

# Uniformization techniques for stochastic simulation of chemical reaction networks

Cite as: J. Chem. Phys. **150**, 154107 (2019); <https://doi.org/10.1063/1.5081043>

Submitted: 13 November 2018 . Accepted: 24 March 2019 . Published Online: 16 April 2019

Casper H. L. Beentjes , and Ruth E. Baker 



View Online



Export Citation



CrossMark

## ARTICLES YOU MAY BE INTERESTED IN

[A new class of reaction path based potential energy surfaces enabling accurate black box chemical rate constant calculations](#)

The Journal of Chemical Physics **150**, 154105 (2019); <https://doi.org/10.1063/1.5092589>

[Unsupervised machine learning in atomistic simulations, between predictions and understanding](#)

The Journal of Chemical Physics **150**, 150901 (2019); <https://doi.org/10.1063/1.5091842>

[A novel approach to describe chemical environments in high-dimensional neural network potentials](#)

The Journal of Chemical Physics **150**, 154102 (2019); <https://doi.org/10.1063/1.5086167>

Lock-in Amplifiers  
up to 600 MHz



Watch



# Uniformization techniques for stochastic simulation of chemical reaction networks

Cite as: J. Chem. Phys. 150, 154107 (2019); doi: 10.1063/1.5081043

Submitted: 13 November 2018 • Accepted: 24 March 2019 •

Published Online: 16 April 2019



Casper H. L. Beentjes<sup>a)</sup>  and Ruth E. Baker 

## AFFILIATIONS

Mathematical Institute, University of Oxford, Oxford, United Kingdom

E-mail: beentjes@maths.ox.ac.uk.

## ABSTRACT

This work considers the method of uniformization for continuous-time Markov chains in the context of chemical reaction networks. Previous work in the literature has shown that uniformization can be beneficial in the context of time-inhomogeneous models, such as chemical reaction networks incorporating extrinsic noise. This paper lays focus on the understanding of uniformization from the viewpoint of sample paths of chemical reaction networks. In particular, an efficient pathwise stochastic simulation algorithm for time-homogeneous models is presented which is complexity-wise equal to Gillespie's direct method. This new approach therefore enlarges the class of problems for which the uniformization approach forms a computationally attractive choice. Furthermore, as a new application of the uniformization method, we provide a novel variance reduction method for (raw) moment estimators of chemical reaction networks based upon the combination of stratification and uniformization.

Published under license by AIP Publishing. <https://doi.org/10.1063/1.5081043>

## I. INTRODUCTION

Mathematical modeling has become an indispensable tool in the study of chemical reaction networks. Models often include stochastic effects to account for the intrinsic and extrinsic noise of the chemical system under study. This presence of noise in cellular reaction networks has been observed experimentally,<sup>10,13,18,36</sup> and the addition of stochastic effects can lead to phenomena such as stochastic focusing<sup>39</sup> or resonance-inducing oscillations<sup>29</sup> that are not observed in the deterministic counterpart models.

Commonly, stochastic models take the form of a continuous-time Markov chain (CTMC), and whilst many such models are often deceptively simple to write down, analytic exploration is only possible for a marginal class of models.<sup>30,41</sup> In order to study these stochastic models, one therefore frequently has to rely either on approximate methods,<sup>25,34,38</sup> which can work well for low-dimensional problems, or on stochastic simulation if the problem dimension is too high for the approximate methods to be feasible. Stochastic simulation involves generating realizations of the process underlying a model by use of a stochastic simulation algorithm (SSA). The sample paths generated by these SSAs are then used to calculate Monte Carlo (MC) summary statistics to interrogate the system dynamics. These MC methods, however, suffer from an

important drawback, namely, the convergence rate is of the order  $N^{-1/2}$ , where  $N$  is the number of MC trajectories simulated. This slow convergence can render the exploration of models with high variance in copy numbers computationally intractable with existing simulation methods.

In order to enlarge the class of problems that can be studied computationally, work has therefore focused either on “speeding up” the SSAs in order to generate more sample paths in a given time span or on reducing the variance of the summary statistics employed. The former approach has been relatively successful and has led to a range of variations on the gold standard Gillespie's direct method (DM)<sup>1,15,20,22,37,56</sup> that are widely employed. Variance reduction techniques, on the other hand, have traditionally received less attention in the field of stochastic simulation of chemical reaction networks. However, with the recent advent of multilevel Monte Carlo (MLMC)<sup>2,9,33</sup> and quasi-Monte Carlo (QMC) methods,<sup>8,26</sup> it was shown that it is possible to use variance reduction ideas from related fields in the context of chemical reaction modeling, which can lead to large computational gains.

In this work, we focus on the application of uniformization of CTMCs to the stochastic simulation of chemical reaction networks. This method has been used previously in different forms in Refs. 5, 26, 54, and 57, and we show here how one can improve

the computational efficiency of SSAs based on uniformization. This makes stochastic simulation with uniformization competitive with the standard SSAs currently available. In addition to this, we also provide a new application of the uniformization technique; we combine uniformization with stratification to create a new variance reduction technique, which, importantly, can be applied at no extra computational cost.


## A. Outline

This paper is structured as follows. In Sec. II, we review the basic aspects of stochastic modeling and simulation of chemical reaction networks. Section III introduces the uniformization method and shows how it can be applied to the stochastic simulation of chemical reaction networks in an efficient way. Building on the uniformization method, Sec. IV introduces a new variance reduction method combining uniformization and stratification. The applicability of this new application of uniformization is demonstrated by examples. Finally, the conclusions can be found in Sec. V.

## II. STOCHASTIC SIMULATION

In this work, we consider stochastic simulation of a model describing the temporal evolution of molecule copy numbers in a well-mixed system of volume  $V$ , i.e., we ignore any spatial information and assume that the molecules are homogeneously distributed within  $V$ . There exists a natural extension of this model that includes spatial information, and we refer to Refs. 19 and 23 for more detail on how to include the spatial movement of molecules in the framework described here.

We consider a collection of  $n$  types of chemical species  $X_1, \dots, X_n$  and define  $\mathbf{X}(t)$  to be the state vector, describing the evolution of all the species as time evolves, i.e.,  $X_i(t)$  is the number of molecules  $X_i$  at time  $t$ . We assume that the system starts with  $\mathbf{X}(0) = \mathbf{X}_0$  molecules, though it is also possible to describe an initial distribution of particles. These species  $X_i$  can interact via  $K$  different types of reactions  $R_1, \dots, R_K$ , referred to as reaction channels. In general form, such a reaction  $R_k$  can be written as



where  $\alpha_{i,k}, \beta_{i,k} \in \mathbb{N}$  and  $c_k$  is the reaction rate constant for reaction channel  $R_k$ . For each reaction  $R_k$ , we define the stoichiometric vector  $\zeta_k$ , which defines the change in the copy number when reaction  $R_k$  fires, i.e., the  $i$ th component of  $\zeta_k$  equals  $\beta_{i,k} - \alpha_{i,k}$ . If the system is in state  $\mathbf{X}(t)$  and reaction  $R_k$  fires at  $t + \tau$ , the system jumps to the new state  $\mathbf{X}(t + \tau) = \mathbf{X}(t) + \zeta_k$ .

The temporal evolution of the chemical species can now be described as

$$\mathbf{X}(t) = \mathbf{X}_0 + \sum_{k=1}^K N_k(t) \zeta_k, \quad (1)$$

where  $N_k(t)$  denotes the number of times that reaction channel  $R_k$  fires in the time interval  $[0, t)$ . We associate every reaction channel,  $R_k$ , with a propensity function,  $a_k(\mathbf{X}(t))$ , which describes the

probability that the reaction channel fires in the infinitesimal time interval  $[t, t + dt)$  as follows:

$$\mathbb{P}(R_k \text{ fires in } [t, t + dt)) = a_k(\mathbf{X}(t)) dt. \quad (2)$$

We also define the total reaction propensity  $a_0(\mathbf{X}(t)) = \sum_{k=1}^K a_k(\mathbf{X}(t))$ . For this work, we primarily use “mass action kinetics” under which the propensity for reaction  $R_k$  to fire is proportional to the number of possible combinations of reactants; see, for example, Ref. 27. However, note that the selection of this function is a modeling choice and not essential for anything that follows.

The modeling choice for the counting functions,  $N_k(t)$  for  $k = 1, \dots, K$ , that we follow here is such that the model describing the evolution of  $\mathbf{X}(t)$  is a CTMC. This can be described by the Kurtz random time change representation (RTCR),<sup>4</sup> where the number of times reaction  $R_k$  fires is described by an inhomogeneous Poisson counting process  $Y_k$  so that

$$\mathbf{X}(t) = \mathbf{X}_0 + \sum_{k=1}^K Y_k \left( \int_0^t a_k(\mathbf{X}(s)) ds \right) \zeta_k. \quad (3)$$

Here, the counting processes  $Y_k$  for  $k = 1, \dots, K$  are all independent, reflecting the independence of the firing of reactions. The RTCR (3) yields a pathwise representation of the dynamics of the species in the system. An alternative description of the same dynamics is by use of the chemical master equation (CME), which comprises a (high-dimensional) system of ordinary differential equations (ODEs). The CME fully describes the evolution in time of the distribution of the occupation probability of the state space of  $\mathbf{X}(t)$ . Only for a small class of problems does the CME have a known analytical solution,<sup>30,41</sup> and one therefore often has to rely on approximate methods or stochastic simulation via Monte Carlo methods to explore the behavior of the system. Here, we focus on the stochastic simulation approach and we refer to Ref. 43 for an overview of alternative methods that build on (computational) approximations to the CME, such as the finite state projection method,<sup>24,35,38</sup> moment-closure methods,<sup>12,44</sup> and system-size expansions.<sup>52</sup>

One of the most widely used methods to generate exact realizations of the RTCR (3) is Gillespie’s DM,<sup>22</sup> as depicted in Algorithm 1. The DM constructs sample trajectories which are sampled from the exact distribution described by the CME.

Given an ensemble of simulated sample paths  $\mathbf{X}^{(i)}(t)$  for  $i = 1, \dots, N$ , we want to estimate some desired summary statistic of the system of interest

$$Q = \mathbb{E}[f(\mathbf{X}(t))]. \quad (4)$$

This can be achieved by the usual Monte Carlo estimate combining the sample paths

$$\hat{Q} = \frac{1}{N} \sum_{i=1}^N f(\mathbf{X}^{(i)}(t)). \quad (5)$$

The individual realizations of  $\hat{Q}$  will in general not be equal to  $Q$  and are random variates which therefore have an inherent uncertainty related to them. This uncertainty can be quantified by the mean squared error (MSE)

$$\text{MSE}[\hat{Q}] = \mathbb{V}[\hat{Q}] + \text{Bias}^2[\hat{Q}]. \quad (6)$$

**Algorithm 1.** Gillespie's direct method (DM).

---

**Input:** Initial data  $\mathbf{X}_0$   
**Input:** Stoichiometric matrix  $\zeta$   
**Input:** Propensity functions  $a_k(\mathbf{x})$   
**Input:** Final time  $T$

```
1:  $\mathbf{X} \leftarrow \mathbf{X}_0$ 
2:  $t \leftarrow 0$ 
3: while  $t < T$  do
4:   Generate  $u_1, u_2 \sim \mathcal{U}(0, 1)$ 
5:    $\mathbf{A}_k \leftarrow a_k(\mathbf{X})$ 
6:    $a_0 \leftarrow \sum_k \mathbf{A}_k$ 
7:    $\tau \leftarrow -\log u_1 / a_0$ 
8:    $t \leftarrow t + \tau$ 
9:   if  $t \geq T$  then
10:     $t \leftarrow T$ 
11:    break
12:   end if
13:   Find  $p$  such that  $\sum_{k=1}^{p-1} \mathbf{A}_k < a_0 u_2 \leq \sum_{k=1}^p \mathbf{A}_k$ 
14:    $\mathbf{X} \leftarrow \mathbf{X} + \zeta_p$ 
15: end while
16: return  $\mathbf{X}$ 
```

---

▷ Calculate the propensities.  
▷ Calculate the total reaction propensity.  
▷ Calculate the next reaction time.

