

This item is likely protected under Title 17 of the U.S. Copyright Law. Unless on a Creative Commons license, for uses protected by Copyright Law, contact the copyright holder or the author.

Access to this work was provided by the University of Maryland, Baltimore County (UMBC) ScholarWorks@UMBC digital repository on the Maryland Shared Open Access (MD-SOAR) platform.

Please provide feedback

Please support the ScholarWorks@UMBC repository by emailing scholarworks-group@umbc.edu and telling us what having access to this work means to you and why it's important to you. Thank you.

Stochastic Filtering of Reaction Networks Partially Observed in Time Snapshots

Muruhan Rathinam¹ and Mingkai Yu²

¹Department of Mathematics and Statistics, University of Maryland
Baltimore County

²Department of Mathematics and Statistics, University of Maryland
Baltimore County

August 1, 2023

Abstract

Stochastic reaction network models arise in intracellular chemical reactions, epidemiological models and other population process models, and are a class of continuous time Markov chains which have the nonnegative integer lattice as state space. We consider the problem of estimating the conditional probability distribution of a stochastic reaction network given exact partial state observations in time snapshots. We propose a particle filtering method called the *targeting method*. Our approach takes into account that the reaction counts in between two observation snapshots satisfy linear constraints and also uses inhomogeneous Poisson processes as proposals for the reaction counts to facilitate exact interpolation. We provide rigorous analysis as well as numerical examples to illustrate our method and compare it with other alternatives.

1 Introduction

Stochastic reaction network models describe a wide array of phenomena such as intracellular biochemical reactions, epidemiological models, predator prey systems etc. These models are a particular form of continuous time Markov chains

(CTMCs) $Z(t)$ that have the multidimensional nonnegative integer lattice \mathbb{Z}_+^n as state space and have finitely many types of “reaction” events $j = 1, \dots, m$ which change the state by a fixed amount $v_j \in \mathbb{Z}^n$. Given the wide applicability of these models, the problem of inferring the state as well as parameters from partial observations is of critical interest.

In this paper, we are focused on stochastic reaction networks where some of the states, say $Y(t)$ where $Z(t) = (X(t), Y(t))$, are observed exactly at time snapshots $t_1 < t_2 < \dots < t_{T_s}$. We are interested in estimating the conditional probability

$$\pi(t, z) = P\{Z(t) = z | Y(t_l) = y_l, l = 1, \dots, T_s\}$$

via a suitable filtering algorithm. More generally, our methods can be used to estimate the conditional expectation of a path functional given a sequence of observations. Starting from the celebrated Kalman filter [20] for discrete time linear systems with Gaussian noise, several filtering methods have been developed for various dynamic models [21, 7, 6, 12, 17, 18, 5, 9]. In linear and Gaussian models the conditional distributions themselves are Gaussian and are characterized by the mean and the covariance. This means that it is possible to obtain (deterministic) evolution equations for the conditional mean and covariance which makes the computations relatively straightforward. However, when the dynamic models are nonlinear or if the noise in the dynamics are non-Gaussian, the situation is no longer simple. While deterministic evolution equations may be obtained for the conditional probabilities, these equations are infinite dimensional or very high dimensional. This necessitates Monte Carlo methods. Particle filtering, also known as Sequential Monte Carlo (SMC), are usually the most practical approach for such problems since these methods enable recursive computations as new observations are added [7, 6, 2]. Particle filters can also be used for Bayesian estimation of parameters by treating parameters as states that are constant in time, thereby providing a unified approach for state and parameter estimation. As such, our focus in this paper is on the state estimation. If the goal is only to estimate parameters, one may approach the problem in a non-Bayesian setting and several methods exist in the literature [22, 19].

Particle filtering methods involve simulating identically distributed proxies $\tilde{Z}^{(i)}(t)$ (which we call the “filter processes”) of the reaction network along with importance weights $W^{(i)}(t)$ for $i = 1, \dots, N_s$. The fidelity of the filter estimate depends on N_s as well as the variance of $W^{(i)}(t)$. Usually, the variance of $W^{(i)}(t)$ tends to increase in time, leading to poor fidelity of the estimate. Thus most methods are concerned with coming up with ways to slow down the growth of variance.

We briefly describe some of the results in the literature for particle filtering of stochastic reaction networks. The work in [16] considers partial state observations in discrete time with added Gaussian noise and proposes particle filtering methods where the proxies $\tilde{Z}^{(i)}$ are simulated according to the conditional probability of a Gaussian approximation of Z . [8, 9] are concerned with situations where some molecular copy numbers are very large so that a hybrid approximation of the process Z may be employed and provide rigorous mathematical analysis to justify the method. The first of these is concerned with discrete time observations while the second is concerned with both discrete time as well as continuous time observations. A related work [10] studies the convergence of a regularized form of the particle filter. All of the above works are concerned with noisy observations where the observation noise is additive Gaussian even though the reaction network model is not.

Our previous work [25] provides a particle filter for a reaction network where exact (noiseless) partial state observations are made in continuous time. To our best knowledge this is the first such filter for a CTMC. In this paper, we provide new particle filtering methods for the case of exact partial state observations in discrete time snapshots. Since we are concerned with noiseless observations, the filters mentioned previously (which apply to observations with Gaussian noise) are not suitable. At first glance, noiseless observations may seem unrealistic. However, noisy observation processes themselves could be included in the reaction network model as additional reactions. For instance, if we observe a fluorescently tagged molecular species S , then the actual observations count the copy number of photons P emitted by S . Hence this measurement process may be modeled by adding a reaction $S \rightarrow S + P$ and we consider P as the observed species.

The conceptually simplest approach to filtering for discrete time observations (with or without noise) involves the “prediction/correction” approach. If we denote by $\pi_{l,j}(z)$ the conditional probability distribution of $Z(t_l)$ given observations up to time point t_j , then one computes $\pi_{l,l}$ from $\pi_{l-1,l-1}$ in two steps. The first step is the prediction step which simply evolves the proxy process \tilde{Z} according to the Markovian evolution of Z to obtain weighted samples from $\pi_{l,l-1}$ and then a correction is performed by altering the weights appropriately so that the weighted sample reflects $\pi_{l,l}$. In the noiseless case, this approach will result in most of the weights set to zero. Even if one modifies the law of the proxy \tilde{Z} via Girsanov changes of intensity, the probability of satisfying the observation at the next observation time t_l may be negligibly small. The method we propose in this paper forces the process \tilde{Z} to satisfy the next observation and we call this the *targeting method*. In order to target the next observation we work with a Girsanov change of

measure under which the reaction counts are inhomogeneous Poissons. We provide rigorous justification of our method and illustrate its superior performance compared to other alternatives.

The rest of the paper is organized as follows. Section 2 provides a brief review of stochastic reaction networks and states the filtering problem addressed in this paper. In Section 3 we discuss the mathematical setup, motivate our method which we call the *targeting algorithm* as well as discuss other ideas explored in the literature. Section 5 first describes the reaction network examples that are later used in Section 6 to illustrate various choices within our method and compare with other methods. Finally, in Section 7, we offer some concluding remarks.

2 Reaction Networks and The Filtering Problem

We describe the widely used stochastic model of a *reaction network*. A reaction network consists of $n \in \mathbb{N}$ *species* and $m \in \mathbb{N}$ *reaction channels*. The copy number of the species i ($1 \leq i \leq n$) at time $t \geq 0$ is denoted by $Z_i(t) \in \mathbb{Z}_+$ and the number of firings of reaction channel j ($1 \leq j \leq m$) during $(0, t]$ is denoted by $R_j(t) \in \mathbb{Z}_+$. The vector valued processes X and R are defined by $Z(t) = (Z_1(t), \dots, Z_n(t))$ and $R(t) = (R_1(t), \dots, R_m(t))$ and Z is assumed to be a continuous time Markov chain (CTMC) with statespace \mathbb{Z}_+^n . The characteristics of a reaction network are uniquely defined by the *stoichiometric vectors* $v_j \in \mathbb{Z}^n$ and the *propensity functions* $a_j : \mathbb{Z}_+^n \rightarrow [0, \infty)$ for $j = 1, \dots, m$. If the reaction channel j fires at time t , the state Z of the process is changed by v_j . We assume that the process Z is *cadlag*, that is, right continuous with left hand limits. We shall define the matrix $v \in \mathbb{Z}^{n \times m}$ so that its j th column equals v_j . The (i, j) th entry of v is denoted v_{ij} which is the i th entry of the column vector v_j . Thus we may write

$$Z(t) = Z(0) + vR(t). \quad (1)$$

The propensity function a_j has the interpretation that given $Z(t) = z$ (for $z \in \mathbb{Z}_+^n$), the probability that the reaction channel j fires during $(t, t+h]$ is given by $a_j(z)h + o(h)$ as $h \rightarrow 0+$. It follows that R is an m -variate *counting process* with *stochastic intensity* $a(Z(t-))$ [4]. We assume that the processes Z and R are carried by a probability space (Ω, \mathcal{F}, P) .

