Stiffness detection and reduction in discrete stochastic simulation of biochemical systems

Yang Pu,^{1,a)} Layne T. Watson,^{1,2,b)} and Yang Cao^{1,c)}

¹*Department of Computer Science, Virginia Polytechnic Institute and State University, Blacksburg, Virginia 24061, USA.*

²*Department of Mathematics, Virginia Polytechnic Institute and State University, Blacksburg, Virginia 24061, USA.*

(Received 7 October 2010; accepted 6 January 2011; published online 1 February 2011)

Typical multiscale biochemical models contain fast-scale and slow-scale reactions, where “fast” reactions fire much more frequently than “slow” ones. This feature often causes stiffness in discrete stochastic simulation methods such as Gillespie’s algorithm and the Tau-Leaping method leading to inefficient simulation. This paper proposes a new strategy to automatically detect stiffness and identify species that cause stiffness for the Tau-Leaping method, as well as two stiffness reduction methods. Numerical results on a stiff decaying dimerization model and a heat shock protein regulation model demonstrate the efficiency and accuracy of the proposed methods for multiscale biochemical systems. © 2011 American Institute of Physics. [doi:10.1063/1.3548838]

I. INTRODUCTION

Over the last couple of decades, remarkable advancements in the field of computational biology have led to better understanding of biological systems at the cellular level and the development of predictive biological models of cellular systems. As biological models become more and more complicated, practical, accurate, and efficient numerical simulation algorithms are in heavy demand, especially for biological systems with a multiscale nature. Multiscale problems arise in many scientific applications and demand special attention for their analysis and computation. This paper focuses on stiffness, one of the most important multiscale problems in discrete stochastic simulation of biochemical systems.

Most available biological models are continuous and deterministic and are formulated as ordinary differential equations (ODEs). However, chemical species with small populations in living cells often result in system behavior that is discrete and stochastic rather than continuous and deterministic.^{1–4} For those systems, it is necessary to consider stochastic modeling and simulation techniques. There are two important groups of discrete stochastic simulation methods for biochemical systems. One includes Gillespie’s stochastic simulation algorithm (SSA),⁵ which follows the probability distribution defined by the chemical master equation (CME). The other includes Tau-Leaping methods,^{5,6,8} a group of approximation schemes intended to gain efficiency while sacrificing accuracy. For both groups of stochastic simulation methods, stiffness presents a major challenge. In classical numerical solution of ODEs,⁷ stiffness refers to the situation where a numerical method’s stepsizes are excessively small, not because of accuracy requirements but due to numerical stability requirements. As a result, the numerical solution

takes an excessively long time. Similar behavior has been observed and studied in the discrete stochastic simulation scenario.^{8–11} Due to stiffness, in SSA most central processing unit (CPU) time is spent on fast reaction events even though the system dynamics are driven mostly by slow reaction events. In Tau-Leaping methods, stiffness causes very small stepsizes. Both cases lead to low computational efficiency. Some efforts^{8–11} have been made toward more efficient simulation, when stiffness is present in the system. However, they are either focused only on SSA, or when applied to Tau-Leaping methods, are based on a partial equilibrium (PE) assumption and cannot work under the quasisteady state (QSS) assumption.

The PE and QSS assumptions are both important multiscale features in biochemical systems. PE refers to the situation, where some reactions fire much faster than others and the corresponding subsystem reaches a partial equilibrium state.^{8,10} QSS refers to the situation, where some state variables fluctuate very quickly around their QSS.⁹ Figure 1 shows two typical cases observed in biological models. Suppose species Y has a relatively small population compared to those of species X_1 , X_2 , and X_3 . When reactions R_1 and R_3 are much faster than the remaining reactions, X_1 and Y form a closed subsystem that reaches partial equilibrium. When reactions R_2 and R_4 are much faster than the other reactions, the mean value of Y quickly reaches a state such that the influx and outflux of Y are balanced. Y is then said to be in QSS.

For biochemical systems, stiffness is almost always related to QSS species. This paper proposes a method to automatically identify the species that are in QSS and lead to stiffness. This goal is achieved by analyzing the tau values required by the accuracy of all the species in the hybrid SSA/Tau-Leaping simulation. The tau values reveal the species that require much smaller stepsizes than others, and then the system structure related to these species is analyzed. If the corresponding species have reached QSS, identify them as stiff variables and apply stiffness reduction techniques

^{a)}Electronic mail: ypu@cs.vt.edu.

^{b)}Electronic mail: ltw@cs.vt.edu.

^{c)}Electronic mail: ycao@cs.vt.edu.

