

Accelerated stochastic simulation of the stiff enzyme-substrate reaction

Yang Cao^{a)}Department of Computer Science, University of California, Santa Barbara,
Santa Barbara, California 93106Daniel T. Gillespie^{b)}

Dan T. Gillespie Consulting, Castaic, California 91384

Linda R. Petzold^{c)}Department of Computer Science, University of California, Santa Barbara,
Santa Barbara, California 93106

(Received 5 May 2005; accepted 11 August 2005; published online 13 October 2005)

The enzyme-catalyzed conversion of a substrate into a product is a common reaction motif in cellular chemical systems. In the three reactions that comprise this process, the intermediate enzyme-substrate complex is usually much more likely to decay into its original constituents than to produce a product molecule. This condition makes the reaction set mathematically “stiff.” We show here how the simulation of this stiff reaction set can be dramatically speeded up relative to the standard stochastic simulation algorithm (SSA) by using a recently introduced procedure called the *slow-scale* SSA. The speedup occurs because the slow-scale SSA explicitly simulates only the relatively rare conversion reactions, skipping over occurrences of the other two less interesting but much more frequent reactions. We describe, explain, and illustrate this simulation procedure for the isolated enzyme-substrate reaction set, and then we show how the procedure extends to the more typical case in which the enzyme-substrate reactions occur together with other reactions and species. Finally, we explain the connection between this slow-scale SSA approach and the Michaelis–Menten [Biochem. Z. **49**, 333 (1913)] formula, which has long been used in deterministic chemical kinetics to describe the enzyme-substrate reaction. © 2005 American Institute of Physics.
[DOI: 10.1063/1.2052596]

I. INTRODUCTION

The simple enzyme-substrate reaction set



describes a common mechanism by which an *enzyme* S_1 catalyzes the conversion of a *substrate* S_2 into a *product* S_4 . The conversion proceeds by way of an unstable *enzyme-substrate complex* S_3 , which decays either into its original constituents S_1 and S_2 , or into its converted constituents S_1 and S_4 . Often, the former outcome is much more likely than the latter, in consequence of the condition

$$c_2 \gg c_3, \quad (2)$$

and we shall assume that this condition holds in what follows. But we make no assumptions about the molecular population levels of the various species; in particular, we expressly allow that the average number of S_3 molecules might be small, even less than one, which sometimes happens in practice. Accordingly, our analysis will suppose that the system’s state vector $X(t) \equiv (X_1(t), \dots, X_4(t))$, where $X_i(t)$ is the number of molecules of species S_i in the system at time

t , moves over the non-negative *integer* lattice in a *stochastic* manner. The stochasticity is a consequence of the premise that if $X(t)=x$, then each reaction R_j has probability $a_j(x)dt$ of firing in the next infinitesimal time dt , where a_j is the *propensity function* for reaction R_j .¹ The propensity functions for the three reactions in (1) are

$$a_1(x) = c_1 x_1 x_2, \quad a_2(x) = c_2 x_3, \quad a_3(x) = c_3 x_3. \quad (3)$$

On account of condition (2), the reaction set (1) evolves on two separate time scales: the “fast” time scale of reactions R_1 and R_2 , and the “slow” time scale of reaction R_3 . In the context of a traditional *deterministic* analysis, the system is said to be *stiff*, and numerically solving the associated ordinary differential equations then poses special challenges. From the stochastic viewpoint that we are taking here, the difficulty introduced by the stiffness condition (2) is that the stochastic simulation algorithm (SSA), which simulates the discrete reaction events sequentially in accordance with the propensity functions,¹ is forced to simulate very many R_1 and R_2 reactions in order to simulate each R_3 reaction. Since the former reactions merely “undo” each other while the latter completes the conversion of a substrate molecule into a product molecule, R_1 and R_2 are “frequently occurring unimportant” reactions, while R_3 is a “rarely occurring important” reaction. A more efficient stochastic simulation procedure would skip over the R_1 and R_2 reactions and simulate only the R_3 reactions. An *approximate* way of doing that is

^{a)}Electronic mail: ycao@engineering.ucsb.edu

^{b)}Electronic mail: gillespiedt@mailaps.org

^{c)}Electronic mail: petzold@engineering.ucsb.edu

afforded by the recently introduced *slow-scale* stochastic simulation algorithm.²

In Secs. II–VIII, we show how the slow-scale SSA can be applied to reactions (1) under condition (2) to achieve a substantial gain in simulation speed at practically no cost in simulation accuracy. In Sec. IX we describe how the simulation procedure is to be applied in the more common case in which reactions (1) are occurring together with other reactions and species. Finally, in Sec. X, we connect the slow-scale SSA approach to some well-known deterministic strategies of approximating reactions (1), namely, the quasi-steady-state approximation, the partial (or rapid) equilibrium approximation, and the Michaelis–Menten formula.

System (1) has been addressed previously by Rao and Arkin³ using an approach very similar to the one described here. We believe that our approach represents an improvement over Ref. 3, both in its theoretical perspective and in its algorithmic implementation; however, the development of our approach was materially aided and informed by that earlier work.

II. PARTITIONING THE SYSTEM

A detailed derivation of the slow-scale SSA is given in Ref. 2. We shall not repeat that derivation here, but simply describe how the slow-scale SSA is applied to reactions (1) under condition (2). The first step is to partition the reaction set $\{R_1, R_2, R_3\}$ into fast and slow subsets. In general this partitioning is done *provisionally*, making an educated guess that is later validated or invalidated. For our system, condition (2) clearly implies that R_2 will be a fast reaction and R_3 will be a slow reaction. We shall tentatively classify R_1 as a fast reaction, like R_2 , on the heuristic grounds that those two reactions will approximately equilibrate, so their respective propensity functions $c_1x_1x_2$ and c_2x_3 will become approximately equal. (The “speed” of a reaction is measured by the size of its propensity function, not its rate constant.) We shall see later that classifying R_1 and R_2 as fast and R_3 as slow will, in fact, be acceptable, provided condition (2) is satisfied strongly enough.

The next step is to partition the species set $\{S_1, \dots, S_4\}$ into fast and slow subsets. The rule for doing this is unambiguous: Any species whose population *gets changed* by a *fast* reaction is called a *fast* species, and all other species are called *slow*. So in the present case, S_1 , S_2 , and S_3 will be fast species, and S_4 will be a slow species. The system state vector can now be written $X(t) = (X^f(t), X^s(t))$, where the *fast process* is $X^f(t) = (X_1(t), X_2(t), X_3(t))$, and the *slow process* is $X^s(t) = X_4(t)$. The fact that the only reactant in the slow reaction is a fast species underscores the asymmetric and rather subtle relationship between fast and slow reactions and fast and slow species.

The third step in setting up the slow-scale SSA is to construct what is called the *virtual fast process* $\hat{X}^f(t)$. It is defined to be the fast state variables evolving under *only* the fast reactions; i.e., $\hat{X}^f(t)$ is $X^f(t)$ with all the slow reactions switched off. Since the slow reactions by definition occur only infrequently, we may expect that $\hat{X}^f(t)$ will provide a reasonably good approximation to $X^f(t)$. This approximation

will be useful because $\hat{X}^f(t) = (\hat{X}_1(t), \hat{X}_2(t), \hat{X}_3(t))$ evolving under reactions R_1 and R_2 is mathematically much simpler to analyze than $X^f(t) = (X_1(t), X_2(t), X_3(t))$ evolving under reactions R_1 , R_2 , and R_3 .

Two conditions, which together define what might be called “stochastic stiffness,” are now required to be satisfied in order for the slow-scale SSA to be applicable: First, the virtual fast process $\hat{X}^f(t)$ must be stable, in the sense that it approaches a well-defined time-independent limit $\hat{X}^f(\infty)$ as $t \rightarrow \infty$; this requirement can be thought of as the stochastic equivalent of the well-known deterministic stiffness requirement that the system’s fastest dynamical mode be stable. Second, the approach of $\hat{X}^f(t)$ to $\hat{X}^f(\infty)$ must be accomplished in a time that is small compared to the expected time to the next slow reaction; this is essentially a refinement of condition (2)—a more precise specification of the degree of separation that must exist between the time scales of the fast and slow reactions. If these two conditions are not satisfied, then we must modify our initial partitioning of the reactions into fast and slow subsets. If that fails, we must forego using the slow-scale SSA. We shall later see that, for our problem here, both of these stiffness conditions will be satisfied if the inequality (2) is strong enough.

Our virtual fast process $\hat{X}^f(t)$ obeys two *conservation relations*:

$$\hat{X}_1(t) + \hat{X}_3(t) = x_{71} \quad (\text{const.}), \quad (4a)$$

$$\hat{X}_2(t) + \hat{X}_3(t) = x_{72} \quad (\text{const.}). \quad (4b)$$

These can be understood by viewing an S_3 molecule as the union of one S_1 molecule and one S_2 molecule, so that Eq. (4a) expresses the constancy of the total number of S_1 molecular units, and Eq. (4b) the constancy of the total number of S_2 molecular units, all under the two fast reactions R_1 and R_2 . Equations (4a) and (4b) reduce the number of independent variables making up the virtual fast process from three to one, which is a considerable mathematical simplification. By contrast, the real fast process $X^f(t)$ satisfies Eq. (4a) but *not* Eq. (4b), since $X^f(t)$ is affected by the slow reaction R_3 . $X^f(t)$ thus has two independent variables, and accordingly would be more difficult to analyze than $\hat{X}^f(t)$.

