Adaptive deployment of model reductions for tau-leaping simulation

Sheng Wu, Jin Fu, and Linda R. Petzold

Citation: The Journal of Chemical Physics **142**, 204108 (2015); doi: 10.1063/1.4921638
View online: http://dx.doi.org/10.1063/1.4921638
View Table of Contents: http://scitation.aip.org/content/aip/journal/jcp/142/20?ver=pdfcov
Published by the AIP Publishing

Adaptive deployment of model reductions for tau-leaping simulation

Sheng Wu,[a] Jin Fu,[b] and Linda R. Petzold[c]
Department of Computer Science, University of California Santa Barbara, Santa Barbara, California 93106, USA

(Received 29 January 2015; accepted 13 May 2015; published online 27 May 2015)

Multiple time scales in cellular chemical reaction systems often render the tau-leaping algorithm inefficient. Various model reductions have been proposed to accelerate tau-leaping simulations. However, these are often identified and deployed manually, requiring expert knowledge. This is time-consuming and prone to error. In previous work, we proposed a methodology for automatic identification and validation of model reduction opportunities for tau-leaping simulation. Here, we show how the model reductions can be automatically and adaptively deployed during the time course of a simulation. For multiscale systems, this can result in substantial speedups. © *2015 AIP Publishing LLC.* [http://dx.doi.org/10.1063/1.4921638]

I. INTRODUCTION

In the study of cellular biological systems, discreteness and stochasticity are often observed in the dynamics of the system due to small populations of key reactant species.[1–3] The Stochastic Simulation Algorithm (SSA)[4] and the explicit tau-leaping algorithm[5,6] are often used to simulate the dynamics of such systems.

The SSA is an exact kinect Monte-Carlo method. The explicit tau-leaping algorithm is an approximate method for chemically reacting systems that can often substantially outperform the SSA. However, neither of the two algorithms is efficient for stiff systems, where vastly different time scales are involved.[7,8] Thus, numerous model reductions have been proposed to accelerate both algorithms, including the slow-scale SSA (ssSSA),[7] the nested SSA,[9,10] the stochastic quasi-steady-state approximation (sQSSA),[11–13] the stochastic Michaelis-Menten model reduction (M-M),[11,12,14,15] and the time-dependent solution method.[16]

Among these methods, the ssSSA[7] eliminates the need to simulate many fast reactions by approximating the fast subsystems that reach stochastic partial equilibrium very quickly between two consecutive slow reactions by their partial equilibrium states. The nested SSA applies a similar idea, but instead of directly calculating the partial equilibrium, it uses short SSA simulations of the fast subsystems to approximate partial equilibrium. Both methods are effective in accelerating SSA simulations because the SSA simulates every reaction event. In contrast, the sQSSA[11–13] eliminates certain fast-changing species from the simulation by approximating the fast-changing species by a near-stationary distribution with their quasi-steady state. It is effective in accelerating tau-leaping simulations because the step size of tau-leaping simulations is limited by how fast the populations of species change.[6,13,17] Alternatively, the stochastic M-M approximation is designed for enzyme-substrate systems and can be derived from the sQSSA approach[11,12] or the ssSSA approach[14,15] under appropriate conditions. Recent research has shown[18,19] that the conditions under which stochastic M-M approximation is valid are more restrictive than those for the classical deterministic M-M conditions.[20] Additionally, the time-dependent solution method[16] is an extension of the sQSSA. It gives more accurate results for fast changing species, regardless of whether they are in quasi-steady state or not.

With all of these model reduction techniques, one of the major practical challenges is to efficiently identify the opportunities for model reductions and to automatically deploy them. To meet this challenge, we have proposed an automatic model analysis algorithm that can identify situations where specific model reductions may be deployed safely and efficiently for both SSA and tau-leaping simulations.[21] In related work, an automatic ssSSA algorithm was proposed to adaptively identify and apply the slow-scale SSA approximation for SSA simulations.[22] However, due to the different dynamic features between the SSA and tau-leaping algorithms,[17] there is no significant benefit in applying the automatic ssSSA algorithm to tau-leaping simulations. On the other hand, the sQSSA and stochastic M-M are two of the most efficient model reduction techniques for accelerating tau-leaping simulations because they eliminate the need to directly simulate the reactions that result in the fastest-changing species in a system. Reference 13 proposed a sQSSA algorithm to automatically identify and apply sQSSA for tau-leaping. But the algorithm cannot be generalized to detect model reduction opportunities other than the application of sQSSA to single fast-changing species in tau-leaping (for example, stochastic M-M).

In this paper, we propose a unified framework to automatically and adaptively identify, apply, and deactivate model reductions in tau-leaping simulations. The framework works for both the sQSSA and the stochastic M-M, and can be extended to other model reductions focusing on eliminating reactions that result in elimination of fast-changing species from a simulation.

The outline of this paper is as follows. In Sec. II A, we review the tau-leaping algorithm, especially the implementation

[a]Electronic mail: sheng@cs.ucsb.edu
[b]Electronic mail: fujin@cs.ucsb.edu
[c]Electronic mail: petzold@cs.ucsb.edu

details of the explicit tau-leaping algorithm with adaptive step size selection.[6] In Sec. II B, we review the sQSSA and stochastic M-M model reductions. In Sec. III, we describe our new adaptive framework. In Sec. IV we show how sQSSA and stochastic M-M can be applied in the framework. In Sec. V, we apply the tau-leaping algorithm with automatic sQSSA and stochastic M-M model reductions to three realistic models and demonstrate the efficiency and effectiveness of the algorithm.

II. BACKGROUND

A. Explicit tau-leaping method

The explicit tau-leaping algorithm is a widely used approximate method to accelerate the SSA.[5,23] Suppose we have a well-stirred chemical reaction system with n molecular species $S_1,\ldots,S_n$ and m reaction channels $R_1,\ldots,R_m$, and the volume Ω and temperature of the system are constant. Let $x_i(t)$ denote the population of species S_i at time t. Then, the state of the system at time t is given by the state vector $\mathbf{x}(t) = (x_1(t),\ldots,x_n(t))^T$. Each reaction R_j is characterized by two quantities: the probability $a_j(\mathbf{x})\,dt$ that one R_j reaction will occur in the next infinitesimal time interval $[t,t+dt)$, given $\mathbf{x}(t) = \mathbf{x}$, where $a_j(\mathbf{x})$ is called the reaction's propensity function; and $\boldsymbol{\nu}_j$, the change to the system's state vector if one R_j reaction occurs. $\boldsymbol{\nu}_j$ is called the stoichiometry vector of R_j. Because the system is well-stirred with constant volume and temperature, the reactions' propensity functions depend only on the system state vector. For example, for a bimolecular mass action reaction $S_1 + S_2 \to P$, $a_j(\mathbf{x})$ has the form $c_j x_1 x_2$ (or $c_j \frac{1}{2} x_1(x_1-1)$ if $S_1 = S_2$), where c_j is a constant.