FIG. 1. PE and QSS: Rectangular nodes represent reactions and oval nodes represent species.

through the QSS assumption. Two stiffness reduction methods are proposed here. One method applies model reduction but works only when the reaction subnetwork connected to the stiff variable is very simple, such as those shown in Fig. 1. The other method, QSS Tau-Leaping, modifies the Tau-Leaping method based on the QSS assumption. The latter method is less efficient than the model reduction method, but applies to more general problems. Both methods can reduce the stiffness corresponding to the QSS variable. With the stiffness reduction methods, stepsizes can be much larger with negligible loss of accuracy for nonstiff variables. A stiff decaying dimerization model⁵ and a heat shock protein (HSP) regulation model¹² are used to demonstrate the detection and reduction methods and the significantly improved simulation efficiency.

II. BACKGROUND

A. SSA and Tau-Leaping method

Suppose there is a biochemical system that involves N molecule species $\{S_1, \dots, S_N\}$. The state vector is denoted by $X(t) = [X_1(t), \dots, X_N(t)]$, where $X_i(t)$ is the number of molecules of species S_i at time t . There are totally M reaction channels $\{R_1, R_2, \dots, R_M\}$. For $X(t) = x$, the dynamics of reaction channel R_j is characterized by its propensity function $a_j(x)$ and state change vector $v_j = (v_{1j}, \dots, v_{Nj})$, where $a_j(x)dt$ gives the probability that the reaction R_j will occur in the next infinitesimal time interval $[t, t + dt]$, and v_{ij} gives the change in the S_i molecule population induced by the reaction R_j .

Gillespie's SSA (Ref. 5) simulates every reaction firing event in the system. Let $a_0(x) = \sum_{j=1}^M a_j(x)$. In each step, SSA generates two uniform $(0, 1)$ random numbers r_1 and r_2 . The time for the next reaction to occur is given by $t + \tau$, where $\tau = -\log(r_1)/a_0(x)$. The index j for the next reaction is given by the smallest integer satisfying $\sum_{j'=1}^j a_{j'}(x) > r_2 a_0(x)$. After τ and j are calculated, the system state $X(t)$ is updated by $X(t + \tau) = x + v_j$, and the time t is updated to $t + \tau$. The simulation proceeds until t reaches the final time.

Although SSA is an accurate simulation method, it is not efficient for many biochemical systems because it follows ev-

ery reaction event. The Tau-Leaping method⁵ tries to accelerate the simulation by leaping over many reactions in one step. This is achieved by answering the following question: How many times will each reaction channel fire in a time interval τ ? Define

$$K_j(\tau; x, t) \equiv \text{the number of times, given } X(t) = x, \text{ that reaction channel } R_j \text{ will fire in the time interval } [t, t + \tau) \quad (j = 1, \dots, M). \quad (1)$$

The Tau-Leaping method assumes the leap condition: for the current state x , τ is small enough that the change in the state during $[t, t + \tau)$ will be so small that no propensity function will suffer an appreciable change in its value. $K_j(\tau; x, t)$ is then well approximated by Poisson random variable with mean and variance $a_j(x)\tau$,

$$K_j(\tau; x, t) = P_j(a_j(x)\tau) \quad j = 1, \dots, M. \quad (2)$$

So if $X(t) = x$ and τ satisfies the leap condition, the state update at time $t + \tau$ is

$$X(t + \tau) = x + \sum_{j=1}^M v_j P_j(a_j(x)\tau). \quad (3)$$

B. Hybrid SSA/Tau-Leaping

In order to avoid negative populations that could appear in the Tau-Leaping method for species with low populations, the hybrid SSA/Tau-Leaping algorithm⁸ partitions the whole system into critical (simulated by SSA) and noncritical (simulated by Tau-Leaping) reactions. A positive control parameter n_c is used. Any reaction channel with a positive propensity function that is currently within n_c firings of exhausting one of its reactants is classified as a critical reaction. n_c is usually set around 10. The hybrid algorithm chooses τ in such a way that not more than one firing of all the critical reactions can occur during the leap. It uses SSA to simulate the critical reactions, and the Poisson Tau-Leaping method to simulate the remaining noncritical reactions.

