

A “partitioned leaping” approach for multiscale modeling of chemical reaction dynamics

Leonard A. Harris* and Paulette Clancy†

School of Chemical and Biomolecular Engineering, Cornell University, Ithaca, NY 14853, USA

(Dated: February 2, 2008)

We present a novel multiscale simulation approach for modeling stochasticity in chemical reaction networks. The approach seamlessly integrates exact-stochastic and “leaping” methodologies into a single *partitioned leaping* algorithmic framework. The technique correctly accounts for stochastic noise at significantly reduced computational cost, requires the definition of only three model-independent parameters and is particularly well-suited for simulating systems containing widely disparate species populations. We present the theoretical foundations of partitioned leaping, discuss various options for its practical implementation and demonstrate the utility of the method via illustrative examples.

I. INTRODUCTION

Stochastic simulations of chemical reaction networks have become increasingly popular recently, largely due to the observation that stochastic noise plays a critical role in biological function.^{1,2,3,4,5,6,7,8} The issue is relevant in other scientific fields as well, such as microelectronics processing, where statistical variations in dopant profiles can profoundly affect the performance of nanoscale semiconductor devices.^{9,10,11} Gillespie’s stochastic simulation algorithm (SSA),¹² in particular, has found widespread use in computational biology. The method is a kinetic Monte Carlo approach that produces time-evolution trajectories correctly accounting for the inherent stochasticity associated with molecular interactions. Detailed accuracy is achieved by explicitly simulating every reaction occurrence within a system. The method is computationally expensive as a result, however, and practical application is limited to only very small systems.

Motivated by this fact, considerable effort has been undertaken recently to develop accelerated simulation methods capable of correctly accounting for stochastic noise but at significantly reduced computational cost. Broadly speaking, these endeavors can be divided into three categories: (i) algorithmic advances to increase the efficiency of exact-stochastic methods,^{13,14,15} (ii) “leaping” techniques in which efficiency is enhanced by ignoring the exact moments at which reaction events occur,^{16,17,18,19,20,21,22,23,24} and (iii) “partitioned” methods in which sets of reactions are divided into various classifications, such as “fast” and “slow,” and treated either by applying appropriate numerical techniques to each subset (“methodology coupling” or “hybrid”)^{25,26,27,28,29,30} or by reducing the model to incorporate the *effects* of the fast reactions into the dynamics of the slow (“model reduction”).^{31,32,33,34,35,36}

In this article, we present a new simulation method for modeling chemical reaction dynamics that integrates exact-stochastic, leaping and methodology coupling methods into a single multiscale algorithmic framework. The fundamental theory underlying Gillespie’s τ -leap approach^{16,37} is used to formulate a theoretically

justifiable partitioning scheme. The partitioning is based on the number of reaction firings expected within a calculated time interval and partitioning and time step determination are thus inextricably linked. Based on the value of the time step each reaction is then “classified” at one of various levels of description. Species populations are updated for coarsely classified reactions using the leaping formulas introduced by Gillespie,¹⁶ while rare events are treated via the methods of Gibson and Bruck’s exact-stochastic Next Reaction method.¹³ Overall, the technique efficiently simulates systems containing widely disparate species populations using rigorously derived classification criteria and requiring minimal user intervention.

We begin in Sec. II by reviewing the theoretical foundations of exact-stochastic simulation and τ -leaping. This provides a basis for discussing partitioned leaping in Sec. III as well as various options for its practical implementation. In Sec. IV we present two illustrative examples demonstrating the utility of the method: one inspired by biology and a clustering example relevant to materials and atmospheric sciences. Finally, in Sec. V we summarize the attributes of the method, discuss possible modifications, and draw conclusions regarding its place among the many alternative techniques that have been proposed.

II. BACKGROUND

As is customary, we consider a well-mixed system of N molecular species $\{S_1, \dots, S_N\}$ interacting via M reaction channels $\{R_1, \dots, R_M\}$ in a volume Ω at constant temperature. The state of the system is described by the vector $\mathbf{X}(t)$, where $X_i(t)$ represents the population of species S_i at time t . Each reaction channel R_μ has associated with it a propensity function a_μ and a state-change (or stoichiometric) vector $\mathbf{z}_\mu = (z_{\mu 1}, \dots, z_{\mu N})$.³⁸ The propensity a_μ is defined such that $a_\mu dt$ gives the probability that reaction R_μ will fire once during the infinitesimal time interval dt . In other words, a_μ is the stochastic analog to the deterministic *reaction rate*. As

such, a_μ can be written as the product of a stochastic rate constant c_μ (which is related to the deterministic rate constant via a simple scaling by the system volume)¹² and a combinatorial factor h_μ , which represents the number of distinct ways in which a realization of R_μ can occur. In general, h_μ is a simple function of the reactant species populations for R_μ .¹²

A. Exact stochastic simulation

Given the definitions above, Gillespie has developed a simulation methodology for modeling chemical reaction dynamics that “exactly” accounts for the stochastic nature of the process.¹² The stochastic simulation algorithm, or SSA, is exact in the sense that it produces *possible* time-evolution trajectories that are consistent with the underlying *chemical Master Equation* governing the physical process.³⁹

The SSA operates by generating, at each simulation step, random samples of reaction times (τ) and types (μ) and advancing the system forward in time accordingly. Two alternative, yet equivalent, strategies exist for doing so. The first, dubbed the *Direct method* (DM), generates τ according to¹²

$$\tau = -\ln(r_1)/a_0, \quad (1)$$

while μ is the integer satisfying the relationship

$$\sum_{\nu=1}^{\mu-1} a_\nu < a_0 r_2 \leq \sum_{\nu=1}^{\mu} a_\nu, \quad (2)$$

where $a_0 \equiv \sum_\nu a_\nu$ and r_1 and r_2 are unit-uniform random numbers between 0 and 1.

The second approach considers each reaction in the system on an *individual* basis and asks the question, “When will reaction R_μ next fire *assuming* that no other reactions fire first?” Values of such “tentative” next-reaction times, τ_μ^{ES} ,⁴⁰ are then generated for each reaction. τ is set equal to the smallest of these and μ to the corresponding reaction since this is the only one for which the assumption that no other reactions fire first actually holds. This approach was originally developed by Gillespie¹² and dubbed the *First Reaction method* (FRM). Tentative next-reaction times are calculated via

$$\tau_\mu^{\text{ES}} = -\ln(r_\mu)/a_\mu, \quad (3)$$

where r_μ is a unit-uniform random number.

This approach has recently been modified by Gibson and Bruck.¹³ The Next Reaction method (NRM) is more computationally efficient in that a rigorous transformation formula is employed that significantly reduces the number of random number generations required during the course of a simulation. Once a reaction has fired, a new value of τ_μ^{ES} is generated *for that reaction only* using Eq. (3). For all other reactions

$$\tau_\mu^{\text{ES}} = (a'_\mu/a_\mu)(\tau_\mu^{\text{ES}} - \tau'), \quad (4)$$

is used, where the unprimed and primed quantities signify new and old values, respectively.⁴¹

It should be noted that a recent timing analysis¹⁴ has shown that, while the NRM is certainly more efficient than the FRM, an optimized version of the DM actually performs best in most situations. For our purposes, however, this fact is not important. The *ideas* underlying the FRM are what we will use in the development of our new simulation approach, and the increased efficiency offered by the NRM will be utilized in its implementation.

