Exposition and Streamlined Formulation of Adaptive Explicit-Implicit Tau-Leaping*

Werner Sandmann

January 2009

Abstract

The adaptive explicit-implicit tau-leaping method with automatic tau selection [Y. Cao, D. T. Gillespie and L. R. Petzold, J. Chem. Phys., 126(22):224101 (2007)] is a flexible algorithm for accelerated stochastic simulation of chemically reacting systems. It combines the advantages of different simulation schemes and is particularly useful when a system changes its dynamical behavior over time in the sense that it behaves well in some time periods but possesses stiffness in other time periods. However, the ingredients necessary to fully understand and implement the algorithm are spread over several papers, not always consistent in terminology and notation, which considerably hampers and possibly even prevents accessibility and widespread practical use. We present a streamlined description of the algorithm using a unified terminology and notation and introduce significantly simplified versions of two major ingredients, namely the step size selection and the switching mechanism between the sub-algorithms. This should be greatly helpful for researchers not yet familiar with but interested in tau-leaping methods as well as for practitioners in actually implementing and applying the method.

1 Introduction and Background

Stochastic modeling of chemically reacting systems has become prevalent in physics, chemistry and biology since the stochastic nature of such systems has been more and more often demonstrated [2, 5, 14, 22, 23, 37, 39]. It is evident that the system dynamics are well

*Technical Report, Department of Information Systems and Applied Computer Science, University of Bamberg, Germany

represented by continuous-time Markov chains (CTMCs) governed by the chemical master equation (CME). System analysis is frequently performed by stochastic simulation, either by stochastically exact methods such as Gillespie's [15, 16] stochastic simulation algorithm (SSA) or by approximate accelerated simulation algorithms such as tau-leaping, which exists in diverse versions [1, 8, 9, 11, 12, 18, 19, 20, 26, 30, 29, 38, 40].

The adaptive explicit-implicit tau-leaping method with automatic tau selection [8], which is considered in this paper, combines the SSA with (Poisson) explicit tau-leaping [18], implicit tau-leaping [30], and recent improvements in the step size selection procedure for the tau-leaping methods [7, 9]. The goal is to exploit the advantages and at the same time to overcome major drawbacks of either method. In order to facilitate accessibility, implementation and practical application of the algorithm, we present it in a unified manner. We provide a streamlined formulation and clarify important ingredients of the algorithm. Significantly simplified versions are presented for the step size (tau) selection procedure and the switching mechanism that dynamically chooses either SSA, explicit, or implicit tau-leaping in the course of simulation.

The organization of the paper is as follows. In the remainder of the present section, we provide the mathematical background of stochastic chemical kinetics, the CME, and stochastic simulation. In Section 2, the general principles and an algorithmic framework for multi-step simulation approaches is given, followed by brief descriptions of explicit and implicit tau-leaping. Section 3 addresses the tau selection and contains our unified version, which basically integrates the tau selection for the explicit and the implicit version in such a way that both can be taken as special cases of a more general formula. In Section 4, a concise algorithmic description of the complete adaptive scheme is given, including our simplification of the switching mechanism. Finally, Section 5 concludes the paper discusses topics of further research.

1.1 Stochastic Chemical Kinetics

Consider a well-stirred mixture of $d \in \mathbb{N}_+$ molecular species $S_1, \ldots, S_d$ interacting through $M \in \mathbb{N}_+$ chemical reaction channels $R_1, \ldots, R_M$ in a thermally equilibrated system of fixed volume, where each reaction channel is defined by a corresponding stoichiometric equation

$$R_m: \; s_{m_1} S_{m_1} + \cdots + s_{m_r} S_{m_r} \xrightarrow{c_m} s_{m_{r+1}} S_{m_{r+1}} + \cdots + s_{m_\ell} S_{m_\ell}, \quad r, \ell \in \mathbb{N}_0, \; r \leq \ell \qquad (1)$$

