A constrained approach to multiscale stochastic simulation of chemically reacting systems

Simon L. Cotter, Konstantinos C. Zygalakis, Ioannis G. Kevrekidis, and Radek Erban

Citation: *J. Chem. Phys.* **135**, 094102 (2011); doi: 10.1063/1.3624333

View online: <http://dx.doi.org/10.1063/1.3624333>

View Table of Contents: <http://jcp.aip.org/resource/1/JCPSA6/v135/i9>

Published by the [American Institute of Physics](#).

Additional information on J. Chem. Phys.

Journal Homepage: <http://jcp.aip.org/>

Journal Information: http://jcp.aip.org/about/about_the_journal

Top downloads: http://jcp.aip.org/features/most_downloaded

Information for Authors: <http://jcp.aip.org/authors>

ADVERTISEMENT

**ACCELERATE COMPUTATIONAL CHEMISTRY BY 5X.
TRY IT ON A FREE, REMOTELY-HOSTED CLUSTER.**

[LEARN MORE](#)

A constrained approach to multiscale stochastic simulation of chemically reacting systems

Simon L. Cotter,^{1,a)} Konstantinos C. Zygalakis,¹ Ioannis G. Kevrekidis,²
and Radek Erban^{1,b)}

¹*Mathematical Institute, University of Oxford, 24-29 St. Giles', Oxford OX1 3LB, United Kingdom*

²*Department of Chemical and Biological Engineering and Program in Applied and Computational Mathematics, Princeton University, Princeton, New Jersey 08544, USA*

(Received 6 April 2011; accepted 18 July 2011; published online 1 September 2011)

Stochastic simulation of coupled chemical reactions is often computationally intensive, especially if a chemical system contains reactions occurring on different time scales. In this paper, we introduce a multiscale methodology suitable to address this problem, assuming that the evolution of the slow species in the system is well approximated by a Langevin process. It is based on the conditional stochastic simulation algorithm (CSSA) which samples from the conditional distribution of the suitably defined fast variables, given values for the slow variables. In the constrained multiscale algorithm (CMA) a single realization of the CSSA is then used for each value of the slow variable to approximate the effective drift and diffusion terms, in a similar manner to the constrained mean-force computations in other applications such as molecular dynamics. We then show how using the ensuing Fokker-Planck equation approximation, we can in turn approximate average switching times in stochastic chemical systems. © 2011 American Institute of Physics. [doi:10.1063/1.3624333]

I. INTRODUCTION

Stochastic simulation algorithms (SSAs) for systems of chemically reacting species have recently been used in the modelling of a number of biological systems, including the phage λ lysis-lysogeny decision circuit,¹ circadian rhythms,² and the cell cycle.³ The most widely used SSA was developed by Gillespie⁴ in 1970s. It simulates every single chemical reaction, and samples exactly from the solution to the corresponding chemical master equation. In order to characterise the behaviour of the system, one has to simulate a large number of events and this can become computationally intensive. For suitable classes of chemically reacting systems, one can sometimes use exact algorithms which are equivalent to the Gillespie SSA, and which are less computationally intensive. Examples include the Gibson-Bruck SSA⁵ and the optimized direct method.⁶ Another approach, which does not introduce errors in simulations, is the use of parallel computing, for example, on graphics processing units.^{7,8}

The SSAs mentioned so far are equivalent to the Gillespie SSA, i.e., they stochastically simulate the occurrence of every chemical reaction. This can be computationally very expensive, especially for systems involving abundant species. Also, often a modeller is not interested in every single reaction which takes place in the system. There can be multiple time scales, meaning that one has to simulate a large number of fast reactions, while we are interested in a more slowly evolving quantity.⁹

One approach to accelerate the Gillespie SSA is the τ -leap method.¹⁰ In this algorithm the number of reactions

over a specified time increment τ is sampled from a Poisson distribution for each reaction channel. The problem with using Poisson random variables, however, is that it is possible for a species' population to become negative. One possible way around this problem is the binomial τ -leap algorithm^{11,12} which replaces the Poisson distribution with a binomial. Another way to tackle this problem is the so-called R -leap method,¹³ which fixes the number of reaction firings, instead of the time interval τ , as in the τ -leap method.

Alternatively one can take a coarse grained approach where instead of treating the molecular populations as discrete random variables one describes them in terms of real numbers that change continuously. This can be achieved by the chemical Langevin equation,¹⁴ which is a stochastic differential equation (SDE) acting like a bridge between the discrete stochastic simulation algorithm and the deterministic reaction rate equations. This approach, and its recent reformulations,¹⁵ may be less computationally intensive than the exact method. However, for some chemical systems, the CLE might lead to negative populations because its solutions have a small but positive probability of becoming negative.^{16,17}

The two different approaches described above do not take explicitly into account the (possibly existing) separation of time scales. To explain the algorithms that do, we first formulate the multiscale problem in mathematical terms.

Let us consider a well-mixed chemical system, comprised of N different chemical species $\{X_1, \dots, X_N\}$ interacting through M reaction channels $\{R_1, \dots, R_M\}$. The state of the system is given by $\mathbf{X}(t) = [X_1(t), \dots, X_N(t)]$ where $X_i(t)$, $i = 1, \dots, N$, is the number of molecules of type X_i in the system at time t , where for simplicity we denote both the chemicals and the number of molecules of that chemical by the same symbol.

^{a)}Electronic mail: cotter@maths.ox.ac.uk.

^{b)}Electronic mail: erban@maths.ox.ac.uk.

Following Cao *et al.*,¹⁸ we assume that the reactions can be divided into two groups; *fast reactions* and *slow reactions*. Fast reactions are those which occur many times on a time scale for which it is very unlikely that the slow reactions will occur at all. An example of such a system with slow and fast reactions is given in Sec. II. Cao *et al.*¹⁸ define a *fast species* to be any species whose population gets changed by some fast reaction. Otherwise they are defined to be *slow species*. In this paper, we denote by **F** the fast variables and by **S** the slow variables, where we allow the components of the slow variables to be a linear combination of the species present in the system. Together, the components of the fast and slow species can be used to form a basis for the state space of the system, i.e., the dimension of this basis is equal to the number of linearly independent species (taking into account possible conservation relations).

Algorithms for systems with fast and slow species have been developed in the literature by a number of authors. Haseltine and Rawlings¹⁹ choose to simulate the fast reactions using Langevin dynamics while simulating the slow reactions using the Gillespie algorithm. This approach requires not only the time scale separation but also the volume of the system to be sufficiently large. However, this can be avoided by using the probability densities of the fast species conditioned on the slow species to estimate the effective propensity functions of the slow species.^{18,20,21} Other types of solvers have also been put forward and analysed for differing time scales of interest.^{22,23}

The above approaches build algorithms which can simulate the evolution of slow variables while avoiding explicit simulation of all fast reactions. Another approach to multi-scale computation is to estimate the probability distribution of slow variables.⁹ The main idea is to use short bursts of appropriately initialised stochastic simulations to estimate coefficients of unavailable macroscopic coarse-grained equations.²⁴ This approach has been successfully applied to a number of applications, including models in materials science,²⁵ cell motility²⁶ and social insect behaviour.²⁷ In the case of chemical systems, short bursts of stochastic simulations are used to estimate drift and diffusion coefficients for an approximating Fokker-Planck equation written in terms of the slow variables.⁹