We shall be concerned with the situation where a subset of the species are observed at time snapshots $0 \leq t_1 < t_2 \dots$. We shall take the vector of observed species to consist of the last n_2 components of the state $Z(t)$ and denote it by $Y(t)$. The vector of the rest of the unobserved species shall be denoted by $X(t)$. Thus we

write $Z(t) = (X(t), Y(t))$ where $X(t) \in \mathbb{Z}_+^{n_1}$ and $Y(t) \in \mathbb{Z}_+^{n_2}$ with $n = n_1 + n_2$. We shall write $\mathbf{v}_j = (\mathbf{v}'_j, \mathbf{v}''_j)$ where $\mathbf{v}'_j \in \mathbb{Z}^{n_1}$ and $\mathbf{v}''_j \in \mathbb{Z}^{n_2}$. We shall define $\mathbf{v}' \in \mathbb{Z}^{n_1 \times m}$ and $\mathbf{v}'' \in \mathbb{Z}^{n_2 \times m}$ to consist of the first n_1 rows of \mathbf{v} and the last n_2 rows of \mathbf{v} respectively. Thus we may write

$$\begin{aligned} X(t) &= X_0 + \mathbf{v}'R(t), \\ Y(t) &= Y_0 + \mathbf{v}''R(t), \end{aligned} \tag{2}$$

where $Z_0 = (X_0, Y_0)$ is the random initial state.

The filtering problem that we are concerned with is the numerical computation of conditional probabilities of the form:

$$\pi(t, z) = P\{Z(t) = z | Y(t_i) = y_i, i = 1, \dots, T_s\} \tag{3}$$

where $t \geq 0$, $z \in \mathbb{Z}_+^n$ and T_s is the number of snapshots. Due to the fact that the number of states is either infinity or exceedingly large, Monte Carlo methods are the most feasible way to compute $\pi(t, z)$. In addition to the observation, we assume that we know the distribution of $Z_0 = (X_0, Y_0)$. We note that $\pi(t, z)$ depends also on the observation sequence $\bar{y} = (y_1, \dots, y_{T_s})$ so that sometimes we write $\pi(t, z, \bar{y})$ to emphasize this dependence.

3 The Filtering Algorithm

In this section we present a mathematical setup to describe filtering methods and introduce what we call the Naive Algorithm and describe the concept of *effective sample size*. We also discuss a framework that underlies some filtering algorithms in the literature as well as present our proposed filtering method.

3.1 Mathematical Setup

We shall be concerned with weighted Monte Carlo methods which are also known as *importance sampling* or *particle filters*. We present a mathematical description of a general setup as follows. We let (Ω, \mathcal{F}) be a measurable space consisting of the sample space and a collection of events.

It is useful to think of two sets of processes: one on which observations are made and the other consists of the Monte Carlo simulations. We refer to the former as the *lab processes* and the latter as the *method processes* or *filter processes*. Both

sets of processes are carried by (Ω, \mathcal{F}) . Conditioned on the observations, the lab processes and the method processes are independent.

The lab processes consist of $Z = (X, Y)$ and R , where $Z : [0, T_0] \times \Omega \rightarrow \mathbb{Z}_+^n$ and $R : [0, T_0] \times \Omega \rightarrow \mathbb{Z}_+^m$ are measurable, and $T_0 > 0$ is large enough so that $[0, T_0]$ contains all observation time points t_l .

The method processes consist of $\tilde{Z}^{(i)} = (U^{(i)}, V^{(i)}) : [0, T_0] \times \Omega \rightarrow \mathbb{Z}_+^n$, $I^{(i)} : [0, T_0] \times \Omega \rightarrow \mathbb{Z}_+^m$ and $W^{(i)} : [0, T_0] \times \Omega \rightarrow [0, \infty)$ for $i = 1, \dots, N_s$. Here N_s is the *filter sample size* (number of particles) and the method processes for different i are all identically distributed and for convenience we drop the superscript i when describing them, or one may regard $\tilde{Z} = (U, V)$, I and W as the template method processes. The method processes U, V and I are proxies for X, Y and R respectively. The process W is the weight process that assigns a degree of importance to the paths of U, V and I .

Often more than one probability measure on (Ω, \mathcal{F}) is considered. The lab processes have the given propensities under a probability measure P while under another probability measure P_0 (with $P \ll P_0$) both the lab and filter processes have a convenient description. In general, the quantity of interest $\pi(t, z, \bar{y})$ is given as a function of the conditional expectation of path functionals of the filter processes. Going forward we denote the observation event sequence by $O_{\bar{y}}$.

3.2 The Naive Algorithm

It might be clear from the outset that the most straightforward Monte Carlo algorithm is simply based on the following prediction/correction strategy without any changes of measure. One simulates N_s sample trajectories $\tilde{Z}^{(i)}(t) = (U^{(i)}(t), V^{(i)}(t))$ for $i = 1, \dots, N_s$ according to the given reaction network characteristics via the Gillespie algorithm or a variant [14, 15, 13, 1] and assigns a weight $W^i = 1$ (accept) if and only if $V^{(i)}(t_l) = y_l$ for $l = 1, \dots, T_s$ and $W^i = 0$ (reject) otherwise. We note that $\pi(t, z, \bar{y})$ is given by

$$\pi(t, z, \bar{y}) = \frac{\mathbb{E}[1_{\{z\}}(\tilde{Z}(t)) W \mid O_{\bar{y}}]}{\mathbb{E}[W \mid O_{\bar{y}}]}$$

and one simply estimates π via the filter sample as

$$\hat{\pi}(t, z, \bar{y}) = \frac{\sum_{i=1}^{N_s} 1_{\{z\}}(\tilde{Z}^{(i)}(t)) W^i}{\sum_{i=1}^{N_s} W^i}.$$

We note that the weights W^i depend on \bar{y} .

We shall refer to this method as the *Naive Algorithm*. While it is conceptually simple and easy to implement, in many situations the probability of any given sequence of observation snapshots is very low and hence the “effective sample size” is much smaller than N_s . Hence in order to obtain reliable estimates one needs to simulate impractically large numbers N_s of trajectories. However, the naive algorithm has one advantage in that it is conditionally unbiased:

$$\mathbb{E} \left(\frac{1}{\sum_{i=1}^{N_s} W^i} \sum_{i=1}^{N_s} 1_{\{z\}}(\tilde{Z}^{(i)}(t)) W^i \mid \sum_{i=1}^{N_s} W^i \neq 0, O_{\bar{y}} \right) = \pi(t, z, \bar{y}).$$

3.3 Effective Sample Size

Often the filtering method involves estimating a quantity c given by

$$c = \frac{\mathbb{E}[CW \mid O_{\bar{y}}]}{\mathbb{E}[W \mid O_{\bar{y}}]}$$

where C is a path functional and W is an associated weight. The quantity c is estimated via the empirical mean

$$\hat{c} = \frac{\sum_{i=1}^{N_s} C^i W^i}{\sum_{i=1}^{N_s} W^i}$$

from a weighted sample (C^i, W^i) for $i = 1, \dots, N_s$ identically distributed like (C, W) . The fidelity of this estimate depends on the joint (conditional) distribution of (C^i, W^i) for $i = 1, \dots, N_s$. Law of large numbers (which applies under modest assumptions) states that for sufficiently large N_s , \hat{c} is a good approximation of c . If we assume conditional independence (for different i), when N_s is sufficiently large, under modest assumptions on moments we may use the so-called delta method [3] and obtain the approximation that \hat{c} is (approximately) normally distributed with $\mathbb{E}[\hat{c} \mid O_{\bar{y}}] \approx c$ and

$$\text{Var}(\hat{c} \mid O_{\bar{y}}) \approx \frac{\mathbb{E}^2[CW]}{N_s \mathbb{E}^2[W]} \text{Var} \left(\frac{CW}{\mathbb{E}[CW]} - \frac{W}{\mathbb{E}[W]} \mid O_{\bar{y}} \right).$$

The square of the *relative error* $\text{RE} = \frac{\sqrt{\mathbb{E}[(\hat{c}-c)^2 \mid O_{\bar{y}}]}}{c}$ is then approximated by

$$\text{RE}^2 \approx \frac{1}{N_s} \text{Var} \left(\frac{CW}{\mathbb{E}[CW]} - \frac{W}{\mathbb{E}[W]} \mid O_{\bar{y}} \right).$$

The fidelity of the estimate clearly depends on the path functional C and its conditional covariance with W among other factors. For instance, if one is interested in estimating $\pi(t, z)$, then $C = 1_{\{z\}}(\tilde{Z}(t))$, and the fidelity of the estimate varies with z . In order to have a measure of fidelity that is independent of the path functional C and solely dependent on W , the quantity

$$N_e = \frac{\left(\sum_{i=1}^{N_s} W^i\right)^2}{\sum_{i=1}^{N_s} (W^i)^2} \quad (4)$$

known as *effective sample size* has been proposed in the literature [23]. When N_s is very large

$$N_e \approx N_s \frac{\mathbb{E}^2[W|O_{\bar{y}}]}{\mathbb{E}[W^2|O_{\bar{y}}]}. \quad (5)$$

Thus, for the naive method, $N_e \approx N_s p$ where p is the probability of the sequence of observations. When the molecular copy numbers of the species are in the order of hundreds and if multiple species are observed this probability is likely to be very small.