with an associated stochastic rate constant c_m. Mathematically, the stoichiometry is described by the state change vector $v_m = (v_{m1}, \ldots, v_{md})$, where v_{mk} is the change of molecules of species S_k due to R_m. At any time $t \geq 0$ the system state is given by a discrete d-dimensional random vector $X(t) = (X_1(t), \ldots, X_d(t))$, where a discrete random variable $X_k(t)$ describes the number of molecules of species S_k present at time t. The system's state space is the set $\mathcal{X} \subseteq \mathbb{N}_0^d$ of all possible system states and the conditional

transient (time dependent) probability that the system is in state $x \in \mathcal{X}$ at time t, given that the system starts in an initial state $x_0 \in \mathcal{X}$ at time t_0, is denoted by

$$p^{(t)}(x) := p^{(t)}(x|x_0, t_0) = P\left(X(t) = x \mid X(t_0) = x_0\right). \tag{2}$$

The reaction rate for each R_m is given by a state dependent *propensity function* α_m, where $\alpha_m(x)dt$ is the conditional probability that a reaction of type R_m occurs in the time interval $[t, t + dt)$, given that the system is in state x at time t. That is

$$\alpha_m(x)dt = P\left(R_m \text{ occurs in } [t, t + dt) \mid X(t) = x\right). \tag{3}$$

The propensity function is simply given by c_m times the number of possible combinations of the required reactants and thus computes as

$$\alpha_m(x) = c_m \cdot \prod_{j=1}^{m_r} \binom{x_{m_j}}{s_{m_j}}, \tag{4}$$

where x_{m_j} is the number of molecules of species S_{m_j} present in state x, and s_{m_j} is the stoichiometric coefficient of S_{m_j} according to (1). Because at any time the system's future evolution only depends on the current state, $(X(t))_{t\geq 0}$ is a time-homogeneous CTMC. Therefore, given that the system starts in an initial state $x_0 \in \mathcal{X}$ at time t_0, the system dynamics in terms of the state probabilities' time derivatives are described by the CME

$$\frac{\partial p^{(t)}(x)}{\partial t} = \sum_{m=1}^{M} \left(\alpha_m(x - v_m)p^{(t)}(x - v_m) - \alpha_m(x)p^{(t)}(x)\right), \tag{5}$$

well known as the Kolmogorov differential equations in the general theory of Markov processes. Stochastic chemical kinetics via Markov chains governed by the CME can be traced back to the work by Delbrück [13] in the early 1940s and subsequent work by Singer [36] and Bartholomay [3], see, e.g., [4, 24, 39] for historical surveys. It is important to note that the CME provides differential equations for the state probabilities rather than for molecular concentrations as in classical mass action kinetics and that the Markovian approach is a generalization of the traditional deterministic approach rather than an alternative. In fact, the deterministic ODEs based on the law of mass action (LMA) are the special case of the CME in the thermodynamic limit, when the number of molecules and the volume approach infinity but the concentrations remain finite [21, 25]. In general, the accordance of the CME with the theory of thermodynamics has been formally shown by Gillespie [15, 16, 17]. For small numbers of molecules, deterministic models are inappropriate and Markovian modeling is inevitable in order to properly capture the stochastic nature of the system.

1.2 Stochastic Simulation

In principle, the CME may be solved numerically via ODE solvers. However, though a lot of advanced ODE solvers are available and despite a considerable amount of work

spent on adapting Markov chain solution techniques from other domains, there are still many situations where the underlying system is too complex and the CME is intractable by such approaches. Stochastic simulation constitutes an alternative – often the only viable – approach and is in widespread use for more than three decades. Algorithms for stochastically exact trajectory generation of the underlying continuous-time Markov chain were proposed by Bortz et. al. [6] and similarly by Gillespie [15, 16] who formulated it in the context of stochastic chemical kinetics and the CME. It is today famous, often referred to as the stochastic simulation algorithm (SSA) and works as follows:

Init $t := t_0,\ x := x_0$ and t_{end}

while $t < t_{end}$

1. Compute all $\alpha_m(x)$ and $\alpha_{_0}(x) := \alpha_{_1}(x) + \cdots + \alpha_{_M}(x)$;
2. Generate two random numbers u_1, u_2, uniformly distributed on $(0, 1)$;
3. Generate time τ to next reaction, exp. distributed with mean $1/\alpha_0(x)$: $\tau := -\ln(u_1)/\alpha_0(x)$;
4. Determine reaction channel R_m : $m := \min\{k : \alpha_1(x) + \cdots + \alpha_k(x) > u_2\alpha_{_0}(x)\}$;
5. Set $t := t + \tau;\ x := x + v_m$.