Constrained dynamics approaches have been used to approximate the drift term in other scenarios, particularly with regards to stochastic differential equations relating to molecular dynamics. “Umbrella” sampling^{28,29} in molecular dynamics constrains the system by adding deep potential wells to the dynamics of the variables so that the system never moves far away from its target value. “Blue moon” sampling³⁰ has also been developed to sample rare or long-time events. Naive constraining of dynamics in the molecular dynamics scenario can lead to bias, and so different numerical schemes for the simulation of constrained SDEs have also been proposed.³¹

In this paper, we introduce the conditional stochastic simulation algorithm (CSSA). This is an algorithm that allows us to sample from the distribution of the fast variables conditioned on the slow ones. The effective dynamics of the slow process can then be approximated using the conditional distribution either as a Markov jump process^{18,20,21} or as an SDE.⁹

Our approach is a coarse-grained one, i.e., we choose to approach the effective dynamics as an SDE. We estimate the coefficients of the effective SDE on the fly during the sampling of the conditional distribution in the spirit of other multiscale approaches.²²

The paper is organised as follows. In Sec. II, we introduce an illustrative example which is used throughout the paper. The conditional stochastic simulation algorithm is presented in Sec. III. In Sec. IV, we introduce the constrained multi-scale algorithm (CMA), based on the CSSA, allowing us to constrain the Gillespie algorithm to one particular state of the slow variables. At the same time, the algorithm estimates the drift and diffusion terms of that slow variable, analogously to the constrained mean force calculations that some have used in molecular dynamics simulations.³² We show that it is sufficient to use a single realization of the CSSA for each set of values of the slow variable **S** for which we wish to estimate the drift and diffusion coefficients. Illustrative numerical results for the example from Sec. II are presented in Sec. V. In Sec. VI, we show how these methods can be used to explore the change in the dynamics of the system with variation in the parameters, including the probability density and the average switching time in a bistable system. Finally, in Sec. VII, we briefly discuss issues regarding the length of our simulations with respect to the accuracy in our estimates of the effective dynamics of the slow variable.

II. AN ILLUSTRATIVE EXAMPLE

Our goal is to introduce a general multiscale algorithm for stochastic simulation of N chemical species X_i , $i = 1, 2, \dots, N$, which are subject to M reaction channels R_j , $j = 1, 2, \dots, M$. We will explain it using the following illustrative example:

In what follows, we label the reactions $R_1, \dots, R_5$ corresponding to the reaction rate subscripts. Each reaction R_j has associated with it a propensity function⁴ $\alpha_j(t)$. It is defined so that the product $\alpha_j(t)dt$ is the probability that the j th reaction occurs in the infinitesimally small time interval $[t, t + dt)$. In the case of model (1), the propensity functions are defined by

$$\begin{aligned} \alpha_1(t) &= k_1 V, & \alpha_2(t) &= k_2 X_1(t), & \alpha_3(t) &= k_3 X_2(t), \\ \alpha_4(t) &= k_4 X_2(t), & \alpha_5(t) &= k_5 X_3(t), \end{aligned}$$

where V is the volume of the reactor. Depending on the choice of parameters, this system can have several separated time scales. We consider this system with the following set of non-dimensionalised parameters:

$$k_1 V = 100, \quad k_2 = k_3 = 200, \quad k_4 = k_5 = 1. \quad (2)$$

Note that the time variable is dimensionless. In Figure 1, we plot the time evolution of the different species of system (1), along with $(X_1 + X_2)/2$, starting from $X_1(0) = X_2(0) = X_3(0) = 100$. We see that X_1 and X_2 are changing quickly, while X_3 is changing slowly. However, note that if we consider the variable $X_1 + X_2$, this too is changing

FIG. 1. Sample trajectories of all three of the chemical species from the system (1). Parameters are given by Eq. (2).

slowly. In this situation, reactions R_2 and R_3 are occurring on a faster time scale than the reactions R_1 , R_4 , and R_5 , as can be seen in Table I, which shows the relative number of occurrences of each reaction. Therefore, R_1 , R_4 , and R_5 are the slow reactions, and R_2 and R_3 are the fast reactions.¹⁸ In Table II, we show the number of changes of each of the listed variables during a standard Gillespie simulation of the system up to time $T = 10^4$.

On average, during this run, X_1 and X_2 changed over 200 times more often than either X_3 or $X_1 + X_2$. Given such a separation of time scales, slow variables are linear combinations of the chemical species, which are invariant with respect to the fast reactions. In this simple first-order chemical system, the two quantities X_3 and $X_1 + X_2$ are both invariant with respect to R_2 and R_3 , and so are slow quantities. X_1 and X_2 are fast variables, as they change with every occurrence of reaction R_2 or R_3 .

There are several different choices of bases of slow and fast variables for the system. In what follows we consider the following three examples:

$$S = [X_3], \quad \mathbf{F} = [X_1, X_2]; \quad (3)$$

$$S = [X_1 + X_2, X_3], \quad \mathbf{F} = [X_1]; \quad (4)$$

$$S = [X_1 + X_2, X_3], \quad \mathbf{F} = [X_2]. \quad (5)$$

Notice that these systems have a different number of dimensions in the slow variable. In general, identifying which variables are slow and which are fast can be a challenging task in itself, as discussed in Sec. VIII.

TABLE I. Percentages of occurrences of each reaction during a standard Gillespie simulation up to time $T = 10^4$ of the chemical system (1) with parameters given by (2).

Reaction	R_1	R_2	R_3	R_4	R_5
Frequency	0.25%	49.75%	49.50%	0.25%	0.25%

TABLE II. Numbers of changes of each of the listed variables during a standard Gillespie simulation up to time $T = 10^4$ of the chemical system (1) with parameters given by (2).

Variable	X_1	X_2	X_3	$X_1 + X_2$
Changes	4.017×10^8	4.017×10^8	1.997×10^6	1.997×10^6

A. Invariant measure of the illustrative example (1)

Since the chemical system (1) is comprised only of mono-molecular reactions, we can find explicitly the invariant measure of this system.³³ This result will be useful in Sec. V to demonstrate the accuracy of the CSSA and CMA. The invariant probability distribution on all of the variables is a multivariate Poisson distribution:

$$\mathbb{P}(X_1 = n_1, X_2 = n_2, X_3 = n_3) \quad (6)$$

$$= \frac{\bar{\lambda}_1^{n_1} \bar{\lambda}_2^{n_2} \bar{\lambda}_3^{n_3}}{n_1! n_2! n_3!} \exp[-(\bar{\lambda}_1 + \bar{\lambda}_2 + \bar{\lambda}_3)],$$

where, using Eq. (2),

$$\bar{\lambda}_1 = \frac{k_1 V(k_3 + k_4)}{k_2 k_4} = 100.5,$$

$$\bar{\lambda}_2 = \frac{k_1 V}{k_4} = 100, \quad \bar{\lambda}_3 = \frac{k_1 V}{k_5} = 100.$$

III. CONDITIONAL STOCHASTIC SIMULATION ALGORITHM

In this section, we introduce the conditional stochastic simulation algorithm, an algorithm which samples from the conditional distribution $\mathbb{P}(\mathbf{F}|\mathbf{S} = \mathbf{s})$. This is the distribution of the fast variables $\mathbf{F}$ given a value $\mathbf{s}$ of the slow variables $\mathbf{S}$. This distribution is important for integrating correctly the effective dynamics of the slow variables, which is the ultimate aim of this paper. Given a state of the slow variables $\mathbf{S} = \mathbf{s}$, the CSSA is presented in Table III. Initially we set the slow variables $\mathbf{S}$ to the value $\mathbf{s}$ for which we wish to sample $\mathbb{P}(\mathbf{F}|\mathbf{S} = \mathbf{s})$. We also initialise the fast variable $\mathbf{F} = \mathbf{f}$. Ideally we would take $\mathbf{f}$ to be a sample from $\mathbb{P}(\mathbf{F}|\mathbf{S} = \mathbf{s})$ itself. However, since this is the distribution we are actually interested in characterising, we can instead repeatedly run steps [1]–[6] until we are satisfied that the Markov chain has entered equilibrium and is drawing samples from $\mathbb{P}(\mathbf{F}|\mathbf{S} = \mathbf{s})$.