B. τ -leaping

Despite the improved efficiency offered by the NRM¹³ and the optimized version of the DM,¹⁴ the SSA remains limited as to the system size amenable to treatment due to its “one reaction at a time” nature. As a result, Gillespie has recently attempted to move beyond the “exact” approach by introducing approximations regarding the number of times a reaction can be expected to fire within a given time interval. His approach, known as τ -leaping, begins by defining the quantity $K_\mu(\tau)$ as the number of times reaction channel R_μ fires during the time interval $[t, t + \tau)$.^{16,37} In general, $K_\mu(\tau)$ is a complex random variable dependant upon the propensity a_μ and the manner in which it changes during $[t, t + \tau)$. Obtaining a rigorous expression for the probability function governing $K_\mu(\tau)$ is thus tantamount to solving the usually intractable chemical Master Equation.

Gillespie recognized, however, that if a time period exists over which the propensity a_μ remains *essentially* constant then one can approximate $K_\mu(\tau)$ as a *Poisson random variable*,^{16,37}

$$K_\mu(\tau) \approx \mathcal{P}_\mu(a_\mu, \tau). \quad (5)$$

There always exists a value of τ over which this assumption holds; in the extreme case it is the interval between successive reactions. In many cases, however, the interval is likely to span numerous reaction events, especially when the reactant populations are “large.”¹⁶ Gillespie then noted that if the mean value of $\mathcal{P}_\mu(a_\mu, \tau)$, i.e., $a_\mu\tau$, is “large,” then one can approximate the Poisson random variable as a *normal* random variable,^{16,37}

$$\begin{aligned} K_\mu(\tau) &\approx \mathcal{N}_\mu(a_\mu\tau, a_\mu\tau) \\ &= a_\mu\tau + \sqrt{a_\mu\tau} \times \mathcal{N}(0, 1), \end{aligned} \quad (6)$$

where the second equality follows from the linear combination theorem for normal random variables.^{16,37} Equation (6) is essentially a chemical Langevin equation and amounts to a “continuous-stochastic” representation of the reaction dynamics. Finally, Gillespie showed that if the ratio of the “deterministic” term in (6), $a_\mu\tau$, to the “noise” term, $\sqrt{a_\mu\tau}$, is “large,” then the noise term can be neglected, leaving^{16,37}

$$K_\mu(\tau) \approx a_\mu\tau, \quad (7)$$

or a “continuous-deterministic” representation.

The expressions in Eqs. (5)–(7) represent a theoretical “bridge” connecting the discrete-stochastic representation of reaction dynamics and the more familiar continuous-deterministic description. With reactions classified at one of these levels of description, these formulae are used in τ -leaping to determine the number of times each reaction fires within a given time interval. One is thus able to “leap” forward in the temporal evolution of a system multiple reaction firings at a time.

Implementing these ideas algorithmically requires a practical method for determining the time interval over which the Poisson approximation (5) can be expected to hold. τ -selection has been the subject of extensive research^{16,17,22} and has undergone numerous refinements recently. The basic idea is to impose a constraint on the magnitude of the change of an individual reaction propensity,

$$|a_\mu(t + \tau_\mu^{\text{Leap}}) - a_\mu(t)| / \xi = \epsilon, \quad (0 < \epsilon \ll 1), \quad (8)$$

where ξ is an appropriate scaling factor. In applying this constraint, one seeks to identify the time interval τ_μ^{Leap} over which the propensity a_μ for reaction R_μ will remain essentially constant within a factor of ϵ assuming that the propensities for all other reactions also remain essentially constant. One then, in essence, calculates a value of τ_μ^{Leap} for each reaction in the system and sets τ equal to the smallest of these.

We will discuss τ -selection further in Sec. III B. For now, however, note that there is an interesting analogy between the procedure used in τ -leaping and that in the FRM (and NRM by extension). In both cases, a time interval is calculated for a specific reaction channel with assumptions made regarding all other reaction channels. The time step is then set equal to the smallest of the set as it is the only one for which the assumptions actually hold. This suggests that the two methods might be seamlessly merged in some way, and forms the basis of our new simulation approach.

With the time step τ calculated, one then proceeds to determine the number of times each reaction fires in τ using Eqs. (5)–(7). Strictly speaking, the τ -leap method uses only Eq. (5). If the propensities of *all* reactions are such that $\{a_\nu\tau\} \gg 1$, however, then Eq. (6) is employed. In Ref. 16 this is termed the *Langevin method*. Furthermore, if all $\{\sqrt{a_\nu\tau}\} \gg 1$ then Eq. (7) is employed, which is equivalent to the explicit Euler method for solving ordinary differential equations.¹⁶ Finally, a *proviso* is added¹⁶ whereby the SSA is used if $\tau \lesssim 1/a_0$, since $1/a_0$ is the expected time to the next reaction firing in the system.¹²

Additional modifications to τ -leaping have been introduced recently. Primarily, strategies have been developed to prevent the possible occurrence of negative species populations. This possibility arises from the fact that Poisson and normal random variables are positively unbounded and, while unlikely, could produce physically unrealizable numbers of reaction firings that result in the

consumption of more reactant entities than are present in the system. Tian and Burrage²³ and Chatterjee *et al.*²⁴ avoid this problem by using binomial random numbers rather than Poisson. Cao *et al.*²⁰ identify “critical” reactions deemed in danger of exhausting their available reactant populations and treat them using a DM SSA approach. The interested reader is referred to Ref. 22 for further details regarding the current state of τ -leaping methodologies.

III. PARTITIONED LEAPING

A. The Framework

The primary change that we make to τ -leaping involves classifying reactions *individually* once the time step τ has been determined. The expressions (5)–(7) are derived for individual reactions and there is nothing precluding their use in this manner. Thus, sets of reactions can essentially be partitioned into “fast,” “medium,” and “slow” classifications based on the quantities $\{a_\nu\tau\}$. We can apply Gillespie’s *proviso*¹⁶ to each individual reaction as well, essentially classifying a reaction as “very slow” if $\tau \lesssim 1/a_\mu$. A tentative next-reaction time for such a reaction can then be generated from Eqs. (3) or (4) and R_μ deemed to fire if $\tau_\mu^{\text{ES}} \leq \tau$.

This procedure amounts to a theoretically justifiable partitioning scheme in which reactions are classified into four different categories based on their propensity values, the calculated time step τ , and the criteria identified by Gillespie^{16,37} for transitioning between the descriptions (5)–(7). The classifications are made as follows:

- If $a_\mu\tau \lesssim 1 \rightarrow$ *Exact Stochastic* (very slow)
- If $a_\mu\tau > 1$ but $\not\gg 1 \rightarrow$ *Poisson* (slow)
- If $a_\mu\tau \gg 1$ but $\sqrt{a_\mu\tau} \not\gg 1 \rightarrow$ *Langevin* (medium)
- If $\sqrt{a_\mu\tau} \gg 1 \rightarrow$ *Deterministic* (fast)

These classifications constitute the basis of the partitioned leaping approach. At each simulation step, a time step τ is calculated and each reaction classified in the manner outlined above. The numbers of reaction firings are then determined, based on these classifications, using Eqs. (3)–(7) and the system evolved accordingly.