Hence, one such simulation run generates a realization of a Markov chain trajectory and repeated runs provide a statistical sample from which properties such as means, variances or higher moments of the numbers of all molecular species at any time and even complete probability distributions, all along with confidence intervals, can be obtained by statistical analysis. The SSA is particularly appealing for its simplicity as well as for the feature of stochastic exactness. Unfortunately, the latter one is also one of the major drawbacks. In particular, since it simulates every single reaction, the SSA is highly inefficient for stiff systems and accelerated stochastic simulation algorithms become desirable. One such approach is tau-leaping which aims at advancing the simulation by larger time steps and performs a corresponding approximate trajectory generation based on combining stochastic simulation with some ODE solver principles.

2 Tau-Leaping

For $m = 1, \ldots, M$ denote by K_m the random variable describing the number of times that a reaction of type R_m occurs in the time interval interval $[t, t + \tau)$. Then

$$X(t + \tau) = X(t) + \sum_{m=1}^{M} v_m K_m. \tag{6}$$

Accordingly, a general algorithmic framework for approximate trajectory generation where the simulation is advanced by pre-defined time steps instead of simulating every single reaction is as follows:

Init $t := t_0,\ x := x_0$ and t_{end};

while $t < t_{end}$

1. Compute all $\alpha_m(x)$ and $\alpha_0(x) := \alpha_1(x) + \cdots + \alpha_M(x)$;
2. Choose a step size τ according to some appropriate rule;
3. Compute suitable estimates $\hat{k}_1, \ldots, \hat{k}_M$ for $K_1, \ldots, K_M$;
4. Set $t := t + \tau$ and update the system state x according to (6).

If $X(t) = x$ and all propensity functions are constant in the time interval $[t, t + \tau)$, the random variable K_m is Poisson distributed with mean $\tau\alpha_m(x)$, that is for $k \in \mathbb{N}_0$:

$$P(K_m = k) = \frac{(\tau\alpha_m(x))^k \exp(-\tau\alpha_m(x))}{k!}, \qquad m = 1, \ldots, M. \tag{7}$$

With tau-leaping, it is assumed that all propensity functions are approximately constant in $[t, t+\tau)$, referred to as the *leap condition* that needs to be formally specified. Handling all propensity functions as if they were indeed constant gives an appropriate rule for Step 2 of the above algorithm. An essential difference between explicit and implicit tau-leaping lies in computing the estimates $\hat{k}_1, \ldots, \hat{k}_M$. Besides, the choice of the step size τ differs.

2.1 Explicit Tau-Leaping

Explicit tau-leaping proceeds by simply computing the estimates $\hat{k}_1, \ldots, \hat{k}_M$ as realizations of the corresponding Poisson random variables. Obviously, (6) then becomes an explicit deterministic expression for $X(t+\tau)$ as a function of x and obeys similarities to the explicit (forward) Euler method for solving systems of deterministic ODEs. More specifically, if the number of molecules of each species is large and the Poisson random variates are approximated by their means, (6) becomes the explicit Euler formula for the deterministic reaction rate equations. However, explicit ODEs solvers become instable for stiff ODE systems, and the same holds for explicit tau-leaping in the case of stiff Markovian systems.

