Available online at www.sciencedirect.com
ScienceDirect
Journal of Computational Physics 224 (2007) 897–923
JOURNAL OF COMPUTATIONAL PHYSICS
www.elsevier.com/locate/jcp

Reversible-equivalent-monomolecular tau: A leaping method for "small number and stiff" stochastic chemical systems

Muruhan Rathinam [a,*,1], Hana El Samad [b]

[a] Department of Mathematics and Statistics, University of Maryland, Baltimore County, 1000 Hilltop Circle, Baltimore, MD 21250, United States
[b] Department of Biochemistry and Biophysics, California Institute for Quantitative Biomedical Research, University of California, San Francisco, 1700, 4th Street, San Francisco, CA 94143-2542, United States

Received 27 October 2005; received in revised form 20 October 2006; accepted 24 October 2006
Available online 11 January 2007

Abstract

Leaping methods provide for efficient and approximate time stepping of chemical reaction systems modeled by continuous time discrete state stochastic dynamics. We investigate the application of leaping methods for "small number and stiff" systems, i.e. systems whose dynamics involve different time scales and have some molecular species present in very small numbers, specifically in the range 0 to 10. We propose a new explicit leaping scheme, reversible-equivalent-monomolecular tau (REMM-τ), which shows considerable promise in the simulation of such systems. The REMM-τ scheme is based on the fact that the exact solution of the two prototypical monomolecular reversible reactions $S_1 \leftrightarrow S_2$ and $S \leftrightarrow 0$ as a function of time takes a simple form involving binomial and/or Poisson random variables. The REMM-τ method involves approximating bimolecular reversible reactions by suitable monomolecular reversible reactions as well as considering each reversible pair of reactions in the system to be operating in isolation during the time step τ. We illustrate the use of the REMM-τ method through a number of biologically motivated examples and compare its performance to those of the implicit-τ and trapezoidal implicit-τ algorithms. In most cases considered, REMM-τ appears to perform better than these two methods while having the important advantage of being computationally faster due to the explicit nature of the method. Furthermore when stepsize τ is increased the REMM-τ exhibits a more robust performance than the implicit-τ or the trapezoidal implicit-τ for small number stiff problems.

Keywords: Tau-leaping; Stochastic chemical kinetics; Chemical master equation

1. Introduction

Gene expression and regulation involve cellular events that are stochastic and discrete in nature. The stochasticity of these events, often hidden by population measurements, is increasingly being captured by novel

* Corresponding author. Tel.: +1 410 455 2423; fax: +1 410 455 1066.
E-mail addresses: muruhan@math.umbc.edu (M. Rathinam), helsamad@biochem.ucsf.edu (H. El Samad).
[1] The research of this author was supported by Grants NSF DMS 0309647 and NSF DMS 0610013.

0021-9991/$ - see front matter © 2006 Elsevier Inc. All rights reserved.
doi:10.1016/j.jcp.2006.10.034

experimental techniques that probe cells and organisms at the single cell and sometimes single molecule level [1]. The resolution of such techniques has allowed researchers to identify purely stochastic phenomena, not the least of them being population heterogeneity [2,3]. At the same time, it has been suggested that the accurate stochastic mathematical modeling of cellular dynamics is an important tool that can complement and guide experimentation in identifying the spectrum of biological noise-induced phenomena. The success of such an endeavor, however, crucially hinges on the development of modeling and simulation methods and algorithms that are mathematically rigorous, numerically accurate and computationally tractable for realistic biological systems.

The traditional description of a chemically reacting set of molecules relies on the formulation of the Chemical Master Equation (CME), a differential equation for the time evolution of probabilities [4,5]. The CME corresponds to a continuous time discrete state Markov process model, the sample paths of which can be simulated exactly using a simple Monte carlo procedure known as the stochastic simulation algorithm (SSA) [6,7]. The SSA has been used with great success for the study of a number of cellular networks. However, it has been repeatedly pointed out that the SSA can become prohibitively inefficient when reactions occur so frequently that accounting for every event causes a dramatic slow down in simulation time. This effect is exacerbated by the presence of stiffness, i.e. the coexistence of drastically different time scales for the occurrence of the chemical reactions. In such situations, most of the computational time is dedicated to the accurate tracking of the fast reactions, such as the reversible binding and unbinding of molecules to each other. In most practical circumstances, such binding and dissociation events are of lower importance compared to production or degradation of molecules, typically much less frequent events.

A number of approaches have been proposed to improve the computational efficiency of the SSA. These approaches are roughly divided into two classes: leaping methods and singular perturbation based methods. Both classes involve generating approximate sample paths of the process described by CME. Direct numerical solution of the CME to obtain the time evolution of probabilities is computationally expensive for large systems as the effective number of states grows exponentially with the number of distinct species. However, methods have been proposed to find reduced order models of CME by a finite state projection approach [8].

The leaping methods, more widely known as the tau-leaping methods, were first introduced by Gillespie [9]. Rather than accounting for the time of occurrence of every molecular reaction, leaping techniques proceed by generating an approximation for the number of reactions of each type that occur during a certain time step and then updating the state of the system accordingly. Devising an appropriate approximation for the number of reactions occurring during a time step τ as well as a suitable strategy for the selection of τ has been the subject of active investigation recently and several leaping methods have been proposed in the literature. The explicit, implicit and trapezoidal implicit leaping methods [9–11] all use Poisson random variables. Since Poisson random numbers are unbounded, there is a non-zero probability that either of these three methods will produce a negative state. In order to address this problem, leaping methods based on the Binomial distribution and sequential updating of states have been proposed [12,13]. In [14] a bounding procedure was proposed to deal with the same problem. Implicit or semi-implicit methods such as the implicit tau and the trapezoidal implicit tau lead to non-integer states which are not physically meaningful and rounding was proposed in [10] to remedy this problem. The leaping methods have their counterparts in the deterministic chemical kinetics, namely the numerous time stepping schemes such as Runge–Kutta, backward differencing formula (BDF) etc. that have been developed for the numerical solution of ordinary differential equations (ODEs).

In addition to leaping methods other forms of approximations based on the quasi-steady-state assumption or the partial equilibrium assumption have also been proposed to speed up the SSA [15–19]. These methods are intended for stiff systems that exhibit a clear scale separation between very fast reactions that are stable and very slow reactions. These methods in principle should work well for ''small number and stiff'' systems provided there is a wide separation of scales. Just like leaping methods, these methods have deterministic counterparts as well. The theoretical justification of the quasi-steady-state or the partial equilibrium approximation involves singular perturbation theory [20]. Singular perturbation theory makes use of the existence of a very small parameter $\epsilon > 0$ in the model description and attempts to expand the exact solution in terms of an asymptotic series in powers of ϵ. See [21] for a survey paper that applies singular perturbation theory to stochastic systems. In the case of chemical reaction systems, the small parameter ϵ is the ratio between the time

scales of the slow and the fast reactions. The quasi-steady-state as well as the partial equilibrium approximations roughly correspond to the leading order term in the asymptotic expansion.

A key difference between the singular perturbation inspired approximations mentioned above and the leaping methods is that in the latter, we can control the stepsize parameter τ while in the former we do not have any control over the parameter ϵ, rather it is given to us in the form of system parameters. Thus the accuracy of these methods is not under our control, unless we choose to include the higher order terms from the asymptotic expansion. If one were to learn lessons from the ODE counterparts, the flexibility of the time stepping paradigm is clearly evident from the wide range of software systems that mostly produce a robust solution. Thus the leaping methods have the appeal that they are potentially amenable to a general software implementation in which the time steps τ are chosen to keep local errors within a certain tolerance and maintain stability just as it is done in the case of time stepping of ODEs. Local error analysis of leaping methods may be carried out in a relatively straight forward manner as it is shown in [14] where convergence is also proven. Thus we already have some of the essential tools necessary for adaptive stepsize selection in place. However some key problems remain to be addressed, and it is yet to be decided if the leaping paradigm will prove to be as fruitful as its counterpart did for the ODEs. Some issues that need to be addressed include rounding which becomes necessary for implicit or semi-implicit leap methods, as well as negativity of states mentioned earlier.

While the implicit and semi-implicit methods work reasonably well for stiff problems, when the number of molecules is small, the effects of rounding and bounding become prominent, and our investigations in this paper show that their performance is impaired severely by the rounding procedure. One possible solution is to not implement rounding; this however leads to non-integer states which are difficult to interpret physically. Thus we propose an alternative new method that holds substantial promise in tackling small number and stiff situations. This method is based on the idea of decomposing a system of chemically reacting species into ''motifs'' (or commonly occurring subsystems) and then generating appropriate approximations for the individual subsystems. While all leaping methods in the literature consider every reaction in isolation and seek approximation for the number of firings of this reaction during a simulation time step, our method currently considers reversible pairs of bimolecular or monomolecular reactions as motifs. We present a derivation of exact time solutions for two common types of monomolecular reversible reactions. Since bimolecular reversible reactions are not amenable to a simple analytical solution, we approximate them by suitable monomolecular reversible reactions. We then use the exact solution of these monomolecular reactions to produce the leaping approximation. We call our new method the reversible-equivalent-monomolecular tau, REMM-τ in short. Our method leads to both binomial and Poisson random variables with parameters that differ from those used by other leaping methods. Our method has the great advantage of being explicit and hence computationally less demanding, and yet works well for small number and stiff systems. It is first order consistent by construction and always results in integer states, and therefore avoids the error introduced by rounding present in implicit or semi-implicit methods. We describe two different implementations of the method, parallel and sequential. The parallel updating procedure may produce negative states which we correct by a bounding procedure. The sequential updating is always guaranteed to produce nonnegative states. However the sequential updating may be less desirable in systems where race conditions exist (see Remark 5).

This paper is organized as follows. In Section 2 we review the stochastic chemical model behind the CME as well as review some of the existing leaping methods. In Section 3 we motivate and derive the REMM-τ method. In Section 4, we numerically investigate the behavior of REMM-τ, implicit tau and trapezoidal implicit tau when applied to simple test problems. We particularly focus on stiff behavior in the presence of small number of molecules where bounding and rounding procedures play a critical role. We finally consider a stiff genetic circuit example with small number of molecules and show that REMM-τ performs very well.

2. The stochastic chemical kinetics model

In this section we describe the discrete state, continuous time Markov process model for well-stirred chemical reaction systems, as well as an exact simulation algorithm for this model known as the *Stochastic Simulation Algorithm* (SSA) [6,7]. We also describe the tau-leaping schemes, explicit, implicit and trapezoidal-implicit, which were suggested for efficient and approximate simulation of this model [9,22,10]. Throughout this paper, $\mathbb{Z}_+$ denotes the set of nonnegative integers and $\mathbb{R}_+$ the set of nonnegative real numbers.

The formulation we consider consists of a system of well-stirred chemical reactions with N molecular species. We use the state $X(t) \in \mathbb{Z}_+^N$ to denote the vector whose elements $X_i(t)$ are the number of molecules of the ith species at time t. If there are M elementary chemical reactions R_j $(j = 1, \ldots, M)$ that can occur among these N species, then we associate with each reaction R_j a nonnegative *propensity function* $a_j : \mathbb{Z}_+^N \to \mathbb{R}_+$ defined such that $a_j(X(t))\tau + o(\tau)$ is the probability that reaction R_j will happen in the next small time interval $(t, t+\tau]$, as $\tau \to 0$. The form of the propensities $a_j(x)$ may be derived from fundamental principles under certain assumptions, and the a_j turn out to be polynomials [6]. Furthermore, occurrence of a reaction R_j leads to a change of $v_j \in \mathbb{Z}^N \setminus \{0\}$ for the state X. The vectors v_j are the stoichiometric changes to the reactant species due to a reaction R_j, and are therefore independent of the state and time. Based on these premises, it can be shown that the probability density function for the waiting time τ for the next reaction is given by $a_0(x)e^{-a_0(x)\tau}$, where x is the current state and $a_0(x) = \sum_{j=1}^{M} a_j(x)$. Also, the probability that the next reaction is R_j is given by $\frac{a_j(x)}{a_0(x)}$ and is independent of τ. Knowing these two probability densities for the next reaction time and type, we can simulate the system one reaction event at a time. This method is known as the *Stochastic Simulation Algorithm* (SSA) [7].

2.1. The explicit, implicit and trapezoidal implicit tau-leaping methods

Since the SSA accounts for every reaction that occurs in the system, it can become very computationally expensive, and sometimes prohibitively slow in realistic biological systems whose dynamics evolve at different times scales and involve species whose molecular count can differ by orders of magnitude. Leaping methods, whose main rationale is to leap in simulation time over a number of reactions, have been proposed to speed up SSA. Generally stated, the leaping methods try to accelerate the simulation by asking the question: How many times does each reaction channel fire in each subinterval? If one defines

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

then this question can be reformulated as: What is the best approximation for $K_j(\tau; x, t)$? To answer this question, several approximations have been suggested, including the explicit, implicit and trapezoidal implicit tau-leaping [9,22]. Below, we give a brief account of these methods.

Suppose $X^{(\mathrm{et})}(t) = x$ is the current state (the superfix "et" stands for explicit tau). Then for a time step of $\tau > 0$, the update equation for the state at $t + \tau$ is given by

$$X^{(\mathrm{et})}(t+\tau) = x + \sum_{j=1}^{M} v_j K_j^{(\mathrm{et})}(x, \tau). \tag{2}$$

In this method, $K_j^{(\mathrm{et})}(x, \tau) = \mathscr{P}_j(a_j(x), \tau)$, for $j = 1, \ldots, M$, are independent Poisson random variables with mean and variance $a_j(x)\tau$. The explicit tau-leaping method morphs to the explicit Euler method when the SSA can be approximated by the chemical Langevin equation or even the reaction rate equations.