Various technical issues must be considered, however, in order to correctly implement this approach. The first involves the inclusion of the “Exact Stochastic” (ES) classification and the random nature of τ_μ^{ES} . Specifically, if $\tau_\mu^{\text{ES}} < \tau$, and the clock is subsequently advanced by τ , then the possibility of R_μ firing again within the interval $(\tau - \tau_\mu^{\text{ES}})$ is precluded. To overcome this complication we employ an iterative procedure for determining τ and classifying reactions. Once the reaction classifications are made τ_μ^{ES} values are calculated for all ES reactions. Some of these may be smaller than τ , and τ is thus adjusted

to the smallest of these. Decreasing τ will result in decreased values of $\{a_\nu\tau\}$ and each reaction must thus be reclassified. Some reclassified reactions may become ES which previously were not and τ_μ^{ES} values must thus be calculated for all of these reactions. Again, these may be smaller than τ and the procedure is thus repeated until no further adjustments are necessary.⁴²

Situations may also arise in which all values of τ_μ^{ES} are larger than τ . In this case, we leave τ unchanged *unless* all reactions in the system are classified as ES. This is because when all reactions are ES it is always safe to allow τ to increase to the time at which the next reaction will fire in the system. When more coarsely classified reactions are present, however, this is not always so. In situations where the coarse reactions are independent of the ES reactions it *is* safe to increase τ . Automatically and efficiently differentiating between this situation and that in which it is not acceptable to increase τ is not trivial, however, especially when considering large, complex reaction networks.

Another issue to consider concerns the proper use of Eq. (4) in our algorithm. The expression in (4) is a transformation formula in which a new value of τ_μ^{ES} is calculated using the new value of a_μ , the old value of a_μ , the old value of τ_μ^{ES} , and the old time step. Under normal circumstances, the “old” values are those from the previous simulation step. As discussed in note 14 of Ref. 13, however, this is not always the case. Specifically, when a reaction becomes inactive, $a_\mu = 0$ and $\tau_\mu^{\text{ES}} = \infty$. Upon reactivation, these values clearly cannot be used in Eq. (4). Thus, the values of a_μ , τ_μ^{ES} and τ are used from the last simulation step at which R_μ was active. Values of $\{a'_\nu(\tau'_\nu{}^{\text{ES}} - \tau')\}$ are stored for all reactions that did not fire at the current step and used in Eq. (4) at the next step at which the corresponding reactions are active. Usually this is the subsequent step, but sometimes it is not.

In our approach, we must account for the fact that reactions classified as ES can change their classification in the subsequent step and then return to an ES classification at a later step. Although the reaction does not become inactive in this case, the same procedure of “carrying over” the values of a_μ , τ_μ^{ES} and τ is used. We simply extend the procedure described above, therefore, and use the stored values of $\{a'_\nu(\tau'_\nu{}^{\text{ES}} - \tau')\}$ at the next step at which the corresponding reactions are active *and* are classified as ES.

We also consider the problem of negative populations. We avoid this problem by simply tracking the state of the system and reversing the population updates if any populations are found to be negative after all reaction firings have been accounted for. The value of τ is then reduced (which is always acceptable) and all reactions reclassified. We reduce τ by 50%, but any amount should suffice. A reduction in τ will result in smaller numbers of reaction firings and, hopefully, alleviation of all negative populations. If not, then the procedure is repeated until no further adjustments are necessary.⁴²

This approach is equivalent to the simple “try again” procedure discussed and implemented in Ref. 20 as a second layer of protection against the occurrence of negative populations. This is used in conjunction with the primary strategy of identifying critical reactions. Note, however, that the critical reaction strategy essentially amounts to a partitioning of reactions into ES and “Poisson” classifications with the intent of avoiding negative populations by maintaining a fine-level description of reactions with small reactant populations. Our method already accomplishes this via inclusion of the ES classification and thus avoids the need to introduce additional calculations or parameters (see Sec. V for a further discussion).

Finally, in determining the set of reaction firings, use of Eqs. (6) and (7) for “Langevin” and “Deterministic” reactions, respectively, will result in values that are real numbers rather than integers. Since it is difficult to determine at what point a continuous population description is acceptable in lieu of an integer description, we choose to round these values before updating the species populations. In Ref. 16 it was argued that use of Eq. (6) as opposed to (5) is an improvement computationally since generation of normal random numbers is faster than Poisson random numbers. Some of this improvement is clearly negated, therefore, by including a subsequent rounding operation, although we have yet to quantitatively investigate the extent of this effect. While the same argument holds for “Deterministic” reactions, the elimination of the random number generation operation should more than compensate for the added rounding procedure.

With these issues accounted for, the partitioned leapfrog algorithm (PLA) is presented as follows:

1. Initialize (species populations, rate constants, define $\approx 1, \gg 1, \epsilon$, etc.).
2. Determine the *initial* value of τ (see Sec. III B).
3. Classify all reactions (not already classified as ES) using the criteria presented above.
4. For all (newly classified) ES reactions, calculate tentative next-reaction times, τ_μ^{ES} , using Eqs. (3) and (4).
- 5a. If $\text{Min}\{\tau_\nu^{\text{ES}}\} \neq \tau$ and all reactions are ES, set $\tau = \text{Min}\{\tau_\nu^{\text{ES}}\}$.
- 5b. If $\text{Min}\{\tau_\nu^{\text{ES}}\} < \tau$, set $\tau = \text{Min}\{\tau_\nu^{\text{ES}}\}$ and return to step 3.
6. Determine the set of reaction firings $\{k_\nu(\tau)\}$ using the appropriate formulas⁴³ and update the species populations.
7. If any $X_i(t+\tau) < 0$, reverse all population updates, set $\tau = \tau/2$ and return to step 3. If not, advance the clock by τ and return to step 2 if stopping criterion not met.

It is important to recognize the minimal user intervention required to implement this approach. Once the reaction network is defined and the associated rate parameters set, one need only define three model-*independent* parameters quantifying the concepts ≈ 1 (for ES classification), $\gg 1$ (for coarse classifications) and $\ll 1$ (defining ϵ). The algorithm then automatically and dynamically partitions the reactions into various subsets, correctly accounting for stochastic noise and “leaping” over unimportant reaction events.

B. Determining the Initial Time Step

In practice, there are two alternate approaches for carrying out Step 2 of the PLA. The first is a reaction-based approach in which values of τ_μ^{Leap} are calculated for each reaction in the system using the constraint expression in (8) directly.^{16,17,22} τ is then set equal to $\text{Min}\{\tau_\nu^{\text{Leap}}\}$. Older versions did this with $\xi \equiv a_0$ in (8),^{16,17} but the current approach uses $\xi \equiv a_\mu$.²² In this approach, τ_μ^{Leap} is given by²²

$$\begin{aligned} \tau_\mu^{\text{Leap}} &= \text{Min} \left\{ \frac{\epsilon_\mu}{|m_\mu(t)|}, \frac{\epsilon_\mu^2}{\sigma_\mu^2(t)} \right\}, & (9) \\ \epsilon_\mu &\equiv \text{Max}\{\epsilon a_\mu, \beta_\mu\} \\ m_\mu(t) &\equiv \sum_{\nu=1}^M f_{\mu\nu}(t) a_\nu(t), \\ \sigma_\mu^2(t) &\equiv \sum_{\nu=1}^M f_{\mu\nu}^2(t) a_\nu(t), \\ f_{\mu\nu}(t) &\equiv \sum_{j=1}^N z_{\nu j} \frac{\partial a_\mu(t)}{\partial X_j}. \end{aligned}$$

Here, β_μ represents the *minimum* possible change in the propensity a_μ . This quantity is used to overcome problems associated with small reactant populations. In Ref. 22, $\beta_\mu \equiv c_\mu$, the stochastic rate constant for R_μ . For first-order reactions this definition is exact. For higher-order reactions, however, this is a lower bound: If the propensity changes at all it will do so by an amount $\geq c_\mu$.²² Thus, this approach is simple but may also be restrictive. We propose an alternative: An *approximate* but less restrictive value of β_μ is the smallest *non-zero* value of $\{\partial a_\mu / \partial X_j\}$, where j indexes all reactant species involved in R_μ . If all values of $\{\partial a_\mu / \partial X_j\}$ are zero, however, then $\beta_\mu = c_\mu$.

