Scaling methods for accelerating kinetic Monte Carlo simulations of chemical reaction networks

Cite as: J. Chem. Phys. 150, 244101 (2019); doi: 10.1063/1.5096774

Submitted: 19 March 2019 • Accepted: 31 May 2019 •

Published Online: 24 June 2019

Yen Ting Lin,^{a),b)} Song Feng (冯松),^{a),c)} and William S. Hlavacek^{d)}

AFFILIATIONS

Center for Nonlinear Studies and Theoretical Biology and Biophysics Group, Theoretical Division,
Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA

^{a)} **Contributions:** Y. T. Lin and S. Feng contributed equally to this work.

^{b)} **Present address:** Information Sciences Group, Computer, Computational and Statistical Sciences Division,
Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA. **Electronic mail:** yentingl@lanl.gov.

^{c)} **Present address:** Biosciences Center, National Renewable Energy Laboratory, Golden, CO 80401, USA.

Electronic mail: song.feng@outlook.com.

^{d)} **Electronic mail:** wish@lanl.gov.

ABSTRACT

Various kinetic Monte Carlo algorithms become inefficient when some of the population sizes in a system are large, which gives rise to a large number of reaction events per unit time. Here, we present a new acceleration algorithm based on adaptive and heterogeneous scaling of reaction rates and stoichiometric coefficients. The algorithm is conceptually related to the commonly used idea of accelerating a stochastic simulation by considering a subvolume $\lambda\Omega$ ($0 < \lambda < 1$) within a system of interest, which reduces the number of reaction events per unit time occurring in a simulation by a factor $1/\lambda$ at the cost of greater error in unbiased estimates of first moments and biased overestimates of second moments. Our new approach offers two unique benefits. First, scaling is adaptive and heterogeneous, which eliminates the pitfall of overaggressive scaling. Second, there is no need for an *a priori* classification of populations as discrete or continuous (as in a hybrid method), which is problematic when discreteness of a chemical species changes during a simulation. The method requires specification of only a single algorithmic parameter, N_c , a global critical population size above which populations are effectively scaled down to increase simulation efficiency. The method, which we term partial scaling, is implemented in the open-source BioNetGen software package. We demonstrate that partial scaling can significantly accelerate simulations without significant loss of accuracy for several published models of biological systems. These models characterize activation of the mitogen-activated protein kinase ERK, prion protein aggregation, and T-cell receptor signaling.

Published under license by AIP Publishing. <https://doi.org/10.1063/1.5096774>

I. INTRODUCTION

Kinetic Monte Carlo (KMC) procedures,¹ such as the well-known direct and next-reaction methods of Gillespie,² are commonly used to study stochastic chemical kinetics, especially in biochemical systems. These procedures have been widely used to study systems in which population fluctuations arise from intrinsic molecular noise and the discreteness of the chemical species of interest.³ Population fluctuations are characteristic of many genetic regulatory circuits, as populations in these systems tend to be discrete.

KMC procedures are also useful for studying systems, such as cell signaling networks, which have large state spaces⁴ arising from the combinatorial number of chemical species that can be

generated by biomolecular interactions of interest.⁵ For such systems, it may be impracticable, even with the aid of a computer, to enumerate the chemical species that are potentially populated. Nonetheless, simulations can be performed by formulating rules to represent biomolecular interactions⁵ and then using these rules as event generators in a so-called network-free simulation algorithm,⁴ such as that implemented in the NFsim software package.^{6,7} In a network-free simulation (of cell signaling dynamics), stochastic effects may be negligible.⁸ Populations in cell signaling networks tend to be large such that population densities are nearly continuous variables.

In any exact KMC procedure, the system state is updated (and time is advanced) only one reaction event at a time, which can be a severe limitation when the number of events per unit time is large. In

many applications of KMC aimed at studying noisy system behavior, this difficulty is not encountered because populations of interest are small, which makes for efficient simulation. However, inefficiency may arise if some population sizes are large and/or some reactions are fast. Both of these features of a reaction system can introduce a large, unmanageable number of reaction events per unit time. A variety of approaches have therefore been developed for addressing this problem.

In the near-continuum limit where populations of chemical species are uniformly large (but not infinite such that concentrations can be appropriately treated as continuous variables), the diffusive stochastic differential equation (SDE) obtained from the diffusion approximation^{9,10} provides an efficient solution for simulating stochastic chemical kinetics. Gillespie's τ -leaping method¹¹ provides an efficient solution for another regime of behavior, where the discreteness of populations is still relevant. In this method, a time window τ is prescribed. Importantly, τ should be (1) large enough such that multiple events occur within this window and (2) small enough such that the system's configuration and state transition rates (also called reaction propensity functions) do not change significantly as a result of the events occurring within the window. Unfortunately, these opposing requirements limit the applicability of the method. Finally, hybrid methods have been developed and applied to solve specific problems.^{12–21} A characteristic feature of these methods is the division of state variables into two distinct sets, the discrete variables and the continuous variables. The continuous variables are evolved forward in time via numerical integration of ordinary differential equations (ODEs), while the discrete variables are simultaneously evolved forward in time via KMC. A general framework for hybrid simulation is provided by the formalism of discretely switching Langevin dynamics.²²

All of the abovementioned methods are designed to trade accuracy for efficiency. In other words, acceleration gains are obtained through simplifying assumptions that introduce approximations. The diffusive SDE approach is based on the assumption that only the first and second moments of distributions are important; higher order moments are ignored. The τ -leaping approach is based on the assumption of Poissonian statistics in the time window τ . Hybrid approaches are based on the assumption that fluctuations in certain populations can be ignored entirely over the full time window of simulation. For most, if not all, of these methods, there is no way to determine *a priori* if their approximations are tolerable, and unacceptable errors can arise unexpectedly during the course of a simulation. SDE and hybrid approaches become inappropriate if populations of all or a subset of chemical species, respectively, are not sufficiently large throughout a simulation. As we will see, for some systems, population sizes are distributed across multiple scales, and the size of a population can change qualitatively (e.g., from large to discrete to large again) during the course of a simulation, as when a system exhibits oscillatory behavior. The τ -leaping method is also problematic: for a given system, there is no guarantee that a time window τ having the necessary properties exists. For all of these methods, approximation accuracy depends in a nonobvious way on parameters (initial conditions and rate constants). Thus, their use in any analysis involving variation of parameter values, such as a curve fitting procedure, requires careful consideration.

An alternative, simpler idea for accelerating stochastic simulation is to consider only a subvolume of the system of interest, taking this subvolume to be representative of the whole system. Given a system with volume Ω , if we simulate reactions in only a subvolume $\lambda\Omega$, where $0 < \lambda < 1$, we can expect to reduce simulation costs by the factor $1/\lambda$ because the number of reaction events per unit time is reduced by this factor. As we will discuss later in detail, this homogeneous and static scaling approach can speed up estimation of first moments of the stochastic dynamics, and estimates are unbiased so long as populations are not scaled too aggressively to a size of order 1 or below. A drawback of scaling is that, with achievement of any acceleration, there is an unavoidable bias in estimation of higher moments. Another drawback is that the acceleration attainable with reasonable accuracy may be limited by the small population size of one or more critical chemical species. If system behavior is influenced by a chemical species with a small population size, as in stochastic switching,²³ the setting for λ cannot be such that the population of the chemical species falls near or below 1, which would fundamentally change or eliminate its influence on behavior.

Scaling is commonly used as part of a (network-free) stochastic simulation approach when the dynamics of interest are nearly deterministic, as recommended, for example, by Faeder *et al.*²⁴ In such cases, the goal of simulation is simply to calculate first moments. The inability of scaling to yield unbiased estimates of higher moments is of little concern. However, in these applications, there remains a need to avoid overaggressive scaling of critical populations. As we will see through examples, this constraint can severely limit the usefulness of scaling via the standard (homogeneous and static) approach.

Here, we present a new approach for accelerating stochastic simulation through scaling, which we term partial scaling. In partial scaling, for each enumerated individual reaction considered in a system, its reaction rate is scaled by a factor that effectively makes the smallest reactant or product population, if greater than a critical population size N_c , equivalent or close to this critical population size. Stoichiometric coefficients are scaled by the inverse of the scaling factor for the reaction rate. No scaling is applied if the smallest reactant or product population is already at or below N_c . The scaling is heterogeneous (vs homogeneous) because each reaction has its own scaling factor, which is determined by N_c and the smallest population of the chemical species involved in the reaction. Furthermore, reaction rates are scaled *on the fly* as population sizes change over time, i.e., scaling is adaptive (vs static). Because scaling is not performed for a reaction when the smallest population of the participating chemical species is smaller than N_c , partial scaling permits not only unbiased estimation of first moments but also less biased estimation of second moments compared to the standard scaling method. Although partial scaling entails heterogeneous and adaptive scaling, whereas standard scaling entails homogeneous and static scaling, both methods involve only a single algorithmic parameter: λ in the case of standard scaling and N_c in the case of partial scaling.

The remainder of this report is organized as follows. In Sec. II, we introduce notation. In Sec. III, we review how scaling is currently used to accelerate stochastic simulation. Through analysis of a simple reaction network, we explain why this approach yields unbiased estimates of first moments while improving simulation efficiency at

the expense of biased estimates of higher moments (e.g., overestimates of variance). We also illustrate the consequences of overaggressive scaling. In Sec. IV, to address the limitations of standard scaling, we introduce partial scaling. To complement the description of the method, we provide pseudocode for a partial scaling algorithm. In Sec. V, using published models for three complex biochemical reaction networks,^{25–27} we evaluate the performance of partial scaling relative to standard scaling. In Sec. VI, we call attention to our general-purpose implementation of partial scaling in the BioNetGen software package.²⁸ We conclude with a discussion of results and possible future directions in Sec. VII.

II. NOTATION

We consider a closed, well-mixed, and isothermal reaction system having constant volume²⁹ Ω . The system contains a (dilute) mixture of up to M chemical species, $X_1, \dots, X_M$, that react within a network of R reactions. We assume that each reaction $r \in \{1, \dots, R\}$ obeys mass-action kinetics with rate constant κ_r . We use N_i to denote the non-negative and discrete population of X_i . The discrete state of the system is defined by the vector $\mathbf{N} \equiv (N_1, \dots, N_M)$. We use $\mathbb{P}_{\mathbf{N}}(t)$ to denote the probability that the system is in state $\mathbf{N}$ at time t . In the continuum limit, reached as $\Omega \rightarrow \infty$ and $N_i \rightarrow \infty$ while N_i/Ω remains constant for all i , the concentration of each X_i , defined as $n_i \equiv N_i/\Omega$, is a continuous variable, and the system state is defined by the vector $\mathbf{n} \equiv (n_1, \dots, n_M)$. We use $\rho(\mathbf{n}, t)$ to denote the probability density of state $\mathbf{n}$ at time t .

