Guided proposals for efficient weighted stochastic simulation

Colin S. Gillespie*

Andrew Golightly

School of Mathematics, Statistics and Physics, Newcastle University,
Newcastle upon Tyne, NE1 7RU, UK

Abstract

Rare event probabilities play an important role in the understanding of the behaviour of biochemical systems. Due to the intractability of the most natural Markov jump process representation of a system of interest, rare event probabilities are typically estimated using importance sampling. While the resulting algorithm is reasonably well developed, the problem of choosing a suitable importance density is far from straightforward. We therefore leverage recent developments on simulation of conditioned jump processes to propose an importance density that is simple to implement and requires no tuning. Our results demonstrate superior performance over some existing approaches.

Keywords: Rare event simulation; stochastic simulation algorithm; importance sampling.

1 Introduction

Simulating discrete biochemical networks via stochastic simulations in recent years has become routine. The default algorithm for many practitioners was developed over four decades ago by Gillespie (1977). This algorithm simulates a single stochastic realisation from the Markov jump process (MJP) representation of a system of interest. Typically, multiple realisations are generated in order to estimate state probabilities. Given the intractability of the MJP, the importance of the stochastic simulation algorithm (SSA) cannot be understated.

Multiple improvements have been proposed to the *default*. For example, Gibson and Bruck (2000) use efficient data structures, whereas other authors McCollum et al. (2006); Slepoy et al. (2008) have leveraged clever sorting strategies. It is also possible to speed up a single stochastic simulation by grouping sets of reactions together and simulating on a multi-core CPU (Gillespie, 2012).

Despite these improvements in computational efficiency, the problem of estimating probabilities of rare events remains challenging. While an ensemble simulation approach can in principle be routinely applied, the number of simulations required to attain a certain level of accuracy can be computationally prohibitive.

The weighted stochastic simulation algorithm (wSSA) was developed by Kuwahara and Mura (2008) to alleviate the computational burden associated with calculating rare event probabilities, that is, estimating $p(x_0, x'; t')$, the probability of reaching x' before time t' , given x_0 at time $t = 0$.

*email: colin.gillespie@ncl.ac.uk

The general idea is to carefully *weight* or bias key reactions, which will reduce the total number of simulations required (relative to a vanilla Monte Carlo approach based on the SSA) to give an estimator that attains a given level of accuracy. As an aside, while these simulations can be used to estimate the probability of interest, they can not be regarded as valid realisations from the system of interest. The wSSA can be seen as an importance sampler, and this connection was made explicitly by Gillespie et al. (2009), who further extended the wSSA by proposing to select suitable biasing parameters by *minimising* the variance of the resulting importance sampling estimator. The need for state-dependent biasing parameters was discussed by Roh et al. (2010) who proposed to partition each reaction depending on whether it should be encouraged or not, and specified state-dependent propensity functions for each reaction group. An automated approach to biasing parameter selection was presented by Daigle et al. (2011) with a subsequent extension to state-dependent bias in Roh et al. (2011). In the latter two contributions, a cross entropy method is used to continually refine biasing parameters. Consequently, previous work in this area either requires the specification of additional tuning parameters or additional simulations.

Our novel approach is the specification of an importance density that is simple to implement and requires no additional tuning. We leverage recent developments on simulation of conditioned jump processes to derive a propensity function that takes into account the rare event of interest. We then use this propensity inside the SSA to generate trajectories that are guided towards the rare event x' , thereby increasing the likelihood of hitting x' before time t' . We find that the proposed approach outperforms competing methods in a variety of settings. Limitations are also discussed.

The remainder of this article is organised as follows. In Section 2, we discuss some recent developments around the wSSA algorithms. Section 3 will consider a simple two-state model, where it is possible to derive analytical expressions for the optimal weights. Following these insights, we will derive our new method in 4 and compare its performance with other methods in Section 5 using three examples of increasing complexity.

2 The wSSA Algorithm

Consider a reaction network involving u species $\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_u$ and v reaction channels $\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_v$ such that

where the p_{ij} and q_{ij} are non-negative integers known as stoichiometric coefficients. Let $X_{j,t}$ denote the (discrete) number of species $\mathcal{X}_j$ at time t , and let X_t be the u -vector $X_t = (X_{1,t}, X_{2,t}, \dots, X_{u,t})^T$. The time evolution of X_t is most naturally described by a Markov jump process (MJP), so that for an infinitesimal time increment dt and a propensity function (or instantaneous hazard) $h_i(x_t)$, the probability of a type i reaction occurring in the time interval $(t, t + dt]$ is $h_i(x_t)dt$. Under the standard assumption of mass action kinetics, h_i is proportional to a product of binomial coefficients. Specifically