The second approach for determining τ constrains the relative changes in the *species populations* in such a way that (8) is satisfied for all reactions.²² τ is then set equal to $\text{Min}\{T_j^{\text{Leap}}\}$, where²²

$$\begin{aligned} T_i^{\text{Leap}} &= \text{Min} \left\{ \frac{e_i}{|\hat{m}_i(t)|}, \frac{e_i^2}{\hat{\sigma}_i^2(t)} \right\}, & (10) \\ e_i &\equiv \text{Max}\{\epsilon X_i / g_i, 1\} \end{aligned}$$

$$\begin{aligned} \hat{m}_i(t) &\equiv \sum_{\nu=1}^M z_{\nu i} a_\nu(t), \\ \hat{\sigma}_i^2(t) &\equiv \sum_{\nu=1}^M z_{\nu i}^2 a_\nu(t). \end{aligned}$$

The parameter g_i depends on the highest-order reaction that species S_i is involved in and can be determined by simple inspection during initialization.²² The expressions in (10) require fewer computational operations than those in (9) and each T_i^{Leap} calculation should thus be significantly faster than each τ_μ^{Leap} .

In Ref. 22, expressions for g_i are presented for all reaction types up to third order. For reactions involving more than one entity of a certain species, such as $2S_i \rightarrow \text{products}$, g_i is a function of X_i . These expressions have clear upper and lower bounds, however. Thus, in order to simplify the calculations we use the more restrictive, but computationally simpler, upper bounds. We determine g_i as follows:

If S_i is a reactant in ...	then $g_i = \dots$
$3S_i \rightarrow \text{products}$,	11/2, <i>else ...</i>
$2S_i + S_j \rightarrow \text{products}$,	9/2, <i>else ...</i>
$S_i + 2S_j \rightarrow \text{products}$,	3, <i>else ...</i>
$S_i + S_j + S_k \rightarrow \text{products}$,	
$2S_i \rightarrow \text{products}$,	2, <i>else ...</i>
$S_i + S_j \rightarrow \text{products}$,	
$S_i \rightarrow \text{products}$,	1, <i>else ...</i>
zero reactions,	0

Note that the last item will result in $T_i^{\text{Leap}} = \infty$, which will never be $\text{Min}\{T_j^{\text{Leap}}\}$.

IV. EXAMPLES


With the presentation of our simulation approach now complete, we will demonstrate in this section the utility of the method, in terms of efficiency and accuracy, via two illustrative examples. We will begin by considering a simple model of clustering that illustrates the algorithm’s ability to treat systems in which species populations vary over many orders of magnitude. A biologically inspired model system will then be considered that illustrates how the stochastic process of gene expression can be accurately and efficiently simulated in conjunction with reactions involving large reactant populations (e.g., metabolic processes).

A. Simple clustering

Clustering processes are inherently multiscale since large numbers of small clusters generally coexist within a system with small numbers of large clusters. As such,

clustering provides an ideal way to demonstrate the ability of the PLA to treat systems in which species populations vary over many orders of magnitude.

We have considered a simple clustering model comprised of the following nine reactions:



For the sake of simplicity we have neglected dissociation reactions and assume that monomers (S_1) are the only mobile species in the system. Furthermore, in order to confine the multiscale effects to variations in the species populations alone, we have chosen deterministic rate constants such that their stochastic counterparts are equivalent for all reactions. For R_1 we choose $3.0 \times 10^6 \text{ M}^{-1} \text{ s}^{-1}$, and for all other reactions $6.0 \times 10^6 \text{ M}^{-1} \text{ s}^{-1}$. We set the initial monomer concentration to $1.66 \times 10^{-6} \text{ M}$ and consider various system volumes Ω ranging from 10^{-15} to 10^{-9} L . This corresponds to initial monomer populations $X_1(0) = 10^3$ to 10^9 and stochastic rate constants $\{c_\nu\} = 10^{-2}$ to 10^{-8} s^{-1} . All simulations were run until consumption of all monomers was complete.

In Fig. 1, we compare average numbers of simulation steps and CPU times required for PLA and SSA simulations of (11) at all system sizes considered. In general, the CPU times follow the same trends as the simulation steps and the reaction-based calculations (PLA-RB) require more steps, and more CPU time, than the species-based (PLA-SB). We also see that for the smallest system size the SSA and PLA give identical results, meaning that all reactions were classified as ES at all steps of the PLA simulations. As the system size increases, however, increased amounts of leaping are observed. The effect is modest for system sizes of 10^{-14} and 10^{-13} L , but increases dramatically beyond that. Moreover, the number of steps actually *decreases* for PLA simulations of systems larger than $\sim 10^{-11} \text{ L}$. The effects of leaping thus overtake the system size effects for these large systems.

In Fig. 2, we show the classifications achieved for each reaction in (11) at each simulation step of a single PLA-RB-3% simulation at a system size of 10^{-9} L . These results show leaping in action and illustrate the inherent multiscale nature of the reaction network. Reactions involving the smallest clusters (i.e., R_1 – R_4) experience extensive amounts of leaping for much of the simulation, with classifications varying rapidly between ES, ‘‘Poisson,’’ ‘‘Langevin’’ and, at times, ‘‘Deterministic.’’ Conversely, reactions involving larger clusters experience much less leaping. Reactions R_7 – R_9 , for example, experience only small amounts of leaping up into the ‘‘Langevin’’ regime.

As for accuracy, smoothed frequency histograms (see Appendix) were generated from 10 000 simulation runs

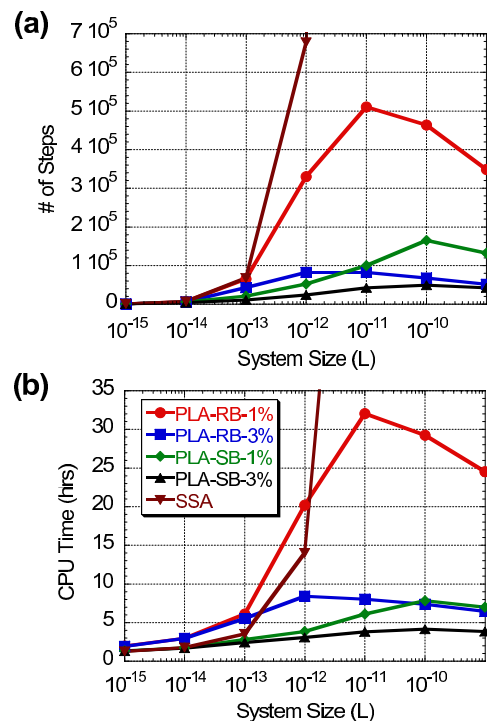


FIG. 1: Average numbers of steps (a) and CPU times (b) required for 10 000 PLA and SSA simulation runs of the simple clustering model (11) [PLA-RB-1% means reaction-based τ -selection with $\epsilon = 0.01$, etc.]. Reaction classifications were made in the PLA runs using $\approx 1 = 3$ and $\gg 1 = 100$. All simulations were performed on a 1.80 GHz Athlon processor.

of the SSA and the PLA for various cluster sizes (the populations of which vary by as many as four orders of magnitude) at various system sizes. Modified versions of the histogram distance D and the self distance $D_{\text{SSA}}^{\text{self}}$, introduced by Cao and Petzold,⁴⁴ were then used to make quantitative comparisons. In all cases, the calculated histogram distances were smaller than the SSA self distances, indicating excellent accuracy (see the Appendix for further details).