TABLE III. Conditional stochastic simulation algorithm (CSSA).

[1]	Calculate propensity functions $\alpha_i(t)$, $i = 1, 2, \dots, M$.
[2]	Waiting time to the next reaction is given by Eq. (7).
[3]	Choose one $j \in \{1, 2, \dots, M\}$, with probability α_j/α_0 , and perform reaction R_j .
[4]	If $\mathbf{S} \neq \mathbf{s}$ due to reaction j occurring, then reset $\mathbf{S} = \mathbf{s}$ while not changing the value of $\mathbf{F}$.
[5]	If $X_i < 0$ for any i , then revert to the state of the system before the reaction j occurred.
[6]	Continue with step [1] with time $t = t + \tau$.

In step [1], as in the original Gillespie SSA, we calculate the values of the propensity functions. As the algorithm progresses, it is not necessary to recalculate all of these, as not all of them will change after each reaction has occurred. We can make use of the system's dependency graph⁵ in order to ascertain which propensities are changed by each reaction. Steps [2] and [3] are the same as the original Gillespie algorithm. The waiting time to the next reaction is given by the following formula:

$$\tau = -\frac{\log(u)}{\alpha_0}, \quad \text{where} \quad \alpha_0 = \sum_{i=1}^M \alpha_i(t), \quad (7)$$

and $u \sim U(0, 1)$ is a sample from the uniform distribution on the unit interval $(0, 1)$. The reaction R_j is chosen in step [3] using another $U(0, 1)$ random variable.

Step [4] is where this algorithm diverges from the original Gillespie SSA. Here, if a reaction has occurred which changes the value of the slow variable $\mathbf{S}$, then we reset $\mathbf{S}$, thus constraining it to the desired value. (In some special cases, this resetting may result in non integer values for some of the chemical species, but this is not within the class of problems we consider in this paper.) Note that there is a unique way in how the reset is made, since the components of the variables $\mathbf{S}$ and $\mathbf{F}$ together form a basis of the state space, and we reset $\mathbf{S}$ in such a way that $\mathbf{F}$ is not altered.

Since we have effectively altered the stoichiometric matrix so that the slow variable cannot change, we are not always guaranteed to naturally preserve the positivity of all of the chemical reactants, in the case where one of the slow variables is a linear combination of two or more chemical species. Step [5] ensures that the states with negative copy numbers of one/more chemical species have zero probability mass. Without this step, the CSSA can run into difficulties if the reset of $\mathbf{S}$ causes one of the species to have a negative copy number. This step helps to inform our choice of fast variables, since some choices can naturally preserve this non-negativity condition without step [5], as we show in Sec. III A.

Step [5], on the face of it, could appear to cause a small amount of bias away from the true distribution. However, this entire approach is only valid for systems for which the slow dynamics of $\mathbf{S}$ can be accurately approximated by an effective drift-diffusion process. This is only the case if the chances of approaching low populations are small. Therefore, any bias which is incurred in this region due to step [5] will be minimal.

A. Choice of fast and slow variables for the CSSA

In Sec. II, we identified several candidates of fast and slow variables for the system (1). Let us consider the choice of variables (4). In this case, if reaction R_1 was to occur, the change in $\mathbf{S}$ is $[1, 0]$. Since we must not alter the fast variable in step [4] of the CSSA, we must therefore reduce X_2 by one to reset $\mathbf{S}$. However, it is possible for reaction R_1 to occur when $X_2 = 0$, meaning that in this case X_2 could go below zero.

In many scenarios, however, these problems can be overcome by a better choice of fast variables, meaning that step [5]

of the CSSA is no longer necessary. If we choose setup (5), then we are in just such a scenario, and this is certainly preferable. If, however, there was also a source reaction for X_2 (i.e., $\emptyset \rightarrow X_2$), then no such choice of fast variables is possible, and step [5] becomes necessary.

IV. CONSTRAINED MULTISCALE ALGORITHM

The CSSA described in Sec. III is able to sample from the probability distribution $\mathbb{P}(\mathbf{F}|\mathbf{S} = \mathbf{s})$. However, a modeller is often more interested in the effective dynamics of the slow variable $\mathbf{S}$. In this instance, we will attempt to approximate these dynamics by a continuous stochastic differential equation, which has corresponding Fokker-Planck equation given by

$$\frac{\partial p}{\partial t}(\mathbf{s}, t) = \nabla \cdot (\nabla[D(\mathbf{s}) p(\mathbf{s}, t)] - V(\mathbf{s}) p(\mathbf{s}, t)). \quad (8)$$

The challenge now is to estimate the effective drift V and diffusion D of $\mathbf{S}$. In order to estimate this we can run the CSSA for a range of values of the slow variable. For each value $\mathbf{s}$ of $\mathbf{S}$, we use a CSSA simulation of length $T(\mathbf{s})$ to compute the following quantities:

$$V(\mathbf{s}) = \frac{1}{T(\mathbf{s})} \sum_{m=1}^{Q(T(\mathbf{s}))} \delta \mathbf{S}_m, \quad (9)$$

$$D(\mathbf{s}) = \frac{1}{2T(\mathbf{s})} \sum_{m=1}^{Q(T(\mathbf{s}))} \delta \mathbf{S}_m \delta \mathbf{S}_m^T, \quad (10)$$

where $\delta \mathbf{S}_m$ is the change in $\mathbf{S}$ due to the m th iteration of the CSSA before the reset is made in step [4], $T(\mathbf{s})$ is the chosen length of CSSA simulation, and $Q(T(\mathbf{s}))$ is the number of iterations of the CSSA that are made up to time $T(\mathbf{s})$. These quantities are estimators for the effective drift V and diffusion D of the slow variable $\mathbf{S}$. The formulae (9) and (10) are derived in the Appendix.

Having estimated $V(\mathbf{s})$ and $D(\mathbf{s})$ by Eqs. (9) and (10), we compute the approximation π_{CMA} of the invariant measure of the slow variables by solving

$$\nabla \cdot (\nabla[D(\mathbf{s}) \pi_{\text{CMA}}(\mathbf{s})] - V(\mathbf{s}) \pi_{\text{CMA}}(\mathbf{s})) = 0, \quad (11)$$

which is the steady state equation corresponding to Eq. (8). Equation (11) is solved by standard numerical approaches.³⁴ It is worth noting that for a one-dimensional slow variable, the solution to Eq. (11) can also be given analytically.

Since the calculation [Eqs. (9) and (10)] of the drift and diffusion terms boil down to two different statistics on the size/number of resets of $\mathbf{S}$, the additional cost to estimate these quantities is negligible in comparison to the cost of running the CSSA. Furthermore, every single reaction that occurs which changes $\mathbf{S}$ (and therefore requires a reset of this variable in step [4] of the CSSA) has an effect on the final output, so that we are garnering all the possible information that is available from the stochastic simulation.