III. THE SLOW-SCALE APPROXIMATION

The condition that $\hat{X}^f(t)$ approaches $\hat{X}^f(\infty)$ in a time that is small compared to the expected time to the next slow reaction—a condition that we shall quantify later—sets the stage for invoking a result called the slow-scale approximation.² This approximation forms the theoretical basis for the slow-scale SSA. It says, in essence, that we can approximately simulate reactions (1) one R_3 reaction at a time if we *replace* the R_3 propensity function $a_3(x) = c_3x_3$ with its average with respect to the asymptotic virtual fast process:

$$\bar{a}_3(x) = c_3 \langle \hat{X}_3(\infty) \rangle. \quad (5)$$

This surrogate propensity function for R_3 is called its “slow-scale propensity function.” To use it, we obviously must be able to compute $\langle \hat{X}_3(\infty) \rangle$. If we want to allow reactions (1) to occur along with other slow reactions besides R_3 , which might have as reactants any of the fast species in any combination, we would also need to be able to compute $\langle \hat{X}_i(\infty) \rangle$ and $\langle \hat{X}_i^2(\infty) \rangle$ for $i=1, 2, 3$, since these values might be needed to construct the slow-scale propensity functions for those other slow reactions. Finally, if we want our simulation to show the trajectories of the fast species as well as the trajectories of slow species, as we usually do, then we will need to be able to efficiently generate *random samples* of $\hat{X}_1(\infty)$, $\hat{X}_2(\infty)$, and $\hat{X}_3(\infty)$, since that is how the fast species trajectories get constructed in the slow-scale SSA.

In Secs. IV and V we shall show how all these computations involving the asymptotic virtual fast process $\hat{X}^f(\infty)$ can be done.

IV. THE MASTER EQUATION FOR THE VIRTUAL FAST PROCESS

We begin our analysis of the virtual fast process by using the conservation relations (4) to eliminate $\hat{X}_1(t)$ and $\hat{X}_2(t)$ in favor of $\hat{X}_3(t)$. Once that is done, $\hat{X}_3(t)$ takes the form of what is known as a “birth-death Markov process”⁴ with “stepping functions”

$$W_-(x_3) = c_2 x_3, \quad W_+(x_3) = c_1(x_{T1} - x_3)(x_{T2} - x_3). \quad (6)$$

Here, $W_{\pm}(x_3)dt$ gives the probability, given $\hat{X}_3(t)=x_3$, that $\hat{X}_3(t+dt)$ will equal $x_3 \pm 1$. The master equation for $\hat{X}_3(t)$ reads

$$\begin{aligned} \frac{\partial \hat{P}(x_3, t | x_{T1}, x_{T2}, x_3^{(0)})}{\partial t} &= W_-(x_3 + 1) \hat{P}(x_3 + 1, t | x_{T1}, x_{T2}, x_3^{(0)}) \\ &\quad - W_-(x_3) \hat{P}(x_3, t | x_{T1}, x_{T2}, x_3^{(0)}) \\ &\quad + W_+(x_3 - 1) \hat{P}(x_3 - 1, t | x_{T1}, x_{T2}, x_3^{(0)}) \\ &\quad - W_+(x_3) \hat{P}(x_3, t | x_{T1}, x_{T2}, x_3^{(0)}), \end{aligned} \quad (7)$$

where $\hat{P}(x_3, t | x_{T1}, x_{T2}, x_3^{(0)})$ is the probability that $\hat{X}_3(t \geq 0) = x_3$, given that $\hat{X}_3(0) = x_3^{(0)}$, $\hat{X}_1(0) = x_{T1} - x_3^{(0)}$, and $\hat{X}_2(0) = x_{T2} - x_3^{(0)}$.

It follows from Eqs. (4a) and (4b) that $\hat{X}_3(t)$ will be bounded above by

$$x_{3 \max} = \min(x_{T1}, x_{T2}). \quad (8)$$

Now, a *bounded* birth-death Markov process is always *stable*, meaning in this case that $\hat{P}(x_3, t | x_{T1}, x_{T2}, x_3^{(0)})$ will approach, as $t \rightarrow \infty$, a well-behaved stationary form, $\hat{P}(x_3, \infty | x_{T1}, x_{T2})$. This of course is the probability density function of $\hat{X}_3(\infty)$. Being a time-independent solution of the master equation (7), it satisfies

$$\begin{aligned} 0 &= W_-(x_3 + 1) \hat{P}(x_3 + 1, \infty | x_{T1}, x_{T2}) \\ &\quad - W_-(x_3) \hat{P}(x_3, \infty | x_{T1}, x_{T2}) \\ &\quad + W_+(x_3 - 1) \hat{P}(x_3 - 1, \infty | x_{T1}, x_{T2}) \\ &\quad - W_+(x_3) \hat{P}(x_3, \infty | x_{T1}, x_{T2}). \end{aligned}$$

Rearranging this equation reveals that the quantity

$$\begin{aligned} W_-(x_3) \hat{P}(x_3, \infty | x_{T1}, x_{T2}) - W_+(x_3 - 1) \hat{P}(x_3 - 1, \infty | x_{T1}, x_{T2}) \end{aligned}$$

must be a constant (independent of x_3), and a consideration of the case $x_3=0$ shows that the constant must be zero; thus we obtain the *detailed balance relation*

$$W_-(x_3) \hat{P}(x_3, \infty | x_{T1}, x_{T2}) = W_+(x_3 - 1) \hat{P}(x_3 - 1, \infty | x_{T1}, x_{T2}). \quad (9)$$

A simple rearrangement of (9) yields a recursion relation for $\hat{P}(x_3, \infty | x_{T1}, x_{T2})$,

$$\begin{aligned} \hat{P}(x_3, \infty | x_{T1}, x_{T2}) &= \frac{W_+(x_3 - 1)}{W_-(x_3)} \hat{P}(x_3 - 1, \infty | x_{T1}, x_{T2}) \\ &\quad (x_3 = 1, \dots, x_{3 \max}). \end{aligned} \quad (10a)$$

The initial value $\hat{P}(0, \infty | x_{T1}, x_{T2})$ for this recursion is chosen to satisfy the normalization condition

$$\sum_{x_3=0}^{x_{3 \max}} \hat{P}(x_3, \infty | x_{T1}, x_{T2}) = 1. \quad (10b)$$

Once $\hat{P}(x_3, \infty | x_{T1}, x_{T2})$ has been computed from Eqs. (10a) and (10b), the first two moments of $\hat{X}_3(\infty)$ can then be computed as

$$\langle \hat{X}_3^k(\infty) \rangle = \sum_{x_3=0}^{x_{3 \max}} x_3^k \hat{P}(x_3, \infty | x_{T1}, x_{T2}) \quad (k = 1, 2). \quad (11)$$

Another approach to computing these moments is as follows: First, sum the detailed balance relation (9) over all x_3 to obtain

$$\langle W_-(\hat{X}_3(\infty)) \rangle = \langle W_+(\hat{X}_3(\infty)) \rangle. \quad (12a)$$

Next, multiply Eq. (9) by x_3^n , and then sum the result over all x_3 to obtain

$$\langle \hat{X}_3(\infty)^n W_-(\hat{X}_3(\infty)) \rangle = \langle (\hat{X}_3(\infty) + 1)^n W_+(\hat{X}_3(\infty)) \rangle \quad (n = 1, 2, \dots). \quad (12b)$$

When the polynomial functions $W_{\pm}(x_3)$ in Eqs. (6) are inserted, Eqs. (12a) and (12b) become an infinite set of equations that interrelate all the moments of $\hat{X}_3(\infty)$.

For the stepping functions (6), the calculations indicated in Eqs. (10a), (10b), and (11) can all be carried out numerically, as detailed in the Appendix, *unless* $x_{3 \max}$ is too large for that to be practical. Analytical solutions are more problematic. Solving Eqs. (12a) and (12b) directly is not generally possible because the quadratic form of the function $W_+(x_3)$ in Eqs. (6) causes Eq. (12a) to contain $\langle \hat{X}_3^2(\infty) \rangle$ as

well as $\langle \hat{X}_3(\infty) \rangle$, and the $n=1$ version of Eq. (12b) to contain $\langle \hat{X}_3^3(\infty) \rangle$, etc., so there is always one more unknown moment than there are interrelating equations. One *approximate* way around this difficulty is to assume that $\hat{X}_3(\infty)$ is a normal random variable; this closes the first two equations (since the third moment of a normal random variable is expressible in terms of its first two moments), but solving the resulting two equations for those two moments must be done implicitly. Darvey *et al.*⁵ have analyzed reactions R_1 and R_2 in great detail using the generating function approach, and they derive *exact, explicit* formulas for the first two moments in terms of confluent hypergeometric functions. They also derive an exact formula for the generating function $G(s, \infty) \equiv \sum_{x_3} s^{x_3} \hat{P}(x_3, \infty)$ of the probability function $\hat{P}(x_3, \infty)$; however, it would appear that deconvolving $G(s, \infty)$ to get $\hat{P}(x_3, \infty)$, which is required to generate exact random samples of $\hat{X}_3(\infty)$, will involve at least as many computational steps as the procedure for directly iterating Eq. (10a) that is outlined in the Appendix.