The approximation in the implicit tau (unrounded version) method, on the other hand, proceeds as follows [10]. Given that $X^{(\mathrm{it})}(t) = x$ is the current state, the state at time $t + \tau$ $(\tau > 0)$, is taken to be

$$X^{(\mathrm{it})}(t+\tau) = x + \sum_{j=1}^{M} v_j(\mathscr{P}_j(a_j(x), \tau) - a_j(x)\tau + a_j(X^{(\mathrm{it})}(t+\tau))\tau). \tag{3}$$

Here, the superfix "it" stands for implicit tau. The estimate for K_j is then given by

$$K_j^{(\mathrm{it})} = (\mathscr{P}_j(a_j(x), \tau) - a_j(x)\tau + a_j(X^{(\mathrm{it})}(t+\tau))\tau), \quad j = 1, \ldots, M. \tag{4}$$

Newton's method is used to solve (3). It has been demonstrated [10] that the implicit tau method allows much larger stepsizes than the explicit tau method, when applied to stiff problems. Convergence proofs for the explicit and implicit tau-leaping methods are given in [14]. However, the unrounded implicit tau has the disadvantage that it leads to state values that are not integers. In order to circumvent this problem, the rounded implicit tau was proposed in [10]. It may be described as follows.

Suppose that at time t we have the state $X^{(\mathrm{itr})}(t) = x$ (the superfix "itr" stands for rounded implicit tau). First, compute the intermediate state value X' according to (3). Thus

$$X' = x + \sum_{j=1}^{M} v_j a_j(X')\tau + \sum_{j=1}^{M} v_j(\mathscr{P}_j(a_j(x),\tau) - a_j(x)\tau). \tag{5}$$

Then approximate the actual number of firings $K_j(x,\tau)$ of reaction channel R_j in the time interval $(t, t+\tau]$ by the *integer-valued* random variable $K_j^{(\mathrm{itr})}(x,\tau)$ defined by

$$K_j^{(\mathrm{itr})}(x,\tau) = [a_j(X')\tau + \mathscr{P}_j(a_j(x),\tau) - a_j(x)\tau]. \tag{6}$$

Here the $\mathscr{P}_j(a_j(x),\tau)$ for $j = 1,\ldots,M$ are the *same numbers* used in Eq. (5), and $[z]$ denotes the nearest nonnegative integer to z.

Finally, take the state at time $t+\tau$ to be

$$X^{(\mathrm{itr})}(t+\tau) = x + \sum_{j=1}^{M} v_j K_j^{(\mathrm{itr})}(x,\tau). \tag{7}$$

If $X^{(\mathrm{itr})}(t) = x$ is an integer vector, then so is $X^{(\mathrm{itr})}(t+\tau)$.

The trapezoidal implicit tau method, has been defined in [11] (also see [23]). The trapezoidal method generates the update equation

$$X^{(tr)}(t+\tau) = x + \sum_{j=1}^{M} v_j\left(\mathscr{P}_j(a_j(x),\tau) - \frac{\tau}{2}a_j(x) + \frac{\tau}{2}a_j(X^{(tr)}(t+\tau))\right) \tag{8}$$

for the unrounded version. The rounded version is implemented as follows. Given current state x first solve for X' as in the unrounded trapezoidal implicit tau:

$$X' = x + \sum_{j=1}^{M} v_j \frac{\tau}{2}a_j(X') + \sum_{j=1}^{M} v_j\left(\mathscr{P}_j(a_j(x),\tau) - \frac{\tau}{2}a_j(x)\right). \tag{9}$$

Then approximate the actual number of firings $K_j(x,\tau)$ of reaction channel R_j in the time interval $(t, t+\tau]$ by the *integer-valued* random variable $K_j^{(\mathrm{trr})}(x,\tau)$ defined by

$$K_j^{(\mathrm{trr})}(x,\tau) = \left[\frac{\tau}{2}a_j(X') + \mathscr{P}_j(a_j(x),\tau) - \frac{\tau}{2}a_j(x)\right]. \tag{10}$$

Here the $\mathscr{P}_j(a_j(x),\tau)$ for $j = 1,\ldots,M$ are the *same numbers* used in Eq. (9), and $[z]$ denotes the nearest nonnegative integer to z.

Finally, take the state at time $t+\tau$ to be

$$X^{(\mathrm{trr})}(t+\tau) = x + \sum_{j=1}^{M} v_j K_j^{(\mathrm{trr})}(x,\tau). \tag{11}$$

We note that the unrounded trapezoidal tau method tends to the well known trapezoidal method for ODEs in the deterministic regime.

2.2. The bounding procedure

Since all of the above methods explicit, (rounded) implicit and (rounded) trapezoidal implicit taus may lead to negative integer states, it was proposed in [14] that a bounding procedure be applied whenever a negative state is encountered. The bounding procedure is described below.

Suppose the state update corresponding to stepsize τ results in $\tilde{x} = \widehat{X}(t+\tau)$ which has negative components. Then execute the following loop to obtain a new state which is nonnegative:

```
while x̃ has negative components
  for i = 1 : N
    while x̃_i < 0
      for j = 1 : M
        if (v_j)_i < 0, K̂_j ← K̂_j − 1 end if
      end for
    end while
  end for
end while
```
(12)

Upon exiting the loop set $\widehat{X}(t+\tau) \leftarrow \tilde{x}$.

Throughout this paper we shall refer to the Algorithm (12) as the *bounding procedure*. We like to remark that the bounding procedure (12) is guaranteed to terminate and produce a nonnegative state $\widehat{X}(t+\tau)$. It is shown in [14] that this bounding procedure does not affect the consistency of the leaping method since it will only modify events with probabilities of order $O(\tau^2)$ or higher.

3. Reversible-equivalent-monomolecular tau

3.1. The general idea

The proposed REMM-τ method is based on a sequence of three ideas.

The first idea is the decomposition of the system of reactions into subsystems that are then considered in isolation. The number of firings of reaction channels forming such subsystems during a time step τ is approximated by the exact number of firings that would have occurred in the absence of all other reaction subsystems. The smaller the step size τ, the more accurate we expect this approximation to be.

To see the intuition behind this idea let us consider the case where each subsystem consists of a single reaction channel. We denote by $X(t)$ the exact solution of a stochastic chemical system and denote by $\widehat{X}(t)$ the leaping approximation. An important requirement of a leaping approximation is *consistency*. It was shown in [14] that the explicit and implicit (Poisson) tau methods are first order consistent in the sense that

$$E(g(\widehat{X}(t+\tau)) - g(X(t+\tau))|X(t) = \widehat{X}(t) = x) = O(\tau^2), \quad \tau \to 0, \tag{13}$$

where g is any scalar valued multivariate polynomial function of the state. Let us denote by $K_j(x,\tau)$ the actual (random) number of firings of reaction channel j during the interval $(t, t+\tau]$ given $X(t) = x$. The analysis in [14] shows that to order $O(\tau)$, $K_1, \ldots, K_M$ are independent random variables. Thus if $\widehat{K}_j(x,\tau)$ represent the actual number of firings of reaction channel j in the absence of all other reaction channels, then all moments of $\widehat{K}_j(x,\tau)$ and $K_j(x,\tau)$ will agree to order $O(\tau)$

$$E(\widehat{K}_j^r) - E(K_j^r) = O(\tau^2), \quad \tau \to 0. \tag{14}$$

Then it follows from the state update formulae:

$$X(t+\tau) = x + \sum_{j=1}^{M} v_j K_j(x,\tau)$$

and

$$\widehat{X}(t+\tau) = x + \sum_{j=1}^{M} v_j \widehat{K}_j(x,\tau)$$

that (13) holds for any scalar valued multivariate polynomial function g. This may be easily shown by the (finite) Taylor expansion of g and using (14).

The second idea of our method is based on the observation that several chemical reaction systems contain reversible pairs of reaction channels that are often fast and are a source of stiffness. Thus our method treats each reversible pair as a whole but in isolation from the other reaction channels. Such reversible pairs can be monomolecular or bimolecular, in the latter case a further approximation is necessary.

The third idea is to approximate reversible bimolecular reactions by suitable reversible monomolecular reactions. This is necessary since no known simple analytical solutions for bimolecular reversible reactions are available. In the next section, we derive exact analytical solutions for two different types of reversible monomolecular reactions that will be used extensively. We then elaborate on the bimolecular to monomolecular approximations in Section 3.3.

Remark 1. It must be noted upfront that the REMM-τ method produces a Markovian approximation to a Markov process (see Remark 7). The notion of equivalence between a monomolecular reversible reaction pair and a bimolecular reversible reaction pair to be presented is not an exact one; but rather an approximate one. Approximating one Markov process by another leads to "loss of memory" in the sense that the error made in each step is propagated very much in the same manner as it happens in one-step methods such as Runge–Kutta for ODE solvers. However this approximation applies only over a time interval τ and a new approximation is computed based on the new state just as the vector field is recomputed after each time step in Runge–Kutta type methods. Thus the "loss of memory" is limited in the sense that after every time step τ a new approximation is computed based on the bimolecular nature of the system. Thus the sample trajectories produced by the REMM-τ method do not correspond to those of any single monomolecular system.

3.2. Reversible monomolecular reactions

The considerations outlined in the previous subsection would require that we know the exact solution for the two types of monomolecular reversible reactions: $S_1 \leftrightarrow S_2$ (Type 1) and $S \leftrightarrow 0$ (Type 2).

3.2.1. Reversible monomolecular reaction: Type 1

The Type 1 reversible monomolecular reaction is also known as the reversible isomerization reaction and is given by

$$\begin{aligned} S_1 &\xrightarrow{c_1} S_2, \\ S_2 &\xrightarrow{c_2} S_1. \end{aligned} \tag{15}$$

Suppose we denote by $X_i(t)$ the total number of molecules of species S_i for $i = 1, 2$ then we may write

$$\begin{aligned} X_1(t) &= X_1(0) - Y_{12}(t) + Y_{21}(t), \\ X_2(t) &= X_2(0) - Y_{21}(t) + Y_{12}(t), \end{aligned} \tag{16}$$

where $Y_{ij}(t)$ for $i = 1, 2$ and $j = 1, 2$ is the number of S_i molecules at time $t = 0$ that *are observed to be* S_j molecules at a later time t. Note that the definition of Y_{ij} is justified because the monomolecular nature of the system allows us to think of each molecule in the system to be changing from being S_1 to being S_2 and vice versa *independently* of all other molecules. It may be shown (see Appendix A) that $Y_{12}(t)$ and $Y_{21}(t)$ are independent and that

$$Y_{12}(t) \sim \mathscr{B}(N_1, p_1(t)), \quad Y_{21}(t) \sim \mathscr{B}(N_2, p_2(t)), \tag{17}$$

where $\mathscr{B}(N, p)$ denotes the binomial distribution with parameters N and p [24], $N_1 = X_1(0)$, $N_2 = X_2(0)$ and the probabilities $p_j(t)$ are given by

$$p_j(t) = \frac{c_j}{c_1 + c_2}(1 - \mathrm{e}^{-(c_1+c_2)t}). \tag{18}$$

By definition, $Y_{ij}(t)$ satisfy

$$\begin{aligned} Y_{11}(t) + Y_{12}(t) &= X_1(0), \\ Y_{21}(t) + Y_{22}(t) &= X_2(0). \end{aligned} \tag{19}$$

Using (19), we may rewrite (16) as

$$\begin{aligned} X_1(t) &= Y_{11}(t) + Y_{21}(t), \\ X_2(t) &= Y_{22}(t) + Y_{12}(t), \end{aligned}$$

establishing that the solution at all times is the sum of two independent binomial random variables.

Let us define $K_1(t)$ and $K_2(t)$ to be the number of firings of reaction channels 1 and 2, respectively during the interval $[0, t]$ and *notice that they are different from* $Y_{12}(t)$ and $Y_{21}(t)$. However, the relation between these quantities is such that

$$K_1(t) - K_2(t) = Y_{12}(t) - Y_{21}(t).$$

Since the changes in $X_i(t)$ only depend on $K_1(t) - K_2(t)$, it suffices to approximate $K_1(t) - K_2(t)$. Let us define $\widehat{K}_1$ and $\widehat{K}_2$ to be $\widehat{K}_1(t) = Y_{12}(t)$ and $\widehat{K}_2(t) = Y_{21}(t)$.

Thus in vector form we may write the state update for reaction (15) as

$$X(t+\tau) = X(t) + v_1\widehat{K}_1 + v_2\widehat{K}_2,$$

where $v_1 = (-1, 1)^{\mathrm{T}}$, $v_2 = (1, -1)^{\mathrm{T}}$ and

$$\widehat{K}_j \sim \mathscr{B}(N_j, p_j(\tau)), \quad j = 1, 2,$$

with $N_j = X_j(t)$ and $p_j(\tau)$ given by (18). This update method is exact for this particular reaction. Remark that, as one would expect, if $c_2 = 0$ then we only have the forward reaction $S_1 \to S_2$ and

$$\begin{aligned} N_1 &= X_1(0), \quad N_2 = X_2(0) = 0, \\ p_1(\tau) &= 1 - \mathrm{e}^{-c_1\tau}, \quad p_2(\tau) = 0. \end{aligned}$$

Remark 2. It must be emphasized that $\widehat{K}_1(\tau)$ and $\widehat{K}_2(\tau)$ are not meant to approximate $K_1(\tau)$ and $K_2(\tau)$ separately, but rather meant to yield a result which gives the correct update for the state. Since we are leaping in the state space of $X(t)$ values rather than the reaction count space of $K(t)$ values this is sufficient for our purposes. It is critical to observe that attempting to generate approximations for $K_j(t)$ themselves is unnecessary and may even be a hindrance in trying to generate good leaping approximations for $X(t)$.