III. STATIC AND HOMOGENEOUS SCALING

In this section, which can be bypassed by readers who are expert in stochastic modeling, we provide a brief review of scaling as a tool to accelerate stochastic simulation, mainly through the exercise of analyzing a simple reaction system. Our goal is to introduce background facts important for appreciating the benefits and limitations of scaling, which should strictly be considered only for a system with large populations such that (all) concentrations are nearly continuous. As we will see later, many biochemical systems have concentrations that are nearly continuous. The background facts of concern here are as follows. First, the number of events generated per unit time in an exact stochastic simulation is an extensive quantity. Accordingly, the efficiency of an exact stochastic simulation algorithm (SSA) scales with system size. Second, if after scaling, the system of interest remains near the continuum limit, exact simulations of the scaled system yield unbiased estimates of the first moments of the random variables that are being sampled in the simulations. Third, scaling can introduce discreteness in populations that is inappropriate. In other words, overaggressive scaling can introduce systematic errors.

An exact SSA generates sample paths of the stochastic dynamics of a (well-mixed) reaction system by executing a series of reaction events (in a Monte Carlo fashion) that each brings about discrete changes of the populations of the reacting chemical species. State transition rates are characterized by propensity functions that, by construction, are extensive quantities.^{9,30} Thus, as we will see shortly for a specific example, the values of these functions scale linearly with system size Ω , and as Ω increases, an SSA generates more

events per unit time. Events occur at an overall rate proportional to Ω and inversely proportional to the average waiting time between events. Consequently, as Ω increases, execution of an SSA eventually becomes too computationally expensive to be practical.

When system size is large, the joint probability distribution of the discrete-state random process sampled by an SSA can be approximated by various methods. Ordered by degree of granularity of approximation, these methods include (1) τ -leaping, which approximates the state-dependent transition kernel of an SSA by a constant (Poissonian) transition kernel, which is used over a small time window τ ;^{11,31,32} (2) the diffusion approximation, which coarse-grains the discrete state space into a continuous state space via the Kramers–Moyal expansion, the end result of which is a Fokker–Planck equation describing the joint probability density of the system's state characterized in terms of intensive state variables, such as concentrations or population densities;^{9,33–39} and (3) the linear-noise approximation (LNA), which linearizes the multiplicative noise kernel of a Fokker–Planck equation such that the resulting dynamics involves only Gaussian white noise.^{9,40}

Another approach for coping with large system size is to use an SSA to simulate the stochastic dynamics not in the whole system of interest but rather in a suitably small subvolume $\lambda\Omega$, where $0 < \lambda < 1$. The dynamics of the subvolume are taken to represent the dynamics of the whole system. This approach, which we will refer to as *static and homogeneous scaling* or standard scaling, is seemingly justified for a well-mixed system, and from our discussion above, it can be expected to reduce the cost of simulation by the factor $1/\lambda$: the smaller the value of λ , the greater the efficiency gain. However, the approach can produce erroneous and misleading results if scaling is overaggressive, as when scaling reduces the population of a critical chemical species (i.e., one that influences overall system behavior even if its population size is relatively small) to 0 or $\mathcal{O}(1)$. In these cases, the dynamics in the subvolume will differ qualitatively from the dynamics in the whole system. For example, scaling to a population size ~ 1 will introduce bursty behavior, which has unique statistical properties.^{17,18,21} Thus, scaling should be limited, but it is not clear how to impose an appropriate lower bound on λ . A constraint such as $\lambda^{-1} \geq \min_{1 \leq i \leq M} N_i(t)$ requires a trial-and-error procedure to select λ because the bounding term is dependent on both time and parameters (initial conditions and rate constants). Another concern, even with an appropriate setting for λ , is how scaling affects the statistics of sample paths. As we will see, estimates for first moments are unbiased, but scaling leads to biased overestimates of second moments.

A. A simple reaction network

To illustrate features of static and homogeneous scaling, we will consider the following reaction network, consisting of zeroth-, first-, and second-order elementary reactions:

Here, $M = 2$ and $R = 3$. The chemical species X_1 is injected via a zeroth-order reaction into the system (volume Ω) with rate constant κ_1 . In a second-order reaction, two copies of X_1 react irreversibly to form X_2 with rate constant κ_2 . The chemical species X_2 is removed from the system via a first-order reaction with rate constant κ_3 .

B. Deterministic and stochastic chemical kinetics

For the reaction network of Eqs. (1), we can write the following ODEs for mass-action kinetics:

$$\frac{dn_1}{dt} = \kappa_1 - \kappa_2 n_1^2, \quad (2a)$$

$$\frac{dn_2}{dt} = \frac{1}{2} \kappa_2 n_1^2 - \kappa_3 n_2, \quad (2b)$$

where $n_1(t)$ is the time-dependent concentration of X_1 and $n_2(t)$ is the time-dependent concentration of X_2 . It should be noted that the $1/2$ factor in Eq. (2b) accounts for the symmetry of the left-hand side of Eq. (1b)⁴¹ and that all terms on the right-hand side of Eqs. (2) incorporate the stoichiometric coefficients of the reactions in Eqs. (1). The terms κ_1 , $\kappa_1 n_1^2$, $(1/2)\kappa_2 n_1^2$, and $\kappa_3 n_2$ are the mass-action rate laws that follow from reactions $r = 1, 2, 3$ in Eqs. (1); these terms give the rates of X_1 injection, X_1 consumption, X_2 generation, and X_2 removal, respectively. The ODEs of Eqs. (2) provide a deterministic description of the chemical kinetics, with each n_i corresponding to N_i/Ω in the continuum limit.

To obtain a stochastic description of the chemical kinetics,^{9,30} let us enumerate the possible state transitions when the system is in state $\mathbf{N} = (N_1, N_2)$. From Eqs. (1), the possible state transitions and their stochastic transition rates (given above the arrows) are as follows:

Here, we use $h_r(\mathbf{N})$ to denote the number of ways reaction r can take place and $\bar{\kappa}_r$ to denote the transition rate for reaction r when this reaction can take place in one and only one way.⁴² For reaction $r = 1$ [Eq. (1a)], there is only one way to inject a copy of X_1 into the system. Thus, $h_1 = 1$. For reaction $r = 2$ [Eq. (1b)], the number of ways that two copies of X_1 can react to form X_2 is N_1 choose 2. Thus, $h_2 = C_2^{N_1} = N_1(N_1 - 1)/2$. Finally, for reaction $r = 3$ [Eq. (1c)], there are N_2 choose 1 ways to remove a copy of X_2 from the system. Thus, $h_3 = C_1^{N_2} = N_2$.

The transition rates in Eqs. (3) are related to the mass-action rate laws of Eqs. (2). The relationships are revealed by considering the continuum limit, where Eqs. (3) must be consistent with Eqs. (2). Equation (3a) indicates that injection of X_1 increases the concentration of X_1 (i.e., N_1/Ω), on average, by $\bar{\kappa}_1/\Omega$ per unit time, which is the rate characterized by the rate law κ_1 in Eq. (2a). Thus, we find that $\bar{\kappa}_1 = \kappa_1 \Omega$. Equation (3b) indicates that conversion of X_1 to X_2 decreases the concentration of X_1 (i.e., N_1/Ω), on average, by $2 \times (1/2)\bar{\kappa}_2 N_1(N_1 - 1)/\Omega$ per unit time [noting that each firing of the reaction in Eq. (3b) consumes two copies of X_1], which is the rate

characterized by the rate law $\kappa_2 n_1^2$ in Eq. (2a). Thus, taking $N_1 \ll N_1^2$, which is appropriate in the continuum limit, we find that $\bar{\kappa}_2 = \kappa_2/\Omega$. Equation (3c) indicates that removal of X_2 decreases the concentration of X_2 (i.e., N_2/Ω), on average, by $\bar{\kappa}_3 N_2/\Omega$ per unit time, which is the rate characterized by the rate law $\kappa_3 n_2$ in Eq. (2b). Thus, we find that $\bar{\kappa}_3 = \kappa_3$. In general,⁹ for a k th order reaction r , $\bar{\kappa}_r$ is equal to $\kappa_r \Omega^{1-k}$.

C. Efficiency, accuracy, and precision of standard scaling

The utility of scaling derives from the fact that transition rates (i.e., propensity functions) are extensive, meaning that they scale with system size. To verify this claim, consider the continuum limit, which is approached as $\Omega \rightarrow \infty$ and $N_i \rightarrow \infty$ for all i , but with each N_i/Ω held constant. Thus, in the finite but large population limit, $N_i \propto \Omega \forall i$. Accordingly, the stochastic transition rate of any k th order reaction r has magnitude given by the expression $\bar{\kappa}_r \mathcal{O}(\Omega^k)$. However, from the relationship discussed above, $\bar{\kappa}_r \propto \kappa_r \Omega^{1-k}$, we can see that the expression simplifies to $\mathcal{O}(\Omega)$, indicating that the number of reaction events per unit time is proportional to the system size. By reducing system size $\Omega \rightarrow \lambda \Omega$, a speed up by a factor of $1/\lambda$ can be obtained as events occur $1/\lambda$ times less frequently.

A key idea of scaling is to take the stochastic dynamics within a subvolume of a system, or scaled system, to be representative of the dynamics in the whole system, or unscaled system. What this means in practice is as follows. In the simulation of a scaled system, when a reaction event causes a change in the population of species X_i by ΔN_i , we take this change to correspond to a change of $\Delta N_i/\lambda$ in the unscaled state space. Thus, the method of static and homogeneous scaling can be seen as scaling down the values of all stochastic transition rates by a factor λ and scaling up the values of all stoichiometric coefficients by a factor $1/\lambda$. As we will see later, this interpretation can be generalized to make the scaling of transition rates and stoichiometric coefficients adaptive and heterogeneous. However, let us first establish the accuracy and precision of static and homogeneous scaling, which we will pursue through an analysis of the reaction network of Eqs. (1).

We start by considering the chemical master equation^{9,35} (CME) for the reaction network of Eqs. (1). For arbitrary N_1 and N_2 , the CME includes an ODE of the following form:

$$\begin{aligned} \frac{d}{dt} \mathbb{P}_{N_1, N_2} = & -\Omega \kappa_1 [\mathbb{P}_{N_1, N_2} - \mathbb{P}_{N_1-1, N_2}] - \frac{\kappa_2}{\Omega} \frac{N_1(N_1-1)}{2} \mathbb{P}_{N_1, N_2} \\ & + \frac{\kappa_2}{\Omega} \frac{(N_1+2)(N_1+1)}{2} \mathbb{P}_{N_1+2, N_2-1} \\ & - \kappa_3 [N_2 \mathbb{P}_{N_1, N_2} - (N_2+1) \mathbb{P}_{N_1, N_2+1}]. \end{aligned} \quad (4)$$

Recall that $\mathbb{P}_{\mathbf{N}}(t)$, written here as $\mathbb{P}_{N_1, N_2}$, denotes the probability of system state $\mathbf{N} = (N_1, N_2)$ at time t .