▷ Choose next reaction to fire.  
▷ Update state vector.

---

The sample paths from the DM are unbiased, and the MSE is therefore purely determined by the variance of the estimator  $\hat{Q}$ . Given a sample variance  $\sigma^2$  of the desired summary statistic, we know that  $\mathbb{V}[\hat{Q}] = \sigma^2/N$  with  $N$  the number of sample paths (Ref. 32, Chap. 1).

For a fixed amount of computational resource, we would like to have a MSE that is as small as possible. A drawback of the DM is that it generates each and every reaction in the system individually, which can incur large computational costs for systems with many particles and/or many reactions, or that it runs for a long time. In that particular case, the computational budget is only sufficient for a low number of sample paths,  $N$ , resulting in a high variance of the estimator  $\hat{Q}$ . There exists a variety of (computational) modifications to the DM<sup>1,15,20,37,56</sup> which aim to improve the runtime by optimizing certain steps, predominantly steps 5 and 8, of Algorithm 1 without affecting the statistically exact sampling. These methods therefore result in a smaller MSE for a given computational budget by increasing the number of sample paths,  $N$ , that can be generated in fixed time. An alternative approach is to allow the sampling to be only approximately correct in order to speed up computations, and this yields methods such as Gillespie's  $\tau$ -leaping<sup>21</sup> and  $R$ -leaping,<sup>6</sup> which introduce a bias into the MSE by virtue of only being approximate methods. A computational complexity analysis of some of the standard SSAs such as Gillespie's DM and  $\tau$ -leaping in the classical population scaling can be found in Ref. 3.

As an alternative approach, we can attempt to generate sample paths using a different method which yields a smaller sample variance,  $\sigma^2$ . This class of techniques is known as variance reduction methods (Ref. 32, Chap. 4). In the context of the CME, several variance reduction methods exist, such as QMC methods<sup>8,26</sup> and MLMC methods.<sup>2,9,33</sup> In this work, we will demonstrate how

the application of the uniformization technique can lead to a new variance reduction method for the simulation of chemical reaction networks.


### III. UNIFORMIZATION

The uniformization method is a well-known method in probability theory, and it can be used to convert a CTMC into a discrete-time Markov chain (DTMC). For an overview of uniformization from a probability perspective, we refer to Ref. 51. Here, however, we describe how to employ the uniformization technique in the context of the simulation of exact sample paths according to the RTCR, which has previously appeared, albeit in slightly different forms and with different motivation, in Refs. 5, 26, 42, 46, 54, and 57. For simplicity, we make the assumption that the reaction propensities,  $a_k$ , have no explicit time-dependence and only depend on time through the state vector  $\mathbf{X}(t)$ . This assumption restricts the class of problems to time-homogeneous Markov chains and thus excludes, for example, time-dependent reaction rates. It is, however, possible to extend most of the results in this work by relaxing this condition, as we show in Appendix A.

We start with the original system from Sec. II containing  $K$  reaction channels and consider an extension of this system by adding a new reaction channel  $R_{K+1}$  which takes the trivial form



and we denote this as a virtual reaction. It should be clear that the addition of this virtual reaction does not change the dynamics of  $\mathbf{X}$  because the new channel has none of the species as the reactant or product. We note that, as a result, the statistics from this new extended system are equal to that of the original system and therefore sample paths for the extended system are exact realizations of



**FIG. 1.** Illustration of reactions  $R_1, \dots, R_K$  firing in the original system and the uniformized/extended system, indicated by dots (•). We observe extra virtual reactions firing (•) in the uniformized system.

the original RTCR (3). This observation is independent of the reaction propensity  $a_{K+1}(\mathbf{X}(t))$  that we choose for the virtual reaction. Given some  $\bar{a} > 0$ , we therefore use this liberty to set

$$a_{K+1}(\mathbf{X}(t)) = \bar{a} - a_0(\mathbf{X}(t)) = \bar{a} - \sum_{k=1}^K a_k(\mathbf{X}(t)). \quad (8)$$

In order for (8) to constitute a well-defined reaction propensity, we need  $\bar{a} \geq a_0(t)$  for all  $t \in [0, T]$ , the time interval of interest. We will come back to discuss this assumption later in this section.

The choice (8) for the reaction propensity of the reaction channel  $R_{K+1}$  might seem peculiar at first, but we note that in the new extended system, the total propensity of a reaction happening is given by  $a_0 + a_{K+1} = \bar{a}$ , which is therefore independent of the particular state  $\mathbf{X}(t)$  the system is in. This is in contrast to the original system where the total reaction propensity,  $a_0$ , is generally state dependent. We therefore end up with a uniform total reaction propensity  $\bar{a}$  which we will call the uniformization rate, and we will denote the extended system consisting of reactions  $R_1, \dots, R_K, R_{K+1}$  from now on as the uniformized system. This uniformization of the system has a few implications which we will discuss next.

First, in the uniformized system (trivially), at least as many reactions fire as in the original system because of the addition of the

independent virtual reaction channel, as illustrated in Fig. 1. Naïve application of the DM to the uniformized system means we explicitly simulate every virtual reaction firing in addition to the real reactions present in the original system. As a result of the extra time taken to simulate the virtual reaction channel firing, such naïve application of the DM to the uniformized system yields slower run-times for exactly the same level of statistical accuracy compared to the DM for the original system.

Second, the reaction times in the uniformized system are independent and identically distributed (i.i.d.) exponential random variables with parameter  $\bar{a}$ . As a simple consequence of this, we observe that the number of reactions,  $M$ , firing in a time interval  $[t, t + \tau)$  is Poisson distributed with parameter  $\bar{a}\tau$ , which we will denote as  $M \sim \mathcal{P}(\bar{a}\tau)$ . Note that this is similar to the  $\tau$ -leaping method, where one makes the approximation that the reaction propensities of the original system stay constant within a time interval  $[t, t + \tau)$ . For the uniformized system, however, this expression for the number of firing reactions is instead exact and leads to a uniformized version of the DM (UDM), as shown in Algorithm 2.

In the UDM, there is no explicit computation of the next reaction time (cf. step 7 in Algorithm 1). For every reaction firing, the UDM only needs to generate a single uniform random variate to determine which reaction takes place and, therefore, involves fewer computations. As a result, the UDM will run faster than the DM applied to the uniformized system. It is possible to generate the reaction times, if required, by noting that we can condition on the fact that  $M$  reactions happen in  $[0, T)$  with a uniform rate. The reaction times are then distributed as the order statistics of  $M$  uniform random variables in  $[0, T)$ .

We observe, however, that in the current form, the UDM still suffers from two drawbacks. First, the simulation will in general involve the firing of virtual reactions that do not contribute to the dynamics of the original system. This can therefore be thought of as computational waste and, as a result, the runtime of the UDM in its current form will still be comparatively larger than the runtime of the DM applied to the original system for time-homogeneous systems. Second, we mentioned that (8) needs to be a well-defined propensity

---

**Algorithm 2.** Uniformized direct method (UDM).

---

**Input:** Initial data  $\mathbf{X}_0$

**Input:** Stoichiometric matrix  $\zeta$

**Input:** Propensity functions  $a_k(\mathbf{x})$

**Input:** Uniformization rate  $\bar{a}$

**Input:** Final time  $T$

```

1: Generate  $M \sim \mathcal{P}(\bar{a}T)$ 
2: for  $m = 1, \dots, M$  do
3:   Generate  $u_1 \sim \mathcal{U}(0, 1)$ 
4:    $\mathbf{A}_k \leftarrow a_k(\mathbf{X})$ 
5:    $\mathbf{A}_{K+1} \leftarrow \bar{a} - \sum_{k=1}^K \mathbf{A}_k$ 
6:   Find  $p$  such that  $\sum_{k=1}^{p-1} \mathbf{A}_k < \bar{a}u_1 \leq \sum_{k=1}^p \mathbf{A}_k$ 
7:   if  $p \in \{1, \dots, K\}$  then
8:      $\mathbf{X} \leftarrow \mathbf{X} + \zeta_p$ 
9:   end if
10: end for
```

---

▷ Total number of reactions that fire in  $[0, T)$ .

▷ Calculate real reaction propensities.

▷ Calculate virtual reaction propensity.

▷ Choose next reaction to fire.

▷ Only need to fire real reactions.

▷ Update state vector.

---

function, i.e.,  $\bar{a} \geq a_0(t)$  needs to hold for all  $t \in [0, T]$ . It is not clear *a priori* whether such a uniformization rate  $\bar{a}$  exists<sup>58</sup> or what happens when  $a_0$  becomes greater than  $\bar{a}$  in the course of a simulation. These two issues are discussed next, and we will show that the UDM can be adapted to be at least as fast as the DM applied to the original system.

### A. Firing virtual reactions

In order to get around the issue of potentially slowing down the simulation by having to fire many virtual reactions, we look at the distribution of the number of virtual reactions firing in between real reactions in the uniformized system.

Suppose the system is in a state with propensity  $a_0$  for the real reactions to fire. If the system is uniformized with rate  $\bar{a} > a_0$ , this means that the probability that the next reaction firing belongs to one of the  $K$  real reaction channels is given by  $a_0/\bar{a}$  and, equivalently, the probability a virtual reaction will fire next is  $1 - a_0/\bar{a}$ . Note that when a virtual reaction fires, none of the propensities of the real reactions change because none of the species  $X_i$  are changed in the reaction. The repeated firing of the virtual reaction channel before a real reaction fires can therefore be viewed as a series of Bernoulli trials with probability  $a_0/\bar{a}$  of success (firing a real reaction) and  $1 - a_0/\bar{a}$  of failure (firing a virtual reaction). In this scenario, we are interested in the number of failures until the first success, i.e., the number of consecutive virtual reactions firing before a real reaction fires. This quantity has a well-known distribution, namely, the geometric distribution, so that we have

$$\mathbb{P}(r \text{ consecutive virtual reactions before next real reaction fires}) = \left(1 - \frac{a_0}{\bar{a}}\right)^r \frac{a_0}{\bar{a}}. \quad (9)$$

As a result, it is possible to fire all consecutive virtual reactions at once by sampling a single geometric random variable. This can be done efficiently by generating a uniform random variable  $u \sim \mathcal{U}(0, 1)$  and calculating  $\lceil \ln(u) / \ln(1 - a_0/\bar{a}) \rceil$ , akin to the sampling of an exponential random variable. This observation leads to an improved version of the UDM, as depicted in Algorithm 3.