3.2.2. Reversible monomolecular reaction: Type 2

The Type 2 reversible monomolecular reaction we consider here is given by

$$\begin{aligned} S_1 &\xrightarrow{c_1} 0, \\ 0 &\xrightarrow{a_2} S_1, \end{aligned} \tag{20}$$

where we have intentionally used the notation a_2 (instead of c_2) for reasons that will become clear.

Note that 0 represents the void, and thus $S_1 \to 0$ stands for disintegration of a molecule S_1, and $0 \to S_1$ stands for the production or inflow of a molecule S_1. The first reaction is also called the irreversible isomerization reaction and in the absence of the second reaction has the binomial distribution $\mathscr{B}(N_1, 1 - \mathrm{e}^{-c_1 t})$ for all times. However in the presence of the second reaction, the pair of reactions represent a birth–death process [24,25], which reaches a stationary distribution that is Poisson distributed with mean a_2/c_1 [24]. Here, we are additionally interested in the time evolution for the distribution of $X_1(t)$ (the number of S_1 molecules for all t). We present a derivation of an expression for this distribution by regarding the reversible reaction (20) as a parametric limit of the reversible reaction (15) in which $N_2 = X_2(0) \to \infty$ while $N_2 c_2 = a_2$ is held constant. This can be done because one may think of S_1 in (20) as becoming a "ghost" molecule S_2 in the forward reaction, with the ghost molecule S_2 becoming S_1 in the backward reaction, conditional on assuming an infi-

nite supply of ghost molecules S_2 in order for the backward reaction to have a constant propensity. Thus we take $a_2 = c_2 N_2$ to be held fixed, while we take the limit $N_2 \to \infty$. This means that $c_2 \to 0$. From Eq. (18) we obtain the limits

$$p_1 \to 1 - e^{-c_1 t}, \quad p_2 \to 0$$

and

$$N_2 p_2 \to \frac{a_2}{c_1}(1 - e^{-c_1 t}).$$

It is a very well known property of binomial distributions [24,25] that in the limit that $N_2 \to \infty$ with $N_2 p_2 = \mu_2$ held constant, the binomial distribution $\mathscr{B}(N_2, p_2)$ limits to the Poisson distribution $\mathscr{P}(\mu_2)$ with mean and variance μ_2. For us, μ_2 is given by

$$\mu_2 = N_2 p_2 = \frac{a_2}{c_1}(1 - e^{-c_1 t}).$$

If $Y_{12}(t)$ and $Y_{21}(t)$ are as defined in Section 3.2.1, then in the above limit Eq. (17) tends to

$$Y_{12}(t) \sim \mathscr{B}(N_1, p_1), \quad Y_{21}(t) \sim \mathscr{P}(\mu_2).$$

However, since for all times the solution $X_1(t)$ is given by

$$X_1(t) = Y_{11}(t) + Y_{21}(t),$$

it is consequently the sum of two independent random variables: one with a binomial distribution $Y_{11}(t) \sim \mathscr{B}(N_1, 1 - p_1(t))$ and the other with a Poisson distribution $Y_{21}(t) \sim \mathscr{P}(\mu_2(t))$. As $t \to \infty$, it follows that $p_1(t) \to 1$ and hence $Y_{11} \to 0$ with probability 1. Also as $t \to \infty$, we obtain that

$$\mu_2 \to \frac{a_2}{c_1},$$

recovering the asymptotic stationary Poisson distribution for X_1.

As in Section 3.2.1 it follows that

$$K_1 - K_2 = Y_{12} - Y_{21}.$$

Thus we may write the state update for reaction (20) as

$$X_1(t + \tau) = X_1(t) + \nu_1 \widehat{K}_1 + \nu_2 \widehat{K}_2,$$

where $\nu_1 = -1$, $\nu_2 = 1$ and

$$\widehat{K}_1 \sim \mathscr{B}(N_1, 1 - e^{-c_1 \tau}), \quad \widehat{K}_2 \sim \mathscr{P}\left(\frac{a_2}{c_1}(1 - e^{-c_1 \tau})\right),$$

with $N_1 = X_1(t)$. Once again, this update method is exact for this particular reaction.

3.3. Reversible bimolecular reactions

After having described the exact solutions for the Type 1 and Type 2 reversible monomolecular reactions, we are now in a position to consider general pairs of reversible bimolecular reactions. We distinguish between five different types of bimolecular reaction pairs, and provide a way to approximate each one of them by a suitable monomolecular reversible reaction.

3.3.1. Reversible bimolecular reaction: Type 1

The first reversible bimolecular reaction that we consider is given by

$$\begin{aligned} S_1 + S_2 &\xrightarrow{c_1} S_3, \\ S_3 &\xrightarrow{c_2} S_1 + S_2. \end{aligned} \tag{21}$$

The first thing to note about the dynamics of (21) is that the reaction satisfies the conservation relations

$$X_1(t) + X_3(t) = X_1(0) + X_3(0),$$
$$X_2(t) + X_3(t) = X_2(0) + X_3(0)$$

for all times. The following conservation relation:

$$X_2(t) - X_1(t) = X_2(0) - X_1(0),$$

also holds for all times. Thus there are two independent conserved quantities $X_T = X_1 + X_3$ and $X_e = X_2 - X_1$. In the following discussion we assume without loss of generality that $X_e \geqslant 0$; which implies that $X_1(t) \leqslant X_2(t)$, and hence S_1 is the *limiting species* for the forward reaction. For the backward reaction, S_3 is always the limiting species. The notion of the limiting species is captured by the inequalities:

$$-X_3(t) \leqslant K_1(X(t),\tau) - K_2(X(t),\tau) \leqslant \min\{X_1(t), X_2(t)\} = X_1(t),$$

where K_1 and K_2 are the actual number of firings of the forward and backward reaction channels.

Now, based on the limiting species, we approximate the system (21) by the monomolecular reversible system

$$\begin{aligned} S_1 &\xrightarrow{\tilde{c}_1} S_3, \\ S_3 &\xrightarrow{\tilde{c}_2} S_1, \end{aligned} \tag{22}$$

where the *effective monomolecular rate constants* $\tilde{c}_j$ need to be found. Assuming that we have suitable values for $\tilde{c}_j$ and that $\widehat{X}(t) = x$, we update the state according to

$$\widehat{X}(t+\tau) = x + \nu_1 \widehat{K}_1 + \nu_2 \widehat{K}_2, \tag{23}$$

where $\nu_1 = (-1,-1,1)$, $\nu_2 = (1,1,-1)$,

$$\widehat{K_1} \sim \mathscr{B}(N_1, p_1), \quad \widehat{K_2} \sim \mathscr{B}(N_2, p_2),$$

where $N_1 = x_1$, $N_2 = x_3$ (number of limiting species) and

$$\begin{aligned} p_1 &= \frac{\tilde{c}_1}{\tilde{c}_1 + \tilde{c}_2}(1 - \mathrm{e}^{-(\tilde{c}_1+\tilde{c}_2)\tau}), \\ p_2 &= \frac{\tilde{c}_2}{\tilde{c}_1 + \tilde{c}_2}(1 - \mathrm{e}^{-(\tilde{c}_1+\tilde{c}_2)\tau}). \end{aligned} \tag{24}$$

At this point, we need to derive appropriate values for $\tilde{c}_j$. For this purpose, we first require that the propensity of the forward or backward reaction in (22) must be equal to the propensity of the forward or backward reaction in (21) respectively. This ensures that the method is *consistent* in the sense of numerical analysis terminology [26,14], keeping in mind that in general $\tilde{c}_j$ may depend on the current state $\widehat{X}(t) = x$.

Since the backward reactions are the same for both (22) and (21) it is natural to choose $\tilde{c}_2 = c_2$, and this choice keeps the propensities the same.

Equating the propensities of the forward reactions for both (22) and (21), leads to the condition that

$$\tilde{c}_1 x_1 = a_1(x) = c_1 x_1 x_2$$

for all values of x_1 and x_2. If $x_1 \neq 0$ this implies that

$$\tilde{c}_1 = c_1 x_2.$$

If $x_1 = 0$, both propensities are zero and consistency condition alone does not give us any useful information. One may still use the formula $\tilde{c}_1 = c_1 x_2$ to hold for all values of x_1 and x_2, and this may seem like the simplest choice. But this however does not lead to an accurate formula, especially when $x_2 = 0$ and τ is large. To see this, consider the case when $x_1 = x_2 = 0$. The formula $\tilde{c}_1 = c_1 x_2$ would dictate that $\tilde{c}_1 = 0$ and hence $p_1 = 0$ according to (24). This implies that $\widehat{K}_1 = 0$. i.e. The forward reaction never fires in the interval $[t, t+\tau]$. Which is reasonable when $\tau \to 0$, but unreasonable when τ is large.

In order to deal with the $x_1 = 0$ case, it is instructive to consider the specific small number situation in which $X_T = 1$. This means $X_1(t) + X_3(t) = 1$, and in this instance the bimolecular system (21) behaves exactly like the

monomolecular system $S_1 \leftrightarrow S_3$ with forward and backward rate parameters $(X_e + 1)c_1$ and c_2, respectively. Then the best choice for $\tilde{c}_1$ in the system (22) would be $\tilde{c}_1 = c_1(x_2 + 1)$ when $x_1 = 0$ (because then $X_e = x_2$). The "simpler" choice $\tilde{c}_1 = c_1 x_2$ would lead to a monomolecular system with forward rate $X_e c_1$ which differs from $(X_e + 1)c_1$. This difference is most dramatic when $X_e = 0$, the situation described in the previous paragraph. In Section 4.3 we illustrate this point via a numerical example.

For the general situation (either S_1 or S_2 could be the limiting species) we summarise our choice for $\tilde{c}_1$ as follows:

$$\begin{aligned} &\tilde{c}_1 = (\max\{x_1, x_2\} + 1)c_1 \quad \text{if } \min\{x_1, x_2\} = 0, \\ &\tilde{c}_1 = \max\{x_1, x_2\}c_1 \quad \text{otherwise.} \end{aligned} \tag{25}$$

Remark 3. The equivalent monomolecular approximation chosen here is consistent as $\tau \to 0$, but for large τ we may not expect it to be very accurate in general. Intuitively, we expect it to be more accurate when X_e is large, since the relative fluctuations in the "excess species" will be small when X_T/X_e is small. Clearly the worst case is when $X_e = 0$ which corresponds to $X_1 = X_2$. In Section 4.3 this point is explored via a numerical example.

In summary the REMM-τ applied to the reversible reaction (21) is given by

$$\widehat{X}(t + \tau) = x + v_1 \widehat{K}_1 + v_2 \widehat{K}_2, \tag{26}$$

where $x = \widehat{X}(t)$,

$$\begin{aligned} &\widehat{K_1} \sim \mathscr{B}(N_1, p_1), \quad \widehat{K_2} \sim \mathscr{B}(N_2, p_2), \\ &N_1 = \min\{x_1, x_2\}, \quad N_2 = x_3, \end{aligned}$$

and the p_j are given by (24), $\tilde{c}_1$ by (25) and $\tilde{c}_2 = c_2$.

Remark 4. It is important to note that the v_j in (26) are the stoichiometric vectors of the bimolecular reactions and not of the monomolecular approximation; in other words

$$v_1 = (-1, -1, 1)^T, \quad v_2 = (1, 1, -1)^T.$$

Thus the monomolecular system is only used to generate approximations $\widehat{K}_j$ for the K_j, but not in the updating of the state.

3.3.2. Reversible bimolecular reaction: Type 2

In the same spirit, the slightly more complex bimolecular reversible reaction

$$\begin{aligned} &S_1 + S_2 \xrightarrow{c_1} S_3 + S_4, \\ &S_3 + S_4 \xrightarrow{c_2} S_1 + S_2 \end{aligned} \tag{27}$$

may be approximated in the obvious way. We obtain

$$\begin{aligned} &N_1 = \min\{x_1, x_2\}, \quad N_2 = \min\{x_3, x_4\}, \\ &\tilde{c}_1 = (\max\{x_1, x_2\} + 1)c_1 \quad \text{if } \min\{x_1, x_2\} = 0, \\ &\tilde{c}_1 = \max\{x_1, x_2\}c_1 \quad \text{otherwise;} \\ &\tilde{c}_2 = (\max\{x_3, x_4\} + 1)c_2 \quad \text{if } \min\{x_3, x_4\} = 0, \\ &\tilde{c}_2 = \max\{x_3, x_4\}c_2, \quad \text{otherwise;} \end{aligned}$$

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

3.3.3. Reversible bimolecular reaction: Types 3 and 4

Now, we consider the reversible reaction where two molecules of the same species come together

$$\begin{aligned} &S_1 + S_1 \xrightarrow{c_1} S_2, \\ &S_2 \xrightarrow{c_2} S_1 + S_1. \end{aligned} \tag{28}$$

The REMM-τ update for this reversible pair is still based on the Type 1 reversible monomolecular reaction, and hence is similar to the reversible pair (21), except for the choices for N_1 and $\tilde{c}_1$, which are given by (with $x = \widehat{X}(t)$)

$$\begin{aligned} N_1 = \text{floor}(x_1/2), \quad \tilde{c}_1 &= a_1(x)/N_1, \quad \text{when } N_1 \neq 0, \\ \tilde{c}_1 &= c_1, \quad \text{when } x_1 = 0, \\ \tilde{c}_1 &= 3c_1, \quad \text{when } x_1 = 1, \end{aligned}$$

where floor(z) denotes the largest integer less than or equal to z and $a_1(x) = c_1 x_1(x_1 - 1)/2$ is the propensity of the forward reaction. Note that we obtain $N_1 = 0$ when either $x_1 = 0$ or $x_1 = 1$. Our choice of $\tilde{c}_1$ is exact for specific types of initial conditions for which the bimolecular system behaves exactly like a monomolecular system. One is for the initial condition $X(0) = (2,0)$ (or equivalently $X(0) = (0,1)$) in which case the system behaves identical to the monomolecular system $S_1 \leftrightarrow S_2$ with parameters c_1 and c_2. The other is for the initial condition $X(0) = (3,0)$ (or equivalently $X(0) = (1,1)$) in which case the system behaves identical to the monomolecular system $S_1 \leftrightarrow S_2$ with parameters $3c_1$ and c_2.