In Sec. V, we describe a practical, approximate way of computing the needed properties of $\hat{X}_3(\infty)$ which appears to be acceptably accurate and reasonably fast. We do not discount the possibility that other approaches, such as the generating function approach of Darvey *et al.*,⁵ might be able to compute some of those properties even more efficiently, or might be more easily extended to other virtual fast processes.

V. PRACTICAL APPROXIMATE COMPUTATIONS OF $\hat{X}_3(\infty)$

In this section we shall describe practical approximate ways to compute the first two moments of $\hat{X}_3(\infty)$ and to generate random samples of $\hat{X}_3(\infty)$. We shall investigate when these approximate methods are satisfactory, and how to proceed alternatively when they are not. Finally, we shall derive a quantitative test—essentially a refinement of condition (2)—that will tell us when the slow-scale SSA is applicable to reactions (1).

A. The first two moments

Let us first focus on estimating $\langle \hat{X}_3(\infty) \rangle$ and $\langle \hat{X}_3^2(\infty) \rangle$. Note that a knowledge of these two moments will enable us to easily compute asymptotic first and second moments involving any of the other fast species S_1 and S_2 , for use in computing the slow-scale propensity functions of any slow reactions besides R_3 . Thus, from the conservation relations (4a) and (4b) we have

$$\hat{X}_1(\infty) = x_{T1} - \hat{X}_3(\infty), \quad \hat{X}_2(\infty) = x_{T2} - \hat{X}_3(\infty), \quad (13)$$

and so it follows that

$$\langle \hat{X}_1(\infty) \rangle = x_{T1} - \langle \hat{X}_3(\infty) \rangle, \quad (14a)$$

$$\langle \hat{X}_1^2(\infty) \rangle = x_{T1}^2 - 2x_{T1} \langle \hat{X}_3(\infty) \rangle + \langle \hat{X}_3^2(\infty) \rangle, \quad (14b)$$

$$\langle \hat{X}_1(\infty) \hat{X}_3(\infty) \rangle = x_{T1} \langle \hat{X}_3(\infty) \rangle - \langle \hat{X}_3^2(\infty) \rangle, \quad (14c)$$

and so forth. The corresponding variances can then be calculated from

$$\text{var}\{\hat{X}_i(\infty)\} = \langle \hat{X}_i^2(\infty) \rangle - \langle \hat{X}_i(\infty) \rangle^2 \quad (i = 1, 2, 3), \quad (15)$$

and it is not hard to show from Eqs. (13) that these three variances will be equal.

The simplest approximation to the *first* moment $\langle \hat{X}_3(\infty) \rangle$ is the stationary (asymptotic) solution to the *reaction-rate equation* (RRE),

$$\frac{d\hat{X}_3(t)}{dt} = W_-(\hat{X}_3(t)) - W_+(\hat{X}_3(t)), \quad (16)$$

wherein $\hat{X}_3(t)$ is now regarded as a *sure* (nonrandom), *continuous* (real) variable. The stationary solution $\hat{X}_3(t \rightarrow \infty) \equiv \hat{X}_3^{\text{RRE}}$ to Eq. (16) evidently satisfies

$$W_-(\hat{X}_3^{\text{RRE}}) - W_+(\hat{X}_3^{\text{RRE}}) = 0. \quad (17)$$

Notice that this is the same as Eq. (12a), except that $\langle \hat{X}_3^2(\infty) \rangle$ has now been approximated by $\langle \hat{X}_3(\infty) \rangle^2$. Inserting the explicit formulas (6) for the functions W_{\pm} and then applying the quadratic formula, we find that the only solution to Eq. (17) in the interval $[0, x_{3 \text{ max}}]$ is

$$\hat{X}_3^{\text{RRE}} = \frac{1}{2} \left\{ \left(x_{T1} + x_{T2} + \frac{c_2}{c_1} \right) - \sqrt{\left(x_{T1} + x_{T2} + \frac{c_2}{c_1} \right)^2 - 4x_{T1}x_{T2}} \right\}. \quad (18)$$

As we shall see later, for this single steady-state system \hat{X}_3^{RRE} practically always turns out to be an acceptably accurate approximation to $\langle \hat{X}_3(\infty) \rangle$. But this deterministic approach of approximating $\langle \hat{X}_3^2(\infty) \rangle \approx \langle \hat{X}_3(\infty) \rangle^2$ inevitably also approximates $\text{var}\{\hat{X}_3(\infty)\} \approx 0$, and that will usually not be adequate. A better approximation for $\langle \hat{X}_3^2(\infty) \rangle$ can be developed by using two results from the general theory of birth-death Markov processes, both of which follow directly from the recursion relation (10a). First is the result that *any relative maximum* of $\hat{P}(x_3, \infty | x_{T1}, x_{T2})$ will be the greatest integer in a *down-going root* of the function⁴

$$\alpha(x_3) \triangleq W_+(x_3 - 1) - W_-(x_3). \quad (19)$$

For the stepping functions (6), this function is

$$\begin{aligned} \alpha(x_3) &= c_1(x_{T1} - (x_3 - 1))(x_{T2} - (x_3 - 1)) - c_2x_3 \\ &= c_1x_3^2 - [c_1(x_{T1} + x_{T2} + 2) + c_2]x_3 \\ &\quad + c_1(x_{T1} + 1)(x_{T2} + 1). \end{aligned} \quad (20)$$

Since this is the equation for a concave up parabola, it will have at most one down-going root, x_3^{dgr} , and the quadratic formula reveals that root to be

$$x_3^{\text{dgr}} = \frac{1}{2} \left\{ \left(x_{T1} + x_{T2} + 2 + \frac{c_2}{c_1} \right) - \sqrt{\left(x_{T1} + x_{T2} + 2 + \frac{c_2}{c_1} \right)^2 - 4(x_{T1} + 1)(x_{T2} + 1)} \right\}. \quad (21)$$

This value therefore locates the single maximum of the function $\hat{P}(x_3, \infty | x_{T1}, x_{T2})$ in the interval $[0, x_{3 \text{ max}}]$. Comparing Eqs. (21) and (18) reveals that the difference between x_3^{dgr} and \hat{X}_3^{RRE} will be small if x_{T1} and x_{T2} are both $\gg 1$. It can further be shown from the recursion relation (10a) that the variance of the “best Gaussian fit” to the function $\hat{P}(x_3, \infty | x_{T1}, x_{T2})$ in the vicinity of the maximum at x_3^{dgr} is given by the simple formula⁴

$$\sigma_G^2(x_3^{\text{dgr}}) = \frac{W_-(x_3^{\text{dgr}})}{-\alpha'(x_3^{\text{dgr}})}, \quad (22)$$

where the prime on α denotes the derivative. Using Eqs. (6) and (20), this works out to

$$\sigma_G^2(x_3^{\text{dgr}}) = \frac{c_2 x_3^{\text{dgr}}}{-2c_1 x_3^{\text{dgr}} + c_1(x_{T1} + x_{T2} + 2) + c_2}. \quad (23)$$

Since the function $\hat{P}(x_3, \infty | x_{T1}, x_{T2})$ has only one peak, and since $\sigma_G^2(x_3^{\text{dgr}})$ gives the variance of the best Gaussian fit to that peak, then it should be reasonable to approximate $\text{var}\{\hat{X}_3(\infty)\}$ by $\sigma_G^2(x_3^{\text{dgr}})$.

Thus we are led to the following approximate formulas for the first two moments of $\hat{X}_3(\infty)$:

$$\langle \hat{X}_3(\infty) \rangle \approx \hat{X}_3^{\text{RRE}}, \quad (24a)$$

$$\langle \hat{X}_3^2(\infty) \rangle \approx \sigma_G^2(x_3^{\text{dgr}}) + \langle \hat{X}_3(\infty) \rangle^2. \quad (24b)$$

Here, Eq. (24a) refers to formula (18), and Eq. (24b) refers to formulas (23) and (21). Any first- or second-order moments involving the other fast species S_1 and S_2 , which might be needed to evaluate slow-scale propensity functions of other slow reactions besides R_3 , can be computed from these estimates by using Eqs. (14a) and (14b), etc.

B. Condition for applying the slow-scale SSA

Another useful result from the general theory of stable birth-death Markov processes concerns the *relaxation time*, or equivalently the *fluctuation time*, of the process. It characterizes the time it takes the process to reach its stationary asymptotic form, or equivalently, the time it takes the process to fully explore the region under the peak of its stationary probability density function. The theory tells us that for the virtual fast process $\hat{X}_3(t)$, this time t^f can be estimated by the formula^{2,4}

$$t^f \approx \frac{1}{-\alpha'(x_3^{\text{dgr}})} = \frac{1}{-2c_1 x_3^{\text{dgr}} + c_1(x_{T1} + x_{T2} + 2) + c_2}. \quad (25)$$

It is shown in the proof of the slow-scale approximation in Ref. 2 that the following condition must hold in order for the slow-scale SSA to be applicable: the relaxation time of

the virtual fast process must be comfortably less than the expected time to the next slow reaction. In this case, the former time is estimated by (25). The latter time, the mean time to the next R_3 reaction, can be estimated as the reciprocal of the R_3 propensity function $a_3(x) = c_3 x_3$, with x_3 replaced by its most likely value x_3^{dgr} . Thus, the condition for applying the slow-scale SSA to reactions (1) is

$$-2c_1 x_3^{\text{dgr}} + c_1(x_{T1} + x_{T2} + 2) + c_2 \gg c_3 x_3^{\text{dgr}}. \quad (26)$$

Since $x_3^{\text{dgr}} \leq x_{3 \text{ max}} = \min(x_{T1}, x_{T2})$, the magnitude of the first term on the left is always less than the magnitude of the second term; therefore, (26) will be satisfied if the inequality (2) is strong enough that c_2 is “comfortably larger” than c_3 times the average number of S_3 molecules. Condition (26) is the earlier mentioned “refinement” of condition (2). The stronger the inequality (26) is, the more accurate the slow-scale SSA will be.