3.4 Conditional Propensity Function

Recall that the propensity $a_j(z)$ is defined by the prescription that the conditional probability of one reaction event of j occurs during $(t, t+h]$ conditioned on $Z(t) = z$ is given by $a_j(z)h + o(h)$ as $h \rightarrow 0+$. We may define the *conditional propensity* $b_j(z, t, t_1, y_1 \dots)$ of the reaction channel j by replacing the conditioned event $\{Z(t) = z\}$ by the event

$$\{Z(t) = z, Y(t_l) = y_l \ l = 1, \dots, T_s\}$$

in the above definition so that it includes the observations. See [11, 16] for instance. Suppose that $t_{l^*-1} < t \leq t_{l^*}$. Then using the Markov property, it can be shown that

$$b_j(z, t) = a_j(z) \frac{P(Y(t_l) = y_l, l \geq l^* | Z(t) = z + v_j)}{P(Y(t_l) = y_l, l \geq l^* | Z(t) = z)}, \quad (6)$$

where we have suppressed the dependence on the observation sequence (t_l, y_l) . We note that due to the explicit dependence on time t , the (conditional) wait time between reaction events is no longer exponential. If an explicit formula for $b_j(z, t)$

is available and the wait times have an analytically tractable formula (as we shall see in a specific example) it is indeed possible to simulate the reaction network according to these conditional propensities. In that event, all trajectories will satisfy the observations.

In general, an explicit formula for b_j is not available and hence it is impossible to draw samples from the conditional distribution. However, if a good approximation of b_j is proposed, which we denote by $\hat{a}_j(z, t)$, then one may simulate the reaction network according to these “proposal propensities”. Since the proposal propensities may not be exact, a weight $\tilde{L}_{\bar{y}}(T)$ given by the Girsanov change of measure that corresponds to the change in propensity from \hat{a}_j to a_j (for the method process I) is computed (from the trajectory \tilde{Z}) and applied. Additionally, the indicator weight W as in the naive method is applied depending on whether the observations are satisfied or not, so that the overall weight is $W\tilde{L}_{\bar{y}}(T)$. We note that the change of measure weight is observation dependent and hence the subscript \bar{y} denoting the observation sequence.

If the proposal propensities \hat{a}_j are exactly equal to the conditional propensities b_j , then with probability one, all trajectories \tilde{Z} will satisfy the observations and $W\tilde{L}(T)$ will be a constant, making the effective sample size N_s . If \hat{a}_j are close enough to b_j , then one expects the effective sample size to be close to N_s . This is the strategy that underlies the methods in [16].

3.5 The Targeting Problem and the Targeting Algorithm

In order to better motivate the new method proposed in this paper, we focus on the problem of a single observation snapshot which we take to be $Y(T) = y$ at $t = T > 0$. So we have a random initial state Z_0 with a known distribution μ and we are given the partial final state condition $Y(T) = y$. Suppose our goal is to construct a process $\tilde{Z} = (U, V)$ which is proxy for $Z = (X, Y)$ such that $V(T) = y$. That is, \tilde{Z} always satisfies the partial terminal condition $V(T) = y$. This ensures that the indicator weight W is always 1, thereby reducing an important source of variance. See Remark 1 below. We call this approach “targeting” since the filter process $\tilde{Z} = (U, V)$ always reaches the target $V(T) = y$.

Our approach to the targeting problem is based on two key ideas. The first is that the partial final state constraint $V(T) = y$ can be translated into reaction counts constraints for some of the reactions. The second is that for a (possibly inhomogeneous) Poisson process if the value at time $T > 0$ is known, then one may simulate exact conditional realizations of its path on $[0, T]$. We explain below in detail.

After reordering and splitting the reactions $R(t) = (R'(t), R''(t))$, and possible removal of dependent observable species (corresponding to dependent rows of v'') we may write the modified v'' as

$$v'' = [A \ B] \quad (7)$$

where B is invertible. This yields

$$y = Y_0 + AR'(T) + BR''(T)$$

and hence

$$R''(T) = B^{-1}(y - Y_0 - AR'(T)). \quad (8)$$

Thus, we see that $R''(T)$ is determined by Y_0, y and $R'(T)$. We shall make our proxy processes \tilde{Z} and I satisfy this condition. Since we want $V(T) = y$, we shall require

$$y = V_0 + AI'(T) + BI''(T),$$

so

$$I''(T) = B^{-1}(y - V_0 - AI'(T)). \quad (9)$$

If we choose $I''(T)$ according to (9) we automatically satisfy the target observation! This motivates our key strategy as outlined below. In what follows, we define $K = I(T)$.

1. Given $Y(T) = y$, generate a sample for $\tilde{Z}_0 = (U_0, V_0)$ as well as $K = I(T)$ so that \tilde{Z}_0 is distributed like μ , and an associated weight $W > 0$ so that

$$P_0\{R(T) = k, Z_0 = z_0 | Y(T) = y\} = \mathbb{E}_0[1_{(z_0, k)}(\tilde{Z}_0, K) W] / \mathbb{E}_0[W],$$

where P_0 is the “simulation probability measure”. This is accomplished by the following procedure which can be thought of as consisting of prediction and correction steps.

- (a) Generate a sample for \tilde{Z}_0 and $K' = I'(T)$ (prediction).
- (b) With the knowledge of $V_0, V(T) = y$ and $K' = I'(T)$ compute $K'' = I''(T)$ via (9). If $I''(T) \in \mathbb{Z}_+^{m_2}$ accept, otherwise reject and go back to the previous step.
- (c) Set the weight W to reflect the probability of the assigned value for $I''(T)$. (Correction).

2. Having obtained a value for \tilde{Z}_0 and $I(T)$ along with a weight $W > 0$, obtain $I(t)$ for $0 \leq t \leq T$ by “interpolation”.

Generating a sample from $I'(T)$ can in principle be done by Gillespie method simulation. However, computing the weight W as the probability of the value of $I''(T)$ given by (9) in closed form is not simple. Furthermore, the interpolation of $I(T)$ for t values in $[0, T]$ is not easy either.

However, the above steps can be solved with relative ease if the m -variate counting process R is assumed to be an inhomogeneous Poisson process. Hence our strategy is to take R as an m -variate inhomogeneous Poisson process (when conditioned on Z_0) with a time varying intensity that possibly depends on Z_0 and the observed value y of $Y(T)$ under a y dependent probability measure P_0^y . With this strategy it can be shown that

$$P_0^y\{Z_0 = z_0, R(T) = k | Y(T) = y\} = \frac{\mathbb{E}_0^y[1_{\{(z_0, k)\}}(\tilde{Z}_0, K) W]}{\mathbb{E}_0^y[W]}.$$

Then a y dependent Girsanov transformation is performed via an associated weight $L_y(T)$ so that under the resulting probability measure P^y given by $dP^y = L_y(T)dP_0^y$ the process Z has the given propensities a_j . An analogous Girsanov weight $\tilde{L}_y(T)$ needs to be applied to the paths of the filter process \tilde{Z} . Then, it can be shown that for $0 \leq t \leq T$

$$P^y\{Z_0 = z_0, R(t) = k | Y(T) = y\} = \frac{\mathbb{E}_0^y[1_{\{(z_0, k)\}}(\tilde{Z}_0, K) W \tilde{L}_y(T)]}{\mathbb{E}_0^y[W \tilde{L}_y(T)]}.$$

In fact, in Section 4 we show slightly stronger results. The targeting algorithm is summarized in Algorithm 1.