In this form, the improved UDM is comparable with the DM applied to the original system, i.e., the gold standard in the field of simulation chemical reaction networks. Both methods now fire an equal number of reactions and need two random numbers per reaction firing. In addition to this, we note that many improvements that have been made to the DM, related to speed-ups in the choice of the next reaction firing and the update of the propensities such as in Refs. 15, 20, 37, and 48 and/or the reuse of random numbers,<sup>56</sup> can be equally well applied to the improved UDM. Although it is not possible to use the well-known  $\tau$ -leaping approach in combination with the (improved) UDM, it is trivial to apply the  $R$ -leap approximation<sup>6</sup> to the (improved) UDM. The improved version of the UDM is also insensitive to the uniformization rate,  $\bar{a}$ , as can be seen in Fig. 2. This can be intuitively understood from the observation that increasing or decreasing  $\bar{a}$  only changes the number of consecutive virtual reactions firing which is taken care of in a single step in the improved UDM.

---

#### Algorithm 3. Improved uniformized direct method.

---

```

Input: Initial data  $\mathbf{X}_0$ 
Input: Stoichiometric matrix  $\boldsymbol{\zeta}$ 
Input: Propensity functions  $a_k(\mathbf{x})$ 
Input: Uniformization rate  $\bar{a}$ 
Input: Final time  $T$ 
1: Generate  $M \sim \mathcal{P}(\bar{a}T)$ .
2:  $m \leftarrow 0$ 
3: while  $m < M$  do
4:   Generate  $u_1 \sim \mathcal{U}(0, 1)$ .
5:    $\mathbf{A}_k \leftarrow a_k(\mathbf{X})$ 
6:    $a_0 \leftarrow \sum_{k=1}^p \mathbf{A}_k$ 
7:   Generate  $m_{\text{virtual}} \sim \text{Geometric}(a_0/\bar{a})$ 
8:    $m \leftarrow m + m_{\text{virtual}}$ 
9:   if  $m \geq M$  then
10:     $m \leftarrow M$ 
11:    break
12:   end if
13:   Find  $p$  such that  $\sum_{k=1}^{p-1} \mathbf{A}_k < a_0 u_1 \leq \sum_{k=1}^p \mathbf{A}_k$ 
14:    $\mathbf{X} \leftarrow \mathbf{X} + \boldsymbol{\zeta}_p$ 
15:    $m \leftarrow m + 1$ 
16: end while


```

---

▷ Total number of reactions that fire in  $[0, T]$ .  
 ▷ Counter for number of reactions that have fired.  
 ▷ Calculate real reaction propensities.  
 ▷ Calculate the total real reaction propensity.  
 ▷ Number of virtual reactions firing consecutively.  
 ▷ Choose next reaction to fire.

---






**FIG. 2.** CPU time to generate  $N = 10^5$  sample paths with the UDM and improved UDM compared with the DM for different uniformization rates  $\bar{a}$  relative to the initial total propensity  $a_0$ . Timings are for the example in Sec. IV C 1 with  $c_1 = 0.2$ ,  $c_2 = 0.1$  and  $\mathbf{X}_0 = (10, 0)$ , which was run until  $T = 10$ . Timing experiments were performed using MATLAB R2017b on an Ubuntu desktop PC with a 3.40 GHz Intel Core i7-2600K CPU and 16 GB of random access memory.

We note from Fig. 2 and the description of the improved UDM that the computational run-time of simulations of the uniformized system can be brought back to the speed of the DM applied to the original system. However, there is no intrinsic reason to expect the improved UDM to be faster than the DM for the original system<sup>59</sup> and this is also observed in Fig. 2. A computational speed-up, as such, is therefore not a sufficient motivation to employ the uniformization technique in the stochastic simulation of chemical reaction networks. However, as mentioned earlier, the usage of uniformization allows one to consider new applications, such as variance reduction methods, that are not possible under the standard SSAs for the original system.

## B. Adapting the uniformization rate

In the exposition so far, the uniformization rate has been a free parameter and we mentioned the caveat that, by expression (8),  $\bar{a} \geq a_0(t)$  should hold for all  $t \in [0, T)$  in order for the system to remain a well-defined chemical reaction model. We therefore discuss here the ramifications of a simulation in which the total propensity  $a_0(t)$  exceeds  $\bar{a}$ .

A first approach to the situation  $a_0 > \bar{a}$  in the course of a simulation might be to classify such sample paths as invalid and generate a new, independent sample path for each of the invalid paths until we end up with the desired number of sample paths that all satisfy  $a_0 \leq \bar{a}$  across  $[0, T)$ . This approach, unfortunately, comes at a cost because it introduces a bias due to the new sampling strategy. Effectively, this approach only samples from a subset of the distribution defined by the CME and it is clear from Fig. 3 that the effect of the resulting bias can be dramatic and should be avoided if possible. We therefore propose a way to keep samples for which the uniformization rate is breached, whilst keeping the estimator unbiased at the same time.




**FIG. 3.** Bias effect due to invalid samples of the uniformized system for different uniformization rates  $\bar{a}$  relative to the initial total propensity  $a_0$ . The model system used is the example in Sec. IV C 1 with  $c_1 = 0.4$ ,  $c_2 = 0.1$  and  $\mathbf{X}_0 = (0, 20)$ , which was run until  $T = 10$ . The summary statistic is the number of  $X_1$  molecules at final time  $T$ , and estimators are shown with 99.5% confidence interval for  $N = 10^3$  samples.

To do so, we start with a uniformized system with uniformization rate  $\bar{a}$  and sample the number of reactions  $M \sim \mathcal{P}(\bar{a}T)$  firing in  $[0, T)$ . Now, suppose that we observe that  $a_0 > \bar{a}$  after  $M^* \leq M$  reactions have fired. Rather than discarding this sample path completely, we can ask at what time  $T^* \in [0, T)$  reaction  $M^*$  took place, i.e., at what time did  $a_0$  become larger than  $\bar{a}$ . As mentioned earlier, given that  $M$  reactions fire in  $[0, T)$ , the reaction times  $t_1, \dots, t_M$  are uniformly distributed on  $[0, T)$ . The reaction time  $T^* = t_{M^*}$  of the  $M^*$ th reaction out of  $M$  reactions therefore follows the distribution of the  $M^*$ th order statistic of a collection of  $M$  uniform random variables on  $[0, T)$ . This is a well-known distribution and leads to  $T^*/T \sim \text{Beta}(M^*, (M - M^*) + 1)$ .

We can therefore sample  $T^*$  and use the Markov property to restart the simulation from the state after  $M^*$  reactions for the remaining time interval  $[T^*, T)$  with a new uniformization rate,  $\bar{a}_{\text{new}}$ . This procedure to describe  $T^*$  conditional on  $M$  reactions in  $[0, T)$  leads to the correct unconditional distribution for  $T^*$  and therefore does not introduce a bias; see Appendix B for more detail. If we choose this new uniformization rate such that  $\bar{a}_{\text{new}} \geq a_0(T^*)$ , the simulation for  $[T^*, T)$  can be done in exactly the same manner with the (improved) UDM, and this approach is illustrated in Fig. 4.

The approach of changing the uniformization rate,  $\bar{a}$ , as described in Fig. 4, can be applied throughout the simulation and yields a method akin to adaptive uniformization.<sup>53</sup> We note that the adaptation of the uniformization rate is not restricted to time points  $T^*$  at which  $a_0(T^*) > \bar{a}$  is observed; adapting the uniformization rate can be done at any point in the simulation, as shown in Algorithm 4 and yields unbiased sample paths. For example, one could have a system that initially has a high total propensity,  $a_0$ , requiring a high uniformization rate,  $\bar{a}$ , but after  $M^*$  reactions,  $a_0$  dips considerably. In this scenario, one might want to adapt the uniformization rate and make it smaller. Such an approach is similar to the adaptive



**FIG. 4.** Illustration of the procedure to adapt the uniformization rate. (a) With uniformization rate  $\bar{a}$ , we sample  $M$  reactions to fire in  $[0, T)$ . (b) After  $M^*$  reactions, we want to adapt the uniformization rate, for example, because the uniformization rate is breached,  $a_0 > \bar{a}$ . (c) We sample the time  $T^*$  at which to adapt the uniformization rate and restart the simulation at  $T^*$  with a new uniformization rate  $\bar{a}_{\text{new}}$  and  $M_{\text{new}} \sim \mathcal{P}(\bar{a}_{\text{new}}(T - T^*))$ .

fluctuation intervals as used in the rejection-based SSAs.<sup>47,48,50</sup> An even more extreme choice might be to update the uniformization rate after every reaction, which would yield a procedure akin to the Extrande-method.<sup>54</sup>

We finally mention two observations regarding the adaptive uniformization method. First, the sample paths in this framework are again exact samples according to the distribution described by the CME and can therefore be used to create unbiased summary statistics. Second, from the formulation of the improved UDM in Sec. III A, we can see that the choice of  $\bar{a}$  ultimately does not influence the run-time of the simulation. One might therefore be tempted to take  $\bar{a}$  very large in the hope that  $\bar{a} > a_0(t)$  for all  $t \in [0, T)$ , and we never have to adapt the uniformization rate  $\bar{a}$ . However, for some choices of the reaction propensities, such as mass action kinetics, it is not always possible to select such an upper bound for  $a_0$  as a uniformization rate,  $\bar{a}$ , because the reaction propensities are unbounded. This is, for example, the case in a system where at least one of the species involved in mass action kinetics is unbounded itself, a situation that is often encountered. In Ref. 26, it was suggested to run a few presimulations to determine an empirical upper bound to  $a_0$  on  $[0, T)$  as a workaround if we do not know an analytical expression for an upper bound. Alternatively, one might have other motives to not choose the uniformization rate much higher than  $a_0$ ; an example of such a situation will be given in Sec. IV. To create unbiased estimators with the (improved) UDM, it can therefore be necessary to employ an adaptive uniformization method, as described in this section.

In conclusion, we can create an (improved) UDM that is equally as versatile and applicable as the DM. The main benefit of uniformization, however, lies in the fact that it opens the door to

**Algorithm 4.** Adaptive improved uniformized direct method.

---

**Input:** Initial data  $\mathbf{X}_0$   
**Input:** Stoichiometric matrix  $\zeta$   
**Input:** Propensity functions  $a_k(\mathbf{x})$   
**Input:** Uniformization rate  $\bar{a}$   
**Input:** Final time  $T$

- 1: Generate  $M \sim \mathcal{P}(\bar{a}T)$  ▷ Total number of reactions that fire in  $[0, T)$ .
- 2:  $m \leftarrow 0$  ▷ Counter for number of reactions that have fired.
- 3: **while**  $m < M$  **do**
- 4:   Generate  $u_1 \sim \mathcal{U}(0, 1)$
- 5:    $\mathbf{A}_k \leftarrow a_k(\mathbf{X})$  ▷ Calculate real reaction propensities.
- 6:    $a_0 \leftarrow \sum_k \mathbf{A}_k$  ▷ Calculate the total real reaction propensity.
- 7:   Set `adapt_flag` ▷ For example, when  $a_0 > \bar{a}$  we set `adapt_flag = true`.
- 8:   **if** `adapt_flag` **then** ▷ Option to adapt the uniformization rate.
- 9:     Generate  $T^*/T \sim \text{Beta}(m, (M - m) + 1)$
- 10:    Adapt the uniformization rate  $\bar{a}$
- 11:    Generate  $\tilde{M} \sim \mathcal{P}(\bar{a}(T - T^*))$  ▷ Total number of reactions that fire in  $[T^*, T)$ .
- 12:     $M \leftarrow m + \tilde{M}$
- 13:   **end if**
- 14:   Generate  $m_{\text{virtual}} \sim \text{Geometric}(a_0/\bar{a})$  ▷ Number of virtual reactions firing consecutively.
- 15:    $m \leftarrow m + m_{\text{virtual}}$
- 16:   **if**  $m \geq M$  **then**
- 17:      $m \leftarrow M$
- 18:     **break**
- 19:   **end if**
- 20:   Find  $p$  such that  $\sum_{k=1}^{p-1} \mathbf{A}_k < a_0 u_1 \leq \sum_{k=1}^p \mathbf{A}_k$  ▷ Choose next reaction to fire.
- 21:    $\mathbf{X} \leftarrow \mathbf{X} + \zeta_p$
- 22:    $m \leftarrow m + 1$
- 23: **end while**

---



new applications that are not possible in the DM framework, as we will see next.