2.2 Implicit Tau-Leaping

Implicit tau-leaping is inspired by the implicit (backward) Euler method which is known to be well suited for stiff ODE systems. Unfortunately, a completely implicit version of tau-leaping would require to generate random variates according to the Poisson distribution

with means $\tau\alpha_m(X(t+\tau)), m = 1, \ldots, M$, which depend on the unknown random state $X(t+\tau)$. Instead, a partially implicit version is considered. Rewriting the random variables K_m as $K_m - \tau\alpha_m(X(t)) + \tau\alpha_m(X(t))$ and evaluating all propensity functions in the last term at $X(t+\tau)$ instead of $X(t)$ yields

$$X(t+\tau) = X(t) + \sum_{m=1}^{M} v_m \Big(K_m - \tau\alpha_m(X(t)) + \tau\alpha_m(X(t+\tau)) \Big). \tag{8}$$

Then, in a first step, all K_m are approximated by computing realizations of Poisson distributed random variables as with explicit tau-leaping. Once these realizations, denoted by $k_1, \ldots, k_M$, have been generated and given $X(t) = x$, (8) becomes an implicit deterministic equation that is solved by, e.g., Newton iteration. Typically, the resulting estimate $\hat{x}(t+\tau)$ for $X(t+\tau)$ is not integer-valued. Therefore, in practice, the estimates to be used for the updating in Step 4 of the above algorithm are obtained by rounding the corresponding term in (8) to the nearest integer. That is

$$\hat{k}_m = \text{round}\Big(k_m - \tau\alpha_m(x) + \tau\alpha_m(\hat{x}(t+\tau)) \Big). \tag{9}$$

It has been empirically demonstrated that implicit tau-leaping significantly speeds up the simulation of *some* stiff systems. As an alternative to (8), motivated by the properties of the trapezoidal rule for solving systems of deterministic ODEs, [9] proposed to substitute (8) by the *trapezoidal tau-leaping formula*

$$X(t+\tau) = X(t) + \sum_{m=1}^{M} v_m \Big(K_m - \frac{\tau}{2}\alpha_m(X(t)) + \frac{\tau}{2}\alpha_m(X(t+\tau)) \Big), \tag{10}$$

which sometimes yields higher accuracy. However, it depends on the specific problem at hand whether (8) or (10) should be preferred.

3 Tau Selection

As the accuracy of both explicit and implicit tau-leaping relies on the leap condition of approximately constant propensity functions in the time interval $[t, t+\tau)$, it is important to select an appropriate step size τ. First of all, the leap condition must be formally specified, that is, it must be mathematically defined what is meant by approximately constant propensity functions. Accordingly, a selection procedure must be developed providing a step size τ preferably much larger than with the SSA and efficiently computed such that the computational overhead is small and does not spoil the potential simulation acceleration compared to the SSA. Besides, it is obviously possible that the updating step taken without care can yield a negative number of molecules because certain critical reactions may exhaust one or more of its reactants. Such a tentative updating step must

be rejected and in order to be efficient such rejections should be rare. In practice, this is handled by keeping the probabilities of critical reactions small. A reaction is taken as critical in state x if it exhausts at least one of its reactants when it occurs n_c times, where n_c is a threshold and constitutes a parameter that has to be specified in the algorithm implementation. Accordingly, we can define the set of indices of critical reactions:

$$\mathcal{C} := \left\{ m \in \{1, \ldots, M\} : \alpha_m(x) > 0 \wedge \min_{i:v_{im}<0} \left\lfloor \frac{x_i}{|v_{im}|} \right\rfloor < n_c \right\}. \tag{11}$$

The tau selection procedure has been successively improved over the years. In early versions the goal was to assure the leap condition by bounding for every reaction the expected change in its propensity function during a time step of size τ by $\epsilon\alpha_0(x)$, where $\epsilon > 0$ is an error control parameter. According to [19] the largest value of τ that satisfies this condition can be obtained by bounding the mean and the standard deviation of the expected change in the propensity function of each reaction by $\epsilon\alpha_0(x)$. It was also recognized that it is more appropriate to bound the change in the propensity function individually for every reaction R_m by $\epsilon\alpha_m(x)$, which corresponds to bounding the relative changes in each propensity function by ϵ. Strictly applied, this implies that τ becomes zero and the simulation does not advance at all if any of the propensity functions evaluated at state x is very small. But, as noted in [7], if α_m changes, then according to Equation (4) it changes by at least c_m such that a change of less than c_m does not make sense. Therefore, the change in α_m can be bounded by the maximum of $\epsilon\alpha_m(x)$ and c_m. Furthermore, [7] presented a procedure that approximately enforces this bound, which is much faster than estimating the mean and the standard deviation according to [19]. Essentially, instead of directly considering propensity functions, the relative changes in populations of certain molecular species are bounded such that the relative changes in the propensity functions will be all approximately bounded by ϵ. The details of the derivation can be found in [7]. With implicit tau-leaping for stiff systems the step size can be often chosen much larger than suggested by [7] when neglecting some of the species [8]. In the following, we present this current state of the art in tau selection for explicit and implicit tau-leaping in an integrated manner and further streamline some ingredients, which results in a compact form that unifies and significantly simplifies the formulae in [7, 8, 9].

