Numerical solution of stochastic models of biochemical kinetics*

SILVANA ILIE[a] WAYNE H. ENRIGHT[b] KENNETH R. JACKSON[b]

[a]Department of Mathematics, Ryerson University, Toronto, ON, M5B 2K3, Canada
email: silvana@ryerson.ca
[b]Department of Computer Science, University of Toronto, Toronto, ON, M5S 3G4, Canada
emails: {enright,krj}@cs.toronto.edu

Abstract

Cellular processes are typically viewed as systems of chemical reactions. Often such processes involve some species with low population numbers, and where a traditional deterministic model of classical chemical kinetics fails to accurately capture the dynamics of the system. In this case, stochastic models are needed to account for the random fluctuations observed at the level of a single cell. We survey the stochastic models of well-stirred biochemical systems and discuss important recent advances in the development of numerical methods for simulating them. Finally, we identify some key topics for future research.

Keywords: Chemical Master Equation, Chemical Langevin Equation, Gillespie algorithm, tau-leaping, stiff systems, multi-scale simulations, hybrid methods, stochastic biochemical modeling.

1 Introduction

Modeling and simulation of biochemical systems has become an area of intense research in recent years [34, 38, 39]. Cellular processes are typically represented as systems of chemical reactions. The evolution of these systems has been traditionally modeled by deterministic reaction rate equations. However, at the level of a single cell it is often the case that some key reactants are present in low molecular numbers (eg., only few copies of a gene or of some important regulatory molecules). Therefore a continuous model may no longer be employed [63, 67]. Moreover, such a system behaves stochastically rather than deterministically [5, 6, 15, 16, 18, 46, 47, 49, 50, 59, 61]. The refined model of stochastic chemical kinetics, the Chemical Master Equation, was developed decades ago by Gillespie [22], who also proposed an exact algorithm to simulate it. In spite of this, it was not until recently that this model has established itself as a standard model for a wide variety of biological processes. The Chemical Master Equation has since been the subject of intense research and it has been successfully applied to numerous biochemical systems, even when the well-mixed assumption is not valid (such as in the cell). The first application of Gillespie's algorithm to a biological system is due to McAdams & Arkin. They showed that stochasticity plays a critical role in the lysis/lysogeny decision of the bacteria λ-phage [46]. Samoilov et al. [58] demonstrated that noise can induce bi-stability in an otherwise monostable system.

However, stochastic models are computationally much more challenging than deterministic models. Furthermore, biochemical systems are generally very complex. They involve reactant species with a wide range of molecular numbers and/or reactions with multiple time-scales. Also, the network of interactions between the reactant species can be quite complicated. Gillespie's algorithm becomes prohibitively expensive on these systems. All these challenges have renewed the interest in developing effective numerical methods for a stochastic

*Research partially supported by grants from the Natural Sciences and Engineering Research Council of Canada (NSERC).

model, cable of coping with the level of complexity characteristic to biological systems. An enhanced exact algorithm was proposed by Gibson & Bruck [20]. Despite this improvement, the high computational cost of exact methods on realistic applications motivated the search for more promising approaches. Several approximate algorithms have been developed and this continues to be a very active research area [32]. One strategy to reduce the computational cost is due to Gillespie [26], who proposed the tau-leaping method. More refined algorithms based on various improvements of the tau-leaping methods were explored by Rathinam et al. [54], Cao, Gillespie & Petzold [10], Tian & Burrage [62], Chatterjee et al. [11]. Several theoretical studies of tau-leaping methods were considered, addressing topics such as consistency and stability by Rathinam et al. [55], higher-order methods by Li [42], adaptive time-stepping by Anderson [3].

Moreover, the tau-leaping method has been shown to be a theoretical bridge between a microscopic, stochastic and discrete model of well-stirred biochemical kinetics, the Chemical Master Equation [23] and a macroscopic, stochastic and continuous model, the Chemical Langevin Equation [25]. Langevin type equations, which are stochastic differential equations (SDE), have received considerable attention, not only in systems biology [38], but also in a wide range of practical applications in physics, chemistry and biology [19, 35]. There exist numerous studies on numerical solution of stochastic differential equations. We suggest as reading on this topic the introductory treatment by Higham [31] and the more comprehensive reference by Kloeden & Platen [40], respectively.

Another strategy to reduce the computational cost when simulating more challenging biochemical systems is to use a combination of models and/or a combination of exact and approximate algorithms for them. Hybrid methods were designed and applied for models that span multiple scales in space and/or time. Among them are the methods of Alfonsi et al. [2], Cao et al. [9], Haseltine & Rawlings [27], Hellander & Lötstedt [30], Kiehl et al. [36], Mattheyses & Simmons [45], Puchalka & Kierzek [51], Rao & Arkin [52], Salis & Kaznessis [56], Samant & Vlachos [57], Weinan et al. [66].

In this paper we review the most important advances in the mathematical modeling and simulation of well-stirred biochemical reaction systems. We provide the background information for the most relevant existing mathematical models for biochemical systems. Our aim is to give a detailed description of the some of the main numerical methods developed for them. Since this area of research has grown beyond the scope of any single survey, it was not possible to cover all existing methods. An excellent introductory reference for modeling and simulations of biochemical kinetics is due to Higham [32]. Wilkinson [67] gives an accessible introduction to probability theory and stochastic modeling of biological systems.

The paper is organized as follows. In section 2 we give an overview of the hierarchy of mathematical models for isothermal well-stirred biochemical kinetics, from the most refined, the Chemical Master Equation, to the intermediate model of the Chemical Langevin Equation and to the less accurate model based on reaction rate equations. In section 3, we identify the key issues in approximating the solution of these models and some of the most significant contributions to algorithm development for them. In section 4 we give some numerical results and we conclude with a description of key challenges and opportunities for future research.

2 Stochastic chemical kinetics

Stochastic chemical kinetics of well-stirred systems is most accurately described by the Chemical Master Equation. The theoretical justification of this model was given by Gillespie [23]. Let us consider a well-stirred system which contains N biochemical species $S_1, \cdots, S_N$ involved in M reactions $R_1, \cdots, R_M$. The system is assumed at thermal equilibrium with a constant volume V. The dynamical system is described by the vector of states $\mathbf{X}(t) = (X_1(t), \ldots, X_N(t))^T$, where $X_i(t)$ is the number of S_i molecules at time t. The vector $\mathbf{X}(t)$ is a discrete (jump) Markov process. The aim is to find the state vector $\mathbf{X}(t)$, given that at the initial time, $t = t_0$, the system was in the state $\mathbf{X}(t_0) = \mathbf{x}_0$.

Each reaction R_j produces a change in the system given by the state-change vector $\boldsymbol{\nu}_j \equiv (\nu_{1j}, \ldots, \nu_{Nj})^T$. Here we denoted by ν_{ij} the change in the number of S_i molecules caused by one reaction R_j. The matrix $\boldsymbol{\Sigma} = \{\nu_{ij}\}_{1 \le i \le N, 1 \le j \le M}$ is called the *stoichiometric matrix*. In addition, a reaction R_j can be described by its

propensity function a_j, which is defined by

$$a_j(\mathbf{x})dt = \text{ the probability of a single occurence of } R_j \text{ in the infinitesimal interval } [t, t+dt) \ ,$$

where $\mathbf{X}(t) = \mathbf{x}$.

For a *second-order reaction*,

$$S_\ell + S_k \xrightarrow{c_j} \text{ 'reaction products'} \tag{1}$$

with $\ell \neq k$ the propensity function has the form $a_j(\mathbf{x}) = c_j x_\ell x_k$. Intuitively, this means that the probability that this reaction takes place is proportional to the number of S_ℓ and S_k molecules present.

Similarly, for a *first-order reaction*,

$$S_k \xrightarrow{c_j} \text{ 'reaction products'} \tag{2}$$

the propensity function has the form $a_j(\mathbf{x}) = c_j x_k$. A similar argument is used: the probability that a reaction of this type takes place is proportional to the number of molecules of type S_k.

Finally, for a *dimerization*,

$$S_k + S_k \xrightarrow{c_j} \text{ 'reaction products'} \tag{3}$$

the propensity function has the form $a_j(\mathbf{x}) = c_j x_k(x_k - 1)/2$. This is a consequence of the observation that the probability that this reaction takes place is proportional to the number of ways two molecules of type S_k can be grouped as unordered pairs.

These formulae for the propensity functions are rigourously derived from the theory of molecular physics and kinetic theory [23]. Making further approximating assumptions can lead to the reaction rates in the deterministic chemical kinetics, as will be discussed later.