Achieving high efficiency while maintaining a prescribed accuracy requires a procedure to calculate the largest value of τ that is compatible with the leap condition. Gillespie originally proposed that the leap condition can be satisfied if the expected change in each propensity function $a_j(x)$ during the leap were bounded by $\epsilon a_0(x)$, where ϵ is an error control parameter ($0 < \epsilon \ll 1$). Later, Gillespie and Petzold⁶ showed that the largest value of τ that satisfies the leap condition should be estimated by requiring that the mean and standard deviation of the expected change in $a_j(x)$ in the time period τ be bounded by $\epsilon a_0(x)$ for all j . Cao *et al.*⁸ improved the τ -selection formula for the hybrid SSA/Tau-Leaping algorithm by taking

$$\tau = \min_{i \in I_{\text{RS}}} \left\{ \frac{\max\{\epsilon x_i / g_i, 1\}}{|\mu_i(x)|}, \frac{\max\{\epsilon x_i / g_i, 1\}^2}{\sigma_i^2(x)} \right\}, \quad (4)$$

where I_{RS} is the set of indices of all reactant species, g_i is given by a formula that guarantees that bounding the relative

change of states is sufficient for bounding the relative change of propensity functions, and

$$\mu_i(x) \equiv \sum_{j \in J_{\text{ncr}}} v_{ij} a_j(x), \quad \forall i \in I_{\text{rs}}, \quad (5a)$$

$$\sigma_i^2(x) \equiv \sum_{j \in J_{\text{ncr}}} v_{ij}^2 a_j(x), \quad \forall i \in I_{\text{rs}}, \quad (5b)$$

where J_{ncr} is the set of indices of all *noncritical reactions*.

III. STIFFNESS IN THE STOCHASTIC REGIME

Stiffness has been studied for decades in the regime of ODEs.^{7,13,14} Briefly speaking, an ODE is called stiff if an explicit method, such as the explicit Euler method, is inefficient because stability restrictions force the stepsize h to be extremely small. Such a restriction is undesirable because the stepsize h usually has to be much smaller than the accuracy requirement. Similarly in the stochastic regime, we consider stiffness as the situation when a numerical method has to take excessively smaller stepsizes than what is necessary for the accuracy requirement to follow the system dynamics.

A. Stiffness in the SSA method

Consider the following system:

Assume the pair of reversible reaction channels is much faster than the third reaction channel. SSA simulation follows every firing reaction. Since the propensities of the two fast reaction channels are much larger than that of the third reaction

channel, SSA spends most of the CPU time on the simulation of the pair of fast reversible reactions although this pair often cancel the effects of each other and does not change the system dynamics at all. Since the computational cost of the SSA is proportional to the number of firing reactions, the large numbers of reversible reaction firing naturally make the SSA simulation very inefficient. Such a situation occurs very often in practical problems. For example, in the stochastic simulation of the heat shock response (HSR) model,^{15,16} the number of firings for all 61 reaction channels is plotted in log scale in Fig. 2. The scale differences are obvious. Three pairs of reaction channels take 94% of the total number of firings. In other words, 94% of the CPU time is used to simulate these three pairs of reactions.

To avoid stiffness in SSA simulation the slow-scale SSA method^{8,17} was proposed. This method assumes that the fast reaction channels quickly reach PE and simulates only the slow reaction channels. The dynamics of fast reaction channels are maintained by solving algebraic equations⁸ given by the PE assumption and conservation laws or by sampling in a small time window of the fast reaction channels.¹⁷ For example, the system (6) can be reduced to

where S_t is the sum of S_1 and S_2 , and $\hat{c}_3$ is given by

$$\hat{c}_3 = \frac{c_1 c_3}{c_1 + c_2}. \quad (8)$$

Equation (8) is obtained by directly solving the algebraic equations given by the PE assumption. It can be proved⁸ that the slow-scale system (7) with the new reaction rate (8) reproduces the stochastic dynamics of the c_3 channel in Eq. (6) very well if $c_1 + c_2 \gg c_3$. For the HSR model, by selecting six pairs of fast reactions,⁸ the slow-scale SSA

FIG. 2. The numbers of firing for all reaction channels in the HSR model.

method reproduces the system dynamics reasonably well with a computational cost of only 1% of the original SSA.

B. Stiffness in the Tau-Leaping method

Ideally, the Tau-Leaping method can leap over many reactions in one simulation step for the example system (6), if the firing of the two fast reactions could cancel each other exactly. However, the number of firings for each reaction channel is approximated by a Poisson random number K_j . Due to the stochastic fluctuation of the K_j 's, exact cancellation can hardly be achieved. Instead the τ values have to be limited by Eq. (4). For the example system (6), τ can be estimated by

$$\min \left\{ \frac{\max(\epsilon \min(x_1, x_2), 1)}{|c_1 x_1 - c_2 x_2|}, \frac{[\max(\epsilon \min(x_1, x_2), 1)]^2}{|c_1 x_1 + c_2 x_2|} \right\}. \quad (9)$$

For the fast reaction pair, c_1 and c_2 are very large. Equation (4) leads to a very small τ . This limitation can also be explained by numerical stability.¹⁸ It was proved¹⁸ that, to maintain numerical stability, τ should satisfy

$$\tau < \frac{2}{c_1 + c_2}. \quad (10)$$

For large values of c_1 and c_2 , this stability restriction forces τ to be very small.