In either case, it suffices to consider reactant species. Denote by $\mathcal{R}$ the set of indices of all reactant species and define for all $i \in \mathcal{R}$ and an arbitrary set $\mathcal{M}$ of indices of reactions

$$\hat{\mu}_{i,\mathcal{M}}(x) := \sum_{m\in\mathcal{M}} v_{mi}\alpha_m(x), \qquad \hat{\sigma}^2_{i,\mathcal{M}}(x) := \sum_{m\in\mathcal{M}} v^2_{mi}\alpha_m(x). \tag{12}$$

Then a step size dependent on $\mathcal{M}$ can be expressed by

$$\tau_{\mathcal{M}} = \min_{i\in\mathcal{R}} \left(\frac{\max(\epsilon x_i/g_i(x), 1)}{|\hat{\mu}_{i,\mathcal{M}}(x)|}, \frac{\max(\epsilon x_i/g_i(x), 1)^2}{|\hat{\sigma}^2_{i,\mathcal{M}}(x)|} \right) \tag{13}$$

where g_i is a function defined in order to guarantee that bounding the relative change of states is sufficient for bounding the relative change of propensity functions. This function g_i appears rather complicated in [7] as it involves a couple of case differentiations. We have reasonably simplified it and present it in a quite artless closed form. Denote by $h(i)$ the highest order of reactions in which species S_i appears as reactant and by $n(i)$ the maximum number of S_i molecules required by any of the highest-order reactions. Then

$$g_i(x) = h(i) + \frac{h(i)}{n(i)} \sum_{j=1}^{n(i)-1} \frac{j}{x_i - j}. \tag{14}$$

The step size depends on $\mathcal{M}$ only through $\hat{\mu}_{i,\mathcal{M}}(x)$ and $\hat{\sigma}^2_{i,\mathcal{M}}(x)$, and the only difference in the step size selection for explicit and implicit tau-leaping is in the choice of $\mathcal{M}$ which defines exactly those reactions that are considered in the step size selection.

For explicit tau-leaping, these are simply the non-critical reactions. Hence,

$$\tau^{(expl)} = \tau_{\{1,\ldots,M\}\setminus\mathcal{C}}. \tag{15}$$

For implicit tau-leaping, the principle of the partial equilibrium assumption (see, e.g., [27, 28, 32, 33, 34, 35] for detailed descriptions) is exploited. As it is difficult to identify all reactions that are in partial equilibrium, only reversible reactions are checked for partial equilibrium, which means that their propensity functions evaluated at state x must be approximately equal. More specifically, [8] assume that two reactions R_{m_1}, R_{m_2} where R_{m_1} reverses R_{m_2} and vice versa are in partial equilibrium if the difference of their propensity functions, both evaluated at state x, is less than δ times the minimum of the propensity functions, where δ is a parameter to be specified in the algorithm implementation:

$$|\alpha_{m_1}(x) - \alpha_{m_2}(x)| \leq \delta \min(\alpha_{m_1}(x), \alpha_{m_2}(x)), \qquad \delta > 0. \tag{16}$$

Then in the step size selection for implicit tau-leaping only reactions that are neither critical nor in partial equilibrium (for reversible reactions) are considered. Hence, denoting by $\mathcal{E}$ the set of indices of (reversible) partially equilibrated reactions, we have

$$\tau^{(impl)} = \tau_{\{1,\ldots,M\}\setminus\mathcal{C}\setminus\mathcal{E}}. \tag{17}$$

4 Adaptive Tau-Leaping

In practice, when a system has to be simulated it is not always clear in advance whether or not it possesses stiffness. Quite often, this even changes along the trajectory, i.e. in some time periods the system "behaves well" but in other periods it appears highly stiff. In some periods it might be even the case that neither explicit nor implicit tau-leaping provide an acceleration compared to SSA if the leap condition implies a step size τ that is

so small that the computational overhead eliminates the speed-up. Then, of course, SSA should be executed. This reasoning motivates an adaptive stochastic simulation algorithm that switches between the methods depending on which of them is most efficient in the current region.