The reversible pair

$$\begin{aligned} &S_1 + S_1 \xrightarrow{c_1} S_2 + S_2, \\ &S_2 + S_2 \xrightarrow{c_2} S_1 + S_1 \end{aligned} \tag{29}$$

may be dealt with in the obvious way.

3.3.4. Reversible bimolecular reaction: Type 5

Up to this point, all the reversible bimolecular examples we considered were treated using the Type 1 reversible monomolecular reaction as given by (15). The next bimolecular example that we consider can be best approximated using the Type 2 reversible monomolecular reaction given by (20). The reaction pair we consider here is given by

$$\begin{aligned} &S_1 + S_2 \xrightarrow{c_1} S_2, \\ &S_2 \xrightarrow{c_2} S_1 + S_2. \end{aligned} \tag{30}$$

In this reaction pair S_2 is never altered and the limiting species in the forward reaction is always S_1. The backward reaction has no limiting species. Hence the equivalent monomolecular system is not of the form (22) but rather of the form

$$\begin{aligned} &S_1 \xrightarrow{\tilde{c}_1} 0, \\ &0 \xrightarrow{\tilde{a}_2} S_1. \end{aligned} \tag{31}$$

Consistency considerations would require that

$$\tilde{c}_1 = c_1 x_2, \quad \tilde{a}_2 = c_2 x_2,$$

where $x = \widehat{X}(t)$. In fact since $X_2(t)$ is constant the equivalent monomolecular reaction (31) is identical to the original (30). The presence of S_2 merely scales the constants c_1 and c_2. This leads to the leaping approximation (in fact the exact solution)

$$\widehat{K}_1 \sim \mathscr{B}(N_1, p_1), \quad \widehat{K}_2 \sim \mathscr{P}(\tilde{\mu}_2),$$

where

$$N_1 = x_1, \quad p_1 = 1 - \mathrm{e}^{-\tilde{c}_1 \tau}, \quad \tilde{\mu}_2 = \frac{c_2}{c_1}(1 - \mathrm{e}^{-c_1 \tau}).$$

3.3.5. Reversible bimolecular reaction: Type 6

Here we consider the reaction pair

$$\begin{aligned} &S_1 + S_2 \xrightarrow{c_1} S_1 + S_3, \\ &S_1 + S_3 \xrightarrow{c_2} S_1 + S_2. \end{aligned} \tag{32}$$

In the absence of other reactions, this reaction pair behaves identically to $S_2 \leftrightarrow S_3$ with parameters $\tilde{c}_1 = c_1 x_1$ and $\tilde{c}_2 = c_2 x_1$, where x_1 is the (constant) number of S_1 molecules. This is clear because the number of S_1 molecules is unchanged by the two reactions. For the same reason, the limiting species are always S_2 and S_3 for the forward and backward reactions, respectively.

This leads to the leaping approximation (in fact the exact solution)

$$\widehat{K}_1 \sim \mathscr{B}(N_1, p_1), \quad \widehat{K}_2 \sim \mathscr{B}(N_2, p_2),$$

where

$$N_1 = x_2, \quad N_2 = x_3$$

and

$$\begin{aligned} p_1 &= \frac{\tilde{c}_1}{\tilde{c}_1 + \tilde{c}_2}(1 - \mathrm{e}^{-(\tilde{c}_1+\tilde{c}_2)\tau}), \\ p_2 &= \frac{\tilde{c}_2}{\tilde{c}_1 + \tilde{c}_2}(1 - \mathrm{e}^{-(\tilde{c}_1+\tilde{c}_2)\tau}). \end{aligned}$$

3.4. The REMM-τ algorithm: parallel versus sequential updating

At this point, we have derived and presented the main steps of the REMM-τ method as it applies to subsystems consisting of reversible pair of reactions. As outlined in Section 3.1, a general algorithm to apply REMM-τ method to a large system of chemical reactions can be devised by considering each such reversible reaction in isolation, then updating according to the methods described in Sections 3.2 and 3.3. If a reaction is not reversible but is of the form of any of the reactions in the reversible pairs considered in Sections 3.2 and 3.3, then one can still apply the method with the rate constant of the backward reaction set to zero.

A general important requirement for any leap method is that starting with a state $\widehat{X}(t)$ with nonnegative components, the resulting state $\widehat{X}(t+\tau)$ after one application of the method should also have nonnegative components. As pointed out in [13,12] there are two main situations that lead to a negative state.

To see this, we first note that given a current state x with nonnegative components, for each reaction j one may define a *limiting number of firings* N_j, by the criterion that N_j is the largest nonnegative integer such that $x + v_j N_j$ does not have negative components. We also define the limiting species corresponding to j (and x) to be any species i such that $x_i + (N_j + 1)v_{ji} < 0$. The first situation of negative states described in [13,12] may occur when for a given reaction with index j, the leaping approximation $\widehat{K}_j$ exceeds N_j, i.e. $\widehat{K}_j > N_j$. This situation is circumvented by the construction of REMM-τ, since the method is designed to satisfy the condition $\widehat{K}_j \leqslant N_j$. It is important to note that the condition $\widehat{K}_j \leqslant N_j$ is neither necessary nor sufficient to avoid negative states in general. However in systems where each species is decreased by at most one reaction channel (particular examples being a single reversible pair of reactions), this is indeed a *sufficient condition*.

To illustrate the second situation that leads to negative states, consider a species with index i which is diminished by two different reactions with indices j and l with limiting numbers N_j and N_l. Thus we have assumed $v_{ji} < 0$ and $v_{li} < 0$. It is possible that $x_i + v_{ji}\widehat{K}_j + v_{li}\widehat{K}_l < 0$ even if $\widehat{K}_j \leqslant N_j$ and $\widehat{K}_l \leqslant N_l$. i.e. The combined decrease in the species i due to the approximated simultaneous firings of reactions j and l may lead to a negative state. The REMM-τ as described so far does not necessarily avoid this second situation of negative states. We suggest two ways to deal with such negative states. These two algorithms, the parallel update and the sequential update are described below.

3.4.1. REMM-τ with parallel updating

(1) Given current state $x = \widehat{X}(t)$ and having chosen a stepsize τ, consider each reversible reaction pair in isolation from all other reaction channels and generate samples for $\widehat{K}_j$ according to the prescriptions set forth in Sections 3.2 and 3.3. Note that conditioned on $x = \widehat{X}(t)$ these $\widehat{K}_j$ are all independent random variables, some are binomial and some are Poisson depending on the reaction.
(2) Update the state according to the general state update formula

$$\widehat{X}(t+\tau) = x + \sum_{j=1}^{M} v_j \widehat{K}_j.$$

(3) If $\widehat{X}(t+\tau)$ has negative components, apply the bounding procedure (12) to obtain a new nonnegative value for $\widehat{X}(t+\tau)$.

3.4.2. REMM-τ with sequential updating

Instead of updating the state of the system based on the collective firing of all reaction channels, the idea here is to sequentially update based on individual reversible reaction pairs. The main idea is the same as the one proposed in [12], except that we update reversible pairs simultaneously.

The algorithm proceeds as follows. Suppose the reaction channels are ordered so that channels with indices $j = 2l-1$ and $2l$ for $l = 1, \ldots, L$ form a reversible pair. Thus we are assuming there are $2L$ reaction channels. If a reaction channel with index $2l-1$ does not have a reversible counterpart we simply take $c_{2l} = 0$.

(1) Given current state $x = \widehat{X}(t)$ and a step size τ set $x^{(0)} \leftarrow x$.
(2) Execute the following loop. for $l = 1{:}L$
   (a) Generate samples for $\widehat{K}_{2l-1}$ and $\widehat{K}_{2l}$ using the prescriptions outlined in Section 3.3 taking x to be $x^{(l-1)}$ and using the same time step τ.
   (b) Compute the lth intermediate state $x^{(l)}$ using

$$x^{(l)} = x^{(l-1)} + v_{2l-1}\widehat{K}_{2l-1} + v_{2l}\widehat{K}_{2l}.$$

   end for
(3) Set $\widehat{X}(t+\tau) \leftarrow x^{(L)}$.

The sequentially updated REMM-τ is guaranteed to produce nonnegative states without having to apply the bounding procedure.

Remark 5. The choice of ordering $l = 1, \ldots, L$ of the reversible reaction pairs will be an important subject to be studied further. It is intuitive to choose an ordering based on the total propensity of a reversible pair; updating the fastest pair first and the slowest pair last. Naturally this ordering will in general depend on the current state $X(t) = x$. It is also possible to choose the ordering $l = 1, \ldots, L$ of the reversible reaction pairs in a random manner at each time step τ as suggested in [12]. In this case, it is still intuitive to choose the probabilities of the ordering such that faster pairs are more likely to get updated earlier. It is also important to note that in the presence of "race conditions", i.e. when the system is in a state where small fluctuations may lead to drastically different future states, the ordering of the reactions in the sequential updating may bias the probabilities of the future outcomes. Typically this situation arises in bimodal systems (also known as bistable) where the asymptotic distribution has two modes (peaks). The deterministic analogue of this situation is a system with two stable equilibria with their own basins of attraction and the sensitive states are the ones close to the boundary between the two regions of attraction. Intuitively one would expect that the random ordering of the reaction pairs based on propensities may be better since all reaction pairs are given an opportunity to fire first based on their propensities. But it is also important to note that when the system is in such a sensitive state, it is desirable that the stepsize τ be chosen to be small. Thus adaptive stepsize selection is again an important issue. We do not investigate the topic of ordering the reactions in this paper, but consider it to be an important topic of future investigation.

Remark 6. In some examples it is natural to group three reactions into a subsystem (motif) which is exactly solvable in isolation. See the last example presented in this paper for such a situation. We believe that future development of REMM-τ is likely to involve groupings of several reactions as a motif.

Remark 7. All leap methods proposed in literature as well as the REMM-τ (both the sequential and parallel versions) lead to discrete time Markov processes. In particular, in the sequential updating, the successive intermediate states $x^{(l)}$ are not part of the final numerical solution. Thus given $\widehat{X}(t)$ for a fixed τ there is a well defined Markov transition function

$$P(\widehat{X}(t+\tau) = x' | \widehat{X}(t) = x),$$

where $x, x' \in \mathbb{Z}_+^N$ are any pair of states. This is also true of random order of updating.

4. Examples

In this section we illustrate the use of the REMM-τ method and compare its performance to those of the implicit tau and the trapezoidal implicit tau using a number of test examples. The first two examples we consider involve reversible reactions that are themselves very fast, in the absence of any slow non-reversible reaction. The rationale for starting with these examples is that such systems are often "fast subsystems" of a bigger system where the rest of the reactions occur on a much slower time scale. It is therefore desirable that a leaping method produces reasonable approximates for the asymptotic probability distribution for the "fast species" when applied with a stepsize that is small in comparison to the time scale of the rest of the reactions but large in comparison with the time scale of the fast subsystem. The third example precisely illustrates such a situation in a biologically motivated context.

In the analysis of these examples, we mostly focus on the *stiff behavior* of the methods. We use the term *stiff behavior* to denote the asymptotic behavior of a method when applied with a constant stepsize τ that is not necessarily smaller than the time scale of the system. Such aspect of any leaping method is of great important since the aim is to generate accurate approximations of the real process with the largest possible leaping stepsize τ. Therefore, we specifically consider the stiff behavior as a function of τ and investigate if the methods being compared provide a robust performance as τ is increased. Since our simple test examples are meant to be fast subsystems of a larger system, it would be typical to use stepsizes τ that may be much larger than the time scale of the fast subsystem. As a reference for the comparison, we use the true asymptotic probability distribution of the test systems computed through exact SSA simulations or accurate numerical integration of the CME.

To estimate the time scales of a system, two approaches are possible. The first approach is to consider the reaction rate equations (RRE) and compute the eigenvalues of the (time varying) Jacobian matrix along "typical trajectory values". We use this method for linear propensity systems where the Jacobian is constant and therefore choice of trajectories is irrelevant. In this case, it may be easily shown that the time evolution of the moments of the system follow linear systems of ODEs whose eigenvalues are the same (for the mean) or are an integer multiple (for higher order moments) of the eigenvalues of the corresponding RRE. Alternatively, one may consider the eigenvalues of the chemical master equations (CME). The latter approach is practical when the system is closed, and hence the CME constitutes a finite dimensional linear time invariant system. When the system is not closed, it is still possible to truncate the infinite system generated by the CME and consider the eigenvalues of the approximate truncated system.

4.1. Linear test example $S_1 \leftrightarrow S_2$

The simplest linear propensity system whose state does not decay to zero is the reversible pair of reactions $S_1 \leftrightarrow S_2$. As shown in previous sections, the time solution of this reversible reaction pair is the sum of two

binomials, and the REMM-τ method is exact by construction. The stiff behavior of the explicit, implicit and trapezoidal implicit τ methods when applied to this test system were studied in [11]. However the analysis in [11] ignores rounding and bounding, whose effect can become very noticeable if the molecular species are present in small numbers as we demonstrate here.