If the system has large populations (i.e., the probability mass concentrates in regions where both N_1 and N_2 are large), we can perform the Kramers–Moyal expansion^{9,33,34} to transform the CME into a partial differential equation describing a diffusive process. This procedure begins by replacing extensive variables (discrete populations) by intensive ones (concentrations/population densities). We introduce the intensive variables $n_1 \equiv N_1/\Omega$ and $n_2 \equiv N_2/\Omega$. We also

introduce $dn \equiv dn_1 \equiv dn_2 \equiv 1/\Omega$, which is the size of the grid in the state space of n_1 and n_2 . Finally, we replace discrete probabilities with probability densities such that $\mathbb{P}_{N_1, N_2}(t) = \rho(n_1, n_2, t) dn_1 dn_2 = \rho(n_1, n_2, t)/\Omega^2$. Thus, after changing variables, we obtain

$$\begin{aligned} \partial_t \rho(n_1, n_2, t) = & -\Omega \kappa_1 [\rho(n_1, n_2, t) - \rho(n_1 - dn, n_2, t)] \\ & - \kappa_2 \Omega \frac{n_1(n_1 - dn)}{2} \rho(n_1, n_2, t) \\ & + \kappa_2 \Omega \frac{(n_1 + 2dn)(n_1 + dn)}{2} \rho(n_1 + 2dn, n_2 - dn, t) \\ & - \kappa_3 \Omega [n_2 \rho(n_1, n_2, t) - (n_2 + dn) \rho(n_1, n_2 + dn, t)]. \end{aligned} \quad (5)$$

At this stage, Eq. (5) is simply a rewritten form of Eq. (4), obtained after changing from extensive to intensive variables. As such, $\rho(n_1, n_2, t)$ can only be evaluated on a lattice $(n_1, n_2) \in \{0, 1/\Omega, 2/\Omega, \dots\}^2$.

The first key idea of the Kramers–Moyal expansion for the CME is to extend the domain of the function ρ from $(n_1, n_2) \in \{0, 1/\Omega, 2/\Omega, \dots\}^2$ to $\mathbb{R}_{\geq 0}^2$ while assuming (5) holds true for any non-negative and real-valued n_1 and n_2 . This transforms Eq. (5) into a difference equation for ρ . The second key idea is that under the assumption that $\Omega \gg 1$, the difference equation can be approximated by local derivatives.

Under the assumption that $\Omega \gg 1$ and $dn \ll 1$ in Eq. (5), which is viewed now as a difference equation, the Kramers–Moyal expansion procedure continues by Taylor expanding the terms $\rho(n_1 \pm dn, n_2 \pm dn)$ in Eq. (5). After all terms in each Taylor expansion higher than second order are dropped,⁴⁷ we obtain the following Fokker–Planck equation:

$$\begin{aligned} \partial_t \rho = & -\partial_{n_1} \left[\left(\kappa_1 - \kappa_2 n_1^2 \right) \rho \right] - \partial_{n_2} \left[\left(\frac{\kappa_2}{2} n_1^2 - \kappa_3 n_2 \right) \rho \right] \\ & + \frac{1}{2\Omega} \left\{ \partial_{n_1}^2 \left[\left(\kappa_1 + \kappa_2 n_1^2 \right) \rho \right] + \partial_{n_2}^2 \left[\left(\frac{\kappa_2}{2} n_1^2 + \kappa_3 n_2 \right) \rho \right] \right. \\ & \left. + 2\partial_{n_1} \partial_{n_2} \left[\left(\kappa_2 n_1^2 \right) \rho \right] \right\}. \end{aligned} \quad (6)$$

The solution to Eq. (6) approximates the stochastic dynamics of the system of interest [Eqs. (1)] when Ω is large but finite (i.e., $\Omega \gg 1$).

It is worth pointing out that in the continuum limit $\Omega \rightarrow \infty$, the Fokker–Planck equation (6) reduces to a Liouville equation,

$$\partial_t \rho = -\partial_{n_1} \left[\left(\kappa_1 - \kappa_2 n_1^2 \right) \rho \right] - \partial_{n_2} \left[\left(\frac{\kappa_2 n_1^2}{2} - \kappa_3 n_2 \right) \rho \right]. \quad (7)$$

It can be shown that when the initial condition of the probability density function is a Dirac δ -distribution, $\rho(x, y, t=0) = \delta(x - x_0)\delta(y - y_0)$, the temporal evolution of the probability density $\rho(x, y, t)$ remains as a δ -distribution and the peak of the distribution corresponds to the solution of Eqs. (2) for any given time $t > 0$. Thus, Eq. (7) is a slightly generalized version of the equations for mass-action kinetics [Eqs. (2)]. Equation (7) captures the same information as Eqs. (2), but Eq. (7) permits a probabilistic distribution as an initial condition.

When Ω is large but finite, as is well-known, the dynamics described by a Fokker–Planck equation such as Eq. (6) are largely determined by the equation’s advection or drift terms, i.e., the terms in Eq. (6) that are preserved in the Liouville equation [Eq. (7)]. However, additional terms, i.e., the terms with the operators $\partial_{n_1}^2$, $\partial_{n_2}^2$, and $\partial_{n_1} \partial_{n_2}$, are present in Eq. (6) and these terms introduce diffusion in the state space (n_1, n_2) . It is these terms that characterize density fluctuations that arise from stochastic and discrete reaction events. Because the scale of diffusion is of order $\mathcal{O}(1/\Omega)$ (i.e., the variance), sample paths $(n_1(t), n_2(t))$ exhibit fluctuations of order $\mathcal{O}(1/\Omega^{1/2})$ (i.e., the standard deviation).

Now let us consider the scaled system, i.e., let us replace Ω with $\lambda\Omega$ in the above analysis. If $\lambda\Omega$ remains much larger than 1, and we perform the Kramers–Moyal expansion of the CME for the scaled system, we obtain

$$\begin{aligned} \partial_t \rho = & -\partial_{n_1} \left[\left(\kappa_1 - \kappa_2 n_1^2 \right) \rho \right] - \partial_{n_2} \left[\left(\frac{\kappa_2}{2} n_1^2 - \kappa_3 n_2 \right) \rho \right] \\ & + \frac{1}{2\Omega\lambda} \left\{ \partial_{n_1}^2 \left[\left(\kappa_1 + \kappa_2 n_1^2 \right) \rho \right] + \partial_{n_2}^2 \left[\left(\frac{\kappa_2}{2} n_1^2 + \kappa_3 n_2 \right) \rho \right] \right. \\ & \left. + 2\partial_{n_1} \partial_{n_2} \left[\left(\kappa_2 n_1^2 \right) \rho \right] \right\}. \end{aligned} \quad (8)$$

This equation is nearly the same as Eq. (6). The only difference is the factor multiplying the diffusive terms. The factor is $1/(2\Omega)$ in Eq. (6) and $1/(2\Omega\lambda)$ in Eq. (8). This result generalizes to any chemical reaction system with large species populations (i.e., $N_i \gg 1$ for all i).

The significance of our comparison of Eqs. (6)–(8) is threefold. First, it is known that first moments are dominated by drift (i.e., the terms with the operators ∂_{n_1} and ∂_{n_2}) when $\Omega \gg 1$, and drift is identical in Eqs. (6)–(8). Thus, we can expect simulations of scaled and unscaled systems to yield comparable estimates of first moments when Ω and $\lambda\Omega$ are large enough such that effects of multiplicative noise can be ignored. Second, we can expect scaling to yield overestimates of second moments because diffusivity (i.e., variance) is enhanced by a factor of $1/\lambda$ in the scaled system relative to that in the unscaled system. Third, given that the standard error of the mean (SEM) is calculated as $\sigma/\sqrt{N}$, where σ is the sample standard deviation and N is the number of samples, and given that we can expect the sample variance σ^2 to be amplified by a factor of $1/\lambda$ in estimates of first moments based on scaled simulations, the SEM for each first moment estimate is amplified by a factor of $1/\sqrt{\lambda}$ when using the same number of sample paths from scaled simulations as for estimates based on unscaled simulations. This loss of precision can be overcome by increasing the number of sample paths used. However, the number of sample paths must be increased by a factor of $1/\lambda$ to match the precision of estimates from unscaled simulations. This factor exactly matches the efficiency gain of a single scaled simulation. Thus, scaling yields acceleration *only* when a loss of precision in first moment estimates is permitted.

D. Limitation by small but critical populations

The Kramers–Moyal expansion of the scaled system is only valid when $\mathcal{O}(\lambda N_i) \gg 1$ for each i and $dn \ll 1$. In the context of

FIG. 1. Simulations based on Eqs. (1). In the six panels at left (a)–(f), we compare stochastic sample paths found using an SSA to the corresponding deterministic trajectories found by numerically integrating the ODEs in Eqs. (2). We consider a range of system sizes (without changing intensive parameters): $\Omega = 10^4$ [(a) and (b)], 10^2 [(c) and (d)], and 3 [(e) and (f)]. The model parameters are $\kappa_1 = 3$, $\kappa_2 = 300$, and $\kappa_3 = 1.5$. In the top panels [(a), (c), and (e)], we compare the random variable N_1/Ω to the deterministic quantity n_1 . Similarly, in the bottom panels [(b), (d), and (f)], we compare N_2/Ω to n_2 . For panels (a)–(f), as indicated by the legend of each panel, we plot a single stochastic sample path, the first moment estimated by averaging over 500 sample paths, and the corresponding deterministic trajectory. In panels [(g) and (h)], we plot the marginal stationary distributions for N_1/Ω and N_2/Ω , each measured from 10^5 sample paths. Note that, when $\Omega = 3$, the probability distribution of N_1/Ω is skewed to the right (g), leading to a first moment estimation that deviates noticeably from the continuum limit (e). Time and volume each have arbitrary units (a.u.).

scaling, these conditions are satisfied only when $\lambda \gg 1/\min_{1 \leq i \leq M}\{N_i\}$. When λ is too small, one or more populations become discrete. Discrete populations have distinct dynamics, which are not captured by a Fokker–Planck equation.^{17,18,21} Thus, the analysis of Sec. III C

eventually becomes invalid and errors are introduced as λ decreases and becomes too small. For example, for the system of Eqs. (1), $\lambda = 3 \times 10^{-4}$ is evidently too small, as indicated by the errors in the estimates of the first moments that can be seen in Figs. 1(e)

FIG. 2. The stochastic dynamics of models for (a) activation of the mitogen-activated protein kinase (MAPK) ERK,²⁵ (b) prion protein (PrP^{Sc}) aggregation,²⁶ and (c) T-cell receptor (TCR) signaling.²⁷ 500 independent sample paths were generated through continuous-time Markov chain simulation,^{30,43} and the first moments of all chemical species considered in each model were each estimated by averaging over the generated sample paths. Plots in each panel show these first moment estimates. As can be seen, for each model, there is a broad spectrum of population scales, which spans several orders of magnitude. The models and simulations considered here are defined in BioNetGen input files (also called BNGL files) available online within the RuleHub repository.^{44–46} The BNGL file for the ERK activation model⁴⁴ is an edited version of the BNGL file provided by Kočańczyk *et al.*²⁵ The BNGL file for the prion protein aggregation model⁴⁵ is new; we confirmed that this formulation of the model is consistent with the simulation results reported by Rubenstein *et al.*²⁶ The BNGL file for the TCR signaling model⁴⁶ is new. We rewrote the BNGL file originally provided by Lipniacki *et al.*²⁷ for compatibility with current BNGL conventions; we confirmed that the new formulation of the model is consistent with the original formulation.