In state x, the expected time to the next reaction is $1/\alpha_0(x)$. Consequently, if a candidate step size is less than $n_a/\alpha_0(x)$ it is considered inefficient and n_b single reactions are simulated according to the standard direct method, where both n_a and n_b are parameters to be specified. With adaptive tau-leaping, at each updating step during the simulation either explicit or implicit tau-leaping is chosen dynamically. Hence, a decision rule is necessary. Adaptive tau-leaping applies the simple rule that the system is considered to be stiff if the tentative step size for explicit tau-leaping is more than n_d times smaller than the tentative step size for implicit tau-leaping, which introduces another parameter to be specified.

Now, we have all ingredients we need for formulating an algorithm that, given state x at time t and step size selection procedures for explicit and for implicit tau-leaping, dynamically chooses one of the two methods with an appropriate step size and resorts to SSA if both step sizes are too small to provide simulation acceleration. Our formulation here streamlines that in Section 4 of [8] and is much more concise but equivalent to it.

1. Define the set $\mathcal{C}$ of indices of critical reactions according to Equation (11);
2. Compute candidate step sizes $\tau^{(expl)}, \tau^{(impl)}$ for explicit and implicit tau-leaping;
3. If $\tau^{(expl)} < n_a/\alpha_0(x) \wedge \tau^{(impl)} < n_a/\alpha_0(x)$
   then simulate n_b single reactions, update t and x, and goto 1;
4. Compute candidate step size $\tilde{\tau}$ as expected time to next critical reaction:
$$\text{Generate } \tilde{\tau} \sim \text{Exponential}\left(\sum_{m\in\mathcal{C}} \alpha_m(x)\right);$$
5. If $\tau^{(expl)} > \min(\tau^{(impl)}/n_d; \tilde{\tau})$
   then use explicit tau-leaping with $\tau := \min(\tau^{(expl)}; \tilde{\tau})$;
   else use implicit tau-leaping with $\tau := \min(\tau^{(impl)}; \tilde{\tau})$;
6. If $x + \sum_m \hat{k}_m v_m$ has negative components
   then reduce $\tau^{(expl)}$ and $\tau^{(impl)}$, and goto 3.

Note that the last step is required because there is still a positive probability of generating negative population sizes though this probability should be small for appropriately chosen parameters. Hence, altogether the step size selection procedure can be interpreted as an acceptance-rejection method. The inventors more specifically reduce $\tau^{(expl)}$ and $\tau^{(impl)}$ by

half, but this seems rather arbitrary and may be subject to changes. In any case, the reduction factor is one more parameter to be specified in the algorithm implementation.

5 Conclusion and Discussion

We have presented a unified description of the adaptive explicit-implicit tau-leaping method with automatic tau selection. Thereby, some inconsistencies in earlier papers have been clarified, and simplified versions of the definition of the function g_i which is central for the tau selection and of the switching mechanism for adaptively choosing either SSA, explicit or implicit tau-leaping have been provided. We believe that this contributes to the accessibility of the method by researchers and practitioners and facilitates practical implementation of the algorithm. Therefore, also some necessary theoretical background has been briefly addressed in order to make the description in some sense self-contained and the algorithm understandable without consulting too much other material.