2.1 Chemical Master Equation

For $t \geq t_0$ we can define the probability that the state vector at time t is $\mathbf{X}(t) = \mathbf{x}$ given that at time t_0 it was $\mathbf{X}(t_0) = \mathbf{x}_0$:

$$P(\mathbf{x}, t \,|\, \mathbf{x}_0, t_0) = \mathrm{Prob}\{\mathbf{X}(t) = \mathbf{x}, \text{ given } \mathbf{X}(t_0) = \mathbf{x}_0\} \ .$$

To compute this probability, we first derive the probability $P(\mathbf{x}, t+dt \,|\, \mathbf{x}_0, t_0)$ where dt is small enough that no more than one elementary reaction occurs in the interval $[t, t+dt)$. The system will be in state $\mathbf{x}$ at time $t+dt$ if one of the following events took place: either the system was in this state at time t and no reaction occurred in $[t, t+dt)$ or, for some $1 \leq j \leq M$, the system was in a state $\mathbf{x} - \boldsymbol{\nu}_j$ at time t and exactly one reaction R_j occurred in the interval $[t, t+dt)$. There are thus $M+1$ such events which lead to the system being in state $\mathbf{x}$ at time $t+dt$, and these events are disjoint and exhaustive. Since the probability that one reaction R_j fires in $[t, t+dt)$ when at time t the system was in state $\mathbf{x} - \boldsymbol{\nu}_j$ is, by definition, $a_j(\mathbf{x} - \boldsymbol{\nu}_j)dt$ and the probability that no reaction occurs in $[t, t+dt)$ when at time t the system was in state $\mathbf{x}$ is $(1 - \sum_{j=1}^{M} a_j(\mathbf{x})dt)$, from the laws of probability, we derive that

$$P(\mathbf{x}, t+dt \,|\, \mathbf{x}_0, t_0) = P(\mathbf{x}, t \,|\, \mathbf{x}_0, t_0) \left(1 - \sum_{j=1}^{M} a_j(\mathbf{x})dt\right) + \sum_{j=1}^{M} P(\mathbf{x} - \boldsymbol{\nu}_j, t \,|\, \mathbf{x}_0, t_0) a_j(\mathbf{x} - \boldsymbol{\nu}_j)dt \ .$$

By rearranging the terms, dividing by dt and taking the limit as $dt \to 0$, we obtain the following equation

$$\frac{d}{dt} P(\mathbf{x}, t \,|\, \mathbf{x}_0, t_0) = \sum_{j=1}^{M} \left(P(\mathbf{x} - \boldsymbol{\nu}_j, t \,|\, \mathbf{x}_0, t_0) a_j(\mathbf{x} - \boldsymbol{\nu}_j) - P(\mathbf{x}, t \,|\, \mathbf{x}_0, t_0) a_j(\mathbf{x}) \right) \ , \tag{4}$$

known as the *Chemical Master Equation*. It is a coupled system of ordinary differential equations with size equal to the number of all possible states of the system subject to the M reaction channels. This *discrete* and

stochastic model is the most refined mathematical model of well-stirred chemical reaction systems. Unfortunately, it is intractable for most realistic applications. For example, in the Arkin et al. [5] model of λ-phage, realistic population numbers lead to approximately 10^{70} possible states in the Chemical Master Equation. Thus other methods for simulation are needed for such problems. However, we should mention that recent progress has been made in dealing directly with the Chemical Master Equation [48, 33].

2.2 Chemical Langevin Equation

Under certain assumptions, the Chemical Master equation can be approximated by a less refined model, the Chemical Langevin Equation, which is more efficient to approximate numerically. Below, we outline the conditions for which this approximation is justified.

Let $\tau > 0$ be some stepsize. For each $1 \leq j \leq M$, we define the number of reactions R_j that fire in an interval of length τ by

$$K_j(\tau|\mathbf{x}, t) = \text{ the number of reactions of type } R_j \text{ that fire in } [t, t+\tau), \text{ given that } \mathbf{X}(t) = \mathbf{x}\,, \tag{5}$$

and we wish to find a good approximation for $K_j(\tau|\mathbf{x}, t)$. Clearly

$$\mathbf{X}(t+\tau) = \mathbf{x} + \sum_{j=1}^{M} \boldsymbol{\nu}_j K_j(\tau|\mathbf{x}, t)\,.$$

Let us assume that there exists a time $\tau > 0$ such that the following *Leap condition* is satisfied: $a_j(\mathbf{X}(t))$ remains almost constant on $[t, t+\tau)$ for all $1 \leq j \leq M$. For $a_j(\mathbf{X}(\cdot))$ constant in this interval, the probability that one reaction R_j occurs in $[t, t+\tau)$ is $a_j(\mathbf{x})\tau$ where $\mathbf{X}(t) = \mathbf{x}$. Consequently, the number of reactions R_j that fire in the interval has a Poisson distribution with parameter $a_j(\mathbf{x})\tau$. That is, $K_j(\tau|\mathbf{x}, t) \approx P_j(a_j(\mathbf{x}), \tau)$. Under the Leap condition assumption one can then approximate

$$\mathbf{X}(t+\tau) = \mathbf{x} + \sum_{j=1}^{M} \boldsymbol{\nu}_j P_j(a_j(\mathbf{x}), \tau)\,. \tag{6}$$

This is called the (explicit) *tau-leaping method* [26]. If $\tau > 0$ is both small enough such that it satisfies the leap condition, but also large enough such that

$$a_j(\mathbf{x})\tau \gg 1 \text{ for all } 1 \leq j \leq M\,,$$

then the Poisson random variable $P_j(a_j(\mathbf{x}), \tau)$, with mean and variance $a_j(\mathbf{x})\tau$, can be approximated by a normal random variable with the same mean and variance

$$P_j(a_j(\mathbf{x}), \tau) \approx a_j(\mathbf{x})\tau + \sqrt{a_j(\mathbf{x})\tau}\, N_j(0,1)\,, \tag{7}$$

where $N_j(0,1)$ with $1 \leq j \leq M$ are statistically independent normal random variables with mean zero and variance one. This approximation holds when all reactant species have sufficiently large population numbers. Substituting the approximation (7) into the tau-leaping approximation (6) we obtain

$$\mathbf{X}(t+\tau) = \mathbf{x} + \sum_{j=1}^{M} \boldsymbol{\nu}_j a_j(\mathbf{x})\tau + \sum_{j=1}^{M} \boldsymbol{\nu}_j \sqrt{a_j(\mathbf{x})}\sqrt{\tau} N_j(0,1) \tag{8}$$

We recognize in (8) the Euler-Maruyama numerical approximation of the following stochastic differential equation

$$d\,\mathbf{X}(t) = \sum_{j=1}^{M} \boldsymbol{\nu}_j\, a_j(\mathbf{X}(t))dt + \sum_{j=1}^{M} \boldsymbol{\nu}_j\, \sqrt{a_j(\mathbf{X}(t))}\, dW_j(t) \tag{9}$$

where W_j for $1 \leq j \leq M$ are independent Wiener processes.

The equation (9) is called the *Chemical Langevin Equation*. The discrete stochastic process is approximated by a continuous stochastic process $\mathbf{X}(t)$ in the new model (9). The conditions under which the Langevin model is valid are those for which the approximation (8) holds: there exists $\tau > 0$ such that (i) each propensity function has a small variation on an interval $[t, t+\tau)$ (the Leap condition) and that (ii) τ is sufficiently large such that each $a_j(\mathbf{x})\tau \gg 1$. As discussed above, these conditions are typically satisfied when molecular numbers of all species are sufficiently large. This Langevin model is a reduction from the Chemical Master Equation, and consists of a system of coupled stochastic differential equations of size equal to the number of reactant species.

2.3 Reaction rate equations

A further reduction of the model of well-stirred chemical kinetics is obtained when very large numbers for each species are present. More precisely, let us consider the *thermodynamic limit*, that is $\mathbf{X}_i(t) \to \infty$ and the volume $V \to \infty$ such that $\mathbf{X}_i(t)/V$ remains constant for all $1 \leq i \leq N$. Hence, the stochastic terms in the Chemical Langevin Equation (9) become much smaller than the deterministic terms, as the former grow as the square root of the system size while the latter grow as the system size. Therefore the stochastic terms can be neglected in a neighborhood of the thermodynamic limit.

In this regime, the mathematical model is typically written in terms of concentrations, rather than in population numbers. We define the vector of concentrations $\mathbf{Z}(t)$ to have components $Z_i(t) = X_i(t)/(VN_A)$ for $1 \leq i \leq N$, where $N_A = 6.02214179 \times 10^{23}\text{mol}^{-1}$ is Avogadro's number and V is the volume. It follows that the concentrations satisfy

$$\frac{d\mathbf{Z}(t)}{dt} = \sum_{j=1}^{M} \boldsymbol{\nu}_j \, \hat{a}_j(\mathbf{Z}(t)) \, . \tag{10}$$

These are the classical *reaction rate equations*. Note that the reaction rates $\hat{a}_j(\mathbf{Z}(t))$ for $1 \leq j \leq M$ in the reaction rate equation correspond to the propensity functions in the Chemical Master Equation.

Indeed, in the case of the second-order reaction (1) the reaction rate is

$$\hat{a}_j(\mathbf{Z}(t)) = k_j Z_\ell(t) Z_k(t) \text{ with } k_j = c_j N_A V \, .$$

The reaction rate of the first-order reaction (2) is

$$\hat{a}_j(\mathbf{Z}(t)) = k_j Z_k(t) \text{ with } k_j = c_j \, .$$

In the case of the dimerization (3), the term $X_k(t)$ in the expression of the propensity function is negligible compared to $X_k^2(t)$, for large population numbers. Therefore the reaction rate in this case is

$$\hat{a}_j(\mathbf{Z}(t)) = k_j Z_k^2(t) \text{ with } k_j = \frac{c_j}{2} N_A V \, .$$

The reaction rate equations (10) are a deterministic and continuous model, consisting of a system of coupled ordinary differential equations of dimension equal to the number of reactive species. It is a simplification of the Chemical Langevin Equation, which is valid in the regime of very large population numbers. An important observation is that, unless all reactions are of first-order, the reaction rate equations are not obtained from averaging over the Chemical Langevin Equation. Indeed, computing the expectation of (9) leads to

$$\frac{dE(\mathbf{X}(t))}{dt} = \sum_{j=1}^{M} \boldsymbol{\nu}_j \, E(a_j(\mathbf{X}(t)))$$

However $E(a_j(\mathbf{X}(t))) \neq a_j(E(\mathbf{X}(t)))$ if the propensity a_j is quadratic, which is the case for dimerization or second-order reactions.