ALGORITHM 1. Partial scaling algorithm (PSA).

This algorithm produces a sample path of the stochastic dynamics of a chemical reaction network. The following inputs are required to execute the algorithm. (1) A chemical reaction network consisting of R reactions among M chemical species. Reactions are indexed by $r \in \{1 \dots R\}$ and chemical species are indexed by $s \in S \equiv \{1 \dots M\}$. (2) The number of reactant and product species participating in each reaction r , $\mu_r \in \mathbb{N}$ and $\nu_r \in \mathbb{N}$, respectively.^a (3) The set of indices of reactant and product species participating in each reaction r , $\{\alpha_{r,k}\}_{k=1}^{\mu_r}$ and $\{\beta_{r,k}\}_{k=1}^{\nu_r}$, where each $\alpha_{r,k}, \beta_{r,k} \in S$. (4) The propensity function of each reaction r , $f_r(\mathbf{N}) \equiv h_r(\mathbf{N})\tilde{\kappa}_r$. (5) The stoichiometric coefficient $\xi_{r,s}$ for each species s participating in each reaction r . We take stoichiometric coefficients to be negative for reactants and positive for products. (6) The initial condition, $\{N_s(t=0)\}_{s=1}^M$. (7) A list of K discrete times $0 \leq t_1 \leq t_2 \leq \dots \leq t_K$ at which system state should be reported. The last input is a critical population scale N_c .

```

1:  $t \leftarrow 0$  ▷ Initialize time  $t$ 
2: for  $s$  in  $\{1, \dots, M\}$  do
3:    $N_s \leftarrow N_s(t=0)$  ▷ Initialize each population  $N_s$ 
4: end for
5: for  $r$  in  $\{1, \dots, R\}$  do
6:    $S_r \leftarrow \{\alpha_{r,1}, \dots, \alpha_{r,\mu_r}\} \cup \{\beta_{r,1}, \dots, \beta_{r,\nu_r}\}$  ▷ Create a list of chemical species that participate in each reaction  $r$ 
7: end for
8: for  $i$  in  $\{1, \dots, K\}$  do ▷ Start continuous-time Markov chain (CTMC)
9:   while  $t < t_i$  do
10:    for  $r$  in  $\{1, \dots, R\}$  do
11:       $\lambda_r \leftarrow 1 / \max\left\{1, \left\lfloor \frac{\min_{s \in S_r} \{N_s(t)\}}{N_c} \right\rfloor\right\}$  ▷ Calculate the scaling factor  $\lambda_r$  for each reaction  $r$ 
12:       $\kappa_r \leftarrow \lambda_r f_r(\mathbf{N})$  ▷ Calculate the scaled rate of reaction for each reaction  $r$ 
13:    end for
14:     $\kappa \leftarrow \sum_{r=1}^R \kappa_r$  ▷ Calculate the scaled overall rate of reaction
15:     $\Delta t \leftarrow \text{Exp}(\kappa)$  ▷ Generate a random waiting time to the next reaction
16:    if  $t + \Delta t < t_i$  then ▷ Check if the advanced time is ahead of the next report time
17:       $t \leftarrow t + \Delta t$  ▷ Increment time
18:       $\phi \leftarrow \kappa \times \text{Unif}(0, 1)$  ▷ Inverse sampling to select the next reaction  $r$ 
19:       $r \leftarrow 1$ 
20:      while  $\sum_{j=1}^r \kappa_j < \phi$  do
21:         $r \leftarrow r + 1$ 
22:      end while
23:      for  $s$  in  $S_r$  do
24:         $N_s \leftarrow N_s + (1/\lambda_r) \times \xi_{r,s}$  ▷ Update the populations in  $S_r$  (using scaled stoichiometric coefficients)
25:      end for
26:    else
27:       $t \leftarrow t_i$  ▷ Advance time to the next report time without updating the system state
28:      Export state vector  $\mathbf{N}(t)$  ▷ Report system state (all populations)
29:    end if
30:  end while
31: end for ▷ Perform next step in CTMC until reports have been made for all specified report times

```

^aFor simplicity, we assume $\mu_r, \nu_r > 0$, but the algorithm can be trivially generalized to account for $\mu_r = 0$ (e.g., creation) or $\nu_r = 0$ (e.g., annihilation).

and 1(g). These errors arise because the scaled system is no longer representative of the unscaled system.

Scaling a system so that the smallest population vanishes (i.e., becomes 0) or reduces to order 1 might be acceptable in practice but requires careful consideration. On the one hand, the population of concern may be of little significance for system behavior, and the erroneous behavior predicted for it can be safely ignored. On the other hand, if the population is critical for system behavior, then scaling will produce seriously misleading results.

Specialized methods are available for coping with discrete populations that influence system dynamics so long as these populations are always discrete and well-separated from other populations.^{14,17,18,21} However, for many systems, populations are not only distributed over a spectrum of scales but also dynamically changing scale, as illustrated in Fig. 2. We suspect that this behavior is generic, at least for biochemical systems, which limits the applicability of static and homogeneous scaling. Below, we present a new scaling approach that overcomes this limitation.

IV. ADAPTIVE AND HETEROGENEOUS SCALING

There are two barriers to applying static and homogeneous scaling. First, population scales may be scattered across a broad spectrum (Fig. 2). Second, populations may evolve dynamically across multiple scales (Fig. 2). To address these problems, we propose a new scaling approach, which we term *partial scaling*, which, in

contrast with standard scaling, is adaptive and heterogeneous, as we will see. Both of these features of partial scaling arise from tying scaling to population sizes, on the fly. Scaling is heterogeneous because population sizes are heterogeneous, and scaling is adaptive because populations are dynamic. A particular, intentionally simple implementation of partial scaling is outlined as pseudocode in Algorithm 1. We call this algorithm the partial scaling algorithm (PSA).

With partial scaling, there is no longer a global scaling factor λ , as in standard scaling. Instead, at any given time t , there is a scaling factor $\lambda_r(t)$ for each reaction $r = 1, \dots, R$. These scaling factors are not set in advance of a simulation nor statically. Rather, they are assigned values according to a dynamic update schedule (e.g., a schedule synchronized with time updates). This update schedule ensures that

$$\lambda_r(t) = \frac{1}{\max\left\{1, \left\lfloor \frac{N_{\min}^r(t)}{N_c} \right\rfloor\right\}}, \quad (9)$$

where N_c is a threshold or critical population—this quantity is the *one* parameter of the method—and $N_{\min}^r(t)$ is the smallest population among those of the reactants and products of reaction r at time t . It should be noted that the brackets around the ratio $N_{\min}^r/N_c$ in Eq. (9) denote the floor function and, furthermore, that there is no scaling (i.e., $\lambda_r = 1$) for reaction r if any reactant or product population is smaller than N_c .

The scaling factors $\{\lambda_r(t)\}_{r=1}^R$ defined by Eq. (9) are used to modify an SSA, such as Gillespie's direct method,³⁰ as follows. In

FIG. 3. Comparison of the standard scaling method [(a), (c), (e), and (g)] with the partial scaling method [(b), (d), (f), and (h)] for simulations based on a model for ERK activation.²⁵ Shown here are the mean populations for four selected chemical species, labeled 8, 12, 14, and 18. These labels are those produced by BioNetGen; the species are defined in the NET file produced as output when BioNetGen processes the BNGL file defining the ERK activation model.⁴⁴ Means are based on 500 sample paths. Results from standard scaling begin to deviate from those obtained for the unscaled system via exact simulation when the scaling factor is still fairly large ($\lambda = 0.1$). The partial scaling method yields reasonable estimates for both first and second moments, even with very aggressive scaling ($N_c = 300$). Compare 300 to the largest population size in Fig. 2(a), which is on the order of 10^6 .

the calculation of stochastic transition rates [e.g., as for Eqs. (3) as we have discussed], the rate for each reaction r is calculated as usual but then scaled (down) by a factor λ_r . When reaction r with rate scaled by $\lambda_r < 1$ is selected to fire, to compensate for its reduced rate of firing, its stoichiometric coefficients are each temporarily scaled (up) by the factor $1/\lambda_r$. The amplified stoichiometric coefficients are then used to update the system state $\mathbf{N}$. In other words, $\mathbf{N}$ is modified in accordance with scaled population changes triggered by the selected reaction event. It should be noted that $1/\lambda_r \in \mathbb{N}$ by construction [Eq. (9)]. Thus, scaled stoichiometric coefficients, like unscaled coefficients, are natural numbers. A specific example of how the scaling factors $\{\lambda_r(t)\}_{r=1}^R$ are used to modify an SSA is provided in Algorithm 1. It should be noted that Algorithm 1 reduces to Gillespie's direct method if we statically assign each λ_r a value of 1. Furthermore, standard scaling is achieved by replacing each reaction-specific adaptive scaling factor λ_r in Algorithm 1 with a universal static scaling factor λ .

With partial scaling, population sizes are not directly scaled but the populations of reactants and products participating in any reaction r for which $\lambda_r < 1$ are effectively $\mathcal{O}(N_c)$ or larger. Thus, a choice of $N_c \gg 1$ guarantees that stochastic dynamics do not become inappropriately bursty. Furthermore, any critical species X_i with a small population $N_i < N_c$ will never be scaled out of the system. In other words, N_i is never effectively scaled such that the population vanishes because $\lambda_r = 1$ for any reaction having a reactant or product population less than N_c [Eq. (9)].

A. Accuracy and precision of partial scaling

As we will see in this section, provided that $N_c \gg 1$, partial scaling preserves means but not variances, similar to standard scaling.

Let us consider a time interval $\Delta t \ll (\sum_{r=1}^R h_r \bar{\kappa}_r)^{-1}$, i.e., an interval that is much shorter than the expected waiting time to the next reaction event. Furthermore, consider an arbitrary species X_i , $i \in \{1, \dots, M\}$, which participates, as a reactant or product, in a set of reactions with indices $I \subseteq \{1, \dots, R\}$. We will use $I' \subseteq I$ to denote the subset of these indices corresponding to reactions associated with a scaling factor less than 1. Thus, for each $r \in I'$, $\lambda_r < 1$, and for each $r \in I \setminus I'$, $\lambda_r = 1$. For an unscaled simulation, we will use $\delta N_{i,r}$ to denote the change in population N_i brought about by a reaction $r \in I$ occurring within the time interval Δt . Given that multiple reaction events can occur within the time window Δt (albeit with low probability, by construction), the total change in N_i over this time window, which we will denote as δN_i , is given by

$$\delta N_i = \sum_{r \in I} \delta N_{i,r}. \quad (10)$$

Similarly, for a (partially) scaled simulation, the total population change is given by

$$\delta N_i^{\text{scaled}} = \sum_{r \in I'} \frac{\delta N_{i,r}}{\lambda_r} + \sum_{r \in I \setminus I'} \delta N_{i,r}. \quad (11)$$

FIG. 4. Comparison of the standard scaling method [top row: panels (a), (c), (e), and (g)] with the partial scaling method [bottom row: panels (b), (d), (f), and (h)] for simulations based on a model for prion protein aggregation.²⁶ Mean abundances of four selected chemical species, labeled 1, 2, 15, and 30, are calculated from 500 sample paths. The species labels are those produced by BioNetGen; the species are defined in the NET file produced as output when BioNetGen processes the BNGL file defining the prion protein aggregation model.⁴⁵ The standard scaling method produces results that deviate from those generated for the unscaled system via exact simulation even when the scaling factor is fairly large ($\lambda = 0.3$). This finding is explained by the discreteness of the species seeding prion protein aggregation. The partial scaling method produces reasonable estimates of both first and second moments even with aggressive scaling ($N_c = 300$). Compare 300 with the largest population size in Fig. 2(b), which is on the order of 10^3 – 10^4 .

