

Integral tau methods for stiff stochastic chemical systems

Yushu Yang,^{a)} Muruhan Rathinam,^{b)} and Jinglai Shen^{c)}

Department of Mathematics and Statistics, University of Maryland, Baltimore County, Baltimore, Maryland, 21250 USA

(Received 10 June 2010; accepted 8 December 2010; published online 28 January 2011)

Tau leaping methods enable efficient simulation of discrete stochastic chemical systems. Stiff stochastic systems are particularly challenging since implicit methods, which are good for stiffness, result in noninteger states. The occurrence of negative states is also a common problem in tau leaping. In this paper, we introduce the implicit Minkowski–Weyl tau (IMW- τ) methods. Two updating schemes of the IMW- τ methods are presented: implicit Minkowski–Weyl sequential (IMW-S) and implicit Minkowski–Weyl parallel (IMW-P). The main desirable feature of these methods is that they are designed for stiff stochastic systems with molecular copy numbers ranging from small to large and that they produce integer states without rounding. This is accomplished by the use of a split step where the first part is implicit and computes the mean update while the second part is explicit and generates a random update with the mean computed in the first part. We illustrate the IMW-S and IMW-P methods by some numerical examples, and compare them with existing tau methods. For most cases, the IMW-S and IMW-P methods perform favorably. © 2011 American Institute of Physics. [doi:10.1063/1.3532768]

I. INTRODUCTION

Chemical reactions occurring at the intracellular level often involve certain molecular species present only in small copy numbers. Such systems are best described by a discrete state and continuous in time Markov process model where the components of the state vector are integers that describe the nonnegative copy number of the different molecular species.^{1–3} Probabilistically correct realizations of sample paths of such systems can be generated by the *stochastic simulation algorithm* (SSA).^{1,2} It also follows that the probability distribution as a function of time satisfies the *chemical master equation* (CME).³

When the copy numbers of all the molecular species are very large, such systems behave nearly deterministically. In the large copy number limit, the chemical reaction systems can be modeled by the familiar *reaction rate equations* (RRE) which are ordinary differential equations (ODEs). The transition from the discrete stochastic model to the continuous and deterministic model is explained in Ref. 4. A rigorous derivation using the law of large numbers and a correction using the central limit theorem may be found in Refs. 5 and 6. This limiting behavior is known as the *thermodynamic limit* or the *fluid limit*.

Numerical simulation of such stochastic chemical systems falls into two broad categories. One approach is to directly compute the probabilities via the CME. This is often prohibitive due to the fact that the number of possible states grows exponentially with the number of distinct molecular species. Nevertheless methods have been devised to improve the efficiency of these computations.^{7,8}

The second approach is to generate sample trajectories via SSA. This approach does not suffer from an exponential growth in complexity with increase in the number of species. However, even this approach is computationally intensive in many practical examples as the reaction events are often too many. One major reason is stiffness, which is the presence of multiple time scales. Another reason is the presence of some species in large copy numbers. Approximate methods have been devised to speed up SSA. These fall into two classes. One being the tau leap methods which are analogous to the time stepping methods such as Runge–Kutta for ODEs and is the subject of this paper. The second approach is inspired by singular perturbation techniques. When there is a clear and vast separation between the time scales of a fast group of reactions and those of the other (slow) reactions, these methods are most appropriate. The slow-scale SSA, partial equilibrium approach, nested SSA, and the quasi-steady-state approach belong to this category.^{9–12} In this context, a comprehensive rigorous framework utilizing the functional law of large numbers and functional central limit theorem to obtain various approximations may be found in Ref. 13.

The tau leap methods involve advancing the system trajectory by leaping over several reaction events at each time step. Since the probability distribution for the number of reaction events is generally not known, this involves utilizing some criteria to generate suitable approximations. Examples of tau leap methods in literature include the explicit tau,¹⁴ the implicit tau,¹⁵ the trapezoidal implicit tau,¹⁶ and the REMM tau¹⁷ to name a few. The explicit tau method uses the simplest approximation criterion in that it freezes the propensities (probabilistic rates) of all the reaction events over the interval of the time step, leading to the result that the number of firings are independent Poissons. In the fluid limit (i.e., in the large copy number limit), the explicit tau method becomes the well known explicit Euler method for ODEs. Some variations on

^{a)}Electronic mail: yushu1@umbc.edu.

^{b)}Electronic mail: muruhan@umbc.edu.

^{c)}Electronic mail: shenj@umbc.edu.

the explicit tau where Poisson random variables are replaced by binomial random variables may be found in Refs. 18 and 19. A higher order accurate explicit tau method may be found in Ref. 20. The implicit and the trapezoidal implicit tau methods were developed in order to deal with stiff systems which are ubiquitous in chemical kinetics and in the fluid limit they become the well known implicit Euler and trapezoidal methods for ODEs. However, these tau methods suffer from the fact that they do not produce integer states. The *reversible equivalent monomolecular tau* (REMM tau) method was proposed to overcome this difficulty. However, in the fluid limit the behavior of REMM tau is not fully understood as it yields an unknown method. Error analysis of tau methods may be found in Refs. 21–23.

Two common issues with tau leap methods have been the occurrence of negative or noninteger states. The noninteger states can be rounded to yield integers, but when the numbers are small this results in unacceptable errors.¹⁷ The negative states, when they occur, can be reset to suitable nonnegative states and the error involved depends on the probability of occurrence of negative states.

In this paper, we propose two variants of a new tau leap method which takes into account four key issues: stiffness, integrality and nonnegativity of states, and the behavior at fluid limit. To obtain desirable behavior for stiff systems in the fluid limit, we aim to design the tau leap method to yield the implicit Euler as its fluid limit. This is accomplished by the use of a split step. The first part of the step involves an implicit Euler step to compute the mean update. The second part involves generating random variables with the mean computed in the first part. To deal with negativity, we use the *Minkowski–Weyl decomposition* to describe the polyhedral region in the reaction count space that corresponds to the set of feasible reaction counts. We have not found a method that addresses all the issues in a completely satisfactory manner while remaining computationally tractable. The methods we propose in this paper reflect a compromise among these various issues. Both methods are called *implicit Minkowski–Weyl tau* method (IMW- τ) and involve partitioning the set of reactions into groups in such a manner that the Minkowski–Weyl decomposition is always carried out in a one- or two-dimensional space. One variant is the *implicit Minkowski–Weyl sequential tau* (IMW-S) and the other is the *implicit Minkowski–Weyl Parallel tau* (IMW-P). The IMW-S tau method produces integer and nonnegative states and remains stable for stiff systems. However, in the fluid limit, it becomes the *sequentially updated implicit Euler*, which suffers certain drawbacks when applied to stiff systems. The IMW-P tau method also produces integer states and in the fluid limit becomes implicit Euler. However, it suffers from the fact that it has nonzero probability of producing negative states and hence a bounding procedure is used. Additionally both methods IMW-S and IMW-P are designed to be first order consistent.

The outline of the paper is as follows. We review stochastic chemical kinetics and some of the existing tau leap methods and discuss concerning issues in Sec. II. In Sec. III, we provide a description of the Minkowski–Weyl decompo-

sition of convex polyhedral regions and motivate the general approach behind the proposed IMW- τ methods. Section IV describes all the different types of feasible regions in one and two dimensions that are relevant for the IMW- τ methods. The two IMW- τ methods proposed are described in detail in Sec. V. In Sec. VI, we provide numerical examples to illustrate these methods. Conclusions are presented in Sec. VII.

II. OVERVIEW OF STOCHASTIC CHEMICAL SYSTEMS AND TAU LEAPING METHODS

A. Stochastic chemical model and SSA

Stochastic chemical reaction systems involved with small number of molecules have a dynamic behavior that is discrete and stochastic rather than continuous and deterministic. We describe the standard well-stirred chemical model here,^{3,24} which is a Markov process in continuous time with state space \mathbb{Z}_+^N , the set of nonnegative integer vectors.

Suppose there is a well-stirred mixture of N molecular species $\{S_1, \dots, S_N\}$ interacting through M chemical reaction channels $\{R_1, \dots, R_M\}$. The state of the system is described by $[X_1(t), \dots, X_N(t)]$, where $X_i(t)$ is the number of molecules S_i at time t . For each $j = 1, \dots, M$, $a_j(x)h + o(h)$ is the probability, given $X(t) = x$, that reaction R_j will occur during $(t, t + h]$, where $a_j(x)$ is the *propensity function* of the reaction channel R_j . Vector v_j for $j = 1, \dots, M$ is the *stoichiometric vector*, whose i th component v_{ij} is the change in the number of S_i molecules produced by one occurrence of reaction R_j . Since $X(t)$ is a continuous time Markov process on a multidimensional integer lattice, it can be simulated exactly by the SSA.^{1,2}

B. Thermodynamic or fluid limit

When all the molecular species are present in large numbers and under certain additional assumptions on the propensity functions, the chemical system is well approximated by a deterministic ODE model known as the RRE. This equation can be thought of as a limit in which the system volume V approaches ∞ with the initial number of species $X(0)$ also growing proportional to V thus keeping the concentration $X(0)/V$ fixed.^{6,25} This limit is known as the *thermodynamic limit* in chemical literature and is also known as the *fluid limit* in queuing theory.

We describe the fluid limit in mathematical terms following Ethier and Kurtz.⁶ Consider a system with initial state $X(0) = x_0 \in \mathbb{Z}_+^N$ and volume V_0 . Denote by $z_0 = x_0/V_0$, the initial concentration. Let us consider a family of related systems with different volumes V and initial states $X_V(0) = Vz_0 = (V/V_0)x_0$, so that they have the same initial concentration. Let the solution trajectory for system with volume V be denoted by $X_V(t)$. Note that our original system has a trajectory $X(t) = X_{V_0}(t)$. Thus the concentrations are $Z_V(t) = X_V(t)/V$. Additionally we assume that the propensities $a_j(x)$ depend on volume V in such a manner that as $V \rightarrow \infty$, $a_j(x, V)/V$ approaches a limit $\bar{a}_j(x)$, which is true in the standard model of stochastic chemical kinetics. It is shown in Ref. 6 (theorem on page 456) that for each fixed

$t \geq 0$, $Z_V(t)$ converges with probability 1 to the deterministic quantity $\bar{Z}(t)$ which is the unique solution of RRE

$$\dot{\bar{Z}}(t) = \sum_{j=1}^M v_j \bar{a}_j[\bar{Z}(t)], \quad (1)$$

with initial condition $\bar{Z}(0) = z_0$. See Appendix D for more details.

Since we are considering systems with both large and small number of molecules in this paper, it is important that the tau leap methods proposed also have appropriate thermodynamic or fluid limit. In this context the work in Ref. 23 provides an error analysis of certain explicit tau leap methods in the large volume regime.

C. Tau leaping methods

The SSA is very computationally expensive because it simulates one reaction event each time. The tau leaping methods¹⁴⁻¹⁷ were proposed to accelerate the chemical reaction simulations. The principle of tau leaping methods is to propose a time step τ and leap over a number of reactions with a reasonable loss of accuracy.

Mathematically, the tau leaping methods proceed as follows. First, a time step τ is chosen. Given $X(t) = x$, define $R_j(x, \tau)$ to be the (random) number of times that j th reaction channel will fire during the time interval $(t, t + \tau]$, for $j = 1, \dots, M$. Then

$$X(t + \tau) = x + \sum_{j=1}^M v_j R_j(x, \tau). \quad (2)$$

In general, the distribution of $R_j(x, \tau)$ is not known. In a tau leap method, an approximation $K_j(x, \tau)$ of $R_j(x, \tau)$ is computed. The *explicit tau* method¹⁴ chooses $K_j(x, \tau)$ for $j = 1, \dots, M$ to be independent Poisson random variables with mean $a_j(x)\tau$, i.e., $K_j^{(et)}(x, \tau) \sim \mathcal{P}(a_j(x)\tau)$, where $\mathcal{P}(\lambda)$ denotes a Poisson random variable with mean λ .

The *implicit tau* method¹⁵ is given by

$$X^{(it)}(t + \tau) = x + \sum_{j=1}^M v_j \{P_j - a_j(x)\tau + a_j[X^{(it)}(t + \tau)]\tau\}, \quad (3)$$

where $P_j \sim \mathcal{P}(a_j(x)\tau)$ for $j = 1, \dots, M$ are independent. Thus $R_j(x, \tau)$ is approximated by

$$K_j^{(it)}(x, \tau) = P_j - a_j(x)\tau + a_j[X^{(it)}(t + \tau)]\tau.$$

Newton's method is applied to solve Eq. (3). Note that $X^{(it)}(t + \tau)$ is not an integer vector any more and $K_j^{(it)}(x, \tau)$ is not an integer either. This does not make physical sense for a chemical reaction system. One way to avoid noninteger states is by modifying the implicit tau method, which yields the following *rounded implicit tau* method: First, solve $X' = X^{(it)}(t + \tau)$ according to Eq. (3). Then approximate the number of firings $R_j(x, \tau)$ by the integer-valued random variable $K_j^{(itr)}(x, \tau)$, defined by $K_j^{(itr)}(x, \tau) = [K_j^{(it)}(x, \tau)]$ and update

$$X^{(itr)}(t + \tau) = x + \sum_{j=1}^M v_j K_j^{(itr)}. \quad (4)$$

Here $[z]$ denotes the nearest nonnegative integer corresponding to a real number z .

The *trapezoidal implicit tau* method¹⁶ generates the update equation by

$$X^{(tr)}(t + \tau) = x + \sum_{j=1}^M v_j \left(P_j - \frac{\tau}{2} a_j(x) + \frac{\tau}{2} a_j[X^{(tr)}(t + \tau)] \right), \quad (5)$$

where $P_j \sim \mathcal{P}(a_j(x)\tau)$ are independent. Thus $R_j(x, \tau)$ is approximated by

$$K_j^{(tr)}(x, \tau) = P_j - \frac{\tau}{2} a_j(x) + \frac{\tau}{2} a_j[X^{(tr)}(t + \tau)].$$

It still gives noninteger states for both $X^{(tr)}(t + \tau)$ and $K_j^{(tr)}(x, \tau)$. The *rounded trapezoidal implicit tau* solves $X' = X^{(tr)}(t + \tau)$ from Eq. (5) and approximates the actual number of firings $R_j(x, \tau)$ by $K_j^{(trr)}(x, \tau) = [K_j^{(tr)}(x, \tau)]$. It updates states by

$$X(t + \tau) = x + \sum_{j=1}^M v_j K_j^{(trr)}(x, \tau). \quad (6)$$

The REMM tau¹⁷ is an explicit leaping scheme based on the exact solutions of the two prototypes of reversible monomolecular reactions $S_1 \leftrightarrow S_2$ and $S \leftrightarrow 0$. This method approximates all bimolecular reversible reaction pairs by suitable monomolecular reversible reactions and then updates the system based on the exact solutions of these monomolecular reversible pairs. The REMM tau is stated in parallel and sequential forms, both generate integer-valued states with Poisson and binomial random variables. The sequential version of REMM tau avoids nonnegative states without any bounding procedures. It has been shown¹⁷ that the REMM tau exhibits a more robust performance than the implicit tau and trapezoidal tau methods for “small number and stiff” problems because of the inaccuracies in the latter methods due to rounding.

Ideally a tau leap method should “naturally” produce nonnegative and integer states while maintaining a robust performance when applied to stiff systems. Additionally in the thermodynamic or fluid limit, we wish the tau leap method to behave like a “good stiff” solver for ODEs. In this context, we note that the fluid limit of the explicit tau is the explicit Euler while for the implicit tau it is the implicit Euler. The fluid limit of REMM tau is not a known ODE solver for ODE systems, and as such its performance in the fluid limit for stiff systems is yet to be investigated thoroughly. On the other hand, it is advantageous to devise a tau leap method that has implicit Euler as its fluid limit since the robust behavior of implicit Euler for stiff ODE systems is well established. Table I summarizes some properties of the methods discussed here. In Sec. III, we propose a new framework and new tau methods motivated by addressing these issues.