Remark 1 *We note that, we have exchanged the indicator weight discussed in the previous section with what we shall call the Poisson weight (also denoted by W). This is the price to be paid for targeting (that is ensuring that \tilde{Z} lands in a state consistent with the observation). However, the Poisson weight has lower variance than the indicator weight. See Appendix A, where we obtain a lower bound for the effective sample fraction (ESF) ($\mathbb{E}_0^2[W]/\mathbb{E}_0[W^2]$) of the Poisson weight in terms of the ESF of the indicator weight.*

3.6 Multiple Observation Snapshots

In Section 3.5 we sketched a (weighted) Monte Carlo algorithm for the *targeting problem*. That is, to compute $\pi(t, z)$ for $0 \leq t \leq T$ in the context of a single

snapshot observation at time $T > 0$ given the observation $Y(T) = y$ and initial distribution μ for Z_0 . This algorithm produces as output an iid sample of size N_s

Suppose we have multiple snapshots: $Y(t_l) = y_l$ for $l = 1, \dots, T_s$ with $0 \leq t_1 < t_2 < \dots < t_{T_s}$. We repeatedly apply the targeting algorithm between successive snapshots t_{l-1} and t_l . We describe in more detail here.

If $t_1 > 0$, let us set $l = 1$, let $t_{l-1} = t_0 = 0$. Otherwise $t_1 = 0$, so we set $l = 2$. We apply the targeting algorithm for the span $[t_{l-1}, t_l]$. We take the initial probability measure μ as the known probability distribution of Z_0 . If $t_1 = 0$, the known observation $Y(0) = y_0$ is incorporated into the knowledge of μ .

For a general l , the algorithm results in weighted sample $(Z^{(i)}(t_l), W^{(i)}(t_l))$ which gives us an empirical measure $\hat{\mu}_l$ defined by

$$\hat{\mu}_l = \frac{\sum_{j=1}^{N_s} W^{(j)}(t_l) \delta_{\tilde{Z}^{(j)}(t_l)}}{\sum_{j=1}^{N_s} W^{(j)}(t_l)}.$$

Here δ_z is the Dirac measure concentrated at z . We use $\hat{\mu}_l$ as an approximation of the conditional probability distribution of $Z(t_l)$ given the observations $\{Y(t_1) = y_1, \dots, Y(t_l) = y_l\}$. We proceed to the next interval $[t_l, t_{l+1}]$ where we ideally need the knowledge of the true conditional distribution of $Z(t_l)$ given the observations $Y(t_1) = y_1, \dots, Y(t_l) = y_l$. However, we simply use $\hat{\mu}_l$ in its place as an approximation.

Algorithm 1 Targeting Algorithm

- 1: **Input:** Initial distribution $\mu(z_0)$, T , filter sample size N_s , observation y .
 - 2: Make a judicious choice of intensity function $\lambda(t, \tilde{z}_0, y)$ that is strictly positive.
 - 3: **for** $i = 1$ to N_s **do**
 - 4: Generate a sample from μ_0 resulting in $\tilde{Z}_0^{(i)} = (U_0^{(i)}, V_0^{(i)})$ and an m -variate Poisson random variable $K'^{(i)}$ with mean $\int_0^T \lambda(t, \tilde{Z}_0, y) dt$. Set $K''^{(i)} = G(y - V_0^{(i)}, K'^{(i)})$. If $K''^{(i)} \in \mathbb{Z}_+^{m_2}$ accept and go to next step. Otherwise repeat.
 - 5: Set $W^{(i)} = \rho(K''^{(i)}, \tilde{Z}_0^{(i)}, y)$.
 - 6: Generate interpolated trajectories for $\tilde{Z}^{(i)}(t)$ and $I^{(i)}(t)$ for $t \in [0, T]$.
 - 7: Compute the likelihood ratio $\tilde{L}^{(i)}(T)$.
 - 8: Compute the path functional $C^{(i)} = \phi(Z^{(i)})$.
 - 9: **end for**
 - 10: **Output:** Weighted sample $(C^{(i)}, W^{(i)} \tilde{L}^{(i)}(T))$ for $i = 1, \dots, N_s$.
-

Algorithm 2 Poisson Interpolation Algorithm

- 1: **Input:** Initial state z_0 , final time T , total reaction count $k \in \mathbb{Z}_+^m$, intensity matrix $\bar{\lambda} \in \mathbb{R}_+^{m \times N}$ where $\bar{\lambda}_{ij} = \lambda_i(t)$ for $t \in [t_{j-1}, t_j]$.
 - 2: Initialize $t \leftarrow 0$, $I(t) \leftarrow 0$, $\tilde{Z}(0) \leftarrow z_0$.
 - 3: Compute $\mu_i = \sum_{j=1}^N \bar{\lambda}_{i,j} \Delta t$ for $i = 1, \dots, m$.
 - 4: **for** $j = 1 : N$ **do**
 - 5: Generate reaction count $r_{i,j}$ in $[t_{j-1}, t_j) \sim \text{Bino}(k_i, \frac{\bar{\lambda}_{i,j} \Delta t}{\mu_i})$ for $i = 1, \dots, m$.
 - 6: Compute $r_j = \sum_{i=1}^m r_{i,j}$
 - 7: Generate t_l^i uniformly distributed in $[t_{j-1}, t_j)$ for $i = 1, \dots, m$, $l = 1, \dots, r_{i,j}$.
 - 8: Sort t_l^i so that $t_{(1)} < t_{(2)} \dots < t_{(r_j)}$, where $t_{(k)}$ is the k th smallest number among (t_l^i) , and record the corresponding (t_l^i, i, k) for $k = 1, \dots, r_j$.
 - 9: **for** $k = 1 : r_j$ **do**
 - 10: $t \leftarrow t_{(k)} = t_l^i$, $I_i(t) \leftarrow I_i(t-) + 1$, $\tilde{Z}(t) \leftarrow \tilde{Z}(t-) + v_i$.
 - 11: **end for**
 - 12: $\mu_i \leftarrow \mu_i - \bar{\lambda}_{i,j} \Delta t$, $k_i \leftarrow k_i - r_{i,j}$ for $i = 1, \dots, m$
 - 13: **end for**
 - 14: **Output:** $I(t)$ and $\tilde{Z}(t)$ for $t \in [0, T]$
-

4 Rigorous Treatment of the Targeting Algorithm

The targeting problem for a reaction network is stated as follows. We are given the following information about the lab process Z .

1. An observed set of events O_p from the past (before time $t = 0$ which is referred to as “present”) with $P(O_p) > 0$.
2. The conditional probability μ of the lab process Z at time $t = 0$. Thus, for $z_0 \in \mathbb{Z}_+^n$

$$\mu(z_0) = P\{Z_0 = z_0 \mid O_p\}.$$

3. Future time $T > 0$ and future observation $Y(T) = y$ where $y \in \mathbb{Z}_+^{n_2}$.

The goal is to compute the conditional distribution $\pi(t, z)$ where for $t \in [0, T]$ and $z \in \mathbb{Z}_+^n$

$$\pi(t, z) = P\{Z(T) = z \mid O_p, Y(T) = y\}.$$

The targeting algorithm consists of using an alternate probability measure P_0 such that under P_0 and under conditioning on O_p and $\{Z_0 = z_0\}$, the lab process

R for reaction counts is an inhomogeneous Poisson with a nonvanishing intensity $\lambda(t, z_0, y) > 0$ on $[0, T]$. Then the filter processes \tilde{Z} and I (proxies for Z and R) are created such that, under P_0 and when conditioned on O_p and $\{\tilde{Z}_0 = z_0\}$, I is an inhomogeneous Poisson with a nonvanishing intensity $\lambda(t, z_0, y) > 0$ on $[0, T]$.

We remark that it is conceptually easier to think of the role of the future observation y as fixed. Thus, for a given targeting problem, y is fixed and the probability measure P_0 depends on y , thus $P_0 = P_0^y$. We shall suppress the superscript y . The desired probability measure $P = P^y$ under which R has the correct intensity is accomplished by a Girsanov change of measure: $dP = L(T)dP_0$. Since the filter process \tilde{Z} will be used for the estimation, the probability measure to use shall $d\tilde{P} = \tilde{L}(T)dP_0$.

4.1 The Lab Process

We start with $(\Omega, \mathcal{F}, P_0)$, a probability space. Let $T > 0$, $X_0 : \Omega \rightarrow \mathbb{Z}_+^{n_1}$ and $Y_0 : \Omega \rightarrow \mathbb{Z}_+^{n_2}$ be random variables. Define $Z_0 = (X_0, Y_0)$. Let $v \in \mathbb{Z}^{n \times m}$.