While the SSA simulates every reaction event, the explicit tau-leaping method selects a time interval τ and fires multiple reactions during the time interval. The basic idea of the tau-leaping method is to approximate the number of firings of each reaction during a time interval $[t,t+\tau)$ by a Poisson random number $\mathcal{P}(a_j(\mathbf{x})\tau)$. The approximation is valid only if the propensity functions of all the reactions are nearly constant during the time interval. Then, the state of the system can be advanced by the formula

$$\mathbf{x}(t+\tau) \approx \mathbf{x}(t) + \sum_{j=1}^{m} \mathcal{P}(a_j(\mathbf{x})\tau)\,\boldsymbol{\nu}_j. \tag{1}$$

The requirement that the propensities are nearly constant during the time interval is called the leap condition: for some $\varepsilon \ll 1$,

$$|\Delta_\tau a_j(\mathbf{x})/a_j(\mathbf{x})| \le \varepsilon, \quad \text{for all } j = 1,\ldots,m, \tag{2}$$

where $\Delta_\tau a_j$ is the change of a_j during the time interval $[t,t+\tau)$. The step size τ must be chosen carefully to satisfy the leap condition. The most widely used strategy for step size selection of mass action systems is due to Cao *et al.*[6] In that strategy, the leap condition is expressed in terms of the changes in species' populations,

$$\tau = \min_i \left\{ \frac{\max\{\varepsilon x_i/g_i, 1\}}{|\sum_j \nu_{ij} a_j(x)|}, \frac{\max\{\varepsilon x_i/g_i, 1\}^2}{|\sum_j \nu_{ij}^2 a_j(x)|} \right\}, \tag{3}$$

where $\varepsilon \ll 1$ is the preset accuracy control parameter, ν_{ij} are the stoichiometric coefficients, and g_i is the highest order of reaction in which species S_i appears as a reactant. In this way, the changes of all the propensities are constrained in a uniform manner.

The explicit tau-leaping algorithm with adaptive step size selection is given as follows.[6]

0. Initialization. Initialize t and $\mathbf{x}$, and calculate the propensities $\mathbf{a}(\mathbf{x})$.
1. Identify currently critical reactions. A reaction R_j is critical if $a_j > 0$ and the maximum number of times L_j that R_j can fire before one of its reactants is exhausted is less than a threshold n_c. (We set $n_c = 10$.)
2. Select step size τ.
   - 2.1 Divide all the reactions into two sets based on the result of step 1: critical reaction set and non-critical reaction set.
   - 2.2 Calculate non-critical step size τ' by using the leap condition (3) with only the non-critical reaction set.
   - 2.3 Calculate critical step size τ'' by generating a SSA step size with only the critical reaction set.
   - 2.4 The actual step size τ is the smallest of τ', τ'', or $t_{end} - t$.
3. Calculate the number of firings of each reaction in the step.
   - 3.1 For a non-critical reaction $R_{j'}$, set the number of firings $k_{j'}$ of $R_{j'}$ to be a sample of the Poisson random variable with mean $a_{j'}(\mathbf{x})\tau$.
   - 3.2 For critical reactions, at most one of them will fire at most one time. If $\tau \ne \tau''$, set the number of firings of all of them to 0. Otherwise ($\tau = \tau''$), choose one critical reaction $R_{j''}$ by applying a SSA reaction selection step with only the critical reaction set, set the firing number $k_{j''}$ of $R_{j''}$ to 1 and the firing number of all other critical reactions to 0.
4. Check for negative populations. If there is a negative component in $\mathbf{x} + \sum_j k_j \boldsymbol{\nu}_j$, reduce τ' by half and go to Step 2.4. Otherwise, execute the step by setting $t \leftarrow t + \tau$ and $\mathbf{x} \leftarrow \mathbf{x} + \sum_j k_j \boldsymbol{\nu}_j$, and update $\mathbf{a}(\mathbf{x})$ accordingly.
5. Check for efficiency. If the total number of reaction firings in the step is less than a threshold (we set it to be 10), execute a modest number of SSA steps (we set it to be 100).
6. Stop if $t = t_{end}$. Otherwise return to step 1.

Note that critical reactions are introduced to resolve the negativity problem, which is due to the fact that unbounded Poisson random variables can easily result in critical reactions to fire so many times that the low populations of some reactants are driven negative.[24] Alternative methods have also been proposed to solve the negativity problem.[25–27] Both Refs. 25 and 26 proposed to use bounded binomial random variables to approximate unbounded Poisson random variables. But they did not solve the problem that multiple reaction channels could contribute to the depletion of low population species, which is usually the cause for negative population. Reference 27 used a similar critical threshold idea but suggested to use a "confidence level" to adaptively adjust the critical threshold of population of species. For the problems we tested, the adaptive

threshold method performs at the same level with the aforementioned fixed threshold method. The performance of both methods is similarly restricted by fast-changing critical species that could be reduced by model reduction techniques. Thus, we opted for the fixed threshold method for its simplicity.

B. Model reductions for tau-leaping

Although the tau-leaping algorithm is able to speed up stochastic simulation in many cases, sometimes it is not efficient for models with vastly different time scales, especially the ones with rapidly changing low-population species. Various model reductions have been proposed to accelerate stochastic simulation of systems with different dynamic features. Some of the model reduction methods are more effective for tau-leaping simulation because they focus on separating the time scales of the rates of change of species populations.[17,21] For our purposes, we will review the sQSSA[11–13] and the stochastic M-M approximation.[11,12,14,15]

The sQSSA accelerates stochastic simulation by approximating fast-changing species with their stochastic quasi-steady-state distribution.[11–13] A species S_i is in stochastic quasi-steady-state if

$$\frac{dP(x_i|\mathbf{x}_s)}{dt} \approx 0, \tag{4}$$

where $\mathbf{x}_s$ are the species that are not in quasi-steady-state. Assuming that $x_i|\mathbf{x}_s$ with fixed $\mathbf{x}_s$ is Markovian, then by applying quasi-steady-state approximation (4), the steady-state distribution of $P(x_i|\mathbf{x}_s)$ with fixed $\mathbf{x}_s$ can be used to approximate the distribution of x_i.