TABLE I. Comparison of the existing tau leaping methods: explicit tau, implicit tau, trapezoidal tau, and REMM tau. The star (*) represents that the fluid limit of REMM tau is not a known ODE solver.

Methods/Issues	Fluid limit	Integer states	Nonnegativity
Explicit tau	Explicit Euler	YES	NO
Implicit tau	Implicit Euler	NO	NO
Trapezoidal tau	Trapezoidal Euler	NO	NO
REMM tau (parallel)	*	YES	NO
REMM tau (sequential)	*	YES	YES

III. GENERAL FRAMEWORK AND MOTIVATIONS

The central decision at each step in a tau leap update is the choice of a joint distribution for $K = (K_1, \dots, K_M)^T$. We would like a distribution that satisfies the following conditions:

1. K_j satisfies $O(\tau)$ consistency [meaning errors in one step are $O(\tau^2)$].
2. K_j is integer valued and nonnegative.
3. $x + \nu K \geq 0$ with probability 1.
4. The fluid limit of the resulting method is the implicit Euler.
5. Generating samples for K should be computationally tractable.

The first condition is important to ensure that making the step size smaller guarantees greater accuracy. The second and third conditions ensure integer and nonnegative states, while the fourth condition ensures stable behavior at least in the fluid limit in the case of stiff systems.

Implicit step and the fluid limit of the method: In order to ensure that the tau method in the fluid limit becomes implicit Euler, we use a split step approach where the first part involves computing an intermediate state X' using the implicit Euler:

$$X' = x + \sum_{j=1}^M \nu_j a_j(X')\tau. \quad (7)$$

Then we choose an integer-valued distribution for K such that $E(K_j) = a_j(X')\tau$. Heuristically, in the fluid limit, since K_j will be nearly deterministic, $K_j \approx E(K_j)$ and the updated state $X \approx X'$. See Appendix D where this is discussed in detail.

The Minkowski–Weyl decomposition: The major idea we propose in dealing with negative states is to have a convenient description of the region in K space, i.e., the region in the reaction count vector space, that corresponds to nonnegative integer values for the updated states. We use the *Minkowski–Weyl decomposition* in the description of this region.²⁶ Thus we shall use the term Implicit Minkowski–Weyl tau or IMW- τ in short to describe the family of methods proposed in this paper.

In a single step of the tau method, the following linear inequality is obtained by the nonnegativity of the population state:

$$X = x + \nu K \geq 0, \quad (8)$$

where $x \in \mathbb{Z}_+^N$ and $\nu \in \mathbb{Z}^{N \times M}$ are given and $K \in \mathbb{Z}_+^M$ is the random unknown vector of reaction counts. Note that throughout this paper we write $X \geq 0$ for a vector X to mean that each component is greater than or equal to 0. Let $\mathcal{P} = \{K \in \mathbb{Z}_+^M \mid K \geq 0, x + \nu K \geq 0\}$ be the set of physically feasible values of K such that the resulting state X is nonnegative. We wish to have a convenient description of the set \mathcal{P} . Relaxing the domain of the set \mathcal{P} from \mathbb{Z}_+^M to \mathbb{R}^M , we obtain $\mathcal{F} = \{K \in \mathbb{R}^M \mid K \geq 0, x + \nu K \geq 0\}$ which is a convex polyhedral region.

Two examples are shown here to illustrate the Minkowski–Weyl decomposition. The first example is the reversible monomolecular reaction $S_1 \leftrightarrow S_2$. Equation (8) gives

$$x_1 - K_1 + K_2 \geq 0, \quad x_2 + K_1 - K_2 \geq 0,$$

where

$$(K_1, K_2)^T \geq 0.$$

The feasible region of K values are shown by the shaded region in Fig. 1. We note that the feasible region is a convex polyhedral region and any point in it can be expressed as the sum of two vectors, one representing a point inside the triangle with vertices $(0, 0)^T$, $(x_1, 0)^T$, $(0, x_2)^T$, and the other a vector that is a positive multiple of $(1, 1)^T$. In general, the Minkowski–Weyl theorem states that any point in a convex polyhedron (in a finite dimensional space) can be represented as the sum of a point in a convex hull formed by finite number of points (known as extreme points or vertices) and a vector in a positive cone spanned by finite number of direction vectors. In this example, the convex hull is the triangle and the positive cone is the set of all vectors that are positive multiples of $(1, 1)^T$. Any point inside the

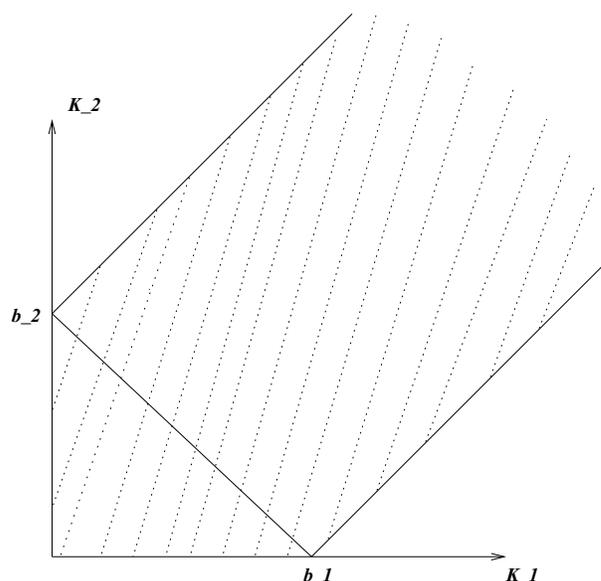


FIG. 1. Type 1 feasible region of K values are shown by the shaded region, which consists of a convex hull and a positive cone. The convex hull is the triangle with vertices $(0, 0)^T$, $(b_1, 0)^T$, and $(0, b_2)^T$, and the positive cone is the set of all vectors that are all positive multiples of $(1, 1)^T$.

triangle is of the form $(0, 0)^T \alpha_0 + (x_1, 0)^T \alpha_1 + (0, x_2)^T \alpha_2$ with $\alpha_0 + \alpha_1 + \alpha_2 = 1$ and $\alpha_i \geq 0$. Since $(0, 0)^T \alpha_0$ is the zero vector we may omit that term. Thus a point $K = (K_1, K_2)^T$ in the feasible region can be expressed by

$$\begin{pmatrix} K_1 \\ K_2 \end{pmatrix} = \begin{pmatrix} x_1 \\ 0 \end{pmatrix} \alpha_1 + \begin{pmatrix} 0 \\ x_2 \end{pmatrix} \alpha_2 + \begin{pmatrix} 1 \\ 1 \end{pmatrix} \beta,$$

where $\alpha_1 + \alpha_2 \leq 1$, $\alpha_1 \geq 0$, $\alpha_2 \geq 0$, and $\beta \geq 0$. It must be noted that the mapping from α and β to K is not one to one. But it is onto the set of all K values in the feasible region. We finally observe, that in this two-dimensional example the Minkowski–Weyl decomposition was easy to obtain from visual observation.

Next we consider the three-dimensional example $0 \rightarrow S_1 \rightarrow S_2 \rightarrow 0$. We obtain the following inequalities for K_j

$$x_1 + K_1 - K_2 \geq 0, \quad x_2 + K_2 - K_3 \geq 0,$$

where

$$(K_1, K_2, K_3)^T \geq 0.$$

Since the feasible region is a region in three dimensions, unlike the earlier example, it is harder to visualize. Nevertheless the Minkowski–Weyl theorem asserts the existence of a similar decomposition of the feasible region into the sum of a convex hull and a positive cone. Additionally, there exists an algorithm to calculate the vertices of the convex hull and a set of vectors that span the positive cone.²⁷

In this case this algorithm shows (we do not show this calculation here as we shall always work in one or two dimensions) that the convex hull has vertices $(0, 0, 0)^T$, $(0, x_1, 0)^T$, $(0, 0, x_2)^T$, and $(0, x_1, x_1 + x_2)^T$, and the cone is formed by the positive linear combination of directions $(1, 1, 0)^T$, $(1, 1, 1)^T$, and $(1, 0, 0)^T$. Thus any feasible $K = (K_1, K_2, K_3)^T$ can be written as

$$\begin{pmatrix} K_1 \\ K_2 \\ K_3 \end{pmatrix} = \begin{pmatrix} 0 \\ x_1 \\ 0 \end{pmatrix} \alpha_1 + \begin{pmatrix} 0 \\ 0 \\ x_2 \end{pmatrix} \alpha_2 + \begin{pmatrix} 0 \\ x_1 \\ x_1 + x_2 \end{pmatrix} \alpha_3 \\ + \begin{pmatrix} 1 \\ 1 \\ 0 \end{pmatrix} \beta_1 + \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} \beta_2 + \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \beta_3,$$

where $\alpha_1 + \alpha_2 + \alpha_3 \leq 1$, $\alpha_i \geq 0$, $\alpha_2 \geq 0$, $\alpha_3 \geq 0$, and $\beta_1 \geq 0$, $\beta_2 \geq 0$, $\beta_3 \geq 0$.

It is clear from the discussion above that in general, by virtue of the Minkowski–Weyl theorem there exist matrices B and D such that, $K \in \mathcal{F}$ if and only if K can be written in the form,

$$K = B\alpha + D\beta,$$

where α and β are arbitrary real vectors subject to the conditions $\alpha \geq 0$, $\mathbf{1}^T \alpha \leq 1$, and $\beta \geq 0$, where $\mathbf{1}$ is the (column) vector whose components are all 1. The conditions on α take this particular form because as in the two examples above, the origin in K space always forms one extreme point of the convex hull associated with the decomposition. Thus B is the

matrix whose column vectors are the extreme points (except the origin), and the columns of the matrix D form the extreme directions that span the positive cone.

Additionally we make an important observation. It can be proven that if we scale the initial state x by a scalar $V > 0$, the feasible region \mathcal{F} changes in such a way that the resulting new convex hull is simply the original convex hull scaled by V , while the positive cone remains unchanged. Thus under the scaling of x by $V > 0$, the matrix B is $O(V)$ while D is independent of V .

Unfortunately the complexity of the computation of Minkowski–Weyl decomposition increases rapidly with the dimensionality of K . This motivates us to limit our algorithms to only work in one or two dimensions of K space at a time. This can be done by partitioning the set of reactions into groups of one or two. We shall provide details of the partitioning method in Sec. V and details of one- and two-dimensional feasible regions in Sec. IV.

Distributions for α and β : Having obtained such a decomposition for K , our problem is transformed into finding suitable distributions for α and β . It must be noted that the number of components of α and β together are in general more than the number of components of K . However, unlike the components of K the components of α and β are subject to the simpler constraints $\alpha \geq 0$, $\mathbf{1}^T \alpha \leq 1$, and $\beta \geq 0$.

Let us denote $E(K) = \lambda$. If $K \in \mathcal{F}$ with probability 1, then by convexity of \mathcal{F} it follows that $E(K) = \lambda \in \mathcal{F}$. If $\lambda \in \mathcal{F}$, then there exist vectors p and q such that $\lambda = Bp + Dq$ with $p \geq 0$, $\mathbf{1}^T p \leq 1$, and $q \geq 0$. Additionally, taking expectations on both sides of the relation $K = B\alpha + D\beta$, we see that $p = E(\alpha)$ and $q = E(\beta)$. Thus the second step of the tau method involves finding p and q from λ . Given $\lambda = E(K)$ the choice of (p, q) is not unique. One may impose an additional constraint to make this choice unique. We shall describe this step in detail in Sec. IV. We have found that the complexity of the computations involved in finding p and q also increases rapidly with the dimensionality of K . This is yet another factor that motivates us to partition the reactions into groups that result in one- or two-dimensional regions.

Once p and q are chosen, the next step is to generate (vector valued) random variables α and β with respective means $E(\alpha) = p$ and $E(\beta) = q$. The conditions on the distribution of K laid down at the beginning of this section imply that the distribution for α and β must satisfy the following conditions:

1. The resulting distribution for K must satisfy $O(\tau)$ consistency condition.
2. The values of α and β are such that the resulting values of $K = B\alpha + D\beta$ encompass all integer values in \mathcal{F} .
3. $\mathbf{1}^T \alpha \leq 1$, $\alpha \geq 0$, and $\beta \geq 0$.
4. $\text{Cov}(\alpha) \rightarrow 0$ and $\text{Cov}(\beta) = O(V)$ as $V \rightarrow \infty$, when the initial state x is scaled according to $x = Vz$ keeping z constant. As explained in Appendix D this ensures that as $V \rightarrow \infty$ the updated state X becomes deterministic and thus equals the implicit Euler solution.
5. Generating a sample from the distribution must be computationally tractable.

We have not been able to find a “natural family of distributions” for the variables α and β such that $B\alpha + D\beta$ takes all the integer values in the convex polyhedron \mathcal{F} in arbitrary dimensions of K space. If we only look for *real valued distributions* taking values in the convex polytope spanned by columns of B , then there is a natural family of distributions, namely the *Dirichlet distributions*.²⁸ However, then $K = B\alpha + D\beta$ will not be integer any more, and we will have to round. Rounding results in errors that are difficult to study in the case of small numbers, and thus we abandon this approach.

In this paper, we propose two specific tau leap methods that follow from the above general approach. Both methods first partition the set of reactions into groups such that the Minkowski–Weyl decomposition is always carried out in a K space of dimension two or one. We shall use scaled binomial distributions for α and Poisson distributions for β . We note that this choice of these distributions is consistent with the above five conditions.

In the first method, each group of reactions is updated in a sequential manner. We call this method as the *implicit Minkowski–Weyl sequential* method and is described in detail in Sec. V A. Since nonnegativity of the intermediate state is maintained after each group is updated, this method guarantees nonnegativity of the updated next state. The major drawback of this method is that in the fluid limit it does not become the implicit Euler, but rather the “sequentially updated implicit Euler.” This latter scheme (which is almost never used to solve ODEs) is not as good a method as the implicit Euler is for stiff systems as we explain later. The second method attempts to rectify the drawback of the first method but as a compromise nonnegativity is no longer guaranteed. This method simultaneously updates each group of reactions *independently of the other groups*. This method may lead to negative states and when a negative state is encountered a bounding procedure is applied to obtain a “nearby” nonnegative state. This method is called the *implicit Minkowski–Weyl parallel* method and is described in Sec. V B.

IV. FEASIBLE REGIONS IN ONE AND TWO DIMENSIONS

In this section we describe all possible one-dimensional regions and all possible two-dimensional regions corresponding to (stoichiometrically) reversible pairs of reactions. For each type of feasible region, we describe the algorithm to generate a sample for K subject to the constraint $E(K) = \lambda$, where λ is assumed computed before.

A. One-dimensional feasible regions in K

For one-dimensional feasible region, the linear inequality condition on K falls into two categories: K is either bounded between 0 and a positive integer b , or K is unbounded and nonnegative. Our algorithm proceeds as follows. Let the state prior to update be x and suppose $\lambda = E(K)$ is given. If the feasible K values are bounded by an integer b , we choose $K \sim \mathcal{B}(b, p)$, where $p = \lambda/b$. If the feasible K values are

unbounded, we choose $K \sim \mathcal{P}(\lambda)$. We illustrate via some examples below.

Example 1, $S_1 \xrightarrow{c} S_2$: The updated state in this example follows the inequalities:

$$x_1 - K \geq 0, \quad x_2 + K \geq 0, \quad \text{where } K \geq 0.$$

Hence, $0 \leq K \leq x_1$. In terms of the Minkowski–Weyl decomposition, we may write $K = x_1\alpha$ with $0 \leq \alpha \leq 1$.