For the case we consider, the parameter values are set to $c_1 = 1$ and $c_2 = 2$ for the forward and backward reactions, respectively. First, we set the total number of molecules in the system to $x_T = 1$, the smallest possible. With this choice of x_T, there are only two possible states $X = 0$ or $X = 1$, where we have used X to denote the number of S_1 molecules. If we let $p = (p_1, p_2)$ be a row vector, where $p_1 = P(X = 0)$ and $p_2 = P(X = 1)$, then the CME is given by

$$\dot{p} = pQ,$$

where the 2×2 matrix Q is given by

$$Q = \begin{bmatrix} -2 & 2 \\ 1 & -1 \end{bmatrix}.$$

The matrix Q has eigenvalues at 0 and -3. Note that the eigenvalue -3 is the same as that of the deterministic RRE. Hence the time scale of the system is considered to be $1/3 \approx 0.33$. The asymptotic distribution of this system is characterized by one quantity, for instance $P(X = 1)$, which is computed by finding the left nullvector of Q. Note that the Markov transition probability matrix $P(\tau)$ defined by

$$P_{ij}(\tau) = P\{X(t + \tau) = j | X(t) = i\}$$

may be computed by matrix exponential $P(\tau) = e^{\tau Q}$. While the exact model of the chemical system is a discrete state and continuous in time Markov process, the result of any leaping approximation with constant time step τ may be regarded as a discrete state and discrete time Markov process. (See Appendix B). The Markov transition matrices $\widehat{P}(\tau)$ corresponding to the implicit tau and trapezoidal implicit tau methods, as well as the stationary probability distributions (if any) corresponding to these matrices, can also be computed for various τ. (See Appendix B for some details.) Once $\widehat{P}(\tau)$ is computed, the asymptotic distribution is given by the left eigenvector of $P(\tau)$ with eigenvalue 1 provided it is unique. We refer the reader to [24,25] for basics on discrete state Markov processes (with both continuous and discrete time cases) and their relevant properties.

Fig. 1 shows the asymptotic probability $P(X = 1)$ as a function of τ for both these methods. For τ values much smaller than 0.33 convergence is seen, and both methods coincide due to rounding as pointed out in [14]. In fact for small τ, both methods coincide with explicit-τ due to rounding. For very large values of τ, in par-

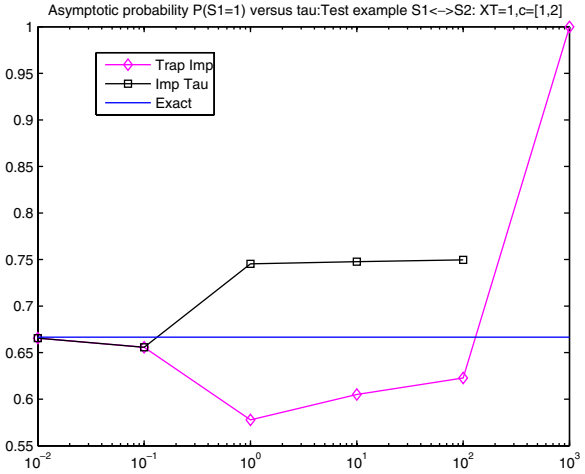

Fig. 1. Linear test example $S_1 \leftrightarrow S_2$ with total number of molecules $x_T = 1$. The plot shows the asymptotic probability $P(X = 1)$ where X is the number of S_1 versus stepsize τ. Square – implicit-τ, diamond – trapezoidal-implicit-τ. The exact value (CME) is shown by the horizontal line. For $\tau = 1000$ the implicit-τ does not lead to an ergodic chain, so it is omitted.

ticular $\tau = 1000$ or greater, both methods become very poor in estimating the asymptotic probability. In fact for very large τ both methods produce a transition matrix very close to the identity matrix. This results in a Markov chain which stays in the initial state with probability almost equal to 1, a behavior far from that of the true system. A noticeable behavior for the trapezoidal implicit-τ in this example is a non-monotonic dependence of the accuracy on τ. Indeed, for intermediate values of the stepsize, the accuracy of the trapezoidal method deteriorates, then approaches the exact solution before grossly overestimating the probability in the limit where the transition matrix tends to one.

We repeated the above computations for the same parameters, except that we used a total number of molecules $x_T = 3$. This resulted in a 4×4 Q matrix for the CME with eigenvalues at 0, -3, -6, -9. Hence the time scale is $1/3 \approx 0.33$, the same as before. The asymptotic distribution in this case is no longer characterized by a single number. We therefore plot the total variation error between the asymptotic distribution resulting from the trapezoidal and implicit τ methods and that computed from the CME. Fig. 2 shows the result, where again divergence is seen for large τ values.

Remark 8. For the case where $x_T = 3$, the asymptotic variance of the trapezoidal implicit-τ for large τ values was much smaller than the actual asymptotic variance of the process. This does not contradict the analysis in [11], which did not include effects of rounding and bounding. In fact in this example, for large values τ the probability of bounding (equivalently encountering a negative state) was almost zero. Thus we attribute this deviation from the analysis in [11] primarily to rounding rather than bounding. It is possible that a suitable modification of the rounding as well as bounding may lead to better performance. However a detailed investigation of this topic is beyond the scope of this paper.

4.2. Linear test example $S_2 \leftrightarrow S_1 \leftrightarrow S_3$

A more complex situation where the REMM-τ method is no longer exact arises when S_1 is modified through two reaction pairs. Specifically, we consider the test system

$$S_1 \leftrightarrow S_2, \quad S_1 \leftrightarrow S_3.$$

This system remains bounded and has one conserved quantity

$$X_1(t) + X_2(t) + X_3(t) = X_T.$$

When $X_T = 10$ and the parameter values are taken to be

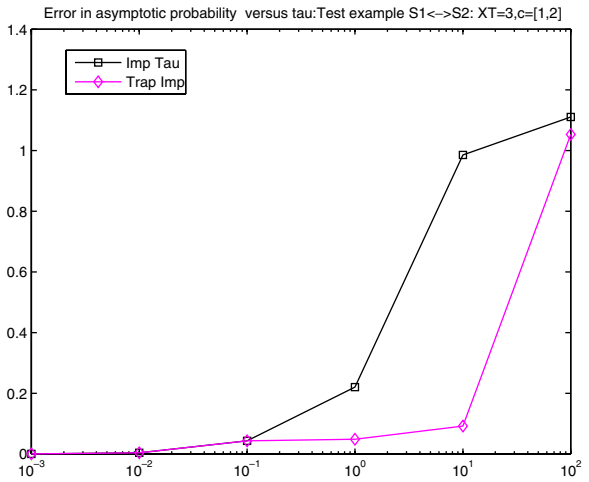

Fig. 2. Linear test example $S_1 \leftrightarrow S_2$ with total number of molecules $x_T = 3$. The plot shows error of the asymptotic probability distribution versus stepsize τ. The error is measured in terms of total variation norm. Square – implicit-τ, diamond – trapezoidal-implicit-τ.

$$c_{12} = 1, \quad c_{21} = 2, \quad c_{13} = 3, \quad c_{31} = 4,$$

the eigenvalues of the Jacobian matrix corresponding to the RRE are 0, −2.3542 and −7.6458. The time scale of the system is approximately given by $1/2.3542 \approx 0.5$. The system is started at the initial conditions $X(0) = (10, 0, 0)^T$, and the simulation interval is taken to be $[0, 1000]$ at the end of which the system has settled to its asymptotic distribution. Fig. 3(a) and (b) shows the estimated probability density function (PDF) at time $T = 1000$ for S_1 and S_2, respectively, computed using the SSA along with PDFs generated by the REMM-τ, implicit-τ and trapezoidal-implicit-τ with a stepsize $\tau = 10$. We deliberately chose the stepsize $\tau = 10$ to be much larger than the slowest time scale ≈0.5. It is clear from these figures that while none of the methods accurately captures the asymptotic PDF for S_1, REMM-τ captures the PDF of S_2 rather accurately. Furthermore, REMM-τ which is a fully explicit method performs similar to trapezoidal implicit-τ and much better than the fully implicit-τ.

In a simulation of 10,000 sample trajectories each with 100 time steps, the REMM-τ encountered 9611 steps (out of 10^6 steps which is about 0.96%) in which negative states occurred and a bounding procedure had to be performed. For the trapezoidal implicit-τ this number was 53,360 (corresponding to 5.34%) and for the implicit-τ method this number was 0.

To investigate the dependence of these results on the choice of τ, we repeated the same computations for a larger time step, specifically $\tau = 100$. We took 100 time steps so as to allow the numerical methods to settle down to their asymptotic distribution. Fig. 4(a) and (b) shows the PDF for S_1 while Fig. 4(c) and (d) show the PDF for S_2 using REMM-τ and trapezoidal implicit-τ for $\tau = 10$ and $\tau = 100$. It is apparent that the behavior of REMM-τ does not change much as the step size is varied, illustrating the robustness of this method to the choice of the time step. In contrast, panels (b) and (d) of Fig. 4 clearly indicate that the behavior of the trapezoidal-implicit-tau deteriorates when τ is changed from $\tau = 10$ to $\tau = 100$. The performance of the implicit-tau was worse and is not shown.

4.3. Nonlinear test example

The first simple nonlinear test example that we investigate is the bimolecular reversible reaction of Type 1 considered in (21) and given by

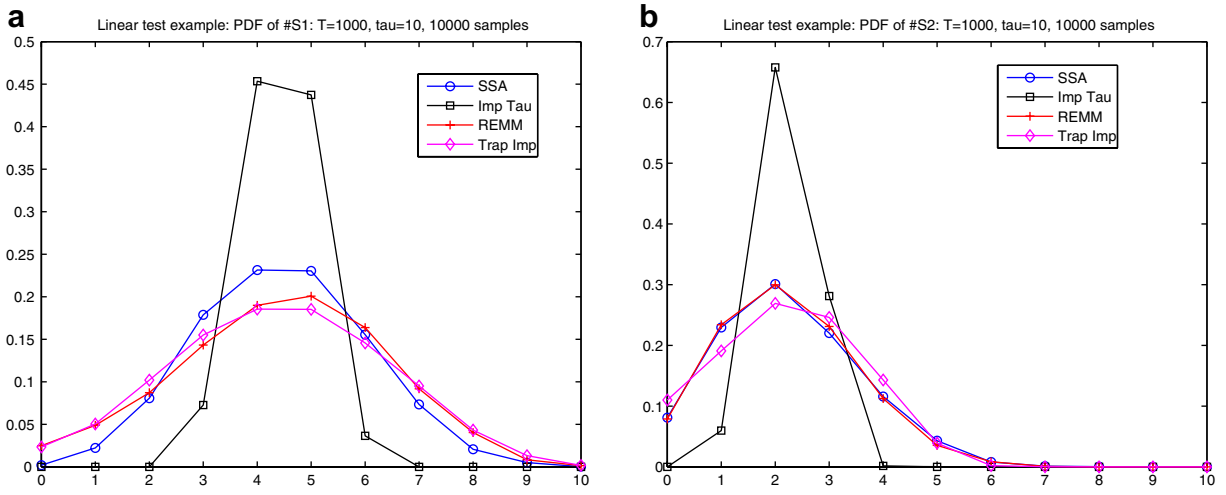

Fig. 3. Linear test example $S_2 \leftrightarrow S_1 \leftrightarrow S_3$: (a) Estimated PDF of $X_1(1000)$. (b) Estimated PDF of $X_2(1000)$. Circle – SSA, plus – REMM-τ, square – implicit-τ, diamond – trapezoidal implicit-τ. Stepsize $\tau = 10$ for all methods. PDF was estimated using Monte Carlo simulation of sample size 10,000.