$$h_i(x_t) = k_i \prod_{j=1}^u \binom{x_{j,t}}{p_{ij}}$$

and we suppress explicit dependence of h_i on the rate constant k_i . Values for $k = (k_1, k_2, \dots, k_v)^T$, the initial system state $X_0 = x_0$ and the $u \times v$ stoichiometry matrix S whose (i, j) th element is given by $q_{ji} - p_{ji}$, complete specification of the Markov process. Despite the intractability of the probability mass function governing the state of the system at any time t (as satisfies the chemical master equation), generating exact realisations of the MJP is straightforward via a technique known in this context as the *stochastic simulation algorithm* (SSA) (Gillespie, 1977). In brief, if

the current time and state of the system are t and X_t respectively, then the time to the next event will be exponential with rate parameter

$$h_0(x_t) = \sum_{i=1}^v h_i(x_t),$$

and the event will be a reaction of type $\mathcal{R}_i$ with probability $h_i(x_t)/h_0(x_t)$ independently of the inter-event time.

We suppose that interest lies in $p(x_0, x'; t')$, that is, the probability that the system starting at the initial value x_0 at time 0 will first reach the state x' at some time t in the interval $(0, t']$. This probability can be written as the expectation

$$\mathbb{E} \{ I(x' \in X_{(0, \mathcal{T}]}) \} \quad (1)$$

where the indicator function $I(x' \in X_{(0, \mathcal{T}]})$ takes the value 1 if $t \leq \mathcal{T}$ and $X_{\mathcal{T}} = x'$ at some stopping time $\mathcal{T} \in (0, t']$. It takes the value 0 otherwise. Note that the expectation is with respect to the probability law of the MJP over the interval $(0, t']$, denoted $p(x_{(0, t']} | x_0)$. Then, a simple Monte Carlo procedure for (unbiasedly) estimating the expectation in (1) is to generate (say) N independent draws $x^{(1)}, \dots, x^{(N)}$ of $x_{(0, \mathcal{T}]}$ via the SSA and set

$$\hat{p}(x_0, x'; t') = \frac{1}{N} \sum_{j=1}^N I(x' \in x^{(j)}).$$

However, it is well known that when dealing with rare events, the corresponding estimator will have high variance (for a given N), since relatively few draws of the MJP will attain the condition for setting the indicator equal to 1 (see Kuwahara and Mura, 2008).

The weighted SSA (wSSA) constructs an *importance sampling* estimate of $p(x_0, x'; t')$. A key ingredient of this approach is a suitable *importance density* $q(x_{(0, t']} | x_0)$. Given N independent draws $x^{(1)}, \dots, x^{(N)}$ of $x_{(0, \mathcal{T}]}$ from $q(x_{(0, \mathcal{T}]} | x_0)$, the importance sampling estimate is

$$\hat{p}(x_0, x'; t') = \frac{1}{N} \sum_{j=1}^N w_{\mathcal{T}}^{(j)} I(x' \in x^{(j)}) \quad (2)$$

where the *weight* function is

$$w_{\mathcal{T}}^{(j)} = \frac{p(x_{(0, \mathcal{T}]}^{(j)} | x_0)}{q(x_{(0, \mathcal{T}]}^{(j)} | x_0)}. \quad (3)$$

The importance density should be chosen so that samples drawn from it are “more likely” to meet the attainment condition. Hence, the choice of importance density is crucial for devising an estimator with low variance. In what follows, we take $q(x_{(0, t']}^{(j)} | x_0)$ to be the probability law associated with an MJP whose propensity function is $\tilde{h}_i(x_t)$ for reaction $i = 1, \dots, v$. In this case, the weight function can be explicitly calculated as follows. We let r_j denote the number of reaction events of type $\mathcal{R}_j$, $j = 1, \dots, v$, and define $n_r = \sum_{j=1}^v r_j$ as the total number of reaction events over the interval $(0, \mathcal{T}]$. Reaction times (assumed to be in increasing order) and types are denoted by (t_i, ν_i) , $i = 1, \dots, n_r$, $\nu_i \in \{1, \dots, v\}$ and we take $t_0 = 0$ and $t_{n_r} = \mathcal{T}$. The *complete-data likelihood* over $(0, \mathcal{T}]$ is

$$p(x_{(0, \mathcal{T}]} | x_0) = \left\{ \prod_{i=1}^{n_r} h_{\nu_i}(x_{t_{i-1}}) \right\} \times \exp \left\{ - \sum_{i=1}^{n_r} h_0(x_{t_{i-1}}) \Delta t_i \right\}$$

where $\Delta t_i = t_i - t_{i-1}$ is the dwell time preceding the i th reaction firing (see Wilkinson, 2012). An expression for $q(x_{(0,\tau]}|x_0)$ is obtained similarly. Hence the weight function is

$$w_{\mathcal{T}} = \left\{ \prod_{i=1}^{n_r} \frac{h_{\nu_i}(x_{t_{i-1}})}{\tilde{h}_{\nu_i}(x_{t_{i-1}})} \right\} \times \exp \left\{ - \sum_{i=1}^{n_r} \left[h_0(x_{t_{i-1}}) - \tilde{h}_0(x_{t_{i-1}}) \right] \Delta t_i \right\}. \quad (4)$$

Computation of an estimate $\hat{p}(x_0, x'; t')$ is then straightforward by performing N independent runs of the SSA with $h_i(x_t)$ replaced by $\tilde{h}_i(x_t)$. The number M_N of runs that reach state x' prior to time t' is obtained, and each such run is assigned a weight using (4). The weights are then used as in (2) to estimate $p(x_0, x'; t')$.