We choose K to be binomial bounded by $b = x_1$ with mean λ . Thus $K = x_1\alpha \sim \mathcal{B}(x_1, p)$, where $p = \lambda/x_1$.

The feasible region of K in this example is also applicable to the example $S_1 + S_2 \xrightarrow{c} S_3$. In this case K satisfies

$$x_1 - K \geq 0, \quad x_2 - K \geq 0, \quad x_3 + K \geq 0, \quad \text{where } K \geq 0.$$

Therefore, $0 \leq K \leq b = \min\{x_1, x_2\}$. Thus we may write $K = \min\{x_1, x_2\}\alpha$ with $0 \leq \alpha \leq 1$. We generate K according to $K \sim \mathcal{B}(\min\{x_1, x_2\}, p)$, where $p = \lambda/\min\{x_1, x_2\}$.

Example 2, $0 \xrightarrow{c} S_1$: This reaction stands for the production of a molecule S_1 . The inequalities for K are

$$x_1 + K \geq 0, \quad K \geq 0.$$

Here K has no upper bound. We choose $K \sim \mathcal{P}(\lambda)$.

It can be verified that the method described above satisfies $O(\tau)$ consistency. Moreover, it can also be verified that $\text{Cov}(K/V) \rightarrow 0$ to ensure the desired fluid limit. We do not show these calculations here, but we comment that they follow from reasoning similar to the ones given for the case of Type 1 region in two dimensions (see Appendix E and Sec. IV B).

B. Two-dimensional feasible region: Type 1

The Type 1 example of polyhedral region is the shaded region shown in Fig. 1. This region corresponds to the pair of inequalities

$$-b_2 \leq K_1 - K_2 \leq b_1.$$

The corresponding convex hull is the triangle with vertices $(0, 0)^T$, $(b_1, 0)^T$, and $(0, b_2)^T$, and the corresponding positive cone is the set of all vectors that are all positive multiples of $(1, 1)^T$. Here b_1 and b_2 depend on the initial state x . Thus the Minkowski–Weyl decomposition for K has the following form:

$$\begin{pmatrix} K_1 \\ K_2 \end{pmatrix} = \begin{pmatrix} b_1 & 0 \\ 0 & b_2 \end{pmatrix} \alpha + \begin{pmatrix} 1 \\ 1 \end{pmatrix} \beta. \quad (9)$$

Here $\alpha = (\alpha_1, \alpha_2)^T$, while β is scalar valued.

The simplest example with Type 1 region is the reversible monomolecular reaction given by $S_1 \leftrightarrow S_2$. This example was already discussed in Sec. III and in this case $b_1 = x_1$ and $b_2 = x_2$.

Type 1 region generally arises corresponding to reversible pairs of reactions where each reaction contributes to a decrease in at least one species. Here we describe a few common examples.

1. $S_1 + S_2 \leftrightarrow S_3$: The constraints on $K = (K_1, K_2)^T$ are given by

$$-x_3 \leq K_1 - K_2 \leq \min\{x_1, x_2\}.$$

Hence $b_1 = \min\{x_1, x_2\}$ and $b_2 = x_3$.

2. $S_1 + S_2 \leftrightarrow S_3 + S_4$: The constraints on $K = (K_1, K_2)^T$ are given by

$$-\min\{x_3, x_4\} \leq K_1 - K_2 \leq \min\{x_1, x_2\}.$$

Therefore, $b_1 = \min\{x_1, x_2\}$, $b_2 = \min\{x_3, x_4\}$.

3. $2S_1 \leftrightarrow S_2$: The constraints on $K = (K_1, K_2)^T$ are given by

$$-x_2 \leq K_1 - K_2 \leq x_1/2.$$

We obtain $b_1 = x_1/2$, $b_2 = x_2$. We note that, if we use $b_1 = \lfloor x_1/2 \rfloor$ (which denotes the largest integer less than or equal to $x_1/2$) instead of $b_1 = x_1/2$, we obtain a smaller feasible region. It can be shown that the ‘‘lost region’’ does not contain any integers and hence we do not miss any valid points.

4. $2S_1 \leftrightarrow 2S_2$: The constraints on $K = (K_1, K_2)^T$ are given by

$$-x_2/2 \leq K_1 - K_2 \leq x_1/2.$$

Here $b_1 = x_1/2$, and $b_2 = x_2/2$. As explained above, we may equivalently use $b_1 = \lfloor x_1/2 \rfloor$ and $b_2 = \lfloor x_2/2 \rfloor$.

5. $S_1 + S_2 \leftrightarrow S_1 + S_3$: Note that the number of species S_1 remains unchanged. The inequality conditions for $K = (K_1, K_2)^T$ are given by $\alpha = (\alpha_1, \alpha_2)^T$, and $p = (p_1, p_2)^T$ while β and q are scalar valued.

$$-x_3 \leq K_1 - K_2 \leq x_2.$$

Hence, $b_1 = x_2$ and $b_2 = x_3$.

We note that in each case, b scales with x in the following way: $b(Vx) = Vb(x)$ for all $V > 0$.

We now describe how the IMW- τ method is applied to a pair of reactions with Type 1 region. Suppose the number of molecules prior to updating this pair of reactions is x . First we compute b from x as described by the examples above. As before we denote $E(K) = \lambda$, $E(\alpha) = p$, and $E(\beta) = q$. Taking the mean of Eq. (9), we get

$$\begin{pmatrix} \lambda_1 \\ \lambda_2 \end{pmatrix} = \begin{pmatrix} b_1 & 0 \\ 0 & b_2 \end{pmatrix} p + \begin{pmatrix} 1 \\ 1 \end{pmatrix} q. \quad (10)$$

Here $p = (p_1, p_2)^T$ while q is scalar valued. The conditions on α and β are $\alpha_1 + \alpha_2 \leq 1$, $\alpha_1 \geq 0$, $\alpha_2 \geq 0$, and $\beta \geq 0$. Suppose $p_1 + p_2 = \bar{p}$. Combining with Eq. (10), we obtain a linear system

$$b_1 p_1 + q = \lambda_1, \quad b_2 p_2 + q = \lambda_2, \quad p_1 + p_2 = \bar{p}. \quad (11)$$

The conditions on α and β yield

$$\bar{p} \leq 1, \quad q \geq 0, \quad p_1 \geq 0, \quad p_2 \geq 0. \quad (12)$$

Note that this is an underdetermined system of equations for p_1 , p_2 , and q . First let us consider the case where both

b_1 and b_2 are nonzero. In this case one may verify that the feasible values of \bar{p} lie between upper bound \bar{p}_U and lower bound \bar{p}_L , given by

$$\begin{aligned} \bar{p}_U &= \min \left\{ 1, \frac{\lambda_1}{b_1} + \frac{\lambda_2}{b_2} \right\}, \\ \bar{p}_L &= \max \left\{ \frac{\lambda_2 - \lambda_1}{b_2}, \frac{\lambda_1 - \lambda_2}{b_1} \right\}. \end{aligned} \quad (13)$$

Our method chooses a combination of \bar{p}_U and \bar{p}_L given by

$$\bar{p} = r \bar{p}_L + (1 - r) \bar{p}_U, \quad r = r_0 [1 - e^{-(\lambda_1 + \lambda_2)}], \quad (14)$$

where we choose $r_0 = 0.5$ in our simulations. The term $1 - e^{-(\lambda_1 + \lambda_2)}$ ensures that r is $O(\tau)$, which in turn ensures the consistency of the method (see Appendix E).

After \bar{p} is computed, p_1 , p_2 , and q are computed by

$$\begin{aligned} q &= \frac{\lambda_1 b_2 + \lambda_2 b_1 - b_1 b_2 \bar{p}}{b_1 + b_2}, \\ p_1 &= \frac{\lambda_1 - q}{b_1}, \quad p_2 = \frac{\lambda_2 - q}{b_2}. \end{aligned} \quad (15)$$

Suppose $b_1 = 0$ then Eq. (11) has unique solutions for p_2 and q . One has freedom in the choice of p_1 , but it will not be used as $b_1 \alpha_1 = 0$ regardless of the choice of α_1 . Similar comment applies when $b_2 = 0$. If both $b_1 = b_2 = 0$, then the feasible region contains only the diagonal where $K_1 = K_2$, and the updated state is always x .

Having found p_1 , p_2 , and q the next step is to choose appropriate distributions for α_1 , α_2 , and β subject to the constraints that $E(\alpha_i) = p_i$, $E(\beta) = q$, $\alpha_i \geq 0$, $\beta \geq 0$, and $\alpha_1 + \alpha_2 \leq 1$. Since $b_i \alpha_i$ and β must be integer valued, $b_i \alpha_i$ is bounded, and β is unbounded, we pick $b_i \alpha_i$ to be binomial with parameters b_i and p_i , and β to be Poisson with parameter q . The reason for the choice of these distributions is motivated by the fact that they are relatively inexpensive to generate and also that their covariances scale in the appropriate way to provide the correct fluid limit as seen below. It is, however, difficult to ensure that $\alpha_1 + \alpha_2 \leq 1$. Given the geometry of the Type 1 region (see Fig. 1), allowing α_1, α_2 to be independent and taking values between 0 and 1 still results in points inside the polyhedral region. Thus, on the whole, we choose α_1, α_2 , and β to be independent and having distributions given by, $b_1 \alpha_1 \sim \mathcal{B}(b_1, p_1)$, $b_2 \alpha_2 \sim \mathcal{B}(b_2, p_2)$, and $\beta \sim \mathcal{P}(q)$, where \mathcal{B} and \mathcal{P} are binomial and Poisson random variables with parameters shown inside the parentheses. Then we set $K_1 = b_1 \alpha_1 + \beta$, $K_2 = b_2 \alpha_2 + \beta$, and the system is updated by $X = x + \nu K$.

The proof for $O(\tau)$ consistency of this update method is shown in Appendix E. In order to see that the fluid limit behaves appropriately, we need to verify that $\text{Cov}(K^V/V) \rightarrow 0$ when $V \rightarrow \infty$ as stated in Sec. III and Appendix D. We set $x = Vz$ and let $V \rightarrow \infty$ keeping z fixed. By the scaling property of the dependence of b on x , it follows that $b = O(V)$. First we note that as $V \rightarrow \infty$, the propensity function $a(V, zV)$ is $O(V)$ and hence λ is $O(V)$. From the expressions for \bar{p}_U , \bar{p}_L , and \bar{p} we can verify that \bar{p} is $O(1)$. Furthermore, Eq. (15) gives $q = O(V)$, $p_1 = p_2 = O(1)$

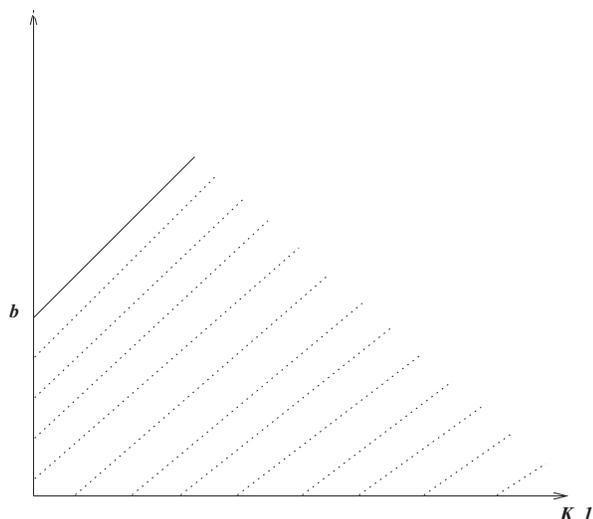


FIG. 2. Type 2 feasible region of K values are shown by the shaded region, which consists of a convex hull and a positive cone. The convex hull is the line segment joining $(0, 0)^T$ and $(0, b)^T$, and the positive cone is the set of all vectors that are all positive multiples of $(1, 1)^T$ and $(1, 0)^T$.

as $V \rightarrow \infty$. Hence $\text{Var}(\alpha_i) = p_i(1 - p_i)/b_i = O(1/V)$,

$$\text{Cov}(\alpha) = \begin{pmatrix} \frac{p_1(1-p_1)}{b_1} & 0 \\ 0 & \frac{p_2(1-p_2)}{b_2} \end{pmatrix} = O(1/V),$$

$$\text{Var}(\beta) = q = O(V),$$

as desired.

C. Two-dimensional feasible region: Type 2

The Type 2 example of polyhedral region is the shaded region shown in Fig. 2, which corresponds to the inequality constraint

$$K_2 - K_1 \leq b.$$

The corresponding convex hull is the line segment joining $(0, b)^T$ and $(0, 0)^T$, and the corresponding positive cone is the set of all vectors that are positive multiples of $(1, 1)^T$ and $(1, 0)^T$. Here b depends on the initial state x . The Minkowski–Weyl decomposition has the following form:

$$\begin{pmatrix} K_1 \\ K_2 \end{pmatrix} = \begin{pmatrix} 0 \\ b \end{pmatrix} \alpha + \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} \beta_1 \\ \beta_2 \end{pmatrix}, \quad (16)$$

where $0 \leq \alpha \leq 1$, $\beta_1 \geq 0$, and $\beta_2 \geq 0$. The simplest example with Type 2 region is the reversible monomolecular reaction $0 \leftrightarrow S_1$, where $K = (K_1, K_2)^T$ satisfies

$$K_2 - K_1 \leq x_1.$$

Thus in this case $b = x_1$. We describe a few more common examples with Type 2 region.

1. $0 \leftrightarrow S_1 + S_1$: The constraint on $K = (K_1, K_2)^T$ is given by

$$K_2 - K_1 \leq x_1/2.$$

We obtain $b = x_1/2$, and we may equivalently use $b = \lfloor x_1/2 \rfloor$ as explained before.

2. $S_2 \leftrightarrow S_1 + S_2$: The inequality condition for $K = (K_1, K_2)^T$ is given by

$$K_2 - K_1 \leq x_1.$$

Hence $b = x_1$.

We now describe how the IMW- τ method is applied to a pair of reactions with Type 2 region. Suppose the number of molecules prior to updating this pair of reactions is x . First compute b from x as described by the above examples. We denote $E(K) = \lambda$, $E(\alpha) = p$, and $E(\beta) = q$ as before. Taking the mean of Eq. (16), we get

$$\begin{pmatrix} \lambda_1 \\ \lambda_2 \end{pmatrix} = \begin{pmatrix} 0 \\ b \end{pmatrix} p + \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} q_1 \\ q_2 \end{pmatrix}. \quad (17)$$

We obtain the following linear system from Eq. (17):

$$q_1 + q_2 = \lambda_1, \quad bp + q_1 = \lambda_2, \quad (18)$$

with conditions

$$q_1 \geq 0, \quad q_2 \geq 0, \quad 0 \leq p \leq 1,$$

which follow from the inequality conditions on α and β .

We consider the case when $b \neq 0$, and we can find an upper bound \bar{p}_U and a lower bound \bar{p}_L for p as

$$\bar{p}_U = \min \left\{ \frac{\lambda_2}{b}, 1 \right\}, \quad \bar{p}_L = \max \left\{ \frac{\lambda_2 - \lambda_1}{b}, 0 \right\}.$$

We choose

$$p = r\bar{p}_L + (1-r)\bar{p}_U, \quad r = 0.5[1 - e^{-(\lambda_1 + \lambda_2)}].$$

The choice of r is to ensure consistency as in the Type 1 case.

We obtain that

$$q_1 = \lambda_2 - bp, \quad q_2 = \lambda_1 - q_1.$$

For the case $b = 0$, Eq. (18) has unique solutions for q_1 and q_2 , and $b\alpha = 0$.