Let $R : [0, T] \times \Omega \rightarrow \mathbb{Z}_+^m$ be an m -variate counting process which is broken into $R = (R', R'')$ where R' and R'' are m_1 -variate and m_2 -variate respectively. This means for $i = 1, \dots, m$, R_i is a *cadlag* process which is non-decreasing, takes values in \mathbb{Z}_+ and $R_i(0) = 0$.

Define processes X, Y by

$$\begin{aligned} X(t) &= X_0 + v'R(t), \\ Y(t) &= Y_0 + v''R(t). \end{aligned}$$

The \mathbb{Z}^n valued process Z is defined by $Z(t) = (X(t), Y(t))$. Thus we may write

$$Z(t) = Z_0 + vR(t).$$

We let $\mathcal{Y}_p \subset \mathcal{F}$ denote a sub σ -algebra of observation events that “happened prior” to time $t = 0$ and we assume $O_p \in \mathcal{Y}_p$ is a set of “past observations” with $P_0(O_p) > 0$. We define

$$\mathcal{F}_t = \mathcal{Y}_p \bigvee \sigma(Z(s), R(s) \mid 0 \leq s \leq t).$$

Recall that we may assume that with suitable splitting of the reactions $R(t) = (R'(t), R''(t))$, we may write v'' as

$$v'' = [A \ B] \tag{10}$$

where B is invertible. This yields

$$Y(T) = Y_0 + AR'(T) + BR''(T)$$

and hence

$$R''(T) = B^{-1}(Y(T) - Y_0 - AR'(T)).$$

For notational convenience we define $G : \mathbb{Z}^{n_2} \times \mathbb{Z}_+^{m_1} \rightarrow \mathbb{Z}_+^{m_2}$ by

$$G(\Delta y, k') = B^{-1}(\Delta y - Ak')$$

for $\Delta y \in \mathbb{Z}^{n_2}$ and $k' \in \mathbb{Z}_+^{m_1}$. For $\Delta y \in \mathbb{Z}^{n_2}$ we also define the subset $S_{\Delta y} \subset \mathbb{Z}_+^{m_1}$ by

$$S_{\Delta y} = \{k' \in \mathbb{Z}_+^{m_1} \mid B^{-1}(\Delta y - Ak') \in \mathbb{Z}_+^{m_2}\}.$$

For $\Delta y \in \mathbb{Z}^{n_2}$ and $k = (k', k'') \in \mathbb{Z}_+^{m_1} \times \mathbb{Z}_+^{m_2}$ the following equivalence holds:

$$\Delta y = v''k \iff k' \in S_{\Delta y} \text{ and } k'' = G(\Delta y, k').$$

Since $Y(T) = Y_0 + v''R(T)$, we note that for $y \in \mathbb{Z}^{n_2}$

$$\{Y(T) - Y_0 = y\} \subset \{R'(T) \in S_y\}.$$

We are given $\lambda : [0, T] \times \mathbb{Z}_+^n \times \mathbb{Z}_+^{n_2} \rightarrow (0, \infty)^m$, a positive function. We are also given μ , be a probability measure on \mathbb{Z}_+^n .

We suppose that the following hold under P_0 :

1. Conditioned on O_p , $Z_0 = (X_0, Y_0)$ has the distribution μ .
2. Conditioned on O_p and on $Z_0 = z_0 = (x_0, y_0)$, R is an m -variate inhomogeneous Poisson process with intensity $\lambda(t, z_0, y)$.

We also split λ into $\lambda = (\lambda', \lambda'')$ where λ' is \mathbb{R}^{m_1} valued and λ'' is \mathbb{R}^{m_2} valued.

Later on we shall consider another probability measure P on (Ω, \mathcal{F}) under which R is an m -variate counting process with \mathcal{F}_t -intensity $a(Z(t-))$. Thus, under P , the processes R and Z represent the reaction network. Since λ is strictly positive, P will be absolutely continuous with respect to P_0 .

Remark 2 *When we refer to an m -variate Poisson process or an m -variate Poisson random variable with an m -vector valued intensity or mean, we assume that the components are independent and have intensity or mean equal to the corresponding component.*

4.2 Construction of filter random variables $\tilde{Z}_0 = (U_0, V_0)$, K and W

On $(\Omega, \mathcal{F}, P_0)$ we create a sequence $(U_{0,i}, V_{0,i}, K'_i)$ which is $\mathbb{Z}_+^{n_1} \times \mathbb{Z}_+^{n_2} \times \mathbb{Z}_+^{m_1}$ with the following properties (under P_0).

1. Conditioned on O_p , (X_0, Y_0) and $(U_{0,i}, V_{0,i})$ (for $i \in \mathbb{N}$) are all independent, and have the same distribution μ .
2. Conditioned on O_p , K'_i for $i \in \mathbb{N}$ is an iid sequence that is independent of X_0, Y_0 and R .
3. For each i , when conditioned on O_p and on $(U_{0,i}, V_{0,i}) = (x_0, y_0)$, K'_i is an m_1 -variate Poisson random variable with mean $\int_0^T \lambda'(t, x_0, y_0, y) dt$. Thus for $k' \in \mathbb{Z}_+^{m_1}$ and $(x_0, y_0) \in \mathbb{Z}_+^n$

$$P_0\{K'_i = k' \mid U_{i,0} = x_0, V_{i,0} = y_0, O_p\} = P_0\{R'(T) = k' \mid X_0 = x_0, Y_0 = y_0, O_p\}.$$

Define $\rho : \mathbb{Z}_+^{m_2} \times \mathbb{Z}_+^{n_1} \times \mathbb{Z}_+^{n_2} \times \mathbb{Z}_+^{n_2} \rightarrow \mathbb{R}$ as follows. For each x_0, y_0, y , $\rho(\cdot, x_0, y_0, y)$ is the p.m.f. of an m_2 -variate Poisson random variable with mean $\int_0^T \lambda''(t, x_0, y_0, y) dt \in \mathbb{R}^{m_2}$.

Define the random variables N, K', K'', U_0, V_0 and W as follows. Define N by

$$N = \min\{i \in \mathbb{N} \mid K'_i \in S_{y-V_{0,i}}\}.$$

Throughout this section we make the following (sensible) assumption.

Assumption 1 $P\{Y(T) = y\} > 0$.

Since P is absolutely continuous with respect to P_0 , it follows that $P_0\{Y(T) = y\} > 0$ and hence there exists a y_0 with $P_0\{Y_0 = y_0\} > 0$ such that S_{y-y_0} is nonempty. Therefore

$$P_0\{R'(T) \in S_{y-y_0}\} > 0.$$

Hence, it follows that

$$\begin{aligned} P_0\{R'_i \in S_{y-V_{0,i}}\} &= \sum_{\tilde{y}_0} P_0\{R'_{0,i}(T) \in S_{y-\tilde{y}_0} \mid V_0 = \tilde{y}_0\} P_0\{V_0 = \tilde{y}_0\} \\ &> P_0\{R'(T) \in S_{y-y_0} \mid Y_0 = y_0\} P_0\{Y_0 = y_0\} > 0 \end{aligned}$$

and thus

$$P_0\{N < \infty\} = 1.$$

Thus under Assumption 1, N is finite and K' , K'' and W below are well defined:
 $K' = R'_N$, $K'' = G(y - V_{0,N}, K')$, $U_0 = U_{0,N}$, $V_0 = V_{0,N}$, $W = \rho(K'', U_0, V_0, y)$. Let
 $K = (K', K'')$. We let $\tilde{Z}_0 = (U_0, V_0)$.

We shall denote the observed event

$$O_y = \{Y(T) = y\},$$

and define $O = O_y \cap O_p$ (past and future observations). We note that if there is an observation $Y_0 = y_0$ at $t = 0$ then it is automatically captured in this analysis by taking the $\sum_{x_0} \mu(x_0, y_0) = 1$ and $\mu(x_0, \tilde{y}_0) = 0$ for $\tilde{y}_0 \neq y_0$.