It remains that we specify a suitable form of $\tilde{h}_i(x_t)$ to be used in the importance sampler. Previous approaches (Kuwahara and Mura, 2008; Gillespie et al., 2009) have used $h_0(x_t)$ to update the inter-event times and $\tilde{h}_i(x_t) = \alpha h_i(x_t)$ to update the reaction events, for some tuning parameter α . Roh et al. (2011) use $\alpha h_i(x_t)$ to update both the inter-event times and reaction events. In the following section we motivate the need for the scaling parameter α to be time-dependent before presenting our novel approach in Section 4.

3 Discrete two state model

Consider a simple, *discrete-time*, two state model. Suppose X can only take the values 0 or 1. The probability of remaining in state 0 is p_0 . The probability of moving from state 0 to 1 is $p_1 = 1 - p_0$. The state $X = 1$, is an absorbing state.

The system commences with $X = 0$, at time $t = 0$. Hence, the probability that the system will first reach the state 1 at some time t in the interval $(0, t']$ is

$$p(0, 1; t') = 1 - p_0^{t'}.$$

Suppose we wish to encourage absorption, i.e. reaching state 1 by time t' . An importance sampling strategy is to weight the probability of a transition from 0 to 1, by α .

Hence at time t , assuming we have not reached state 1, we draw X_t according to the probability mass function with

$$Pr(X_t = 0) = q_0 = \frac{p_0}{p_0 + \alpha p_1}$$

and

$$Pr(X_t = 1) = q_1 = \frac{\alpha p_1}{p_0 + \alpha p_1}.$$

Then, the associated weight should be multiplied by one of

$$w_0 = \frac{p_0}{q_0} \quad \text{or} \quad w_1 = \frac{p_1}{q_1}$$

depending on the transition type. The expected weight at time t' is given by

$$\begin{aligned} E[W_{t'}] &= q_1 w_1 + q_1 q_0 w_1 w_0 + q_1 q_0^2 w_1 w_0^2 + \dots \\ &= q_1 w_1 \sum_{i=0}^{t'-1} w_0^i q_0^i \\ &= 1 - p_0^{t'}. \end{aligned} \quad (5)$$

Hence, the importance weights can be used to unbiasedly estimate $p(0, 1; t')$ for any $\alpha > 0$. However, as Gillespie et al. (2009) showed, the efficiency of the method can vary dramatically

with the choice of α . To optimise the choice α , we minimise the variance of the estimate. For this model, we can derive an analytical expression for the second moment, namely

$$\begin{aligned} E[W_{t'}^2] &= q_1 w_1^2 + q_1 q_0 (w_1 w_0)^2 + q_1 q_0^2 (w_1 w_0^2)^2 + \dots \\ &= q_1 w_1^2 \sum_{i=0}^{t'-1} q_0^i w_0^{2i} \\ &= \frac{p_1^2}{q_1} \sum_{i=0}^{t'-1} \frac{p_0^{2i}}{q_0^i} \\ &= \frac{(1-p_0)^2}{1-q_0} \frac{q_0(1-p_0^{2t'} q_0^{-t'})}{q_0 - p_0^2}. \end{aligned} \quad (6)$$

To obtain an optimal value of α , we need to minimise

$$Var(W_{t'}) = E[W_{t'}^2] - E[W_{t'}]^2. \quad (7)$$

Since $E[W_{t'}]$ does not depend on α , minimising $Var(W_{t'})$ is equivalent to minimising $E[W_{t'}^2]$. The crucial insight from equation 6, is that $E[W_{t'}^2]$ explicitly depends on t' . This implies that the optimal value of α varies over time. This highlights that the methods proposed by Kuwahara and Mura (2008); Gillespie et al. (2009); Roh et al. (2010) are not optimal, since they assume a fixed value of α for all times $t \in [0, t']$.

For the case when $t' \rightarrow \infty$, the expectation simplifies to

$$E[W_{\infty}^2] = \frac{(1-p_0)^2(q_0-p_0)(p_0+q_0)}{(q_0-1)^2(q_0-p_0^2)^2}.$$

This is minimised by $q_0 = p_0$, i.e. setting $\alpha = 1$. This makes intuitive sense. Essentially as t' increases, the probability of hitting the target approaches 1. Hence, every simulation will eventually hit the target and so adjusting the probabilities is not needed.