3 Methods

3.1 Exact methods

As mentioned before, solving directly the Chemical Master Equation (4) is intractable for most realistic applications. Fortunately, there exist alternative methods to deal with this difficulty. Instead of solving the Master Equation to generate trajectories of all possible states, one could simulate one correct trajectory at a time. It is possible to simulate such trajectories, by specifying reactions and times of these reactions with the *exact probability distribution*, consistent with the probability distribution associated with the Chemical Master Equation. Then by running many such trajectories, one can recover the correct statistics for the solution of the Master Equation. Using this approach, Gillespie [21, 22] gave two exact stochastic simulation algorithms to solve the Chemical Master Equation, the Direct Method and the First Reaction Method. The algorithms and their derivation are describe below. Following Gillespie [21] we define

$$p(\tau, j|\mathbf{x}, t) \quad = \quad \text{the probability that the next reaction will occur in the interval } [t+\tau, t+\tau+d\tau) \text{ and this reaction will be } R_j, \text{ given that } \mathbf{X}(t) = \mathbf{x} \, .$$

In addition, we denote by

$$P_0(\tau|\mathbf{x}, t) = \text{ the probability than no reaction occurs in } [t, t+\tau) \text{ given that } \mathbf{X}(t) = \mathbf{x} \, . \tag{11}$$

The two events, that no reactions takes place in $[t, t+\tau)$ and that one reaction R_j occurs in $[t+\tau, t+\tau+d\tau)$, are independent. Their joint probability density function is $P_0(\tau|\mathbf{x}, t)$, multiplied by the probability that a reaction R_j occurs over $[t+\tau, t+\tau+d\tau)$, which is $a_j(\mathbf{x})d\tau$. Therefore

$$p(\tau, j|\mathbf{x}, t)d\tau = P_0(\tau|\mathbf{x}, t) \times a_j(\mathbf{x})d\tau \, . \tag{12}$$

So, we need to compute the probability $P_0(\tau|\mathbf{x}, t)$. Following the definition (11), we observe that no reaction occurs in $[t, t+\tau+d\tau)$ if (i) no reaction occurs in $[t, t+\tau)$ and (ii) no reaction occurs in $[t+\tau, t+\tau+d\tau)$. These two events are independent. Their joint probability is the product of the probability of the event (i), which is $P_0(\tau|\mathbf{x}, t)$, and the probability of the event (ii), which is $(1 - \sum_{k=1}^{M} a_k(\mathbf{x})d\tau)$:

$$P_0(\tau + d\tau|\mathbf{x}, t) = P_0(\tau|\mathbf{x}, t) \times \left(1 - \sum_{k=1}^{M} a_k(\mathbf{x})d\tau\right) \, .$$

By taking the limit $d\tau \to 0$ we obtain a differential equation for the desired probability, whose solution is

$$P_0(\tau|\mathbf{x}, t) = \exp(-a_0(\mathbf{x})\tau) \tag{13}$$

where

$$a_0(\mathbf{x}) = \sum_{k=1}^{M} a_k(\mathbf{x}) \, .$$

Equations (12) & (13) lead to

$$p(\tau, j|\mathbf{x}, t) = \exp(-a_0(\mathbf{x})\tau) \times a_j(\mathbf{x}) \, ,$$

or, equivalently,

$$p(\tau, j|\mathbf{x}, t) = (a_0(\mathbf{x}) \exp(-a_0(\mathbf{x})\tau)) \times \frac{a_j(\mathbf{x})}{a_0(\mathbf{x})} \, . \tag{14}$$

This is the joint density function for the time to the next reaction and of the index of the next reaction. It can be viewed as splitting into two density functions for

(i) τ, the *time* to the next reaction, of density $a_0(\mathbf{x}) \exp(-a_0(\mathbf{x})\tau)$,

(ii) j, the *index* of the next reaction, of density $a_j(\mathbf{x})/a_0(\mathbf{x})$.

These results lead to the following Monte Carlo method of Gillespie [21, 22] for generating sample paths with the correct probabilities.

Gillespie's Direct Method This method computes directly the time to the next reaction and the index of this reaction, according to their correct distributions. First, the system is initialized at time $t = t_0$ by $\mathbf{X}(t_0) = \mathbf{x}_0$. Then, it proceeds with the following steps:

(1). Calculate the propensity functions, $a_k(\mathbf{x})$, for $1 \leq k \leq M$, for the current state of the system, $\mathbf{X}(t) = \mathbf{x}$, and the sum of all propensities, $a_0(\mathbf{x}) = \sum_{k=1}^{M} a_k(\mathbf{x})$.

(2). Generate two independent unit-interval uniform random numbers r_1 and r_2.

(3). Calculate the time to the next reaction by $\tau = (1/a_0(\mathbf{x})) \ln(1/r_1)$.

(4). Calculate the index of the next reaction, as the integer j satisfying

$$\sum_{k=1}^{j-1} a_k(\mathbf{x}) < r_2 a_0(\mathbf{x}) \leq \sum_{k=1}^{j} a_k(\mathbf{x}) .$$

(5). Update the state of the system to reflect that a reaction R_j occurred, $\mathbf{X}(t+\tau) = \mathbf{X}(t) + \boldsymbol{\nu}_j$, set $t = t+\tau$, then return to step (1) or stop.

Gillespie's First Reaction Method Gillespie's [21] second exact algorithm computes the possible time τ_k at which each reaction could occur, if no other reaction takes place. Then it finds the index j of the first reaction, that is with the smallest time. After the initialization, at time $t = t_0$ by $\mathbf{X}(t_0) = \mathbf{x}_0$, the algorithm consists of:

(1). Calculate the propensity functions, $a_k(\mathbf{x})$ for $1 \leq k \leq M$, for the current state of the system, $\mathbf{X}(t) = \mathbf{x}$, and the the sum of all propensities, $a_0(\mathbf{x}) = \sum_{k=1}^{M} a_k(\mathbf{x})$.

(2). For each $1 \leq k \leq M$, calculate the time, τ_k, when the reaction R_k occurs, according to the exponential distribution with parameter $a_k(\mathbf{x})$, where $\mathbf{x}$ is the current state vector. Each $\tau_k = (1/a_k(\mathbf{x})) \ln(1/r_k)$, where $r_1, r_2, \ldots, r_M$ are independent unit-interval uniform random numbers.

(3). Calculate j for which $\tau_j = \min_{1 \leq k \leq M}\{\tau_k\}$ and set $\tau = \tau_j$.

(4). Update the state of the system to reflect that a reaction R_j occurred, $\mathbf{X}(t+\tau) = \mathbf{X}(t) + \boldsymbol{\nu}_j$, set $t = t+\tau$, then return to step (1) or stop.

These two algorithms are equivalent, using the same probability distributions for τ and j, derived above. Each of these methods requires (per iteration) a time proportional to M, the total number of reactions. The Direct Method uses two random numbers per iteration. The First Reaction Method uses M random numbers per iteration and hence it is less efficient than the Direct Method. For this reason we subsequently consider only the Direct Method, which is the faster of the two methods, and call it Gillespie's algorithm or the *stochastic simulation algorithm* (SSA).

Gibson-Bruck Algorithm: The Next Reaction Method Gibson & Bruck [20] modified Gillespie's First Reaction Method to give an exact algorithm which requires, on each iteration, a computational time proportional to the logarithm of the number of reactions, $\log M$. The method does so by constructing a dependency graph from the set of reactions and by using an appropriate data structure to store all the propensities a_k and the possible times τ_k. This structure is an indexed priority graph (also known as a heap). The algorithm can also be applied to systems with time-dependent propensity functions.

Algorithm The initialization consists of setting $\mathbf{X}(t_0) = \mathbf{x}_0$ at time $t = t_0$ and generating a dependency graph G. Then it calculates the propensity functions, $a_k(\mathbf{x})$ for $1 \leq k \leq M$, for the given state of the system. For each $1 \leq k \leq M$, it computes the possible time, τ_k, when the reaction R_k occurs, according to the exponential distribution with parameter $a_k(\mathbf{x})$, for the given state vector. It stores τ_k in an indexed priority queue P and follows the steps:

(1). Find, in the indexed priority queue P, the index j of the reaction for which the possible time τ_j is the smallest and set $\tau = \tau_j$.

(2). Update the state of the system to reflect that a reaction R_j occurred, $\mathbf{X}(t + \tau) = \mathbf{X}(t) + \boldsymbol{\nu}_j$.

(3). For each edge (j, k) in the dependency graph G do

(i). Update a_k.

(ii). For $k \neq j$ take $\tau_k = (a_{k,old}/a_{k,new})(\tau_k - t) + t$.

(iii). For $k = j$ generate one unit-interval uniform random number r and compute $\tau_j = (1/a_j(\mathbf{x})) \ln(1/r) + t$

(iv). Update the values τ_k in the indexed priority queue P and set $t = t + \tau$.

(4). Return to step (1).

Further details on this method can be found in [20]. The Gibson-Bruck algorithm has the potential of being more efficient than Gillespie's algorithm for systems with many species and many reactions. Also, this algorithm is an improvement over Gillespie's algorithm when the system is large and is not strongly coupled, that is firing of one reaction does not affect many other reactions. On the other hand, for small systems, the cost to maintain the data structures dominates the simulation, thus the Next Reaction Method loses it's advantage over Gillespie's algorithm.