Lemma 1 For $i \in \mathbb{N}$, $z_0 = (x_0, y_0) \in \mathbb{Z}_+^n$ and $k' \in \mathbb{Z}_+^{m_1}$

$$P_0\{U_{0,i} = x_0, V_{0,i} = y_0, K'_i = k' \mid N = i, O\} = 1_{S_{y-y_0}}(k') \frac{P_0\{X_0 = x_0, Y_0 = y_0, R'(T) = k' \mid O_p\}}{P_0\{R'(T) \in S_{y-y_0} \mid O_p\}}.$$

Proof If $k' \notin S_{y-y_0}$ then both sides of the equation in the lemma are zero. Suppose $k' \in S_{y-y_0}$. Then we have that

$$\begin{aligned} & P_0\{U_{0,i} = x_0, V_{0,i} = y_0, K'_i = k' \mid N = i, O\} \\ &= P_0\{U_{0,i} = x_0, V_{0,i} = y_0, K'_i = k' \mid K'_i \in S_{y-V_{0,i}}, K'_j \notin S_{y-V_{0,j}} \text{ for } j < i, O_y, O_p\} \\ &= \frac{P_0\{U_{0,i} = x_0, V_{0,i} = y_0, K'_i = k', K'_i \in S_{y-y_0}, K'_j \notin S_{y-V_{0,j}} \text{ for } j < i, O_y \mid O_p\}}{P_0\{K'_i \in S_{y-V_{0,i}}, K'_j \notin S_{y-V_{0,j}} \text{ for } j < i, O_y \mid O_p\}} \\ &= \frac{P_0\{U_{0,i} = x_0, V_{0,i} = y_0, K'_i = k' \mid O_p\}}{P_0\{K'_i \in S_{y-V_{0,i}} \mid O_p\}}. \end{aligned}$$

We note that the third equality follows since the other events in the numerator are independent of $Y(T)$ (and thus of O_y) when conditioned on O_p . Moreover, when conditioned on O_p , $(U_{0,i}, V_{0,i})$ is distributed like (X_0, Y_0) , and conditioned on $(U_{0,i}, V_{0,i}) = (x_0, y_0)$ and O_p , K'_i is distributed like $R'(T)$ when conditioned on $(X_0, Y_0) = (x_0, y_0)$ and O_p . Thus, we have

$$P_0\{U_{0,i} = x_0, V_{0,i} = y_0, K'_i = k' \mid O_p\} = P_0\{X_0 = x_0, Y_0 = y_0, R'(T) = k' \mid O_p\}.$$

Also, $P_0\{K'_i \in S_{y-V_{0,i}} \mid O_p\} = P_0\{R'(T) \in S_{y-y_0} \mid O_p\}$. These observations clinch the result. \blacksquare

Proposition 1 For $z_0 = (x_0, y_0) \in \mathbb{Z}^n, k' \in \mathbb{Z}_+^{m_1}$ we have that

$$P_0\{\tilde{Z}_0 = z_0, K' = k' \mid O\} = 1_{S_{y-y_0}}(k') \frac{P_0\{Z_0 = z_0, R'(T) = k' \mid O_p\}}{P_0\{R'(T) \in S_{y-y_0} \mid O_p\}}.$$

Proof We may write

$$P_0\{U_0 = x_0, V_0 = y_0, K' = k' \mid O\} = \sum_{i \in \mathbb{N}} P_0\{U_{0,i} = x_0, V_{0,i} = y_0, K' = k' \mid N = i, O\} P_0\{N = i \mid O\}.$$

Applying Lemma 1 yields the result. ■

Proposition 2 For $z_0 = (x_0, y_0) \in \mathbb{Z}^n, k \in \mathbb{Z}_+^m$ we have that

$$P_0\{Z_0 = z_0, R(T) = k \mid O\} = \frac{\mathbb{E}_0[1_{\{(z_0, k)\}}(\tilde{Z}_0, K) W \mid O]}{\mathbb{E}_0[W \mid O]}.$$

Proof Let $k = (k', k'')$ where $k' \in \mathbb{Z}_+^{m_1}, k'' \in \mathbb{Z}_+^{m_2}$. If $y \neq y_0 + v''k$ then both sides of the equation in Proposition 2 are zero.

Suppose $y = y_0 + v''k$. Then $k'' = G(y - y_0, k')$, $k' \in S_{y-y_0}$ and the event equality $\{K = k\} = \{K' = k'\}$ holds. Hence

$$\mathbb{E}_0[1_{\{(z_0, k)\}}(\tilde{Z}_0, K) W \mid O] = \frac{\mathbb{E}_0[W; \tilde{Z}_0 = z_0, K' = k', O]}{P_0\{O\}}.$$

We note that on $\{K' = k', \tilde{Z}_0 = z_0\}$

$$W = \rho(k'', x_0, y_0, y) = P_0\{R''(T) = k'' \mid Z_0 = z_0\} = P_0\{R'(T) = k' \mid Z_0 = z_0, O_p\}.$$

Hence

$$\begin{aligned} & \mathbb{E}_0[1_{\{(z_0, k)\}}(\tilde{Z}_0, K) W \mid O] \\ &= P_0\{R''(T) = k'' \mid Z_0 = z_0\} P_0\{\tilde{Z}_0 = z_0, K' = k' \mid O\} \\ &= P_0\{R''(T) = k'' \mid Z_0 = z_0\} \frac{P_0\{Z_0 = z_0, R'(T) = k' \mid O_p\}}{P_0\{R'(T) \in S_{y-y_0} \mid O_p\}} \\ &= P_0\{R''(T) = k'' \mid Z_0 = z_0\} \frac{P_0\{R'(T) = k' \mid Z_0 = z_0\} P_0\{Z_0 = z_0 \mid O_p\}}{P_0\{R'(T) \in S_{y-y_0} \mid O_p\}} \\ &= \frac{P_0\{R'(T) = k', R''(T) = k'' \mid Z_0 = z_0\} P_0\{Z_0 = z_0 \mid O_p\}}{P_0\{R'(T) \in S_{y-y_0} \mid O_p\}} \\ &= \frac{P_0\{R(T) = k, Z_0 = z_0 \mid O_p\}}{P_0\{R'(T) \in S_{y-y_0} \mid O_p\}} \end{aligned}$$

where the second equality follows from Proposition 1 (noting that $k' \in S_{y-y_0}$) and the fourth equality follows from the conditional independence of $R'(T)$ and

$R''(T)$ given Z_0 . We have also used the independence of $(R'(T), R''(T))$ from O_p when conditioned on $Z_0 = z_0$.

Thus for all $k \in \mathbb{Z}_+^m$ we may write

$$\mathbb{E}_0[1_{\{(z_0, k)\}}(\tilde{Z}_0, K) W | O] = 1_{\{y=y_0\}}(v''k) \frac{P_0\{Z_0 = z_0, R(T) = k | O_p\}}{P_0\{R'(T) \in S_{y-y_0} | O_p\}}.$$

Since

$$1_{\{y=y_0\}}(v''k) 1_{\{z_0\}}(Z_0) 1_{\{k\}}(K) = 1_{\{y\}}(Y(T)) 1_{\{z_0\}}(Z_0) 1_{\{k\}}(K),$$

it follows that

$$\mathbb{E}_0[1_{\{(z_0, k)\}}(\tilde{Z}_0, K) W | O] = \frac{P_0\{Z_0 = z_0, R(T) = k, Y(T) = y | O_p\}}{P_0\{R'(T) \in S_{y-y_0} | O_p\}}. \quad (11)$$

Summing over (z_0, k) we obtain

$$\mathbb{E}_0[W | O] = \frac{P_0\{Y(T) = y | O_p\}}{P_0\{R'(T) \in S_{y-y_0} | O_p\}}. \quad (12)$$

Dividing (11) by (12) we obtain

$$\begin{aligned} \frac{\mathbb{E}_0[1_{\{(z_0, k)\}}(\tilde{Z}_0, K) W | O]}{\mathbb{E}_0[W | O]} &= \frac{P_0\{Z_0 = z_0, R(T) = k, O_y | O_p\}}{P_0\{O_y | O_p\}} \\ &= P_0\{Z_0 = z_0, R(T) = k | O\}. \end{aligned}$$

■

4.3 Construction of the proxy processes $I, \tilde{Z} = (U, V)$

Having constructed $\tilde{Z}_0 = (U_0, V_0)$ and K , we define the process $I : [0, T] \times \Omega \rightarrow \mathbb{Z}_+^m$ as follows. Conditioned on $\tilde{Z}_0 = z_0$, we let $\eta_i(t) = \int_0^t \lambda(s, z_0, y) ds$ for $t \in [0, T]$. Note that η_i are strictly increasing functions. Let ξ_l^i for $i = 1, 2, \dots, m$ and $l = 1, 2, \dots$ be an i.i.d. collection of random variables uniformly distributed on $[0, \eta_i(T)]$ and independent of O_p, Z_0, R, \tilde{Z}_0 and K . Let $t_l^i = \eta_i^{-1}(\xi_l^i)$. Then we set

$$I_i(t) = \sum_{l=1}^{K_i} 1_{[t_l^i, \infty)}(t).$$

The process I is a proxy for R and by construction $I(T) = K$. We define $\tilde{Z} = (U, V)$ by

$$\tilde{Z}(t) = \tilde{Z}_0 + \nu I(t).$$

Let us denote the restriction of R to $[0, T]$ also by R . We note that the conditional law of R given $R(T) = k, Z_0 = z_0, O$ is the same as the conditional law of I given $I(T) = K = k, \tilde{Z}_0 = z_0, O$, since both will be “inhomogeneous Poisson bridge” processes with intensity $\lambda(t, z_0, y)$ under the said conditioning. We summarize this as the following proposition.