#### IV. VARIANCE REDUCTION VIA STRATIFICATION

Using the (improved) UDM will yield exact sample paths from the distribution defined by the CME, just like the DM. The resulting MSE of the standard Monte Carlo estimates based on these sample paths will therefore behave exactly like the MSE for the DM. In this section, however, we will show that it is possible to create a new (unbiased) estimator that has a lower sample variance, and therefore MSE, by using a stratification strategy in combination with uniformization.

##### A. Stratification basics

Suppose we are interested in estimating a summary statistic  $Q = \mathbb{E}[f(\mathbf{Y})]$  for some process  $\mathbf{Y} \sim p$ . For example, this could be a simple mean, in which case we consider  $f(x) = x$ , or the probability distribution evaluated at a point  $y$ , for which we would take  $f_y(x) = \mathbb{1}_{\{x=y\}}$ . Rather than directly sampling from  $p$ , we first divide the state space  $\Omega$  of  $\mathbf{Y}$  in a set of disjoint sets such that  $\Omega = \bigcup_{j=1}^J \Omega_j$ , where  $J$  can be infinite. These sets will be denoted as strata. We will use the notation  $\omega_{\mathcal{D}_j} = \mathbb{P}(\mathbf{Y} \in \Omega_j)$  and define  $\mathcal{D}_j$  to be the event that  $\mathbf{Y} \in \Omega_j$ . In the case that the distribution of  $\mathbf{Y}$  is time-dependent, i.e.,  $\mathbf{Y} \sim p(t)$ , note that  $\omega_{\mathcal{D}_j}(t)$  will also be time-dependent. With this notation and using the fact that the sets are disjoint, we can use conditional expectation and the law of total expectation to write

$$Q = \mathbb{E}[f(\mathbf{Y})] = \sum_{j=1}^J \omega_{\mathcal{D}_j} \mathbb{E}[f(\mathbf{Y})|\mathcal{D}_j] = \sum_{j=1}^J \omega_{\mathcal{D}_j} Q_{\mathcal{D}_j}, \quad (10)$$

where  $Q_{\mathcal{D}_j} = \mathbb{E}[f(\mathbf{Y})|\mathcal{D}_j]$  denotes the summary statistic conditional on the event  $\mathcal{D}_j$ . If  $\omega_{\mathcal{D}_j}$  are known, this suggests a new way to construct an estimator of  $Q$

$$\hat{Q}_{\text{str}} = \sum_{j=1}^J \omega_{\mathcal{D}_j} \hat{Q}_{\mathcal{D}_j}, \quad (11)$$

where  $\hat{Q}_{\mathcal{D}_j}$  are estimators of  $Q_{\mathcal{D}_j}$ . To construct these conditional estimators, we denote  $\mathbf{Y}^{(i,j)}$  to be the  $i$ th sample drawn from the conditional distribution of  $(\mathbf{Y}|\mathcal{D}_j)$  which yields

$$\hat{Q}_{\mathcal{D}_j} = \frac{1}{N_j} \sum_{i=1}^{N_j} f(\mathbf{Y}^{(i,j)}). \quad (12)$$

Note that if conditional sampling from  $(\mathbf{Y}|\mathcal{D}_j)$  can be carried out exactly,  $\hat{Q}_{\mathcal{D}_j}$  are unbiased estimators of  $Q_{\mathcal{D}_j}$  and, as a result,  $\hat{Q}_{\text{str}}$  is an unbiased estimator of the summary statistic  $Q$ .

The main benefit of this stratified sampling approach, however, lies in the fact that by a judicious choice of the number of samples,  $N_j$ , per stratum, we can make sure that  $\hat{Q}_{\text{str}}$  has a lower variance than the standard Monte Carlo estimator  $\hat{Q}$ . To see this, note that the law of total variance decomposes the variance of our summary statistic of interest,  $f(\mathbf{Y})$ , as

$$\sigma^2 = \mathbb{V}[f(\mathbf{Y})] = \mathbb{E}[\mathbb{V}[f(\mathbf{Y})|\mathcal{D}_j]] + \mathbb{V}[\mathbb{E}[f(\mathbf{Y})|\mathcal{D}_j]]. \quad (13)$$

If we introduce  $\sigma_j^2$  as the conditional variance of  $f(\mathbf{Y})$  given  $\mathcal{D}_j$ , we can rewrite (13) as

$$\sigma^2 = \sum_{j=1}^J \omega_{\mathcal{D}_j} \sigma_j^2 + \sum_{j=1}^J \omega_{\mathcal{D}_j} (Q - Q_{\mathcal{D}_j})^2. \quad (14)$$

As a result, we see that the sample variance of the standard Monte Carlo estimator is given by

$$\mathbb{V}[\hat{Q}] = \frac{1}{N} \sum_{j=1}^J \omega_{\mathcal{D}_j} \sigma_j^2 + \frac{1}{N} \sum_{j=1}^J \omega_{\mathcal{D}_j} (Q - Q_{\mathcal{D}_j})^2, \quad (15)$$

where  $N$  again denotes the number of sample paths used. For example, this would be the sample variance using samples from the DM to calculate the estimate of the summary statistic given by (5). On the other hand, for the stratified estimator, we have sample variance

$$\mathbb{V}[\hat{Q}_{\text{str}}] = \sum_{j=1}^J \omega_{\mathcal{D}_j}^2 \frac{\sigma_j^2}{N_j}. \quad (16)$$

This expression demonstrates the fact that in order to reduce the sample variance of a stratified estimator, we need to carefully specify how many samples will be used per stratum,  $\mathcal{D}_j$ . Perhaps the simplest and most common choice is “proportional allocation,” i.e., given a budget of  $N$  samples in total, we set  $N_j = \omega_{\mathcal{D}_j} N$ . For this choice, the stratified estimator has a sample variance that is guaranteed to be at least as small as the standard Monte Carlo sample variance

$$\begin{aligned} \mathbb{V}[\hat{Q}_{\text{prop}}] &= \sum_{j=1}^J \omega_{\mathcal{D}_j}^2 \frac{\sigma_j^2}{N \omega_{\mathcal{D}_j}} \leq \frac{1}{N} \sum_{j=1}^J \omega_{\mathcal{D}_j} \sigma_j^2 + \frac{1}{N} \sum_{j=1}^J \omega_{\mathcal{D}_j} (Q - Q_{\mathcal{D}_j})^2 \\ &= \mathbb{V}[\hat{Q}], \end{aligned} \quad (17)$$

where  $\hat{Q}_{\text{prop}}$  is given by (11) and (12) with  $N_j = \omega_{\mathcal{D}_j} N$ . It is therefore clear that a judicious stratification strategy, such as proportional allocation, can in fact be a variance reduction technique. Other choices of sample allocation can be made such as optimal allocation and poststratification, leading to different sample variances. We will not discuss such strategies further here but refer the reader to Ref. 32, Chap. 4.

As a final note, we provide a method for estimating the sample variance of the stratified estimator (11), which can be used to construct confidence intervals. Given the conditional estimators as in (12), we can write down the unbiased conditional sample variance estimator

$$s_j^2 = \frac{1}{N_j - 1} \sum_{i=1}^{N_j} (f(\mathbf{Y}^{(i,j)}) - \hat{Q}_{\mathcal{D}_j})^2. \quad (18)$$

An unbiased estimator of the sample variance of the stratified estimator of the summary statistic  $\hat{Q}_{\text{str}}$  is then given by

$$s^2 = \sum_{j=1}^J \omega_{\mathcal{D}_j}^2 \frac{s_j^2}{N_j}. \quad (19)$$

##### B. Stratification of the number of reactions

Stratification is a technique that can be used to incorporate some exact knowledge of part of the stochastic process into the required estimator so as to lower the sample variance. In the case

of chemical reaction networks, we can apply stratification ideas if we first uniformize the system with some rate  $\bar{a}$ . We can think of the uniformized system in terms of an extended state space of the form  $\mathbf{Y} = (\mathbf{X}, M)$  with  $\mathbf{X}$  the state of the species after  $M$  reactions have fired. The distribution of  $M$  is known *a priori* (as shown in Sec. III), and therefore, we can stratify with respect to the number of reactions,  $M$ , that fire in the time interval of interest,  $[0, T]$ . Because  $M \in \mathbb{N}_{\geq 0}$ , all we need to specify is a division of the non-negative integers into disjoint sets. For this work, we consider the following stratification strategy:  $\mathcal{D}_j = \{M_j \leq M < M_{j+1}\}$  for some collection of non-negative integers  $M_1 < M_2 < \dots < M_{J+1}$ . This divides the state space of  $\mathbf{Y} = (\mathbf{X}, M)$  by  $\Omega = \cup_{j=1}^J \Omega_j$  with  $\Omega_j = \{(\mathbf{X}, M) \text{ s.t. } M_j \leq M < M_{j+1}\}$ , i.e.,  $\Omega_j$  represents all sample paths where the total number of reactions that have fired lies in the prescribed range  $[M_j, M_{j+1})$ . Note that for this choice, the strata probabilities are given by a simple sum of Poisson probabilities

$$\omega_{\mathcal{D}_j}(t) = \sum_{m=M_j}^{M_{j+1}-1} \frac{(\bar{a}t)^m}{m!} e^{-\bar{a}t}. \quad (20)$$

To construct the conditional estimators  $\hat{Q}_{\mathcal{D}_j}$ , we need to be able to construct sample paths conditional on  $\mathcal{D}_j$ . This can be easily achieved by replacing step 1 of the (improved) UDM with drawing  $M$  from the truncated Poisson distribution on  $[M_j, M_{j+1})$

rather than drawing  $M \sim \mathcal{P}(\bar{a}T)$ , as illustrated in Algorithm 5.<sup>60</sup> Therefore, there is no extra cost associated with this stratification strategy compared to the original (improved) UDM and, as a result, the computational cost for the stratified estimator is comparable to that of the (improved) UDM and the DM. However, at the same time, it is guaranteed to have a sample variance that is at least as small as that of the standard estimator from the DM and (improved) UDM. Before applying this new estimator to some examples, we mention a few caveats and observations.