A couple of issues for further investigation arise quite naturally, most of which are concerned with concrete implementation details or related practical as well as theoretical studies of the algorithm's efficiency and the robustness of corresponding estimators. In the inventing papers it has been empirically demonstrated that tau-leaping can significantly accelerate the simulation of chemically reacting systems. The consistency and stability is formally addressed in [31]. However, in particular the choice of the many parameters calls for systematic studies. First of all, the choice of $\epsilon, \delta, n_a, n_b, n_c, n_d$ is currently rather informal and they are specified heuristically. The inventors state that they take "normally" or "usually" ϵ in the range 0.03 to 0.05, $n_a = 10, n_b = 10$, if the previous step uses implicit tau-leaping and $n_b = 100$ otherwise, $n_c = 10, n_d = 100$ and δ "around" 0.05. This seems to rely mainly on empirical comparisons of tau-leaping with SSA. Obviously, in practically relevant cases of huge models, such a comparison will not be feasible as the SSA does not provide accurate results in reasonable time. Similarly, the statistical accuracy of the estimates was most often only shown by comparison to SSA. The first formal approach given in [10] considers the Kolmogorov distance, the histogram distance and the newly introduced self distance for probability measures. Another issue of interest is how the probability of rejecting a candidate step size depends on the parameters mentioned above. As mentioned before, the reduction factor in the case of negative components of the tentative state vector is proposed as 2 to reduce the candidate time steps by half, but this should be also further investigated. In addition to and in support of these issues, it would be particularly desirable to see the algorithm in action, working on practically relevant models from, e.g., systems biology or neuroscience.

References

[1] D. F. Anderson. Incorporating postleap checks in tau-leaping. *Journal of Chemical Physics*, 128(5):054103, 2008.

[2] A. Arkin, J. Ross, and H. H. McAdams. Stochastic kinetic analysis of developmental pathway bifurcation in phage λ-infected escherichia coli cells. *Genetics*, 149:1633–1648, 1998.

[3] A. F. Bartholomay. *A Stochastic Approach to Chemical Reaction Kinetics*. Phd thesis, Harvard University, 1957.

[4] A. T. Bharucha-Reid. *Elements of the Theory of Markov Processes and Their Applications*. McGraw Hill, 1960.

[5] W. J. Blake, M. Kaern, C. R. Cantor, and J. J. Collins. Noise in eukaryotic gene expression. *Nature*, 422:633–637, 2003.

[6] A. B. Bortz, M. H. Kalos, and J. L. Lebowitz. A new algorithm for Monte Carlo simulation of Ising spin systems. *Journal of Computational Physics*, 17:10–18, 1975.

[7] Y. Cao, D. T. Gillespie, and L. R. Petzold. Efficient stepsize selection for the tau-leaping simulation. *Journal of Chemical Physics*, 124:044109–144109–11, 2006.

[8] Y. Cao, D. T. Gillespie, and L. R. Petzold. The adaptive explicit-implicit tau-leaping method with automatic tau selection. *Journal of Chemical Physics*, 126(22):224101–224101–9, 2007.

[9] Y. Cao and L. R. Petzold. Trapezoidal tau-leaping formula for the stochastic simulation of biochemical systems. In *Proceedings of the 1st Conference on Foundations of Systems Biology in Engineering*, pages 149–152, 2005.

[10] Y. Cao and L. R. Petzold. Accuracy limitations and the measurement of errors in the stochastic simulation of chemically reacting systems. *Journal of Computational Physics*, 212(1):6–24, 2006.

[11] Y. Cao and L. R. Petzold. Slow-scale tau-leaping method. *Computer Methods in Applied Mechanics and Engineering*, 197(43–44):3472–3479, 2008.

[12] A. Chatterjee, D. G. Vlachos, and M. A. Katsoulakis. Binomial distribution based τ-leap accelerated stochastic simulation. *Journal of Chemical Physics*, 122:024112, 2005.

[13] M. Delbrück. Statistical fluctuations in autocatalytic reactions. *Journal of Chemical Physics*, 8:120–124, 1940.

[14] N. Fedoroff and W. Fontana. Small numbers of big molecules. *Science*, 297:1129–1131, 2002.

[15] D. T. Gillespie. A general method for numerically simulating the time evolution of coupled chemical reactions. *Journal of Computational Physics*, 22:403–434, 1976.

[16] D. T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *Journal of Physical Chemistry*, 71(25):2340–2361, 1977.

[17] D. T. Gillespie. A rigorous derivation of the chemical master equation. *Physica A*, 188:404–425, 1992.