Proposition 3 *Let $\phi : D_{\mathbb{R}^m}[0, T] \rightarrow [0, \infty)$ be a nonnegative measurable map. Then for $k \in \mathbb{Z}_+^m, z_0 \in \mathbb{Z}^n$, we have that*

$$\mathbb{E}_0[\phi(R) | R(T) = k, Z_0 = z_0, O] = \mathbb{E}_0[\phi(I) | I(T) = k, \tilde{Z}_0 = z_0, O].$$

Proposition 4 *Let $\phi : D_{\mathbb{R}^m}[0, T] \times \mathbb{Z}^n \rightarrow [0, \infty)$ be a nonnegative measurable map. Then*

$$\mathbb{E}_0[\phi(R, Z_0) | O] = \frac{\mathbb{E}_0[\phi(I, \tilde{Z}_0) W | O]}{\mathbb{E}_0[W | O]}.$$

Proof

$$\begin{aligned} & \mathbb{E}_0[\phi(R, Z_0) | O] \\ &= \sum_{z_0 \in \mathbb{Z}^n} \sum_{k \in \mathbb{Z}_+^m} \mathbb{E}_0[\phi(R, z_0) | R(T) = k, Z_0 = z_0, O] P_0\{R(T) = k, Z_0 = z_0 | O\} \\ &= \sum_{k, z_0} \mathbb{E}_0[\phi(I, z_0) | K = k, \tilde{Z}_0 = z_0, O] \frac{\mathbb{E}_0[1_{\{(z_0, k)\}}(\tilde{Z}_0, K) W | O]}{\mathbb{E}_0[W | O]}, \end{aligned}$$

where in the second equality we have used Propositions 2 and 3. Now

$$\begin{aligned} & \frac{\mathbb{E}_0[1_{\{(z_0, k)\}}(\tilde{Z}_0, K) W | O]}{\mathbb{E}_0[W | O]} = \frac{\mathbb{E}_0[W; K = k, \tilde{Z}_0 = z_0, O]}{\mathbb{E}_0[W; O]} \\ &= \frac{\mathbb{E}_0[W | K = k, \tilde{Z}_0 = z_0, O] P_0\{K = k, \tilde{Z}_0 = z_0, | O\}}{\mathbb{E}_0[W | O]}. \end{aligned}$$

Since the process I and W are independent when conditioned on $I(T) = k, \tilde{Z}_0 = z_0, O$,

$$\begin{aligned} & \mathbb{E}_0[\phi(I, z_0) | K = k, \tilde{Z}_0 = z_0, O] \mathbb{E}_0[W | K = k, \tilde{Z}_0 = z_0, O] \\ &= \mathbb{E}_0[\phi(I, z_0) W | K = k, \tilde{Z}_0 = z_0, O]. \end{aligned}$$

Hence

$$\begin{aligned}
& \mathbb{E}_0[\phi(R, Z_0) | O] \\
&= \frac{\sum_{k, z_0} \mathbb{E}_0[\phi(I, z_0) W | K = k, \tilde{Z}_0 = z_0, O] P_0\{K = k, \tilde{Z}_0 = z_0 | O\}}{E_0[W | O]} \\
&= \frac{\mathbb{E}_0[\phi(I, \tilde{Z}_0) W | O]}{E_0[W | O]}.
\end{aligned}$$

■

4.4 Girsanov change of measure

Following [4], we define the processes $L_j, \tilde{L}_j : [0, T] \times \Omega \rightarrow \mathbb{R}$ for $j = 1, \dots, m$ by

$$\begin{aligned}
L_j(t) &= \left(\prod_{i \geq 1} \frac{a_j(Z(T_i^j -))}{\lambda_j(T_i^j -, Z_0, y)} 1_{\{T_i^j \leq t\}} \right) \exp \left(\int_0^t (\lambda_j(s, Z_0, y) - a_j(Z(s))) ds \right), \\
\tilde{L}_j(t) &= \left(\prod_{i \geq 1} \frac{a_j(\tilde{Z}(\tilde{T}_i^j -))}{\lambda_j(\tilde{T}_i^j -, \tilde{Z}_0, y)} 1_{\{\tilde{T}_i^j \leq t\}} \right) \exp \left(\int_0^t (\lambda_j(s, \tilde{Z}_0, y) - a_j(\tilde{Z}(s))) ds \right),
\end{aligned} \tag{13}$$

where T_i^j and \tilde{T}_i^j for $i = 1, 2, \dots$ are the i th jump time of R_j and I_j respectively. Define the processes L, \tilde{L} by

$$L(t) = L_1(t) \dots L_m(t), \quad \tilde{L}(t) = \tilde{L}_1(t) \dots \tilde{L}_m(t). \tag{14}$$

We note that for each $t \in [0, T]$, we can write $L(t)$ and $\tilde{L}(t)$ as

$$L(t) = \psi_t(R, Z_0) \text{ and } \tilde{L}(t) = \psi_t(I, \tilde{Z}_0),$$

where $\psi_t : D_{\mathbb{R}^m}[0, T] \times \mathbb{Z}^n \rightarrow \mathbb{R}$ is a measurable map.

Under certain integrability conditions we may assume that $\mathbb{E}_0(L(T)) = 1$. Let $\mathcal{F}_t = \sigma(Z_0, R(s); 0 \leq s \leq t)$. It follows that L is a $(\{\mathcal{F}_t\}, P_0)$ martingale and $\mathbb{E}_0(L(t)) = 1$ for all $t \in [0, T]$. We define the probability measure P on (Ω, \mathcal{F}) by $P = L(T) P_0$. It follows that the P intensity of R is $a(Z_0 + \nu R(t-)) = a(Z(t-))$, which is the desired intensity [4].

Theorem 1 *Let $\phi : D_{\mathbb{R}^m}[0, T] \times \mathbb{Z}_+^n \rightarrow \mathbb{R}$ be nonnegative measurable. Then*

$$\mathbb{E}[\phi(R, Z_0) | O] = \frac{\mathbb{E}_0[\phi(I, \tilde{Z}_0) \tilde{L}(T) W | O]}{\mathbb{E}_0[\tilde{L}(T) W | O]}$$

Proof

$$\begin{aligned}\mathbb{E}[\phi(R, Z_0) | O] &= \frac{\mathbb{E}[\phi(R, Z_0); O]}{P\{O\}} \\ &= \frac{\mathbb{E}_0[\phi(R, Z_0)L(T); O]}{\mathbb{E}_0[L(T); O]} = \frac{\mathbb{E}_0[\phi(R, Z_0)L(T) | O]}{\mathbb{E}_0[L(T) | O]}.\end{aligned}$$

Since on O , $L(T) = \psi_T(R, Z_0)$ and $\tilde{L}(T) = \psi_T(I, Z_0)$, by Proposition 4, it follows that

$$\mathbb{E}_0[\phi(R, Z_0)L(T) | O] = \frac{\mathbb{E}_0[\phi(I, \tilde{Z}_0)\tilde{L}(T)W | O]}{\mathbb{E}_0[W | O]},$$

and

$$\mathbb{E}_0[L(T) | O] = \frac{\mathbb{E}_0[\tilde{L}(T)W | O]}{\mathbb{E}_0[W | O]}.$$

From the last three equations the result follows. ■

Corollary 1 *Let $z_0 \in \mathbb{Z}_+^n, k \in \mathbb{Z}_+^m$. Then for $t \in [0, T]$*

$$\mathbb{E}[1_{\{(z_0, k)\}}(Z_0, R(t)) | O] = \frac{\mathbb{E}_0[1_{\{(z_0, k)\}}(\tilde{Z}_0, I(t))\tilde{L}(T)W | O]}{\mathbb{E}_0[\tilde{L}(T)W | O]}$$

Corollary 2 *For $z \in \mathbb{Z}_+^n$ and $t \in [0, T]$*

$$\mathbb{E}[1_{\{z\}}(Z(t)) | O] = \frac{\mathbb{E}_0[1_{\{z\}}(\tilde{Z}(t))\tilde{L}(T)W | O]}{\mathbb{E}_0[\tilde{L}(T)W | O]}$$

5 Examples

In this section, we briefly describe the reaction systems that we shall use as numerical examples. The first two are monomolecular and hence have linear propensities while the third example has a bimolecular reaction making it nonlinear. These examples are chosen so that it is possible to compute the conditional probability of interest either analytically or via deterministic computations such as ODE solvers. This way, we can verify the accuracy of our Monte Carlo methods. We describe how the conditional probabilities may be computed for each example. We also note that the exact conditional propensity may be computed if the transition probabilities $P\{Z(T) = z_T | Z(t) = z_t\}$ ($0 \leq t \leq T$) can be computed.