Any exact method simulates all reactions, one at a time. Since most realistic biochemical systems have some reactions evolving on very fast time scales, the exact methods become computationally very intense on these practical applications. Typically, Monte Carlo simulations require tens of thousands or hundreds of thousands of individual trajectories to get an accurate estimation of the probability distributions. Thus the efficiency of these simulations is very important. As a result, finding a better trade-off between speed and accuracy of the numerical methods for approximating the solution of the Chemical Master Equation is essential.

3.2 Tau-leaping methods

A speed-up over the exact methods could be obtained by employing *approximate* schemes, such as *tau-leaping*. A tau-leaping method advances the system by leaping with an a priori chosen time-step τ, rather than by stepping from one reaction to the next with the correct probability distribution. To be faster than Gillespie's algorithm, tau-leaping should take a larger step-size, to allow for more reactions to fire within this time-step. τ should also be chosen to satisfy the Leap condition, that is each propensity function changes only by a "small" amount over the time interval $[t, t + \tau)$. Several conditions to ensure that each propensity function does not vary significantly have been proposed. Among them, the approach described in [10] is currently widely used. This requires that the relative changes in each propensity is uniformly bounded by a small accuracy parameter, $\varepsilon \ll 1$,

$$|a_j(\mathbf{X}(t + \tau)) - a_j(\mathbf{x})| \leq \max(\varepsilon a_j(\mathbf{x}), c_j) \text{ for each } 1 \leq j \leq M \ . \tag{15}$$

Since $(a_j(\mathbf{X}(t + \tau)) - a_j(\mathbf{x}))$ is a random variable, both its mean and standard deviation should satisfy the condition (15). This leads to a procedure to determine the maximum step-size for the desired accuracy [10].

Explicit tau-leaping The (explicit) tau-leaping method (6) is due to Gillespie [26]. The algorithm consists of initialization of the system at time $t = t_0$ by $\mathbf{X}(t_0) = \mathbf{x}_0$, followed by

(1). For the current state of the system, $\mathbf{x}$ at time t, calculate the propensity functions, $a_k(\mathbf{x})$ with $1 \leq k \leq M$, and the step-size τ which satisfies the Leap condition.

(2). For each $1 \leq j \leq M$, generate the number k_j of reactions R_j that occur in the time interval $[t, t+\tau)$ from the Poisson random variable $\mathcal{P}_j(a_j(\mathbf{x}), \tau)$.

(3). Update the state of the system to reflect that k_j reactions R_j occurred, $1 \leq j \leq M$, $\mathbf{X}(t+\tau) = \mathbf{X}(t) + \sum_{j=1}^{M} k_j \boldsymbol{\nu}_j$. Set $t = t + \tau$, then return to step (1) or stop.

The main difficulty with this method is that biochemical systems are almost always stiff, exhibiting both fast and slow dynamics. As in the case of stiff deterministic systems, explicit schemes become impractical when applied to stiff stochastic systems. They restrict the step-size to the system's fastest mode. To improve the efficiency and the accuracy of the simulations, further tau-leaping schemes more suitable for stiff stochastic systems were investigated.

Implicit tau-leaping To overcome the step-size limitation due to the different time scales, implicit versions of the tau-leaping method (6) were proposed in [54]. The Poisson random variable which appears in the explicit tau-leaping method, $P_j(a_j(\mathbf{x}), \tau)$, has mean $a_j(\mathbf{x})\tau$. If the deterministic term, $a_j(\mathbf{X})\tau$, is evaluated at the end of the step, while the stochastic term of zero mean, $(P_j(a_j(\mathbf{X}), \tau) - a_j(\mathbf{X})\tau)$, is evaluated at the beginning of the step, then we derive the *implicit tau-leaping* method

$$\mathbf{X}(t+\tau) = \mathbf{x} + \sum_{j=1}^{M} \left(\tau a_j(X(t+\tau)) + \mathcal{P}_j(a_j(\mathbf{x}), \tau) - \tau a_j(\mathbf{x})\right) \boldsymbol{\nu}_j \tag{16}$$

As in the deterministic case, this equation is normally solved by a variant of Newton's method to determine $\mathbf{X}(t+\tau)$. In the reaction rate equations regime, of very large population numbers, the explicit tau-leaping method (6) reduces to the explicit Euler's method while the implicit tau-leaping method (16) reduces to the implicit Euler's method.

A drawback of the implicit tau-leaping method is that it damps the noise excessively: the variance in the fast components of $\mathbf{X}(t)$ is reduced when large step-sizes are employed. To reduce this effect, a combination of steps with the implicit tau-leaping and steps with Gillespie's algorithm can be taken in order to recover the correct probability distributions for the fast variables [54].

Other tau-leaping methods In an attempt to improve on the convergence and stability properties of tau-leaping methods, other leaping strategies have been proposed. For example, based on the midpoint Runge-Kutta method for ordinary differential equations, a midpoint version of the tau-leaping method has been introduced in [26]. For a τ which satisfies the Leap condition, the predicted state at the midpoint $(t + \tau/2)$ is given by $\mathbf{x}' = \mathbf{x} + [\frac{1}{2}\tau \sum_{j=1}^{M} a_j(\mathbf{x})\boldsymbol{\nu}_j]$ where $[\cdot]$ is the integer part. Then one generates sample values from the Poisson random variable $P_j(a_j(\mathbf{x}'), \tau)$ for each $1 \leq j \leq M$. The predicted state at time $(t+\tau)$ is

$$\mathbf{X}(t+\tau) = \mathbf{x} + \sum_{j=1}^{M} P_j(a_j(\mathbf{x}'), \tau)\boldsymbol{\nu}_j \tag{17}$$

which gives the *midpoint tau-leaping* method.

Burrage & Tian [8] introduced a class of Poisson-Runge-Kutta methods for simulating chemical reaction systems. These methods are similar to the Runge-Kutta methods for stochastic differential equations driven by Wiener processes [7]. The Chemical Langevin Equations are stochastic differential equations (SDE) obtained

by simplifying the Chemical Master Equation, via the tau-leaping method. The idea that higher-order numerical methods for the Langevin model can be generalized to higher order tau-leaping methods for the Chemical Master Equation seems promising, but it has been observed that the order of the Runge-Kutta methods for the Langevin model was not inherited by the Poisson Runge-Kutta methods for the discrete stochastic model. This is due to the lower order of convergence for the stochastic component.

Consistency and stability of tau-leaping A theory of local errors, valid for small step-size τ, for both the explicit and implicit tau-leaping methods is developed in Rathinam et at. [55]. This theory gives the framework for constructing higher-order tau-leaping schemes. In addition, the error estimates may be used for the design of adaptive time-stepping strategies. While the paper [55] shows first-order consistency of the explicit and implicit tau-leaping methods for general chemical reaction systems, it only proves their 0-stability and convergence of all the moments in the particular case when all propensity functions are linear. The systems with linear propensities consist only of first-order reactions. For them, both the explicit and implicit methods are shown to be of weak order of convergence 1. More precisely, for a general chemical reaction system, the r-th conditional moment of the increment $(\mathbf{X}(t+\tau) - \mathbf{X}(t))$ is

$$E[(\mathbf{X}(t+\tau) - \mathbf{X}(t))^r | \mathbf{X}(t) = \mathbf{x}] = \tau \sum_{j=1}^{M} \boldsymbol{\nu}_j^r \, a_j(\mathbf{x}) + \mathcal{O}(\tau^2) \, .$$

Then for the explicit tau-leaping approximation $\hat{\mathbf{X}}$ the following weak consistency result can be derived: for any multivariate polynomial $g : \mathbb{R}^N \to \mathbb{R}$ and initial state $\mathbf{x}$ there exist $C > 0$ and $\delta > 0$ such that for all $\tau \in [0, \delta]$

$$\left| E\Big(g(\hat{\mathbf{X}}(t+\tau)) - g(\mathbf{X}(t+\tau)) | \hat{\mathbf{X}}(t) = \mathbf{X}(t) = \mathbf{x} \Big) \right| < C\tau^2$$

Li [42] generalized this result by showing that, provided that all propensity functions are locally Lipschitz, the explicit tau-leaping method has strong order of convergence $1/2$ in the L^2-norm and weak order of convergence 1. Specifically, suppose that the tau-leaping scheme is applied on a mesh $0 = t_0 < t_1 < \ldots < t_n = T$ of the interval $[0, T]$. Then there exists $C > 0$ such that for $\tau = \max\limits_{0 \le \ell \le n-1} (t_{\ell+1} - t_\ell)$ the following global results hold

$$\sup_{\ell \le n} E\left(\left| \hat{\mathbf{X}}(t_\ell) - \mathbf{X}(t_\ell) \right|^2 \right) \le C\tau$$

and

$$\left| E(g(\hat{\mathbf{X}}(t_n))) - E(g(\mathbf{X}(t_n))) \right| \le C\tau$$

These findings are consistent with those for Euler's (also known as Euler-Maruyama's) method for SDE: it is a scheme of strong order of convergence $1/2$ and weak order of convergence 1. We refer the reader interested in more details on numerical methods for SDE to [40].

Still, none of the tau-leaping methods presented here has higher-order convergency with respect to τ than the explicit tau-leaping method. This behavior is different in the ordinary differential equation setting, where e.g., the midpoint scheme, has higher order than Euler's scheme. For an error analysis of tau-leaping methods with respect to other parameters we refer the interested reader to [4].