C. Generating random samples

A simple approximate way to generate a random sample of $\hat{X}_3(\infty)$ is to take a cue from the normal approximation that was used in deriving formula (22), and assume that $\hat{X}_3(\infty)$ can be approximated as a *normal* random variable with mean m and variance σ^2 ,

$$\hat{X}_3(\infty) \approx \mathcal{N}(m, \sigma^2). \quad (27)$$

Since $\mathcal{N}(m, \sigma^2) = m + \sigma \mathcal{N}(0, 1)$, we can generate a random sample x_3 of $\hat{X}_3(\infty)$ under this approximation by drawing a random sample n of the “unit normal” random variable $\mathcal{N}(0, 1)$, and then making use of the estimates (24a) and (24b) for the mean and variance of $\hat{X}_3(\infty)$. More specifically, we generate the random sample as

$$x_3 = \text{rounded} \{ \hat{X}_3^{\text{RRE}} + \sqrt{\sigma_G^2(x_3^{\text{dgr}})} n \} \quad \text{in } [0, x_{3 \text{ max}}], \quad (28)$$

where the real number in braces is to be rounded to the nearest integer in the interval $[0, x_{3 \text{ max}}]$. The associated random samples of the populations of other two fast species S_1 and S_2 can then be computed from this value through the conservation relations (4a) and (4b):

$$x_1 = x_{T1} - x_3, \quad x_2 = x_{T2} - x_3. \quad (29)$$

D. Condition for using these approximations

To see under what conditions the approximations (24a), (24b), and (27) should be acceptable, let us compare the numerical predictions of these formulas to the exact results for some specific numerical examples. The exact results will be obtained by evaluating the function $\hat{P}(x_3, \infty | x_{T1}, x_{T2})$ using Eqs. (10a) and (10b).

For our first example, we take the reaction parameter values

$$c_1 = 10^{-4}, \quad c_2 = 1, \quad x_{T1} = 2200, \quad x_{T2} = 3000. \quad (30)$$

For these values, \hat{X}_3^{RRE} in Eq. (18) is found to be 50.114, which agrees extremely well with the exact value $\langle \hat{X}_3(\infty) \rangle = 50.117$ obtained from Eq. (11). The value of $\sigma_G^2(x_3^{\text{dgr}})$ in

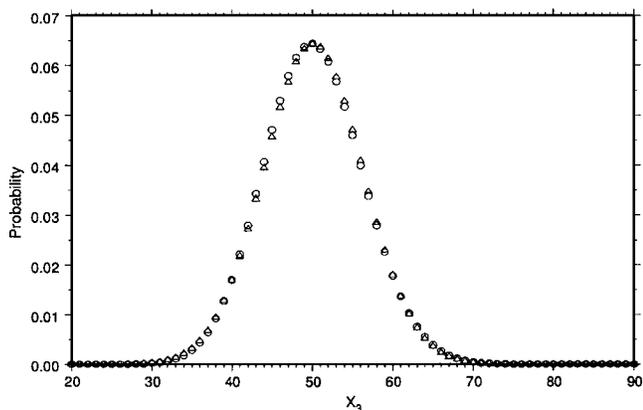


FIG. 1. The probability function $\hat{P}(x_3, \infty | x_{T1}, x_{T2})$ for the parameter values (30). The circles show the *exact* values as computed from the recursion relation (10a), while the triangles show the *normal* distribution with mean \hat{X}_3^{RRE} and variance $\sigma_G^2(x_3^{\text{dgr}})$. The normal approximation actually differs from the exact values by orders of magnitude in the extreme tails, but those differences will be unimportant for our purposes.

Eq. (23) is 38.37, which is only 0.4% higher than the exact value $\text{var}\{\hat{X}_3(\infty)\}=38.20$ obtained from Eq. (11). Figure 1 compares the exact density function $\hat{P}(x_3, \infty | x_{T1}, x_{T2})$ computed from Eqs. (10a) and (10b) with the density function for the normal random variable with mean \hat{X}_3^{RRE} and variance $\sigma_G^2(x_3^{\text{dgr}})$, which is the basis for the sampling approximation (28). On the linear scale of this graph, the two curves appear to be practically the same. A semilog plot, however, would reveal significant differences in the far tails, from 20 down to 0 and from 80 up to 220; e.g., for $x_3=20$ the normal approximation is too large by a factor of 9.3, and for $x_3=80$ it is too small by a factor of 0.36. Further out in the tails the differences are in orders of magnitude. These differences will not be important for generating random samples of $\hat{X}_3(\infty)$, because in those tails both curves are extremely small compared to 1; e.g., the Gaussian value at $x_3=20$ is 4.8×10^{-7} , and at $x_3=80$ it is 5.7×10^{-7} . For these parameter values, the approximations (24a), (24b), and (27) should be entirely satisfactory.

Consider, however, the situation for the parameter values

$$c_1 = 10^{-4}, \quad c_2 = 6, \quad x_{T1} = 10, \quad x_{T2} = 3000. \quad (31)$$

The value of \hat{X}_3^{RRE} in this case is actually quite good, disagreeing with the exact value $\langle \hat{X}_3(\infty) \rangle = 0.47613$ only in the last digit. The value of $\sigma_G^2(x_3^{\text{dgr}}) = 0.499$, however, is about 10% higher than the exact value $\text{var}\{\hat{X}_3(\infty)\} = 0.453$. Figure 2 shows that the true density function $\hat{P}(x_3, \infty | x_{T1}, x_{T2})$ is not very well approximated by that of the normal random variable with mean \hat{X}_3^{RRE} and variance $\sigma_G^2(x_3^{\text{dgr}})$.

Other tests of this kind point to the following general conclusion: If the location x_3^{dgr} of the peak of the function $\hat{P}(x_3, \infty | x_{T1}, x_{T2})$ is at least four standard deviations away from both the upper and lower limits of the interval $[0, x_{3 \text{ max}}]$, then approximating that function by a normal density function with mean \hat{X}_3^{RRE} and variance $\sigma_G^2(x_3^{\text{dgr}})$ should be quite accurate. Therefore, we shall use the approximations (24a), (24b), and (27) only when

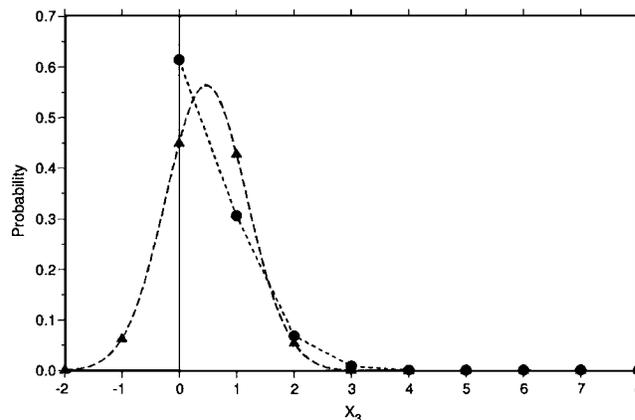


FIG. 2. The probability function $\hat{P}(x_3, \infty | x_{T1}, x_{T2})$ for the parameter values (31). The circles show the *exact* values as computed from the recursion relation (10a), while the triangles show the *normal* distribution with mean \hat{X}_3^{RRE} and variance $\sigma_G^2(x_3^{\text{dgr}})$. The normal approximation in this case would not be adequate.

$$4\sqrt{\sigma_G^2(x_3^{\text{dgr}})} \leq x_3^{\text{dgr}} \leq x_{3 \text{ max}} - 4\sqrt{\sigma_G^2(x_3^{\text{dgr}})}. \quad (32)$$

E. When condition (32) is not satisfied

Numerical tests suggest that cases that do not satisfy condition (32) all share a common feature: the interval $[0, x_{3 \text{ max}}]$, or at least the subinterval of that interval where the function $\hat{P}(x_3, \infty | x_{T1}, x_{T2})$ is “significantly different from zero,” is *small*, usually less than about 15 (cf. Fig. 2). In such circumstances, it becomes feasible to evaluate $\hat{P}(x_3, \infty | x_{T1}, x_{T2})$ using the exact formulas (10a) and (10b), and then compute the first two moments exactly and generate a random sample exactly. This is what we propose to do when condition (32) is not satisfied. In the Appendix, we outline an efficient way to carry out these exact calculations.

VI. THE SLOW-SCALE SSA FOR REACTIONS (1)

We are now in a position to describe the slow-scale stochastic simulation algorithm for simulating reactions (1). After presenting the algorithm in full, we will discuss each of its steps in detail.