The length of simulation $T(\mathbf{s})$ is in general state dependent. We will discuss this state dependence in Sec. VII. Another issue which will have an impact on the efficiency of the method is the choice of initialisation of the fast variable in

step [1] of the CSSA. One might expect, for well behaved systems, the distributions $\mathbb{P}(\mathbf{F}|\mathbf{S} = \mathbf{s})$ and $\mathbb{P}(\mathbf{F}|\mathbf{S} = \mathbf{s} + \boldsymbol{\varepsilon})$ would be very similar for small $\boldsymbol{\varepsilon} \in \mathbb{Z}^d$, where $d \in \mathbb{N}$ is the dimension of the slow variable $\mathbf{S}$. Therefore, once we have characterised the distribution $\mathbb{P}(\mathbf{F}|\mathbf{S} = \mathbf{s})$ (using, for example, an initial condition $\mathbf{F}$ from a “burn-in”³⁵), we could use a sample from this distribution to initialise the fast variable for the CSSA for the neighbouring values of $\mathbf{s}$. This is in some sense similar to the numerical continuation methods used to compute steady states branches in bifurcation analysis.

V. APPLICATION OF CMA TO MODEL (1)

In this section, we apply the CSSA and CMA to the illustrative example (1) presented in Sec. II. First, we show that the CSSA is sampling correctly from the conditional distribution $\mathbb{P}(\mathbf{F}|\mathbf{S} = \mathbf{s})$. Since we have the explicit analytical steady state solution (6) of the chemical master equation, we can check the answers that we get using the CSSA (to ascertain the distribution of the fast variables). In particular, we have an explicit formula for $\mathbb{P}(\mathbf{F}|\mathbf{S} = \mathbf{s})$ which can be used in the error computation defined by

$$\text{Error} = \sum_{\mathbf{f} \in \mathbb{N}^{3-d}} |\mathbb{P}(\mathbf{F} = \mathbf{f}|\mathbf{S} = \mathbf{s}) - \nu_{\text{CSSA}}(\mathbf{f})|, \quad (12)$$

where ν_{CSSA} is the conditional probability distribution estimated using the CSSA, and $d \in \mathbb{N}$ is the dimension of the slow variable vector $\mathbf{S}$.

In Figure 2, we present estimations of error of the CSSA algorithm for two choices of fast and slow variables given by Eqs. (3) and (5). We run simulations for a fixed value of $T(\mathbf{s})$ and compute the error using Eq. (12). As can be seen in comparison with the black dashed line (which has a decay proportional to $1/\sqrt{T}$), both of the choices of slow variable leave us with error which is approximately $\mathcal{O}(1/\sqrt{T})$. This is to be expected, since this is the rate of decay of the standard Monte Carlo error. However, the error incurred by using Eq. (3) is consistently larger than that for Eq. (5). This is because there exists an alternative basis to $[X_1, X_2]$, given by $[X_1, X_1 + X_2]$, which has a slow and a fast variable. The speed of the convergence of the algorithm is dependent on the speed of the slowest of the fast variables. Including a slow variable such as $X_1 + X_2$ in the span of the fast variables leads to slower convergence, and therefore more error for the same length of simulation $T(\mathbf{s})$.

Next, we apply the CMA to the illustrative model (1), for both of the choices for the fast and slow variables, given by Eqs. (3) and (5).

A. One-dimensional slow variable: Choice of variables (3)

The slower rate of convergence of the CSSA shown in Figure 2 due to the poor choice of slow/fast variables will naturally have a knock-on effect on the rate of convergence of the CMA approximation of the effective diffusion of the slow variable $S = X_3$. However, we can still accurately approximate the effective dynamics of the slow variable.

In Figure 3, we present π_{CMA} computed by Eq. (11). In this case, the time $T(\mathbf{s})$ is defined implicitly as the time it takes

FIG. 2. Error curve of the approximation of $\mathbb{P}(\mathbf{F}|\mathbf{S} = 100)$ by the CSSA with both choices of variables given by Eq. (3) (blue diamonds) and Eq. (5) (red squares). Black dashed line denotes decay proportional to $1/\sqrt{T}$.

for L changes to occur in S for a given value of s , where we choose $L = 10^5$. Since we have expression (6) for the whole of the probability mass function, we can integrate out the fast variable $\mathbf{F} = [X_1, X_2]$, leaving us with the marginal probability distribution of the slow variable $S = X_3$, given by

$$\begin{aligned} \mathbb{P}(S = s) &= \sum_{f_1, f_2=0}^{\infty} \mathbb{P}(X_1 = f_1, X_2 = f_2, X_3 = s) \\ &= \frac{100^s}{s!} \exp(-100). \end{aligned} \quad (13)$$

This formula is plotted in Figure 3 as the blue histogram, and the visual comparison is excellent.

In Figure 3, we defined the time $T(\mathbf{s})$ implicitly by counting up to L changes in S . In Figure 4, we show how the results vary with L . We plot the error of this method as a function of L . We define the error to be

$$\text{Error} = \sum_{s \in \mathbb{N}} |\mathbb{P}(S = s) - \pi_{\text{CMA}}(s)|. \quad (14)$$

FIG. 3. Approximation π_{CMA} of the probability density of $S = X_3$ by the CMA. Drift and diffusion functions are estimated from CMA runs with 10^5 changes in S for each s . The blue histogram is an exact solution of the stationary chemical master equation, given by Eq. (13).

FIG. 4. Error of the CMA, given by Eq. (14), as a function of simulation length L . Black dashed line denotes decay proportional to $1/\sqrt{L}$.

Figure 4 shows that the CMA is approximately converging at a rate $\mathcal{O}(1/\sqrt{L})$, as in the CSSA examples earlier in this section. We will return to this issue in more detail in Sec. VII, where we look at how many changes in S are required for a given accuracy of the CMA as a function of the variable S .

B. Two-dimensional slow variable: Choice of variables (5)

In this case, the slow variable is two dimensional, which means that Eq. (11) must be approximated using numerical methods, such as finite elements.³⁴ Figure 5 shows π_{CMA} computed by solving Eq. (11) in domain $[0, 400] \times [0, 200]$ with no flux boundary conditions. Here we use a fixed length of simulation $T(s) = 10^3$ for each value $s = [s_1, s_2]$ of S . The resulting two-dimensional probability density function is indistinguishable from the steady state solution of the chemical master equation, given by

$$\mathbb{P}(X_1 + X_2 = s_1, X_3 = s_2) = \sum_{a=0}^{s_1} \frac{100^{s_1+s_2} 201^a}{a!(s_1-a)!s_2!200^a} \exp(-300.5). \quad (15)$$

FIG. 5. Approximation π_{CMA} of $\mathbb{P}(S = [s_1, s_2] = [X_1 + X_2, X_3])$ by the CMA where we used CSSA simulations of length $T(s) \equiv 10^3$.

The disadvantage of this setup (5) in comparison with setup (3), is that if we wish to retain the same grid size for our approximation of the functions D and V as we did before, then the number of points for which we need to run the CSSA is squared. This is alleviated by the faster convergence for each value s of S as demonstrated in Figure 2.

One might also ask what would happen to the accuracy of the algorithm if the gap in time scales between the slow and fast variables was reduced, for instance by reducing the values of the parameters k_2, k_3 . In this case, the method is only accurate if the effective dynamics of the “slow” variables can still be well approximated by an effective drift-diffusion process. If this approximation is appropriate, then any change in reaction parameters will only affect the rate of convergence of the estimates of the drift and diffusion terms.