Avoiding negative populations When Poisson tau-leaping methods are applied to chemical reacting systems with some species in small population numbers, a large step-size may lead to negative population numbers. Therefore careful step-size selection strategies should be employed to avoid such physically unrealistic predictions [10]. An alternative approach for tackling this problem was proposed by Tian and Burrage [62] and independently by Chatterjee et al. [11]. They considered a *binomial tau-leaping* method, in which the Poisson random variables are replaced by binomial random variables. Since a binomial random variable has a finite

range of sample values, the parameters in the binomial random variables can be chosen such that no molecular population becomes negative over a step.

If some populations are driven negative due to a large step-size with the Poisson tau-leaping method, then the step is rejected. However, rejection of steps may bias the statistics of the sample paths and so must be handled carefully. Anderson [3] develops a new Poisson tau-leaping procedure which allows for performing of post-leap checks to ensure that an accuracy requirement is satisfied, but without biasing the statistics. It does so by storing the information generated during one leap and using this history information to preserve the correct trajectory. First, the reaction times are represented as the firing times of some independent, unit-rate Poisson processes, P_j with $1 \leq j \leq M$ so that we can write

$$\mathbf{X}(t) = \mathbf{X}(0) + \sum_{j=1}^{M} P_j \left(\int_0^t a_j(\mathbf{X}(s)) ds \right) \boldsymbol{\nu}_j \ ,$$

Then, the author proves that: if $P(t)$ is a Poisson process with intensity λ and $0 \leq s < u < t$, then $(P(u) - P(s))$ conditioned on $P(s)$ has a binomial $(P(t) - P(s), \alpha)$ distribution with $\alpha = (u-s)/(t-s)$. This theoretical result is used to construct an adaptive step-size strategy which allows step rejections while ensuring that the statistics of the sample paths are not skewed.

3.3 Methods for stochastic quasi-steady-state or partial equilibrium approximations

A different approach to dealing with stiffness in stochastic biochemical systems is due to Rao & Arkin [52] and is based on the idea of elimination of the fast components by a quasi-steady-state approximation. In the deterministic setting, a quasi steady-state approximation assumes that, for the time-scale of interest, the instantaneous rates of change for the intermediate species are almost zero. Thus for the deterministic kinetics, the differential equations corresponding to the intermediate species are eliminated, by equating their rate of change to zero. For stochastic kinetics, the species are partitioned into primary species $\mathbf{x}^s$ and intermediate or ephemeral species $\mathbf{x}^f$. Thus the state vector can be written as $\mathbf{x} = (\mathbf{x}^s, \mathbf{x}^f)$

Denote by $P(\mathbf{x}^s, \mathbf{x}^f; t)$ the probability density of the system. This joint probability can be represented in terms of conditional probabilities as

$$P(\mathbf{x}^s, \mathbf{x}^f; t) = P(\mathbf{x}^f | \mathbf{x}^s; t) P(\mathbf{x}^s; t) \tag{18}$$

Moreover, we assume that $\mathbf{x}^f$ conditional to $\mathbf{x}^s$ is Markovian. So, for a fixed $\mathbf{x}^s$, the conditional probability distribution $P(\mathbf{x}^f | \mathbf{x}^s; t)$ approximately satisfies a (dynamic) master equation.

The *quasi-steady-state assumption* in the stochastic kinetics setting assumes that the rate of change of the conditional probability distribution $P(\mathbf{x}^f | \mathbf{x}^s; t)$ is almost zero

$$\frac{d}{dt} P(\mathbf{x}^f | \mathbf{x}^s; t) = 0 \ .$$

Consequently, we derive that

$$P(\mathbf{x}^f | \mathbf{x}^s; t) = P(\mathbf{x}^f | \mathbf{x}^s) \ .$$

Therefore, we approximate the conditional probability distribution $P(\mathbf{x}^f | \mathbf{x}^s)$ by a steady-state master equation. Finally, an approximate Chemical Master Equation can be derived in terms of the primary species $\mathbf{x}^s$ only. One can then apply Gillespie's algorithm to this reduced Chemical Master Equation. Simulating a reduced system will speed up the performance of Gillespie's algorithm, compared to simulating the whole system.

While the quasi-steady-state assumption deals with the state variables, the partial equilibrium assumption deals with the reactions in the system. Nonetheless, the two assumptions are quite similar and sometimes equivalent. The partial equilibrium approximation assumes that the fast reaction are in equilibrium. This

assumption constitute the basis for the *slow-scale stochastic simulation algorithm* (ssSSA) [9]. A detailed description of this method is given below.

First, the set of all reactions is partitioned into the set of slow reactions $\mathbf{R}^s$ (with M_s elements) and the set of fast reactions $\mathbf{R}^f$ (with M_f elements) depending on the magnitude of their propensity functions. Next, the state vector is partitioned into the state vector of slow species, $\mathbf{X}^s(t)$, and that of fast species, $\mathbf{X}^f(t)$, such that $\mathbf{X}(t) = (\mathbf{X}^s(t), \mathbf{X}^f(t))$. The following criteria for partitioning of species is used: fast species are those involved in fast reactions, while the remaining species are slow.

A new virtual fast process is introduced, $\hat{\mathbf{X}}^f(t)$. This process is obtained from $\mathbf{X}^f(t)$ by turning all the slow reactions off. The slow species are set constant, $\mathbf{X}^s(t) = \mathbf{x}_0^s$. We define

$$\hat{P}(\mathbf{x}^f, t|\mathbf{x}_0, t_0) = \mathrm{Prob}\{\hat{\mathbf{X}}^f(t) = \mathbf{x}^f, \text{ given } \mathbf{X}(t_0) = \mathbf{x}_0\}$$

The new process $\hat{\mathbf{X}}^f(t)$ is Markovian and satisfies a Chemical Master Equation with the slow reactions turned off. Two conditions should be satisfied for the ssSSA to apply. (i) The first condition is that the virtual fast process must be stable. This reduces to requiring that $\hat{P}(\mathbf{x}^f, t|\,\mathbf{x}_0, t_0) \to \hat{P}(\mathbf{x}^f, \infty|\,\mathbf{x}_0, t_0)$ as $t \to \infty$. (ii) The second condition is that the relaxation of the virtual fast process $\hat{\mathbf{X}}^f(t)$ to its stationary asymptotic limit occurs much faster than the expected time to the next slow reaction. This entails a separation of the time-scales of the fast and slow reactions. The conditional probability $\hat{P}(\mathbf{x}^f, \infty|\,\mathbf{x}_0, t_0)$ satisfies a steady-state master equation

$$0 = \sum_{j=1}^{M_f} \left(\hat{P}(\mathbf{x}^f - \boldsymbol{\nu}_j^f, \infty|\,\mathbf{x}_0, t_0) a_j^f(\mathbf{x}^f - \boldsymbol{\nu}_j^f, \mathbf{x}_0^s) - \hat{P}(\mathbf{x}^f, \infty|\,\mathbf{x}_0, t_0) a_j^f(\mathbf{x}^f, \mathbf{x}_0^s) \right) . \tag{19}$$

Finally, the slow-scale approximation means that the fast reactions may be ignored and only the slow reactions are simulated. The propensity functions of the slow reactions are approximated by an average with respect to the asymptotic virtual fast process. Thus the propensity of a slow reactions R_j can be approximated on the time-scale of the slow reactions by

$$\hat{a}_j^s(\mathbf{x}^f, \mathbf{x}^s) = \sum_{\mathbf{z}^f} \hat{P}(\mathbf{z}^f, \infty|\mathbf{x}^f, \mathbf{x}^s) a_j^s(\mathbf{z}^f, \mathbf{x}^s) . \tag{20}$$

Now the system dynamics can be simulated by applying Gillespie's algorithm for the slow reactions only, while using the approximate propensities (20). More details can be found in [9].

The slow-scale stochastic simulation algorithm The system is partitioned into fast and slow reactions and it is initialized at time $t = t_0$ by $\mathbf{X}(t_0) = (\mathbf{x}_0^f, x_0^s)$. The virtual fast process is found and the stationary probability $\hat{P}(\mathbf{x}^f, \infty|\,\mathbf{x}_0, t_0)$ is computed according to (19).

(1). At time t, calculate the approximate propensity functions, $\hat{a}_j^s(\mathbf{x}^f, \mathbf{x}^s)$, according to (20) for the current state of the system. Also calculate the sum of all approximate slow propensities, $\hat{a}_0^s(\mathbf{x}^f, \mathbf{x}^s) = \sum_{k=1}^{M_s} \hat{a}_k^s(\mathbf{x}^f, \mathbf{x}^s)$.

(2). Generate two independent unit-interval uniform random numbers r_1 and r_2.

(3). Calculate the time to the next slow reaction by $\tau = (1/\hat{a}_0^s(\mathbf{x}^f, \mathbf{x}^s)) \ln(1/r_1)$.

(4). Calculate the index j of the next slow reaction as the integer j satisfying

$$\sum_{k=1}^{j-1} \hat{a}_k^s(\mathbf{x}^f, \mathbf{x}^s) < r_2 \hat{a}_0^s(\mathbf{x}^f, \mathbf{x}^s) \leq \sum_{k=1}^{j} \hat{a}_k^s(\mathbf{x}^f, \mathbf{x}^s) .$$

(5). Update the state of the system to reflect that a slow reaction R_j occurred, $\mathbf{X}^s(t+\tau) = \mathbf{X}^s(t) + \boldsymbol{\nu}_j^s$ and $\mathbf{X}^f(t+\tau) =$ sample of $\hat{P}(\mathbf{z}^f, \infty | \mathbf{x}^f, \mathbf{x}^s)$. Set $t = t + \tau$. Return to step (1) or stop.