Having found p , q_1 , and q_2 , we choose independent α , β_1 , and β_2 where $b\alpha \sim \mathcal{B}(b, p)$, $\beta_1 \sim \mathcal{P}(q_1)$, and $\beta_2 \sim \mathcal{P}(q_2)$. We set $K_1 = \beta_1 + \beta_2$, $K_2 = b\alpha + \beta_1$. The system is updated by $X = x + \nu K$.

A calculation similar to the one in Appendix E shows that $O(\tau)$ consistency is satisfied. For the sake of brevity, we do not show it in this paper. Moreover, we can verify that $\text{Cov}(\alpha) = O(1/V)$, and $\text{Cov}(\beta) = O(V)$ to show that fluid limit for this example behaves appropriately.

V. THE IMW- τ METHODS

In this section, we describe in detail the *implicit Minkowski–Weyl sequential* method (IMW-S) and the *implicit Minkowski–Weyl Parallel* method, mentioned in Sec. IV. Both methods involve partitioning the set of reactions into groups according to the *partitioning criterion* which states that *two different reactions are in the same group if and only if their stoichiometric vectors are either equal or the negative of each other*.

To understand the rationale behind the partitioning criterion, let us consider the example consisting of three reactions: (1) $S_1 \rightarrow S_1 + S_3$, (2) $S_2 \rightarrow S_2 + S_3$, and (3) $S_3 \rightarrow 0$. Note that reactions (1) and (2) both have the same stoichiometric vector $(0, 0, 1)^T$ and reaction (3) has the negative of this stoichiometric vector, namely $(0, 0, -1)^T$. Thus stoichiometrically both reactions (1) and (2) can be regarded as the reversal of the reaction (3). Additionally, since we are only interested in the updated state, it is adequate to know $K_1 + K_2$ without knowing K_1 and K_2 separately. If we set $K_{12} = K_1 + K_2$, we obtain the inequality condition for K_{12} and K_3 given by

$$K_3 - K_{12} \leq x_3.$$

In this problem, even though there are three reactions, there are only two different stoichiometric vectors involved. Consequently, this becomes a two-dimensional problem. In fact this group of three reactions can be handled by the Type 2 two-dimensional region.

In general, the above partitioning algorithm results in groups of reactions where the update problem for each group can be described by a one- or two-dimensional region of the type introduced in Sec. IV.

For ease of exposition, we shall assume throughout the rest of this section that no two reactions have the same stoichiometric vector or equivalently that the reaction counts K_j corresponding to reaction channels with identical stoichiometric vectors have been merged into one count K_j as in the example given above. Suppose there are M reactions partitioned into L groups: J_1, J_2, \dots, J_L , and $J_l \subset \{1, 2, \dots, M\}$. For $j \in J_l$ we denote by $v^{(l)}$ the matrix with the stoichiometric vectors v_j as column vectors, denote by $a^{(l)}$ the column vector with components a_j , and denote by $K^{(l)}$ the column vector with components K_j . Thus J_l contains a single reaction or a (stoichiometrically) reversible pair.

A. Implicit Minkowski–Weyl sequential method

In this section, we describe the implicit Minkowski–Weyl sequential method, where the reaction groups are updated sequentially.

IMW-S algorithm: Suppose the state at time t is x and we wish to compute the state X corresponding to time $t + \tau$, where τ is a chosen step size. We execute the following algorithm.

1. Set $X^{(0)} \leftarrow x$.
2. For $l = 1 : L$, execute the following loop.
 - (a) Compute X' from

$$X' = X^{(l-1)} + v^{(l)} a^{(l)} (X') \tau.$$

Let $\lambda^{(l)} = a^{(l)} (X') \tau$.

- (b) Generate samples $K^{(l)}$ with mean $\lambda^{(l)}$, using the methods in Sec. IV.
- (c) Update the states according to

$$X^{(l)} = X^{(l-1)} + v^{(l)} K^{(l)}.$$

- (1) Set the next state $X \leftarrow X^{(L)}$.

It is clear from Step 2 that if $X^{(l-1)} \geq 0$ then it follows that $X^{(l)} \geq 0$ as well. Thus, if $X^{(0)} = x \geq 0$, by mathematical induction we see that $X = X^{(L)} \geq 0$ as well.

It can be shown that the fluid limit of the IMW-S method is the sequential implicit Euler method for RRE. We now describe the difference between implicit Euler and sequential implicit Euler method applied to RRE as follows.

Recall that one step of the implicit Euler method for solving the RRE is given by

$$Y_{n+1} = Y_n + v \bar{a}(Y_{n+1}) \tau.$$

One step of the sequential implicit Euler method for the RRE with reactions partitioned as above is given by

$$\begin{aligned} Y_{n+1}^{(0)} &= Y_n, \\ Y_{n+1}^{(l)} &= Y_{n+1}^{(l-1)} + v^{(l)} \bar{a}^{(l)}(Y_{n+1}^{(l)}) \tau, \quad \text{for } l = 1, \dots, L, \\ Y_{n+1} &= Y_{n+1}^{(L)}. \end{aligned} \quad (19)$$

The sequential implicit Euler to our knowledge is not used in solving ODEs. As we explain later, it has some undesirable properties.

B. Implicit Minkowski–Weyl parallel method

In this section, we describe IMW-P, where we update reaction groups simultaneously and independently. We denote $\mathcal{P} = \{K \in \mathbb{Z}_+^M \mid x + v K \geq 0\}$ as before. Let $E(K^{(l)}) = \lambda^{(l)}$. We denote $K = (K^{(1)}, K^{(2)}, \dots, K^{(L)})^T$, and $\lambda = (\lambda^{(1)}, \lambda^{(2)}, \dots, \lambda^{(L)})^T$.

We formulate polyhedral regions $\mathcal{P}^{(1)}, \mathcal{P}^{(2)}, \dots, \mathcal{P}^{(L)}$ with one or two dimensions, defined by

$$\mathcal{P}^{(l)} = \{K^{(l)} \in \mathbb{Z}_+^{M_l} \mid x^{(l)} + v^{(l)} K^{(l)} \geq 0\},$$

where $M_l = 1$ or $M_l = 2$. Here $x^{(l)} \in \mathbb{Z}_+^N$ are to be chosen appropriately.

Note that when $K^{(1)}, K^{(2)}, \dots, K^{(L)}$ are computed independently and if $K = (K^{(1)}, K^{(2)}, \dots, K^{(L)})^T$, then $K \in \hat{\mathcal{P}} = \mathcal{P}^{(1)} \times \mathcal{P}^{(2)} \dots \times \mathcal{P}^{(L)}$. In general $\hat{\mathcal{P}} \neq \mathcal{P}$, and nonnegativity is not guaranteed. However, in order to satisfy $E(K) = \lambda$, it follows that $E(K^{(l)}) = \lambda^{(l)}$, and thus $\lambda^{(l)}$ must be in $\mathcal{P}^{(l)}$. We choose $x^{(l)}$ in the following manner to guarantee this.

Let $x^{(l)} = x + y_-^{(l)}$, where $y^{(l)} = x + v^{(l)} \lambda^{(l)}$. Note that $y_-^{(l)} = \max\{-y^{(l)}, 0\}$ is the negative part of $y^{(l)}$, and $y_+^{(l)} = \max\{y^{(l)}, 0\}$ is the positive part of $y^{(l)}$. It is known that $y^{(l)} = y_+^{(l)} - y_-^{(l)}$, where $y_+^{(l)}$ and $y_-^{(l)}$ are nonnegative. By this choice, $x^{(l)} + v^{(l)} \lambda^{(l)} = x + y_-^{(l)} + (y^{(l)} - x) = y_-^{(l)} + y^{(l)} = y_+^{(l)} \geq 0$.

As $\mathcal{P} \neq \hat{\mathcal{P}}$ in general, the IMW-P method may lead to the negative states of the updated state $X(t + \tau)$, and a *bounding procedure*²¹ is applied whenever a negative state is encountered.

IMW-P algorithm: Suppose the state at time t is x and we wish to compute the state X corresponding to time $t + \tau$, where τ is a chosen step size. We execute the following algorithm.

1. Given current state x and a step size τ , compute X' from $X' = x + v a(X') \tau$. Let $\lambda = a(X') \tau$.
2. For $l = 1 : L$, set $y^{(l)} \leftarrow x + v^{(l)} \lambda^{(l)}$, and $x^{(l)} \leftarrow x + y_-^{(l)}$. Generate samples $K^{(l)}$ with mean $\lambda^{(l)}$ and in the region $\mathcal{P}^{(l)} = \{K^{(l)} \in \mathbb{Z}_+^{M_l} \mid x^{(l)} + v^{(l)} K^{(l)} \geq 0\}$, using methods in Sec. IV.

TABLE II. Comparison of the proposed IMW- τ methods: IMW-S and IMW-P.

Methods/Issues	Fluid limit	Integer states	Nonnegativity
IMW-S	Sequential implicit Euler	YES	YES
IMW-P	Implicit Euler	YES	NO

- Update the states according to $X = x + \nu K$, where $K = (K^{(1)}, K^{(2)}, \dots, K^{(L)})^T$.
- Apply the bounding procedure²¹ if X is negative.

C. The issues with IMW- τ : Sequential versus parallel updating schemes

We proposed two different methods IMW-S and IMW-P, and we compare them in Table II. In Sec. III, we discussed the conditions that an ideal tau leap method satisfies, one being that the fluid limit should be a good stiff ODE solver. The

sequential implicit Euler, unfortunately, has some drawbacks as a stiff ODE solver. First, it does not always preserve the fixed points of an ODE.

There are special situations under which a fixed point of the RRE is also a fixed point for the sequential implicit Euler. This specifically happens if the following holds:

$$\nu a(X^*) = 0, \quad \text{iff} \quad \nu_j a_j(X^*) = 0, \quad (20)$$

for each $j \in J_l$.

In other words, when the equilibrium of the overall system is also an equilibrium within each group of reactions.

Another issue is that the IMW-S is often slower during the transients than the actual system unless smaller step sizes are used (see Sec. VID). This leads to lack of efficiency compared with the IMW-P. The IMW-P, on the other hand, behaves like the implicit Euler in the fluid limit and seems to overcome the above shortcomings of IMW-S. However, negative states may occur with IMW-P and one has to bound the negative states and this leads to errors.

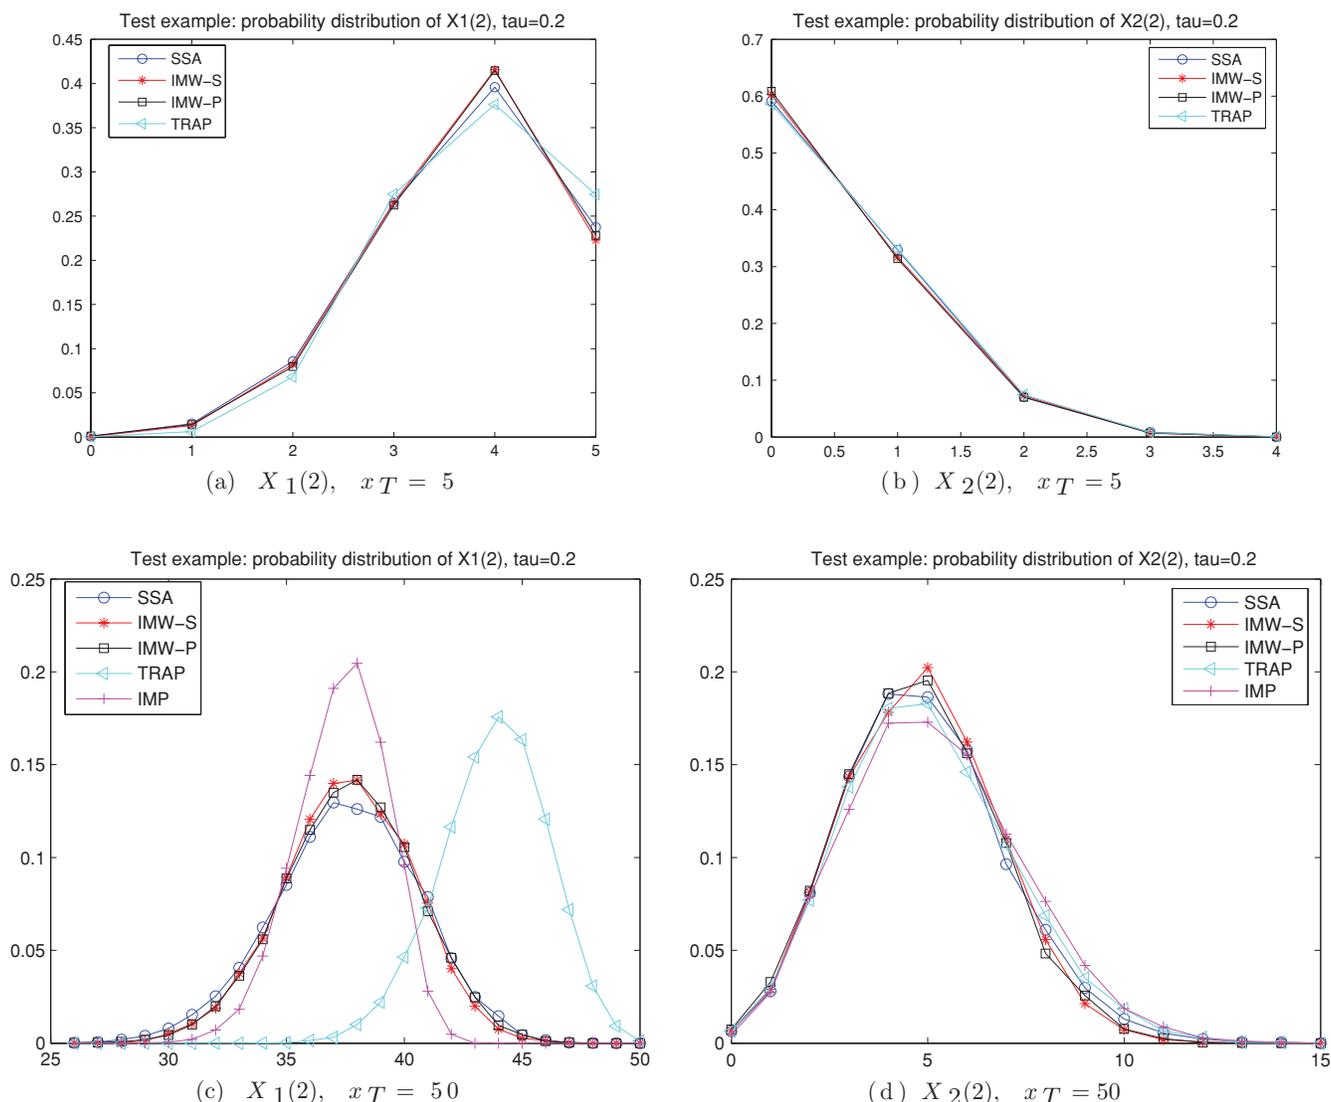


FIG. 3. Test example $S_2 \leftrightarrow S_1 \leftrightarrow S_3$: comparison of probability distributions (10 000 sample trajectories) of $X_1(2)$ and $X_2(2)$ obtained by the SSA (circle), IMW-S (star), IMW-P (square), trapezoidal tau (triangle), and implicit tau (plus). Here $\tau = 0.2$, $T = 2$, $x_T = 5$ for (a) and (b) and $x_T = 50$ for (c) and (d).