[18] D. T. Gillespie. Approximate accelerated stochastic simulation of chemically reacting systems. *Journal of Chemical Physics*, 115:1716–1732, 2001.

[19] D. T. Gillespie and L. R. Petzold. Improved leap-size selection for accelerated stochastic simulation. *Journal of Chemical Physics*, 119:8229–8234, 2003.

[20] L. A. Harris and P. Clancy. A partitioned leaping approach for multiscale modeling of chemical reaction dynamics. *Journal of Chemical Physics*, 125(14):144107, 2006.

[21] T. G. Kurtz. The relationship between stochastic and deterministic models for chemical reactions. *Journal of Chemical Physics*, 57(7):2976–2978, 1972.

[22] H. H. McAdams and A. Arkin. Stochastic mechanisms in gene expression. *Proceedings of the National Academy of Science USA*, 94:814–819, 1997.

[23] H. H. McAdams and A. Arkin. It's a noisy business! *Trends in Genetics*, 15(2):65–69, 1999.

[24] D. A. McQuarrie. Stochastic approach to chemical kinetics. *Journal of Applied Probability*, 4:413–478, 1967.

[25] I. Oppenheim, K. E. Shuler, and G. H. Weiss. Stochastic and deterministic formulation of chemical rate equations. *Journal of Chemical Physics*, 50(1):460–466, 1969.

[26] M. F. Pettigrew and H. Resat. Multinomial tau-leaping for stochastic kinetic simulations. *Journal of Chemical Physics*, 126(8):084101, 2007.

[27] J. D. Ramshaw. Partial chemical equilibrium in fluid dynamics. *Physics of Fluids*, 23(4):675–680, 1980.

[28] J. D. Ramshaw and L. D. Cloutman. Numerical method for partial equilibrium flow. *Journal of Computational Physics*, 39(2):405–417, 1981.

[29] M. Rathinam and H. El Samad. Reversible-equivalent-monomolecular tau: A leaping method for small number and stiff stochastic chemical systems. *Journal of Computational Physics*, 224:897–923, 2007.

[30] M. Rathinam, L. R. Petzold, Y. Cao, and D. T. Gillespie. Stiffness in stochastic chemically reacting systems: The implicit tau-leaping method. *Journal of Chemical Physics*, 119:12784–12794, 2003.

[31] M. Rathinam, L. R. Petzold, Y. Cao, and D. T. Gillespie. Consistency and stability of tau-leaping schemes for chemical reaction systems. *Multiscale Modeling and Simulation*, 4(3):867–895, 2005.

[32] M. Rein. The partial-equilibrium approximation in reacting flows. *Physics of Fluids A*, 4(5):873–886, 1992.

[33] M. Schauer and R. Heinrich. Analysis of the quasi-steady state approximation for an enzymatic one-substrate reaction. *Journal of Theoretical Biology*, 79:425–442, 1979.

[34] L. A. Segel. On the validity of the steady-state assumption of enzyme kinetics. *Bulletin of Mathematical Biology*, 50:579–593, 1988.

[35] L. A. Segel and M. Slemrod. The quasi-steady state assumption: A case study in perturbation. *SIAM Review*, 31(3):446–477, 1989.

[36] K. Singer. Application of the theory of stochastic processes to the study of irreproducible chemical reactions and nucleation processes. *Journal of the Royal Statistical Society, Series B*, 15(1):92–106, 1953.

[37] R. Srivastava, L. You, J. Summers, and J. Yin. Stochastic vs. deterministic modeling of intracellular viral kinetics. *Journal of Theoretical Biology*, 218:309–321, 2002.

[38] T. Tian and K. Burrage. Binomial leap methods for simulating stochastic chemical kinetics. *Journal of Chemical Physics*, 121:10356–10364, 2004.

[39] N. van Kampen. *Stochastic Processes in Physics and Chemistry*. Elsevier, 1992.

[40] Z. Xu and W. Cai. Unbiased tau-leaping methods for stochastic simulation of chemically reacting systems. *Journal of Chemical Physics*, 128(15):154112, 2008.