The algorithms described in this section are applicable if there is a separation in time-scales in the dynamics of the system. It is worth noting that important species should be simulated very accurately. If such species are involved in fast reactions, then the above algorithms would not be appropriate.

3.4 Hybrid methods

Hybrid methods were developed in an attempt to speed-up the simulations of systems with multiple scales in the molecular population numbers and widely disparate time-scales. These numerical methods typically deal with a combination of microscopic and macroscopic models. The system is partitioned in subsets: one set consists of species with low molecular numbers which need to be simulated with exact algorithms and other set with species in large molecular numbers, which are simulated with some approximate algorithms.

Discrete stochastic–continuous deterministic models Kiehl et al. [36] partition the system into two regimes: continuous and discrete. The discrete regime is represented with a stochastic model, while the continuous regime is modeled with the reaction rate equations. Examples of reactions modeled with the Master equation include transcription, translation and molecular signaling. However, a theoretically justified criteria for automatically partitioning the system is needed. A similar approach was taken by Takahashi et al. [60], and Vasudeva & Bhalla [64]. The deterministic model suppresses the intrinsic noise in the continuous variables and this may impact on the overall behavior of the system.

Discrete exact–discrete approximate stochastic algorithms Rather than combining different mathematical models of biochemical kinetics, Puchalka & Kierzek [51] developed a hybrid method for simulating the Chemical Master Equation. Their method, called the *maximal time-step algorithm* uses the exact Gibson-Bruck scheme, for the set of slow reactions and the (Poisson) tau-leaping method for the set of fast reactions. The system is advanced with a time-step which is the minimum of the next reaction time and a user selected maximal time-step. In addition, the method dynamically partitions the reactions set into slow and fast reactions. The algorithm, while being capable of accurately capturing the systems' dynamics for small time-steps, may not be practical for large systems with reaction rates varying over multiple time-scales.

Discrete stochastic–continuous stochastic models An improvement over the above hybrid models was originally proposed by Haseltine & Rawlings [27]. The authors partitioned the reactions into the set of slow and the set of fast reactions, based on the magnitude of their propensity functions and recommended at least two orders of magnitude difference between the values of partitioned reaction probabilities. The slow reactions are modeled with the Chemical Master Equation and are simulated using Gillespie's Direct Method. The fast reactions are modeled either with the Langevin equation simulated using the Euler-Maruyama method, or with the reaction rate equations. Since biochemical systems are often stiff, an explicit simulation method such as Euler-Maruyama could become quite expensive.

Starting from the approach of Haseltine & Rawlings, Salis & Kaznessis [56] proposed an improved, dynamic partitioning of the system into slow and fast reactions and a hybrid model which couples the discrete stochastic and the Langevin regimes. The method they investigated, called the *next reaction hybrid*, employed the Next Reaction Method for the discrete model. The authors recognized the importance of employing efficient adaptive, higher-order and possibly implicit methods for solving the SDE [40], but they employed the low-order Euler-Maruyama method for simulating the Langevin model.

In the dynamic partitioning, a reaction R_j is classifies as fast if it satisfies both

$$a_j(t)\Delta t \geq \lambda \geq 1 \tag{21}$$

and

$$X_i(t) > \varepsilon|\nu_{ji}| \text{ with } i = \{\text{reactant or product of the } R_j \text{ reaction}\}\ . \tag{22}$$

The suggested values for the parameters are $\lambda = 10$ and $\varepsilon = 100$.

Next Reaction Hybrid Algorithm We describe below a simplified version of the next reaction hybrid algorithm. First, the system is initialized at time $t = t_0$ by $\mathbf{X}(t) = \mathbf{x}_0$. Next, these steps are followed:

(1). Classify the reactions into fast and slow according to (21) and (22). Calculate the propensities of the fast and slow reactions, $a^f(t)$ and $a^s(t)$.

(2). Numerically approximate the solution of the Langevin equation over $[t, t + \Delta t)$ using only $a^f(t)$, and obtain the path of integration.

(3). Based on $a^s(t)$, decide if a slow reaction occurred during Δt.

(i). If no slow reaction occurred, update $t = t + \Delta t$ and approximate the fast variables, $\mathbf{X}^f(t + \Delta t)$.

(ii). If only one slow reaction, R_j, occurred, find the next time τ_j at which it occurred and update $t = t + \tau_j$. Integrate the continuous variables on the correct paths based on **step** (2). Then set $\mathbf{X} = \mathbf{X}(t + \tau_j) + \boldsymbol{\nu}_j$.

(iii). If more than one slow reaction occurred, reduce Δt and return to step (2).

(4). Return to step (1) or stop.

4 Numerical experiments

Bi-stability The Schlögl model [43, 55] is a remarkable example of a reaction network which exhibits bi-stability. For the deterministic model represented by Figure 1 and the reactions in Table 1 a solution converges to one of the two stable states, and stays in the neighborhood of that solution after a finite time. However, for the stochastic models, a trajectory of the Chemical Master Equation (Figure 2, left) or of the Chemical Langevin Equation (Figure 2, right) may spontaneously switch between the two stable states, due to the intrinsic noise of the system. This spontaneous transition between the two stable states is not possible for the reaction rate equations, motivating the need for stochastic modeling.

The set of reactions for the Schlögl reaction network and their corresponding propensities are presented in Table 1. The stochastic reaction rate parameters we employed, which lead to the bistable behavior, are also given in Table 1. The molecular numbers for the species A and B are kept at constant values, $A = 10^5$ and $B = 2 \times 10^5$. In the stochastic models, the initial condition for the number of molecules of species X is $X(0) = 250$. To obtain the two stable states in the deterministic state, we took the initial condition $X(0) = 248$ for the lower stable state (represented in blue) and $X(0) = 249$ for the upper stable state (represented in red).

The state vector is $\mathbf{X} = (X, A, B)$ and the state-change vectors for reactions R_1, R_2, R_3 and R_4 are, respectively,

$$\boldsymbol{\nu}_1 = \begin{pmatrix} 1 \\ -1 \\ 0 \end{pmatrix},\ \boldsymbol{\nu}_2 = \begin{pmatrix} -1 \\ 1 \\ 0 \end{pmatrix},\ \boldsymbol{\nu}_3 = \begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix},\ \boldsymbol{\nu}_4 = \begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix}\ .$$

Other examples of interesting qualitative behavior include noise-induced bistable systems which are monostable in the deterministic setting [58] or noise-induced oscillations in systems which are otherwise non-oscillatory [65].

Table 1: The Schlögl model.

	Reactions	Propensities	Reaction rates
R_1	$A + 2X \xrightarrow{k_1} 3X$	$a_1(X) = k_1 AX(X-1)/2$	$k_1 = 3 \times 10^{-7}$
R_2	$3X \xrightarrow{k_2} A + 2X$	$a_2(X) = k_2 X(X-1)(X-2)/6$	$k_2 = 10^{-4}$
R_3	$B \xrightarrow{k_3} X$	$a_3(X) = k_3 B$	$k_3 = 10^{-3}$
R_4	$X \xrightarrow{k_3} B$	$a_4(X) = k_4 X$	$k_4 = 3.5$

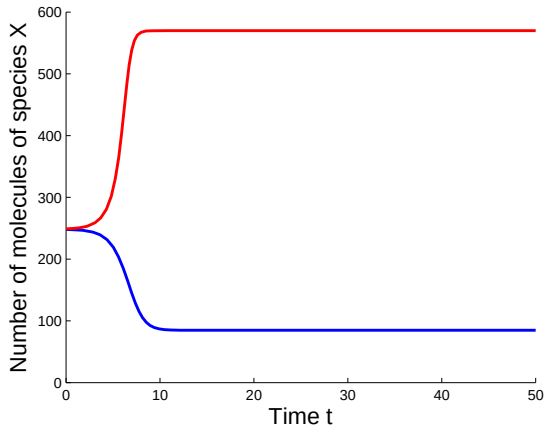

Figure 1: The Schlögl model: Reaction rate equation model

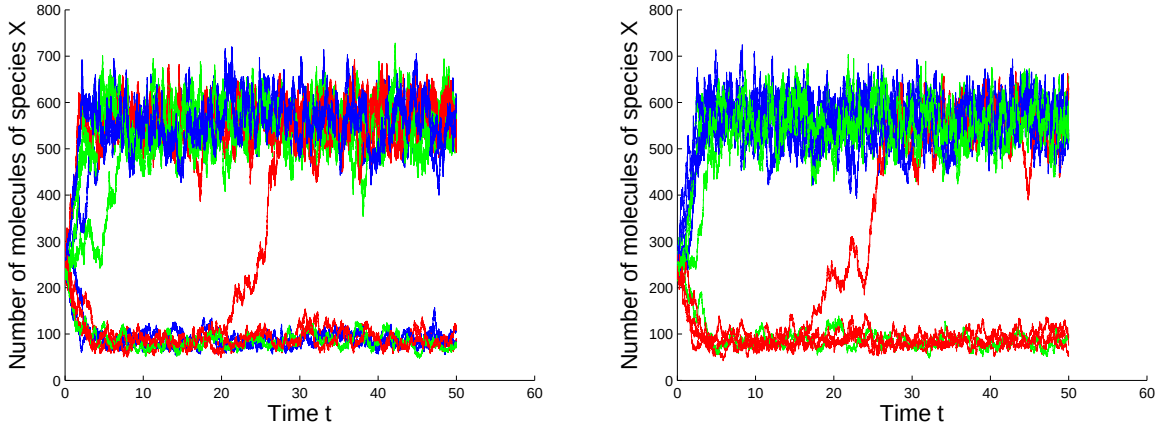

Figure 2: The Schlögl model: Chemical Master Equation model (left), Chemical Langevin Equation model (right) . Only 12 trajectories are shown.