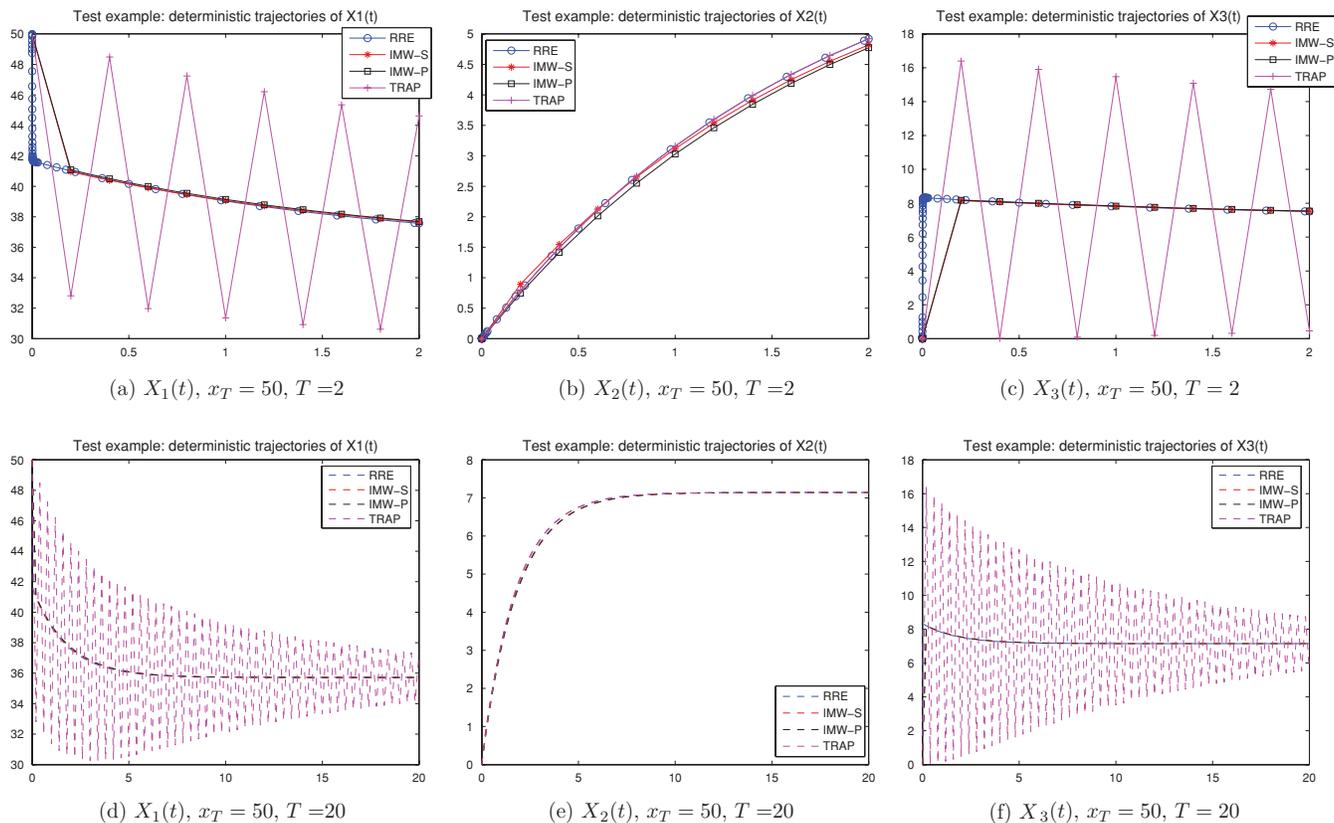


FIG. 4. Test example $S_2 \leftrightarrow S_1 \leftrightarrow S_3$: comparison of deterministic trajectories of $X_i(t)$ ($i = 1, 2, 3$) obtained by the RRE (blue circle), IMW-S (red star), IMW-P (black square), and trapezoidal tau (magenta plus). Here $T = 2$ for (a)–(c) and $T = 20$ for (d)–(f).

VI. NUMERICAL EXAMPLES

In this section, we illustrate the IMW-S and IMW-P methods by giving numerical results with several examples. First is the test example $S_2 \leftrightarrow S_1 \leftrightarrow S_3$, and the second is the test example $0 \leftrightarrow S_1 \rightarrow S_2 \rightarrow S_3 \rightarrow 0$. The other two are more complex biological examples which are introduced and explained later in this section.

We calculate the time scales from the RRE, represented by the eigenvalues of the Jacobian matrix estimated at the final time. The large range of eigenvalues exhibits the stiffness of the system. We choose the step size τ to be small

in comparison with the slowest time scale, but large in comparison with the fastest time scale. We compared the IMW-S and IMW-P methods with the exact simulation by SSA, the implicit tau method, the trapezoidal tau, and the REMM tau (parallel version) methods. We chose the parameters and initial conditions so as to obtain a range of copy numbers from small to medium in order to obtain a comprehensive analysis. The bounding procedure was applied to the IMW-P method, implicit tau, trapezoidal tau, and REMM tau methods when necessary. The IMW-S method naturally preserves the nonnegative states and does not require the bounding procedure.

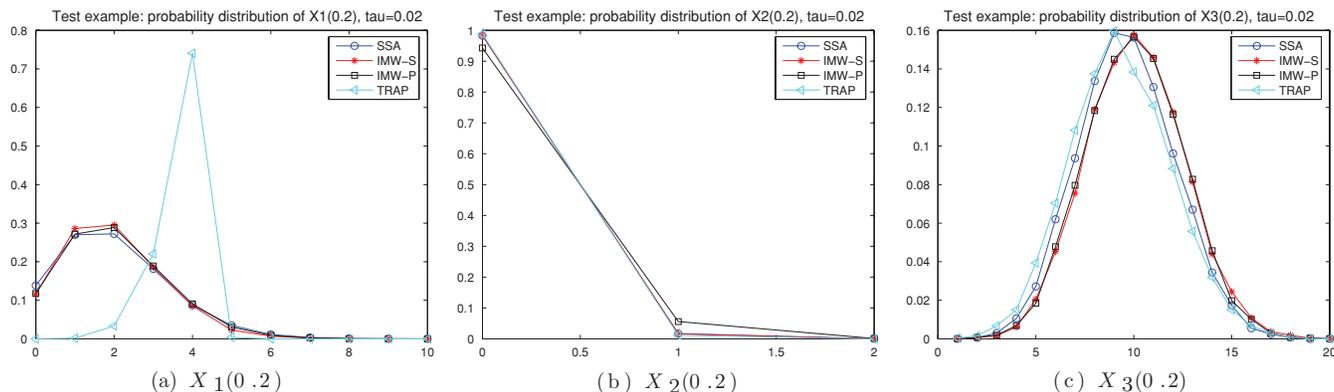


FIG. 5. Test example $0 \leftrightarrow S_1 \rightarrow S_2 \rightarrow S_3 \rightarrow 0$: comparison of probability distributions (10 000 sample trajectories) of $X_1(0.2)$, $X_2(0.2)$, and $X_3(0.2)$ obtained by the SSA (circle), IMW-S (star), IMW-P (square), trapezoidal tau (triangle). Here $\tau = 0.02$, and $T = 0.2$.

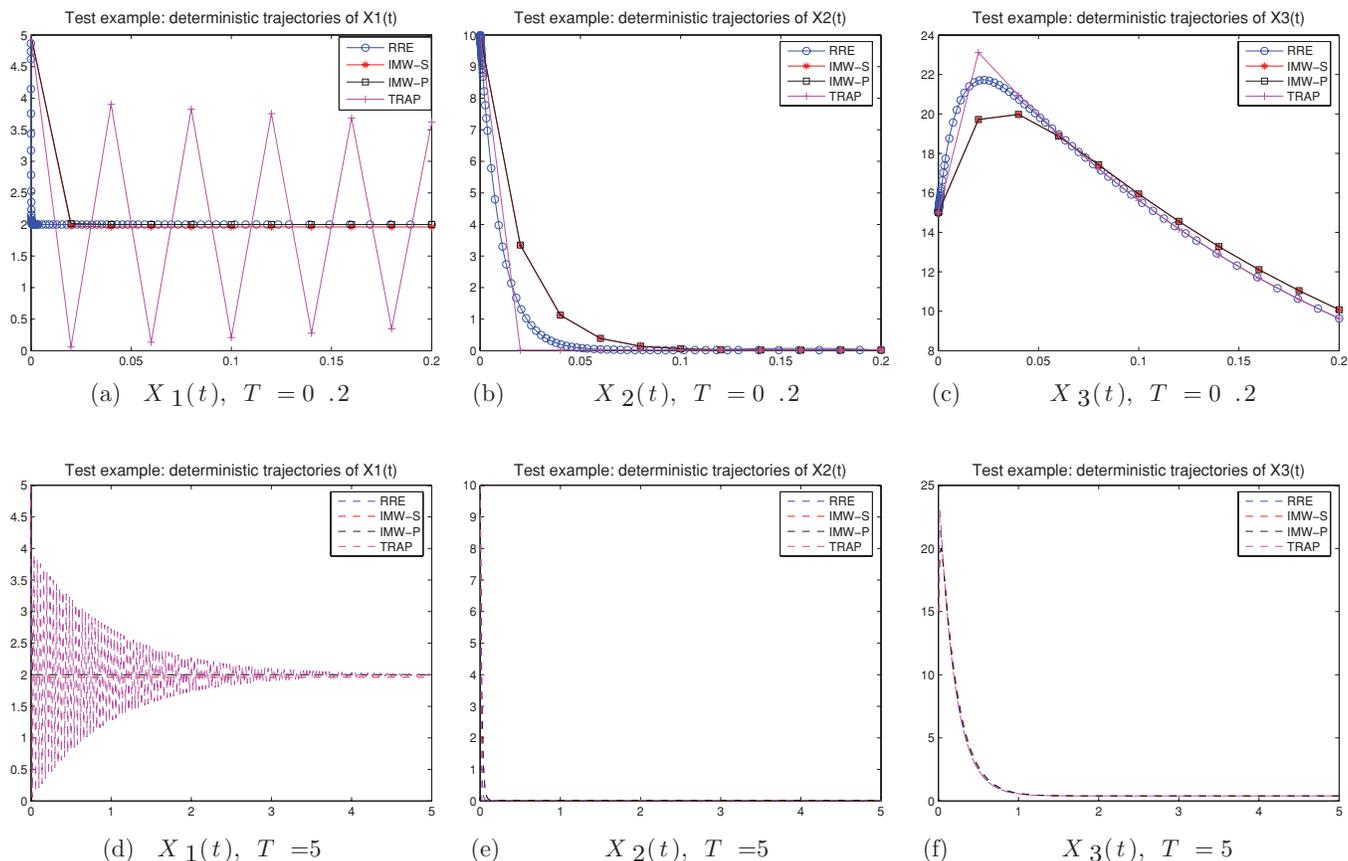
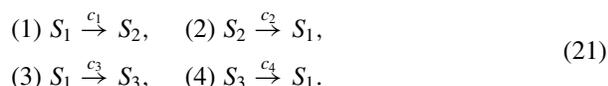


FIG. 6. Test example $0 \leftrightarrow S_1 \rightarrow S_2 \rightarrow S_3 \rightarrow 0$: comparison of deterministic trajectories of $X_i(t)$ ($i = 1, 2, 3$) obtained by the RRE (blue circle), IMW-S (red star), IMW-P (black square), and trapezoidal tau (magenta plus). Here $T = 0.2$ for (a)–(c) and $T = 5$ for (d)–(f).

A. Test example: $S_2 \leftrightarrow S_1 \leftrightarrow S_3$

We consider the following linear example Eq. (21) with two reversible pairs of reactions and they both have Type 1 Minkowski–Weyl decomposition. For sequential updating, the system contains two groups and each group consists of a reversible pair $\{(1), (2)\}$ and $\{(3), (4)\}$.



We chose two initial values: $X(0) = (5, 0, 0)^T$ and $X(0) = (50, 0, 0)^T$. The system has a conserved quantity $X_1(t) + X_2(t) + X_3(t) = X_1(0) + X_2(0) + X_3(0) = x_T$, where $x_T = 5$ and $x_T = 50$ corresponding to these initial values. We set $c_1 = 0.1$, $c_2 = 0.5$, $c_3 = 200$, and $c_4 = 1000$. The eigenvalues of the Jacobian matrix corresponding to the RRE are $(-1200, 0, 0.6)^T$ with the slowest time scale $1/0.6 \approx 1.667$, and the fastest time scale $1/1200 \approx 0.001$. We chose final time $T = 2$. The step size was $\tau = 0.2$.

The comparison of the probability distribution in a simulation of 10 000 trajectories for each method is shown in Fig. 3. The IMW-S and IMW-P perform better than the implicit tau and trapezoidal tau for $x_T = 50$. The trapezoidal tau does not capture the correct mean at $x_T = 50$.

One major issue with the trapezoidal tau method is that it performs poorly for very stiff problems since the transients

of the method decay slower than those of the true solution. This can be best understood by examining the deterministic part of this method applied to this example and compare it against the true solution of the RRE. See Fig. 4 where the RRE solution is compared with the approximate solutions obtained by applying the deterministic part of the various tau methods for the case $x_T = 50$. By the deterministic part of a tau method, we mean that at every time step, we update the state by the expected value of the update due to the tau method with no rounding applied. First, we observe that for a linear propensity system as in this example, applying the deterministic part of the tau method results in computing the mean of the tau leap solution at each step. The plots shown in Fig. 4 indicate that the means computed by IMW-P and IMW-S methods follow the RRE solution reasonably well while the mean computed by the trapezoidal tau method oscillates about the RRE trajectory for two of the components. This oscillation is explained by considering the *amplification factor* for the mean of the trapezoidal method which is given by $R = (2 + \lambda\tau)/(2 - \lambda\tau)$ for time step τ and an eigenvalue λ . In this example, the relevant eigenvalue is $\lambda = -1200$ and $\tau = 0.2$. For this choice $R = -0.9835$ and after ten time steps $R^{10} \approx 0.85$ which is much larger than the decay factor $e^{10\lambda\tau} \approx 0$ of the true solution. Plots in Fig. 4 parts (d) and (f) show that these oscillations remain even for a larger final time of $T = 20$ at which the system would have reached stationarity.

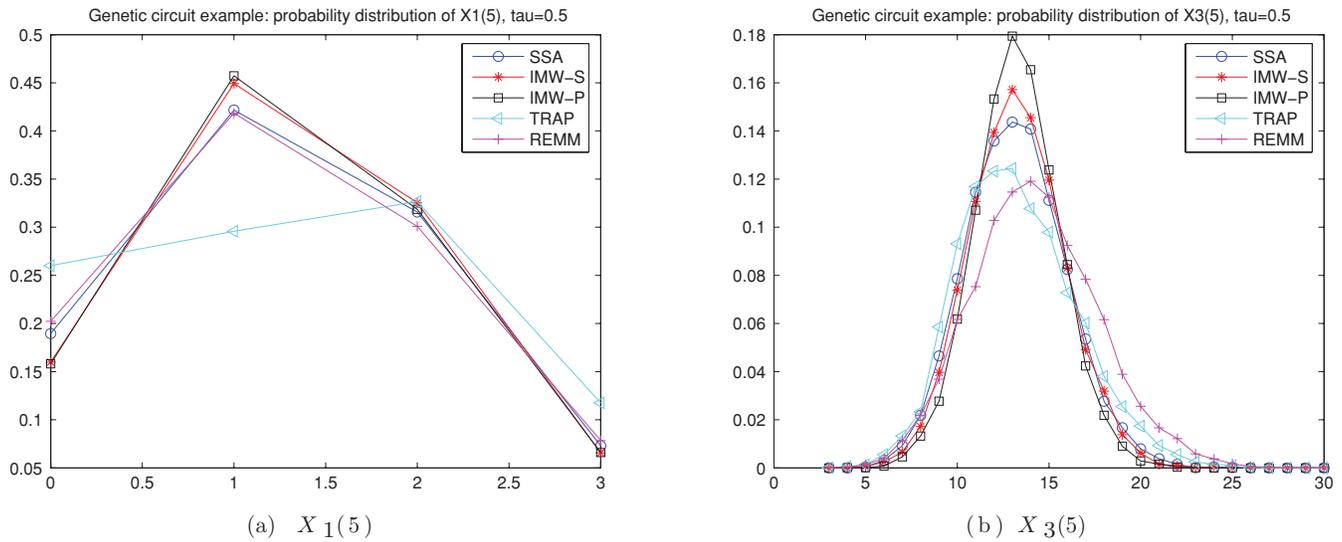
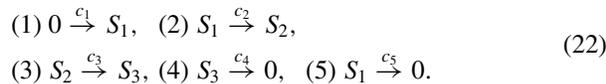


FIG. 7. Genetic circuit example: comparison of probability distributions (10 000 sample trajectories) of $X_1(5)$ and $X_3(5)$ obtained by the SSA (circle), IMW-S (star), IMW-P (square), trapezoidal tau (triangle), REMM tau (plus). Here $\tau = 0.5$, and $T = 5$.