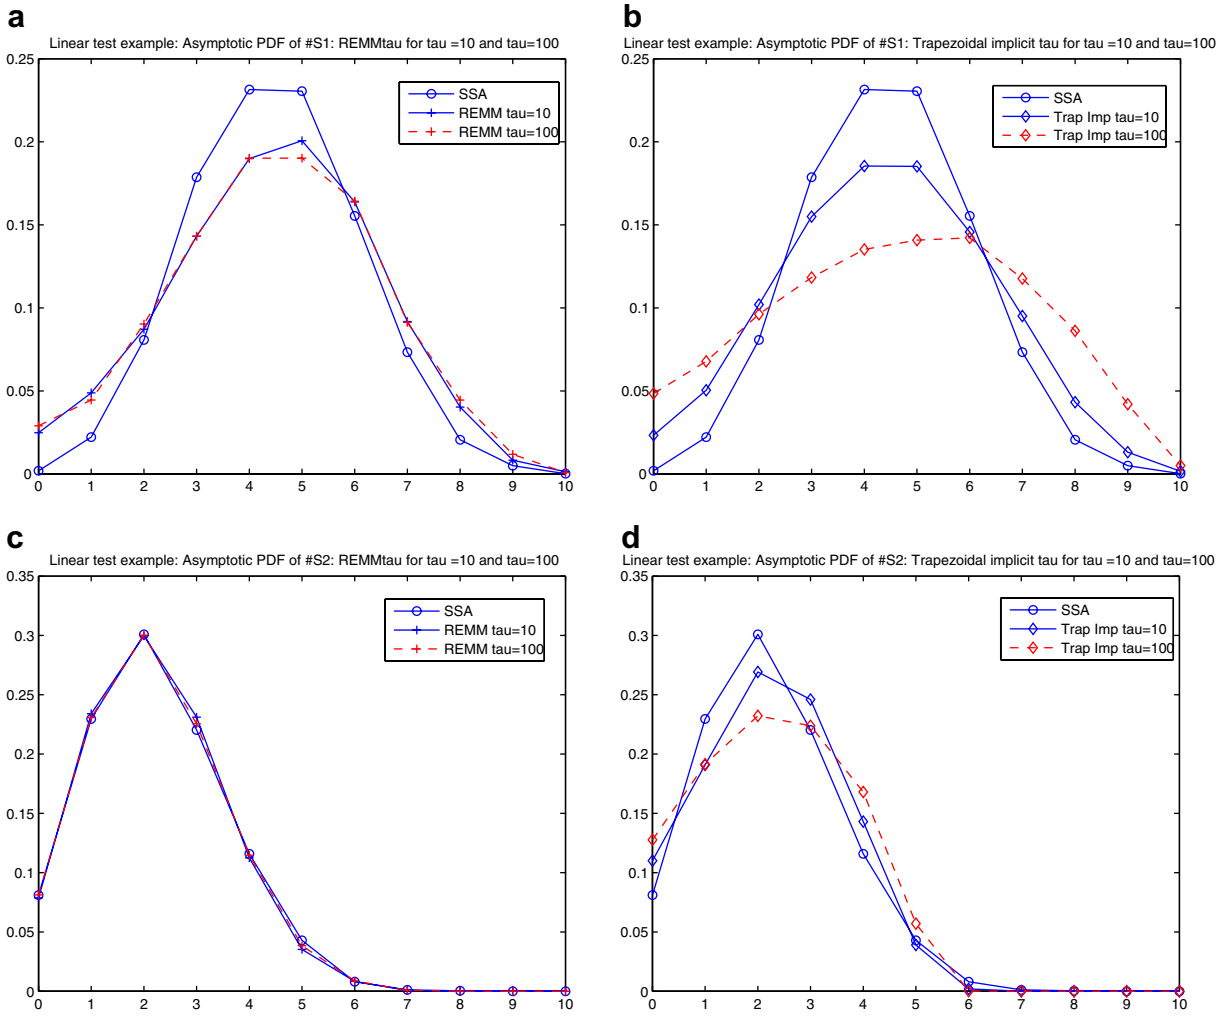

Fig. 4. Linear test example $S_2 \leftrightarrow S_1 \leftrightarrow S_3$: Estimated asymptotic PDF of X_1. (a) S_1 PDF using REMM-τ Circle – SSA, plus with solid line – REMM-τ with $\tau = 10$ and 100 time steps, plus with dashed line – REMM-τ with $\tau = 100$ and 100 time steps. (b) S_1 PDF using trapezoidal implicit τ Circle – SSA, diamond with solid line – trapezoidal-implicit-τ with $\tau = 10$ and 100 time steps, diamond with dashed line – trapezoidal-implicit-τ with $\tau = 100$ and 100 time steps. PDF was estimated using 10,000 trajectories. (c) and (d) are PDF for S_2 with the same conventions. The implicit-tau performed worse and is not shown here.

$$S_1 + S_2 \xrightarrow{c_1} S_3,$$
$$S_3 \xrightarrow{c_2} S_1 + S_2.$$

Let X_i, $i = 1,2,3$ denote the number of molecules of S_i. This reaction system has two independent conserved quantities $X_1 + X_3$ and $X_2 + X_3$. Thus $X_2 - X_1$ is also a conserved quantity. Defining $X_T = X_1 + X_3$, $X_e = X_2 - X_1$, we may essentially treat the system as consisting of a single state $X = X_1$. Without loss of generality we shall take $X_e \geqslant 0$.

It must be noted that when $X_T = 1$ the system behaves identical to a monomolecular system $S_1 \leftrightarrow S_3$ in which the parameters are $(X_e + 1)c_1$ for the forward and c_2 for the backward reactions. Letting $p = (p_1, p_2)$, where $p_1 = P(X_1 = 0)$ and $p_2 = P(X_1 = 1)$, the CME describing the system is given by

$$\dot{p} = pQ,$$

where the 2×2 matrix Q assumes the form

$$Q = \begin{bmatrix} -c_2 & c_2 \\ (X_e+1)c_1 & -(X_e+1)c_1 \end{bmatrix}.$$

In this case, the REMM-τ method with $\tilde{c}_1$ chosen according to (25) gives the correct Markov transition matrix for all τ.

To illustrate this numerically, we set $c_1 = c_2 = 100$, $X_e = 0$ and $X_T = 1$. Then the transition matrix $P(\tau)$ with $\tau = 1$ for the exact process is given by

$$P = \begin{bmatrix} 0.5 & 0.5 \\ 0.5 & 0.5 \end{bmatrix}$$

(note that $P = e^{Q\tau}$), and this coincides with the transition matrix for the REMM-τ method with the choice (25). Note that this leads to the asymptotic probability $P(X_1 = 1) = 1/2$. However, if we choose $\tilde{c}_1 = \max\{x_1, x_2\}c_1 = x_2 c_1$ (we assumed $x_1 \leqslant x_2$) instead of (25) in the REMM method, we obtain the transition matrix

$$\begin{bmatrix} 0 & 1 \\ 0.5 & 0.5 \end{bmatrix}$$

which is clearly wrong and leads to the wrong asymptotic probability $P(X_1 = 1) = 2/3$.

Note that when $X_T > 1$, the bimolecular system is truly nonlinear and the REMM method is no longer exact. We choose $X_T = 3$, $X_e = 5$, $c_1 = 100$, $c_2 = 1000$ and compare the performance of the REMM-τ method against the implicit tau and the trapezoidal implicit tau. For this choice of parameters, the matrix Q has real negative eigenvalues (except one zero), the smallest (in magnitude) of which is ≈ -1830. Hence, the time scale of the system is around $1/1830 \approx 0.001$. Proceeding as with the first test example, we can compute the 4×4 Markov transition matrices corresponding to the implicit-τ, trapezoidal implicit-τ as well as the REMM-τ methods for various τ values and determine the stationary probability distribution corresponding to these matrices.

The result of these computations is shown in Fig. 5 where the asymptotic probability distributions of the methods REMM-τ, implicit-τ and trapezoidal implicit-τ are compared against the exact asymptotic distribution. Our computations show that for small τ values ($\tau \sim 0.0001$) convergence is seen for all methods. However, as τ increases the behavior of implicit-τ and trapezoidal implicit-τ become more erratic. In contrast the REMM-τ method is more robust to variations in τ and shows a modest increase in the error which then stabilizes for τ values greater than 0.01. The plots are shown only for the values $\tau = 0.0001, 0.01, 1$. It can be seen that the REMM-τ does not change noticably for $\tau \geqslant 0.01$.

Remark 9. This robust behavior of the REMM-τ does not come as a surprise since all the probabilities $\widehat{P}_{ij}(\tau)$ in the Markov transition matrix corresponding to the REMM-τ method have a dependence on τ of the form

$$1 - e^{-C_{ij}\tau},$$

where C_{ij} are integer linear combinations of the constants c_1 and c_2 . A careful scrutiny of the formulation of the REMM-τ method will show that the C_{ij} are always greater than or equal to $c_1 + c_2$. A similar statement holds true in a more general context than this specific example, hence one can expect that in general, the stiff behavior of REMM-τ to be robust with respect to changes in τ values.

Finally we explored the behavior of REMM-τ for the larger value $X_T = 100$ and for two different choices for X_e; $X_e = 0$ and $X_e = 10$. See Fig. 6. As mentioned in Remark 3, we expect the performance of REMM-τ to be the worst when $X_e = 0$ for fixed values of X_T, c_1, c_2. Our computations show that when $X_e = 0$ the result of REMM-τ deteriorates for large stepsizes. In fact trapezoidal-implicit tau does better in this case. However even for the small value $X_e = 10$ of excess species (10% of X_T) the performance of REMM-τ quickly recovers. We also observed (plots not shown here) that when $X_e = 0$ and for smaller values of X_T such as 10 the error is less dramatic.

4.4. Biological example: genetic circuit

The final example we present builds on the previous simpler nonlinear test example, to present a genetic transcription module that has important biological significance. This example consists of a gene D_A that encodes for protein A. A itself can bind to its own gene promoter and act as a repressor for its own production.

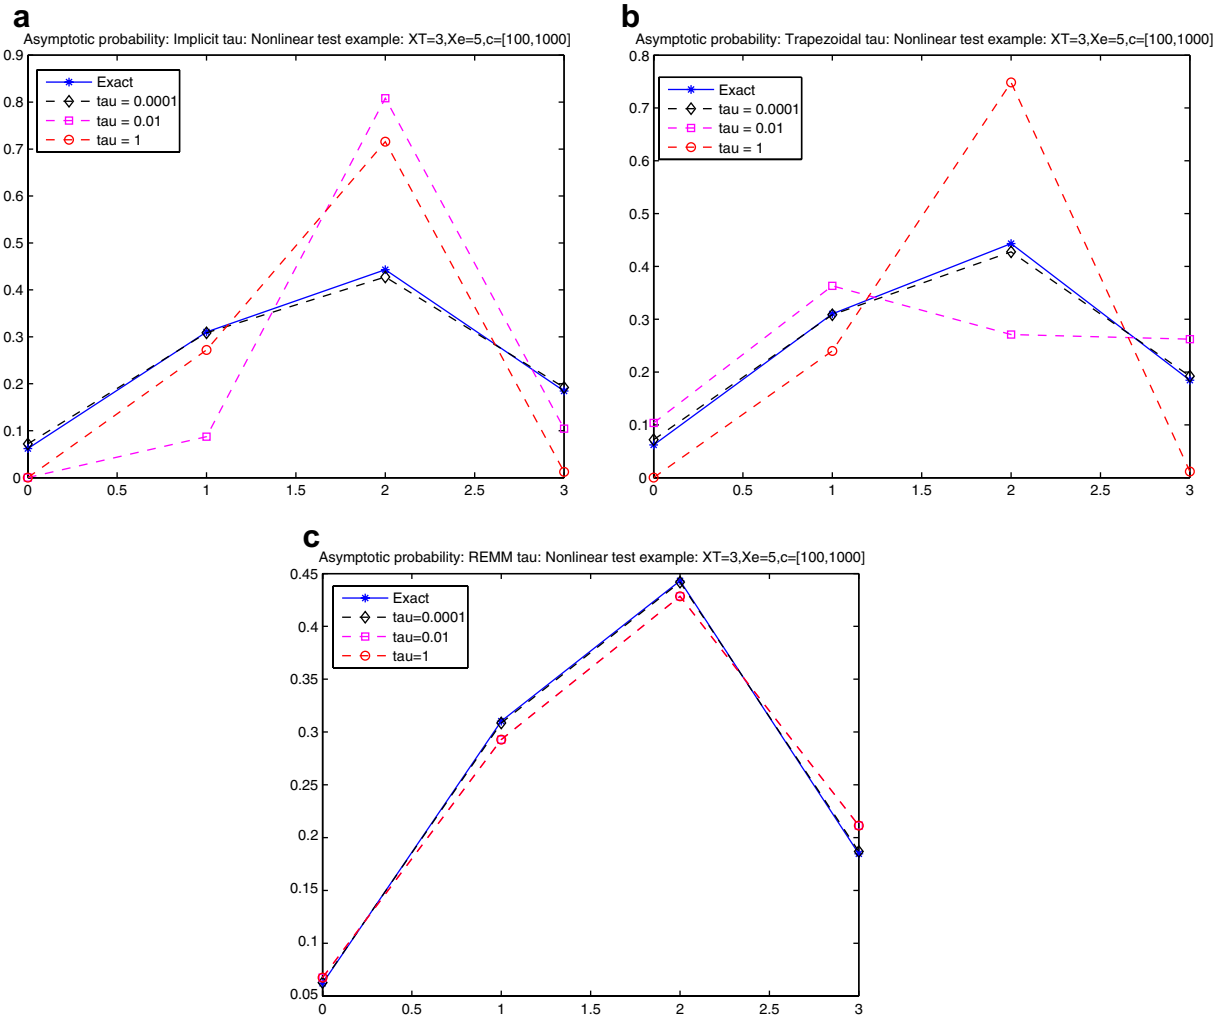

Fig. 5. Nonlinear test example with $X_e = 5$ and $X_T = 3$. The plots show the asymptotic probability distributions corresponding to the leap methods (dashed lines) against the exact asymptotic distribution (star with solid line), for three different stepsizes (diamond – $\tau = 0.0001$, square – $\tau = 0.01$, circle – $\tau = 1$). (a) implicit tau, (b) trapezoidal implicit tau, (c) REMM-τ.