This expression is the same as that given above except for the introduction of scaling factors for stoichiometric coefficients.

The population changes considered above [i.e., the δN_i and $\delta N_{i,r}$ terms in Eqs. (10) and (11)] are random variables. Because Δt is small (relative to the expected waiting time to the next reaction), we can take each of these random variables to be drawn from a Poisson distribution. Thus,

$$\delta N_{i,r} \sim \text{Poisson}(h_r \bar{\kappa}_r \Delta t), \quad (12a)$$

$$\delta N_{i,r}^* \sim \lambda_r^{-1} \text{Poisson}(\lambda_r h_r \bar{\kappa}_r \Delta t), \quad (12b)$$

where $\delta N_{i,r}^* \equiv \delta N_{i,r} / \lambda_r$ denotes a population change triggered by firing of a reaction r with a scaled rate $\lambda_r h_r \bar{\kappa}_r$.

For constants $c, \mu > 0$, recall that $\mathbb{E}[c \cdot \text{Poisson}(\mu)] = c\mu$ and $\text{var}[c \cdot \text{Poisson}(\mu)] = c^2\mu$. Thus, for chemical species X_i and reaction $r \in I'$,

$$\mathbb{E}[\delta N_{i,r}^*] = \frac{\lambda_r h_r \bar{\kappa}_r \Delta t}{\lambda_r} = h_r \bar{\kappa}_r \Delta t = \mathbb{E}[\delta N_{i,r}], \quad (13a)$$

$$\text{var}[\delta N_{i,r}^*] = \frac{\lambda_r h_r \bar{\kappa}_r \Delta t}{\lambda_r^2} > h_r \bar{\kappa}_r \Delta t = \text{var}[\delta N_{i,r}]. \quad (13b)$$

It should now be clear that partial scaling yields unbiased estimates of first moments and overestimates of second moments (in a time interval $\Delta t \ll 1$), as is also the case for standard scaling. However, the bias in estimates of second moments, overall, is less with partial scaling than with standard scaling because the processes in a (partially) scaled simulation that are unaffected by scaling have the same statistics as the processes in an unscaled simulation.

In partial scaling, each time a reaction r (with scaled rate $\lambda_r h_r \bar{\kappa}_r$) is fired, a *deterministic multiplier*, $1/\lambda_r$, is applied to the reaction's stoichiometric coefficients for the purpose of appropriately updating the system state (Algorithm 1, line 24). As indicated by the analysis above, this multiplier is essential for obtaining unbiased estimates of first moments. Estimates of first moments would still be unbiased if the multiplier were instead chosen randomly with an expected value of $1/\lambda_r$; however, estimates of variance would be more biased for a random multiplier than for a deterministic multiplier.⁴⁸

One may wonder if partial scaling yields unbiased first-moment estimates in the short term but not in the long run. We expect that the answer is model dependent. If a model permits noise-induced phenomena involving chemical species participating

FIG. 5. Comparison of the standard scaling method [top row: panels (a) and (c)] with the partial scaling method [bottom row: panels (b) and (d)] for simulations based on a model for TCR signaling.²⁷ Shown here are the time-dependent marginal probability distributions for two selected chemical species considered in the model, which are labeled 7 and 15. These labels are those produced by BioNetGen; the species are defined in the NET file produced as output when BioNetGen processes the BNGL file defining the TCR signaling model.⁴⁶ The marginal probability distributions have been constructed on the basis of 10^4 sample paths. The standard scaling method produces results that deviate from those for the unscaled system even when the scaling factor is fairly large ($\lambda = 0.3$). By contrast, the partial scaling method generates distributions that are comparable to those calculated on the basis of exact simulations even with very aggressive scaling ($N_c = 100$). Compare 100 with the largest population size in Fig. 2(c), which is on the order of 10^5 – 10^6 .

in reactions with scaled rates, then estimates may become biased because of demographic fluctuations. In other words, systematic errors may arise with partial scaling from behavior that depends on second or higher moments. Below, in part to address this issue, we turn our attention to numerical experiments, including simulations for a model that permits bistable stochastic switching, the TCR model.²⁷ As we will see, partial scaling yields unbiased estimates of first moments (and also higher-order moments), even for this model.

V. COMPARISON OF SCALING METHODS

To facilitate benchmarking of partial scaling, we implemented a variation of Gillespie's direct method, an additional variation that incorporates standard scaling, and a variation of PSA (Algorithm 1). Importantly, we implemented these methods so as to eliminate all unnecessary implementation differences. We used our code to perform stochastic simulations based on three published models^{25–27} with the goal of evaluating the relative efficiency of each of the

three methods (no scaling, standard scaling, and partial scaling) and the relative accuracy of standard and partial scaling. In simulations, we used a model-specific pregenerated dependency graph to aid in updates of propensity functions, but we did not use sorting of the propensity functions to accelerate simulations as in the Gibson-Bruck method.⁴⁹ In our simulations, we considered different settings for the algorithmic parameters λ (the static scaling factor in standard scaling) and N_c (the critical population size used in partial scaling).

The models that we considered, all for biological systems, were chosen because each was originally analyzed using exact stochastic simulation, and each challenges the application of standard scaling. The first of the three models characterizes activation of ERK.²⁵ The dynamics of this model are oscillatory. Populations span multiple orders of magnitude and their scales change dynamically [Fig. 2(a)]. The second of the three models characterizes prion protein aggregation,²⁶ which involves polymerizationlike reactions and seeded nucleation of aggregates. As before, the populations considered in this model evolve over multiple scales,

FIG. 6. Summary of numerical experiments aimed at quantifying the relative accuracy of partial scaling. First-moment errors introduced by standard scaling [left column: panels (a), (c), and (e)] and partial scaling [right column: panels (b), (d), and (f)] in simulations of the ERK activation model [top row: panels (a) and (b)], prion protein aggregation model [middle row: panels (c) and (d)], and TCR signaling model [bottom row: panels (e) and (f)]. As described in the Appendix, temporally averaged errors were calculated for the ensemble averaged trajectories of individual species. Errors are plotted in each panel in order, from the greatest (left) to least (right). The species indices are hidden; the scale of each horizontal axis matches the number of species in the corresponding model. It should be noted that each vertical axis is logarithmic, with the scale varying from model to model. The values of λ_{eff} indicated in the legends at right indicate aggressiveness of scaling (see the Appendix); these values can be compared to the values of λ indicated in the legends at left. Overall first- and second-moment errors for the numerical experiments considered here are reported in Table I.

and the initial abundance of prion protein is discrete [Fig. 2(b)]. The third of the three models characterizes TCR signaling.²⁷ The behavior characterized by this model includes stochastic bistable switching. In other words, for a particular regime of behavior, which we considered in our simulations (as is evident from inspection of individual sample paths), intrinsic noise is capable of driving the system from one stable fixed point to another (and back). This type of behavior has been intensely studied in the context of genetic regulatory circuits.^{14,18,50–55} For the TCR model, unlike for many models of gene regulation that have been considered, many of the populations are large (i.e., near the continuum limit) [Fig. 2(c)].

In Figs. 3 and 4, we plot predicted time courses for selected chemical species considered in the ERK model²⁵ and the prion model,²⁶ respectively. Each point in each time course corresponds to the mean population for the indicated chemical species estimated on the basis of 500 sample paths. (This number of

sample paths was deemed sufficient for accurate estimation of statistical quantities of interest.) In each figure, the top row of panels shows results of calculations based on either exact simulation or standard scaling (for different values of λ), as indicated in the legend at right. Similarly, in each figure, the bottom row of panels shows results of calculations based on either exact simulation or partial scaling (for different values of N_c), as indicated in the legend at right. As can be seen, standard scaling yields estimates that deviate markedly from those based on exact simulation. By contrast, even with aggressive scaling ($N_c = 300$), partial scaling yields first-moment estimates that are barely distinguishable from those based on exact simulation.

In Fig. 5, we plot time-dependent marginal probability distributions (estimated on the basis of 10^4 sample paths) for two selected chemical species considered in the TCR model.²⁷ (We considered 10^4 vs 500 sample paths because here we are estimating distributions vs properties of distributions.) In this figure, we focus on the initial

TABLE I. Efficiency and accuracy of standard scaling and partial scaling methods for stochastic simulation based on models for ERK activation,²⁵ prion protein aggregation,²⁶ and TCR signaling.²⁷ CPU time was measured using the C++ `clock()` function. Summary statistics for first-moment errors, $\langle \epsilon \rangle$, and Van Valen coefficients of variation, $\langle CV_{VV} \rangle$, were calculated as described in the Appendix.