The stochastic M-M approximation[11,12,14,15] replaces the set of three reactions,

$$E + S \underset{c_2}{\overset{c_1}{\rightleftharpoons}} C \xrightarrow{c_3} E + P, \tag{5}$$

with the single M-M reaction,

$$S \xrightarrow{c} P. \tag{6}$$

It has been shown that the stochastic M-M approximation may introduce large errors in the variance of substrate S under some conditions where the deterministic M-M approximation is valid,[18,19] thus, it is important to include $c_2 \gg c_3$ as a validity condition for the stochastic M-M approximation.[18] It has also been shown that the total Quasi-Steady-State (tQSSA) rate of the stochastic M-M approximation can give more accurate results than the QSSA rate.[28,29] Another interesting note is that the M-M reduction speeds up SSA and tau-leaping simulation differently under different conditions.[17]

III. AUTOMATIC MODEL REDUCTION FRAMEWORK FOR TAU-LEAPING

The main components of the original tau-leaping algorithm include step size selection and reaction firing. Model reductions accelerate simulations by enlarging step sizes and computing the number of firings of reactions accordingly. We propose an automatic model reduction framework that consists of three main function modules: step size selection, reaction firing, and model reduction update.

A. Framework structure

To fulfill these functions and keep the framework unified and extensible for different model reductions, the structure of the framework consists of a master controller and individual model reductions. The master controller keeps and updates a list of active model reductions, and serves as an interface between the tau-leaping algorithm and individual model reductions. During each time step, step size and reaction firing information of all active model reductions are combined in the master controller and passed to the tau-leaping algorithm. System state and time information from the tau-leaping algorithm are analyzed in the master controller to update the active model reduction list, as well as to be passed to individual model reductions for the calculation of step size adjustments and the numbers of reaction firings. The diagram of the framework is shown in Fig. 1. In the figure, t, $\mathbf{x}(t)$, and $\mathbf{a}(\mathbf{x}(t))$ are system state and time information, τ is the step size of the original model, and τ^r is the step size of the reduced model. $\{\Delta\tau_i^r\}_l$ is a list of step size adjustments for all the species due to the lth model reduction, while $\{k_j\}_l$ is a list of the number of firings of all the reactions eliminated by the lth model reduction in each step.

Each individual model reduction shares the same interface that can be called by the master controller, including motif detection (to activate model reduction), step size adjustment, reaction firing calculation, and validity criteria (to deactivate the model reduction). To incorporate a new model reduction, one needs only to instantiate these interfaces for the new model reduction.

B. Step size selection

The step size selection module of the framework is a modification of the step size selection of the original tau-leaping algorithm. The original tau-leaping algorithm divides all of the reactions into two sets based on whether a reaction is critical, and calculates two step sizes: a critical step size that is a SSA step size of all critical reactions, and a non-critical step size for each species by using the leap condition (3).

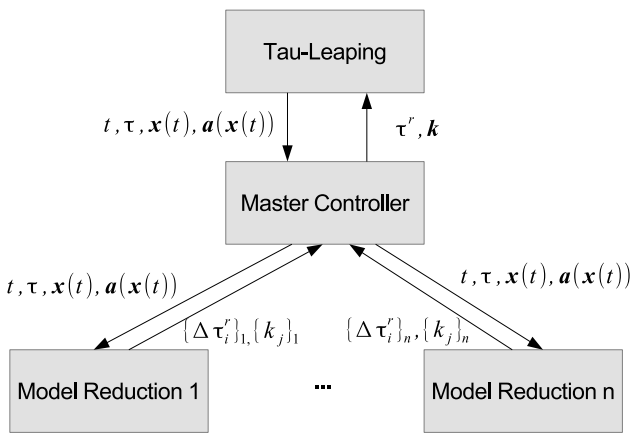

FIG. 1. Structure diagram of the automatic model reduction framework.

When a model reduction is applied to some of the reactions, we can no longer calculate these step sizes based on the original model. Instead, we must consider the reduced model. First, we categorize the reactions in both the original model and the reduced model.

Definition 1. Given a chemically reacting system consisting of reactions $\mathbf{R}_{ori} = \{R_1, \ldots, R_{m1}\}$ *and a set of model reductions* $\{MR_1, \ldots, MR_p\}$. *Let* $\mathbf{R}_{red} = \{R'_1, \ldots, R'_{m2}\}$ *be the list of reactions in the reduced model when the model reductions are applied to the system.*

- *A reaction R is an* ***unreduced reaction*** *if* $R \in \mathbf{R}_{ori}$ *and* $R \in \mathbf{R}_{red}$.
- *A reaction R is an* ***eliminated reaction*** *if* $R \in \mathbf{R}_{ori}$ *and* $R \notin \mathbf{R}_{red}$.
- *A reaction R is a* ***reduced reaction*** *if* $R \notin \mathbf{R}_{ori}$ *and* $R \in \mathbf{R}_{red}$.

In other words, the original model is the union of the unreduced reaction set and the eliminated reaction set, while the reduced model is the union of the unreduced reaction set and the reduced reaction set. Note that while the stochastic M-M is a simple conversion from the eliminated reaction set to the reduced reaction set, not all model reductions are simple conversions. Some model reductions, such as the sQSSA, eliminate corresponding reactions and create rules to calculate the numbers of firings of the eliminated reactions, as opposed to creating a list of reduced reactions.

For example, for enzyme-substrate system (5) with the stochastic M-M reduction applied, none of the reactions is an unreduced reaction, all three reactions in the original model are eliminated reactions, and the only reaction in reduced model (6) is a reduced reaction. However, for the same model (5) but with the sQSSA applied to species C, none of the reactions is an unreduced reaction, all 3 reactions in the original model are eliminated reactions, and there is no reduced reaction in the reduced model. The reduced model in sQSSA in this case is only a set of rules to calculate the number of firings of the eliminated reactions.