Step 0 (initialization):

- Specify values for c_j ($j=1, 2, 3$) and $x_i^{(0)}$ ($i=1, \dots, 4$).
- Set $t \leftarrow 0$ and $x_i \leftarrow x_i^{(0)}$ ($i=1, \dots, 4$).
- Compute $x_{T1} = x_1 + x_3$, $x_{T2} = x_2 + x_3$, and $x_{3 \text{ max}} = \min(x_{T1}, x_{T2})$.
- Evaluate x_3^{dgr} in Eq. (21), and $\sigma_G^2(x_3^{\text{dgr}})$ in Eq. (23). Set $J_{\text{norm}} = \text{TRUE}$ or FALSE according to whether condition (32) is or is not satisfied.
- If $J_{\text{norm}} = \text{TRUE}$, evaluate \hat{X}_3^{RRE} in Eq. (18), and then set $\langle \hat{X}_3(\infty) \rangle = \hat{X}_3^{\text{RRE}}$.
- If $J_{\text{norm}} = \text{FALSE}$, compute f , Z_0 , and Z_1 according to formulas (A2)–(A6), and then set $\langle \hat{X}_3(\infty) \rangle = Z_1/Z_0$.

Step 1 (begin simulation loop): If condition (26) is satisfied, set $\bar{a}_3 = c_3 \langle \hat{X}_3(\infty) \rangle$; otherwise, abort this procedure and use the

exact SSA or tau-leaping to advance the system in time.

Step 2 (find the next slow reaction): Draw a unit-interval uniform random number r , and set $\tau=(1/\bar{a}_3)\ln(1/r)$.

Step 3 (actualize the next slow reaction):

- (a) $t \leftarrow t + \tau$,
- (b) $x_3 \leftarrow x_3 - 1$, $x_1 \leftarrow x_1 + 1$, $x_4 \leftarrow x_4 + 1$,
- (c) $x_{T2} \leftarrow x_2 + x_3$, $x_{3 \max} \leftarrow \min(x_{T1}, x_{T2})$.

Step 4 (relax the fast variables):

- (a) Evaluate x_3^{dgr} in Eq. (20), and $\sigma_G^2(x_3^{\text{dgr}})$ in Eq. (23). Set $J_{\text{norm}}=\text{TRUE}$ or FALSE according to whether condition (32) is or is not satisfied.
- (b) If $J_{\text{norm}}=\text{TRUE}$, evaluate \hat{X}_3^{RRE} in Eq. (18), and set $\langle \hat{X}_3(\infty) \rangle = \hat{X}_3^{\text{RRE}}$. Then generate a random sample x_3 according to Eq. (28).
- (c) If $J_{\text{norm}}=\text{FALSE}$, compute f , Z_0 , and Z_1 according to formulas (A2)–(A6). Set $\langle \hat{X}_3(\infty) \rangle = Z_1/Z_0$, and generate a random sample x_3 according to Eq. (A8).
- (d) Take $x_1 = x_{T1} - x_3$ and $x_2 = x_{T2} - x_3$.

Step 5 (end loop):

- (a) Plot (t, x_1, \dots, x_4) if desired.
- (b) If $t \geq t_{\text{stop}}$ then stop; otherwise, return to step 1.

Step 0 makes the necessary initializing computations. In (a) we input the parameter values, in (b) we initialize the time and state variables, and in (c) we evaluate the “constants” x_{T1} , x_{T2} , and $x_{3 \max}$. In (d) we compute the location x_3^{dgr} and approximate width $2\sqrt{\sigma_G^2(x_3^{\text{dgr}})}$ of the peak in the probability density function for $\hat{X}_3(\infty)$, and then we check to see if that peak is far enough from the end points of the interval $[0, x_{3 \max}]$ that a normal (Gaussian) approximation would be satisfactory. If so ($J_{\text{norm}}=\text{TRUE}$), we approximate the mean of $\hat{X}_3(\infty)$ by the stationary solution \hat{X}_3^{RRE} of the RRE. If not ($J_{\text{norm}}=\text{FALSE}$), we evaluate the mean of $\hat{X}_3(\infty)$ using the exact procedure described in the Appendix.

Step 1 checks to see if the relaxation time of the virtual fast process is less than the expected time to the next slow (R_3) reaction. If it is, we prepare to step to the next R_3 reaction by evaluating its slow-scale propensity function, \bar{a}_3 . If it is not, we should not use the slow-scale SSA.

Step 2 generates the time to the next R_3 reaction. We use the standard SSA time step formula, except the actual R_3 propensity function a_3 is replaced by the slow-scale R_3 propensity function \bar{a}_3 , in accordance with the slow-scale approximation (see Sec. III).

Step 3 actualizes the next R_3 reaction by (a) updating the time variable, (b) changing the species populations according to the stoichiometry of reaction R_3 , and (c) computing the postreaction values of the parameters x_{T2} and $x_{3 \max}$. In connection with the last, note that the state change (b) decreases the value of parameter $x_{T2} = x_2 + x_3$ by 1 but leaves the value of parameter $x_{T1} = x_1 + x_3$ unchanged.

Step 4 “relaxes” the fast state variables by sampling them at a time after the R_3 reaction that is typically short

compared to the time step τ just taken, but long enough that the virtual fast process will have relaxed to its asymptotic form. The relaxed values for the fast state variables are generated by first picking x_3 as a random sample of $\hat{X}_3(\infty)$, using the postreaction parameter values just computed in step 3(c), and then choosing x_1 and x_2 in accordance with the conservation relations. Generating the sample value x_3 is the key here. To do that, we begin in (a) by computing the location x_3^{dgr} and approximate width $2\sqrt{\sigma_G^2(x_3^{\text{dgr}})}$ of the peak in the probability density function for $\hat{X}_3(\infty)$, and checking to see if that peak is far enough from the end points of the interval $[0, x_{3 \max}]$ that a normal (Gaussian) approximation is warranted. If so ($J_{\text{norm}}=\text{TRUE}$), we approximate $\hat{X}_3(\infty)$ as a normal random variable with mean \hat{X}_3^{RRE} and variance $\sigma_G^2(x_3^{\text{dgr}})$, and we generate the random sample x_3 accordingly. If not ($J_{\text{norm}}=\text{FALSE}$), we generate the random sample of $\hat{X}_3(\infty)$ exactly, using the procedure described in the Appendix, and in the process we also calculate $\langle \hat{X}_3(\infty) \rangle$ exactly. Either way, we leave step 4 with a postreaction value for $\langle \hat{X}_3(\infty) \rangle$ that can be used to evaluate \bar{a}_3 on the next pass through step 1.

Step 5 reads out the current state, and then either loops back to simulate the next R_3 reaction or else stops.

VII. NUMERICAL EXAMPLES

To illustrate the slow-scale SSA for reactions (1), we first take the parameter values

$$c_1 = 10^{-4}, \quad c_2 = 1, \quad c_3 = 10^{-4};$$

$$x_1^{(0)} = 220, \quad x_2^{(0)} = 3000, \quad x_3^{(0)} = x_4^{(0)} = 0. \quad (33)$$

Figure 3 shows the results of an exact SSA run made using these values, with the molecular populations of all the species being plotted out immediately after the occurrence of each R_3 reaction. After an initial rapid transient that brings the free enzyme population X_1 down from 220 and the enzyme-substrate complex population X_3 up from 0, a quasi-equilibrium is established between those two species. Thereafter, the substrate population X_2 slowly decreases, and the product population X_4 slowly increases, until all 3000 of the initial substrate molecules have been converted into product molecules, and the free enzyme population X_1 has returned to its initial value 220. The 3000th R_3 reaction event actually occurred at time 4.25×10^6 in this run. This exact simulation of the 3000 R_3 reactions required simulating 58.4×10^6 reactions in all, so successive trajectory dots here are separated by, on average, about 19 000 fast reactions (R_1 or R_2).

For the parameter values (33), the left-hand side of condition (26) is found to be about 260 times larger than the right-hand side, and that favorable imbalance only improves as the simulation proceeds since x_3^{dgr} , which approximately locates the center of the X_3 trajectory in Fig. 3(a), decreases. The slow-scale SSA should therefore be applicable. Figure 4 shows the results of a slow-scale SSA run, with the molecular populations again plotted out immediately after each R_3 reaction. The trajectories in Fig. 4 are statistically indistinguishable from those in Fig. 3, at least on the scale of these

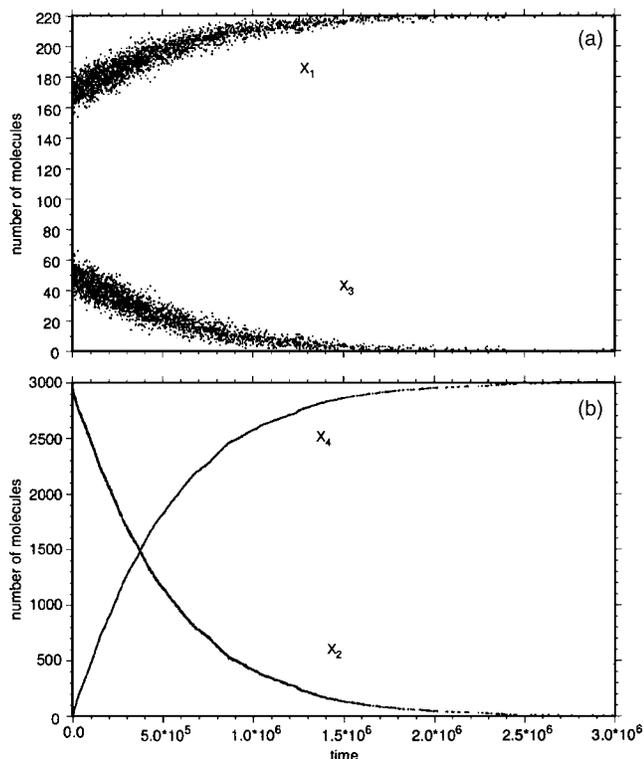


FIG. 3. Trajectories obtained in an *exact* SSA run of reactions (1) for the parameter values (33). A total of over 58×10^6 reactions were simulated here, but the species populations are plotted out only after each occurrence of a slow (R_3) reaction, of which there were 3000.

figures, which means that the slow-scale SSA is performing very well. Whereas the exact SSA run in Fig. 3 had to simulate 58.4×10^6 reactions, the approximate slow-scale SSA run in Fig. 4 simulated only 3000 reactions. The run time for Fig. 4 was about 1.75 s, whereas the run time for Fig. 3 was 27.75 min, giving a real-time speedup factor for the slow-scale SSA of about 950. By monitoring the value of J_{norm} in the slow-scale SSA run, it was determined that step 4(b) was used to generate the X_3 trajectory from time 0 until time 7.8×10^5 , and thereafter step 4(c) was used; we note that the transition between those two computational methods for generating the X_3 trajectory cannot be detected in the data.