VI. STOCHASTIC “BIFURCATION DIAGRAMS”

In this section, we show how the CMA can aid us in understanding how the behaviour of the slow variables change as we alter certain system parameters. We will illustrate this using a simple bistable chemical system written in terms of two chemical species X_1 and X_2 :

As before, we denote the reactions $R_1, R_2, \dots, R_6$ according to the reaction rate subscripts. The propensity functions of reactions (16) are given by

$$\begin{aligned} \alpha_1(t) &= k_1 X_2(t), & \alpha_2(t) &= \frac{k_2 X_1(t) X_2(t)}{V}, \\ \alpha_3(t) &= k_3 V, & \alpha_4(t) &= k_4 X_1(t), \\ \alpha_5(t) &= \frac{k_5 X_1(t)(X_1(t) - 1)}{V}, & \alpha_6(t) &= k_6 X_2(t), \end{aligned}$$

where V is the system volume. In Figure 6, we show an illustrative trajectory of the system (16) computed with the Gillespie algorithm, for dimensionless parameters:

$$\begin{aligned} k_1 &= 32, & \frac{k_2}{V} &= 0.04, & k_3 V &= 1475, \\ k_4 &= 19.75, & \frac{k_5}{V} &= 10, & k_6 &= 4000. \end{aligned} \quad (17)$$

It is clear from these trajectories that the system has two favourable states around $[X_1, X_2] = [142, 50]$ and $[X_1, X_2] = [393, 386]$. In the parameter regime (17), reactions R_5 and R_6 are occurring on a much faster time scale than the other reactions. This leads us to the choice of slow variable $S = X_1 + 2X_2$ which is invariant with respect to these reactions, as confirmed in Figure 7. Choosing $F = X_2$ means that the CSSA naturally preserves the non-negativity of all of the chemical species of the system, i.e., step [5] of the CSSA in Table III is not needed.

A. CMA for the bistable example

In Figure 8, we plot π_{CMA} computed for the parameter values (17), with CSSA simulations of length $T(s)$

FIG. 6. Trajectories of the chemical system (16), showing switching between two favourable states, where the system parameters are given by Eq. (17).

$= 10^3$ for each value s of the slow variable $S = X_1 + 2X_2$. To demonstrate the accuracy of the CMA, we can compare π_{CMA} with the solution to the stationary chemical master equation. This can be approximated by making the assumption that given a big enough domain $\Omega = \{[X_1, X_2] \in [0, U_1] \times [0, U_2]\}$, the probability mass of the invariant distribution outside of Ω is negligible. Let us denote by $p_{n,m}(t)$ the probability that $[X_1, X_2] = [n, m]$ at time t . It evolves according to the chemical master equation for system (16), which is given by

$$\begin{aligned} \frac{dp_{n,m}}{dt} = & [k_1 m + k_3 V] p_{n-1,m} + k_6(m+1) p_{n-2,m+1} \\ & + \left[\frac{k_2(n+1)m}{V} + k_4(n+1) \right] p_{n+1,m} \\ & + \frac{k_5(n+2)(n+1)p_{n+2,m-1}}{V} - \left[k_1 m + \frac{k_2 nm}{V} \right. \\ & \left. + k_3 V + k_4 n + \frac{k_5 n(n-1)}{V} + k_6 m \right] p_{n,m}. \end{aligned} \quad (18)$$

Letting $p_{n,m} = 0$ for all $[n, m] \notin \Omega$, and setting the left hand side of Eq. (18) to zero, we are left with a linear system.

FIG. 7. Trajectories of the chemical system (16), showing slow behaviour of the variable $S = X_1 + 2X_2$, where the system parameters are given by Eq. (17).

FIG. 8. The estimated probability density of $S = X_1 + 2X_2$, both by the CMA (red solid line, using CSSA simulations of length $T = 10^3$ for each value s of S), and by approximation of the steady state of the chemical master equation (18), on $\Omega = [0, 700] \times [0, 1225]$ (blue bars).

Finding the eigenvector of the corresponding matrix, with eigenvalue zero, gives us an approximate solution of the stationary chemical master equation. Summing over the possible values of the fast variable, we then obtain the probability distribution on the slow variable S , which is plotted in Figure 8 as the blue bars. The CMA approximation is visually indistinguishable from the approximation of the stationary CME.

B. Analysis of parameter dependence

One motivation for computing estimates of the effective dynamics of the slow variable, is to try to understand how the behaviour of the system is altered by a change in reaction rate parameters. In this section, we show that our SDE approximation of the effective dynamics of the slow variable can be used to compute a stochastic version of a bifurcation diagram for system (16). In the deterministic case (using mass action kinetics), the chemical system (16) is described by a system of two ordinary differential equations (ODEs):

$$\begin{aligned} \frac{dx_1}{dt} &= k_1 x_2 - k_2 x_1 x_2 + k_3 - k_4 x_1 - 2k_5 x_1^2 + 2k_6 x_2, \\ \frac{dx_2}{dt} &= k_5 x_1^2 - k_6 x_2, \end{aligned} \quad (19)$$

where x_1 and x_2 are concentrations of chemical species X_1 and X_2 , respectively. To make a better comparison with the stochastic model, we introduce variables $\bar{X}_1 = x_1 V$ and $\bar{X}_2 = x_2 V$, which are approximations of the number of molecules of chemical species X_1 and X_2 as predicted by the ODE system (19). Multiplying Eq. (19) by V , we obtain the following system of ODEs for $\bar{X}_1$ and $\bar{X}_2$:

$$\begin{aligned} \frac{d\bar{X}_1}{dt} &= k_1 \bar{X}_2 - \frac{k_2}{V} \bar{X}_1 \bar{X}_2 + k_3 V \\ &\quad - k_4 \bar{X}_1 - 2 \frac{k_5}{V} \bar{X}_1^2 + 2k_6 \bar{X}_2, \\ \frac{d\bar{X}_2}{dt} &= \frac{k_5}{V} \bar{X}_1^2 - k_6 \bar{X}_2. \end{aligned} \quad (20)$$

Let us consider the system (16) with parameters given by Eq. (17), but now with $k_3 V \in [1350, 1600]$. In this interval,

FIG. 9. Stochastic bifurcation of the chemical system (16) with parameters given by Eq. (17) with $k_3V \in [1350, 1600]$, estimated using the CMA with length of simulation $T = 10$. Blue circles (resp. red squares) denote the states which are estimated to be located at local maxima (resp. minima) of probability density. The solid (resp.) lines denote stable (resp. unstable) steady states of the ODE mean field approximation (20).

the mean field ODE approximation (20) of the dynamics of the system undergoes two fold bifurcations as it is shown in the corresponding bifurcation diagram plotted in Figure 9, denoted by the black line. Here a solid line indicates a stable steady state, and the dashed line indicates an unstable steady state. To make a direct comparison with the stochastic model, we visualise the bifurcation diagram using the variable $S = \bar{X}_1 + 2\bar{X}_2$.