Next, we add the notation of critical/non-critical reactions to resolve the possible negativity problem. A critical reaction is defined in the same way as in the original tau-leaping algorithm: a reaction R_j is critical if $a_j > 0$ and the maximum number of times L_j that R_j can fire before one of its reactants is exhausted is less than a threshold n_c. (We set $n_c = 10$.) A reaction that is not a critical reaction is a non-critical reaction. This results in 6 sets of reactions: critical unreduced, non-critical unreduced, critical eliminated, non-critical eliminated, critical reduced, and non-critical reduced.

The tau-leaping algorithm will generate a non-critical step size for each species using the leap condition on the original model. The master controller gets the step size information from the tau-leaping algorithm and sends that information along with the system state information to all active model reductions. Each active model reduction adjusts the non-critical step sizes and marks the eliminated reactions or reduced reactions that need to be included in the critical step size calculation. The master controller then combines the results from all individual model reductions and sends them to the tau-leaping algorithm. The tau-leaping algorithm calculates the non-critical step size $\tau^{r\prime}$ and critical step size $\tau^{r\prime\prime}$ with adjustments from the model reduction module.

Different types of model reductions have different adjustment rules. The general rule is to remove the non-critical eliminated reaction terms from the non-critical step size calculation in (3), and/or to remove the critical eliminated reaction terms from the critical step size calculation, and/or to add the reduced reaction terms to the non-critical or critical step size calculation. There can be exceptions to conform with the calculation of the numbers of reaction firings. We will give an example in Sec. IV A.

Finally, the step size τ^r for the reduced model is set to be the minimum of the adjusted critical step size $\tau^{r\prime\prime}$, the adjusted non-critical step size $\tau^{r\prime}$, and $t_{end} - t$.

The category information for the unreduced reactions is updated by the tau-leaping algorithm, while the category information for the eliminated or reduced reactions is updated by the corresponding individual model reduction as a result of model reduction activation, deactivation, or system state change. A copy of the category information for all of the reactions is kept in the master controller for step size selection.

C. Reaction firing

The reaction firing module of the framework is a modification to the reaction firing module of the original tau-leaping algorithm. In the original tau-leaping algorithm, the numbers of firings of non-critical reactions are samples of Poisson random variables, while at most one of the critical reactions can fire at most one time in a SSA fashion. To be exact, if $\tau = \tau''$, exactly one of the critical reactions will fire once. Otherwise ($\tau \neq \tau''$), none of the critical reactions will fire.

When a model reduction is applied to some of the reactions, we need to consider the roles of the eliminated reactions and/or the reduced reactions. We calculate the number of firings of the unreduced reactions and the eliminated reactions. There are two reasons for this. First, the rules of some model reductions, such as sQSSA, do not produce reduced reactions but rather are directly applied to the eliminated reactions.[13] The second reason is that some model reductions, such as stochastic M-M, may eliminate species in the reduced model. By calculating the numbers of reaction firings in the original model, we are able to update the populations of all species, including the eliminated ones.

For the unreduced non-critical reactions, the calculation of the numbers of reaction firings is handled in the same way (Poisson random samples) as in the original tau-leaping algorithm with the updated step size τ^r. The unreduced critical reactions need to be grouped together with the reduced critical reactions, and sometimes even with the eliminated critical reactions, depending on the rules of the corresponding model reduction. An example of including an eliminated critical reaction in critical reaction selection will be given in Sec. IV A. If $\tau^r = \tau^{r\prime\prime}$, exactly one of the critical reactions will fire once, including unreduced, reduced, and possibly eliminated reactions. Otherwise ($\tau^r \neq \tau^{r\prime\prime}$), none of the critical reactions will fire. For all of the eliminated reactions that

are not involved in the critical reaction selection, including eliminated non-critical reactions and sometimes eliminated critical reactions, the numbers of reaction firings are calculated in each individual model reduction module.

The numbers of firings of unreduced non-critical reactions are calculated in the original tau-leaping algorithm. The numbers of firings of critical reactions are calculated in the master controller. And the numbers of firings of all other eliminated reactions are calculated in each individual model reduction module.

D. Model reduction update

The model reduction update module is a collection of functions that the master controller executes to keep the active model reduction list up to date. It includes an initialization module and an in-step update module.

The initialization module consists of all the processes that need to be executed only once, in the initialization stage of the simulation. We also include some static analysis in this stage to reduce the overhead of the in-step update module. The processes include the initialization of the static Petri net and the species connection graph (without weights). This module also detects sub Petri net motifs that some model reductions require. For example, the stochastic M-M model reduction requires a 4-species 3-reaction subnetwork as shown in (5).

The in-step update module includes 3 main components.

1. Update the fastness function of species, i.e., the moving time average of τ_i^{-1}'s, as defined in Ref. 21.
2. Check to see if all active model reductions are still valid and deactivate any invalid model reductions.
3. If the number of consecutive steps for which the detection algorithm has not been executed is above a given threshold, execute the detection algorithm in Ref. 21 to look for new active model reductions. (We set the threshold to 100 in our experiments.)

E. Algorithm

The tau-leaping algorithm with automatic model reduction is specified as follows.