Our second example illustrates a case in which the average number of molecules of the enzyme-substrate complex S_3 remains less than one for an extended period of time. The plots in Fig. 5 show the results of an exact SSA run for the parameter values

$$c_1 = 10^{-4}, \quad c_2 = 6, \quad c_3 = 10^{-3};$$

$$x_1^{(0)} = 10, \quad x_2^{(0)} = 3000, \quad x_3^{(0)} = x_4^{(0)} = 0. \quad (34)$$

Figure 5(a) shows the steady decline of the substrate (S_2) population, and the steady increase in the product (S_4) population, with the data points being taken immediately after each R_3 reaction. There are 1172 R_3 reaction events in the time span shown here, and a total of 13.7×10^6 reaction events in all, so over 11 000 fast reaction events (R_1 and R_2) typically occur between successive points in Fig. 5(a). Figure 5(b) shows the population of the enzyme-substrate complex (S_3) taken at these same times. Note that most of the time X_3

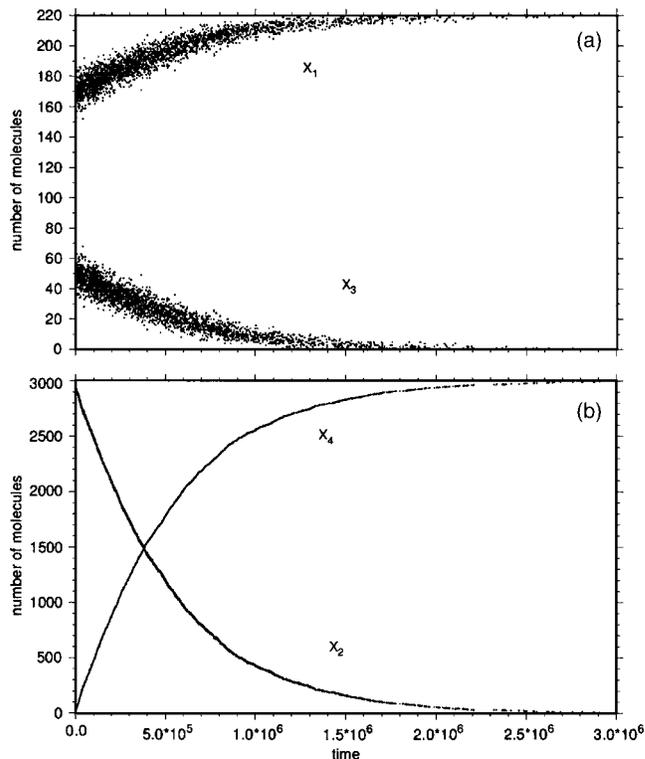


FIG. 4. Trajectories obtained in an *approximate slow-scale* SSA run of reactions (1) for the parameter values (33). Only R_3 reactions were explicitly simulated here. The trajectories are statistically indistinguishable from the exact SSA trajectories generated in Fig. 3, but this simulation ran about 900 times faster.

is either 0 or 1, but it is practically never greater than 2.

Figure 5(c) shows the S_3 population for the *same* SSA run, but now taken at *equally spaced time intervals* of $\Delta t = 2.56 \times 10^3$, a value chosen to give the same total number of plotted points as in Figs. 5(a) and 5(b). The slight difference between the X_3 plots in Figs. 5(b) and 5(c) [note, for instance, the higher number of $X_3=3$ events in Fig. 5(c)] reflects the fact that the X_3 population is not statistically independent of the times at which R_3 reactions occur. That becomes clear when we recognize that an R_3 reaction is twice as likely to occur when $X_3=2$ as when $X_3=1$, and an R_3 reaction cannot occur at all when $X_3=0$. The plot in Fig. 5(c) provides a sampling of the X_3 population at times that are *uncorrelated* with the times of the R_3 reactions, and it arguably represents a more “natural” sampling than that shown in Fig. 5(b). Of course, both plots are “correct.”

For the parameter values (34), the left-hand side of condition (26) is found to be larger than the right-hand side by about four orders of magnitude, and that favorable imbalance persists for the time span simulated since x_3^{dgr} stays fairly constant (at about 0.5) during this time. The slow-scale SSA should therefore be applicable. Figure 6 shows the results of a slow-scale SSA run. The populations are plotted after each simulated (R_3) reaction, and there were 1165 such reactions in the interval shown. We see that the X_2 and X_4 plots in Fig. 6(a) are statistically indistinguishable from the SSA-generated plots in Fig. 5(a). We also see that the X_3 plot in Fig. 6(b) more closely resembles the equal-time plot in Fig. 5(c) than the R_3 reaction-time plot in Fig. 5(b). This is to be

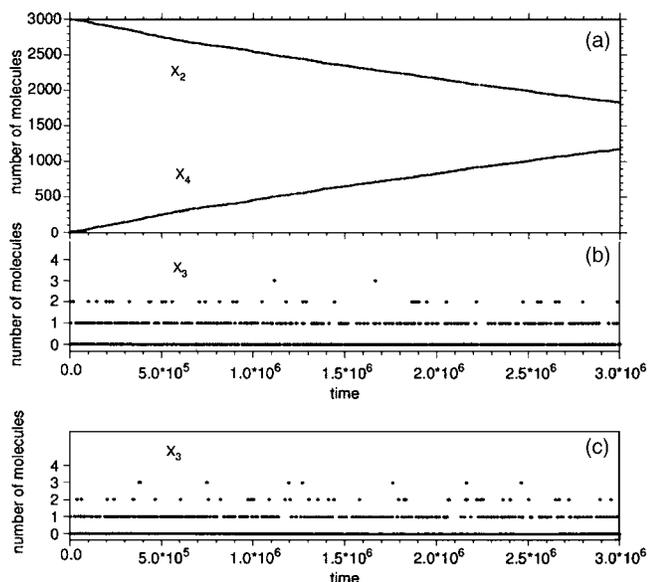


FIG. 5. Trajectories obtained in an *exact* SSA run of reactions (1) for the parameter values (34). A total of 13.7×10^6 reactions were simulated, but the species populations in (a) and (b) are plotted out only after each occurrence of a slow (R_3) reaction, of which there were 1172. The X_3 plot in (c) is from the same run, but here the points are plotted at *equally spaced time intervals*, with the interval spacing chosen to give the same total number of plotted points. The explanation for the slight but statistically significant differences between the (b) and (c) plots is discussed in the text.

expected from the way in which the slow-scale SSA relaxes the fast variables after each R_3 reaction before sampling them. For this particular slow-scale SSA run, the X_3 values were always generated using the exact method in step 4(c), rather than the faster approximate method in step 4(b); nevertheless, the actual speedup factor of the slow-scale SSA run relative to the x_4 exact SSA run here was found to be about 400.

VIII. STREAMLINING THE ALGORITHM

A close examination of the computational procedure described in Sec. VI will reveal that the three fast variables x_1 , x_2 , and x_3 affect the slow reaction R_3 only through the values of the two parameters $x_{T1} \equiv x_1 + x_3$ and $x_{T2} \equiv x_2 + x_3$, and the values of those two parameters do not get changed by the relaxation step 4. Therefore, if we are content to track the status of the three fast variables only through the two variables x_{T1} and x_{T2} , we can speed up the simulation procedure by *omitting* the random samplings of $\hat{X}_3(\infty)$ in steps 4(b) and 4(c).

The specific changes required to “streamline” the algorithm in this way are as follows: In step 3(b), delete the two updates to x_3 and x_1 . In step 3(c), change the x_{T2} update to read $x_{T2} \leftarrow x_{T2} - 1$. In steps 4(b) and 4(c), delete the “generate” instructions. Omit step 4(d) entirely. Finally, in step 5(b), plot (t, x_{T1}, x_{T2}, x_4) .