If the partitioning can be efficiently implemented, the slow-scale SSA can be extended to develop the slow-scale Tau-Leaping method.¹⁰ In that case, fast reactions will not appear in the τ -selection formula and a larger τ value can be chosen. Another solution is the implicit Tau-Leaping method,¹⁹ given by

$$X^{(it)}(t + \tau) = x + \sum_{j=1}^M v_j [P(a_j(x), \tau) - a_j(x)\tau + a_j(X^{(it)}(t + \tau))\tau]. \quad (11)$$

For the example system (6), the implicit Tau-Leaping method is numerically stable regardless of c_1 and c_2 . Thus no matter how large c_1 and c_2 are, the implicit Tau-Leaping method can still choose a relatively large τ value. However, the implicit Tau-Leaping method has to solve a nonlinear equation in every step. The computational cost is much higher than the explicit Tau-Leaping method. It is thus highly desirable to detect and reduce stiffness for the explicit Tau-Leaping method.

IV. STIFFNESS DETECTION AND REDUCTION RELATED TO QSS

A. QSS assumption

When a species is produced and consumed with balanced rates and if these two rates are much faster than the changing rates of other species, there may exist a time interval where other species do not change or change very little while this species reaches a dynamical balance. Such a state is called the QSS. In stochastic simulation QSS indicates that the influx,

measured by the sum of propensities of all reactions that produce this species, is assumed to be equal to the outflux, measured by the sum of propensities of all reactions that consume this species. Call this assumption the quasi steady state assumption (QSSA), though it is also referred to as Bodenstein–Semenov kinetics⁹ or the pseudosteady state assumption.²⁰

The QSS assumption's application in SSA has been studied,⁹ while its application in the tau-leaping context has not. To see the difference between SSA and Tau-Leaping when the system is in QSS, consider the system (6) with parameters $c_3 = 10^4$, $c_1 = c_2 = 1$, $x_1 = 10^4$, and $x_2 = 1$. For this system S_2 is in QSS. The system can be reduced to (7) with a modified $\hat{c}_3$,¹¹

$$\hat{c}_3 = \frac{c_1 c_3}{c_1 + c_2 + c_3}. \quad (12)$$

For SSA, this reduction leads to an average stepsize around 10^{-4} while the average stepsize for the original system is around 5×10^{-5} , only a twofold improvement. However, for the Tau-Leaping method with $\epsilon = 0.05$ and the stepsize from the τ -selection formula (4), the reduced system allows a stepsize around 0.05, while the original system only allows a stepsize around 5×10^{-5} , the same as the average stepsize for SSA. Thus the efficiency gain through the reduction can be as high as 10^3 for the QSS assumption coupled with the Tau-Leaping method. This observation motivates improving the Tau-Leaping method based on the QSS assumption.

Why is the tau-leaping stepsize for the original system so small? Consider the τ values for all state variables in this small system; S_3 is not a reactant. The leap condition for system (6) only applies to S_1 and S_2 . For S_1 , the τ value is limited by [see Eq. (4)]

$$\tau_1 = \min \left\{ \frac{\epsilon x_1}{|c_1 x_1 - c_2 x_2|}, \frac{(\epsilon x_1)^2}{c_1 x_1 + c_2 x_2} \right\} \approx 0.05, \quad (13)$$

while the accuracy of S_2 limits the τ value by

$$\tau_2 = \min \left\{ \frac{1}{|c_1 x_1 - (c_2 + c_3)x_2|}, \frac{1}{c_1 x_1 + (c_2 + c_3)x_2} \right\} \approx 5 \times 10^{-5}. \quad (14)$$

From Eqs. (13) and (14), the required small stepsize comes from the accuracy requirement for S_2 . The question is: Is this accuracy requirement necessary?

Consider the general case. These τ values measure the time scale for which the population of chemical species will not have a noticeable change. If the tau value τ_0 required for one species S is much smaller than the tau values τ_i required for other species, then during the next time interval $[t, t + \hat{\tau}]$, where $\tau_0 < \hat{\tau} < \min_i \tau_i$, the other species will not have a substantial change and their populations can be treated as constants. Thus the reactions related to this species S can be simplified as

where A denotes a species with a constant population. This system has been studied in Gillespie's handbook²¹ as a

“payroll process.” The propensities for both reactions are given as

$$a_1(x) = cA, \quad a_2(x) = \lambda x, \quad (16)$$

where x is the population for S . For this system, $1/\lambda$ gives the relaxation time, the time scale for the system to reach a steady state, and X follows a Poisson distribution with mean and variance cA/λ .²¹ Around this steady state, the number of reactions that produce S and the number of reactions that consume S are balanced. Moreover, the covariance of X at different times $t_1 \leq t_2$ is estimated by

$$\text{cov}\{X(t_1), X(t_2)\} \approx \frac{cA}{\lambda} e^{-\lambda(t_2-t_1)}. \quad (17)$$