First, in order for the stratification method to work, it is important that the uniformization rate  $\bar{a}$  is in fact a valid uniformization rate over the interval of interest  $[0, T]$ , i.e.,  $\bar{a} \geq a_0(t)$  for all  $t \in [0, T]$ . If this is not the case and we are required to adapt the uniformization rate, the samples  $(\mathbf{X}, M)$  are not all drawn from the same distribution and cannot be combined to yield the conditional estimators  $\hat{Q}_{\mathcal{D}_j}$ . It is therefore necessary either to know an upper bound to the total propensity *a priori* or to run some presimulations to generate an empirical upper bound.

Second, the effectiveness of stratification will hinge on the choice of strata. If we use the proportional allocation strategy, we see that the variance reduction benefit of stratification over standard Monte Carlo comes from the  $\frac{1}{N} \sum_{j=1}^J \omega_{\mathcal{D}_j} (Q - Q_{\mathcal{D}_j})^2$  term that is lacking in the sample variance. This term represents the inter-strata variance, i.e., how much the conditional summary statistics

**Algorithm 5.** Stratification with the improved uniformized direct method.

---

**Input:** Initial data  $\mathbf{X}_0$   
**Input:** Stoichiometric matrix  $\zeta$   
**Input:** Propensity functions  $a_k(\mathbf{x})$   
**Input:** Uniformization rate  $\bar{a}$   
**Input:** Final time  $T$   
**Input:** Number of samples  $N$   
**Input:** Stratum definitions  $M_1, \dots, M_{J+1}$

- 1: **for**  $j = 1, \dots, J$  **do**
- 2:    $N_j \leftarrow \lceil \omega_{\mathcal{D}_j}(T)N \rceil$  ▷ Proportional allocation for number of samples per stratum
- 3:   **for**  $n = 1, \dots, N_j$  **do**
- 4:     Generate  $M$  from the truncated distribution  $\mathcal{P}(\bar{a}T)$  on  $[M_j, M_{j+1})$ .
- 5:      $m \leftarrow 0$
- 6:     **while**  $m < M$  **do**
- 7:       Generate  $u_1 \sim \mathcal{U}(0, 1)$
- 8:        $\mathbf{A}_k \leftarrow a_k(\mathbf{X})$
- 9:        $a_0 \leftarrow \sum_k \mathbf{A}_k$
- 10:       Generate  $m_{\text{virtual}} \sim \text{Geometric}(a_0/\bar{a})$
- 11:        $m \leftarrow m + m_{\text{virtual}}$
- 12:       **if**  $m \geq M$  **then**
- 13:           $m \leftarrow M$
- 14:       **break**
- 15:     **end if**
- 16:     Find  $p$  such that  $\sum_{k=1}^{p-1} \mathbf{A}_k < a_0 u_1 \leq \sum_{k=1}^p \mathbf{A}_k$
- 17:      $\mathbf{X} \leftarrow \mathbf{X} + \zeta_p$ .
- 18:      $m \leftarrow m + 1$
- 19:   **end while**
- 20: **end for**
- 21: **end for**

---

deviate from the overall summary statistic. The larger this interstrata variance is, the bigger the variance reduction gain is. Unfortunately, it is in general not possible to know how to choose the strata so as to attain a relatively large interstrata variance.

In choosing the strata, we have to make a choice as to the number of strata. One might be tempted to take as many strata as possible as this will in theory yield the largest variance reduction. Note that theoretically, we could use an infinite number of strata if we use the strata boundaries  $M_j = (j - 1)$  for all  $j \in \mathbb{N}_{\geq 0}$ . However, in order to get accurate estimators for every stratum, we need at least two samples per stratum, and ideally more. Therefore, increasing the number of strata will also increase the number of sample paths needed to get accurate estimates of, for example, the sample variance. It has been observed that increasing the number of strata beyond six, under some mild assumptions, yields very little extra benefit (Ref. 16, Sec. 5A.8). We will show that this is also true for the simulation of chemical reaction networks in an example in Sec. IV C.

Finally, there is the choice of the strata boundaries. Without knowing the joint distribution of  $(\mathbf{X}, M)$ , we simply choose to use the quantiles of the distribution of  $M$  to define the strata. For example, if we require four strata, we use the 4-quantiles  $q_1, q_2, q_3$  to define  $M_1 = 0, M_2 = q_1, M_3 = q_2, M_4 = q_3$ , and  $M_5 = \infty$ . Note that this choice roughly allocates an equal proportion of the total  $N$  samples to each stratum because the weights  $w_{\mathcal{D}_j}$  are equal.<sup>61</sup> An alternative choice, reminiscent of work in Ref. 26, is to define  $M_1$  and  $M_{J+1}$  such that  $\mathbb{P}(M < M_1 \text{ or } M > M_{J+1}) \leq \varepsilon$ . We can use either  $M_j = M_1 + (j - 1)$  or some other stratification for  $M_1 \leq M \leq M_{J+1}$ . This will yield a biased estimator because the  $\varepsilon$ -tails of the distribution of  $M$  are neglected. However, if  $\varepsilon$  is sufficiently small, this bias will be negligible by construction.

## C. Examples

In this part of the work, we consider two examples to test the variance reduction effects of using stratification in combination with uniformization for chemical reaction networks. The first example is a linear system involving two species, and therefore, it is possible to show some analytical results to complement numerical results. The second example is a mitogen-activated protein kinase (MAPK) cascade model comprising eight species interacting via ten reactions with nonlinear reaction propensities. The results presented here are given only in terms of the variance reduction factor (VRF) and not based on the relative efficiency of the methods because this will depend heavily on the implementation details and the particular computing system used. The variance reduction factor, however, is implementation independent.

### 1. Isomerization

In this example, we consider the simple reversible isomerization of species  $X_1$  into  $X_2$




Due to the linear nature of the reactions, we can write down an analytical expression for the probability distribution  $\mathbb{P}(\mathbf{X}, t)$ .<sup>30</sup> We note that the state space of  $\mathbf{X}$  is bounded and as a result there exists a uniformization rate valid for all  $T > 0$ , e.g.,  $\bar{a} = \max(c_1, c_2) \cdot (X_1(0) + X_2(0))$ . Furthermore, because it is a simple linear system,

we can find the joint distribution of  $(X, M)$  numerically. This is achieved by writing down the transition matrix for the DTMC of the uniformized system. This allows one to numerically evaluate (14) and therefore give an analytical value for the variance reduction that can be attained. We will use the following notation to denote the variance reduction factor (VRF):

$$\alpha = \frac{\mathbb{V}[\hat{Q}]}{\mathbb{V}[\hat{Q}_{\text{str}}]}. \quad (22)$$

We first look at the influence of the uniformization rate  $\bar{a}$  on the VRF. Because the system is linear and effectively monomolecular, we can calculate the VRF for both proportional allocation and optimal allocation. The latter method maximizes the VRF but for general systems is not a practical method. It does, however, provide an upper bound for the VRF for any stratification strategy and is therefore included here. As can be seen in Fig. 5, the VRF is moderate and decreases monotonically as a function of  $\bar{a}$ . Furthermore, the VRF depends on the number of strata,  $J$ , used and, as expected, more strata entail a larger VRF. Note, however, that as expected, the VRF is always larger than unity, meaning that regardless of  $\bar{a}$ , there is always a (slight) variance reduction. One might be tempted to take a lower value of  $\bar{a}$  than that drawn in the hope of getting larger VRFs. Unfortunately, there is a balance between the uniformization rate needing to be large enough so as to not bias the simulations and keeping  $\bar{a}$  small enough to see a more significant VRF.

To see the effect of varying the number of strata, we look at the same system that was used to generate Fig. 5 but fix the uniformization rate  $\bar{a}$  whilst we vary the number of strata,  $J$ . The results are presented in Table I and clearly show that increasing the number of strata beyond a moderate number like  $J = 6$  has only a marginal effect. We can also observe that there seems to be an inverse relationship between the uniformization rate,  $\bar{a}$ , and the variance reduction factor of the form  $\alpha = 1 + A(J)/\bar{a}$  for some constant  $A(J)$  depending on the number of strata.



**FIG. 5.** Variance reduction factor for different uniformization rates  $\bar{a}$  relative to the initial total propensity  $a_0$ . The example in Sec. IV C 1 was run with  $c_1 = 0.3$ ,  $c_2 = 0.1$  and  $\mathbf{X}_0 = (20, 0)$  until  $T = 5$ . The summary statistic is the number molecules  $X_1$  at final time  $T$ .

**TABLE I.** Effect of the number of strata,  $J$ , and the uniformization rate,  $\bar{a}$ , on the VRF. Tabulated is  $\alpha - 1$  (higher is better) for the isomerization model (21) with  $c_1 = 0.3$ ,  $c_2 = 0.1$  and  $\mathbf{X}_0 = (20, 0)$ , run until  $T = 5$ . The last rows,  $J = \infty$ , indicate the maximal VRF that can be achieved for both proportional allocation and optimal allocation, respectively.

Number of strata $J$	Relative uniformization rate $\bar{a}/a_0(\mathbf{X}_0)$		
	1	10	100
2	$7.9 \times 10^{-2}$	$7.7 \times 10^{-3}$	$7.7 \times 10^{-4}$
4	$1.2 \times 10^{-1}$	$1.1 \times 10^{-2}$	$1.0 \times 10^{-3}$
6	$1.3 \times 10^{-1}$	$1.1 \times 10^{-2}$	$1.1 \times 10^{-3}$
12	$1.4 \times 10^{-1}$	$1.2 \times 10^{-2}$	$1.2 \times 10^{-3}$
$\infty$ prop.	$1.5 \times 10^{-1}$	$1.2 \times 10^{-2}$	$1.2 \times 10^{-3}$
$\infty$ opt.	$1.5 \times 10^{-1}$	$1.2 \times 10^{-2}$	$1.2 \times 10^{-3}$

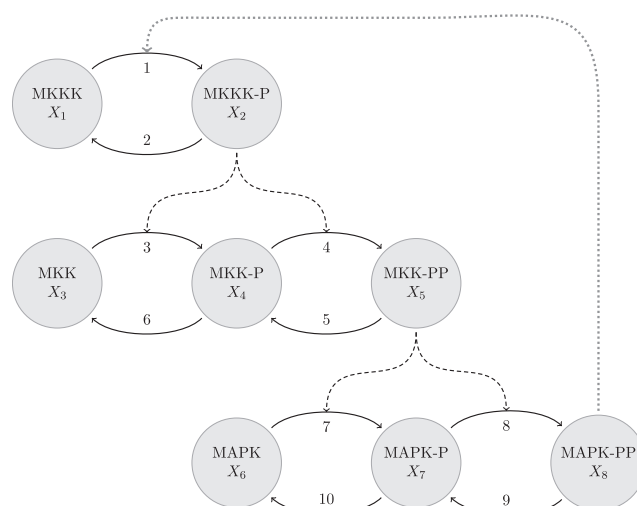
We can understand the relatively small VRF that we observe based on the decomposition of the estimator variances (15) and (16). The difference between the two variances is due to the interstrata variance, and it increases with an increasing interstrata variance. However, by increasing the uniformization rate,  $\bar{a}$ , the virtual reactions become more prevalent. This has the result that most of the reactions in a stratum become virtual reactions and therefore the strata distributions become more similar, diminishing the interstrata variance.