In the stochastic model, we can identify the favourable states by the local maxima of the probability density. To estimate this, we use the CMA with simulation length $T = 10$ for each parameter value k_3V , for each value of the slow variable $S = X_1 + 2X_2$. For a range of parameter values $k_3V \in [1350, 1600]$, we estimate π_{CMA} , as in Figure 8. The results for this are plotted in Figure 10, where the log density is plotted for each parameter value k_3V . We then estimate the positions of the local maxima in the density. The maxima are shown in Figure 9 as blue circles. For a one-dimensional

FIG. 10. Log probability density of the steady state of the chemical system (16) with parameters given by Eq. (17) with $k_3V \in [1350, 1600]$, estimated using the CMA with length of simulation $T = 10$.

slow variable, as in this example, we can also identify local minima in a similar way, and these are plotted in Figure 9 as red squares. In this simple toy example the stochastic “bifurcation diagram” bears a strong resemblance to the bifurcation of the ODE approximation (20); however, other systems exist in which they are vastly different.³⁴

Stochastic modelling also gives rise to additional phenomena which are not captured by mean field ODEs, such as the switching behaviour between the favourable states in this example. Therefore, in the stochastic case, we can also estimate the average length of time that it takes to go from one favourable state to the other, and then back to the first again. Having estimated drift and diffusion coefficients, one can obtain estimates for the average switching time, by solving a suitable elliptic partial differential equation,³⁴ similar to Eq. (11). In the case of a one-dimensional slow variable S , this equation can be solved explicitly.³⁶ Let us suppose that $s_1 < s_2$. The average time to reach state $S = s_2$, given that we started at state $S = s_1$, can be computed by:

$$\tau_{s_1 \rightarrow s_2} = 4 \int_{s_1}^{s_2} \frac{1}{\log(\pi_{\text{CMA}}(y))} \int_0^y \frac{\log(\pi_{\text{CMA}}(z))}{D(z)} dz dy, \quad (21)$$

where π_{CMA} is the approximation of the density $\mathbb{P}(S = s)$. Similarly, the average time it takes for the process to reach state $S = s_1$, given that we started at state $S = s_2$, is given by

$$\tau_{s_1 \leftarrow s_2} = 4 \int_{s_1}^{s_2} \frac{1}{\log(\pi_{\text{CMA}}(y))} \int_y^\infty \frac{\log(\pi_{\text{CMA}}(z))}{D(z)} dz dy. \quad (22)$$

Let us denote the local minima of π_{CMA} by s_u (red squares in Figure 9), and local maxima of π_{CMA} by $s_- < s_+$ (blue circles in Figure 9). Figure 11 shows the plots for $\tau_{s_- \rightarrow s_+}$ and $\tau_{s_+ \leftarrow s_-}$, computed by Eqs. (21) and (22), respectively. Naturally, we truncate the second integral in Eq. (22) to compute it numerically. A modeller is sometimes interested in an average “switching time,” which can be computed as $\tau_{\text{SW}} = \tau_{s_- \rightarrow s_+} + \tau_{s_+ \leftarrow s_-}$, i.e., the average time it takes to go from one state of locally maximal probability density to the other, and back again. This is plotted as the dashed line in Figure 11. Additionally, we plot on this figure the results of a direct computation of average switching time, with $k_3V = 1475$. This value was averaged over 10^3 switches from one well to the other, using the Gillespie algorithm. For this parameter set, the agreement between the CMA approximation of the switching time and the direct stochastic computations is remarkable.

The formulae (21) and (22) can be used for general estimation of average travel times between two points for one-dimensional slow variables. For example in Figure 12, we show the average exit times from each well which are defined by $\tau_{s_- \rightarrow s_u}$ and $\tau_{s_u \leftarrow s_+}$, respectively. If the slow variable is of more than one dimension, then no exact formula such as Eq. (21) exists for the transition times, and the corresponding elliptic PDE needs to be solved numerically.³⁴

FIG. 11. Estimated average switching time for the system (16) with parameters given by Eq. (17) with $k_3 V \in [1350, 1600]$, using the CMA with length of simulation $T(s) = 10$ for each value s of S . Black plot denotes estimation using direct Gillespie estimation of the average switching time for $k_3 V = 1475$. Error bar denotes one standard deviation (approximated by $\sigma/\sqrt{10^3}$, where σ^2 is the sample variance and the average was taken over 10^3 switches).

VII. CHOOSING A SUITABLE LENGTH OF SIMULATION

When implementing the CSSA or CMA, one has to choose the state-dependent length of simulation $T(s)$. There are three natural ways to do this. The simplest approach is to select a fixed $T(s) \equiv \hat{T} \in \mathbb{R}^+$. Alternatively, one might also want to define $T(s)$ indirectly by postulating that we allow $L \in \mathbb{N}$ changes of the slow variable S before stopping the simulation for each value s . The third possibility is to design a stopping condition within the algorithms, to determine when the estimation of the observable we are interested in has converged, up to defined tolerances. This may vary depending on what the aim of the simulation is, whether it be approximation of switching times, or of the stationary probability density, or some other quantity. However a sensible stopping condition with the desired properties may be hard to identify. In this section, we instead look at how choosing $T(s)$ in one of the first two ways, either setting it to be a constant, or through allowing only L changes of the slow variable, over all values s

FIG. 12. Estimated average escape time for the two potential wells in the system (16) with parameters given by Eq. (17) with $k_3 V \in [1350, 1600]$, using the CMA with length of simulation $T(s) = 10$ for each value s of S .

FIG. 13. Error curves for fixed simulation length $\hat{T}$ and for fixed number of changes L of S for each value s , for the system (16), with parameters given by Eq. (17), as a function of total computational cost $\mathbb{Q}$ given by Eq. (23).

of S , can affect the accuracy of the approximation of a given quantity.

Naturally, there is a linear relationship between the total cost of running the CMA, with a given fixed simulation length $T(s) = \hat{T} \in \mathbb{R}^+$ (or fixed number of changes $L \in \mathbb{N}$ of the slow variable) for each value s of S . This was also the case in our numerical simulations. In the following we measure the computational cost $\mathbb{Q}$ in terms of how many individual reactions have been simulated in the algorithm. In the CMA, $Q(T(s))$ denotes the total number of reactions simulated in the algorithm up to time $T(s)$. This implies that the computational cost is

$$\mathbb{Q} \equiv \mathbb{Q}(T(\cdot)) = \sum_s Q(T(s)), \quad (23)$$

where the sum is over the chosen region of estimation of the slow variable in the algorithm. By running the algorithm for different values of $\hat{T}$ (or L in the case of limiting changes of the slow variable) and counting the total number of reaction evaluations, the linear relationship between simulation length parameter $\hat{T}$ (respectively, L) and the total computational cost $\mathbb{Q}$ can then be found. We can then use these approximations to compare the calculated π_{CMA} for the system using a fixed $\hat{T}$ and L , respectively, with roughly the same total computational cost, and then compare these approximations with the approximation to the solution of the stationary chemical master equation (18).

Figure 13 shows the errors (as defined by Eq. (14)) in the approximation π_{CMA} as a function of total computational cost, both for fixed simulation length $\hat{T}$, and for fixed number of changes L of S , for the bistable chemical system (16) with reaction parameters given by Eq. (17). In this example, the simulations with fixed number L of changes of the slow variable, instead of fixed $T(s) = \hat{T}$, led to approximately half the error in the resulting estimation π_{CMA} of the invariant distribution, when given the same amount of computational resource. This result is system dependent, however, since the same error curves computed for the illustrative system (1) gave results which were barely distinguishable for the two simplest methods of choosing the simulation length.

In theory, we could also estimate a suitable length of simulation by calculating the smallest non-zero eigenvalue of the evolution equation describing the system dynamics. This is equivalent to solving the chemical master equation, and thus is computationally expensive or in some cases unfeasible. For simple systems, such as Eq. (1), such a computation is possible. In particular, we find that the smallest negative eigenvalue for the choice of fast and slow variables (3) is smaller in magnitude than the one corresponding to the choice (5), which is consistent with error plots seen in Figure 2.

Another approach would be to pick a time T , and then run the estimations for length T , $2T$, $4T$, etc. until consecutive estimations are within a given tolerance of each other. This approach is much the same as that commonly used when solving PDEs, where we might trust our approximation of the solution if we halve the mesh size and get the same result (up to some tolerance).