When $t_2 - t_1 \gg 1/\lambda$, $X(t_2)$, and $X(t_1)$ are almost not correlated. Thus if $x(t)$ is close to cA/λ , for a time scale larger than $1/\lambda$, $X(t)$ can be treated as a random variable following a Poisson distribution $P(cA/\lambda)$. In other words, the dynamics of $X(t)$ can be ignored and the accuracy requirement on $X(t)$ (restricting the τ value) is *not* necessary at all. This accuracy restriction is the source of stiffness for the Tau-Leaping method.

B. Stiffness detection

Define *related species* of any species S as those species that are involved with a reaction in which S is either a reactant or a product. According to the above analysis of the connection between QSS and stiffness, two conditions are necessary for QSS. One is that the τ value for a QSS species should be much smaller than the τ values for its related species. The other is that the population of this QSS species should be close to its QSS mean. These two conditions underpin the proposed stiffness detection algorithm.

Let the τ value for state variable X_i be

$$\tau_i = \left\{ \frac{\max\{\epsilon x_i/g_i, 1\}}{|\mu_i(x)|}, \frac{\max\{\epsilon x_i/g_i, 1\}^2}{\sigma_i^2(x)} \right\}. \quad (18)$$

Each τ_i represents the accuracy requirement for one state variable. The original stepsize selection method chooses the minimum τ_k of all these values. If the minimum τ_k corresponds to species S_k and τ_k is much smaller than the τ values for all the *related species* of S_k , then check if X_k is in QSS. A state variable is considered to be in QSS if its influx and outflux (measured by the sum of all corresponding propensities) are close to each other. Thus the state variable X_k is in QSS if and only if

$$\sum_{j:v_{kj}>0} a_j(x) \approx \sum_{j:v_{kj}<0} a_j(x). \quad (19)$$

If X_k is in QSS, then X_k is considered a stiff variable and the accuracy requirement for it is removed (τ_k is omitted from the stepsize calculation). Then the detection algorithm continues with the next larger τ value until a nonstiff variable is found or the list is exhausted.

The procedure for stiffness detection is summarized below:

- (1) Calculate all τ values in the hybrid SSA/Tau-Leaping method.

- (2) Sort the τ values and find the index i corresponding to the smallest τ value. If τ_i is much smaller than the τ values for all the *related species* of S_i , go to step 3. Otherwise, stop and record τ_i as the stepsize for the next tau-leaping calculation.
- (3) Measure the influx and outflux for the state variable X_i . If Eq. (19) is satisfied (X_i is considered to be in QSS and is a stiff variable), then remove τ_i from the list, and if the list is not empty, return to step 2. Otherwise, stop and record τ_i as the stepsize for the next tau-leaping calculation.

C. Stiffness reduction

1. Model reduction method

The next step after identifying those stiff variables in QSS is to reduce the stiffness corresponding to these stiff variables. Having already removed the accuracy requirements for the stiff variables during the τ -selection procedure and obtained a much larger stepsize, it is necessary to change the simulation method to maintain the accuracy without accuracy control on the stiff variables. The implicit Tau-Leaping method is a natural, but very expensive choice. Another choice is model reduction based on the QSS assumption. For this, calculate the mean values of the stiff variables from the algebraic equation derived from the QSS assumption,

$$\sum_{j:v_{ij}<0} a_j(x) = \sum_{j:v_{ij}>0} a_j(x). \quad (20)$$

For example, in system (6), if X_2 is detected as the stiff variable, then the system can be reduced to (7). The corresponding QSS condition,

$$(c_2 + c_3)X_2 = c_1X_1, \quad (21)$$

yields

$$X_2 = \frac{c_1}{c_2 + c_3} X_1. \quad (22)$$

Letting $X_t = X_1 + X_2$ yields

$$X_2 = \frac{c_1}{c_1 + c_2 + c_3} X_t. \quad (23)$$

Note that Eq. (23) would be ideal since it automatically leads to the reduced system (7). Equation (22) also approximates Eq. (23) very well if X_2 is detected as a stiff variable, which implies that $X_2 \ll X_1$.

If a large system contains reactions such as (6), each occurrence of (6) can be reduced to the form of (7). The reduced model does not contain the stiff variables. The stiff variables in the original model are updated based on the QSS equations, such as Eq. (23), or more accurately, by Poisson random numbers with mean values given by those equations.