This streamlined stochastic simulation procedure evidently ignores the fast variables x_1 , x_2 , and x_3 (the enzyme, substrate, and enzyme-substrate complex populations), and focuses solely on the slow variable x_4 (the product species population). We shall see in Sec. X that simulating the evolution of the system in this way is essentially equivalent to

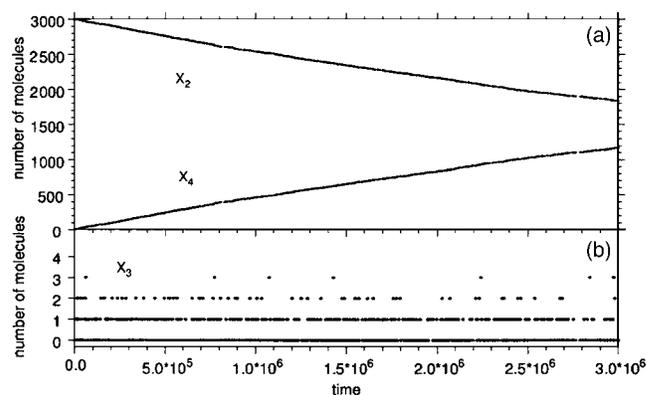


FIG. 6. Trajectories obtained in an *approximate slow-scale* SSA run of reactions (1) for the parameter values (34). Only R_3 reactions were explicitly simulated here, and they numbered 1165. The X_2 and X_4 trajectories in (a) are statistically indistinguishable from the exact SSA trajectories in Fig. 5(a). The X_3 trajectory in (b) more closely resembles the one in Fig. 5(c) than the one in Fig. 5(b), for reasons explained in the text. This slow-scale SSA simulation ran about 400 times faster than the SSA simulation in Fig. 5.

invoking the standard Michaelis–Menten formula for the deterministic rate of the slow reaction. Note that stochastically correct values of the fast variables x_1 , x_2 , and x_3 can be regained at any time, and without any loss of accuracy, by simply restoring steps 4(b)–4(d) to their previous forms.

IX. INCORPORATING ADDITIONAL SLOW REACTION CHANNELS

It is straightforward to generalize the simulation procedure for reactions (1) described in Sec. VI to accommodate K additional slow reactions, R_4, \dots, R_{3+K} . These additional slow reactions might involve any of the species S_1, \dots, S_4 already present in reactions (1), as well as any new (slow) species. For example, a reaction could be added that slowly creates new molecules of the substrate species S_2 , or that slowly consumes existing molecules of the free enzyme species S_1 .

Condition (26) for applying the slow scale SSA, which gets tested in step 1, will have to be modified so that it compares the relaxation time of the virtual fast process to the expected time to the next slow reaction of any kind. The needed modification is the replacement of a_3 on the right-hand side of (26) with $\sum_{j=3}^{3+K} a_j$, evaluated at the current values of the slow species populations and the most likely values of the fast species populations.

The formulas for the slow-scale propensity functions of the new slow reactions are determined as follows: If the true propensity function a_j of slow reaction R_j does not involve any of the fast species variables (x_1, x_2, x_3), then the slow-scale propensity function \bar{a}_j of R_j would be identical to a_j . If a_j depends on a *single* fast species variable, x_1 or x_2 or x_3 , then \bar{a}_j is constructed by simply replacing that fast variable by $x_{T1} - \langle \hat{X}_3(\infty) \rangle$ or $x_{T2} - \langle \hat{X}_3(\infty) \rangle$ or $\langle \hat{X}_3(\infty) \rangle$, the values of which are available in the existing algorithm through the current estimate of $\langle \hat{X}_3(\infty) \rangle$. If a_j depends on the product of two fast variables, then evaluating \bar{a}_j would require knowing also $\langle \hat{X}_3^2(\infty) \rangle$; for example, the product $x_1 x_3$ would get replaced by the expression in Eq. (14c). In that case, we would

need to augment steps 0(e) and 4(b) to also compute $\langle \hat{X}_3^2(\infty) \rangle = \sigma_G^2(x_3^{\text{der}}) + \langle \hat{X}_3(\infty) \rangle^2$, and to augment steps 0(f) and 4(c) to also compute Z_2 and $\langle \hat{X}_3^2(\infty) \rangle = Z_2/Z_0$. Step 1 must then be expanded to evaluate *all* the slow-scale propensity functions $\bar{a}_3, \dots, \bar{a}_{3+K}$ according to these rules, and to also compute their sum

$$\bar{a}_0 \equiv \bar{a}_3 + \dots + \bar{a}_{3+K}.$$

Step 2 must be modified so that it determines not only when the next slow reaction occurs but also which slow reaction that will be. In direct analogy with the standard SSA, the new step 2 would read:

Step 2: Draw two unit-interval uniform random numbers, r_1 and r_2 . Take $\tau = (1/\bar{a}_0)\ln(1/r_1)$, and take j to be the *smallest* integer that satisfies $\sum_{j'=3}^j \bar{a}_{j'} \geq r_2 \bar{a}_0$.

Step 3(b) must be changed so that it implements the state change induced by the particular slow reaction R_j whose index was found in step 2. Step 3(c) must additionally update $x_{T1} \leftarrow x_1 + x_3$ if the occurrence of any of the new slow reactions is capable of altering the value of that parameter.

Finally, the actions in steps 0(a), 0(b), and 5(a) must be extended to all reactions and all species.

The resulting algorithm, which we note leaves the core computations involving the virtual fast process unchanged, will simulate the evolution of the expanded system *one slow reaction at a time*.

X. CONNECTION TO MICHAELIS–MENTEN

The traditional *deterministic* approach to the enzyme-substrate reactions (1) is to invoke the Michaelis–Menten formula for the rate ν at which molecules of product species S_4 are being produced. The deterministic reaction-rate equation for reactions (1) gives ν as

$$\nu = \frac{dX_4}{dt} = c_3 X_3, \quad (35)$$

where X_i is now viewed as a sure (nonrandom) variable giving the instantaneous number of S_i molecules in the system. Deriving the Michaelis–Menten formula for ν evidently requires obtaining an estimate for X_3 , and a quick review of how that is done will enable us to understand the connection between the Michaelis–Menten approach to reactions (1) and the slow-scale SSA approach presented in this paper.

There are actually two ways of obtaining the Michaelis–Menten formula: one can either invoke the partial equilibrium (or rapid equilibrium) approximation, or one can invoke the quasi-steady-state approximation. In the *partial (rapid) equilibrium approximation*, one assumes that the fast reactions R_1 and R_2 are in approximate equilibrium with each other. Equating (approximately) the deterministic rates at which those two reactions occur gives

$$c_1 X_1 X_2 \doteq c_2 X_3. \quad (36)$$

In the *quasi-steady-state approximation*, one assumes that the population of the enzyme-substrate complex species S_3 is approximately constant. Since the reaction-rate equation

gives dX_3/dt as $c_1 X_1 X_2 - c_2 X_3 - c_3 X_3$, then setting that rate to (approximately) zero gives

$$c_1 X_1 X_2 \doteq (c_2 + c_3) X_3.$$

Under condition (2), this relation reduces to Eq. (36). So we see that the partial (rapid) equilibrium approximation and the quasi-steady-state approximation lead to the same result (36) when condition (2) holds.

The Michaelis–Menten formula is a simple consequence of Eq. (36) and the fact that reactions (1) conserve “enzyme units;” i.e., the number of free enzyme molecules ($=X_1$) plus the number of enzyme-substrate complexes ($=X_3$) remains constant [cf. Eq. (4a)]:

$$X_1 + X_3 = x_{T1} \quad (\text{const}). \quad (37)$$

Solving Eq. (37) for X_1 and substituting the result into Eq. (36) gives

$$c_1(x_{T1} - X_3)X_2 \doteq c_2 X_3,$$

and then solving this for X_3 gives

$$X_3 \doteq \frac{x_{T1} X_2}{(c_2/c_1) + X_2}. \quad (38)$$

Putting this into Eq. (35) yields the *Michaelis–Menten formula*⁶

$$\nu = \frac{c_3 x_{T1} X_2}{(c_2/c_1) + X_2}. \quad (39)$$

The connection between this well-known result and our slow-scale SSA can be understood by recognizing that, because of condition (2), reaction R_3 will be occurring at a negligibly small rate compared to reactions R_1 and R_2 . Since the latter two reactions *by themselves* conserve “substrate units,” then condition (2) implies the *approximate* conservation relation [cf. Eq. (4b)]

$$X_2 + X_3 \doteq x_{T2} \quad (\text{const}). \quad (40)$$

Solving this for X_2 and then substituting the result into Eq. (38) gives

$$X_3 \doteq \frac{x_{T1}(x_{T2} - X_3)}{(c_2/c_1) + (x_{T2} - X_3)}.$$

Clearing the fraction and collecting terms produces

$$X_3^2 - \left(\frac{c_2}{c_1} + x_{T1} + x_{T2} \right) X_3 + x_{T1} x_{T2} \doteq 0,$$

and the solution to this quadratic equation for X_3 is precisely the quantity in Eq. (18):

$$X_3 \doteq \hat{X}_3^{\text{RRE}}. \quad (41)$$

By combining this result with the definition of ν in Eq. (35), and the approximate formula for the mean of $\hat{X}_3(\infty)$ in Eq. (24a), and finally the definition (5) of the slow-scale propensity function for reaction R_3 , we conclude that

$$\nu = c_3 X_3 \doteq c_3 \hat{X}_3^{\text{RRE}} \doteq c_3 \langle \hat{X}_3(\infty) \rangle = \bar{a}_3(x). \quad (42)$$

In words, under condition (2), the Michaelis–Menten estimate of the rate of production of product is approximately equal to the slow-scale propensity function for reaction R_3 . Therefore, *the stochastic rate of production of product species in the slow-scale SSA is, for all practical purposes, equal to the deterministic rate that is prescribed by the Michaelis–Menten formula.*

The result (42) suggests that our earlier analysis amounts to a *derivation* of the deterministic Michaelis–Menten result from a purely *stochastic* point of view. More specifically, the slow-scale approximation that was proved in Ref. 2 leads to the conclusion that if condition (26) is satisfied, then the propensity function of reaction R_3 on the *slow time scale* is effectively the value given by the Michaelis–Menten formula. The stochastic approach, in addition to providing the validity criterion (26), also allows us to *simulate* the fluctuations in the fast species S_1 , S_2 , and S_3 , which of course a deterministic analysis cannot do. The “streamlined” stochastic simulation procedure described in Sec. VIII does *not* simulate the populations of the fast species, and in that sense is very close to the usual deterministic Michaelis–Menten treatment. Whereas conventional deterministic derivations of the Michaelis–Menten formula (39) simply *assume* that the fluctuations in the fast variable X_2 can be safely ignored, our stochastic derivation provides *a priori* assurance of that. Finally, the slow-scale SSA approach automatically provides an accurate discrete-stochastic description of the creation of product molecules by reaction R_3 , for those circumstances where that turns out to be important.