B. Stochastic gene expression

The origins and consequences of stochasticity in biological systems has been a subject of intense interest recently.^{1,2,3,4,5,6,7,8} In cellular systems, the primary source of ‘‘intrinsic’’ stochastic noise is gene expression, where the small numbers of regulatory molecules involved in the process result in proteins being produced in ‘‘bursts’’ rather than continuously.^{1,3,7} Other cellular processes, such as metabolism, often involve large numbers of molecules and it has been shown that stochastic fluctuations in gene expression can quantitatively affect these dynamics.²⁷ A fully stochastic treatment of biological systems involving both gene expression and metabolism is infeasible,⁴⁵ however, and thus provides





FIG. 2: Classifications vs. simulation step for each reaction of the simple clustering model (11) at $\Omega = 10^{-9}$ L [i.e., $X_1(0) = 10^9$]. Classifications are: (1) Exact Stochastic, (2) Poisson, (3) Langevin, (4) Deterministic. Results are for a single PLA-RB-3% simulation using $\approx 1=3$ and $\gg 1=100$.

motivation for developing multiscale simulation methods capable of treating systems involving both large- and small-number dynamics.

We have applied the PLA to a simple biologically inspired model system that involves both gene expression and protein-protein interactions. The network that we have considered is as follows:



The first two reactions constitute the gene expression part of the network, where the single gene G spontaneously converts into an active conformation G^* that then produces proteins P in bursts of n . The third and fourth reactions constitute the protein-protein enzymatic part of the network where P interacts with Q to form an enzyme-substrate complex $\{P:Q\}$ that then produces R and regenerates Q . The final reaction models the degradation of R .

Rate constants for the five reactions in (12) were set equal to 750 s^{-1} , 750 s^{-1} , $6.0 \times 10^8 \text{ M}^{-1} \text{ s}^{-1}$, 100 s^{-1} , and 50 s^{-1} , respectively. The initial enzyme concentration $[Q]_0$ was set equal to $1.66 \times 10^{-7} \text{ M}$, n to $0.2X_Q(0)$,⁴⁶ and investigations were carried out for various system sizes ranging from 10^{-15} to 10^{-7} L. This corresponds to initial enzyme populations $X_Q(0)$ ranging from 10^2 to 10^{10} . In all cases, the system began with a single entity of G and

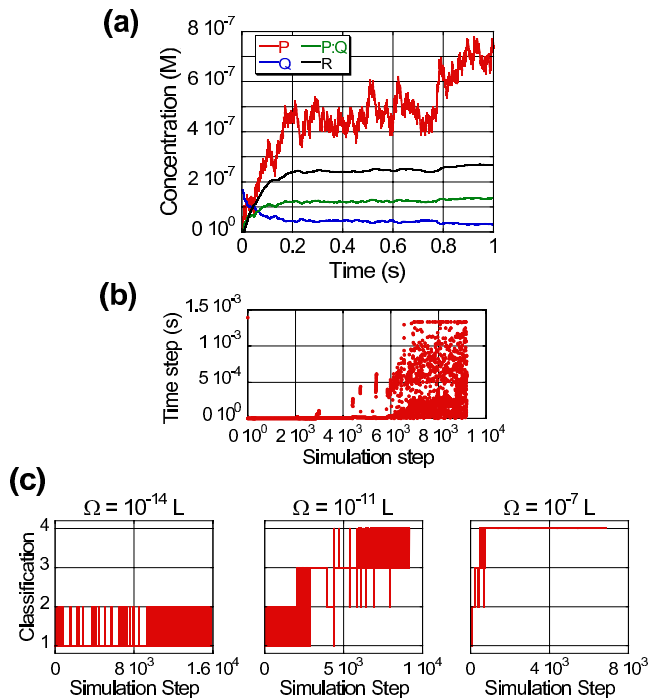


FIG. 3: Results for typical PLA-RB-3% simulations ($\approx 1=3$, $\gg 1=100$) of the simple gene expression model (12): (a) A typical time course at $\Omega = 10^{-11}$ L, (b) value of the time step at each simulation step for the same run as in (a), (c) classifications vs. simulation step for reaction R_3 at various system sizes.

null populations of G^* , $\{P:Q\}$ and R . All simulations were run until $t = 1$ s.

In Fig. 3, we show typical results for PLA-RB-3% simulations of (12). The time course plot in Fig. 3(a), generated for a system size of 10^{-11} L, illustrates the stochastic nature of the gene expression dynamics, apparent in the noisy time evolution of the gene product P . Conversely, the large-number dynamics result in much smoother trajectories for Q , $\{P:Q\}$ and R . The time step plot in Fig. 3(b), taken from the same simulation as in 3(a), illustrates the algorithm’s ability to dynamically adjust as the system evolves. The first 3000 simulation steps correspond to ~ 0.01 s of simulated time, while a time of 0.1 s is reached around step 6400. The entire simulation takes 9192 steps to complete. Thus, simulating the first 1% of the evolved time required $\sim 33\%$ of all simulation steps and the first 10% $\sim 70\%$ of all steps.

In Fig. 3(c), we see how the classifications coarsen with increasing system size. For $\Omega = 10^{-14}$ L the classifications for R_3 never reach beyond “Poisson.” For $\Omega = 10^{-11}$ L the classifications coarsen approximately midway through the simulation, and for $\Omega = 10^{-7}$ L “Deterministic” status is achieved quickly and maintained almost exclusively throughout. Similar results are obtained for reactions R_4 and R_5 (data not shown) while R_1 and R_2 are classified as ES at all steps of the simulation.

In Fig. 4, we compare the average numbers of simula-

tion steps and the CPU times required for PLA and SSA simulations of (12). Again we see that the CPU times generally follow the same trends as the simulation steps but, contrary to what is observed in Sec. IV A, the SB calculations require more steps, and longer CPU times, than the RB. The efficiencies of the two τ -selection procedures thus appear to be system-dependent. This is an interesting result and an area of possible future investigation.

We also see in Fig. 4 that the general trend for the simulation steps is an initial increase with system size, followed by a drop, another increase and then another drop with an eventual leveling off (the only departure from this is for the PLA-RB-3% where the number of steps initially decreases with system size). The initial behavior is the same as that seen for the simple clustering model (see Fig. 1) and can easily be explained in terms of a competition between system size and leaping effects. The second increase is unique to this example, however, and arises from an increasing fraction of simulations requiring large numbers of steps. At 10^{-10} L, for example, 95% of all PLA-RB-1% simulations require between 38 078 and 437 061 steps to complete. At 10^{-12} L, the range is 38 081 and 142 687. The stochastic nature of the gene expression dynamics thus results in a wide variety of possible time-evolution trajectories for systems of size 10^{-11} to 10^{-9} L, some of which require many more simulation steps to complete than others. The ability of the PLA to handle systems in this “medium” size range is a particular strength of the method.

