi.org/10.1063/1.481811>
3. Weidlich W, Haag G. Concepts and models of a quantitative sociology: the dynamics of interacting populations. vol. 14. Springer Science & Business Media; 2012.
4. Fraser D, Kærn M. A chance at survival: gene expression noise and phenotypic diversification strategies. *Mol Microbiol*. 2009; 71(6):1333–1340. <https://doi.org/10.1111/j.1365-2958.2009.06605.x>
5. Eldar A, Elowitz MB. Functional roles for noise in genetic circuits. *Nature*. 2010; 467(7312):167–173. <https://doi.org/10.1038/nature09326>
6. Tonner PD, Darnell CL, Bushell FML, Lund PA, Schmid AK, Schmidler SC. A Bayesian non-parametric mixed-effects model of microbial growth curves. *PLOS Comput Biol*. 2020; 16(10):1–21. <https://doi.org/10.1371/journal.pcbi.1008366>
7. Thomas P. Intrinsic and extrinsic noise of gene expression in lineage trees. *Sci Rep*. 2019; 9(1):1–16. <https://doi.org/10.1038/s41598-018-35927-x>
8. Shu CC, Chatterjee A, Dunny G, Hu WS, Ramkrishna D. Bistability versus Bimodal Distributions in Gene Regulatory Processes from Population Balance. *PLOS Comput Biol*. 2011; 7(8):1–13. <https://doi.org/10.1371/journal.pcbi.1002140>
9. Mora T, Walczak AM. Effect of Phenotypic Selection on Stochastic Gene Expression. *J Phys Chem B*. 2013; 117(42):13194–13205. <https://doi.org/10.1021/jp403231f>
10. Sato K, Kaneko K. On the distribution of state values of reproducing cells. *Phys Biol*. 2006; 3(1):74–82. <https://doi.org/10.1088/1478-3975/3/1/008>
11. Duso L, Zechner C. Stochastic reaction networks in dynamic compartment populations. *Proc Natl Acad Sci USA*. 2020; <https://doi.org/10.1073/pnas.2003734117> PMID: 32868438
12. Perthame B. Transport equations in biology. Birkhäuser Basel; 2006.
13. Friedman N, Cai L, Xie XS. Linking Stochastic Dynamics to Population Distribution: An Analytical Framework of Gene Expression. *Phys Rev Lett*. 2006; 97:168302. <https://doi.org/10.1103/PhysRevLett.97.168302>
14. Bardet JB, Christen A, Guillin A, Malrieu F, Zitt PA. Total variation estimates for the TCP process. *Electron J Probab*. 2013; 18. <https://doi.org/10.1214/EJP.v18-1720>
15. Ghusinga KR, Dennehy JJ, Singh A. First-passage time approach to controlling noise in the timing of intracellular events. *Proc Natl Acad Sci USA*. 2017; 114(4):693–698. <https://doi.org/10.1073/pnas.1609012114>

16. Lunz D. Flips solver library. Version 0.0.2 [software]; 2021. [Available from]: <https://gitlab.inria.fr/dlunz/flips>.
17. Gardiner C. Stochastic Methods: A Handbook for the Natural and Social Sciences. 3rd ed. Springer Series in Synergetics. Springer Berlin Heidelberg; 2009.
18. Risken H. Fokker–Planck Equation. Berlin, Heidelberg: Springer Berlin Heidelberg; 1996.
19. Aldous DJ. Deterministic and Stochastic Models for Coalescence (Aggregation, Coagulation): a Review of the Mean-Field Theory for Probabilists. *Bernoulli*. 1997; 5:3–48. <https://doi.org/10.2307/3318611>
20. Kawano F, Okazaki R, Yazawa M, Sato M. A photoactivatable Cre–loxP recombination system for opto-genetic genome engineering. *Nat Chem Biol*. 2016; 12(12):1059–1064. <https://doi.org/10.1038/nchembio.2205>
21. Gillespie DT. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *J Comput Phys*. 1976; 22(4):403–434. [https://doi.org/10.1016/0021-9991\(76\)90041-3](https://doi.org/10.1016/0021-9991(76)90041-3)
22. Gillespie DT. Exact stochastic simulation of coupled chemical reactions. *J Phys Chem*. 1977; 81(25):2340–2361. <https://doi.org/10.1021/j100540a008>
23. Gibson MA, Bruck J. Efficient Exact Stochastic Simulation of Chemical Systems with Many Species and Many Channels. *J Phys Chem A*. 2000; 104(9):1876–1889. <https://doi.org/10.1021/jp993732q>
24. Gillespie DT. Approximate accelerated stochastic simulation of chemically reacting systems. *J Chem Phys*. 2001; 115(4):1716–1733. <https://doi.org/10.1063/1.1378322>
25. Hellander A, Lötstedt P. Hybrid method for the chemical master equation. *J Comput Phys*. 2007; 227(1):100–122. <https://doi.org/10.1016/j.jcp.2007.07.020>
26. Cao Y, Gillespie DT, Petzold LR. The slow-scale stochastic simulation algorithm. *J Chem Phys*. 2005; 122(1):014116. <https://doi.org/10.1063/1.1824902>
27. Munsky B, Khammash M. The finite state projection algorithm for the solution of the chemical master equation. *J Chem Phys*. 2006; 124(4):044104. <https://doi.org/10.1063/1.2145882>