B. Test example: $0 \leftrightarrow S_1 \rightarrow S_2 \rightarrow S_3 \rightarrow 0$

We consider the following chemical system (22). It contains three species undergoing five chemical reactions. By our method, we group Reactions (1), (5) together, which form a Type 2 region. We partition (2), (3), and (4) separately in three groups.

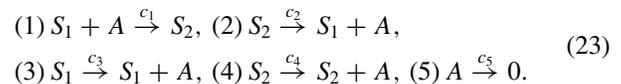


We chose the initial value to be $X(0) = (5, 10, 15)^T$. We set $c_1 = 20\,000$, $c_2 = 1$, $c_3 = 100$, $c_4 = 5$, and $c_5 = 10\,000$. The eigenvalues of the Jacobian matrix corresponding to the RRE are $(-5, -100, -10\,001)^T$ and hence the slowest time scale is 0.2 and the fastest time scale is $1/10\,001 \approx 10^{-4}$. We set $T = 0.2$. We chose $\tau = 0.02$.

The results are shown in Fig. 5 with the comparison of the probability distributions for $X_1(0.2)$, $X_2(0.2)$, and $X_3(0.2)$. The trapezoidal tau method performs poorly for species S_1 . The amplification factor for the mean of the trapezoidal method is $R = (2 + \lambda\tau)/(2 - \lambda\tau) = -0.98$, where $\lambda = -10\,001$ and $\tau = 0.02$. Thus $R^{10} \approx 0.82$ after ten time steps, and it is much larger than the decay factor $e^{10\lambda\tau} \approx 0$. As in the test example of the previous section we apply the deterministic part of the tau leap methods to the RRE and compare with the true solution of RRE. See Fig. 6, where the trapezoidal tau method has the oscillation about the RRE solution for the species S_1 .

C. Biological example 1: Genetic circuit

We consider the following genetic circuit example in Eq. (23), which describes a genetic transcription module with important biological significance.¹⁷ Reactions (1) and (2) correspond to the binding and unbinding, respectively, of the protein A to its own gene promoter S_1 . When the gene promoter is naked, it produces A at a rate c_3 by reaction (3). When A is bound to it, the gene promoter produces A at a rate c_4 by reaction (4). Finally A degenerates at a rate c_5 . In biological systems, if $c_3 < c_4$, it represents a positive feedback loop, and if $c_3 > c_4$, it is a negative feedback loop. Here we consider the situation of negative feedback loop



Let $X_1 = \#S_1$, $X_2 = \#S_2$, $X_3 = \#A$. The reactions (3) and (4) have the same stoichiometric vector which is the negative of that of reaction (5). Thus reactions (3), (4), and (5) form a group, which reduces to a Type 2 region. Thus we may group reactions (3), (4), and (5) together, resulting in a Type 2 region. We also group reactions (1) and (2) together, which form a Type 1 region.

We chose the initial value to be $X(0) = (3, 0, 14)^T$, and we note that $X_1(t) + X_2(t) = 3$ is a conserved quantity. We chose $c_1 = 100$, $c_2 = 1000$, $c_3 = 1$, $c_4 = 0.1$, $c_5 = 0.1$, and $T = 5$. The eigenvalues of the Jacobian at $T = 5$ are $(-2455, 0, -0.14)^T$. The fastest time scale is approximately 4×10^{-4} , and the slowest time scale is $1/0.14 \approx 7$. Here we chose $\tau = 0.5$.

TABLE III. Genetic circuit example: sample means and standard deviations of the state $X_1(5)$ (the sample size is 10 000) as computed by SSA, IMW-S, IMW-P, REMM tau, and trapezoidal tau.

$X_1(5)$ /Methods	SSA	IMW-S	IMW-P	Trapezoidal tau	REMM tau
Sample mean	1.2723	1.2966	1.2925	1.3018	1.2555
Standard deviation	0.8504	0.8120	0.8094	0.9829	0.8669

TABLE IV. Genetic circuit example: sample means and standard deviations of the state $X_3(5)$ (the sample size is 10 000) as computed by SSA, IMW-S, IMW-P, REMM tau, and trapezoidal tau.

$X_3(5)$ /Methods	SSA	IMW-S	IMW-P	Trapezoidal tau	REMM tau
Sample mean	13.1924	13.2558	13.2718	13.2854	14.2235
Standard deviation	2.7156	2.5778	2.3037	3.1997	3.4222

The sample means and standard deviations for each method are also provided for X_1 and X_3 at $T = 5$ for these methods. It is noted from Fig. 7 and Tables III and IV that the IMW-S and IMW-P methods capture the stochasticity better than the trapezoidal tau method.

D. Biological example 2: Genetic positive feedback loop

A more complex and stiff chemical network is considered with the example of the genetic positive feedback loop

in Eq. (24). Here x is the protein monomer, y is the protein dimer, d_0 is the regulatory site unbounded to protein dimer, d_r is the regulatory site bounded to protein dimer, and m is the mRNA. Reactions (1) and (2) describe the reversible reactions involving the dimerization of the protein. Reactions (3) and (4) are the binding and unbinding processes of the dimer to the regulatory site. Reactions (5) and (6) are the processes of transcription, and reaction (7) is the process of translation. Reactions (8) and (9) are the decays of the protein monomers and the mRNA. Reactions (1) – (4) have much faster time

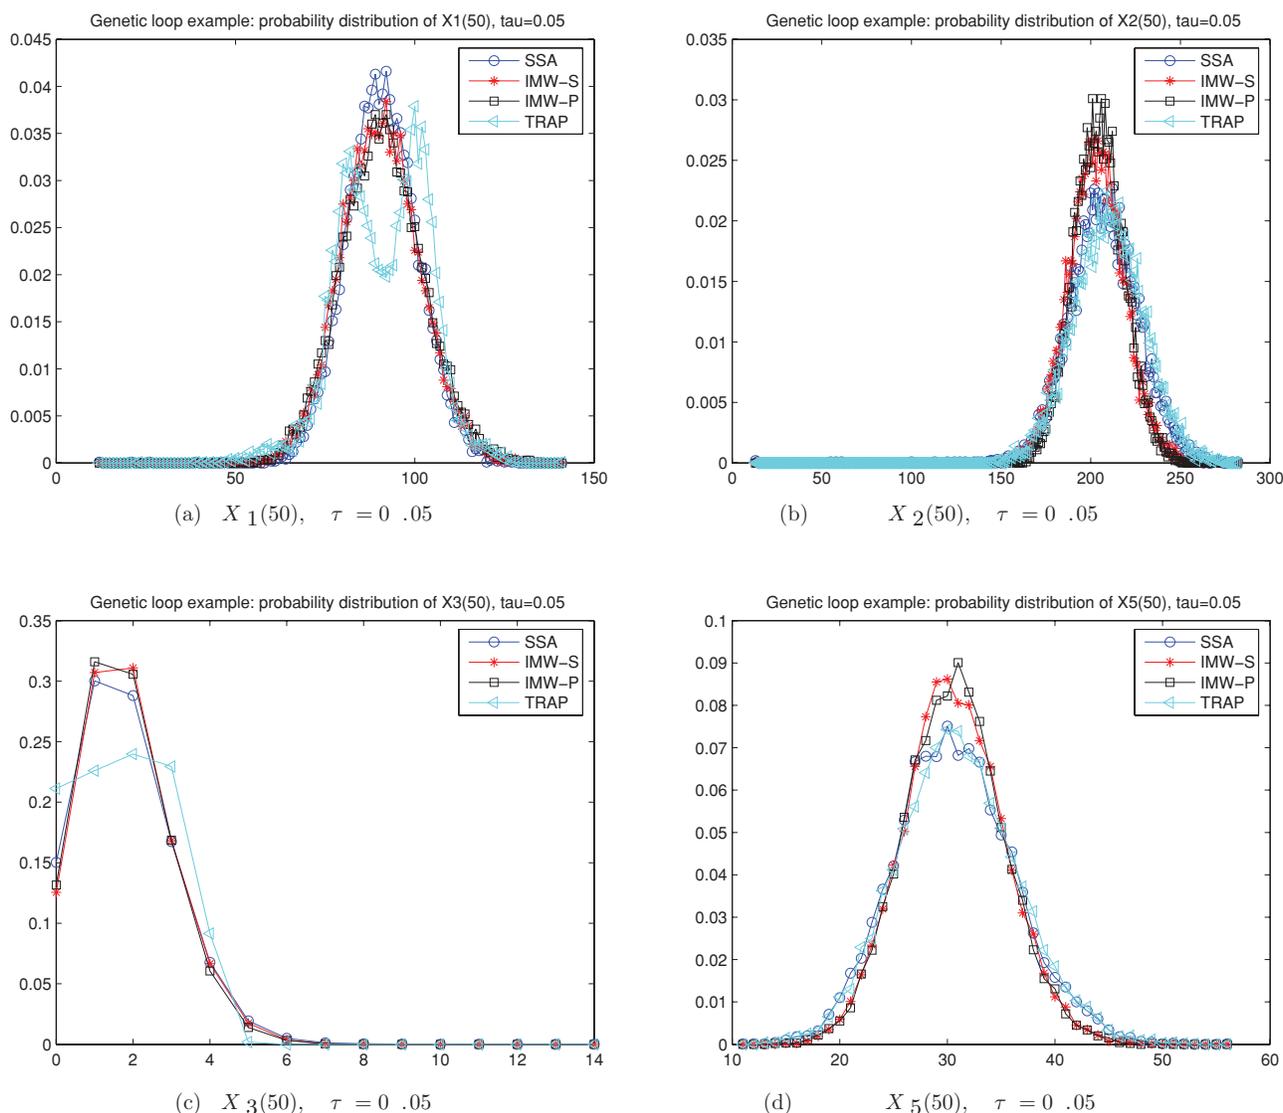


FIG. 8. Genetic loop example: comparison of probability distributions (10 000 sample trajectories) of $X_i(50)$ ($i = 1, 2, 3, 5$) obtained by the SSA (circle), IMW-S (star), IMW-P (square), trapezoidal tau (triangle). Here $\tau = 0.05$ and $T = 50$.

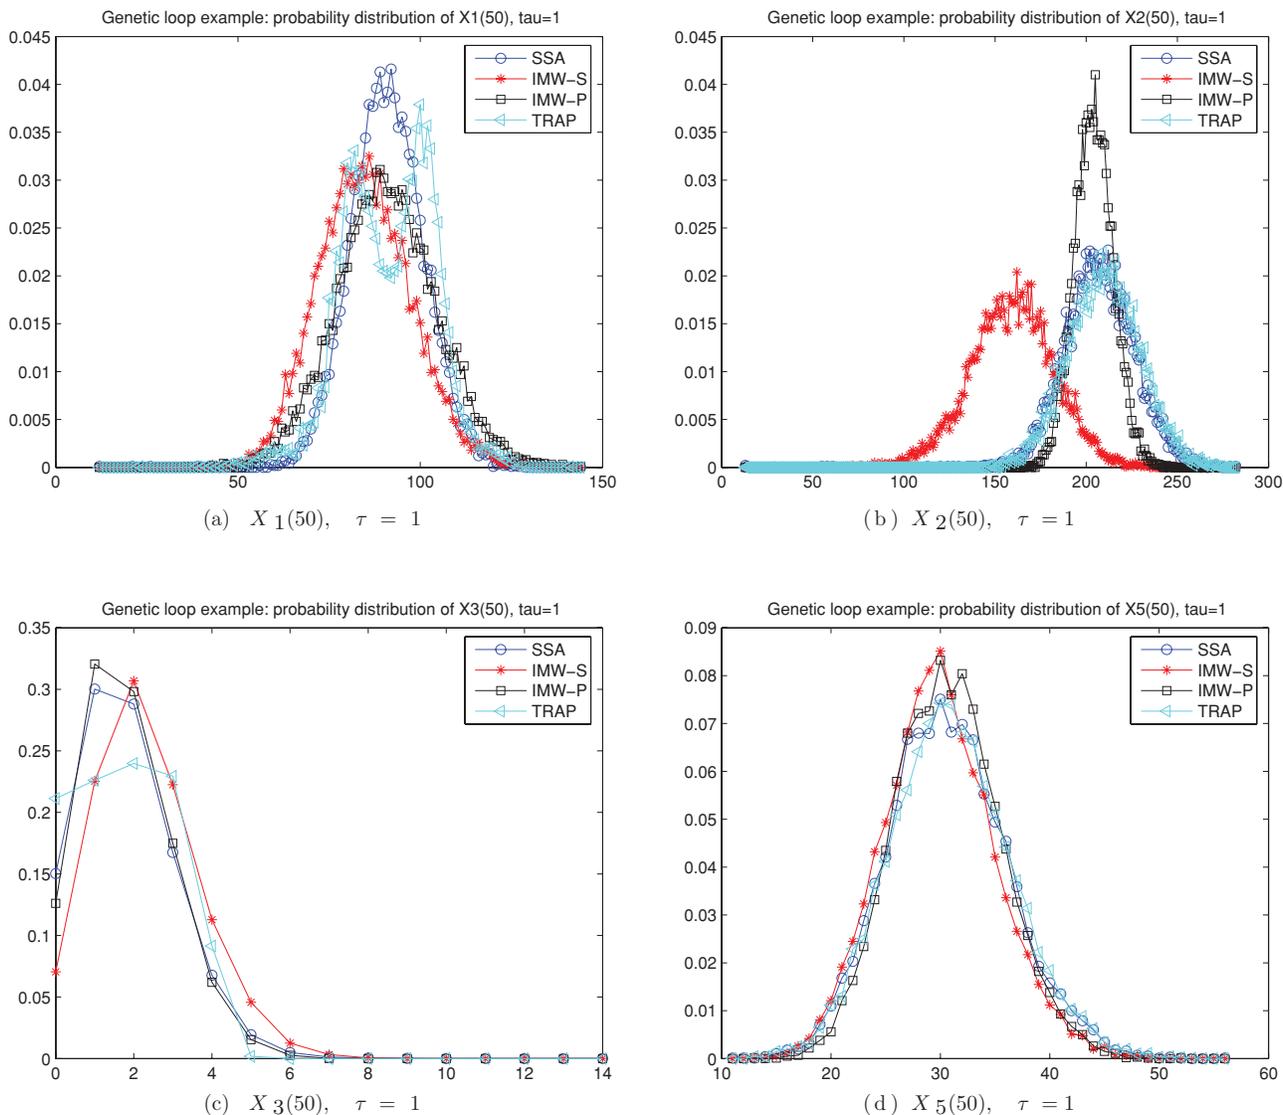
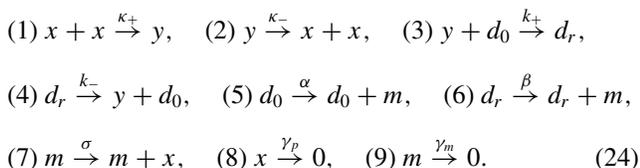


FIG. 9. Genetic loop example: comparison of probability distributions (10 000 sample trajectories) of $X_i(50)$ ($i = 1, 2, 3, 5$) obtained by the SSA (circle), IMW-S (star), IMW-P (square), trapezoidal tau (triangle). Here $\tau = 1$, and $T = 50$.

scale than reactions (5) – (9).²⁹



Let $X_1 = \#x$, $X_2 = \#y$, $X_3 = \#d_0$, $X_4 = \#d_r$, $X_5 = \#m$. The initial value is $X(0) = (10, 0, 20, 0, 0)^T$, and $X_3(t) + X_4(t) = 20$ is a conserved quantity. The reaction parameter values are $\kappa_+ = 50$, $\kappa_- = 1000$, $k_+ = 50$, $k_- = 1000$, $\alpha = 1$, $\beta = 10$, $\sigma = 3$, $\gamma_p = 1$, $\gamma_m = 6$. The final time chosen is $T = 50$. The eigenvalues of the Jacobian corresponding to RRE at $T = 50$ are $(-9979, -11382, -0.083, -6.02, 0)^T$. The fastest and the slowest time scales are $1/11382 \approx 0.0001$ and $1/0.083 = 12$, respectively. Following the partitioning criterion, we group the reactions as $\{(1), (2)\}$, $\{(3), (4)\}$, $\{(5), (6), (9)\}$, $\{(7), (8)\}$.