Finally, in terms of accuracy, we generated smoothed frequency histograms for each species in (12) at various system sizes and again observed excellent agreement between the PLA and SSA. In all cases, the calculated histogram distances were smaller than the SSA self distances.

V. DISCUSSION AND CONCLUSIONS

In this article, we have presented a new multiscale simulation approach, termed the “partitioned leaping algorithm,” for modeling stochasticity in chemical reaction networks. Clearly, this method is closely related to τ -leaping.^{16,17,20,22} The key difference is that reactions are classified individually in the PLA, which leads to various advantages over τ -leaping: The PLA overcomes the problem of negative populations more simply, the SSA is incorporated into the algorithmic framework in a more seamless and natural way and reactions are treated at four levels of description rather than two.

The negative populations problem is avoided in the PLA primarily through the structure of the algorithm. This can be seen most easily by considering the RB τ -selection procedure (9), though the arguments hold for the SB approach as well. If a reaction has a small reactant population, such that the reaction would be deemed “critical” in τ -leaping,^{20,22} then its τ_{μ}^{Leap} value will be

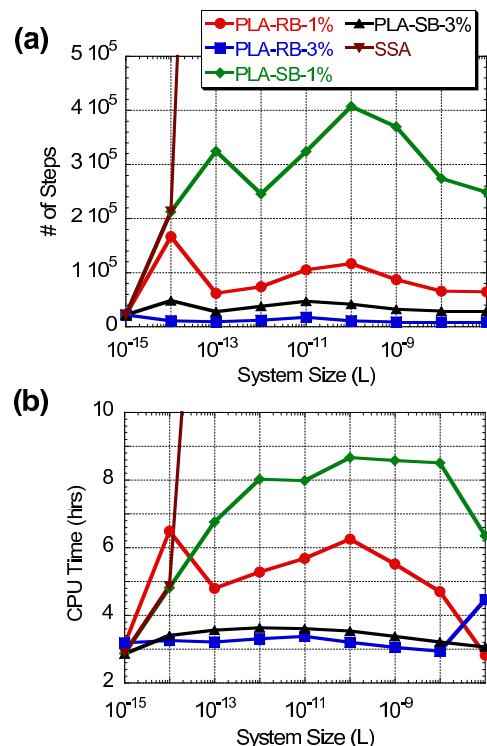


FIG. 4: Average numbers of steps (a) and CPU times (b) required for 10 000 PLA and SSA simulation runs of the simple gene expression model (12). Reaction classifications were made in the PLA runs using $\approx 1 = 3$ and $\gg 1 = 100$. All simulations were performed on a 1.80 GHz Athlon processor.

$\leq 1/a_{\mu}$, the expected time of its next firing. This is because, for very small reactant populations, a single firing will change a_{μ} by a fractional amount $\geq \epsilon$, which will be automatically recognized in the τ -selection process. As a result, the reaction will be classified as ES (since τ is always $\leq \tau_{\mu}^{\text{Leap}}$) and thus prevented from firing more than once, assuring that its reactant populations do not fall below zero. Since this outcome arises as a natural consequence of the design of the PLA, there is no need for additional calculations or parameters, such as those associated with the critical reaction search in τ -leaping.^{20,22}

Another difference between the methods is that in τ -leaping the τ -selection formulas (9) and (10) are applied only to non-critical reactions.^{20,22} In the PLA, these formulas are applied to all reactions. This is merely a superficial difference, however. Although critical reactions are treated separately from non-critical reactions in τ -leaping, they still contribute to τ -selection, just in a different way. In τ -leaping, “candidate” time steps are calculated for the critical and non-critical reaction subsets and τ set equal to the smaller of the two.^{20,22} The critical reaction time step is calculated using Eq. (1) with a_0 replaced by a_0^c , the sum of the critical reaction propensities. Obviously, this will give a time interval on the order of $1/\text{Max}\{a_{\nu}\}$. In the PLA, these same reactions will be treated using Eqs. (9) or (10). As just discussed,

in the PLA-RB scenario the result will be τ_{μ}^{Leap} values on the order of $1/a_{\mu}$. τ is set equal to $\text{Min}\{\tau_{\nu}^{\text{Leap}}\}$, so the overall result is essentially the same: the “critical” reactions contribute to the τ -selection process a time interval on the order of $1/\text{Max}\{a_{\nu}\}$. The randomness associated with these reactions is then added in in Step 4 of the PLA using the methods of the NRM¹³ [Eqs. (3) and (4)], which is the obvious choice given the structure of the algorithm.

The ability of the PLA to treat reactions at the continuous-stochastic and deterministic scales, rather than just the discrete-stochastic, is an important attribute of the method as well; the PLA could be seen as unifying, into one overarching approach, three different types of simulation methods: exact-stochastic,^{12,13} τ -leaping,^{16,22} and methodology coupling, or hybrid, techniques.^{25,26,27,28,29,30} Hybrid methods operate by partitioning reactions into “fast” and “slow” subsets and using, e.g., deterministic reaction rate equations or stochastic differential equations, to describe the fast reactions and the SSA (or modified versions thereof) for the slow. The primary shortcoming of these approaches is the lack of a sound theoretical basis for the partitioning scheme,¹⁹ which calls into question the extent of their utility. The partitioning procedure described in Sec. III A overcomes this problem; partitioning is accomplished via Gillespie’s rigorously derived criteria for transitioning between the descriptions (5)–(7).^{16,37} Of course, as currently formulated, “Langevin” and “Deterministic” reactions are treated in a manner equivalent to the explicit Euler method for solving stochastic and ordinary differential equations, respectively, which may not be satisfactory in all cases.

Recognizing this, Petzold and co-workers have developed various “implicit” τ -leaping methods^{18,19,21} that take into account the values of the propensities at both the beginning *and* end of the time step τ . By doing so, these methods are able to take larger time steps than explicit methods^{16,17} that only consider the propensity at the beginning of the step (the algorithm presented here is explicit). This comes at a cost, however: Implicit methods dampen fluctuations in the species populations and, hence, underestimate the amplitude of the noise.^{18,19,21} Nevertheless, the ability of these methods to maintain stability in situations where explicit methods fail has been demonstrated. Incorporating these ideas into the PLA is an interesting possibility, and an area of possible future investigation.

From a *practical* point of view, it is unclear to what extent the efficiency of the PLA is actually improved by including the “Langevin” and “Deterministic” classifications. As mentioned above, the reason for including these is to presumably improve computational efficiency via faster generation of normal random numbers¹⁶ or by eliminating random number generation altogether. Their inclusion adds complexity to the method, however. To what extent this attenuates the gains in efficiency remains to be seen, and is something we plan on investi-

gating further in the future. Note, however, that a useful feature of the PLA is its ability, via manipulation of the classification parameters, to force a simulation at a particular level of description (e.g., if $\gg 1$, $\approx 1 \equiv \infty$, then all reactions will be classified as ES at all steps of a simulation – See Sec. III A). With the coarse classifications included, one could use this feature (with rounding turned off) to perform, e.g., deterministic simulations of a system for direct comparisons to results from the PLA.

Finally, we must also acknowledge the primary shortcoming of the PLA, namely, in handling systems with widely disparate rate constants. Consider a system that contains a single entity, such as a gene, experiencing rapid ON/OFF or binding/unbinding behavior. Because there exists only one entity these reactions will be flagged as ES, and because of the large rate constants their τ_{μ}^{Leap} values (in the PLA-RB) will be small. The result will be small time steps and a “bogging down” of the PLA. Note that this is a problem in τ -leaping as well.