XI. SUMMARY AND CONCLUSIONS

The enzyme-substrate reaction set (1) is a common motif in cellular systems, where it often satisfies condition (2). We have shown here how dramatic increases in the speed of stochastically simulating reactions (1) under condition (2) can be achieved, without noticeable loss of computational accuracy, by using the slow-scale SSA that was developed in Ref. 2. We described the application of the slow-scale SSA to this system in detail, and we indicated how the procedure can be adapted to the case in which reactions (1) are embedded in a larger set of (slow) reactions. Finally, we showed that the slow-scale SSA approach to reactions (1) is a natural stochastic extension of the deterministic Michaelis–Menten approach, and, in fact, provides a broader perspective on the Michaelis–Menten formula that helps clarify its rationale.

The speedup provided by the slow-scale SSA for reactions (1) generally depends on the strength of the inequality (2), or more accurately, the inequality (26): the stronger those inequalities are, the greater the speedup will be, and the greater the accuracy will be as well. If those inequalities do *not* hold, the three reaction channels in (1) will be firing at comparable rates. In that case the system will not be stiff, and the slow-scale SSA cannot be applied. If at least some of the reactant populations are large, substantial but less dramatic speedups over the exact SSA should still be possible by using the explicit tau-leaping procedure.⁷ If all the reac-

tant populations are small, the SSA itself will probably be the most efficient simulation procedure.

ACKNOWLEDGMENTS

We thank Carol Gillespie for assistance in performing the numerical calculations and creating the figures. This work was supported in part by the U.S. Air Force Office of Scientific Research and the California Institute of Technology under DARPA Award No. F30602-01-2-0558, and in part by the Department of Energy under DOE Award No. DE-FG02-04ER25621. Two of the authors (Y.C. and L.P.) received additional support from the National Science Foundation under NSF Award No. CTS-0205584, and from the Institute for Collaborative Biotechnologies through Grant No. DAAD19-03-D-0004 from the U.S. Army Research Office. Another author (D.G.) received additional support through the Molecular Sciences Institute (Contract No. 244725) from Sandia National Laboratories and the Department of Energy’s “Genomes to Life” Program.

APPENDIX: EXACT COMPUTATIONS FOR $\hat{X}_3(\infty)$

In some circumstances it is feasible to compute the first two moments of $\hat{X}_3(\infty)$, and to generate random samples of $\hat{X}_3(\infty)$, *exactly*. The key to doing this is to compute the probability density function $\hat{P}(x_3, \infty | x_{T1}, x_{T2})$ of $\hat{X}_3(\infty)$ according to the exact equations (10a) and (10b). To that end, we begin by using the current values of X_1 , X_2 , and X_3 to evaluate the parameters for this distributions,

$$x_{T1} = X_1 + X_3, \quad x_{T2} = X_2 + X_3, \quad x_{3 \max} = \min(x_{T1}, x_{T2}). \quad (A1)$$

Next we compute x_3^{dgr} in Eq. (21), which locates the (single) peak in the function $\hat{P}(x_3, \infty | x_{T1}, x_{T2})$; more precisely, $\hat{P}(x_3, \infty | x_{T1}, x_{T2})$ achieves its largest value for *integer* values of x_3 at *the greatest integer in x_3^{dgr}* , which we denote by

$$\hat{x}_3 \equiv [x_3^{\text{dgr}}]. \quad (A2)$$

Although the function $\hat{P}(x_3, \infty | x_{T1}, x_{T2})$ is strictly positive for all $x_3 \in [0, x_{3 \max}]$, it gets dramatically closer and closer to zero as we move further and further away from the peak at x_3 . For numerical work, we can safely confine ourselves to a usually much smaller subinterval $[x_{3L}, x_{3H}]$ of $[0, x_{3 \max}]$ which is defined by the condition that everywhere *outside* this subinterval, $\hat{P}(x_3, \infty | x_{T1}, x_{T2})$ is less than some very small fraction, say 10^{-7} , of its maximum value at \hat{x}_3 .

To capture the information contained in Eqs. (10a) and (10b) for $x_3 \in [x_{3L}, x_{3H}]$, we let $f(x_3)$ denote the *unnormalized* function $\hat{P}(x_3, \infty | x_{T1}, x_{T2})$, and we execute the following computational steps, where the functions $W_-(x_3)$ and $W_+(x_3)$ are as defined in Eqs. (6): First, we set

$$f(\hat{x}_3) = 1. \quad (A3)$$

Then, if $\hat{x}_3 > 0$, we compute for $x_3 = \hat{x}_3 - 1, \hat{x}_3 - 2, \dots$, in that order,

$$f(x_3) = \frac{W_-(x_3 + 1)}{W_+(x_3)} f(x_3 + 1), \quad (\text{A4})$$

until either $x_3=0$ or $f(x_3) \leq 10^{-7}$, after which we set $x_{3L}=x_3$ and stop; or, if $\hat{x}_3=0$, we simply set $x_{3L}=0$. Similarly, if $\hat{x}_3 < x_{3 \max}$, we compute for $x_3=\hat{x}_3+1, \hat{x}_3+2, \dots$, in that order

$$f(x_3) = \frac{W_+(x_3 - 1)}{W_-(x_3)} f(x_3 - 1), \quad (\text{A5})$$

until either $x_3=x_{3 \max}$ or $f(x_3) \leq 10^{-7}$, after which we set $x_{3H}=x_3$ and stop; or, if $\hat{x}_3=x_{3 \max}$, we simply set $x_{3H}=x_{3 \max}$. Concurrent with the computations in (A3)–(A5), we also compute the following three quantities as running sums:

$$Z_k \equiv \sum_{x_3=x_{3L}}^{x_{3H}} x_3^k f(x_3) \quad (k=0, 1, 2). \quad (\text{A6})$$

Since $\hat{P}(x_3, \infty | x_{T1}, x_{T2}) = f(x_3)/Z_0$, we can now compute the first two moments of $\hat{X}_3(\infty)$ according to the exact formula (11) simply by taking

$$\langle \hat{X}_3^k(\infty) \rangle = Z_k/Z_0 \quad (k=1, 2). \quad (\text{A7})$$

If R_3 is the *only* slow reaction being considered, the $k=2$ calculations in Eqs. (A6) and (A7) can be omitted, since the value of $\langle \hat{X}_3^2(\infty) \rangle$ will not be needed.

A random sample x_3 of $\hat{X}_3(\infty)$ can be generated by using the exact “inversion” Monte Carlo generating method, in which one draws a unit-interval uniform random number r and then takes x_3 to be the *smallest* integer that satisfies

$$\sum_{x'_3=x_{3L}}^{x_3} \hat{P}(x'_3, \infty | x_{T1}, x_{T2}) \geq r.$$

Since $\hat{P}(x_3, \infty | x_{T1}, x_{T2}) = f(x_3)/Z_0$, we can therefore compute the random sample x_3 as the *smallest* integer that satisfies

$$\sum_{x'_3=x_{3L}}^{x_3} f(x'_3) \geq rZ_0. \quad (\text{A8})$$

The time required to make all the above computations will be roughly proportional to $(x_{3H}-x_{3L})$, the number of steps in the two do-loops for the computations of (A3)–(A6) and of (A8). For the circumstances in which we are proposing to make these exact computations inside the main simulation loop, $(x_{3H}-x_{3L})$ will usually be less than 15, so these computations ought to be feasible.

¹D. T. Gillespie, J. Comput. Phys. **22**, 403 (1976); J. Phys. Chem. **81**, 2340 (1977); Physica A **188**, 404 (1992).

²Y. Cao, D. T. Gillespie, and L. R. Petzold, J. Chem. Phys. **122**, 014116 (2005).

³C. V. Rao and A. P. Arkin, J. Chem. Phys. **118**, 4999 (2003).

⁴D. T. Gillespie, *Markov Processes: An Introduction for Physical Scientists* (Academic, San Diego, 1992). See especially Chap. 6 and Appendix D.

⁵I. G. Darvey, B. W. Ninham, and P. J. Staff, J. Chem. Phys. **45**, 2145 (1966).

⁶L. Michaelis and M. L. Menten, Biochem. Z. **49**, 333 (1913); K. M. Plowman, *Enzyme Kinetics* (McGraw-Hill, New York, 1971).

⁷D. T. Gillespie, J. Chem. Phys. **115**, 1716 (2001); D. T. Gillespie and L. R. Petzold, *ibid.* **119**, 8229 (2003).