We first chose the step size $\tau = 0.05$. Figure 8 compares the SSA, IMW-S, IMW-P, and trapezoidal tau methods. First, the performance of IMW-S and IMW-P methods appear to be similar. Second, we notice that for some species the IMW-S and IMW-P methods perform better while the trapezoidal tau performs better for the other species. But the performance of the IMW-S and IMW-P methods overall seems to be more robust than that of the trapezoidal tau method.

We also compare these distributions by choosing a larger step size $\tau = 1$. The same initial states and parameter values described before were used. Figure 9 compares the probability distribution (at $T = 50$) for the SSA, IMW-S, IMW-P, and trapezoidal tau methods. We observe that IMW-S fails to reach the correct mean values for X_1 , X_2 , and X_3 . This is due to the fact that the sequential update is slower to catch up during the transient. The IMW-P does not suffer from this problem though it is less accurate than with time step $\tau = 0.05$. However, the IMW-S performs well for $\tau = 0.05$ as shown

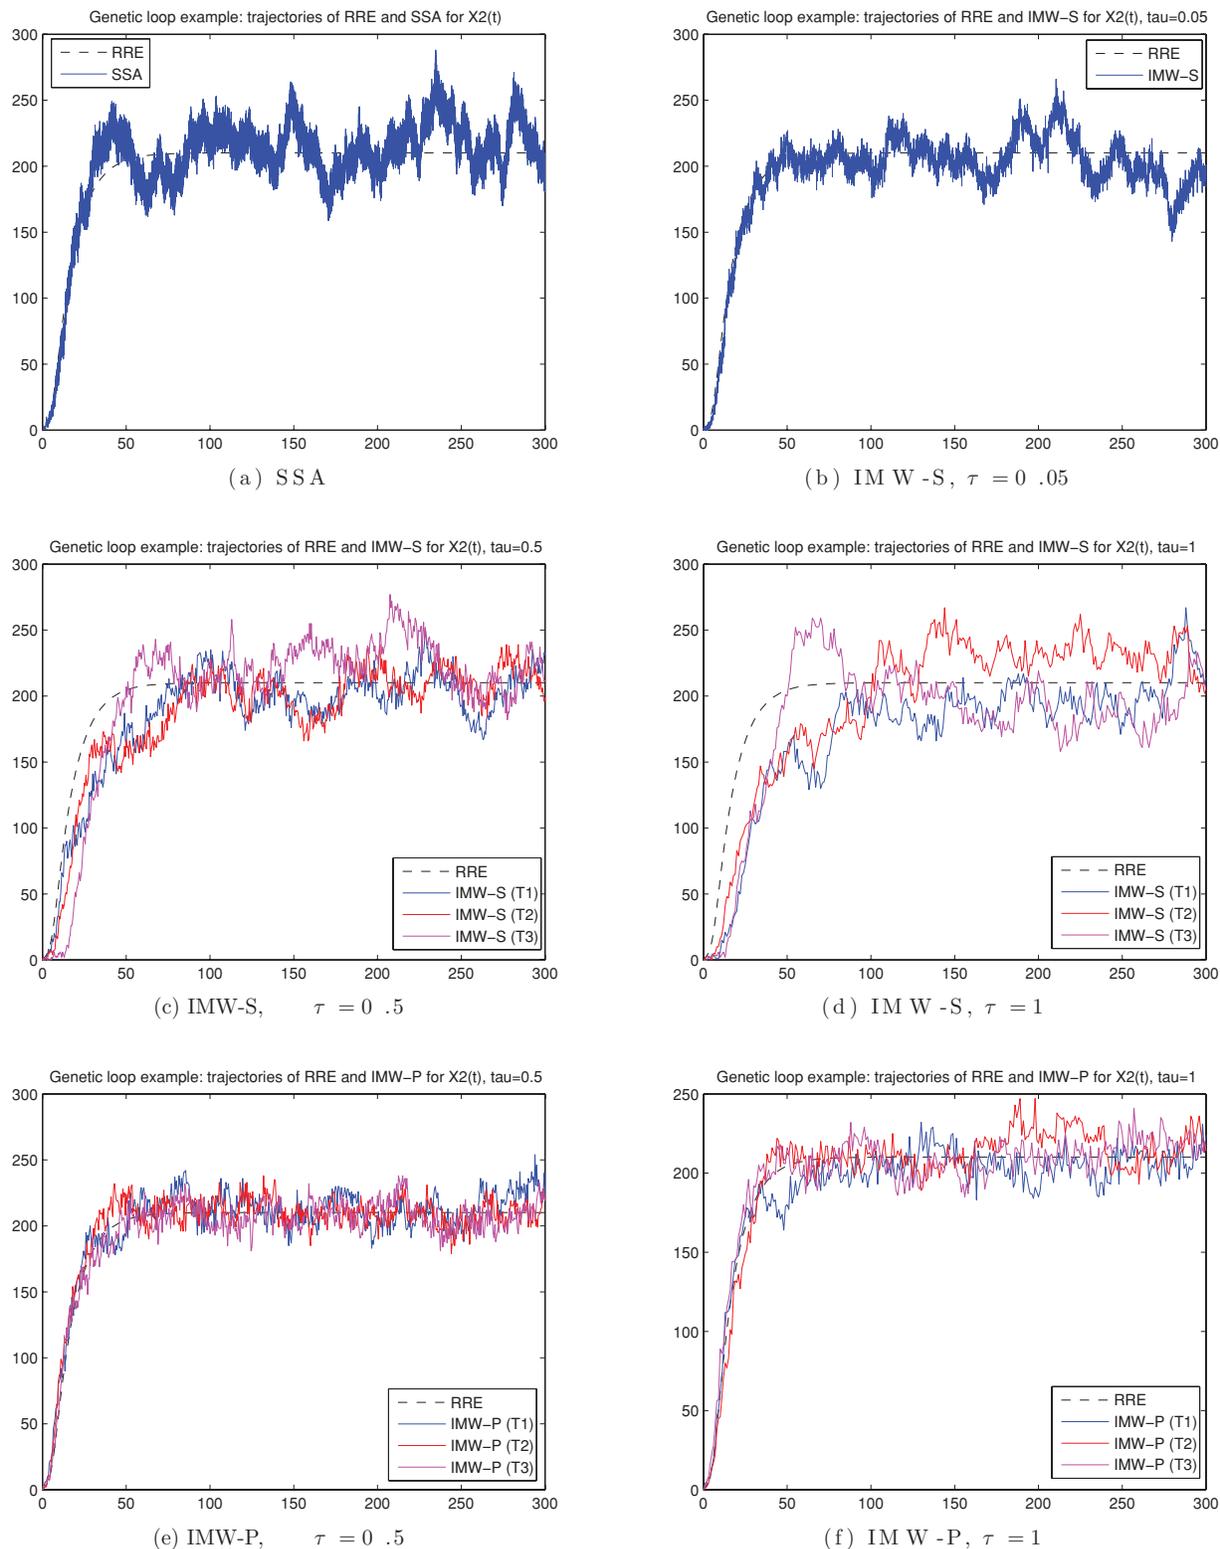


FIG. 10. Genetic loop example: comparison of trajectories $X_2(t)$ obtained by SSA, RRE, IMW-S and IMW-P for $T = 300$ computed with various τ values. Here (a) is the RRE vs SSA, (b)–(d) are RRE vs IMW-S corresponding to $\tau = 0.05$, $\tau = 0.5$, and $\tau = 1$, and (e)–(f) are RRE vs IMW-P corresponding to $\tau = 0.5$ and $\tau = 1$. Three sample trajectories (T1, T2, and T3) are provided for each case in (c)–(f). Note that these individual sample trajectories may not be compared across different methods as they are chosen independently.

earlier, which is still a step size very large compared with the fastest time scale.

We explored the methods more by comparing the trajectories of X_2 for different τ . Figure 10 depicts the trajectories of $X_2(t)$ with the SSA, IMW-S, and IMW-P methods. Since

the trajectories of SSA and the RRE are close during the transient, for ease of comparison, we only compare the methods with RRE in these plots. When $\tau = 0.05$, the IMW-S captures the transient well. For large step size $\tau = 0.5$, and $\tau = 1$, the paths of the method are deviated from RRE, with the method

showing a slower trend to reach the equilibrium states. On the other hand, the choice of the step size for IMW-P does not affect the transient approximation and the speed to reach the equilibrium states.

VII. CONCLUSION

In this paper, we developed two new tau leaping methods, both generally called IMW- τ , for the simulation of stiff stochastic chemical reactions. The IMW- τ methods use the Minkowski–Weyl decomposition to describe the polyhedral region of reaction count vectors that correspond to nonnegative population states. Additionally, they use a split step method where the first part involves computing the mean of the update implicitly and the second part involves generating random variables with the mean computed in the first part. The methods are presented in two versions, the sequential (IMW-S) and parallel (IMW-P) updating schemes, and both lead to integer valued states without rounding. The IMW-S method partitions the sets of reactions into groups, and each group consists of a single reaction or reversible pair. The method updates states according to these groups in a sequential manner. It naturally preserves the nonnegative states, and the fluid limit of IMW-S is the sequential implicit Euler. The IMW-P method maintains the advantage of IMW-S of working in one or two dimensions, but it updates the groups of reactions simultaneously in an independent manner. The fluid limit of IMW-P is the implicit Euler. It may lead to negative states, and the bounding procedure is applied whenever a negative state occurs.

We studied the numerical behavior of the IMW-S and IMW-P methods through a number of biologically motivated examples and compared them with the SSA, trapezoidal tau, and REMM tau (parallel) methods. We demonstrated that both the IMW- τ methods achieve good approximations for stiff systems. However, the IMW-P method was better able to capture the statistics of the trajectories during the transient (i.e., before reaching stationarity) with larger step sizes when compared to IMW-S.

ACKNOWLEDGMENTS

The authors M.R. and J.S. wish to acknowledge financial support from the National Science Foundation under Grant Nos. NSF-DMS-0610013 and NSF-ECCS-0900960, respectively. The authors also thank the anonymous reviewers for their constructive comments.

APPENDIX A: FLUID LIMIT

Consider a stochastic chemical system with M reaction channels and N molecular species. We associate with reaction channel j nonnegative vectors $\mu'_j, \mu_j \in \mathbb{Z}_+^N$, where μ_j is the vector whose i th component counts the number of molecules of i th species appearing as reactants in the reaction while μ'_j is the vector whose i th component counts the number

of molecules of i th species appearing as products in the reaction. For instance if $N = 2$ and if reaction j is given by $S_1 + S_1 \rightarrow S_1 + S_2$ then $\mu_j = (2, 0)^T$ and $\mu'_j = (1, 1)^T$. We denote the reaction propensity constants by c_1, \dots, c_M and the system volume by V . Define $\nu_j = \mu'_j - \mu_j$ to be the stoichiometric vectors. For convenience we define the “combinations” function $k : \mathbb{Z}_+ \times \mathbb{Z}_+ \rightarrow \mathbb{Z}_+$ by

$$k(x, y) = \frac{x!}{y!(x-y)!}, \quad y \leq x, \\ k(x, y) = 0, \quad y > x. \quad (\text{A1})$$

Thus $k(x, y)$ is the number of distinct ways to choose y items out of x items. Note that $k(x, 0) = 1$.

1. Volume dependence of propensity function

Suppose the system has the system volume V . The propensity function a_j of reaction j is given by³

$$a_j(x, V) = c_j \frac{1}{V^{|\mu_j|-1}} \prod_{i=1}^N k(x_i, \mu_{ij}), \quad (\text{A2})$$

where $|\mu_j| = \mu_{1j} + \dots + \mu_{Nj}$. Note that if reaction j is of the form $0 \rightarrow S_1$ then $|\mu_j| = 0$ and above equation gives

$$a_j(x, V) = c_j V.$$

This is a reasonable model since a “pure production” event has a propensity proportional to the volume V .

We can define the concentration $Z(t) = X(t)/V \in \mathbb{R}_+^N$ to be the number of species per volume. If we introduce the change of variable $z = x/V$ in Eq. (A2) and keep z fixed and let $V \rightarrow \infty$ (such that Vz remains integer), we get the asymptotic form

$$a_j(Vz, V) \sim V c_j \prod_{i=1}^N \frac{z_i^{\mu_{ij}}}{\mu_{ij}!}. \quad (\text{A3})$$

This follows because if we let $V \rightarrow \infty$ (such that Vz remains integer) we obtain

$$k(Vz, m) \sim \frac{V^m \prod_{i=0}^{m-1} (z - \frac{i}{V})}{m!} \sim \frac{V^m z^m}{m!}.$$

Note that if $m = 0$, $k(Vz, m) = 1$.

We define the reaction rate function \bar{a}_j of reaction j by

$$\bar{a}_j(z) = \kappa_j \prod_{i=1}^N z_i^{\mu_{ij}}, \quad (\text{A4})$$

where the reaction rate constant κ_j is related to the reaction propensity constant c_j by

$$\kappa_j = \frac{c_j}{\prod_{i=1}^N \mu_{ij}!}. \quad (\text{A5})$$

Thus it follows that as $V \rightarrow \infty$ such that Vz remains integer, we obtain the asymptotic relationship between the propensity function and the reaction rate function,

$$a_j(Vz, V) \sim V \bar{a}_j(z). \quad (\text{A6})$$

2. Fluid limit of the tau leap method

As we discussed in Sec. II B, let $\bar{Z}(t)$ be the unique solution of the RRE,

$$\dot{\bar{Z}}(t) = v\bar{a}[\bar{Z}(t)], \quad (\text{A7})$$

with initial condition $\bar{Z}(0) = z$.