0. Initialization. Initialize t and $\mathbf{x}$, and calculate the propensities $\mathbf{a}(\mathbf{x})$. Initialize the automatic model reduction framework.
1. In-step update of the automatic model reduction module. Update active model reduction list. Identify currently unreduced, reduced, and eliminated reactions. Identify currently critical reactions.
2. Select step size τ^r.
   - 2.1 Calculate non-critical step sizes $\{\tau_i'\}$ for all the species with the original model.
   - 2.2 For each active model reduction MR_l, collect step size information.
     - 2.2.1 Collect adjustments to the non-critical step sizes $\{\Delta\tau_i^{r\prime}\}_l$ by removing the eliminated reaction terms and adding the reduced reaction terms into the leap condition (3).
     - 2.2.2 Collect critical reduced or eliminated reactions that need to be included in the critical reaction set.
   - 2.3 Calculate non-critical step sizes for all the species $\tau_i^{r\prime} = \tau_i' + \sum_l \{\Delta\tau_{i,l}^{r\prime}\}$. Then, set $\tau^{r\prime} = \min_i\{\tau_i^{r\prime}\}$.
   - 2.4 Calculate the critical step size $\tau^{r\prime\prime}$ by generating a SSA step size with the critical reaction set.
   - 2.5 The actual step size τ^r is the smallest of $\tau^{r\prime}$, $\tau^{r\prime\prime}$, or $t_{end} - t$.
3. Calculate the number of firings of each reaction in the step.
   - 3.1 For an unreduced non-critical reaction $R_{j'}$, set the number of firings $k_{j'}$ of $R_{j'}$ to be a sample of the Poisson random variable with mean $a_{j'}(\mathbf{x})\tau^r$.
   - 3.2 For all the critical reactions in the critical reaction set collected in step 2, at most one of them will fire at most one time. If $\tau^r \neq \tau^{r\prime\prime}$, set the number of firings of all of them to 0. Otherwise ($\tau^r = \tau^{r\prime\prime}$), choose one critical reaction $R_{j''}$ by applying a SSA reaction selection step with only the critical reaction set, setting the firing number $k_{j''}$ of $R_{j''}$ to 1 and all others in the critical reaction set to 0.
   - 3.3 For eliminated reactions that are not in the critical reaction set, calculate the numbers of firings according to individual model reduction rules.
4. Check for negative populations. If there is a negative component in $\mathbf{x} + \sum_j k_j\boldsymbol{\nu}_j$, reduce $\tau^{r\prime}$ by half and go to step 2.5. Otherwise, execute the step by setting $t \leftarrow t + \tau^r$ and $\mathbf{x} \leftarrow \mathbf{x} + \sum_j k_j\boldsymbol{\nu}_j$, and update $\mathbf{a}$ accordingly.
5. Check for efficiency. If the total number of reaction firings in the step is less than a threshold (we set it to 10), execute a modest number of SSA steps (we set it to 100). Keep the fastness function of the species updated while executing SSA steps.
6. Stop if $t = t_{end}$. Otherwise, return to step 1.

IV. MODEL REDUCTION EXAMPLES

In this section, we illustrate the application of two of the model reductions, sQSSA and stochastic M-M, in the adaptive framework.

A. sQSSA

The sQSSA[11–13] removes certain fast-changing species in stochastic quasi-steady-state from the step size calculation, and calculates the numbers of firings of the reactions linked to the stochastic quasi-steady-state species, based on the assumption that the influx to the stochastic quasi-steady-state species equals the outflux from the same species. The influx is measured by the sum of the numbers of firings of all the reactions that produce this species, while the outflux is measured by the sum of the numbers of firings of all the reactions that consume this species. Reference 21 derived the identification criteria to determine whether a fast-changing species is in stochastic quasi-steady-state, and if sQSSA can be applied and is beneficial. The same criteria can also be used in each step to determine whether sQSSA is still valid. In this section, we focus on how the sQSSA is applied in the framework. Note

that the calculation of the numbers of reaction firings is very similar to that of the algorithm in Ref. 13.

Suppose that species S_l is in stochastic quasi-steady-state. The reactions generating species S_l are $R_{i_1}, R_{i_2}, \ldots, R_{i_{m1}}$ (stoichiometry coefficients $\nu_{i,l} > 0$, $i = i_1, i_2, \ldots, i_{m1}$). The reactions consuming species S_l are $R_{j_1}, R_{j_2}, \ldots, R_{j_{m2}}$ (stoichiometry coefficients $\nu_{j,l} < 0$, $j = j_1, j_2, \ldots, j_{m2}$). Then for the sQSSA that is applied to species S_l, the eliminated reactions consist of all of the reactions generating or consuming species S_l: $\{R_i\} \cup \{R_j\}$, where $i = i_1, i_2, \ldots, i_{m1}$ and $j = j_1, j_2, \ldots, j_{m2}$. There are no reduced reactions.

For the step size calculation, the sQSSA removes the non-critical step size corresponding to species S_l from the leap condition. It also removes any critical reactions that consume species S_l from the critical reaction set. However, it must keep the critical reactions that generate species S_l in the critical reaction set due to negativity concerns. (As we will show below, the numbers of firings of reactions that generate species S_l will be sampled in the same way as in the original tau-leaping algorithm.)

In the module that calculates the numbers of reaction firings, the sQSSA first calculates the numbers of firings of reactions generating species S_l in the same fashion as in the original tau-leaping algorithm. For a non-critical reaction $R_{j'} \in \{R_i\}$, where $i = i_1, i_2, \ldots, i_{m1}$, the number of firings $k_{j'}$ of $R_{j'}$ is a sample of the Poisson random variable with mean $a_{j'}(\mathbf{x})\tau^r$. The critical reactions are included in the critical reaction firing calculation along with unreduced critical reactions. The total number of firings of reactions generating species S_l is recorded as N_{total}.

The sQSSA calculates the numbers of firings of reactions consuming species S_l by sampling the multinomial random variable. The total number is set to be N_{total}. The probability for each reaction $R_{j''} \in \{R_j\}$, where $j = j_1, j_2, \ldots, j_{m2}$ to fire is set to be proportional to its propensity $a_{j''}$. Note that it is possible that species S_l has a zero population, which results in $a_{j''} = 0$ for all the reactions consuming species S_l. In that case, we use a dummy propensity with $x_l = 1$ to calculate the probabilities.

B. Stochastic M-M

The stochastic M-M[11,12,14,15,18,29] substitutes all the reactions from an enzyme-substrate subsystem (5) with reduced reaction (6) in the step size calculation, and calculates the numbers of firings of eliminated reactions based on the number of firings of the reduced reaction. We have included $c_2 \gg c_3$ along with the identification criteria given in Ref. 21 as the criteria under which to apply the stochastic M-M approximation to ensure accuracy.[18,19] The same criteria are used at each step to determine whether stochastic M-M remains valid.

For the step size calculation, the stochastic M-M removes the terms corresponding to the original reactions R_1, R_2, and R_3 in (5) from both the non-critical and critical step size calculations. It then includes reduced reaction (6) in the step size calculation for both species S and P.