In other terms, when the gene promoter is naked, A is produced at a rate c_3. However, when A is bound, the rate of production of A itself decreases to $c_4 < c_3$, hence implementing a negative feedback loop. If $c_4 > c_3$, then the situation represents a positive feedback loop. Both possibilities exist in biological systems. We consider the negative feedback situation here. The rates of binding and unbinding of A to the gene promoters are c_1 and c_2, respectively. Finally, A naturally decays at a rate c_5. The set of biochemical reactions that give rise to this scheme of gene regulation are given by

$$D_A + A \xrightarrow{c_1} D'_A,$$
$$D'_A \xrightarrow{c_2} D_A + A,$$
$$D_A \xrightarrow{c_3} D_A + A,$$
$$D'_A \xrightarrow{c_4} D'_A + A,$$
$$A \xrightarrow{c_5} 0,$$

where D'_A is the bound form of the gene promoter D_A. The state vector describing the system is given by $x = (x_1, x_2, x_3)$ where $x_1 = \#A$, $x_2 = \#D_A$ and $x_3 = \#D'_A$ and $x_2 + x_3$ is a constant. The propensities are in turn given by

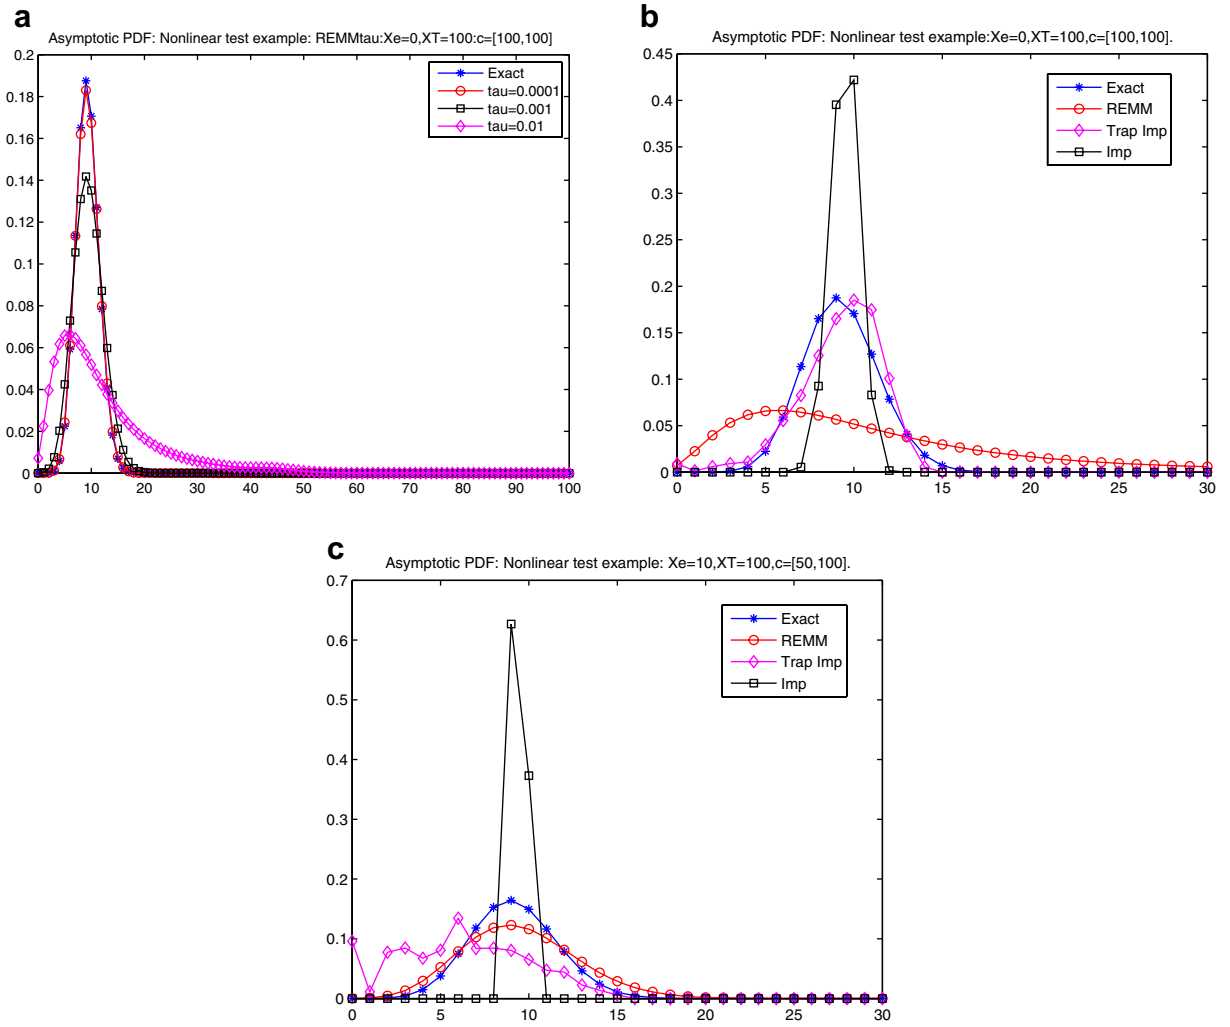

Fig. 6. Nonlinear test example with $X_T = 100$. The plots show the asymptotic probability distributions corresponding to REMM-τ as well as the exact asymptotic distribution for X_1. (a) For the case $X_e = 0$, $c_1 = 100$, $c_2 = 100$ the plots show REMM-τ for time steps $\tau = 0.0001$ (circle), $\tau = 0.001$ (square) and $\tau = 0.01$ (diamond) as well as the exact values (star). For larger values of τ, REMM-τ gives the same results as $\tau = 0.01$. The time scale of the exact system is about 5×10^{-4} (as measured by the smallest eigenvalue). (b) For the same case as in (a) this plots shows REMM-τ (circle), trapezoidal-implicit tau (diamond) and implicit tau (square) for time step $\tau = 0.01$, as well as the exact values (star). The PDFs of the trapezoidal-implicit and implicit tau methods were estimated by running 10,000 samples. The PDFs of REMM and the exact process (SSA) were computed using the Markov transition matrices. (c) For the case $X_e = 10$, $c_1 = 50$, $c_2 = 100$ the plot shows REMM-τ (circle),trapezoidal-implicit tau (diamond) and implicit tau (square) for a very large time step $\tau = 1$ and the exact values (star). The time scale of the exact system is about 6.5×10^{-4}. The PDFs of the trapezoidal-implicit and implicit tau methods were estimated by running 10,000 samples. The PDFs of REMM and the exact process (SSA) were computed using the Markov transition matrices.

$$a_1(x) = c_1 x_1 x_2,$$
$$a_2(x) = c_2 x_3,$$
$$a_3(x) = c_3 x_2,$$
$$a_4(x) = c_4 x_3,$$
$$a_5(x) = c_5 x_1.$$

With the corresponding stoichiometric vectors v_j

$$v_1 = (-1,-1,1)^T, \quad v_2 = (1,1,-1)^T, \quad v_3 = (1,0,0)^T, \quad v_4 = (1,0,0)^T, \quad v_5 = (-1,0,0)^T.$$

In order to apply the REMM-τ method to this system, we first note that the first two reactions form a bimolecular reversible pair of Type 1 described in (21). Thus the first pair forms a subsystem. Contrary to the examples presented so far, we consider the last three reactions together as a subsystem or motif since they reduce to the exactly solvable monomolecular reversible form $0 \leftrightarrow S$ if considered in isolation from the first subsystem. In the absence of the first pair of reactions, $X_2(t)$ and $X_3(t)$ are constants as the last three reactions do not modify the number of D_A or D_A'. Thus the last three reactions together (in the absence of the first two) are equivalent to the exactly solvable monomolecular trio of reactions

$$0 \xrightarrow{\tilde{c}_3} S, \quad 0 \xrightarrow{\tilde{c}_4} S, \quad S \xrightarrow{\tilde{c}_5} 0,$$

where $\tilde{c}_3 = x_2 c_3$, $\tilde{c}_4 = x_3 c_4$ and $\tilde{c}_5 = c_5$. This trio, in turn, is equivalent to the reversible monomolecular reaction $0 \leftrightarrow S$ since the sum of two independent Poisson processes with rates $\tilde{c}_3$ and $\tilde{c}_4$ is equivalent to a single Poisson process with a rate $\tilde{c}_3 + \tilde{c}_4$.

Thus the parallel update REMM-τ when applied to this reaction system starting at state x with stepsize τ leads to the update formula

$$\widehat{X}(t+\tau) = x + \sum_{j=1}^{5} v_j \widehat{K}_j,$$

where the v_j are the stoichiometric vectors of the original bimolecular system (see Remark 4) and the $\widehat{K}_j$ are given by

$$\widehat{K}_1 \sim \mathscr{B}(N_1, p_1), \quad \widehat{K}_2 \sim \mathscr{B}(N_2, p_2), \quad \widehat{K}_3 \sim \mathscr{P}(\mu_3), \quad \widehat{K}_4 \sim \mathscr{P}(\mu_4), \quad \widehat{K}_5 \sim \mathscr{B}(N_5, p_5),$$

with

$$N_1 = \min\{x_1, x_2\}, \quad N_2 = x_2, \quad N_5 = x_1,$$
$$p_1 = \frac{\tilde{c}_1}{\tilde{c}_1 + c_2}(1 - \mathrm{e}^{-(\tilde{c}_1 + c_2)\tau}), \quad p_2 = \frac{c_2}{\tilde{c}_1 + c_2}(1 - \mathrm{e}^{-(\tilde{c}_1 + c_2)\tau}), \quad p_5 = 1 - \mathrm{e}^{-c_5 \tau},$$

and

$$\tilde{c}_1 = (\max\{x_1, x_2\} + 1)c_1 \quad \text{if} \ \ \min\{x_1, x_2\} = 0,$$
$$\tilde{c}_1 = \max\{x_1, x_2\} c_1 \quad \text{otherwise},$$
$$\mu_3 = \frac{a_3(x) p_5}{c_5}, \quad \mu_4 = \frac{a_4(x) p_5}{c_5}.$$

We consider the circuit with three promoters, thus $x_2 + x_3 = 3$. We choose the parameter values

$$c_1 = 100, \quad c_2 = 1000, \quad c_3 = 1, \quad c_4 = 0.1, \quad c_5 = 0.1.$$

The system is started at the initial conditions $x(0) = (14, 3, 0)$ and ran for a time interval of length $T = 200$. This formulation results in a stiff system where the first two reaction channels fire much more often than the rest of the three reaction channels. The mean time step for the SSA when used to simulate the system was about 3×10^{-4}, and we chose a much large time step of $\tau = 1$ for the leap methods. With τ thus chosen, we are leaping on average over 3300 reaction events in one time step.

Fig. 7 shows the PDF of A estimated for X_1(T) using SSA. The same figure shows the PDF computed from the different leaping methods. We used a sample size of 10,000 trajectories for all the methods. It is clear that while the REMM-τ captures the distribution well, both the implicit and trapezoidal implicit methods show a bias in the mean and produce more skewed distributions.

In the sample of 10,000 trajectories each with 200 steps the number of time steps in which a negative state was reached (resulting in the bounding procedure being applied) were 0, 131,174 and 0 for the implicit-τ, trapezoidal implicit-τ and the REMM-τ (parallel update) respectively. In terms of probability these correspond to 0, 0.066 and 0, therefore the error in implicit τ cannot be attributed to bounding. Also, note that the parallel REMM-τ does not guarantee nonnegativity because more than one reaction diminish the number of A molecules. On the other hand, the sequential version of REMM-τ guarantees nonnegative states without having to

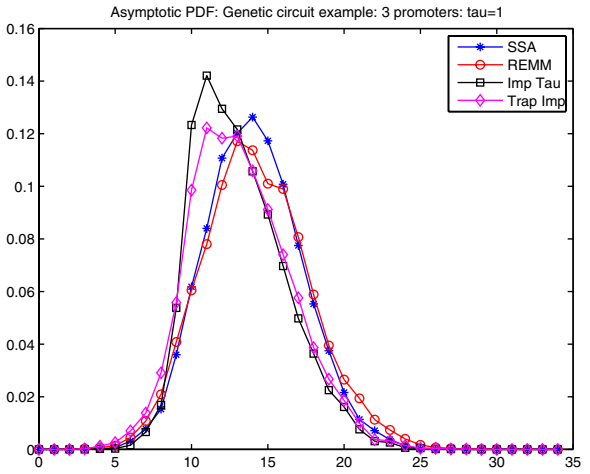

Fig. 7. The estimated final state PDF for $X_1(200)$ for the genetic circuit example with three promoters. The leap methods with stepsize $\tau = 1$ are compared against SSA simulation.

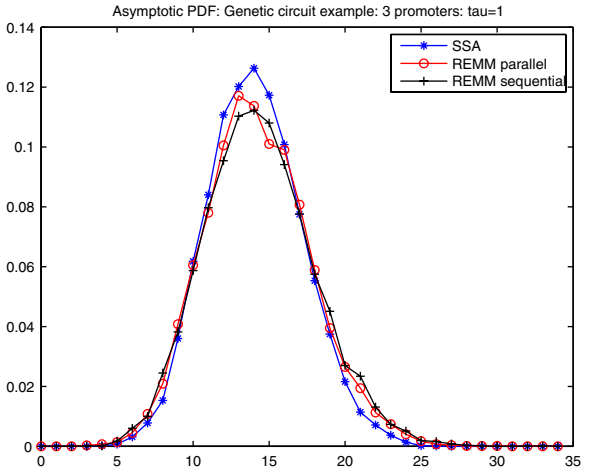

Fig. 8. The estimated final state PDF for $X_1(200)$ for the genetic circuit example with three promoters. The parallel and sequential REMM-τ methods with stepsize $\tau = 1$ are compared against SSA simulation.

perform any bounding procedures. For comparison we ran 10,000 sample trajectories using the sequential REMM-τ as well. In the implementation of the sequential update, we always updated the fast reversible pair first, and then the other slow reactions. Fig. 8 shows the plot comparing the results of parallel REMM-τ and sequential REMM-τ against SSA. Both seem comparable in their performance.