Let \bar{Z}' be one step implicit Euler solution of Eq. (A7) with step size τ ,

$$\bar{Z}' = z + v\bar{a}(\bar{Z}')\tau. \quad (\text{A8})$$

Here we consider one step of the tau method with a fixed step size τ applied to an initial state $x = Vz$, where V is the system volume and z is the initial concentration. The resulting updated state X depends on V and thus we shall use X_V to indicate it. We wish the updated concentration $Z_V = X_V/V$ to approach a deterministic limit as $V \rightarrow \infty$ and we wish this limit to be the result of one step of implicit Euler with step size τ applied to the corresponding fluid limit system governed by the RRE. Now let X_V be one step leap approximation

$$X_V = Vz + vK_V. \quad (\text{A9})$$

Recall that the first step inside each time step of the tau method is the computation of the X'_V given by

$$X'_V = Vz + va(X'_V, V)\tau.$$

Divide by V

$$\frac{X'_V}{V} = z + \frac{va(X'_V, V)\tau}{V}.$$

Let $Z'_V = X'_V/V$ and assume that $\lim_{V \rightarrow \infty} Z'_V$ exists. Taking the limit as $V \rightarrow \infty$, we obtain

$$\lim_{V \rightarrow \infty} \frac{X'_V}{V} = z + v\bar{a}\left(\lim_{V \rightarrow \infty} \frac{X'_V}{V}\right)\tau,$$

and if we assume Eq. (A8) has a unique solution it follows that

$$\lim_{V \rightarrow \infty} Z'_V = \bar{Z}'.$$

Divide X_V by V from Eq. (A9), we get

$$\frac{X_V}{V} = z + \frac{vK_V}{V}.$$

Since $E(K_V) = a(X'_V)\tau$,

$$E\left(\frac{X_V}{V}\right) = z + \frac{va(X'_V)\tau}{V} = \frac{X'_V}{V}.$$

Therefore,

$$\lim_{V \rightarrow \infty} E\left(\frac{X_V}{V}\right) = \lim_{V \rightarrow \infty} \frac{X'_V}{V} = \bar{Z}'.$$

3. Sufficient conditions on α and β

We use the following lemma to ensure that $X_V/V \rightarrow \bar{Z}'$ weakly.

Lemma 1. Suppose $\text{Cov}(Y_V) \rightarrow 0$ and $E(Y_V) \rightarrow y$ as $V \rightarrow \infty$. Then $Y_V \rightarrow y$ weakly as $V \rightarrow \infty$.

Since we have already established that $E(X_V/V) \rightarrow \bar{Z}'$, it is sufficient to ensure that $\text{Cov}(X_V/V) \rightarrow 0$. By the relation $X_V/V = z + vK_V/V$, it is sufficient to ensure that $\text{Cov}(K_V/V) \rightarrow 0$.

Recall that for IMW- τ method, $K = B\alpha + D\beta$. We now find conditions on α and β to satisfy $\text{Cov}(K_V/V) \rightarrow 0$. First notice that D is independent of V while B is linear in V . Thus we can write

$$K_V = B_V\alpha_V + D\beta_V,$$

and

$$\text{Cov}(K_V) = B_V\text{Cov}(\alpha_V)B_V^T + D\text{Cov}(\beta_V)D^T.$$

Therefore,

$$\text{Cov}\left(\frac{K_V}{V}\right) = \frac{B_V\text{Cov}(\alpha_V)B_V^T}{V^2} + \frac{D\text{Cov}(\beta_V)D^T}{V^2}. \quad (\text{A10})$$

As $V \rightarrow \infty$, we may choose $\text{Cov}(\alpha_V) \rightarrow 0$ to ensure $B_V\text{Cov}(\alpha_V)B_V^T/V^2 \rightarrow 0$, and choose $\text{Cov}(\beta_V)$ to be $O(V)$, so that $D\text{Cov}(\beta_V)D^T/V^2 \rightarrow 0$. Our choice of distributions for α and β in the IMW-tau methods presented in this paper satisfies these conditions and thus the fluid limit of the tau method is the implicit Euler applied to the RRE.

APPENDIX B: CONSISTENCY OF THE IMW- τ METHOD FOR TYPE 1

Here we demonstrate $O(\tau)$ consistency of the IMW- τ method for the two-dimensional Type 1 region mentioned in Sec. IV B.

Let $R(\tau)$ be the exact number of reactions in $(t, t + \tau]$, and $K(\tau)$ be the approximated number of reactions for IMW- τ method in $(t, t + \tau]$. We shall establish that $P\{K(\tau) = l\} - P\{R(\tau) = l\} = O(\tau^2)$ for all $l \in \mathbb{Z}_+^2$.

When the IMW- τ method is applied to this type, we have that $K_1 = b_1\alpha_1 + \beta$, $K_2 = b_2\alpha_2 + \beta$, and we choose $r = r_0(1 - e^{-(\lambda_1 + \lambda_2)\tau}) = O(\tau)$. We calculate $P\{K(\tau) = l\} = P\{K_1 = l_1, K_2 = l_2\}$, and write them in terms of ordered powers of τ . We do not discuss the trivial case when $b_1 = b_2 = 0$, since the updated state is always x .

First we shall establish the orders of λ , p_1 , p_2 , and q . Since the implicit Euler solution X' of the RRE satisfies $X' = x + va(X')\tau$, we can verify that $X' = x + O(\tau)$ by the implicit function theorem. Thus $\lambda_j = a_j(X')\tau = a_j(x + O(\tau))\tau$ can be Taylor expanded at x as

$$\lambda_j = a_j(x + O(\tau))\tau = a_j(x)\tau + a'_j(x)o(\tau).$$

Therefore, when b_1 and b_2 are both nonzero, all the elements of x are nonzero, and we can write $\lambda_1 = O(\tau)$ and $\lambda_2 = O(\tau)$. From Eq. (13) we obtain $\bar{p}_U = O(\tau)$, $\bar{p}_L = O(\tau)$, and hence from Eq. (14) we obtain $\bar{p} = O(\tau)$. Moreover, $r = O(\tau)$, from Eq. (15).

For the case $b_1 = 0$ and $b_2 \neq 0$, we can verify that $\lambda_1 = O(\tau^2)$ and $\lambda_2 = O(\tau)$. Thus $q = \lambda_1 = O(\tau^2)$, and p_2

$= \lambda_2 - q/b_2 = O(\tau)$. Likewise, when $b_1 \neq 0$ and $b_2 = 0$, $q = \lambda_2 = O(\tau^2)$, and $p_1 = O(\tau)$. We summarize the orders of p_i and q as follows.

- If $b_1 b_2 \neq 0$, $p_1 = p_2 = O(\tau)$, $q = O(\tau^2)$.
- If $b_1 = 0$, $b_2 \neq 0$, $q = O(\tau^2)$, $p_2 = O(\tau)$.
- If $b_2 = 0$, $b_1 \neq 0$, $q = O(\tau^2)$, $p_1 = O(\tau)$.

We first show that the IMW- τ method is $O(\tau^{l_1+l_2})$ for $(b_1, b_2) \neq (0, 0)^T$. Recall that $b_1 \alpha_1 \sim \mathcal{B}(b_1, p_1)$, $b_2 \alpha_2 \sim \mathcal{B}(b_2, p_2)$, and $\beta \sim \mathcal{P}(q)$, where \mathcal{B} and \mathcal{P} are binomial and Poisson random variables, and $b_1 \alpha_1$, $b_2 \alpha_2$, and β are independent.

- If $b_1 b_2 \neq 0$, then

$$\begin{aligned} P\{K_1 = l_1, K_2 = l_2\} &= \sum_{l=0}^{\min(l_1, l_2)} P\{\beta = l, b_1 \alpha_1 = l_1 - l, b_2 \alpha_2 = l_2 - l\} \\ &= \sum_{l=0}^{\min(l_1, l_2)} P\{\beta = l\} P\{b_1 \alpha_1 = l_1 - l\} \\ &\quad \times P\{b_2 \alpha_2 = l_2 - l\} \\ &= \sum_{l=0}^{\min(l_1, l_2)} \frac{e^{-q} q^l}{l!} \\ &\quad \times \binom{b_1}{l_1-l} p_1^{l_1-l} (1-p_1)^{b_1-l_1+l} \\ &\quad \times \binom{b_2}{l_2-l} p_2^{l_2-l} (1-p_2)^{b_2-l_2+l} \\ &= \sum_{l=0}^{\min(l_1, l_2)} O(\tau^{2l}) \times O(\tau^{l_1-l}) \times O(\tau^{l_2-l}) = O(\tau^{l_1+l_2}). \end{aligned}$$

- Similarly, if $b_1 = 0$ and $b_2 \neq 0$, then $P\{b_1 \alpha_1 = 0\} = 1$, and hence

$$\begin{aligned} P\{K_1 = l_1, K_2 = l_2\} &= P\{\beta = l_1\} P\{b_2 \alpha_2 = l_2 - l_1\} = O(\tau^{l_1+l_2}). \end{aligned}$$

- Likewise, if $b_2 = 0$, and $b_1 \neq 0$, then

$$\begin{aligned} P\{K_1 = l_1, K_2 = l_2\} &= P\{\beta = l_1\} P\{b_1 \alpha_1 = l_2 - l_1\} = O(\tau^{l_1+l_2}). \end{aligned}$$

Thus we have shown that $P\{K(\tau) = l\} = P\{K_1 = l_1, K_2 = l_2\} = O(\tau^{l_1+l_2})$. It is known that $P\{R(\tau) = l\} = P\{R_1 = l_1, R_2 = l_2\} = O(\tau^{l_1+l_2})$.²¹ Thus when $l_1 + l_2 \geq 2$, we can obtain $P\{K(\tau) = l\} - P\{R(\tau) = l\} = O(\tau^2)$.

For the case $l_1 + l_2 = 1$, namely $(K_1 = 0, K_2 = 1)$ or $(K_1 = 1, K_2 = 0)$, in order to guarantee the $O(\tau^2)$ consistency for this case, the coefficient of τ of the IMW- τ method should be the same as the coefficient of τ for true solution. By definition of propensity function of R_2 , $P\{R_1 = 0, R_2 = 1\} = a_2(x)\tau + O(\tau^2) = \lambda_2 + O(\tau^2)$. Similarly, $P\{R_1 = 1, R_2 = 0\} = a_1(x)\tau + O(\tau^2) = \lambda_1 + O(\tau^2)$. For the IMW- τ method, we have the following results.

- If $K_1 = 0, K_2 = 1$, then

$$\begin{aligned} P\{K_1 = 0, K_2 = 1\} &= P\{b_2 \alpha_2 = 1\} = b_2 p_2 (1 - p_2)^{b_2-1} = \lambda_2 + O(\tau^2), \end{aligned}$$

since $p_2 = (\lambda_2 - q)/b_2$.

- Similarly, if $K_1 = 1, K_2 = 0$,

$$P\{K_1 = 1, K_2 = 0\} = P\{b_1 \alpha_1 = 1\} = \lambda_1 + O(\tau^2).$$

We obtain the local error formulae below. Thus we reach the $O(\tau)$ consistency.

1. If $0 < l_1 + l_2 \leq 1$, $P\{K_1 = l_1, K_2 = l_2\}$ and $P\{R_1 = l_1, R_2 = l_2\}$ have the same coefficient for τ , so $O(\tau)$ can be eliminated, then

$$P\{K_1 = l_1, K_2 = l_2\} - P\{R_1 = l_1, R_2 = l_2\} = O(\tau^2). \quad (\text{B1})$$

2. If $l_1 + l_2 \geq 2$, then

$$\begin{aligned} P\{K_1 = l_1, K_2 = l_2\} - P\{R_1 = l_1, R_2 = l_2\} &= O(\tau^{l_1+l_2}). \end{aligned} \quad (\text{B2})$$

3. If $l_1 + l_2 = 0$, namely $l_1 = l_2 = 0$, we apply the statements (B1) and (B2) to obtain,

$$\begin{aligned} P\{K_1 = 0, K_2 = 0\} - P\{R_1 = 0, R_2 = 0\} &= (1 - \sum_{(l_1, l_2) \neq (0,0)} P\{R_1 = l_1, R_2 = l_2\}) \\ &\quad - (1 - \sum_{(l_1, l_2) \neq (0,0)} P\{K_1 = l_1, K_2 = l_2\}) \\ &= \sum_{(l_1, l_2) \neq (0,0)} (P\{K_1 = l_1, K_2 = l_2\} \\ &\quad - P\{R_1 = l_1, R_2 = l_2\}) \\ &= O(\tau^2). \end{aligned}$$

- ¹D. T. Gillespie, *J. Comput. Phys.* **22**, 403 (1976).
- ²D. T. Gillespie, *J. Phys. Chem.* **81**, 2340 (1977).
- ³D. T. Gillespie, *Physica A* **188**, 404 (1992).
- ⁴D. T. Gillespie, *J. Chem. Phys.* **113**, 297 (2000).
- ⁵T. G. Kurtz, *Stochastic Proc. Appl.* **6**, 223 (1978).
- ⁶S. N. Ethier and T. G. Kurtz, *Markov Processes: Characterization and Convergence* (Wiley, New York, 1986).
- ⁷B. Munsky and M. Khammash, *J. Chem. Phys.* **124**, 044104 (2006).
- ⁸J. Zhang, L. T. Waston, and Y. Cao, *Comput. Math. Appl.* **59**, 573 (2010).
- ⁹Y. Cao, D. T. Gillespie, and L. R. Petzold, *J. Chem. Phys.* **122**, 14116 (2005).
- ¹⁰E. L. Haseltine and J. B. Rawlings, *J. Chem. Phys.* **117**, 6959 (2002).
- ¹¹E. Weinan, D. Liu, and E. Vanden-Eijnden, *J. Comput. Phys.* **221**, 158 (2007).
- ¹²C. V. Rao and A. P. Arkin, *J. Chem. Phys.* **118**, 4999 (2003).
- ¹³K. Ball, T. G. Kurtz, L. Popovic, and G. Rempala, *Ann. Appl. Probab.* **16**, 1925 (2006).
- ¹⁴D. T. Gillespie, *J. Chem. Phys.* **115**, 1716 (2001).
- ¹⁵M. Rathinam, L. R. Petzold, Y. Cao, and D. T. Gillespie, *J. Chem. Phys.* **119**, 12784 (2003).
- ¹⁶Y. Cao, L. R. Petzold, M. Rathinam, and D. T. Gillespie, *J. Chem. Phys.* **121**, 12169 (2004).
- ¹⁷M. Rathinam and H. E. Samad, *J. Comput. Phys.* **224**, 897 (2007).
- ¹⁸A. Chatterjee, D. Vlachos, and M. Katsoulakis, *J. Chem. Phys.* **122**, 24112 (2005).
- ¹⁹T. Tian and K. Burrage, *J. Chem. Phys.* **121**, 10356 (2004).
- ²⁰Y. Hu and T. Li, *J. Chem. Phys.* **130**, 124109 (2009).
- ²¹M. Rathinam, L. R. Petzold, Y. Cao, and D. T. Gillespie, *Multiscale Model Simul.* **4**, 867 (2005).
- ²²T. Li, *Multiscale Model Simul.* **6**, 417 (2007).
- ²³D. Anderson, A. Ganguly, and T. G. Kurtz, *Error analysis of the tau-leap simulation method for stochastically modeled chemical reaction systems*, *Ann. Appl. Probab.* (to be published).
- ²⁴N. G. Van Kampen, *Stochastic Processes in Physics and Chemistry* (North Holland, Amsterdam, 1992).
- ²⁵D. T. Gillespie, *Markov Processes: An Introduction for Physical Scientists* (Academic, Philadelphia, PA, 1991).
- ²⁶G. L. Nemhauser, *Integer and Combinatorial Optimization* (Wiley, New York, 1999).
- ²⁷E. Neuman, MATLAB Tutorials: <http://www.math.siu.edu/matlab/tutorial6.pdf>.
- ²⁸R. J. Connor and J. E. Mosimann, *J. Am. Stat. Assoc.* **64**, 194 (1969).
- ²⁹M. R. Bennett, D. Volfson, L. Tsimring, and J. Hasty, *J. Biol. Phys.* **92**, 3501 (2007).