In the module that calculates the numbers of reaction firings, the stochastic M-M first calculates the number of firings j_{red} of reduced reaction (6) in the same way as an unreduced reaction, most likely a Poisson random sample as a non-critical reaction, because it is highly unlikely that species S is critical and stochastic M-M can still be applied. To improve accuracy, the reaction rate we use to calculate the number of firings of the reduced reaction is the tQSSA rate,[28,29]

$$c_{red} = \frac{2c_3 E_T S}{E_T K_m S + \sqrt{(E_T + K_m + S)^2 - 4E_T S}}, \tag{7}$$

where $E_T = E + C$, $K_m = (c_2 + c_3)/c_1$. Then, the population of E at the end of the step is calculated using a random sample of a binomial random variable with parameters $n = x_E + x_C$ and $p = c_1 x_S/(c_1 x_S + c_2 + c_3)$. The number of firings of reaction R_2 will be calculated by sampling a Poisson random variable with mean $a_2(\mathbf{x})\tau^r$. The number of firings of reaction R_3 equal the number of firings of the reduced reaction: $k_3 = k_{red}$. The number of firings of reaction R_1 can then be calculated from the change Δx_E of the population of E and the number of firings of reactions R_2 and R_3: $k_1 = k_2 + k_3 - \Delta x_E$.

V. NUMERICAL RESULTS

In this section, we present several examples that demonstrate the efficiency and accuracy of automatic model reduction for tau-leaping. The tau-leaping algorithm with automatic model reduction is available in StochKit2.0[30] (http://sourceforge.net/projects/stochkit/). The simulations were performed on an Intel i7-2600 Linux workstation with 8 GB RAM.

A. Stiff decaying dimerization model

We first simulated a decaying dimerization model with the adaptive algorithm. This model was proposed in Ref. 5 and used in Ref. 13. The model consists of three species, S_1, S_2, and S_3, and four reactions,

$$\begin{aligned} S_1 &\xrightarrow{c_1} \varnothing, \\ S_1 + S_1 &\underset{c_3}{\overset{c_2}{\rightleftharpoons}} S_2, \\ S_2 &\xrightarrow{c_4} S_3. \end{aligned} \tag{8}$$

In our experiment, the reaction rate constants were set to $c_1 = 0.1, c_2 = 10.0, c_3 = 5.0, c_4 = 0.01$. The initial populations were given by $x_1(0) = 100$, $x_2(0) = 10\,000$, $x_3(0) = 0$.

In the numerical experiment, species S_1 reached stochastic quasi-steady-state very rapidly and the sQSSA module was automatically deployed to apply sQSSA to species S_1. We simulated from $t = 0$ to $t = 100$ for a total of 1000 trajectories. The accuracy control parameter ε was set to 0.03 in both tau-leaping simulations. The original tau-leaping simulation of the model took about 198 s and the tau-leaping simulation with automatic model reduction took only about 0.65 s. The automatic model reduction module resulted in a ~300 times speed-up. We compared the histograms of $x_3(100)$ from both simulations to the result of a SSA simulation under the same conditions using the Euclidean distance and Manhattan distance, which are, respectively, the L^2 norm and L^1 norm of

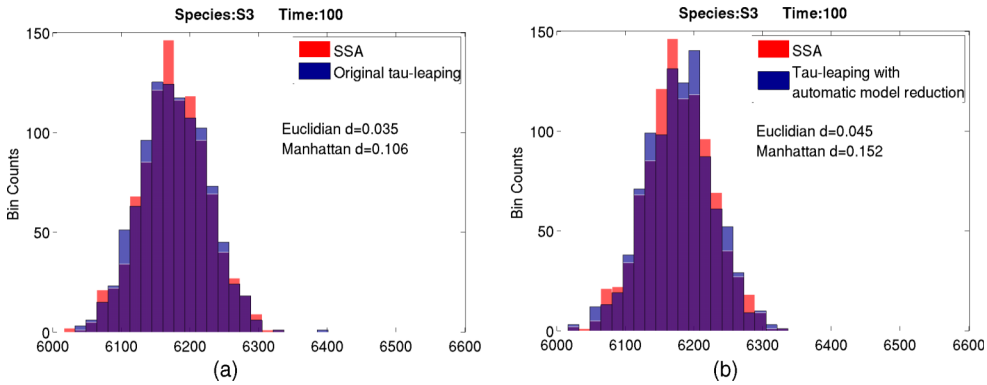

FIG. 2. Comparison of histograms (1000 samples) from tau-leaping simulation and SSA simulation of species S_3 of model (8) at $t = 100$. (a) shows the result using the original tau-leaping algorithm. (b) shows the result using the tau-leaping algorithm with automatic model reduction.

the histogram distance:[31] suppose X and Y are two groups of samples with N samples in X and M samples in Y, and all the sample values are bounded in the interval $I = [x_{min}, x_{max})$. Let $L = x_{max} - x_{min}$. Divide the interval I into K subintervals $I_i = [x_{min} + \frac{(i-1)L}{K}, x_{min} + \frac{iL}{K})$. Then, the histogram distance is given by

$$D_K(X,Y) = \sum_{i=1}^{K} \left| \frac{\sum_{j=1}^{N} \chi(x_j, I_i)}{N} - \frac{\sum_{j=1}^{M} \chi(y_j, I_i)}{M} \right|, \quad (9)$$

where the characteristic function $\chi(x, I_i)$ is defined as

$$\chi(x, I_i) = \begin{cases} 1, & \text{if } x \in I_i, \\ 0, & \text{otherwise.} \end{cases} \quad (10)$$

The comparison results are shown in Fig. 2. The histogram distance results show that the automatic model reduction is accurate for this model.