2. QSS Tau-Leaping method

The above model reduction method can only be applied to models for which the fast reactions can be analytically reduced and thus is not general. If S_2 in Eq. (6) is also related to

FIG. 3. Stiff variable X with influx reactions and outflux reactions.

other complicated reactions and those reactions are not negligible, the model reduction method may not be applicable. To simulate the original system directly, one might simply treat a stiff variable as its mean value and implement the Tau-Leaping method, but that is not right. To see this, consider the system (6) again. Replacing X_2 by its mean value and generating three Poisson random numbers as the numbers of reactions for the three reaction channels, the consumption of S_1 and the production of S_3 will not be equal as the reduced system indicates because these three random variables are independent. The difference can be very large since the system is stiff.

Observe that these three numbers are not independent at all, but are related according to the QSS assumption: the influx and outflux for a QSS variable should be equal. Thus it is reasonable to require that the total number of influx firings be equal to the total number of outflux firings. Thus generating the total number of influx firings first forces a sum constraint on the numbers of the outflux firings. This conditional distribution does not follow Poisson distribution anymore, but a multinomial distribution. This observation leads to a new simulation method (QSS Tau-Leaping) based on the QSS assumption.

Suppose X is a stiff variable as shown in Fig. 3, where the left-hand arrows indicate the influx reactions and the right-hand arrows represent the outflux reactions related to X . According to the QSS assumption, the influx should equal the outflux, a property that the simulation must preserve. The QSS Tau-Leaping method is formulated as follows:

- (1) Let τ be the stepsize decided in the τ selection and stiffness detection procedure. Generate the firing num-

ber $K_j(\tau; x, t)$ according to Poisson random distribution (2) for each reaction channel j except the outflux reaction channels of stiff variables.

- (2) For each stiff variable, calculate the total number of firings Ψ of the influx reactions. Let l be the total number of outflux reaction channels for this variable. Then the number of firings for each outflux reaction j of this stiff variable is approximated by a multinomial random variable that can be generated as

$$K_j(x, t) = B \left(\Psi - \sum_{i=1}^{j-1} K_i(x, t), \frac{a_j(x)}{\sum_{i=j}^l a_i(x)} \right); \quad (24)$$

$$j = 1, \dots, l;$$

where $B(n, p)$ represents a binomial random number.

- (3) Update each stiff variable using the QSS assumption (20).
- (4) Update all other state variables via $X(t + \tau) = x + \sum_{j=1}^M v_j K_j$.

Note that Eq. (20) often leads to a simple linear equation for the stiff variable. However, if not, then Newton's method is used to solve Eq. (20).

V. NUMERICAL EXPERIMENT

Below we present numerical results for two problems. Algorithms were implemented in StochKit, a stochastic simulation toolkit developed by Li *et al.*²² Simulations were performed on a 64 bit Mac OS machine with a 2.4 GHz Intel(R) Core 2 Duo CPU and 4 GB memory.

A. Stiff decaying dimerization model

In order to test the accuracy and efficiency of the proposed methods, we first apply these two methods on a simple decaying dimerization model, which was originally proposed by Gillespie.⁵ The model consists of three species S_1 , S_2 , and

FIG. 4. Histograms of S_3 simulated by the hybrid SSA/Tau-Leaping method, model reduction method, and QSS Tau-Leaping method.

TABLE I. The numbers in the table are simulation times for the decaying dimerization model.

Total runs of simulation	2000	10 000
Hybrid SSA/Tau-Leaping	0m19.536s	1m41.084s
Model reduction	0m0.063s	0m0.305s
QSS Tau-Leaping	0m0.103s	0m0.504s

S_3 and four reaction channels as follows:

$$S_1 \xrightarrow{c_1} 0, \quad (25a)$$

$$S_1 + S_1 \xrightarrow{c_2} S_2, \quad (25b)$$

$$S_2 \xrightarrow{c_3} S_1 + S_1, \quad (25c)$$

$$S_2 \xrightarrow{c_4} 0. \quad (25d)$$

We set the reaction rates as $c_1 = 0.1$, $c_2 = 10$, $c_3 = 5$, and $c_4 = 0.01$. The initial conditions are given by $x_1(0) = 100$, $x_2(0) = 10\,000$, and $x_3(0) = 0$.

In this numerical experiment, the stiff variable corresponds with S_1 . The order of magnitude of τ for S_1 is 10^{-6} , while the smallest τ of the other species is about 8×10^{-3} . The influx and outflux propensities are close to each other for S_1 . For example, at a certain time, the influx and outflux propensities are 5.15×10^4 and 4.99×10^4 , respectively. Table I shows the CPU time for the hybrid SSA/Tau-Leaping method, the model reduction method, and the QSS Tau-Leaping method applied to the decaying dimerization model.