For the general case, we can minimise $E(W_{t'}^2)$ by solving

$$\frac{d}{dq_0} E[W_{t'}^2] = 0 \quad (8)$$

for q_0 . After some differentiation we have that

$$\frac{t' p_0^{2t'} q_0^{-t'}}{(1-q_0)(q_0-p_0^2)} + \frac{q_0(1-p_0^{2t'} q_0^{-t'})}{(1-q_0)^2(q_0-p_0^2)} - \frac{q_0(1-p_0^{2t'} q_0^{-t'})}{(1-q_0)(q_0-p_0^2)^2} + \frac{1-p_0^{2t'} q_0^{-t'}}{(1-q_0)(q_0-p_0^2)} = 0. \quad (9)$$

Although equation (9) cannot be solved analytically, a standard numerical solver can be used to obtain q_0 for a given value of t' . Suppose that $p_0 = 0.9$ and interest lies in hitting $X = 1$ by time $t' = 5$. The optimal value of α (and therefore q_0) can be calculated for $t = 0, 1, \dots, 4$, under the assumption that for each t , we have not yet reached the absorbing state. For example, if at $t = 4$ we have not yet reached state 1, the optimal α is the value that minimises $E[W_{t'}^2]$ with $t' = 5 - t = 1$. Figure 1 shows that as t increases, and approaches $t' = 5$, the optimal value of α increases accordingly. When $t = 0$, $\alpha = 2.2$. This corresponds to $q_1 = 0.20$ ($p_1 = 0.1$). However, when $t = 4$, i.e. the last possible time at which X_t can be pushed towards state 1, $\alpha = 13.51$ and $q_1 = 0.60$.

Figure 1: Optimal values of α against t for the simple two-state model, with $p_0 = 0.9$ and $t' = 1$.

4 Guided wSSA

To obtain an importance density $q(x_{(0,\mathcal{T}]}|x_0)$, a suitable form of $\tilde{h}_i(x_t)$ (for use in the SSA) is required. We appeal to the recent literature on the sampling of jump processes that are conditioned to match some value at a given future time (Golightly and Wilkinson, 2015). Essentially, we seek a tractable approximation of the expected number of reaction events in an interval of the form $(t, t']$. The resulting expression is used as the propensity function in the SSA, giving draws from the corresponding $q(x_{(0,\mathcal{T}]}|x_0)$.

In what follows we assume that

$$x' = F^T x_{t'} \quad (10)$$

for some constant matrix F of dimension $u \times u_o$, with u_o denoting the number of components of x' . Hence, we allow for the rare event of interest to depend on a linear combination of the components of the MJP. For simplicity, consider the generation of a single draw from the importance density $q(x_{(0,\mathcal{T}]}|x_0)$. Suppose that we have simulated as far as time $t \in [0, \mathcal{T})$. Let ΔR_t denote the number of reaction events over the time $t' - t = \Delta t$. We approximate ΔR_t by assuming a constant reaction hazard over Δt . A Gaussian approximation to the corresponding Poisson distribution then gives

$$\Delta R_t \sim N(h(x_t)\Delta t, H(x_t)\Delta t)$$

where $H(x_t) = \text{diag}\{h(x_t)\}$. By linearity of (10) and using $X_t = S\Delta R_t$, we have that

$$X'|X_t = x_t \sim N(F^T(x_t + S h(x_t)\Delta t), F^T S H(x_t) S^T F \Delta t).$$

Hence, the joint distribution of ΔR_t and X' (conditional on x_t) can be obtained approximately as

$$\begin{pmatrix} \Delta R_t \\ X' \end{pmatrix} \sim N \left\{ \begin{pmatrix} h(x_t)\Delta t \\ F^T(x_t + S h(x_t)\Delta t) \end{pmatrix}, \begin{pmatrix} H(x_t)\Delta t & H(x_t)S^T F \Delta t \\ F^T S H(x_t)\Delta t & F^T S H(x_t) S^T F \Delta t \end{pmatrix} \right\}.$$

Taking the expectation of $\Delta R_t|X' = x'$ and dividing the resulting expression by Δt gives an approximate conditioned hazard as

$$\tilde{h}(x_t) = h(x_t) + H(x_t)S^T F (F^T S H(x_t)S^T F \Delta t)^{-1} \times (x' - F^T [x_t + S h(x_t)\Delta t]).$$