Table 2: The Vilar model.

	Reactions	Propensities	Reaction rates
R_1	$Pa \xrightarrow{k_1} Pa + mRNAa$	$a_1(X) = k_1 Pa$	$k_1 = 50$
R_2	$Pa\_A \xrightarrow{k_2} Pa\_A + mRNAa$	$a_2(X) = k_2 Pa\_A$	$k_2 = 500$
R_3	$Pr \xrightarrow{k_3} Pr + mRNAr$	$a_3(X) = k_3 Pr$	$k_3 = 0.01$
R_4	$Pr\_A \xrightarrow{k_4} Pr\_A + mRNAr$	$a_4(X) = k_4 Pr\_A$	$k_4 = 50$
R_5	$mRNAa \xrightarrow{k_5} mRNAa + A$	$a_5(X) = k_5 mRNAa$	$k_5 = 500$
R_6	$mRNAr \xrightarrow{k_6} mRNAr + R$	$a_6(X) = k_6 mRNAr$	$k_6 = 100$
R_7	$A + R \xrightarrow{k_7} A\_R$	$a_7(X) = k_7 AR$	$k_7 = 20$
R_8	$Pa + A \xrightarrow{k_8} Pa\_A$	$a_8(X) = k_8 Pa\ A$	$k_8 = 1$
R_9	$Pr + A \xrightarrow{k_9} Pr\_A$	$a_9(X) = k_9 Pr\ A$	$k_9 = 1$
R_{10}	$A \xrightarrow{k_{10}} \emptyset$	$a_{10}(X) = k_{10} A$	$k_{10} = 1$
R_{11}	$R \xrightarrow{k_{11}} \emptyset$	$a_{11}(X) = k_{11} R$	$k_{11} = 0.2$
R_{12}	$mRNAa \xrightarrow{k_{12}} \emptyset$	$a_{12}(X) = k_{12} mRNAa$	$k_{12} = 10$
R_{13}	$mRNAr \xrightarrow{k_{13}} \emptyset$	$a_{13}(X) = k_{13} mRNAr$	$k_{13} = 0.5$
R_{14}	$A\_R \xrightarrow{k_{14}} R$	$a_{14}(X) = k_{14} A\_R$	$k_{14} = 10$
R_{15}	$A\_R \xrightarrow{k_{15}} A + R$	$a_{15}(X) = k_{15} A\_R$	$k_{15} = 0$
R_{16}	$Pa\_A \xrightarrow{k_{16}} Pa + A$	$a_{16}(X) = k_{16} Pa\_A$	$k_{16} = 50$
R_{17}	$Pr\_A \xrightarrow{k_{17}} Pr + A$	$a_{17}(X) = k_{17} Pr\_A$	$k_{17} = 100$

Genetic oscillator Vilar et al. [65] proposed a circadian clock model. The biochemical system they described has an interesting feature: while, for some values of the kinetic parameters, the Chemical Master Equation model describes a system with sustained oscillations, the reaction rate equation model predicts no oscillations. The set of reactions for the Vilar model, their propensities and their corresponding stochastic reaction rate parameters are given in Table 2. We used the following initial conditions $Pa(0) = Pr(0) = 1$ and all the other molecular numbers are set to zero [1]. In Figure 3, we show the molecular numbers of species R modeled with the Chemical Master Equation (left) and with the reaction rate equations (right). The oscillations are induced by the intrinsic noise.

5 Challenges

This research area is only at the initial stages and many fundamental open questions remain to be answered by the numerical analysis and scientific computing community.

- **Higher order tau-leaping methods.** Finding higher-order stochastic discrete methods has the potential of improving the speed of computation for practical biological applications, as it should permit larger step-sizes while maintaining the accuracy of the simulation.
- **Adaptive time-stepping schemes.** The effect of adaptive step-size strategies on the convergence of the numerical approximation to the correct solution of stochastic discrete models remains a key question.
- **Hybrid methods.** These methods seem very promising in dealing with the multiple scales which are ubiquitous in biochemical systems. Improved strategies to speed-up the dynamic partitioning of the

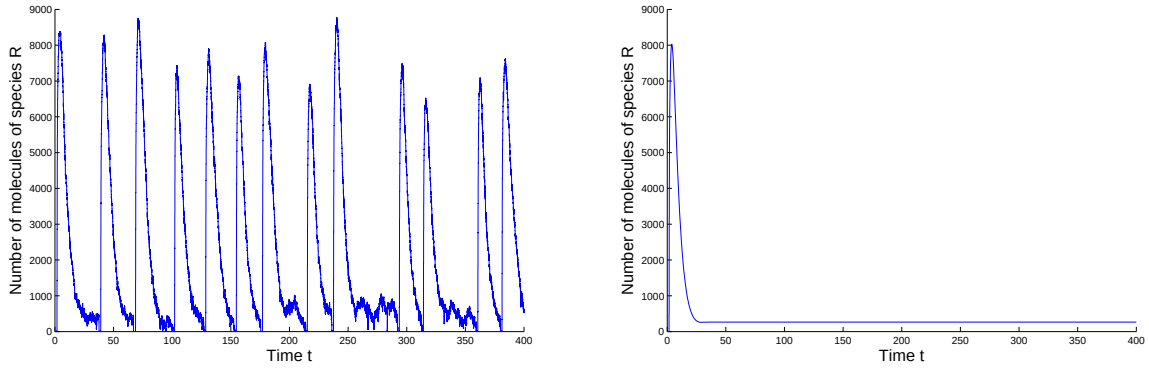

Figure 3: The Vilar et al. [65] model: Chemical Master Equation model (left), reaction rate equation model (right) .

system are needed. Better criteria for partitioning the reaction system is another important issue.

- **Spatially inhomogeneous systems.** The assumption that the biochemical reacting system is homogeneous is not always satisfied. Stochastic models for spatially heterogeneous systems are needed. Efficiency is critical for such models. Existing schemes for molecular-crowding conditions include [14] (where a Monte Carlo method was adapted for the reaction-diffusion Chemical Master Equation), the next volume method, [44] (where a binomial spatial tau-leaping method is developed) and [13] (where a diffusive finite state projection method is introduced).

6 Conclusion

Stochastic modeling and simulation of biological processes are problems of high interest today. The multitude of research opportunities related to the development of effective and reliable simulation tools for these stochastic models as well as for formulating the theoretical foundation to support them, makes this area particularly attractive for numerical analysts. In this paper, we reviewed some of the key achievements in the efficient modeling and simulation of well-stirred biochemical reaction systems and outlined some of the important directions for future research.

References

[1] D. Adalsteinsson, D. McMillen, T.C. Elston, 2004, Biochemical network stochastic simulator (BioNetS): software for stochastic modeling of biochemical networks, *BMC Bioinformatics* **5**(24).

[2] A. Alfonsi, E. Cances, G. Turinici, B. di Ventura, W. Huisinga, 2005, Adaptive simulation of hybrid stochastic and deterministic models for biochemical systems, *ESAIM Proc.* **14**, 1-13.

[3] D.F. Anderson, 2008, Incorporating postleap checks in tau-leaping, *J. Chem. Phys.*, **128**(5): 054103.

[4] D.F. Anderson, A. Ganguly, T.G. Kurtz, 2010, Error analysis of tau-leap simulation methods, *Annals of Applied Probability*, to appear.

[5] A.P. Arkin, J. Ross, H.H. McAdams, 1998, Stochastic kinetic analysis of developmental pathway bifurcation in phage-infected Escherichia coli cells , *Genetics* **149**, 1633.

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

[7] K. Burrage, P.M. Burrage, 2000, Order conditions of stochastic Runge-Kutta methods by B-series, *SIAM J. Num. Anal.* **38**, 1626–1646.

[8] K. Burrage, T. Tian, 2004, Poisson Runge-Kutta methods for chemical reaction systems, *Advances in Scientific Computing and Applications*, Y. Lu et al., eds, Science Press, Beijing/New York, 82–96.

[9] Y. Cao, D.T. Gillespie, L. Petzold, 2005, The slow-scale stochastic simulation algorithm, *J. Chem. Phys.* **122**, 014116.

[10] Y. Cao, D.T. Gillespie, L. Petzold, 2006, Efficient step-size selection for the tau-leaping method, *J. Chem. Phys.* **124**, 044109.

[11] A. Chatterjee, D.G. Vlachos, M.A. Katsoulakis, 2005, Binomial distribution based tau-leap accelerated stochastic simulation, *J. Chem. Phys.* **122**, 024112.

[12] E.J. Crampin, S. Schnell, 2004, New approaches to modelling and analysis of biochemical reactions, pathways and networks, *Progress in Biophysics & Molecular Biology* **86**, 1–4.