Finally, we note that numerical verification of the VRF is computationally intensive because the VRF is small. We used  $N = 2^{14}$  sample paths from the uniformized system with uniformization rate  $\bar{a} = 6$  and  $J = 6$  strata and parameters  $c_1 = 0.3$ ,  $c_2 = 0.1$  and  $\mathbf{X}_0 = (20, 0)$  at time  $T = 5$  to find the estimator variances for the summary statistic  $Q = \mathbb{E}[X_1(T)]$  using the standard Monte Carlo sample variance estimator and (19). By repeating this procedure 256 times, yielding roughly  $4 \times 10^6$  sample paths in total, we can construct confidence intervals for the variances of both estimators. This yields the 99.7% confidence intervals for the standard Monte Carlo method  $\mathbb{V}[\hat{Q}] \in (4.549, 4.568)/N$  and for the stratified estimator  $\mathbb{V}[\hat{Q}_{\text{prop}}] \in (4.035, 4.052)/N$ . Both of these results agree with the theoretical values of  $\mathbb{V}[\hat{Q}] = 4.559/N$  and  $\mathbb{V}[\hat{Q}_{\text{prop}}] = 4.047/N$ , respectively, and show that numerically, there is a variance reduction visible, albeit a small one.

## 2. MAPK-cascade with feedback

The second example is a mitogen-activated protein kinase (MAPK) cascade model from Ref. 31. It consists of eight species linked by ten reaction channels with Michaelis-Menten kinetics and Hill functions. A schematic representation of the network structure is shown in Fig. 6.

The MAPK-cascade network describes phosphorylation of MAPK by a layered process involving several kinases. MAPK, the terminal kinase, is phosphorylated by MAPK kinases (MKKs), depicted by the interaction between the second and the third layer in Fig. 6. In turn, the MKKs are phosphorylated by the MAPK kinase kinases (MKKKs), which is visible as the interaction between the first and the second layer in Figure 6. The activation of MKKK is thought



**FIG. 6.** MAPK-cascade network, adapted from Ref. 31. Reaction channels are depicted by the numbered horizontal arrows. Vertical dashed lines in black and the gray dotted line denote positive and negative feedback effects, respectively.

to be inhibited by MAPK via membrane-bound Ras and is modeled as a negative feedback loop, as can be seen in Fig. 6.

For this system, there exist three elementary conservation laws and every chemical species is involved in one of them

$$X_1(t) + X_2(t) = X_1(0) + X_2(0), \quad (23a)$$

$$X_3(t) + X_4(t) + X_5(t) = X_3(0) + X_4(0) + X_5(0), \quad (23b)$$

$$X_6(t) + X_7(t) + X_8(t) = X_6(0) + X_7(0) + X_8(0). \quad (23c)$$

It is therefore possible to bind the total propensity uniformly. In this case, we uniformize the system with uniformization rate  $\bar{a} = 15$  and apply stratification using  $J = 6$  strata, i.e., the 6-quantiles. We use the parameter values and initial conditions  $[\mathbf{X}_0 = (100, 0, 300, 0, 0, 300, 0, 0)]$  for the model as given in Ref. 31, Table 2. Because of the non-linearity of the reaction dynamics, it is not possible to find analytical values for the estimator sample variances and VRF. We therefore perform 256 repeated stochastic simulations using  $N = 2^{14}$  sample paths per simulation to get estimates for the sample variances. As summary statistics, we use the number of molecules at the final time  $T = 200$  s for the eight different species in the system and the results are tabulated in Table II.

As we can see in Table II, the VRF depends heavily on the chosen summary statistic. The MAPK-cascade is a hierarchical system in the sense that the conversion of the downstream species  $X_6$ ,  $X_7$ , and  $X_8$  depends on other upstream species, in particular,  $X_5$ . The interstrata variation will therefore likely be larger. The lower quantile strata observe fewer reactions and therefore will not see much activation of the downstream species, in contrast to the higher quantile strata, where more reactions fire and we see more conversion of the downstream species. The conditional summary statistics will therefore differ significantly between the different strata, which we know leads to a larger VRF by the decomposition of the variance (14).

**TABLE II.** Sample variance and estimated VRF for the MAPK cascade for both the standard Monte Carlo estimator and the stratified estimator. The system is run until  $T = 200$  with parameter values and initial conditions as given in Ref. 31, Table 2 and using a uniformization rate  $\tilde{a} = 15$ . Values tabulated are the 99.7% confidence intervals.

Sample variance	Summary statistic $Q = \mathbb{E}[X_i]$			
(A) First four species summary statistics				
	$X_1$	$X_2$	$X_3$	$X_4$
$N \cdot \mathbb{V}[\hat{Q}]$	(22.6, 22.7)	(22.6, 22.7)	(202.4, 203.2)	(225.3, 226.2)
$N \cdot \mathbb{V}[\hat{Q}_{\text{prop}}]$	(22.2, 22.3)	(22.2, 22.3)	(193.7, 194.5)	(225.0, 225.9)
VRF	$1.016 \pm 0.003$	$1.016 \pm 0.003$	$1.045 \pm 0.003$	$1.004 \pm 0.003$
(B) Last four species summary statistics				
	$X_5$	$X_6$	$X_7$	$X_8$
$N \cdot \mathbb{V}[\hat{Q}]$	(259.6, 260.7)	(199.8, 201.2)	(347.4, 348.7)	(777.1, 779.9)
$N \cdot \mathbb{V}[\hat{Q}_{\text{prop}}]$	(246.8, 247.9)	(180.0, 181.3)	(312.3, 313.6)	(670.6, 673.4)
VRF	$1.052 \pm 0.003$	$1.110 \pm 0.006$	$1.112 \pm 0.003$	$1.159 \pm 0.003$

We can therefore conclude that the VRF using stratification and uniformization can vary strongly depending on the model system, the summary statistic of choice, the uniformization rate, and the number of strata used. It is, however, also true that for the same amount of computational complexity, a variance at least as small as for the standard DM is observed at no extra cost.

## V. DISCUSSION AND OUTLOOK

Uniformization techniques are well-known tools in probability theory and have previously appeared in the context of chemical reaction networks in various forms.<sup>5,26,46,54,57</sup> In this work, we revisited the uniformization technique in the context of pathwise simulation of realizations of said chemical reaction networks and focused on the question whether pathwise simulations under uniformization can be carried out efficiently and whether there are any benefits of employing uniformization methods.

To simulate a uniformized system, one will introduce computational overhead by virtue of virtual reactions firing, and *a priori* it might therefore be expected that uniformization-based methods are outperformed by more standard approaches like Gillespie's DM. However, the performance benefits of approaches building on uniformization have been observed in the literature, albeit mainly in the context of time-inhomogeneous reaction networks.<sup>54</sup> We therefore, affirmatively, answer the question whether in the case of time-homogeneous systems, which is still the predominant modeling choice of practitioners, it is possible to find a stochastic simulation approach based on uniformization which is as efficient as the aforementioned golden standard, Gillespie's DM. By addressing the issues of efficiency and adaptiveness of the uniformized systems, we arrived at Algorithm 4 which, complexity-wise, is on par with Gillespie's DM. We reiterate that many improvements that have been made to the basic implementation of Gillespie's DM are equally applicable to the uniformized approaches discussed in this work because the approach to choose the reaction channel that is to fire next

is the same as for Gillespie's DM. The work in this paper therefore opens up the exploration of new approaches built on the uniformized systems by removing a previously perceived complexity problem.

Extensions of the standard Gillespie's DM for time-homogeneous models to include delays have been proposed previously.<sup>7,14,48,49</sup> It is also possible to apply uniformization to systems including delay reactions. In order to simulate such uniformized systems, a first approach would be to explicitly generate the reaction times of the reactions as done for the time-inhomogeneous case; see Appendix A. Thereafter, we can proceed as in the delay stochastic simulation algorithm (DSSA) in Ref. 7. We note that this approach also explicitly generates the reaction times of the virtual reactions and therefore has an extra overhead cost relative to the original system. As an additional extension and application of uniformization, we note that one can apply tractable likelihood-based Bayesian inference techniques by using the uniformized system, e.g., Ref. 40. We leave exploration of complexity, delays, inference, and uniformization for future work.

As mentioned in Sec. IV, one can use stratification to improve estimation of summary statistics for a wide variety of functionals  $f$  of the sample path solutions. In the examples of Sec. IV, we focused on the mean species numbers at a fixed time  $T$  but reiterate that the stratification method is not restricted to such use cases only. Following the arguments from Sec. IV, it is clear that for a general class of summary statistics, we can expect a performance at least as good as implementations with Gillespie's DM. At the same time, however, we note that when applying stratification to estimate the one-dimensional marginal distributions of the species in the MAPK-cascade, Sec. IV C 2, a significantly smaller variance reduction, as defined in terms of the mean integrated squared error, will be observed than when looking at the mean number of molecules as the summary statistic of interest. Another issue appears when one is interested in using a summary statistic that inherently



depends on lower order moments of the random variable  $\mathbf{Y}$ , such as the population variance. To use (unbiased) stratified sampling for such summary statistics is a nontrivial task, and a variance reduction is no longer guaranteed; see, for example, Ref. 55. This raises the question as to what kind of problems, both in terms of the chemical reaction network and in terms of the employed summary statistic, result in a substantial variance reduction when combining uniformization and stratification. We aim to investigate this in further work.

Finally, in this work, we were primarily concerned with summary statistics that depend on the value of the paths at some time  $T$ , a commonly encountered choice of summary statistic. However, one might also be interested in path-dependent summary statistics, such as the transient evolution or mean value of a species over a series of time points  $T_1, \dots, T_n$ . A straightforward approach to generate such summary statistics with the standard SSAs would be to run the paths over  $[0, T_1]$  and record the state, then over  $[T_1, T_2]$  and record the state, etc. Such an approach applies equally well to the (improved) UDM, where we can either generate the total number of reactions  $M$  on each interval individually or use the fact that given that  $M$  reactions occur in  $[0, T]$ , the number of reactions in the intervals defined by  $T_1, \dots, T_n$  follows a multinomial distribution. The efficient combination of transient analysis and uniformization based methods is a subject of future work.

## ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers, whose detailed comments improved the clarity of this paper. Casper H. L. Beentjes acknowledges the Clarendon Fund and New College, Oxford, for funding. Ruth E. Baker is a Royal Society Wolfson Research Merit Award holder and a Leverhulme Research Fellow and also acknowledges the BBSRC for funding via Grant No. BB/R000816/1.