Algorithm 1: Guided wSSA

Input: N and x_0

```

1:  $m_N \leftarrow 0$ 
2: for  $n = 1$  to  $N$  do
3:    $t \leftarrow 0, \Delta t \leftarrow t' - t, x \leftarrow x_0, w \leftarrow 1$ 
4:   evaluate all  $h_j(x)$  and calculate  $h_0(x)$ 
5:   while  $t < t'$  do
6:     if  $x' = F^T x$  then
7:        $m_N \leftarrow m_N + w$ 
8:       break out of the while loop
9:     end if
10:    evaluate all  $d_i$  as in (11)
11:    if all  $d_i > 0$  then
12:      set  $\tilde{h}(x) = h(x)$ 
13:    else
14:      set  $\tilde{h}(x)$  using (4)
15:    end if
16:    generate two uniform  $U(0, 1)$  RNs  $r_1$  and  $r_2$ 
17:     $\tau \leftarrow -\tilde{h}_0^{-1}(x) \ln(r_1)$ 
18:     $j \leftarrow$  smallest integer satisfying  $\sum_{i=1}^j \tilde{h}_i(x) \geq r_2 \tilde{h}_0(x)$ 
19:     $w \leftarrow w \times [h_j(x)/\tilde{h}_j(x)] \times \exp\{[h_0(x) - \tilde{h}_0(x)]\tau\}$ 
20:     $t \leftarrow t + \tau, \Delta t \leftarrow t' - t, x \leftarrow x + S^j$ 
21:    update  $h(x)$  and  $h_0(x)$ 
22:  end while
23: end for
24: return  $\hat{p}(x_0, x'; t') = m_N/N$ 

```

Using $\tilde{h}(x_t)$ inside the SSA requires calculation of the conditioned hazard function in (4) after every reaction event. The cost of this calculation will be dictated by the dimension, u_o , of x' , given that a $u_o \times u_o$ matrix must be inverted. We anticipate that it will often be the case that $u_o \ll u$, for example when one gene is of primary interest in a larger system. Nevertheless, to avoid unnecessary computational cost, we suggest that at time t , a *pre-simulation check* is performed. We have approximately that

$$\begin{aligned} e(x_t) &= E(X'|X_t = x_t) \\ &= F^T(x_t + S h(x_t) \Delta t). \end{aligned}$$

Hence, we compare $E(X'|X_t = x_t)$ to x' by computing

$$d_i = (e_i(x_t) - x'_i) \times \text{sgn}(x'_i - [F^T x_0]_i), \quad i = 1, \dots, u_o \quad (11)$$

where, for example, $[F^T x_0]_i$ denotes the i th component of $F^T x_0$. Hence if all values d_i are positive, no additional push towards x' is necessary and we take $\tilde{h}(x_t) = h(x_t)$, thus avoiding calculation of the inverse of a $u_o \times u_o$ matrix. We note that $x'_i - [F^T x_0]_i$ need only be computed once, and that the calculation of $e(x_t)$ should be inexpensive, since it solves simple operations of the pre-specified quantities S and F , and the current hazard $h(x_t)$. Algorithm 1 gives the full wSSA (denoted “guided wSSA”) scheme based on the conditioned hazard function described above. Lines 10–15 describe the pre-simulation check.

The construction of the conditioned hazard is based on an assumption that the hazard function is constant over diminishing time intervals $(t, t']$ and that the number of reactions over this interval

is approximately Gaussian. The performance of the construct is therefore likely to be reduced if applied over time horizons during which the reaction hazards vary substantially. Additionally, the Gaussian assumption is likely to be unreasonable in low count scenarios, although we note that in such scenarios, simply using the SSA is likely to be computationally feasible. In what follows we investigate the performance of the proposed approach in three examples.

5 Examples

We consider three applications of increasing complexity. We compare and contrast the guided wSSA with vanilla use of the SSA, the wSSA of Kuwahara and Mura (2008) and the state-dependent wSSA of Roh et al. (2010). Following Gillespie et al. (2009), we compare the accuracy of each method by computing the sample variance of the unnormalised importance weights. That is the sample variance of

$$W^{(j)} = w_{\mathcal{T}}^{(j)} I(x' \in x^{(j)}), j = 1, \dots, N$$

where $w_{\mathcal{T}}^{(j)}$ is given by (3). We denote the resulting quantity by σ^2 . All methods were coded in R and run on were run on a desktop computer with an Intel Core i7-4770 processor at 3.40GHz

5.1 Single species production and degradation

Our first example is taken from Kuwahara and Mura (2008) (and see also Gillespie et al. (2009) and Roh et al. (2010)). The system contains two reactions

where $k_1 = 1$ and $k_2 = 0.025$. The initial state is $x_0 = (1, 40)^T$ we wish to estimate $p(x_0, F^T x_{t'}; t')$ with $t' = 100$, $F^T = (0, 1)$ and $x' = F^T x_{t'} = 80$. That is, the probability that $X_{2,t}$ reaches $x' = 80$ before time $t' = 100$ given an initial state of $x_0 = (1, 40)^T$ at $t = 0$. This is a well known model and the steady state population of $X_{2,t}$ follows a Poisson distribution with parameter $\lambda = k_1 x_{1,t} / k_2$ where $x_{1,t} = 1$ for all times t . Since the initial state $X_{2,0}$ is set at the equilibrium value, $X_{2,t}$ fluctuates around 40 with a standard deviation of $\sqrt{40}$. Hence, $x' = 80$ is the upper $10^{-6}\%$ quantile of the equilibrium distribution.