B. Heat-shock response model

We then simulated a Heat Shock Response (HSR) model.[32] Initially, most of the heat shock factor 1 (HSF1) binds to the heat shock protein 90 (Hsp90) to synthesize and fold protein to its native state NatP. Some of the NatP may misfold into protein MisP with the help of the reactive oxygen species (ROS). MisP may further degrade or form aggregates AggP. When the cell is exposed to high temperature or other stress, more and more NatP will misfold into MisP. MisP also

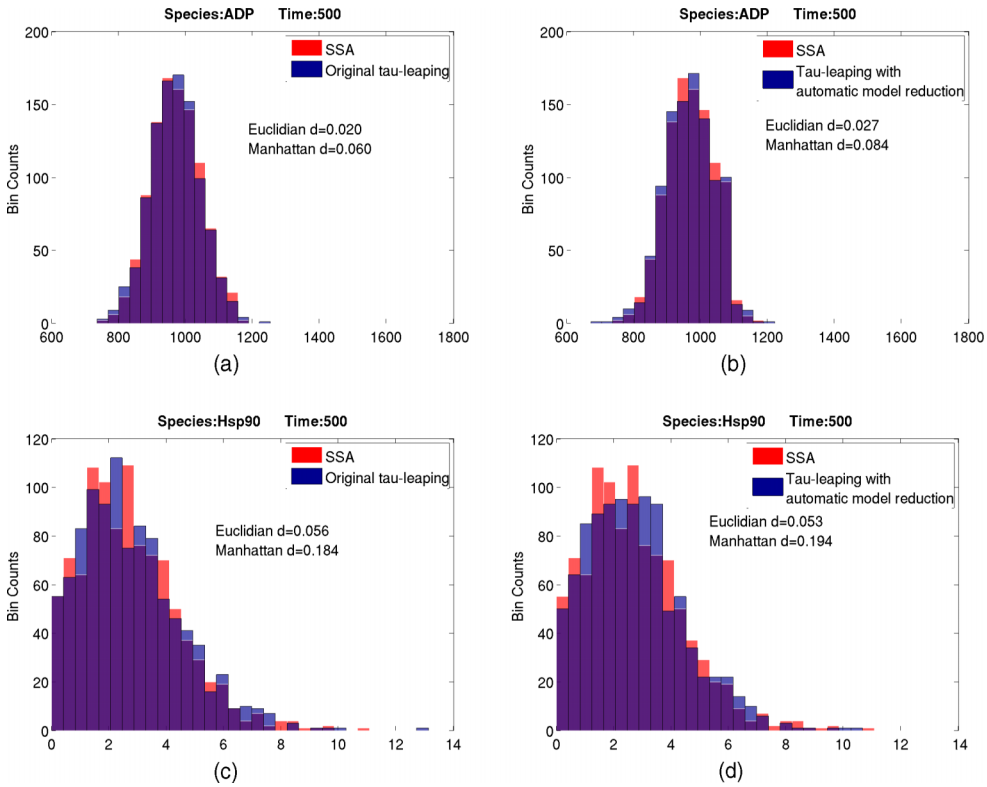

FIG. 3. Comparison of histograms (1000 samples) from tau-leaping simulation and SSA simulation of species adenosine diphosphate (ADP) and Hsp90 of the heat-shock response model at $t = 500$. (a) and (c) show the results using the original tau-leaping algorithm. (b) and (d) show the results using the tau-leaping algorithm with automatic model reduction.

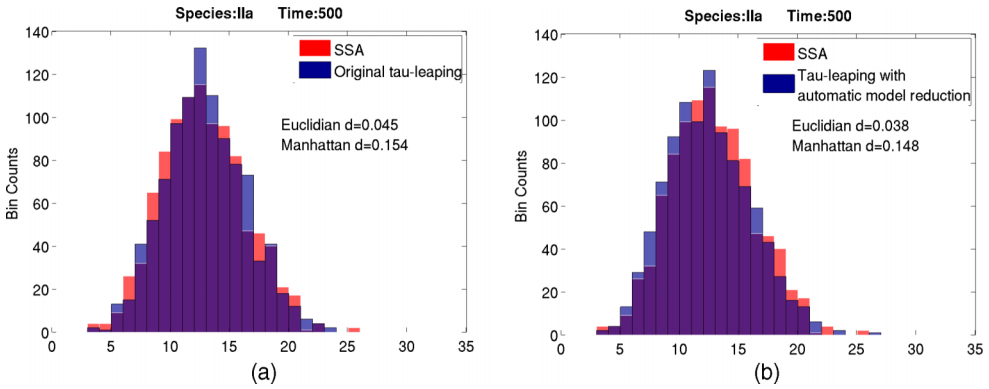

FIG. 4. Comparison of histograms (1000 samples) from tau-leaping simulation and SSA simulation of thrombin (IIa) of the blood coagulation model at $t = 500$. (a) shows the result using the original tau-leaping algorithm. (b) shows the result using the tau-leaping algorithm with automatic model reduction.

binds to Hsp90. An increase in the number of MisP will result in more Hsp90 bound by MisP and not by HSF1. MisP can be refolded with the help of adenosine triphosphate (ATP). In the meantime, HSF1 is free to form dimer DiH and trimer TriH. In addition, TriH can bind to the heat shock element (HSE) to activate the transcription of Hsp90, which results in an increase in the level of Hsp90. More Hsp90 will increase the chance of MisP being correctly refolded. The details of the model for the HSR system can be found in Ref. 32. We use the same stochastic model including initial conditions and rate constants as in Ref. 13, with 14 species participating in 21 chemical reactions.

In the numerical experiment, species MisP and HSF1 were in stochastic quasi-steady-state most of the time. The automatic model reduction module was able to apply sQSSA when necessary. We simulated from $t = 0$ to $t = 500$ for a total of 1000 trajectories. The accuracy control parameter ε was set to 0.03 in both tau-leaping simulations. The original tau-leaping simulation of the model took about 1438 s. The tau-leaping simulation with automatic model reduction took about 274 s. The automatic model reduction module resulted in a ~5 times speed-up. We compared the histograms of two species at $t = 500$ to the results of a SSA simulation under the same conditions. The comparison results are shown in Fig. 3. The histogram distance results show that the automatic model reduction is also accurate for this model.

C. Blood coagulation model

Finally, we tested the automatic model reduction framework on a blood coagulation model.[33] The coagulation model describes the extrinsic blood coagulation system including pro- and anti-coagulants. In this model, one of the most important coagulation factors, thrombin (factor IIa), is formed through a cascade of reactions that are initiated when tissue factor (TF) in the endothelium cells is exposed to blood due to vessel injury. TF activates factor VII. Then, the complex TF_VIIa activates factor X and factor IX. Activated factor Xa in turn activates factor V, factor VIII, and factor IX. The complex factor Xa_Va and the activated Xa activate prothrombin II to thrombin IIa. There are also anti-coagulants such as antithrombin-III (ATIII) and tissue factor pathway inhibitor (TFPI) in the blood that are down-regulating activated pro-coagulants, maintaining the fluidity of the blood. The detailed model can be found in Ref. 33, and involves 34 species and 43 reactions.