Alternatively, the simulation length need not necessarily be set *a priori*, but the simulations could be run until a given degree of convergence has been reached. This can be estimated using standard Markov chain Monte Carlo convergence diagnostics³⁷ (auto-correlation, etc.).

VIII. DISCUSSION

In this paper, we developed a multiscale methodology for analysis of systems with multiple time scales. We demonstrated that this approach can be used to approximate important characteristics of such systems, such as invariant probability density and switching times. One of the limitations of the method is the dimension of the slow variable vector (i.e., number of slow variables). In this paper, we presented examples with one or two slow variables. In biological applications, biochemical systems can be written in terms of tens of variables,^{1–3} often reacting on different time scales. In order to apply this approach to such a system, we must first identify the slowly changing quantities. The direct exploration of a system, similar to that presented in Sec. II, might provide a guidance for the selection of the slow variable(s). However, in some systems, different variables are slow (in comparison with the other variables) in different parts of the state space. In this case, an automated analysis might be needed to identify the slow manifold.^{38–40} If the dimension of the slow variable vector is only one or two, then the CMA can be applied exactly as we have here, no matter what the overall dimension of the system.

The situation becomes more challenging when the system has three or more slow variables. The algorithm we have described here becomes more expensive due to the larger number of points at which one is required to estimate the values of the drift and diffusion terms. Once we have these estimates, it is then more difficult to solve due to the resulting high dimensional PDEs. To tackle this problem, we are currently investigating more efficient approaches for solution of these PDEs, involving adaptive mesh finite element methods. This problem differs from typical problems in numerical analysis, in that we have to optimise the choice of domain in which the drift and diffusion coefficients are estimated. Furthermore, since the estimations are made using stochastic simulations,

they are noisy, and may need to be smoothed in order to obtain good results. We will report on our findings in the future.

In some biological examples, there are variables which can take values as low as in the single digit values (or even zero). For example, the relative abundance of some mRNA molecules in cells can be less than one, on average.^{3,41} Gene states are sometimes represented as Boolean variables (a variable which is zero for a specified gene being “off,” and equal to one when the gene is transcribed).⁹ In some cases, these variables are fast and do not cause any problems.⁴² However, if a slow variable can only have a few discrete values, then the SDE-based approximations can lead to errors, and slow algorithms written in terms of Markov chains are more applicable.^{20,43} To improve efficiency of this approach, one could also couple it with the effective diffusivity approach for areas with larger copy numbers, in a hybrid model. The CMA can also be modified by substituting the Gillespie algorithm by a statistically equivalent method, such as the Gibson-Bruck algorithm, using a suitable dependency graph.⁵

It is worth pointing it out that the CMA shares a common theoretical foundation with the methods proposed by Cao *et al.*,^{18,20} E *et al.*,^{21,43} and other authors.^{19,44} In particular, the main assumption in all these works is that for the chemical system under study there exists a separation of time scales, and that the suitably defined fast variables approach equilibrium faster than the time scale of change of the slow variables. However, unlike the approach taken in Refs. 18, 20, 21, and 43 our focus is not on speeding up the SSA, but on approximation of the evolution of the slow variables by an effective Fokker-Planck equation. This then allows us to estimate important quantities of interest.

The efficiency of the CMA over other existing algorithms is problem dependent. The advantage of this method is that since we are approximating the system by a continuous process, we are able to solve PDEs pertaining to the system on non-regular meshes, where more resolution can be added near to regions of interest, and the resolution can be reduced in others. Furthermore, the efficiency of the algorithm is independent of the number of fast variables, meaning that in theory this method is applicable to large problems in biological applications if the slow variable space is low dimensional.

ACKNOWLEDGMENTS

The research leading to these results has received funding from the European Research Council under the *European Community's* Seventh Framework Programme (FP7/2007-2013)/ERC Grant Agreement No. 239870. This publication was based on work supported in part by Award No. KUK-C1-013-04, made by King Abdullah University of Science and Technology (KAUST). The work of I.G.K. was partially supported by the US Department of Energy. Thanks also to Tomáš Vejchodský for useful conversations regarding the assembly and solution within the finite element method. RE

would also like to thank Somerville College, University of Oxford, for a Fulford Junior Research Fellowship.

APPENDIX: DERIVATION OF DRIFT AND DIFFUSION FORMULAE

Considering the Fokker-Planck equation (8), the drift V and diffusion D are given by the following expressions:

$$V(\mathbf{s}) = \lim_{\Delta t \rightarrow 0} \frac{\langle \mathbf{S}(\Delta t) - \mathbf{s} | \mathbf{S}(0) = \mathbf{s} \rangle}{\Delta t}, \quad (\text{A1})$$

$$D(\mathbf{s}) = \lim_{\Delta t \rightarrow 0} \frac{\langle [\mathbf{S}(\Delta t) - \mathbf{s}][\mathbf{S}(\Delta t) - \mathbf{s}]^T | \mathbf{S}(0) = \mathbf{s} \rangle}{2\Delta t},$$

where the initialization of the fast variables in the computation of averages $\langle \cdot \rangle$ is given by the probability distribution $\mathbb{P}(\mathbf{F} | \mathbf{S} = \mathbf{s})$, which we can sample using the CSSA in Sec. III. The CSSA is just one possible way to determine the conditional distribution on the fast variable given the slow variable. The advantage of this method over others,⁴⁵ is that we are able to estimate the drift and diffusion terms of the slow variables on the fly, while sampling the conditional distribution of the fast variables.

Let us define $\delta \mathbf{S}_j$ to be the change in $\mathbf{S}$ due to the j th iteration of the CSSA. For example, in the case of the illustrative example system (1) for the choice of slow variable $\mathbf{S} = X_1 + X_2$, where the j th reaction to occur was R_i , we have

$$\delta \mathbf{S}_j = \begin{cases} 0 & \text{for } i = 2, 3, 5, \\ 1 & \text{for } i = 1, \\ -1 & \text{for } i = 4. \end{cases}$$

We define the variable $\tilde{\mathbf{S}}$, to be the sum of these elementary changes $\delta \mathbf{S}_j$, i.e.,

$$\tilde{\mathbf{S}}(t) = \sum_{j=1}^{Q(t)} \delta \mathbf{S}_j, \quad (\text{A2})$$

where $Q(t)$ iterations of the CSSA have been simulated up to some time $t > 0$. Since $\mathbf{S}(0) = \mathbf{s}$, formula (A1) implies that

$$V(\mathbf{s}) = \lim_{\Delta t \rightarrow 0} \frac{\langle \mathbf{S}(\Delta t) - \mathbf{s} | \mathbf{S}(0) = \mathbf{s} \rangle}{\Delta t}$$

$$= \lim_{\Delta t \rightarrow 0} \frac{\langle \tilde{\mathbf{S}}(t + \Delta t) - \tilde{\mathbf{S}}(t) \rangle}{\Delta t},$$

for all $t > 0$, due to the resetting in step [4] of the CSSA, in Table III. Therefore, if $T = J\Delta t$ is the length of the CSSA simulation, we get

$$V(\mathbf{s}) = \lim_{\Delta t \rightarrow 0} \lim_{J \rightarrow \infty} \frac{1}{J} \sum_{j=1}^J \frac{\tilde{\mathbf{S}}(j\Delta t) - \tilde{\mathbf{S}}((j-1)\Delta t)}{\Delta t}$$

$$= \lim_{\Delta t \rightarrow 0} \lim_{J \rightarrow \infty} \frac{1}{J\Delta t} (\tilde{\mathbf{S}}(J\Delta t) - \tilde{\mathbf{S}}(0)).$$