Figure 2 shows 3 trajectories that meet the attainment condition, generated using the guided wSSA (Algorithm 1). For this example, the pre-simulation check is performed at time t by computing $e(x_t) - x'$ where

$$e(x_t) = x_{2,t} + (k_1 x_{1,t} - k_2 x_{2,t})(t' - t). \quad (12)$$

If $x' - e(x_t) < 0$, we take $\tilde{h}_1(x_t) = h_1(x_t)$ and $\tilde{h}_2(x_t) = h_2(x_t)$. If $x' - e(x_t) > 0$, we use (4) to obtain

$$\tilde{h}_1(x_t) = h_1(x_t) + h_1(x_t) \frac{x' - e(x_t)}{v(x_t)} \quad (13)$$

and

$$\tilde{h}_2(x_t) = h_2(x_t) - h_2(x_t) \frac{x' - e(x_t)}{v(x_t)} \quad (14)$$

where $v(x_t) = (k_1 x_{1,t} + k_2 x_{2,t})(t' - t)$. In this case, the propensity of reaction $\mathcal{R}_1$ is increased and the propensity of reaction $\mathcal{R}_2$ is decreased.

Based on the trajectories in Figure 2, we computed the ratio $\tilde{h}_1(x_t)/\tilde{h}_2(x_t)$. We additionally computed the same quantity under the wSSA method of Kuwahara and Mura (2008), given by

$$1.85 \times \frac{h_1(x_t)}{h_2(x_t)} = \frac{1.85}{0.025 \times x_{2,t}}$$

corresponding to $\tilde{h}_1(x_t) = 1.85h_1(x_t)$ and $\tilde{h}_2(x_t) = h_2(x_t)$. Hence the Kuwahara/Mura approach is to bias reaction $\mathcal{R}_1$ by increasing its propensity. We also display the ratio of unaltered hazard functions

$$\frac{h_1(x_t)}{h_2(x_t)} = \frac{1}{0.025 \times x_{2,t}}$$

which we refer to as the baseline. We see that the propensity ratios under the baseline and Kuwahara/Mura approach exhibit a decreasing trend, and although the latter is able to increase the propensity of reaction $\mathcal{R}_1$, neither method is able to adapt to changes in x_t . The guided wSSA, on the other hand, is able to adjust the relative propensity depending on the behaviour of x_t . For example, in simulation 2 (middle column of Figure 2), the trajectory appears to be fairly constant between the values 60–65 between times 55–75. The guided wSSA therefore gives a sharp increase in relative propensity in order to meet the attainment condition.

We ran guided wSSA and the wSSA of Kuwahara/Mura for $N = 10^5$ trajectories and computed the sample variance σ^2 of the importance weights. For comparison, we also ran the vanilla SSA approach (using the unaltered relative propensity function given above), which required $N = 10^8$ to give a reasonable number of trajectories that meet the attainment condition. This gave variances of 1.1×10^{-12} for guided wSSA, 1.5×10^{-10} for the Kuwahara and Mura approach and 2.7×10^{-7} for the vanilla SSA approach. Hence, guided wSSA outperforms wSSA by two orders of magnitude and SSA by five orders of magnitude.

5.2 Two-state conformational transition

The following simple reaction network, initially described by Kuwahara and Mura (2008), models two conformational isomers, that is, isomers that can be inter-converted by rotation about single bonds. This model has two reactions

where $k_1 = 0.12$ and $k_2 = 1$. The initial state is set to $x_0 = (100, 0)^T$. For this particular system we wish to estimate $p(x_0, F^T x_{t'}; t')$ with $t' = 10$, $F^T = (0, 1)$ and $x' = F^T x_{t'} = 30$. That is, the probability that given $x_0 = (100, 0)^T$, the value of $x_{2,t}$ reaches $x' = 30$ before time $t' = 10$. The steady state population of $X_{2,t}$ is approximately 11, so the rare event of $x' = 30$ by time $t' = 10$, is around three times larger than its equilibrium value. Since this system is closed, we can numerically calculate this probability to obtain

$$p(x_0, x'; t') = 1.191 \times 10^{-5}.$$

Since the steady state population of $X_{2,t}$ is less than the target value of 30, we can bias the system by either increasing the rate of reaction R_1 or decreasing the rate of reaction R_2 .