As shown in the numerical experiment in Ref. 21, several enzyme-substrate sub-models fit the criteria of stochastic M-M model reduction in different time periods. There were also some species in stochastic quasi-steady-state intermittently. The automatic model reduction module was able to activate stochastic M-M and sQSSA when they were valid and beneficial, and deactivate model reductions when they were no longer valid. We simulated from $t = 0$ to $t = 500$ for a total of 1000 trajectories. The accuracy control parameter ε was set to 0.03 in both tau-leaping simulations. The original tau-leaping simulation took about 1521 s and the tau-leaping simulation with automatic model reduction took about 266 s, resulting in a ~6 times speed-up. We compared the histograms of thrombin (IIa) at $t = 500$ to the result of a SSA simulation under the same conditions to show the accuracy of the automatic model reduction in Fig. 4.

VI. CONCLUSIONS

In this paper, we have developed an adaptive model reduction framework for the explicit tau-leaping algorithm with adaptive step size selection. The framework unifies various model reductions, automatically identifies and activates appropriate model reductions when they are valid and beneficial, and deactivates model reductions when they are no longer valid. It requires no input on the model reduction from users. We showed how to apply sQSSA and stochastic M-M in the framework. We demonstrated in numerical examples the efficiency and accuracy of the automatic model reduction framework. Future work includes implementation of additional model reduction techniques, such as time-dependent solution reduction.[16]

ACKNOWLEDGMENTS

The authors gratefully acknowledge financial support from the U.S. Department of Energy through Grant No. DE-SC0008975, the U.S. Army Research Office through Grant No. W911NF, the Institute for Collaborative Biotechnologies through Grant No. W911NF-09-0001 from the U.S. Army

Research Office, NSF Grant No. DMS-1001012, and NIH Grant No. R01EB014877.

[1]H. McAdams and A. Arkin, Proc. Natl. Acad. Sci. U. S. A. **94**, 814 (1997).
[2]A. Arkin, J. Ross, and H. McAdams, Genetics **149**, 1633 (1998), available online at http://www.genetics.org/content/149/4/1633.
[3]H. McAdams and A. Arkin, Trends Genet. **15**, 65 (1999).
[4]D. T. Gillespie, J. Phys. Chem. **81**, 2340 (1977).
[5]D. T. Gillespie, J. Chem. Phys. **115**, 1716 (2001).
[6]Y. Cao, D. T. Gillespie, and L. R. Petzold, J. Chem. Phys. **124**, 044109 (2006).
[7]Y. Cao, D. T. Gillespie, and L. R. Petzold, J. Chem. Phys. **122**, 014116 (2005).
[8]Y. Cao, D. T. Gillespie, and L. R. Petzold, J. Chem. Phys. **126**, 224101 (2007).
[9]W. E, D. Liu, and E. Vanden-Eijnden, J. Chem. Phys. **123**, 194107 (2005).
[10]W. E, D. Liu, and E. Vanden-Eijnden, J. Comput. Phys. **221**, 158 (2007).
[11]C. Rao and A. P. Arkin, J. Chem. Phys. **118**, 4999 (2003).
[12]E. A. Mastny, E. L. Haseltine, and J. B. Rawlings, J. Chem. Phys. **127**, 094106 (2007).
[13]Y. Pu, L. T. Watson, and Y. Cao, J. Chem. Phys. **134**, 054105 (2011).
[14]D. T. Gillespie, Y. Cao, K. R. Sanft, and L. R. Petzold, J. Chem. Phys. **130**, 064103 (2009).
[15]K. R. Sanft, D. T. Gillespie, and L. R. Petzold, IET Syst. Biol. **5**, 58 (2011).
[16]J. Fu, S. Wu, and L. R. Petzold, J. Comput. Phys. **235**, 446 (2013).
[17]S. Wu, J. Fu, Y. Cao, and L. R. Petzold, J. Chem. Phys. **134**, 134112 (2011).
[18]P. Thomas, A. V. Straube, and R. Grima, J. Chem. Phys. **135**, 181103 (2011).
[19]A. Agarwal, R. Adams, G. C. Castellani, and H. Z. Shouval, J. Chem. Phys. **137**, 044105 (2012).
[20]L. A. Segel and M. Slemrod, SIAM Rev. **31**, 446 (1989).
[21]S. Wu, J. Fu, H. Li, and L. R. Petzold, J. Chem. Phys. **137**, 034106 (2012).
[22]K. R. Sanft, S. Wu, D. T. Gillespie, and L. R. Petzold, "Automatic slow-scale stochastic simulation algorithm with dynamic partitioning and efficient partial equilibrium approximation," J. Comput. Phys. (submitted).
[23]D. T. Gillespie, Annu. Rev. Phys. Chem. **58**, 35 (2007), a review of the stochastic simulation algorithm, chemical master equation, and related topics.
[24]Y. Cao, D. T. Gillespie, and L. R. Petzold, J. Chem. Phys. **123**, 054104 (2005).
[25]T. Tian and K. Burrage, J. Chem. Phys. **121**, 10356 (2004).
[26]A. Chatterjee, D. G. Vlachos, and M. A. Katsoulakis, J. Chem. Phys. **122**, 024112 (2005).
[27]C. A. Yates and K. Burrage, J. Chem. Phys. **134**, 084109 (2011).
[28]M. G. Pedersen, A. M. Bersani, and E. Bersani, Bull. Math. Biol. **69**, 433 (2005).
[29]D. Barik, M. R. Paul, W. T. Baumann, Y. Cao, and J. J. Tyson, Biophys. J. **95**, 3563 (2008).
[30]K. R. Sanft, S. Wu, M. Roh, J. Fu, R. K. Lim, and L. R. Petzold, Bioinformatics **27**, 2457 (2011).
[31]Y. Cao and L. R. Petzold, J. Comput. Phys. **212**, 6 (2006).
[32]C. J. Proctor, C. Sőti, R. J. Boys, C. S. Gillespie, D. P. Shanley, D. J. Wilkinson, and T. B. Kirkwood, Mech. Ageing Dev. **126**, 119 (2005).
[33]M. F. Hockin, K. C. Jones, S. J. Everse, and K. G. Mann, J. Biol. Chem. **277**, 18322 (2002).