The running time for the model reduction method is only 0.3% of that for the hybrid SSA/Tau-Leaping method. The running time for the QSS Tau-Leaping method is slightly longer than that for the model reduction method. The accuracy of the other nonstiff variable S_3 is assessed by comparing the histograms obtained from simulations by the hybrid SSA/Tau-Leaping method, the model reduction method, and the QSS Tau-Leaping method. Figure 4 shows the histogram differences for S_3 simulated by the hybrid SSA/Tau-

TABLE II. The influx and outflux propensities of MisP and HSF1 at a certain step.

Propensities	Influx	Outflux
MisP	$1.14503e + 04$	$1.14504e + 04$
HSF1	$5.99972e + 03$	$5.99983e + 03$

Leaping method, model reduction method, and QSS Tau-Leaping method. The differences among them are small.

B. Heat shock proteins regulation model

We now consider a more complicate model, the HSP regulation model.¹² HSPs are synthesized in increased amounts after brief exposure of cells to an elevated temperature. They also exist in cells at all times and are upregulated by other stress, such as irradiation. Stress-induced transcription of HSPs requires the activation of heat shock factor (HSF). Once the HSF is activated by stress, it will produce more and more HSPs. The HSPs are also involved in many cellular functions, such as the folding of nascent proteins, the refolding of denatured proteins, and the prevention of protein aggregation. Figure 5 shows two pathways in which HSPs are involved, one is the pathway of protein turnover and the other one is the autoregulation of HSPs.

Initially the majority of HSF1 (heat shock factor) is bound to Hsp90 (heat shock proteins). The new proteins are synthesized and correctly folded into its native state NatP. Some of the folded proteins may misfold with the help of ROS (reactive oxygen species). The misfolded proteins (MisP) may further degrade or form aggregates (AggP). When the outside temperature suddenly increases or other stress has been induced, more and more native state NatP will become misfolded MisP. An increase in the number of misfolded proteins leads to Hsp90 disassociating from HSF1 and binding to MisP. Then the misfolded proteins will be refolded with the help of ATP. At the same time, HSF1 is free to form dimer (DiH), and then trimers (TriH), with both reactions being reversible. TriH can bind to heat shock element (HSE) to form a complex (HSETriH), which then activates the transcription of Hsp90 leading to an increase in the level of Hsp90 in the cell. The increasing of Hsp90 level will finally increase the chance of any misfolded protein being correctly refolded. The detailed reactions, kinetic parameters, and initial conditions of the model can be found in Proctor *et al.*¹²

In this numerical experiment, the stiff variables correspond with either MisP or HSF1. The order of magnitude of τ for MisP and HSF1 is 10^{-6} , while the smallest τ of the other species is only about 10^{-4} . Moreover, the influx and outflux

FIG. 5. Heat shock proteins regulation pathway.

TABLE III. The numbers in the table are simulation times for the HSP model.

Total runs of simulation	100	2000
Hybrid SSA/Tau-Leaping	1m57.457s	38m29.432s
Model reduction	0m2.326s	0m46.950s
QSS Tau-Leaping	0m2.897s	0m57.334s

FIG. 6. Histograms of Hsp90 simulated by the hybrid SSA/Tau-Leaping method, model reduction method, and QSS Tau-Leaping method.

FIG. 7. Histograms of ROS simulated by the hybrid SSA/Tau-Leaping method, model reduction method, and QSS Tau-Leaping method.

FIG. 8. Histograms of ADP simulated by the hybrid SSA/Tau-Leaping method, model reduction method, and QSS Tau-Leaping method.

FIG. 9. Histograms of ATP simulated by the hybrid SSA/Tau-Leaping method, model reduction method, and QSS Tau-Leaping method.

propensities are quite close to each other for MisP and HSF1, respectively, as shown in Table II.

For the stiffness reduction methods with the QSS assumption, Table III gives the CPU time for the hybrid SSA/Tau-Leaping method, the model reduction method, and the QSS Tau-Leaping method applied to the HSP model. Figures 6–9 show the histogram differences for species Hsp90, ROS, ADP, and ATP simulated by the hybrid SSA/Tau-Leaping method, model reduction method, and QSS Tau-Leaping method. The differences are also very small among these three methods. It demonstrates the effectiveness of both stiffness reduction methods on stiff systems very well.

Based on the above two numerical experiments, comparing the two stiffness reduction methods, their accuracies are similar while the model reduction method is a little more efficient. The model reduction method can only be applied in certain situations, where the fast reactions can be analytically reduced, whereas the QSS Tau-Leaping can be used for any stiff variables, and thus is completely general.