## APPENDIX A: TIME-INHOMOGENEOUS MODELS

The results presented in this paper are valid for time-homogeneous Markov chains, i.e., models where the propensity functions do not have an inherent time-dependency and only depend on the state of system  $\mathbf{X}$ . This assumption is generally sufficient to model intrinsic noise in a system, but whenever extrinsic noise effects are to be taken into account, one has to relax this assumption. Extrinsic noise can be used to model the influence of external environments on the model behavior and has been shown to influence the model dynamics.<sup>11,28,45,54</sup>

### 1. Stochastic model and simulation

The most common modeling change to account for extrinsic noise is to assume that the reaction rate constants are allowed to be time varying functions, i.e.,  $c_k(t)$ . This dependency can be prescribed as a simple function, e.g.,  $c_k(t) = 10(1 + \sin(t))$ , or be more elaborate, e.g., rates of the form  $c_k(t) = 2 \exp(Z(t))$  with  $Z(t)$  some other stochastic process, which can, for example, be described by a stochastic differential equation (SDE). The result is a time-inhomogeneous Markov chain model, and currently, the best-known SSA for such systems is the Extrande method<sup>54</sup> which

uses the uniformization method to avoid having to (numerically) integrate the propensity functions,  $c_k(t)$ . Methods built on the (numerical) integration of the propensity functions such as the modified next-reaction method<sup>1</sup> can also be constructed, but these methods do not compare favorably with methods based on uniformization in general.<sup>54</sup> The difference between the Extrande method and the UDM described in Sec. III is that the former method, like the DM, generates all the reaction times in a serial fashion. In contrast to the case of time-homogeneous models, where it is not necessary to generate the reaction times explicitly for the UDM, in the case of a time-inhomogeneous model, one has to know the reaction times to account for the explicit time dependence of the reaction propensities. We therefore have to adapt the UDM (Algorithm 2) to account for this. As mentioned in Sec. III, we can sample the reaction times,  $t_i$ , conditional on the knowledge that the number of reactions firing in  $[0, T)$  is  $M$ . We note that the reaction times  $t_1, \dots, t_M$  are then distributed as the order statistics of  $M$  uniform random variables in  $[0, T)$ . To generate these, one can simply generate  $M$  uniform random variables in  $[0, T)$  and sort them in ascending order. Alternative methods to generate sorted uniform random numbers exist with better complexity properties, e.g., Ref. 17, Chap. 3. This extra step can be done prior to the reaction dynamics loop as depicted in Algorithm 6.

We note that if  $a_0 > \bar{a}$  is observed in the course of a sample path, the procedure to adapt the uniformization rate is the same as described in Sec. III B apart from the fact that the time  $T^*$  does not have to be sampled in this case; we can simply set  $T^* = t_{M^*}$  with  $M^*$  the number of reactions that have fired up until the point at which  $a_0 > \bar{a}$ .

The extra computational gain that was achieved by firing virtual reactions consecutively does not apply in this framework. For time-homogeneous models, we used the fact that the distribution of the number of consecutive virtual reactions firing was geometric, but this is no longer valid for time-inhomogeneous models. Due to the explicit dependence on time of the propensity functions, we get an expression for the distribution of the number of consecutive virtual reactions of the form

$$\mathbb{P}(r \text{ consecutive virtual reactions before next real reaction fires}) \\ = \frac{a_0(t_{r+1})}{\bar{a}} \prod_{i=1}^r \left(1 - \frac{a_0(t_i)}{\bar{a}}\right). \quad (\text{A1})$$

This is no longer a single parameter distribution, and although it is analytically tractable, sampling from it will generally involve evaluating the total propensity  $a_0$  at the time points of the virtual reactions. Therefore, such a strategy might potentially negate any speed-ups from generating consecutive virtual reactions in a single step.

As a result, we do not want to have to choose a uniformization rate  $\bar{a}$  that is much larger than  $a_0$  over  $[0, T)$  as every virtual reaction will have to be simulated individually for the time-inhomogeneous case in Algorithm 6. If the total propensity  $a_0$  varies strongly over the course of  $[0, T)$ , it might be beneficial to regularly adapt the rate  $\bar{a}$ , which in the extreme case of adapting the rate after every reaction yields the Extrande method.

### 2. Stratification

The theory for stratification still holds in the case of time-inhomogeneous models. If one can therefore feasibly generate



**Algorithm 6.** Uniformized direct method for time-inhomogeneous models.

---

**Input:** Initial data  $\mathbf{X}_0$   
**Input:** Stoichiometric matrix  $\zeta$   
**Input:** Propensity functions  $a_k(\mathbf{x}, t)$   
**Input:** Uniformization rate  $\bar{a}$   
**Input:** Final time  $T$

- 1: Generate  $M \sim \mathcal{P}(\bar{a}T)$   $\triangleright$  Total number of reactions that fire in  $[0, T)$ .
- 2: Generate  $t_1, \dots, t_M \sim \mathcal{U}(0, T)$  s.t.  $t_1 < t_2 < \dots < t_M$   $\triangleright$  Reaction times.
- 3: **for**  $m = 1, \dots, M$  **do**
- 4:   Generate  $u_1 \sim \mathcal{U}(0, 1)$
- 5:    $\mathbf{A}_k \leftarrow a_k(\mathbf{X}, t_m)$   $\triangleright$  Calculate real reaction propensities.
- 6:    $\mathbf{A}_{K+1} \leftarrow \bar{a} - \sum_{k=1}^K \mathbf{A}_k$   $\triangleright$  Calculate virtual reaction propensity.
- 7:   Find  $p$  such that  $\sum_{k=1}^{p-1} \mathbf{A}_k < \bar{a}u_1 \leq \sum_{k=1}^p \mathbf{A}_k$   $\triangleright$  Choose next reaction to fire.
- 8:   **if**  $p \in \{1, \dots, K\}$  **then**  $\triangleright$  Only need to fire real reactions.
- 9:      $\mathbf{X} \leftarrow \mathbf{X} + \zeta_p$   $\triangleright$  Update state vector.
- 10:   **end if**
- 11: **end for**

---

sample paths using Algorithm 6 with a fixed uniformization rate,  $\bar{a}$ , it would be wise to apply stratification with respect to the number of reactions firing,  $M$ , to achieve a variance reduction at no extra cost. If, however, the propensity  $a_0$  changes noticeably over  $[0, T)$ , one might want to use an adaptive uniformization rate,  $\bar{a}$ , or the Extrande method, in which case the stratification method is no longer valid because the reaction times in the interval  $[0, T)$  are no longer drawn from a single parameter joint distribution that is the same for each sample path.

## APPENDIX B: UNBIASED ADAPTING OF THE UNIFORMIZATION RATE

Here, we provide an alternative view on the sampling of the time  $T^*$  at which one adapts the uniformization rate, as discussed in Sec. III B. At the same time, we show that the conditional sampling of  $T^*$  does in fact yield unbiased samples from the correct distribution.

### 1. Derivation of the conditional distribution of $T^*$

For completeness, we first show that the  $M^*$ th order statistic of  $M$  uniform random variables follows a Beta distribution. We start with the assumption that we know that  $M$  reactions fire in the time interval  $[0, T)$  as sampled from the Poisson distribution with parameter  $\bar{a}T$ . Their reaction times are then uniformly distributed on  $[0, T)$ . To find the distribution of the time  $T^*$  of the  $M^*$ th reaction, we note that in order to have  $T^* \in [t, t + \Delta t)$ , we need exactly  $M^* - 1$  reactions in  $[0, t)$ , one reaction in  $[t, t + \Delta t)$  and  $(M - M^*)$  reactions in  $[t + \Delta t, T)$ . This yields the following expression:

$$\mathbb{P}(T^* \in [t, t + \Delta t)) = \frac{M!}{(M - M^*)!(M^* - 1)!} \left(\frac{t}{T}\right)^{M^* - 1} \times \frac{\Delta t}{T} \left(1 - \frac{t + \Delta t}{T}\right)^{M - M^*}, \quad (\text{B1})$$

where the prefactor stems from the indistinguishable nature of the reactions in  $[0, t)$  and  $[t + \Delta t, T)$ . By letting  $\Delta t \rightarrow 0$ , we get the

probability density function  $f(t)$  for  $T^*$

$$f(t) = \frac{\Gamma(M - 1)}{\Gamma(M - M^* - 1)\Gamma(M^*)} \left(\frac{t}{T}\right)^{M^* - 1} \frac{1}{T} \left(1 - \frac{t}{T}\right)^{M - M^*}. \quad (\text{B2})$$

We now let  $u = t/T$ , i.e., we consider the distribution of  $T^*/T$  and recognize the prefactor as  $1/B(M - M^* + 1, M^*)$ , where  $B$  is the beta-function. This leads to the following probability density function  $g(u)$  for  $T^*/T$ :

$$g(u) = \frac{1}{B(M - M^* + 1, M^*)} (u)^{M^* - 1} (1 - u)^{M - M^*}, \quad (\text{B3})$$

which is readily seen to be equivalent to the Beta distribution on  $[0, 1)$  with parameters  $M^*$  and  $(M - M^*) + 1$ .

### 2. Unconditional distribution of $T^*$

In our approach to adapt the uniformization rate, we sample a time  $T^*$  conditional on the fact that  $M$  reactions within the interval  $[0, T)$  were sampled at the uniformization rate  $\bar{a}$ . For the unconditional distribution of  $T^*$ , the time of the  $M^*$ th reaction and the point at which we adapt the uniformization rate, as in Algorithm 4, we start with the following observation for  $t \in [0, T)$ :

$$\begin{aligned} \mathbb{P}(T^* \leq t) &= \sum_{M=0}^{\infty} \mathbb{P}(T^* \leq t | M \text{ reactions in } [0, T)) \\ &\quad \times \mathbb{P}(M \text{ reactions in } [0, T)) \\ &= \sum_{M=M^*}^{\infty} \mathbb{P}\left(\frac{T^*}{T} \leq \frac{t}{T} \middle| M \text{ reactions in } [0, T)\right) \\ &\quad \times \mathbb{P}(M \text{ reactions in } [0, T)), \end{aligned}$$

where we now recognize the first term in the sum to be described by a Beta distribution (with parameters  $M^*$  and  $(M - M^*) + 1$ ) and the second term by a Poisson distribution (with parameter  $\bar{a}T$ ). Substitution of the relevant expressions for these distributions then yields

$$\begin{aligned}
 \mathbb{P}(T^* \leq t) &= \sum_{M=M^*}^{\infty} \left( \frac{1}{B(M^*, (M-M^*)+1)} \int_0^{t/T} s^{M^*-1} \right. \\
 &\quad \times (1-s)^{M-M^*} ds \Big) \frac{(\bar{a}T)^M}{M!} e^{-\bar{a}T} \\
 &= \frac{1}{\Gamma(M^*)} \int_0^{t/T} s^{M^*-1} (\bar{a}T)^{M^*} e^{-\bar{a}Ts} \\
 &\quad \times \left( \sum_{M=M^*}^{\infty} \frac{(\bar{a}T(1-s))^{M-M^*}}{(M-M^*)!} e^{-\bar{a}T(1-s)} \right) ds \\
 &= \frac{1}{\Gamma(M^*)} \int_0^{\bar{a}t} u^{M^*-1} e^{-u} du,
 \end{aligned}$$

which is the distribution function for the  $\text{Gamma}(M^*, \bar{a})$  distribution. This is exactly the distribution of  $M^*$  i.i.d. exponential random variables with parameter  $\bar{a}$ . As a result,  $T^*$  is distributed as the  $M^*$ th reaction time in a system with constant propensity  $\bar{a}$ . This proves that the conditional sampling of  $T^*$  as described in Sec. III B yields in fact samples from the correct distribution and therefore makes the construction of sample paths with an adaptive uniformization rate unbiased.