The implementation of the guided wSSA follows closely the approach described in Section 5.1. The pre-simulation check calculates $e(x_t) - x'$ where $e(x_t)$ is given by (12). For the non-trivial case of $x' - e(x_t) > 0$, $\tilde{h}_1(x_t)$ and $\tilde{h}_2(x_t)$ are given by (13) and (14) respectively. We compare the output of the guided wSSA against that of the wSSA described by Kuwahara and Mura (2008) and the state-dependent approach of Roh et al. (2010)). For the former, the first reaction R_1 , was encouraged via the biasing parameter γ . This corresponds to

$$\tilde{h}_1(x_t) = \gamma h_1(x_t) \quad \text{and} \quad \tilde{h}_2(x_t) = h_2(x_t).$$

The optimal biasing parameter was estimated Roh et al. (2010) by selecting the value of γ that minimised the variance of the estimator of $p(x_0, x'; t')$. This resulted in $\gamma = 1.4$. The Roh et al. method encourages reaction $\mathcal{R}_1$ with the state-dependent biasing parameter

$$\gamma(\rho(x_t)) = \begin{cases} 1 & \text{if } \rho(x_t) \geq \rho^0 \\ g(\rho(x_t)) & \text{if } \rho(x_t) < \rho^0 \end{cases}$$

Figure 2: Top panel (a): 3 trajectories using the guided SSA that attain $x' = 80$ by $t' = 100$. Bottom panel (b): ratios $\tilde{h}_1(x_t)/\tilde{h}_2(x_t)$ against time t (where x_t is given in the top panel) under the guided wSSA (black) and Kuwahara/Mura approach (green). The ratio of unaltered hazards $h_1(x_t)/h_2(x_t)$ is also shown (red line).

where

$$g(\rho(x_t)) = \frac{\gamma^{max} - 1}{(\rho^0)^2} (\rho^0 - \rho(x_t))^2 + 1$$

and $\rho(x_t) = h_1(x_t)/h_0(x_t)$. As discussed in Roh et al. (2010), $g(\rho(x_t))$ has the desirable properties that $g(0) = \gamma^{max}$, the maximum biasing value allowed, and $g(\rho^0) = 1$, to avoid over-perturbing the system beyond a relative propensity of ρ^0 . For this example, the authors suggest $\rho^0 = 0.5$ and $\gamma^{max} = 20$.

Figure 3 shows the sample variance σ^2 of the importance weights obtained from four runs of each method (and additionally, the vanilla SSA approach) using $N = 10^5$. The guided wSSA and Roh et al. approaches give comparable variances whereas the wSSA of Kuwahara/Mura gives variances that are typically at least two orders of magnitude greater than those exhibited by the former. Simply using the SSA gives importance weights whose variance is around four orders of magnitude greater than the guided approach. We stress the automatic nature of the guided wSSA approach - no tuning parameters are required. For further comparison, we also display results for $t' = 30$.

Figure 3: A plot of σ^2 for the two-state conformational model. Each of the four points is calculated using $N = 10^5$ for the guided wSSA and two competing methods. Upper panel ($t = 10$): The variance associated with estimating $p((100, 0)^T, 30; t' = 10)$ and lower panel ($t = 30$): $p((100, 0)^T, 30; t' = 30)$.

5.3 Motility Regulation

In this section we consider a simplified model of a key cellular decision made by the gram-positive bacterium *Bacillus subtilis* (Sonenshein et al., 2002). This decision is whether or not to grow flagella and become motile (Kearns and Losick, 2005). The *B. subtilis* sigma factor *sigD* is key for the regulation of motility. Many of the genes and operons encoding motility-related proteins are governed by *sigD*, and so understanding its regulation is key to understanding the motility decision. The gene for *sigD* is embedded in a large operon containing several other motility-related genes, known as the *fla/che* operon. Transcription of the operon is strongly repressed by the protein *CodY*, which is encoded upstream of *fla/che*. *CodY* inhibits transcription by binding to the *fla/che* promoter. Since *CodY* is upregulated in good nutrient conditions, this is thought to be a key mechanism for motility regulation. For simplicity we focus here on one gene under the control of *sigD*, that is *hag*, which encodes the protein *flagellin* (or *Hag*), the key building block of the flagella. The reaction network has $u = 9$ species and $v = 12$ reactions and is encoded as

Figure 4: Ten stochastic realisations of the motility regulation model. The rare event of interest is indicated (see bottom left).

follows:

Values of the rate constants are taken to be

$$k = (0.1, 0.0002, 1, 0.0002, 1.0, 0.0002, 0.01, 0.1, 0.02, 0.1, 0.01, 0.1)^T$$

Figure 5: Motility regulation model: expectation (left panel) and variance (right panel) of the importance weights. Each of the eight points is calculated using $N = 10^6$ for guided wSSA $N = 10^8$ for SSA.

and the initial values of

$$(\text{codY}, \text{CodY}, \text{flache}, \text{SigD}, \text{SigD_hag}, \text{hag}, \text{Hag}, \text{CodY_flache}, \text{CodY_hag})$$

are

$$x_0 = (1, 10, 1, 10, 1, 1, 10, 1, 1)^T.$$

Figure 4 gives ten realisations of the system with these settings. Note that trajectories for each species are inherently discrete. Since the guided wSSA is based on a normal approximation of the MJP, we anticipate that these parameter settings provide a challenging scenario for the proposed methodology.