The only way to overcome this problem is to “reduce” the model so that the effects of the fast reactions are accounted for in some approximate way. This is a common approach; examples include quasi-steady-state or Michaelis-Menten reductions. Recently, more advanced model reduction schemes have been proposed that can be applied to generalized reaction networks^{31,32,33,34} and, in some cases, dynamically.^{35,36} Since these methods are specifically designed to overcome problems associated with widely disparate rate constants they could form an effective complement to leaping techniques that tackle the problem of widely disparate species populations. Future techniques combining automatic and dynamic model reduction with partitioned leaping could open the door to computational investigations far beyond the reach of current approaches.

Acknowledgments

F. A. Escobedo, H. Lee, A. A. Quong, A. J. Golumbfskie and C. F. Melius are thanked for useful discussions regarding this work. We acknowledge financial support from the Semiconductor Research Corporation Graduate Fellowship Program.

APPENDIX: HISTOGRAM SMOOTHING, HISTOGRAM DISTANCE AND SELF DISTANCE

For a set of N data points $\{x_1, x_2, \dots, x_N\}$, the total number falling within a discrete interval $[x, x + \Delta)$ can be formally written as $\sum_{i=1}^N \int_x^{x+\Delta} \delta(x_i - x') dx'$, where $\delta(x_i - x')$ is the Dirac delta function and the integral equals unity if x_i lies within $[x, x + \Delta)$ and zero otherwise. A “histogram density” can then be obtained by dividing




FIG. 5: Smoothed population histograms for species S_2 at $t=10$ obtained from 10 000 simulation runs of the “decaying-dimerizing” reaction set¹⁷ using the SSA and the PLA-RB with various values of ϵ . Reaction classifications were made in the PLA runs using $\approx 1=3$ and $\gg 1=100$. All histograms were generated using Eq. (A.3) with $\sigma = 15$.

this quantity by $N\Delta$ and taking the limit as $\Delta \rightarrow 0$,

$$\hat{h}(x) = \lim_{\Delta \rightarrow 0} \frac{1}{N\Delta} \sum_{i=1}^N \int_x^{x+\Delta} \delta(x_i - x') dx'. \quad (\text{A.1})$$

The “hat” in $\hat{h}(x)$ signifies that this quantity is a statistical *estimator* of the true histogram density

$$h(x) = \lim_{N \rightarrow \infty} \hat{h}(x). \quad (\text{A.2})$$

A “smoothed” histogram can then be obtained by approximating the delta function in (A.1) by a finite width Gaussian $\kappa \exp\left(\frac{-(x_i - x')^2}{2\sigma^2}\right)$, where κ is a normalization constant and σ^2 is the (user-defined) variance. Substituting into (A.1), noting that to first order $\int_x^{x+\Delta} e^{-u^2} du = \Delta e^{-x^2}$, and normalizing, gives

$$\hat{h}(x) \cong \frac{1}{\sqrt{2\pi}\sigma N} \sum_{i=1}^N \exp\left(\frac{-(x_i - x)^2}{2\sigma^2}\right). \quad (\text{A.3})$$

All smoothed histograms generated in this work were obtained using this expression. Examples are shown in Fig. 5, obtained from 10 000 PLA and SSA simulations of the “decaying-dimerizing” reaction set using the same rate parameters and initial conditions as in Ref. 17.

In order to quantitatively compare results obtained via different simulation methods (i.e., PLA, SSA, etc.) we use the “histogram distance” discussed by Cao and Petzold.⁴⁴ The quantity δh_x is defined as $h_1(x) - h_2(x)$ and the histogram distance D is then

$$D \equiv \frac{1}{2} \int_x |\delta h_x| dx. \quad (\text{A.4})$$

The factor of 1/2 assures that D lies within $[0,1)$, with 0 constituting a perfect fit and 1 a complete mismatch.

It is important to note that D is defined in (A.4) in terms of the *true* histogram densities $h_1(x)$ and $h_2(x)$. In practice, we only have their estimators and can thus

only calculate an estimated value of D . As a result, a certain amount of statistical uncertainty is associated with the comparison of histograms. In order to quantify this uncertainty Cao and Petzold⁴⁴ introduce the “self distance,” D^{self} , which essentially measures the distance between the estimator $\hat{h}(x)$ and the true density $h(x)$. Expressions for the upper bounds on the mean and variance of D^{self} are presented in Ref. 44 in terms of the number of bin intervals K . However, since we are using Eq. (A.3) rather than a counting procedure to generate $\hat{h}(x)$ we must derive alternate expressions.

We begin by defining

$$D^{\text{self}} \equiv \frac{1}{2} \int_x |\delta h_x^{\text{self}}| dx, \quad (\text{A.5})$$

$$\delta h_x^{\text{self}} \equiv \hat{h}(x) - h(x).$$

Following Cao and Petzold,⁴⁴ we then note that the number of points falling within the interval $[x, x + \Delta)$ is a binomial random variable $B(p_x, N)$, where p_x is the success probability. $\hat{h}(x)$ can then be written as

$$\hat{h}(x) = \lim_{\Delta \rightarrow 0} B(p_x, N)/N\Delta = B(dp_x, N)/Ndx. \quad (\text{A.6})$$

Since the mean $E[B(p_x, N)] = Np_x$ and the variance $\text{Var}[B(p_x, N)] = Np_xq_x$ ($q_x \equiv 1 - p_x$), the mean and variance of $\hat{h}(x)$ are dp_x/dx and $dp_x dq_x / Ndx^2$, respectively. With $dp_x = h(x)dx$ and $dq_x = 1 - h(x)dx$, the mean and variance of δh_x^{self} are then

$$E[\delta h_x^{\text{self}}] = 0, \quad (\text{A.7})$$

$$\text{var}[\delta h_x^{\text{self}}] = [h(x)dx][1 - h(x)dx]/Ndx^2 \approx \hat{h}(x)/Ndx,$$

where the last line utilizes the histogram density estimator $\hat{h}(x)$.

For large N , δh_x^{self} can be approximated as a normal random variable with mean zero and variance $\hat{h}(x)/Ndx$. This means that

$$\frac{\delta h_x^{\text{self}}}{\sqrt{\hat{h}(x)/Ndx}} \quad (\text{A.8})$$

is approximately standard normal and

$$\frac{|\delta h_x^{\text{self}}|}{\sqrt{\hat{h}(x)/Ndx}} \quad (\text{A.9})$$

is approximately chi distributed with one degree of freedom. Since the mean and variance of a chi random variable with one degree of freedom are $\sqrt{2/\pi}$ and $(\pi - 2)/\pi$, respectively, this gives

$$E[|\delta h_x^{\text{self}}|] \approx \sqrt{\frac{2}{N\pi} \frac{\hat{h}(x)}{dx}}, \quad (\text{A.10})$$

$$\text{var}[|\delta h_x^{\text{self}}|] \approx \left(\frac{\pi - 2}{N\pi}\right) \frac{\hat{h}(x)}{dx}.$$

Finally, using the identities⁴⁷

$$\begin{aligned} \mathbb{E} \left[\int_x f(x) dx \right] &= \int_x \mathbb{E}[f(x)] dx, \\ \text{Var} \left[\int_x f(x) dx \right] &\leq \left(\int_x \sqrt{\text{Var}[f(x)]} dx \right)^2, \end{aligned}$$

we get

$$\begin{aligned} \mathbb{E}[D^{\text{self}}] &\approx \frac{1}{2} \sqrt{\frac{2}{N\pi}} \int_x \sqrt{\hat{h}(x)} dx, \quad (\text{A.11}) \\ \text{var}[D^{\text{self}}] &\lesssim \frac{1}{4} \left(\frac{\pi - 2}{N\pi} \right) \left(\int_x \sqrt{\hat{h}(x)} dx \right)^2. \end{aligned}$$

In practice, we calculate $\mathbb{E}[D_{\text{ref}}^{\text{self}}]$, the self distance for a *reference* histogram generally obtained using the SSA. This value then tells us that any histogram with a distance $D < \mathbb{E}[D_{\text{ref}}^{\text{self}}]$ *cannot* be statistically differentiated from the reference histogram. In Fig. 5, we see that only the PLA-RB-1% histogram achieves this level of accuracy. The expression for $\text{var}[D^{\text{self}}]$ is included here for completeness.