[13] B. Drawert, M.J. Lawson, L. Petzold, M. Khammash, 2010, The diffusive finite state projection algorithm for efficient simulation of the stochastic reaction-diffusion master equation, *J. Chem. Phys.***132**, 074101.

[14] J. Elf, M. Ehrenberg, 2004, Spontaneous separation of bi-stable biochemical systems into spatial domains of opposite phases, *Systems Biology, IEE Proceedings* **1**(2), 230–236.

[15] M.B. Elowitz, S. Leibler, 2000, A synthetic oscillatory network of transcriptional regulators, *Nature* **403**, 335–338.

[16] M.B. Elowitz, A.J. Levine, E.D. Siggia, P.S. Swain, 2002, Stochastic gene expression in a single cell, *Science* **297**, 1183–1186.

[17] H. El Samad, M. Khammash, L. Petzold, D.T. Gillespie, 2005, Stochastic modeling of gene regulatory networks, *Int. J. Robust Nonlinear Control* **15**, 691–711.

[18] N. Federoff, W. Fontana, 2002, Small numbers of big molecules, *Science* **297**, 1129–1131.

[19] C.W. Gardiner, 2009, *Stochastic methods: a handbook for the natural and social sciences*, Springer, Berlin.

[20] M.A. Gibson, J. Bruck, 2000, Exact stochastic simulation of chemical systems with many species and many channels, *J. Chem. Phys.* **105**, 1876–1889.

[21] D.T. Gillespie, 1976, A general method for numerically simulating the stochastic time evolution of coupled chemical reactions, *J. Comp. Phys.* **22**, 403–434.

[22] D.T. Gillespie, 1977, Exact stochastic simulation of coupled chemical reactions, *J. Phys. Chem.* **81**, 2340–2361.

[23] D.T. Gillespie, 1992, A rigorous derivation of the chemical master equation, *Physica A*, **188**, 402–425.

[24] D.T. Gillespie, 1992, *Markov processes, an introduction for physical scientists*, Academic Press, New York.

[25] D.T. Gillespie, 2000, The chemical Langevin equations, *J. Phys. Chem.* **113**, 297–306.

[26] D.T. Gillespie, 2001, Approximate accelerated stochastic simulation of chemically reacting systems, *J. Chem. Phys.* **115**, 1716–1733.

[27] E.L. Haseltine, J.B. Rawlings, 2002, Approximate simulation of coupled fast and slow reactions for stochastic chemical kinetics, *J. Chem. Phys.* **117**(15), 6959–6969.

[28] E.L. Haseltine, J.B. Rawlings, 2005, On the origins of approximations for stochastic chemical kinetics, *J. Chem. Phys.* **123**, 164115.

[29] J. Hasty, D. McMillen, F. Isaacs, J.J. Collins, 2001, Computational studies of gene regulatory networks: in numero molecular biology, *Nature Reviews Genetics* **2**, 268–279.

[30] A. Hellander, P. Lötstedt, 2007, Hybrid method for the chemical master equation, *J. Comput. Phys.* **227**(1), 100–122.

[31] D.J. Higham, 2001, An algorithmic introduction to numerical simulation of stochastic differential equations, *SIAM Review* **43**(3), 525–546.

[32] D.J. Higham, 2008, Modeling & simulating chemical reactions, *SIAM Rev.* **50**(2), 347–368.

[33] T. Jahnke, 2010, An adaptive wavelet method for the chemical master equation, *SIAM J. Scient. Comput.* **31**, 4373–4394.

[34] H. de Jong, 2002, Modeling & simulation of genetic regulatory systems, *J. Comput. Biol.* **9**(1), 67–103.

[35] N.G. van Kampen, 2007, *Stochastic processes in physics and chemistry*, Amsterdam, North-Holland.

[36] T.R. Kiehl, R.M. Mattheyses, M.K. Simmons, 2004, Hybrid simulation of cellular behavior, *Bioinformatics* **20**, 316–322.

[37] A.M. Kierzek, 2002, STOCKS: STOChastic Kinetic Simulations of biochemical systems with Gillespie algorithm, *Bioinformatics* **18**(3), 470–481.

[38] H. Kitano, 2002, Systems biology: A brief overview, *Science* **295**, 1662–1664.

[39] H. Kitano, 2002, Computational systems biology, *Nature* **420**, 206–210.

[40] P.E. Kloeden, E. Platen, 1992, *Numerical solution of stochastic differential equations*, Springer-Verlag, Berlin.

[41] V.A. Kuznetsov, G.D. Knott, R.F. Bonner, 2002, General statistics of stochastic process of gene expression in eukaryotic cells, *Genetics* **161**, 1321–1332.

[42] T. Li, 2007, Analysis of explicit tau-leaping schemes for simulating chemically reacting systems, *SIAM Multi. Model. Simul.* **6**, 417-436.

[43] I. Matheson, D.F. Walls, C.W. Gardiner, 1975, Stochastic models of firstorder nonequilibrium phase transitions in chemical reactions, *Journal of Statistical Physics* **12**(1), 21–34.

[44] T. Marquez-Lago, K. Burrage, 2007, Binomial tau-leap spatial stochastic simulation algorithm for applications in chemical kinetics, *J. Chem. Phys.* **127**, 104101.

[45] T. Mattheyses, M. Simmons, 2004, Hybrid simulation of cellular behavior, *Bioinformatics* **20**, 316–322.

[46] H. McAdams, A. Arkin, 1997, Stochastic mechanism in gene expression, *Proc. Natl. Acad. Sci. USA* **94**, 814–819.

[47] H. McAdams, A. Arkin, 1999, It's a noisy business, Genetic regulation at the nanomolar scale, *Trends Genet.* **15**, 65–69.

[48] B. Munsky, M. Khammash, 2006, The finite state projection algorithm for the solution of the chemical master equation, *J. Chem. Phys.* **124**, 044104.

[49] E.M. Ozbudak, M. Thattai, I. Kurtser, A.D. Grossman, A. van Oudenaarden, 2002, Regulation of noise in the expression of a single gene, *Nature Genet.* **31**, 69–73.

[50] J. Paulsson, 2004, Summing up the noise in gene networks, *Nature* **427**, 415–418.

[51] J. Puchalka, A.M. Kierzek, 2004, Bridging the gap between stochastic & deterministic regimes in the kinetic simulations, *Biophys. J.* **86**, 1357–1372.

[52] C.V. Rao, A.P. Arkin, 2003, Stochastic chemical kinetics and the quasi-steady-state assumption: application to the Gillespie algorithm, *J. Chem. Phys.* **118**, 4999–5010.

[53] J.M. Raser, E.K. O'Shea, 2004, Control of stochasticity in eukaryotic gene expression, *Science* **304**, 1811–1814.

[54] M. Rathinam, L. Petzold, Y. Cao, D.T. Gillespie, 2003, Stiffness in stochastic chemically reacting systems: the implicit tau-leaping method, *J. Chem. Phys.* **119**(24) , 12784.

[55] M. Rathinam, L. Petzold, Y. Cao, D.T. Gillespie, 2005, Consistency and stability of tau-leaping schemes for chemical reaction systems, *SIAM Multi. Model. Simul.* **4**, 867-895.

[56] H. Salis, Y. Kaznessis, 2005, Accurate hybrid stochastic simulation of a system of coupled chemical or biochemical reactions, *J. Chem. Phys.* **122**(5), 054103.

[57] A. Samant, D.Vlachos, 2005, Overcoming stiffness in stochastic simulation stemming from partial equilibrium: a multiscale Monte-Carlo algorithm, *J. Chem. Phys.* **123**, 144114.

[58] M. Samoilov, S. Plyasunov, A.P. Arkin, 2004, Stochastic amplification and signaling in enzymatic futile cycles through noise-induced bistability with oscillations, *Proc. Natl. Acad. Sci. USA* **102**(7), 2310–2315.

[59] P.S. Swain, M.B. Elowitz, E.D. Siggia, 2002, Intrinsic and extrinsic contributions to stochasticity in gene expression, *Proc. Natl. Acad. Sci. USA*, **99**, 12795–12800.

[60] K. Takahashi, K. Kaizu, B. Hu, M. Tomita, 2004, A multi-algorithm, multi-timescale method for cell simulation, *Bioinformatics* **20**(4), 538–546.

[61] M. Thattai, A. van Oudenaarden, 2001, Intrinsic noise in gene regulatory networks, *Proc. Natl. Acad. Sci. USA* **98**, 8614–8619.

[62] T. Tian, K. Burrage, 2004, Binomial leap methods for simulating stochastic chemical kinetics, *J. Chem. Phys.*, **121**(1), 10356–10364.

[63] T.E. Turner, S. Schnell, K. Burrage, 2004, Stochastic approaches for modelling in vivo reactions, *Comput. Biol. Chem.* **28**, 165–178.

[64] K. Vasudeva, U.S. Bhalla, 2004, Adaptive stochastic-deterministic chemical kinetic simulations , *Bioinformatics* **20**(1), 78–84.

[65] J.M. Vilar, H.Y. Kueh, N. Barkai, S. Leibler, 2002, Mechanisms of noise-resistance in general oscillations, *Proc. Natl. Acad. Sci. USA* **99**(9), 5988–9592.

[66] E. Weinan, D. Liu, E. Vanden-Eijnden, 2007, Nested stochastic simulation algorithm for chemical kinetic systems with multiple time scales, *J. Chomput. Phys.* **221**, 158–180.

[67] D.J. Wilkinson, 2006, *Stochastic modelling for systems biology*, Chapman & Hall/CRC.