Given the importance of *CodY* in motility regulation, we assume that interest lies in the estimation of $p(x_0, F^T x_{t'}; t')$ with $t' = 10$, $F^T = (0, 1, 0, \dots, 0)$ and $x' = F^T x_{t'} = 20$. That is, the probability that given x_0 as above, the value of *CodY* reaches $x' = 20$ before time $t' = 10$. Given the size of the reaction network, tuning a weighted SSA approach is likely to be difficult. We therefore compare guided wSSA with a vanilla SSA approach, since neither method requires tuning. For guided wSSA, we used $N = 10^6$ and for SSA we used $N = 10^8$, with the latter required to produce a reasonable number of trajectories that met the attainment condition. Our results are reported in Figure 5. The computational cost of guided wSSA versus SSA scales roughly as 2 : 1 for this example. We note however that this is computational effort that is well worth spending, since the guided wSSA gives values of σ^2 (the variance of the importance weights) that are typically three orders of magnitude smaller than when using the SSA.

6 Conclusions

Simulation-based approaches to rare event probability estimation can be computationally prohibitive. The weighted stochastic simulation algorithm (wSSA) aims to avoid wasteful simulation by using an importance sampler to push trajectories towards the rare event of interest. The resulting biased trajectories are appropriately weighted, and the weights are averaged to give an

unbiased estimator of the rare event probability. Constructing the importance sampler usually requires insight into the underlying reaction system, careful tuning or additional simulations.

In this paper we have introduced an importance sampler that requires no tuning. We have leveraged the tractability of a Gaussian approximation to the Markov jump process to construct a propensity function that is conditioned on the rare event occurring at a future time. Use of this propensity inside the SSA gives the guided wSSA. Our experiments suggest that the guided wSSA can provide a practical approach to rare event probability estimation, by comprehensively outperforming the unweighted SSA method, and giving comparable results to other wSSA approaches, without the need for laborious determination of suitable tuning parameters. While the Gaussian approximation (on which the conditioned propensity is based) is likely to be unsatisfactory when species exhibit inherently discrete behaviour, our findings suggest that the guided wSSA is relatively robust in such low-count number scenarios. Nevertheless, improving the efficiency of the guided wSSA (e.g. by leveraging more accurate moment-closure based approximations to the MJP) remain the topic of future work.

References

- Daigle, B. J., M. K. Roh, D. T. Gillespie, and L. R. Petzold (2011, jan). Automated estimation of rare event probabilities in biochemical systems. *The Journal of chemical physics* 134(4), 044110.
- Gibson, M. A. and J. Bruck (2000, March). Efficient Exact Stochastic Simulation of Chemical Systems with Many Species and Many Channels. *The Journal of Physical Chemistry A* 104(9), 1876–1889.
- Gillespie, C. S. (2012). Stochastic simulation of chemically reacting systems using multi-core processors. *The Journal of chemical physics* 136(1), 014101.
- Gillespie, D. T. (1977). Exact stochastic simulation of coupled chemical reactions. *The journal of physical chemistry* 81(25), 2340–2361.
- Gillespie, D. T., M. Roh, and L. R. Petzold (2009). Refining the weighted stochastic simulation algorithm. *Journal of Chemical Physics* 130(17).
- Golightly, A. and D. J. Wilkinson (2015). Bayesian inference for Markov jump processes with informative observations. *Statistical Applications in Genetics and Molecular Biology* 14(2), 169–188.
- Kearns, D. B. and R. Losick (2005). Cell population heterogeneity during growth of *bacillus subtilis*. *Genes & development* 19(24), 3083–3094.
- Kuwahara, H. and I. Mura (2008). An efficient and exact stochastic simulation method to analyze rare events in biochemical systems. *Journal of Chemical Physics* 129(16).
- McCollum, J. M., G. D. Peterson, C. D. Cox, M. L. Simpson, and N. F. Samatova (2006, February). The sorting direct method for stochastic simulation of biochemical systems with varying reaction execution behavior. *Computational Biology and Chemistry* 30(1), 39–49.
- Roh, M. K., B. J. Daigle, D. T. Gillespie, and L. R. Petzold (2011). State-dependent doubly weighted stochastic simulation algorithm for automatic characterization of stochastic biochemical rare events. *The Journal of Chemical Physics* 135(23), 234108.

- Roh, M. K., D. T. Gillespie, and L. R. Petzold (2010). State-dependent biasing method for importance sampling in the weighted stochastic simulation algorithm. *The Journal of Chemical Physics* 133(17), 174106.
- Slepoy, A., A. P. Thompson, and S. J. Plimpton (2008, May). A constant-time kinetic Monte Carlo algorithm for simulation of large biochemical reaction networks. *The Journal of Chemical Physics* 128(20), 205101.
- Sonenshein, A. L., H. J. A., and R. Losick (Eds.) (2002). *Bacillus subtilis and its closest relatives*. ASM Press.
- Wilkinson, D. J. (2012). *Stochastic Modelling for Systems Biology* (2nd ed.). Boca Raton, Florida: Chapman & Hall/CRC Press.