5. Conclusions and future work

In this paper, we developed the reversible-equivalent-monomolecular-tau (REMM-τ), a new leaping method for the simulation of stochastic chemical kinetics. The REMM-τ method considers reversible pairs of reactions in isolation, approximates bimolecular reversible pairs by suitable monomolecular reversible pairs and then advances the state by approximating the actual system by a set of reversible monomolecular pairs of reactions that fire in isolation. The method presented, whether in its parallel or sequential updating form, naturally leads to integer states and hence avoids the error introduced by rounding. Since stiffness manifests itself in a much more complicated way in stochastic systems than in deterministic systems, we defined the notion of

stiff behavior as the behavior of the asymptotic distribution (rather than that of the asymptotic mean and variance alone) generated by a method when applied with a constant stepsize. Using this definition, we compared the stiff behavior of the REMM-τ method to those of the implicit and trapezoidal implicit tau methods for a number of test systems, both linear and nonlinear. We focused on situations which involve a small number of total molecules, since these situations are both frequent in realistic biological problems and challenging for all existing leaping methods. We demonstrated that the stiff behavior of REMM-τ under these conditions is more robust as τ is increased beyond the slowest time scale of the system than those of the implicit and trapezoidal implicit τ methods. The benefits of the accuracy and robustness of REMM-τ are complemented by its simplicity as an explicit method.

In our investigation, we mainly focused on the parallel updating method for REMM-τ. Our future work will focus on the investigation of sequential updating since it guarantees nonnegativity without any bounding procedure and as such is more amenable to analysis. Specifically, a topic of major interest is the choice of the sequential order in which the reversible pairs are updated. In addition the suitability of the sequential updating in the presence of race conditions needs to be investigated.

In addition the performance of leaping methods when applied to stiff oscillatory genetic circuit models will need to be investigated. In oscillatory systems a suitable measure of the periodic nature (such as the power spectrum) rather than the probability distribution at some final time needs to be used as the criterion for comparing the accuracy of leaping methods. Furthermore adaptive stepsize selection will be an important component in a successful implementation of any leaping method when it comes to such oscillatory systems.

While our method is motivated by small number and stiff stochastic chemical models, it will in principle be applicable to other small number and stiff Markov process models such as population dynamics as described in [27]. However the success of REMM-τ in such systems is likely to depend on the presence of fast reversible pair motifs which are common in chemical reaction systems.

Acknowledgment

We like to thank the anonymous reviewers for their valuable comments which helped improve the paper.

Appendix A. Exact solution of $S_1 \leftrightarrow S_2$

Here we show that the exact solution at any finite time t of the chemical system

$$S_1 \xrightarrow{c_1} S_2, \quad S_2 \xrightarrow{c_2} S_1,$$

started with arbitrary initial conditions may be given in terms of two independent binomial random variables.

In [28,29] the solution at time t is derived for general initial conditions by applying transform methods to solve the CME, but the fact that the solution can be expressed in terms of two binomials was not mentioned. Our approach is more combinatorial and is closer to the approach taken in [30] where it is derived that the solution at time t is given by a single binomial random variable when the system is started with initial conditions of a specific form.

Instead of considering the number of S_1 and S_2 molecules in the system, let us focus on a given S_1 or S_2 molecule. Either of these may be considered in isolation from the rest of the molecules to be independently following a dynamic behavior in which an S_1 becomes an S_2 after an exponential waiting time and vice versa. This situation is modeled by a continuous time two state Markov process $Z(t)$, where $Z(t)$ is defined to be $Z(t) = 1$ if the molecule is S_1 at time t, and $Z(t) = 2$ if it is S_2. The Markov process is characterized by the *transition rates* [24] or *propensities* (in the chemical kinetics terminology [9]): c_1 being the rate of transition from S_1 to S_2 and c_2 the rate of transition from S_2 to S_1. Let us define $P_{ij}(t) = P\{Z(t) = j | Z(0) = i\}$; this is the probability that a molecule observed at time $t = 0$ to be S_i will be observed to be an S_j molecule at a later time t. The 2×2 matrix $P(t)$ is the transition probability matrix for the Markov process $Z(t)$ and from elementary Markov process theory [24] it follows that the off diagonal terms are given by

$$P_{ij}(t) = \frac{c_j}{c_1 + c_2} + \frac{c_i}{c_1 + c_2} e^{-(c_1+c_2)t}, \quad i \neq j,$$

and the diagonal terms are given by the condition that the rows sum to 1.

Suppose we have a system in which at time $t = 0$ there are a total of N_1 number of S_1 molecules and N_2 number of S_2 molecules. For $i = 1,2$; $j = 1,2$, denote by $Y_{ij}(t)$ the number of S_i molecules observed at time $t = 0$ that are observed to be S_j molecules at time t. Note that by definition $Y_{i1}(t) + Y_{i2}(t) = N_i$ for $i = 1,2$. It follows that $Y_{ij}(t)$ is the sum of N_i independent identically distributed random variables each of which takes values 0 and 1 with probabilities $1 - P_{ij}(t)$ and $P_{ij}(t)$, respectively. Thus

$$Y_{ij}(t) \sim \mathscr{B}(N_i, P_{ij}(t)),$$

where $\mathscr{B}(N,p)$ denotes the binomial distribution with parameters N and p [24]. It is also clear that $Y_{1i}(t)$ and $Y_{2j}(t)$ are independent for all choices of i, j. Let us denote by $X_i(t)$ for $i = 1,2$ the total number of S_i molecules at time t. Then it follows that $X_1(t) = Y_{11}(t) + Y_{21}(t)$ at any time t, and thus is the sum of two independent binomial random variables. Note that by definition $X_1(0) = N_1$ and $X_2(0) = N_2$. Also note that

$$X_1(t) + X_2(t) = X_1(0) + X_2(0) = N_1 + N_2 = x_T,$$

x_T being the total number of molecules.

Appendix B. Computation of Markov transition matrices for leap methods

All leap methods studied in literature so far are methods that can be described by the general functional form

$$\widehat{X}(t+\tau) = F(x, \widehat{K}(x,\tau)),$$

where $x = \widehat{X}(t)$,

$$\widehat{K}(x,\tau) = (\widehat{K}_1(x,\tau), \ldots, \widehat{K}_M(x,\tau))$$

are nonnegative integer-valued random variables whose distributions depend on x and τ, and F is some function that may sometimes be implicitly defined (as in the case of the implicit and trapezoidal implicit tau methods). For all the explicit methods (including the REMM-τ), $F(x,k)$ is simply

$$F(x) = x + \sum_{j=1}^{M} v_j k_j.$$

If the leaping method is applied with a fixed constant time step τ, then $\widehat{X}(n\tau)$ for $n \in \mathbb{N}$ forms a discrete time Markov process whose transition function $\widehat{P}(x,x';\tau)$ defined by

$$\widehat{P}(x,x';\tau) = P(\widehat{X}(t+\tau) = x' | \widehat{X}(t) = x)$$

may be computed from the distributions of $\widehat{K}_j(x,\tau)$. The reader may refer to [24,25] for general background on Markov processes. Since $\widehat{K}_j$ take nonnegative integer values it follows:

$$\widehat{P}(x,x';\tau) = \sum_{k \in \mathbb{Z}_+^M, x' = \mathbb{F}(x,k)} P(\widehat{K}_1(x,\tau) = k_1, \ldots, \widehat{K}_M(x,\tau) = k_M),$$

where $k = (k_1, \ldots, k_M)$. Thus in principle for any given $x \in \mathbb{Z}_+^N$ and $x' \in \mathbb{Z}_+^N$ (and given τ) knowing the joint probabilities

$$P(\widehat{K}_1(x,\tau) = k_1, \ldots, \widehat{K}_M(x,\tau) = k_M)$$

and F (even implicitly) allows one to compute the $\widehat{P}(x,x';\tau)$. For the parallely updated leap methods (this includes all methods studied in this paper except for the sequentially updated REMM-τ), $K_1(x,\tau), \ldots, K_M(x,\tau)$ are independent random variables and hence

$$P(\widehat{K}_1 = k_1, \ldots, \widehat{K}_M = k_M) = P(\widehat{K}_1 = k_1) P(\widehat{K}_2 = k_2) \cdots P(\widehat{K}_M = k_M).$$

Since $K_j(x,\tau)$ are either Poisson or binomial random variables, the probabilities $P(K_j = k_j)$ are easy to compute. For a given j, if K_j is a binomial random variable the summation over k_j is finite and can be performed

exactly. When K_j is Poisson the summation is infinite. However since the Poisson probabilities $P(K_j = k_j)$ decay for k_j much larger than the mean value of K_j, we may approximate the sum by a suitable truncation.

In general the transition function $\widehat{P}(x, x'; \tau)$ has noncompact support. For chemical systems that are closed such as Sections 4.1–4.3, for a given deterministic initial condition the transition function may be effectively reduced to have compact support, and thus to a finite matrix. When the total number of feasible states is small (such as in Sections 4.1 and 4.3) then the computation of this matrix is quite inexpensive.

References

[1] M. Elowitz, A. Levine, E. Siggia, P. Swain, Stochastic gene expression in a single cell, Science 297 (August) (2002) 1183–1186.
[2] J. Hasty, J. Pradines, M. Dolnik, J. Collins, Noise-based switches and amplifiers for gene expression, PNAS 97 (2000) 2075–2080.
[3] A. Arkin, J. Ross, H. McAdams, Stochastic kinetic analysis of the developmental pathway bifurcation in phase λ-infected *Escehrichia coli* cells, Genetics 149 (4) (1998) 1633–1648.
[4] D.T. Gillespie, A rigorous derivation of the chemical master equation, Physica A 188 (1992) 404–425.
[5] N.V. Kampen, Stochastic Processes in Physics and Chemistry, Elsevier Science Publishing Company, 1992.
[6] D.T. Gillespie, A general method for numerically simulating the stochastic time evolution of coupled chemical reactions, J. Comput. Phys. 22 (1976) 403–434.
[7] D.T. Gillespie, Exact stochastic simulation of coupled chemical reactions, J. Phys. Chem. 81 (1977) 2340–2361.
[8] B. Munsky, M. Khammash, The finite state projection algorithm for the solution of the chemical master equation, J. Chem. Phys. 124 (2006) 044144.
[9] D.T. Gillespie, Approximate accelerated stochastic simulation of chemically reacting systems, J. Chem. Phys. 115 (4) (2001) 1716–1733.
[10] M. Rathinam, L. Petzold, Y. Cao, D. Gillespie, Stiffness in stochastic chemically reacting systems: The implicit tau-leaping method, J. Chem. Phys. 119 (24) (2003) 12784–12794.
[11] Y. Cao, L. Petzold, M. Rathinam, D. Gillespie, The numerical stability of leaping methods for stochastic simulation of chemically reacting systems, J. Chem. Phys. 121 (24) (2004) 12169–12178.
[12] A. Chatterjee, D. Vlachos, M. Katsoulakis, Binomial distribution based τ-leap accelerated stochastic simulation, J. Chem. Phys. 122 (2005) 024112.
[13] T. Tian, K. Burrage, Binomial leap methods for simulating stochastic chemical kinetics, J. Chem. Phys. 121 (21) (2004) 10356–10364.
[14] M. Rathinam, L. Petzold, Y. Cao, D. Gillespie, Consistency and stability of tau leaping schemes for chemical reaction systems, SIAM J. MMS 4 (3) (2005) 867–895.
[15] E. Haseltine, J. Rawlings, Approximate simulation of coupled fast and slow reactions for stochastic chemical kinetics, J. Chem. Phys. 117 (15) (2002) 6959.
[16] C. Rao, A. Arkin, Stochastic chemical kinetics and the quasi-steady-state assumption: application to the Gillespie algorithm, J. Chem. Phys. 118 (11) (2003) 4999–5010.
[17] Y. Cao, D. Gillespie, L. Petzold, Multiscale stochastic simulation algorithm with stochastic partial equilibrium assumption for chemically reacting systems, J. Comput. Phys. 206 (2) (2005) 395–411.
[18] Y. Cao, D. Gillespie, L. Petzold, The slow-scale stochastic simulation algorithm, J. Chem. Phys. 122 (1) (2005) 014116.
[19] J. Goutsias, Quasiequilibrium approximation of fast reaction kinetics in stochastic biochemical systems, J. Chem. Phys. 122 (18) (2005) 184102.
[20] L. Segel, M. Slemrod, The quasi-steady-state assumption: a case study in perturbation, SIAM Rev. 31 (3) (1989) 446–477.
[21] D. Givon, R. Kupferman, A. Stuart, Extracting macroscopic dynamics: model problems and algorithms, Nonlinearity 17 (2004) 55–127.
[22] D.T. Gillespie, L. Petzold, Improved leap-size selection for accelerated stochastic simulation, J. Chem. Phys. 119 (2003) 8229–8234.
[23] Y. Cao, L. Petzold, Trapezoidal tau-leaping formula for the stochastic simulation of biochemical systems, in: Proceedings of Foundations of Systems Biology in Engineering (FOSBE 2005), 2005, pp. 149–152.
[24] S. Ross, Introduction to Probability Models, eighth ed., Academic Press, 2002.
[25] D.T. Gillespie, Markov Processes: An Introduction for Physical Scientists, Academic Press, 1992.
[26] U.M. Ascher, L.R. Petzold, Computer Methods for Ordinary Differential Equations and Differential-Algebraic Equations, SIAM, 1998.
[27] J. Percus, Small population effects in stochastic population dynamics, Bull. Math. Biol. 67 (2005) 1173–1194.
[28] D. McQuarrie, Kinetics of small systems I, J. Chem. Phys. 38 (2) (1963) 433–436.
[29] K.D. Cock, X. Zhang, M.F. Bugallo, P. Djuric, Comment on stiffness in stochastic chemically reacting systems: the implicit tau-leaping method, J. Chem. Phys. 121 (7) (2004) 3347–3348.
[30] C. Jachimowski, M. Russell, Set-theoretic approach to reaction kinetics, J. Chem. Phys. 42 (8) (1965) 2819–2822.