Using $\tilde{\mathbf{S}}(0) = 0$, $T = J\Delta t$, and Eq. (A2), we have that

$$V(\mathbf{s}) = \lim_{T \rightarrow \infty} \frac{1}{T} \tilde{\mathbf{S}}(T) = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{m=1}^{Q(T)} \delta \mathbf{S}_m,$$

which is the formula given by Eq. (9). In a similar way, we can derive Eq. (10), namely,

$$D(\mathbf{s}) = \frac{1}{2} \lim_{\Delta t \rightarrow 0} \frac{\langle [\mathbf{S}(\Delta t) - \mathbf{s}][\mathbf{S}(\Delta t) - \mathbf{s}]^T | \mathbf{S}(0) = \mathbf{s} \rangle}{\Delta t}$$

$$= \frac{1}{2} \lim_{J \rightarrow \infty} \lim_{\Delta t \rightarrow 0} \frac{1}{J\Delta t} \sum_{j=1}^J [\tilde{\mathbf{S}}(j\Delta t) - \tilde{\mathbf{S}}((j-1)\Delta t)]$$

$$\times [\tilde{\mathbf{S}}(j\Delta t) - \tilde{\mathbf{S}}((j-1)\Delta t)]^T = \lim_{T \rightarrow \infty} \frac{1}{2T} \sum_{m=1}^{Q(T)} \delta \mathbf{S}_m \delta \mathbf{S}_m^T.$$

¹A. Arkin, J. Ross, and H. McAdams, *Genetics* **149**, 1633 (1998).

²J. Villar, H. Kueh, N. Barkai, and S. Leibler, *Proc. Natl. Acad. Sci. U.S.A.* **99**, 5988 (2002).

³S. Kar, W. Baumann, M. Paul, and J. Tyson, *Proc. Natl. Acad. Sci. U.S.A.* **16**, 6471 (2009).

⁴D. Gillespie, *J. Phys. Chem.* **81**, 2340 (1977).

⁵M. Gibson and J. Bruck, *J. Phys. Chem. A* **104**, 1876 (2000).

⁶Y. Cao, H. Li, and L. Petzold, *J. Chem. Phys.* **121**, 4059 (2004).

⁷G. Klingbeil, R. Erban, M. Giles, and P. Maini, *IEEE Trans. Parallel Distrib. Syst.* **99** (2011).

⁸G. Klingbeil, R. Erban, M. Giles, and P. Maini, *Bioinformatics* (2011).

⁹R. Erban, I. Kevrekidis, D. Adalsteinsson, and T. Elston, *J. Chem. Phys.* **124**, 084106 (2006).

¹⁰D. Gillespie, *J. Chem. Phys.* **115**, 1716 (2001).

¹¹A. Chatterjee, D. Vlachos, and M. Katsoulakis, *J. Chem. Phys.* **122**, 024112 (2005).

¹²T. Tian and K. Burrage, *J. Chem. Phys.* **121**, 10356 (2004).

¹³A. Auger, P. Chatelain, and P. Koumoutsakos, *J. Chem. Phys.* **125**, 084103 (2006).

¹⁴D. Gillespie, *J. Chem. Phys.* **113**, 297 (2000).

¹⁵B. Mélykúti, K. Burrage, and K. Zygalakis, *J. Chem. Phys.* **132**, 164109 (2010).

¹⁶D. Adalsteinsson, D. McMillen, and T. Elston, *BMC Bioinf.* **5** (2004).

¹⁷J. Wilkie and Y. Wong, *Chem. Phys.* **353**, 132 (2008).

¹⁸Y. Cao, D. Gillespie, and L. Petzold, *J. Comput. Phys.* **206**, 395 (2005).

¹⁹E. Haseltine and J. Rawlings, *J. Chem. Phys.* **117**, 6959 (2002).

²⁰Y. Cao, D. Gillespie, and L. Petzold, *J. Chem. Phys.* **122**, 14116 (2005).

²¹W. E, D. Liu, and E. Vanden-Eijnden, *J. Chem. Phys.* **123**, 194107 (2005).

²²E. Vanden-Eijnden, *Commun. Math. Sci.* **1**, 385 (2003).

²³H. Salis and Y. Kaznessis, *J. Chem. Phys.* **122**, 054103 (2005).

²⁴I. Kevrekidis, C. Gear, J. Hyman, P. Kevrekidis, O. Runborg, and K. Theodoropoulos, *Commun. Math. Sci.* **1**, 715 (2003).

²⁵M. Haataja, D. Srolovitz, and I. Kevrekidis, *Phys. Rev. Lett.* **92**, 160603 (2004).

²⁶R. Erban, I. Kevrekidis, and H. Othmer, *Physica D* **215**, 1 (2006).

²⁷C. Yates, R. Erban, C. Escudero, I. Couzin, J. Buhl, I. Kevrekidis, P. Maini, and D. Sumpter, *Proc. Natl. Acad. Sci. U.S.A.* **106**, 5464 (2009).

²⁸G. Torrie and J. Valleau, *J. Comput. Phys.* **23**, 187 (1977).

²⁹B. Roux, *Comput. Phys. Commun.* **91**, 275 (1995).

³⁰G. Ciccotti, R. Kapral, and E. Vanden-Eijnden, *ChemPhysChem* **6**, 1809 (2005).

³¹T. Lelièvre, C. L. Bris, and E. Vanden-Eijnden, *C. R. Math. Acad. Sci. Paris* **346**, 471 (2008).

³²G. Ciccotti, T. Lelièvre, and E. Vanden-Eijnden, *Commun. Pure Appl. Math.* **61**, 371 (2008).

³³T. Jahnke and W. Huisinga, *J. Math. Biol.* **54**, 1 (2007).

³⁴R. Erban, S. Chapman, I. Kevrekidis, and T. Vejchodský, *SIAM J. Appl. Math.* **70**, 984 (2009).

³⁵“Burn-in” simply means an initial period of stochastic simulation where the results are ignored, in order to allow the process to enter equilibrium.

³⁶C. Gardiner, *Handbook of Stochastic Methods*, 2nd ed., Springer Series in Synergetics Vol. 13 (Springer-Verlag, Berlin, 1985) for physics, chemistry and the natural sciences.

³⁷M. Cowles and B. Carlin, *J. Am. Stat. Assoc.* **91**, 883 (1996).

³⁸R. Erban, T. Frewen, X. Wang, T. Elston, R. Coifman, B. Nadler, and I. Kevrekidis, *J. Chem. Phys.* **126**, 155103 (2007).

- ³⁹A. Singer, R. Erban, I. Kevrekidis, and R. Coifman, *Proc. Natl. Acad. Sci. U.S.A.* **106**, 16090 (2009).
- ⁴⁰C. Siettos, M. Graham, and I. Kevrekidis, *J. Chem. Phys.* **118**, 10149 (2003).
- ⁴¹B. Alberts, A. Johnson, J. Lewis, M. Raff, K. Roberts, and P. Walter, *Molecular Biology of the Cell* (Garland Science, New York, 2008)
- ⁴²T. Kepler and T. Elston, *Biophys. J.* **91**, 3116 (2001).
- ⁴³W. E. D. Liu, and E. Vanden-Eijnden, *J. Comput. Phys.* **221**, 158 (2007).
- ⁴⁴C. Rao and A. Arkin, *J. Chem. Phys.* **118**, 4999 (2003).
- ⁴⁵A. Gelfand and A. Smith, *J. Am. Stat. Assoc.* **85**, 398 (1990).