VI. CONCLUSION AND DISCUSSION

Stiffness causes small stepsizes and computational inefficiency in discrete stochastic simulation. This paper proposed a method to detect stiffness in the discrete stochastic regime using the tau selection formula in the Tau-Leaping method and the QSS assumption. Then the stiff model can either be reduced based on the QSS assumption to a less stiff model or be simulated via the QSS Tau-Leaping method proposed here. Numerical results on a decaying dimerization model and a HSP regulation model show that both stiffness reduction methods can greatly improve the efficiency of the simulation, while the accuracy of the simulation is well maintained. The proposed methods provide a systematic way to automatically detect and reduce stiffness for stochastic simulation of chemical systems. We are aware that more rigorous error analysis is desired for the two stiffness reduction methods. The com-

plexity of that task is beyond the scope of this paper. We plan to discuss that in the future.

ACKNOWLEDGMENTS

This work was supported by the University of California, Santa Barbara under award KK9103, the National Science Foundation under award CCF-0726763, CCF-0953590, AFOSR Grant No. FA9550-09-1-0153, and the National Institutes of Health under award GM073744 and GM078989.

- ¹H. McAdams and A. Arkin, *Proc. Natl. Acad. Sci. U.S.A.* **94**, 814 (1997).
- ²A. Arkin, J. Ross, and H. McAdams, *Genetics* **149**, 1633 (1998).
- ³H. H. McAdams and A. Arkin, *Trends Genet.* **15**(2), 65 (1999).
- ⁴N. Fedorof and W. Fontana, *Science* **297**, 1129 (2002).
- ⁵D. Gillespie, *J. Comput. Phys.* **22**, 403 (1976); *J. Phys. Chem.* **81**, 2340 (1977); *J. Chem. Phys.* **115**, 1716 (2001).
- ⁶D. Gillespie and L. Petzold, *J. Chem. Phys.* **119**, 8229 (2003).
- ⁷U. M. Ascher and L. R. Petzold, *Computer Methods for Ordinary Differential Equations and Differential-Algebraic Equations* (Society for Industrial and Applied Mathematics, Philadelphia, PA, 1998).
- ⁸Y. Cao, D. Gillespie, and L. Petzold, *J. Chem. Phys.* **123**, 054104 (2005); **123**, 144917 (2005); **124**, 044109 (2006); *J. Comput. Phys.* **206**, 395 (2005).
- ⁹C. Rao and A. Arkin, *J. Chem. Phys.* **118**, 4999 (2003).
- ¹⁰Y. Cao and L. Petzold, *Comput. Methods Appl. Mech. Eng.* **197**(43–44), 3472 (2008).
- ¹¹D. Gillespie, Y. Cao, K. Sanft, and L. Petzold, *J. Chem. Phys.* **130**(6), 064103 (2009).
- ¹²C. J. Proctor, C. Soti, R. J. Boys, C. S. Gillespie, D. P. Shanley, D. J. Wilkinson, T. B. L. Kirkwood, *Mech. Ageing Dev.* **126**(1), 119 (2005).
- ¹³G. Dahlquist, *BIT Numer. Math.* **3**(1), 27 (1963).
- ¹⁴E. Hairer and G. Wanner, *Solving Ordinary Differential Equations II: Stiff and Differential-Algebraic Problems*, 2nd ed. (Springer-Verlag, Berlin, 1996).
- ¹⁵H. Kurata, H. El-Samad, T. Yi, M. Khammash, and J. C. Doyle, in *Proceedings of the 40th IEEE Conference on Decision and Control*, pp. 837–842, Orlando, FL, Dec 4–7, 2001 (IEEE Control Systems Society, Piscataway, NJ).
- ¹⁶H. Kurata, M. Khammash, and J. Doyle, *Third International Conference on Systems Biology (ICSB2002)*, Stockholm, Sweden, December 2002.

- ¹⁷W. E., D. Liu, and E. Vanden-Eijnden, *J. Chem. Phys.* **123**(19), 194107 (2005).
- ¹⁸Y. Cao, L. Petzold, M. Rathinam, and D. Gillespie, *J. Chem. Phys.* **121**, 12169 (2004).
- ¹⁹M. Rathinam, L. Petzold, Y. Cao, and D. Gillespie, *J. Chem. Phys.* **119**, 12784 (2003).
- ²⁰L. A. Segel and M. Slemrod, *SIAM Rev.* **31**, 446 (1989).
- ²¹D. Gillespie, *Markov Processes: An Introduction for Physical Scientists* (Academic, 1992).
- ²²H. Li, Y. Cao, L. Petzold, and D. Gillespie, *Biotechnol. Prog.* **24**, 56 (2008).