Model	Scaling method	CPU time (s)	$\langle \epsilon \rangle$	$\langle CV_{VV} \rangle$
ERK	None (Gillespie SSA)	$5.30 \times 10^2 \pm 1.17 \times 10^1$	0.00×10^0	2.98×10^{-2}
ERK	Standard ($\lambda = 0.3$)	$1.57 \times 10^2 \pm 6.10 \times 10^0$	4.36×10^{-1}	5.39×10^{-2}
ERK	Standard ($\lambda = 0.1$)	$5.21 \times 10^1 \pm 2.42 \times 10^0$	1.31×10^0	9.44×10^{-2}
ERK	Standard ($\lambda = 0.03$)	$1.52 \times 10^1 \pm 1.10 \times 10^0$	3.99×10^2	1.77×10^{-1}
ERK	Standard ($\lambda = 0.01$)	$4.63 \times 10^0 \pm 5.12 \times 10^{-1}$	1.20×10^3	2.79×10^{-1}
ERK	Partial ($N_c = 3.0 \times 10^4$)	$4.83 \times 10^2 \pm 1.01 \times 10^1$	3.44×10^{-3}	3.03×10^{-2}
ERK	Partial ($N_c = 1.0 \times 10^4$)	$4.14 \times 10^2 \pm 9.25 \times 10^0$	1.75×10^{-3}	3.17×10^{-2}
ERK	Partial ($N_c = 3.0 \times 10^3$)	$2.82 \times 10^2 \pm 8.42 \times 10^0$	1.47×10^{-2}	3.49×10^{-2}
ERK	Partial ($N_c = 1.0 \times 10^3$)	$9.95 \times 10^1 \pm 4.42 \times 10^0$	1.67×10^{-3}	4.50×10^{-2}
ERK	Partial ($N_c = 3.0 \times 10^2$)	$4.47 \times 10^1 \pm 2.48 \times 10^0$	4.80×10^{-1}	6.94×10^{-2}
Prion	None (Gillespie SSA)	$4.39 \times 10^2 \pm 5.24 \times 10^1$	0.00×10^0	2.03×10^{-2}
Prion	Standard ($\lambda = 0.3$)	$1.25 \times 10^2 \pm 1.38 \times 10^1$	1.22×10^{-1}	3.00×10^{-1}
Prion	Standard ($\lambda = 0.1$)	$2.61 \times 10^1 \pm 7.61 \times 10^{-1}$	2.87×10^{-1}	3.92×10^{-2}
Prion	Standard ($\lambda = 0.03$)	$8.19 \times 10^0 \pm 4.90 \times 10^{-1}$	2.35×10^{-1}	5.73×10^{-2}
Prion	Standard ($\lambda = 0.01$)	$3.08 \times 10^0 \pm 2.68 \times 10^{-1}$	2.39×10^{-1}	5.96×10^{-2}
Prion	Partial ($N_c = 3.0 \times 10^3$)	$6.29 \times 10^2 \pm 6.51 \times 10^0$	1.34×10^{-4}	2.18×10^{-2}
Prion	Partial ($N_c = 1.0 \times 10^3$)	$4.50 \times 10^2 \pm 5.67 \times 10^0$	1.33×10^{-4}	2.79×10^{-2}
Prion	Partial ($N_c = 3.0 \times 10^2$)	$4.18 \times 10^2 \pm 5.56 \times 10^0$	1.32×10^{-4}	3.95×10^{-2}
Prion	Partial ($N_c = 1.0 \times 10^2$)	$3.35 \times 10^2 \pm 1.58 \times 10^1$	1.93×10^{-4}	7.84×10^{-2}
Prion	Partial ($N_c = 3.0 \times 10^1$)	$1.69 \times 10^2 \pm 3.34 \times 10^0$	1.81×10^{-4}	1.08×10^{-1}
Prion	Partial ($N_c = 1.0 \times 10^1$)	$4.76 \times 10^1 \pm 3.57 \times 10^0$	5.44×10^{-4}	1.94×10^{-1}
TCR	None (Gillespie SSA)	$1.04 \times 10^2 \pm 2.27 \times 10^1$	0.00×10^0	3.06×10^{-1}
TCR	Standard ($\lambda = 0.3$)	$2.06 \times 10^1 \pm 3.82 \times 10^0$	2.82×10^{-1}	2.75×10^{-1}
TCR	Standard ($\lambda = 0.1$)	$4.86 \times 10^0 \pm 9.44 \times 10^{-1}$	4.40×10^{-1}	2.29×10^{-1}
TCR	Partial ($N_c = 1.0 \times 10^3$)	$3.62 \times 10^1 \pm 5.02 \times 10^0$	1.19×10^{-3}	3.16×10^{-1}
TCR	Partial ($N_c = 3.0 \times 10^2$)	$1.63 \times 10^1 \pm 1.58 \times 10^0$	1.20×10^{-3}	2.99×10^{-1}
TCR	Partial ($N_c = 1.0 \times 10^2$)	$8.48 \times 10^0 \pm 8.25 \times 10^{-1}$	9.46×10^{-4}	3.10×10^{-1}
TCR	Partial ($N_c = 3.0 \times 10^1$)	$3.23 \times 10^0 \pm 5.42 \times 10^{-1}$	1.63×10^{-3}	3.05×10^{-1}

FIG. 7. Values of $\lambda_r(t)$ in a single stochastic simulation performed using the partial scaling method for every reaction $r \in \{1 \dots R\}$ in the (a) ERK activation model,²⁵ (b) prion protein aggregation model,²⁶ and (c) TCR signaling model.²⁷ In each panel, $N_c = 300$. As can be seen, scaling is adaptive and heterogeneous.

transient shown in Fig. 2(c). The dynamics of this model are inherently stochastic and poorly represented by first moments. As can be seen, standard scaling yields marginalized probability distributions that differ markedly from those obtained from exact simulations (top row), whereas partial scaling yields approximate results that are very close to the exact results (bottom row).

Additional accuracy results, based on the error measures introduced in the Appendix, and timing results for the three benchmark problems are summarized in Fig. 6 and Table I. At least for these problems, partial scaling allows for more significant speed ups than standard scaling with the introduction of much less error in the estimates of first and second moments.

In Fig. 7, for one particular setting for N_c , we have plotted λ_r as a function of time t for every reaction included in each of the three models considered above. These plots show that partial scaling is highly dynamic and multiscale.

VI. IMPLEMENTATION

We added a partial scaling feature to the SSA implemented in BioNetGen,²⁸ an open-source, general-purpose simulation package used by biological modelers. The SSA implemented in BioNetGen is an efficient variation of Gillespie's direct method,² which incorporates various ideas for optimizing simulation efficiency, such as on-the-fly generation (vs pregeneration) of the list of reactions in which (populated) chemical species can participate.^{56,57} BioNetGen is designed for compatibility with models defined using BNGL,²⁴ a language for specifying deterministic and stochastic models for well-mixed (bio)chemical reaction kinetics and for specifying simulations based on such models. BioNetGen also supports models defined using SBML,⁵⁸ such as those available in BioModels Database.⁵⁹ To invoke partial scaling using BioNetGen's `simulate` command, the method argument should be set to `ssa` or `psa` and a new `poplevel` argument, which is used only with partial scaling, should be assigned a non-negative integer value. For accuracy, care should be taken to assign a value much larger than 1. The setting for `poplevel` determines the value of N_c and thereby the aggressiveness

of scaling. Examples of usage are provided in BNGL files that we have made available online.^{44–46} These files define the ERK activation,²⁵ prion protein aggregation,²⁶ and TCR signaling²⁷ models considered in Figs. 3–5. It should be noted that our implementation of partial scaling in BioNetGen (version 2.5 or higher) is a so-called generate-first method, meaning that it requires an enumeration of chemical species and the individual reactions in which these species are able to participate. By contrast, so-called network-free methods do not require an explicit enumeration of chemical species or reactions.⁴

VII. DISCUSSION AND CONCLUSION

The procedure that we have called standard scaling is commonly used for accelerating stochastic simulations of systems with large population sizes. However, this method has limited applicability and it may be challenging to provide an appropriate/optimal setting for the method's one parameter, a universal static scaling factor λ . A simulation speed up can only be attained by sacrificing the ability to calculate unbiased estimates of moments higher than the first, and even first-moment information can be calculated incorrectly with overly aggressive scaling. Furthermore, the speed up that is attainable is constrained by the smallest population size measured over the entire time window of simulation. Thus, in practice, careful application of the method may require trial-and-error numerical experiments to find a suitable best setting for the scaling factor λ . For this reason, the method is especially problematic when one wishes to use it as part of a parameter identification procedure, or any procedure involving parameter variation, because the best λ setting depends in a nonobvious way on parameter values (initial conditions and rate constants).

As we have seen for three nontrivial benchmark problems (Figs. 3–6, Table I), partial scaling significantly outperforms standard scaling in two important ways. First, partial scaling yields greater acceleration for the same or better accuracy (Fig. 6 and Table I). Second, as illustrated in Fig. 5, partial scaling better preserves second-moment information (Table I). Another attractive

feature of partial scaling is the better ability of the user to avoid overly aggressive scaling without any requirement for numerical experiments. With partial scaling, overly aggressive scaling is only possible if the setting for N_c is near or below 1, which is not recommended. By contrast, with standard scaling, any setting for λ must be tested in numerical experiments. As noted earlier, a trial-and-error procedure may be required to find a suitable setting for λ and the suitability of a setting may change with a change in parameter values. Thus, standard scaling seems especially disadvantaged in comparison with partial scaling for simulations within the context of a parameter identification procedure (i.e., a fitting procedure).

Partial scaling is similar to the probability-weighted dynamic Monte Carlo (PW-DMC) method.⁶⁰ In this method, reaction rates are scaled dynamically according to user-specified rules. The rules define scaling factors for rates that fall into specified ranges. Relative to partial scaling, PW-DMC has more parameters: a set of ranges and the associated scaling factors vs partial scaling's single parameter, a critical population size (N_c). Unlike N_c , the PW-DMC parameters must be set through a trial-and-error process. Because rates are scaled without regard to reactant population sizes, PW-DMC can introduce errors where partial scaling does not. These errors arise, for example, when a bimolecular reaction has reactants with disparate population sizes such that one is discrete but the other is large enough to yield a rate that qualifies for scaling. Scaling in this scenario can cause the discrete population to become negative.

We have not directly compared partial scaling against hybrid methods or τ -leaping methods. However, partial scaling has clear advantages over these methods. Partial scaling is far easier to implement than hybrid methods, and partial scaling is more broadly applicable than τ -leaping methods. Partial scaling is useful when population sizes are distributed smoothly across multiple scales (i.e., without clear separation of discrete and continuous populations) because there is no *a priori* requirement for classification of population sizes as either discrete or continuous. Such a requirement can be highly problematic when the classification changes during the time window of simulation, as could be the case for a system that exhibits oscillatory behavior. The chief disadvantage of τ -leaping methods is the lack of guarantee of the existence of a time window τ having the necessary properties. In our experience, hybrid and τ -leaping methods are useful in restricted circumstances. We expect partial scaling to be useful for a wider array of circumstances.

Although partial scaling is an approximate SSA (vs an exact SSA), such methods have important applications. For example, in parameter identification procedures, which typically entail numerous simulation runs, success is highly dependent on simulation efficiency and inexactness is a lesser concern, especially when parameter estimation is based on noisy data. The approximations of partial scaling are very likely to be tolerable, even when statistical distribution data are being used in a fitting procedure. As the optimization algorithm in a fitting procedure converges, numerous samples around a local optimum are inevitably generated. By averaging over a moving window of trial parameter sets, for example, one can, in principle, compensate for the noisy (but unbiased) estimates of first moments when determining goodness of fit.

An attractive future use of partial scaling is acceleration of network-free simulation,⁴ which is an important outstanding problem, as we have discussed elsewhere.⁶¹ One reason to pursue this goal is the high cost of fitting when network-free simulation is necessary,⁶² as is the case for the model of Chylek *et al.*⁶³ Unfortunately, application of partial scaling requires the system state to be tracked in terms of populations and no currently available, general-purpose implementation of a network-free simulation algorithm tracks the system state in this way. Rather, the state is followed in terms of the states of individual biomolecular sites,⁴ which reflects an agent-based or particle-based approach to simulation. However, as discussed by Liu *et al.*,⁶⁴ it is feasible to develop a network-free simulation algorithm in which the system state is tracked in terms of populations, just those of the chemical species with nonzero population sizes. (An exhaustive enumeration of potentially populated chemical species is typically impracticable whenever network-free simulation is under consideration.) Liu *et al.*⁶⁴ presented a specialized algorithm that uses this approach to state tracking, which they termed the full-scale SSA method. However, this method is specialized for a subset of rule-based models—it does not consider systems in which molecules interact to form assemblies (complexes). Thus, an interesting future research direction would be developing a more general version of the full-scale SSA method that is compatible with partial scaling.