28. Peleš S, Munsky B, Khammash M. Reduction and solution of the chemical master equation using time scale separation and finite state projection. *J Chem Phys*. 2006; 125(20):204104. <https://doi.org/10.1063/1.2397685>
29. Wang SW, Georgopoulos PG, Li G, Rabitz H. Condensing Complex Atmospheric Chemistry Mechanisms. 1. The Direct Constrained Approximate Lumping (DCAL) Method Applied to Alkane Photochemistry. *Environ Sci Technol*. 1998; 32(13):2018–2024. <https://doi.org/10.1021/es970967b>
30. Engblom S. Spectral approximation of solutions to the chemical master equation. *J Comput Appl Math*. 2009; 229(1):208–221. <https://doi.org/10.1016/j.cam.2008.10.029>
31. Jahnke T, Udrescu T. Solving chemical master equations by adaptive wavelet compression. *J Comput Phys*. 2010; 229(16):5724–5741. <https://doi.org/10.1016/j.jcp.2010.04.015>
32. Hepp B, Gupta A, Khammash M. Adaptive hybrid simulations for multiscale stochastic reaction networks. *J Chem Phys*. 2015; 142(3):034118. <https://doi.org/10.1063/1.4905196>
33. Sjöberg P, Lötstedt P, Elf J. Fokker–Planck approximation of the master equation in molecular biology. *Comput Vis Sci*. 2009; 12(1):37–50. <https://doi.org/10.1007/s00791-006-0045-6>
34. Ferm L, Lötstedt P, Sjöberg P. Conservative solution of the Fokker–Planck equation for stochastic chemical reactions. *BIT*. 2006; 46(1):61–83. <https://doi.org/10.1007/s10543-006-0082-z>
35. Haseltine EL, Rawlings JB. Approximate simulation of coupled fast and slow reactions for stochastic chemical kinetics. *J Chem Phys*. 2002; 117(15):6959–6969. <https://doi.org/10.1063/1.1505860>
36. Salis H, Kaznessis Y. Accurate hybrid stochastic simulation of a system of coupled chemical or biochemical reactions. *J Chem Phys*. 2005; 122(5):054103. <https://doi.org/10.1063/1.1835951>
37. Safta C, Sargsyan K, Debusschere B, Najm HN. Hybrid discrete/continuum algorithms for stochastic reaction networks. *J Comput Phys*. 2015; 281:177–198. <https://doi.org/10.1016/j.jcp.2014.10.026>
38. Sjöberg P. PDE and Monte Carlo Approaches to Solving the Master Equation Applied to Gene Regulation. Uppsala University; 2007. 2007–028.
39. Kurganov A, Tadmor E. New high-resolution central schemes for nonlinear conservation laws and convection–diffusion equations. *J Comput Phys*. 2000; 160(1):241–282. <https://doi.org/10.1006/jcph.2000.6459>
40. Lin YT, Doering CR. Gene expression dynamics with stochastic bursts: Construction and exact results for a coarse-grained model. *Phys Rev E*. 2016; 93:022409. <https://doi.org/10.1103/PhysRevE.93.022409>
41. Wytock TP, Motter AE. Predicting growth rate from gene expression. *Proc Natl Acad Sci USA*. 2019; 116(2):367–372. <https://doi.org/10.1073/pnas.1808080116>

42. Lunz D. On Continuum Approximations of Discrete-State Markov Processes of Large System Size. *Multiscale Modeling & Simulation*. 2021; 19(1):294–319. <https://doi.org/10.1137/20M1332293>
43. Wehrens M, Ershov D, Rozendaal R, Walker N, Schultz D, Kishony R, et al. Size Laws and Division Ring Dynamics in Filamentous *Escherichia coli* cells. *Current Biology*. 2018; 28(6):972–979.e5. <https://doi.org/10.1016/j.cub.2018.02.006> PMID: 29502951
44. Rooney CM, Griffiths IM, Brunner C, Potter J, Wood-Lee M, Please CP. Dynamics of particle chopping in blenders and food processors. *J Eng Math*. 2018; 112(1):119–135. <https://doi.org/10.1007/s10665-018-9968-4>
45. Cáceres MJ, Cañizo JA, Mischler S. Rate of convergence to an asymptotic profile for the self-similar fragmentation and growth-fragmentation equations. *J Math Pure Appl*. 2011; 96(4):334–362. <https://doi.org/10.1016/j.matpur.2011.01.003>
46. Brunner S, Sauer T, Carotta S, Cotten M, Saltik M, Wagner E. Cell cycle dependence of gene transfer by lipoplex, polyplex and recombinant adenovirus. *Gene Ther*. 2000; 7(5):401–407. <https://doi.org/10.1038/sj.gt.3301102>
47. Tseng WC, Haselton FR, Giorgio TD. Mitosis enhances transgene expression of plasmid delivered by cationic liposomes. *BBA-Gene Struct Expr*. 1999; 1445(1):53–64. [https://doi.org/10.1016/S0167-4781\(99\)00039-1](https://doi.org/10.1016/S0167-4781(99)00039-1)
48. Zopf CJ, Quinn K, Zeidman J, Maheshri N. Cell-Cycle Dependence of Transcription Dominates Noise in Gene Expression. *PLOS Comput Biol*. 2013; 9(7):1–12. <https://doi.org/10.1371/journal.pcbi.1003161>
49. Han J, Jentzen A, Weinan E. Solving high-dimensional partial differential equations using deep learning. *Proc Natl Acad Sci USA*. 2018; 115(34):8505–8510. <https://doi.org/10.1073/pnas.1718942115>