## REFERENCES

- D. F. Anderson, "A modified next reaction method for simulating chemical systems with time dependent propensities and delays," *J. Chem. Phys.* **127**, 214107 (2007).
- D. F. Anderson and D. J. Higham, "Multilevel Monte Carlo for continuous time Markov chains, with applications in biochemical kinetics," *Multiscale Model. Simul.* **10**, 146–179 (2012).
- D. F. Anderson, D. J. Higham, and Y. Sun, "Computational complexity analysis for Monte Carlo approximations of classically scaled population processes," *Multiscale Model. Simul.* **16**, 1206–1226 (2018).
- D. F. Anderson and T. G. Kurtz, *Design and Analysis of Biomolecular Circuits* (Springer New York, New York, 2011), pp. 3–42.
- D. F. Anderson and C. Yuan, "Low variance couplings for stochastic models of intracellular processes with time-dependent rate functions," *Bull. Math. Biol.* (published online).
- A. Auger, P. Chatelain, and P. Koumoutsakos, "R-leaping: Accelerating the stochastic simulation algorithm by reaction leaps," *J. Chem. Phys.* **125**, 084103 (2006).
- M. Barrio, K. Burrage, A. Leier, and T. Tian, "Oscillatory regulation of Hes1: Discrete stochastic delay modelling and simulation," *PLoS Comput. Biol.* **2**, e117 (2006).
- C. H. L. Beentjes and R. E. Baker, "Quasi-Monte Carlo methods applied to tau-leaping in stochastic biological systems," *Bull. Math. Biol.* (published online, 2019).
- C. Ben Hammouda, A. Moraes, and R. Tempone, "Multilevel hybrid split-step implicit tau-leap," *Numer. Algorithms* **74**, 527–560 (2017).
- W. J. Blake, M. Kærn, C. R. Cantor, and J. J. Collins, "Noise in eukaryotic gene expression," *Nature* **422**, 633–637 (2003).
- C. G. Bowsher and P. S. Swain, "Identifying sources of variation and the flow of information in biochemical networks," *Proc. Natl. Acad. Sci. U. S. A.* **109**, E1320–E1328 (2012).
- L. Bronstein and H. Koepl, "A variational approach to moment-closure approximations for the kinetics of biomolecular reaction networks," *J. Chem. Phys.* **148**, 014105 (2018).
- L. Cai, N. Friedman, and X. S. Xie, "Stochastic protein expression in individual cells at the single molecule level," *Nature* **440**, 358–362 (2006).
- X. Cai, "Exact stochastic simulation of coupled chemical reactions with delays," *J. Chem. Phys.* **126**, 124108 (2007).
- Y. Cao, H. Li, and L. Petzold, "Efficient formulation of the stochastic simulation algorithm for chemically reacting systems," *J. Chem. Phys.* **121**, 4059–4067 (2004).
- W. G. Cochran, *Sampling Techniques*, 3rd ed. (John Wiley & Sons, New York, 1977).
- L. Devroye, *Non-Uniform Random Variate Generation* (Springer-Verlag, New York, NY, 1986).
- M. B. Elowitz, A. J. Levine, E. D. Siggia, and P. S. Swain, "Stochastic gene expression in a single cell," *Science* **297**, 1183–1186 (2002).
- R. Erban, S. J. Chapman, and P. K. Maini, "A practical guide to stochastic simulations of reaction-diffusion processes," e-print [arXiv:0704.1908](https://arxiv.org/abs/0704.1908) (2007).
- M. A. Gibson and J. Bruck, "Efficient exact stochastic simulation of chemical systems with many species and many channels," *J. Phys. Chem. A* **104**, 1876–1889 (2000).
- D. T. Gillespie, "Approximate accelerated stochastic simulation of chemically reacting systems," *J. Chem. Phys.* **115**, 1716–1733 (2001).
- D. T. Gillespie, "Exact stochastic simulation of coupled chemical reactions," *J. Phys. Chem.* **81**, 2340–2361 (1977).
- D. T. Gillespie, A. Hellander, and L. R. Petzold, "Perspective: Stochastic algorithms for chemical kinetics," *J. Chem. Phys.* **138**, 170901 (2013).
- A. Gupta, J. Mikelson, and M. Khammash, "A finite state projection algorithm for the stationary solution of the chemical master equation," *J. Chem. Phys.* **147**, 154101 (2017).
- M. Hegland, C. Burden, L. Santoso, S. MacNamara, and H. Booth, "A solver for the stochastic master equation applied to gene regulatory networks," *J. Comput. Appl. Math.* **205**, 708–724 (2007).
- A. Hellander, "Efficient computation of transient solutions of the chemical master equation based on uniformization and quasi-Monte Carlo," *J. Chem. Phys.* **128**, 154109 (2008).
- D. J. Higham, "Modeling and simulating chemical reactions," *SIAM Rev.* **50**, 347–368 (2008).
- A. Hilfinger and J. Paulsson, "Separating intrinsic from extrinsic fluctuations in dynamic biological systems," *Proc. Natl. Acad. Sci. U. S. A.* **108**, 12167–12172 (2011).
- Z. Hou and H. Xin, "Internal noise stochastic resonance in a circadian clock system," *J. Chem. Phys.* **119**, 11508 (2003).
- T. Jahnke and W. Huisinga, "Solving the chemical master equation for monomolecular reaction systems analytically," *J. Math. Biol.* **54**, 1–26 (2007).
- B. N. Kholodenko, "Negative feedback and ultrasensitivity can bring about oscillations in the mitogen-activated protein kinase cascades," *Eur. J. Biochem.* **267**, 1583–1588 (2000).
- C. Lemieux, *Monte Carlo and Quasi-Monte Carlo Sampling* (Springer-Verlag, 2009).
- C. Lester, R. E. Baker, M. B. Giles, and C. A. Yates, "Extending the multi-level method for the simulation of stochastic biological systems," *Bull. Math. Biol.* **78**, 1640–1677 (2016).
- S. MacNamara, A. M. Bersani, K. Burrage, and R. B. Sidje, "Stochastic chemical kinetics and the total quasi-steady-state assumption: Application to the stochastic simulation algorithm and chemical master equation," *J. Chem. Phys.* **129**, 095105 (2008).
- S. MacNamara, K. Burrage, and R. B. Sidje, "Multiscale modeling of chemical kinetics via the master equation," *Multiscale Model. Simul.* **6**, 1146–1168 (2008).
- H. H. McAdams and A. Arkin, "It's a noisy business! Genetic regulation at the nanomolar scale," *Trends Genet.* **15**, 65–69 (1999).
- J. M. McCollum, G. D. Peterson, C. D. Cox, M. L. Simpson, and N. F. Samatova, "The sorting direct method for stochastic simulation of biochemical systems with varying reaction execution behavior," *Comput. Biol. Chem.* **30**, 39–49 (2006).
- B. Munsky and M. Khammash, "The finite state projection algorithm for the solution of the chemical master equation," *J. Chem. Phys.* **124**, 044104 (2006).
- J. Paulsson, O. G. Berg, and M. Ehrenberg, "Stochastic focusing: Fluctuation-enhanced sensitivity of intracellular regulation," *Proc. Natl. Acad. Sci. U. S. A.* **97**, 7148–7153 (2000).

- <sup>40</sup>V. Rao and Y. W. Teh, “Fast MCMC sampling for Markov jump processes and extensions,” *J. Mach. Learn. Res.* **14**, 3295–3320 (2013).
- <sup>41</sup>M. Reis, J. A. Kromer, and E. Klipp, “General solution of the chemical master equation and modality of marginal distributions for hierarchic first-order reaction networks,” *J. Math. Biol.* **77**, 377–419 (2018).
- <sup>42</sup>W. Sandmann, “Discrete-time stochastic modeling and simulation of biochemical networks,” *Comput. Biol. Chem.* **32**, 292–297 (2008).
- <sup>43</sup>D. Schnoerr, G. Sanguinetti, and R. Grima, “Approximation and inference methods for stochastic biochemical kinetics—A tutorial review,” *J. Phys. A: Math. Theor.* **50**, 093001 (2017).
- <sup>44</sup>D. Schnoerr, G. Sanguinetti, and R. Grima, “Comparison of different moment-closure approximations for stochastic chemical kinetics,” *J. Chem. Phys.* **143**, 185101 (2015).
- <sup>45</sup>P. S. Swain, M. B. Elowitz, and E. D. Siggia, “Intrinsic and extrinsic contributions to stochasticity in gene expression,” *Proc. Natl. Acad. Sci. U. S. A.* **99**, 12795–12800 (2002).
- <sup>46</sup>V. H. Thanh, L. Marchetti, F. Reali, and C. Priami, “Incorporating extrinsic noise into the stochastic simulation of biochemical reactions: A comparison of approaches,” *J. Chem. Phys.* **148**, 064111 (2018).
- <sup>47</sup>V. H. Thanh, C. Priami, and R. Zunino, “Accelerating rejection-based simulation of biochemical reactions with bounded acceptance probability,” *J. Chem. Phys.* **144**, 224108 (2016).
- <sup>48</sup>V. H. Thanh, C. Priami, and R. Zunino, “Efficient rejection-based simulation of biochemical reactions with stochastic noise and delays,” *J. Chem. Phys.* **141**, 134116 (2014).
- <sup>49</sup>V. H. Thanh, R. Zunino, and C. Priami, “Efficient stochastic simulation of biochemical reactions with noise and delays,” *J. Chem. Phys.* **146**, 084107 (2017).
- <sup>50</sup>V. H. Thanh, R. Zunino, and C. Priami, “On the rejection-based algorithm for simulation and analysis of large-scale reaction networks,” *J. Chem. Phys.* **142**, 244106 (2015).
- <sup>51</sup>N. M. van Dijk, S. P. J. van Brummelen, and R. J. Boucherie, “Uniformization: Basics, extensions and applications,” *Perform. Eval.* **118**, 8–32 (2018).
- <sup>52</sup>N. Van Kampen, *Stochastic Processes in Physics and Chemistry*, 3rd ed. (Elsevier, Amsterdam, The Netherlands, 2007).
- <sup>53</sup>A. P. Van Moorsel and W. H. Sanders, “Adaptive uniformization,” *Commun. Stat.: Stochastic Models* **10**, 619–647 (1994).
- <sup>54</sup>M. Voliotis, P. Thomas, R. Grima, and C. G. Bowsher, “Stochastic simulation of biomolecular networks in dynamic environments,” *PLoS Comput. Biol.* **12**, e1004923 (2016).
- <sup>55</sup>K. Wakimoto, “On unbiased estimation of the population variance based on the stratified random sample,” *Ann. Inst. Stat. Math.* **22**, 15–26 (1970).
- <sup>56</sup>C. A. Yates and G. Klingbeil, “Recycling random numbers in the stochastic simulation algorithm,” *J. Chem. Phys.* **138**, 094103 (2013).
- <sup>57</sup>J. Zhang, L. T. Watson, and Y. Cao, “A modified uniformization method for the solution of the chemical master equation,” *Comput. Math. Appl.* **59**, 573–584 (2010).
- <sup>58</sup>In the case of a system with a bounded state space, it is (theoretically) possible to find a uniformization rate by taking the maximum of the total propensity over all allowed states. Note, however, that the size of the state space, albeit finite, could be prohibitively large for such an approach to be practically feasible.
- <sup>59</sup>This observation is mainly true for time-homogeneous systems, as considered in this section. For time-inhomogeneous systems, methods based on the DM for the original system can often compare unfavorably with uniformization-based methods, as discussed in [Appendix A](#) and [Ref. 54](#).
- <sup>60</sup>Sampling from the truncated Poisson distribution can be easily carried out via the inverse transform sampling method and the truncated cumulative distribution function (CDF),  $F_{\text{truncated}}(x) = (F(x) - F(M_j)) / (F(M_{j+1}) - F(M_j))$ , where  $F(x)$  is the standard Poisson CDF.
- <sup>61</sup>Due to the discrete nature of the Poisson distribution, it is not possible to define strata that exactly have equal weights.