ACKNOWLEDGMENTS

We thank Danny Perez and Arthur F. Voter for helpful discussions. We also thank the research group of James R. Faeder, which maintains the BioNetGen code base with support from the National Institute of General Medical Sciences (NIGMS) of the National Institutes of Health (NIH) (Grant No. P41GM103712). W.S.H. acknowledges support from NIGMS/NIH (Grant No. R01GM111510). Y.T.L. and S.F. acknowledge support from the Center for Nonlinear Studies (CNLS). CNLS is funded by the Laboratory-Directed Research and Development program at Los Alamos National Laboratory, which is operated by Triad National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy (Contract No. 89233218CNA000001).

APPENDIX: DEFINITION OF ERROR MEASURES

First-moment errors of scaled simulations per species and overall are reported in Fig. 6 and Table I, respectively. These errors were calculated as described here. Below, we also describe the means used to compare the aggressiveness of standard and partial scaling fairly (via calculation of an effective scaling factor, λ_{eff}) and the means used to characterize second moments obtained from either an unscaled or scaled simulation in terms of a multivariate coefficient of variation.

For simulations based on any of the three models of interest involving scaling (either standard scaling or partial scaling), the mean population of chemical species X_i (for $i = 1, \dots, M$) at report time t was calculated from n_s sample paths. Similarly, on the basis of n_s sample paths, we calculated the corresponding mean obtained via exact simulation. In this way, for either scaling scheme of interest (standard scaling with $0 < \lambda < 1$ or partial scaling with $N_c \gg 1$), we obtain a deviation vector,

$$\overline{\mathbf{N}}^{(\lambda, N_c)}(t) \equiv \overline{\mathbf{N}}^{(\lambda, N_c)}(t) - \overline{\mathbf{N}}^{(0)}(t), \quad (\text{A1})$$

where $\overline{\mathbf{N}}^{(\lambda, N_c)}(t)$ denotes the time-dependent M -dimensional vector of sample path means calculated from scaled simulations and $\overline{\mathbf{N}}^{(0)}(t)$ denotes the time-dependent M -dimensional vector of sample path means calculated from exact simulations. In using the notation $\overline{\mathbf{N}}^{(\lambda, N_c)}$, we adopt the convention that replacement of (λ, N_c) with $(1, \infty)$, (λ, ∞) , and $(1, N_c)$ indicates no scaling, standard scaling, and partial scaling, respectively.

To quantify the accuracy of first moments obtained from scaled simulations, we define the following error measure:

$$\varepsilon^i(t) \equiv \begin{cases} \frac{\tilde{N}_i^{(\lambda, N_c)}(t) - \tilde{N}_i^{(0)}(t)}{\tilde{N}_i^{(0)}(t)}, & \text{if } \tilde{N}_i(t) \geq 1 \\ \tilde{N}_i^{(\lambda, N_c)}(t) - \tilde{N}_i^{(0)}(t), & \text{otherwise.} \end{cases} \quad (\text{A2})$$

Here, $\tilde{N}_i(t)$ denotes the averaged population of species i at time t determined from a set of unscaled simulations. Similarly, $\tilde{N}_i^{(\lambda, N_c)}(t)$ denotes the corresponding quantity determined from an equal number of scaled simulations. We note that ε^i is continuous with respect to $\tilde{N}_i^{(\lambda, N_c)}$ for arbitrary $\tilde{N}_i \geq 0$. The derivative of ε^i with respect to $\tilde{N}_i^{(\lambda, N_c)}$ is also continuous. With the definition given above, error does not diverge to infinity as $\tilde{N}_i$ approaches 0.

For each numerical experiment of interest (aimed at quantifying the errors introduced by scaling), we calculated ε^i for each i at a set of evenly spaced report times, from 0 to T . We then calculated the temporal average $\langle \varepsilon^i \rangle$, which is given by

$$\langle \varepsilon^i \rangle = \frac{\Delta t}{T} \sum_{k=0}^{T/\Delta t} \varepsilon^i(k\Delta t), \quad (\text{A3})$$

where Δt is the spacing between report times. Quantities found using this equation are plotted in Fig. 6.

The aggressiveness of partial scaling is related to N_c , and the aggressiveness of standard scaling is related to λ . These quantities are not directly comparable. To fairly compare the aggressiveness of partial and standard scaling, we define λ_{eff} as the normalized value of κ in a scaled simulation (Algorithm 1, line 14) averaged over all discrete report times. We take the normalization constant to be the temporally averaged value of κ in a corresponding unscaled simulation. Thus, in the case of standard scaling, λ_{eff} is exactly λ . In Fig. 6, we indicate the values of λ_{eff} calculated for partially scaled simulations with different values of N_c .

In Table I, for each numerical experiment, we report a summary statistic for overall error, which is defined as

$$\langle \varepsilon \rangle \equiv \sum_{i=1}^M \langle \varepsilon^i \rangle. \quad (\text{A4})$$

To characterize second moments, we calculated the Van Valen multivariate coefficient of variation⁶⁵ of the M -dimensional vector $\mathbf{N}(t)$ describing the species populations at each report time,

$$CV_{VV}(t) \equiv \sqrt{\frac{\text{tr}(\Sigma(t))}{\|\tilde{\mathbf{N}}(t)\|_2^2}}, \quad (\text{A5})$$

where $\|\tilde{\mathbf{N}}(t)\|_2^2$ is the square of the L^2 norm of the averaged populations and $\Sigma(t)$ is the covariance matrix estimated from the samples. To report a second-moment summary statistic, we averaged over report times,

$$\langle CV_{VV} \rangle = \frac{\Delta t}{T} \sum_{k=0}^{T/\Delta t} CV_{VV}(k\Delta t). \quad (\text{A6})$$

Values of $\langle CV_{VV} \rangle$ are reported in Table I for both scaled and unscaled simulations. Second-moment errors attributable to scaling are reflected in the difference between the values of $\langle CV_{VV} \rangle$ given for a scaled simulation and the corresponding unscaled simulation.

The values given in Table I for $\langle \varepsilon \rangle$ and $\langle CV_{VV} \rangle$ are based on the following choices for Δt , T , and n_s , the number of sample paths. For the ERK activation model,²⁵ $\Delta t = 8.64$ s, $T = 8640$ s, and $n_s = 500$. For the prion protein aggregation model,²⁶ $\Delta t = 0.01$ day, $T = 300$ days, and $n_s = 500$. For the TCR signaling model,²⁷ $\Delta t = 10$ s, $T = 10^4$ s, and $n_s = 10^4$.

REFERENCES

- A. F. Voter, "Introduction to the kinetic Monte Carlo method," in *Radiation Effects in Solids* (Springer, 2007), pp. 1–23.
- D. T. Gillespie, "Stochastic simulation of chemical kinetics," *Annu. Rev. Phys. Chem.* **58**, 35–55 (2007).
- H. H. McAdams and A. Arkin, "It's a noisy business! Genetic regulation at the nanomolar scale," *Trends Genet.* **15**, 65–69 (1999).
- R. Suderman, E. D. Mitra, Y. T. Lin, K. E. Erickson, S. Feng, and W. S. Hlavacek, "Generalizing Gillespie's direct method to enable network-free simulations," *Bull. Math. Biol.* (published online 2018).
- L. A. Chylek, L. A. Harris, C.-S. Tung, J. R. Faeder, C. F. Lopez, and W. S. Hlavacek, "Rule-based modeling: A computational approach for studying biomolecular site dynamics in cell signaling systems," *Wiley Interdiscip. Rev.: Syst. Biol. Med.* **6**, 13–36 (2014).
- J. Yang, M. I. Monine, J. R. Faeder, and W. S. Hlavacek, "Kinetic Monte Carlo method for rule-based modeling of biochemical networks," *Phys. Rev. E* **78**, 031910 (2008).
- M. W. Sneddon, J. R. Faeder, and T. Emonet, "Efficient modeling, simulation and coarse-graining of biological complexity with NFsim," *Nat. Methods* **8**, 177–183 (2011).
- M. S. Creamer, E. C. Stites, M. Aziz, J. A. Cahill, C. W. Tan, M. E. Berens, H. Han, K. J. Bussey, D. D. Von Hoff, W. S. Hlavacek, and R. G. Posner, "Specification, annotation, visualization and simulation of a large rule-based model for ERBB receptor signaling," *BMC Syst. Biol.* **6**, 107 (2012).
- N. G. Van Kampen, *Stochastic Processes in Physics and Chemistry* (Elsevier, 1992), Vol. 1.
- H. Risken, "Fokker-Planck equation," in *The Fokker-Planck Equation* (Springer, 1996), pp. 63–95.
- D. T. Gillespie, "Approximate accelerated stochastic simulation of chemically reacting systems," *J. Chem. Phys.* **115**, 1716–1733 (2001).
- H. Salis and Y. Kaznessis, "Accurate hybrid stochastic simulation of a system of coupled chemical or biochemical reactions," *J. Chem. Phys.* **122**, 054103 (2005).
- J. M. Newby, P. C. Bressloff, and J. P. Keener, "Breakdown of fast-slow analysis in an excitable system with channel noise," *Phys. Rev. Lett.* **111**, 128101 (2013).
- P. Bokes, J. R. King, A. T. Wood, and M. Loose, "Transcriptional bursting diversifies the behaviour of a toggle switch: Hybrid simulation of stochastic gene expression," *Bull. Math. Biol.* **75**, 351–371 (2013).
- P. C. Bressloff and J. M. Newby, "Stochastic hybrid model of spontaneous dendritic NMDA spikes," *Phys. Biol.* **11**, 016006 (2014).