-
- * Electronic address: lh64@cornell.edu
 † Electronic address: pqc1@cornell.edu
- ¹ H. H. McAdams and A. Arkin, Proc. Natl. Acad. Sci. USA **94**, 814 (1997).
 - ² A. P. Arkin, J. Ross, and H. H. McAdams, Genetics **149**, 1633 (1998).
 - ³ H. H. McAdams and A. Arkin, Trends Genet. **15**, 65 (1999).
 - ⁴ M. B. Elowitz, A. J. Levine, E. D. Siggia, and P. S. Swain, Science **297**, 1183 (2002).
 - ⁵ N. Fedoroff and W. Fontana, Science **297**, 1129 (2002).
 - ⁶ C. V. Rao, D. M. Wolf, and A. P. Arkin, Nature **420**, 231 (2002).
 - ⁷ M. Kærn, T. C. Elston, W. J. Blake, and J. J. Collins, Nature Rev. Genet. **6**, 451 (2005).
 - ⁸ J. M. Raser and E. K. O’Shea, Science **309**, 2010 (2005).
 - ⁹ J. D. Plummer and P. B. Griffin, Nucl. Instr. Meth. Phys. Res. B **102**, 160 (1995).
 - ¹⁰ S. Roy and A. Asenov, Science **309**, 388 (2005).
 - ¹¹ The International Technology Roadmap for Semiconductors, 2001 Ed., <http://public.itrs.net>.
 - ¹² D. T. Gillespie, J. Comput. Phys. **22**, 403 (1976).
 - ¹³ M. A. Gibson and J. Bruck, J. Phys. Chem. A **104**, 1876 (2000).
 - ¹⁴ Y. Cao, H. Li, and L. Petzold, J. Chem. Phys. **121**, 4059 (2004).
 - ¹⁵ H. Resat, H. S. Wiley, and D. A. Dixon, J. Phys. Chem. B **105**, 11026 (2001).
 - ¹⁶ D. T. Gillespie, J. Chem. Phys. **115**, 1716 (2001).
 - ¹⁷ D. T. Gillespie and L. R. Petzold, J. Chem. Phys. **119**, 8229 (2003).
 - ¹⁸ M. Rathinam, L. R. Petzold, Y. Cao, and D. T. Gillespie, J. Chem. Phys. **119**, 12784 (2003).
 - ¹⁹ Y. Cao, L. R. Petzold, M. Rathinam, and D. T. Gillespie, J. Chem. Phys. **121**, 12169 (2004).
 - ²⁰ Y. Cao, D. T. Gillespie, and L. R. Petzold, J. Chem. Phys. **123**, 054104 (2005).
 - ²¹ M. Rathinam, L. R. Petzold, Y. Cao, and D. T. Gillespie, Multiscale Model. Simul. **4**, 867 (2005).
 - ²² Y. Cao, D. T. Gillespie, and L. R. Petzold, J. Chem. Phys. **124**, 044109 (2006).
 - ²³ T. Tian and K. Burrage, J. Chem. Phys. **121**, 10356 (2004).
 - ²⁴ A. Chatterjee, D. G. Vlachos, and M. A. Katsoulakis, J. Chem. Phys. **122**, 024112 (2005).
 - ²⁵ E. L. Haseltine and J. B. Rawlings, J. Chem. Phys. **117**, 6959 (2002).
 - ²⁶ K. Burrage, T. Tian, and P. Burrage, Prog. Biophys. Mol. Biol. **85**, 217 (2004).
 - ²⁷ J. Puchalka and A. M. Kierzek, Biophys. J. **86**, 1357 (2004).
 - ²⁸ K. Vasudeva and U. S. Bhalla, Bioinformatics **20**, 78 (2004).
 - ²⁹ T. R. Kiehl, R. M. Mattheyses, and M. K. Simmons, Bioinformatics **20**, 316 (2004).
 - ³⁰ H. Salis and Y. Kaznessis, J. Chem. Phys. **122**, 054103 (2005).
 - ³¹ C. V. Rao and A. P. Arkin, J. Chem. Phys. **118**, 4999 (2003).
 - ³² Y. Cao, D. T. Gillespie, and L. R. Petzold, J. Chem. Phys. **122**, 014116 (2005).
 - ³³ Y. Cao, D. Gillespie, and L. Petzold, J. Comput. Phys. **206**, 395 (2005).
 - ³⁴ J. Goutsias, J. Chem. Phys. **122**, 184102 (2005).
 - ³⁵ A. Samant and D. G. Vlachos, J. Chem. Phys. **123**, 144114 (2005).
 - ³⁶ H. Salis and Y. N. Kaznessis, J. Chem. Phys. **123**, 214106 (2005).
 - ³⁷ D. T. Gillespie, J. Chem. Phys. **113**, 297 (2000).
 - ³⁸ The notation used in this article differs from that in Refs. 16,17,18,19,20,21,22 and 37. Here, latin character subscripts are used to refer to species and greek characters to reactions.
 - ³⁹ D. T. Gillespie, Physica A **188**, 404 (1992).
 - ⁴⁰ The superscript “ES” is used because differentiating between exact-stochastic time intervals and “leaping” intervals will be important later.
 - ⁴¹ The expression in Eq. (4) differs slightly from that in Ref. 13 as it is a *relative* time version of the transformation formula.
 - ⁴² The number of iterations required in this procedure is definitively finite. In extreme situations, if τ is continually reduced, at some point *all* reactions will become classified as ES. Reclassifications will then no longer be necessary and the iterative loop will terminate.
 - ⁴³ Standard techniques exist for generating Poisson and normal random deviates.⁴⁸ For ES reactions, if $\tau_{\mu}^{\text{ES}} = \tau$ then $k_{\mu}(\tau) = 1$, otherwise zero.
 - ⁴⁴ Y. Cao and L. Petzold, J. Comput. Phys. **212**, 6 (2006).
 - ⁴⁵ D. Endy and R. Brent, Nature **409**, 391 (2001).

⁴⁶ By allowing the number of proteins produced per expression event to change we are effectively varying the degree of “translational efficiency”⁷.

⁴⁷ D. T. Gillespie, *Markov Processes: An Introduction for Physical Scientists* (Academic, San Diego, 1992).

⁴⁸ See, e.g., W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C, The Art of Scientific Computing, 2nd Ed.* (Cambridge University Press, New York, NY, 1999).