- ¹⁶P. C. Bressloff, "Path-integral methods for analyzing the effects of fluctuations in stochastic hybrid neural networks," *J. Math. Neurosci.* **5**, 4 (2015).
- ¹⁷Y. T. Lin and C. R. Doering, "Gene expression dynamics with stochastic bursts: Construction and exact results for a coarse-grained model," *Phys. Rev. E* **93**, 022409 (2016).
- ¹⁸Y. T. Lin and T. Galla, "Bursting noise in gene expression dynamics: Linking microscopic and mesoscopic models," *J. R. Soc., Interface* **13**, 20150772 (2016).
- ¹⁹P. C. Bressloff, "Stochastic switching in biology: From genotype to phenotype," *J. Phys. A: Math. Theor.* **50**, 133001 (2017).
- ²⁰P. C. Bressloff, "Feynman-Kac formula for stochastic hybrid systems," *Phys. Rev. E* **95**, 012138 (2017).
- ²¹Y. T. Lin and N. E. Buchler, "Efficient analysis of stochastic gene dynamics in the non-adiabatic regime using piecewise deterministic Markov processes," *J. R. Soc., Interface* **15**, 20170804 (2018).
- ²²X. Mao and C. Yuan, *Stochastic Differential Equations with Markovian Switching* (World Scientific, 2006).
- ²³T. B. Kepler and T. C. Elston, "Stochasticity in transcriptional regulation: Origins, consequences, and mathematical representations," *Biophys. J.* **81**, 3116–3136 (2001).
- ²⁴J. R. Faeder, M. L. Blinov, and W. S. Hlavacek, "Rule-based modeling of biochemical systems with BioNetGen," *Methods Mol. Biol.* **500**, 113–167 (2009).
- ²⁵M. Koczańczyk, P. Kocieńewski, E. Kozłowska, J. Jaruszewicz-Błońska, B. Sparta, M. Pargett, J. G. Albeck, W. S. Hlavacek, and T. Lipniacki, "Relaxation oscillations and hierarchy of feedbacks in MAPK signaling," *Sci. Rep.* **7**, 38244 (2017).
- ²⁶R. Rubenstein, P. C. Gray, T. J. Cleland, M. S. Piltch, W. S. Hlavacek, R. M. Roberts, J. Ambrosiano, and J.-I. Kim, "Dynamics of the nucleated polymerization model of prion replication," *Biophys. Chem.* **125**, 360–367 (2007).
- ²⁷T. Lipniacki, B. Hat, J. R. Faeder, and W. S. Hlavacek, "Stochastic effects and bistability in T cell receptor signaling," *J. Theor. Biol.* **254**, 110–122 (2008).
- ²⁸L. A. Harris, J. S. Hogg, J.-J. Tapia, J. A. P. Sekar, S. Gupta, I. Korsunsky, A. Arora, D. Barua, R. P. Sheehan, and J. R. Faeder, "BioNetGen 2.2: Advances in rule-based modeling," *Bioinformatics* **32**, 3366–3368 (2016).
- ²⁹The parameter Ω characterizes system size. Here, we take it to represent volume, but it can alternatively be interpreted as a population that defines a population scale. With this interpretation, (dimensionless) population densities replace concentrations as the state variables in the continuum limit with no change in the mathematical form of the governing equations.
- ³⁰D. T. Gillespie, "Exact stochastic simulation of coupled chemical reactions," *J. Phys. Chem.* **81**, 2340–2361 (1977).
- ³¹Y. Cao, D. T. Gillespie, and L. R. Petzold, "Avoiding negative populations in explicit Poisson tau-leaping," *J. Chem. Phys.* **123**, 054104 (2005).
- ³²Y. Cao, D. T. Gillespie, and L. R. Petzold, "Efficient step size selection for the tau-leaping simulation method," *J. Chem. Phys.* **124**, 044109 (2006).
- ³³H. A. Kramers, "Brownian motion in a field of force and the diffusion model of chemical reactions," *Physica* **7**, 284–304 (1940).
- ³⁴J. Moyal, "Stochastic processes and statistical physics," *J. R. Stat. Soc.: Ser. B* **11**, 150–210 (1949).
- ³⁵C. W. Gardiner *et al.*, *Handbook of Stochastic Methods* (Springer Berlin, 1985), Vol. 3.
- ³⁶C. R. Doering, K. V. Sargsyan, and L. M. Sander, "Extinction times for birth-death processes: Exact results, continuum asymptotics, and the failure of the Fokker-Planck approximation," *Multiscale Model. Simul.* **3**, 283–299 (2005).
- ³⁷Y. T. Lin, H. Kim, and C. R. Doering, "Features of fast living: On the weak selection for longevity in degenerate birth-death processes," *J. Stat. Phys.* **148**, 647–663 (2012).
- ³⁸Y. T. Lin, H. Kim, and C. R. Doering, "Demographic stochasticity and evolution of dispersion I. Spatially homogeneous environments," *J. Math. Biol.* **70**, 647–678 (2015).
- ³⁹Y. T. Lin, H. Kim, and C. R. Doering, "Demographic stochasticity and evolution of dispersion II: Spatially inhomogeneous environments," *J. Math. Biol.* **70**, 679–707 (2015).
- ⁴⁰N. v. Kampen, "A power series expansion of the master equation," *Can. J. Phys.* **39**, 551–567 (1961).
- ⁴¹The symmetry factor is not strictly required in the context of ODE modeling of chemical kinetics (because a constant times a constant is still a constant), but it emerges from the CME description of chemical kinetics.
- ⁴²The limit of $h_r \bar{\kappa}_r \Delta t$, evaluated at time t , as $\Delta t \rightarrow 0$ equals the probability that a reaction r takes place somewhere in the system within a time window of t to $t + \Delta t$. A good approximation of the probability is obtained with finite Δt so long as Δt is small enough such that the probability of two or more reactions of any kind occurring within the time window $[t, t + \Delta t)$ is much smaller than the probability of just one reaction.
- ⁴³R. Schwartz, *Biological Modeling and Simulation: A Survey of Practical Models, Algorithms, and Numerical Methods* (MIT Press, 2008).
- ⁴⁴Y. T. Lin, "A BioNetGen input file for the ERK activation model," https://github.com/RuleWorld/RuleHub/blob/master/Published/Lin2019/ERK_model.bngl; accessed May 8, 2019.
- ⁴⁵Y. T. Lin, "A BioNetGen input file for the prion protein aggregation model," https://github.com/RuleWorld/RuleHub/blob/master/Published/Lin2019/prion_model.bngl; accessed May 8, 2019.
- ⁴⁶Y. T. Lin, "A BioNetGen input file for the TCR signaling model," https://github.com/RuleWorld/RuleHub/blob/master/Published/Lin2019/TCR_model.bngl; accessed May 8, 2019.
- ⁴⁷Terms higher than second order are dropped because of the Pawula theorem,¹⁰ which states that any higher-order truncation ($\mathcal{O}(1/\Omega^2)$) fails to preserve the positivity of p .
- ⁴⁸Let us use XY to denote the population change of a reaction event (in a partially scaled simulation), where X is a possibly random multiplier and $Y \sim \text{Poisson}(\lambda_r \bar{\kappa}_r h_r \Delta t)$. If the multiplier is distributed such that $\mathbb{E}[X] = 1/\lambda_r$ and X and Y are independent, $\mathbb{E}[XY] = h_r \bar{\kappa}_r \Delta t$ (as desired for preservation of first moments) and $\text{var}[XY] = \text{var}[X] \lambda_r h_r \bar{\kappa}_r \Delta t (1 + \lambda_r h_r \bar{\kappa}_r \Delta t) + h_r \bar{\kappa}_r \Delta t / \lambda_r$. Because all terms in this expression are positive, it follows that $\text{var}[XY]$ is least when $\text{var}[X] = 0$, i.e., when X is a deterministic variable. Recall that for two independently distributed random variables X and Y , $\mathbb{E}[XY] = \mathbb{E}[X]\mathbb{E}[Y]$ and $\text{var}[XY] = \text{var}[X] \text{var}[Y] + \text{var}[X] \mathbb{E}^2[Y] + \mathbb{E}^2[X] \text{var}[Y]$.
- ⁴⁹M. A. Gibson and J. Bruck, "Efficient exact stochastic simulation of chemical systems with many species and many channels," *J. Phys. Chem. A* **104**, 1876–1889 (2000).
- ⁵⁰A. M. Walczak, M. Sasai, and P. G. Wolynes, "Self-consistent proteomic field theory of stochastic gene switches," *Biophys. J.* **88**, 828–850 (2005).
- ⁵¹D. M. Roma, R. A. O'Flanagan, A. E. Ruckenstein, A. M. Sengupta, and R. Mukhopadhyay, "Optimal path to epigenetic switching," *Phys. Rev. E* **71**, 011902 (2005).
- ⁵²P. B. Warren and P. R. ten Wolde, "Chemical models of genetic toggle switches," *J. Phys. Chem. B* **109**, 6812–6823 (2005).
- ⁵³M. Assaf, E. Roberts, and Z. Luthey-Schulten, "Determining the stability of genetic switches: Explicitly accounting for mRNA noise," *Phys. Rev. Lett.* **106**, 248102 (2011).
- ⁵⁴M. Strasser, F. J. Theis, and C. Marr, "Stability and multiattractor dynamics of a toggle switch based on a two-stage model of stochastic gene expression," *Biophys. J.* **102**, 19–29 (2012).
- ⁵⁵M. Lu, M. K. Jolly, R. Gomoto, B. Huang, J. Onuchic, and E. Ben-Jacob, "Tristability in cancer-associated microRNA-TF chimera toggle switch," *J. Phys. Chem. B* **117**, 13164–13174 (2013).
- ⁵⁶L. Lok and R. Brent, "Automatic generation of cellular reaction networks with MolecuLizer 1.0," *Nat. Biotechnol.* **23**, 131–136 (2005).
- ⁵⁷J. R. Faeder, M. L. Blinov, B. Goldstein, and W. S. Hlavacek, "Rule-based modeling of biochemical networks," *Complexity* **10**, 22–41 (2005).
- ⁵⁸M. Hucka, F. T. Bergmann, A. Dräger, S. Hoops, S. M. Keating, N. Le Novère, C. J. Myers, B. G. Olivier, S. Sahle, J. C. Schaff, L. P. Smith, D. Waltemath, and D. J. Wilkinson, "The systems biology markup language (SBML): Language specification for level 3 version 2 core," *J. Integr. Bioinf.* **15**, 20170081 (2018).

- ⁵⁹V. Chelliah, C. Laibe, and N. Le Novère, "BioModels database: A repository of mathematical models of biological processes," *Methods Mol. Biol.* **1021**, 188–199 (2013).
- ⁶⁰H. Resat, H. S. Wiley, and D. A. Dixon, "Probability-weighted dynamic Monte Carlo method for reaction kinetics simulations," *J. Phys. Chem. B* **105**, 11026–11034 (2001).
- ⁶¹Y. Lin, L. A. Chylek, N. W. Lemons, and W. S. Hlavacek, "Using equation-free computation to accelerate network-free stochastic simulation of chemical kinetics," *J. Phys. Chem. B* **122**, 6351–6356 (2018).
- ⁶²B. R. Thomas, L. A. Chylek, J. Colvin, S. Sirimulla, A. H. A. Clayton, W. S. Hlavacek, and R. G. Posner, "BioNetFit: A fitting tool compatible with BioNetGen, NFsim and distributed computing environments," *Bioinformatics* **32**, 798–800 (2016).
- ⁶³L. A. Chylek, V. Akimov, J. Dengjel, K. T. G. Rigbolt, B. Hu, W. S. Hlavacek, and B. Blagoev, "Phosphorylation site dynamics of early T-cell receptor signaling," *PLoS One* **9**, e104240 (2014).
- ⁶⁴Z. Liu, U. J. Mobassera, C. A. Shaffer, L. T. Watson, and Y. Cao, "Multistate modeling and simulation for regulatory networks," in *Proceedings of the 2010 Winter Simulation Conference*, edited by B. Johansson, S. Jain, J. Montoya-Torres, J. Hugan, and E. Yücesan (IEEE, 2010), pp. 631–642.
- ⁶⁵L. Van Valen, "Multivariate structural statistics in natural history," *J. Theor. Biol.* **45**, 235–247 (1974).