5.1 Example 1: Pure-death model

$$S \xrightarrow{c} \emptyset \quad (15)$$

with the propensity $a(x) = cx$. This is the simplest example where an analytically tractable conditional propensity in addition to conditional probabilities is available.

The conditional distribution

$$\pi(x_t | x_0, x_T) := P\{X(t) = x_t | X(0) = x_0, X(T) = x_T\}, \quad (16)$$

is given in terms of the transition probabilities as

$$\pi(x_t | x_0, x_T) = \frac{P\{X(T) = x_T | X(t) = x_t\} P\{X(t) = x_t | X(0) = x_0\}}{P\{X(T) = x_T | X(0) = x_0\}}. \quad (17)$$

where transition probability from t_1 to t_2 ($t_1 \leq t_2$) follows a binomial distribution with parameters $n = x_1$ and $p = e^{-c(t_2 - t_1)}$:

$$P\{X(t_2) = x_2 | X(t_1) = x_1\} = \binom{x_1}{x_2} e^{-c(t_2 - t_1)x_2} (1 - e^{-c(t_2 - t_1)})^{x_1 - x_2}. \quad (18)$$

Using equations (17) and (18), we can obtain that

$$\pi(x_t | x_0, x_T) = \binom{x_0 - x_T}{x_0 - x_t} \left(\frac{e^{-ct} - e^{-cT}}{1 - e^{-cT}} \right)^{x_t - x_T} \left(\frac{1 - e^{-ct}}{1 - e^{-cT}} \right)^{x_0 - x_t}. \quad (19)$$

Furthermore, we can also obtain the conditioned propensity as

$$\lambda(x_t | x_T) = a(x_t) \frac{p(x_T | X(t) = x_t - 1)}{p(x_T | X(t) = x_t)} = \frac{c(x_t - x_T)}{1 - e^{-c(T-t)}}. \quad (20)$$

We note that the conditional propensity depends explicitly on time t and hence the waiting times between reaction events according to the conditional propensity are not exponential. Yet, the waiting time has an analytically tractable form. Let us denote the wait time at time t for the next reaction to occur by τ and let

$$G(s) = P\{\tau > s | X(t) = x_t, X(T) = x_T\}. \quad (21)$$

Then it can be shown that

$$G(s) = \left(\frac{e^{-cs} - e^{-c(T-t)}}{1 - e^{-c(T-t)}} \right)^{x_t - x_T}. \quad (22)$$

For Monte Carlo simulation, a sample for τ can be obtained by

$$\tau = \min\{s : G(s) \leq 1 - u\}, \quad (23)$$

where u is a uniform random variable in $[0, 1]$.

5.2 Example 2: Reversible Reaction $S_1 \longleftrightarrow S_2$

We consider the system



Let $Z(t) = (\#S_1(t), \#S_2(t))$. Suppose the propensities take mass-action form $a_1(z) = c_1 z_1$ and $a_2(z) = c_2 z_2$. Denote the initial state $z_0 = (z_{01}, z_{02})$.

To find the evolution of the probability distribution, we consider an associated mono-molecular system [24] which is described by a random walk $M(t)$ between two states. This is a continuous-time Markov chain with the rate matrix

$$Q = \begin{pmatrix} -c_1 & c_1 \\ c_2 & -c_2 \end{pmatrix}. \quad (25)$$


We can obtain its transition probability matrix $P(t)$ by Kolmogorov's backward equation $P'(t) = QP(t)$, where $P_{i,j}(t) = \text{Prob}\{M(t) = j | M(0) = i\}$ for $i, j = 1, 2$. We note the conservation relation $Z_1(t) + Z_2(t) = z_{01} + z_{02}$. It follows that

$$\begin{aligned} & P\{Z_1(t) = z_t, Z_2(t) = z_{01} + z_{02} - z_t | Z_1(0) = z_{01}, Z_2(0) = z_{02}\} \\ &= \sum_{k=0}^{z_{01}+z_{02}} \text{Binopdf}(k; z_{01}, P_{11}(t)) \text{Binopdf}(z_t - z_{01} - k; z_{02}, P_{21}(t)). \end{aligned} \quad (26)$$

Using equations (26) and (17), we could obtain an analytical formula for the conditional distributions. We also note that since the transition probability can be computed exactly from (26), it is possible to compute the conditional propensities exactly. However, unlike in Example 1, here the wait times according to the conditional propensities are not analytically tractable. Nevertheless, one may use an approximate conditional propensity which is held constant in between jumps as in [16].

5.3 Example 3: $S_1 \longleftrightarrow S_2, \quad S_1 + S_2 \longleftrightarrow S_3$

We consider



Let $Z(t) = (\#S_1(t), \#S_2(t), \#S_3(t))$. Suppose we observe species S_3 , so $X(t) = (\#S_1(t), \#S_2(t))$ and $Y(t) = \#S_3(t)$. Suppose $a_1(z) = c_1 z_1$, $a_2(z) = c_2 z_2$, $a_3(z) = c_3 z_1 z_2$ and $a_4(z) = c_4 z_3$.

The conditional distribution $\pi(t, z)$ is easy to obtain if we take $t = T$. Let $Z_0 = (z_{10}, z_{20}, z_{30})$, due to the conservation relation $Z_1(t) + Z_2(t) + 2Z_3(t) = z_{10} + z_{20} + 2z_{30}$, the state space is finite and thus we can solve Kolmogorov's forward equation numerically via an ODE solver to obtain the unconditional p.m.f. $p(T, z)$. Given $Y(T) = \#S_3(T) = y$, and $p(T, z)$ for the joint distribution of (Z_1, Z_2, Z_3) , we can easily compute the conditional distribution $\pi(T, z)$.

6 Numerical experiments

We shall use the three examples described in Section 5 to investigate the effectiveness of the various Monte Carlo algorithms. We shall focus on the targeting problem. That is, we know the initial distribution μ_0 which we usually take to be Dirac and we have a partial state observation $Y(T) = y$ at time $T > 0$.

We use the mean total variation error (TVE) to characterize the error in the estimated conditional distribution $\hat{\pi}(t, z)$:

$$\text{TVE} = \mathbb{E} \left[\sum_z |\hat{\pi}(t, z) - \pi(t, z)| \mid Y(T) = y \right].$$

We also record the effective sample fraction (ESF) defined by the ratio of effective sample size N_e over the filter sample size N_s

$$\text{ESF} = \frac{N_e}{N_s} = \frac{1}{N_s} \frac{\left(\sum_{i=1}^{N_s} W^i \right)^2}{\sum_{i=1}^{N_s} (W^i)^2}.$$

The subsections are organized so that each subsection explores a certain choice made within our targeting algorithm or compares our targeting algorithm with other methods.

6.1 Choice of inhomogeneous Poisson intensities λ

Our targeting method involves choosing inhomogeneous Poisson intensities λ for $R(t)$. It is most convenient to interpolate process I when λ is piecewise constant in time. We evenly divide $[0, T]$ into N subintervals $[t_j, t_{j+1}]$ with $t_{j+1} - t_j = \Delta t = \frac{T}{N}$ and let $\lambda(t) = \lambda(t_j)$ for $t \in [t_j, t_{j+1})$. We use a matrix representation for this piecewise constant intensity $\bar{\lambda} \in \mathbb{R}_+^{m \times N}$, where $\bar{\lambda}_{i,j} = \lambda_i(t_j)$. Here we present two strategies for choosing the piecewise constant intensities that we used in our

numerical experiments. Since the efficiency of the Monte Carlo method depends on having a high effective sample size, we ideally want the weight $W\tilde{L}(T)$ to have low variance. Recall that W is the “Poisson weight” that appears in the early steps of the targeting method (see Algorithm 1) while $\tilde{L}(T)$ is the weight due to the Girsanov transformation.

The first idea is to focus only on the Girsanov weight $\tilde{L}(T)$ since this appears to be more variable in our numerical experiments. Intuitively, the closer $\lambda(t)$ is to the propensity $a(Z(t))$, the better. Since $\lambda(t)$ is state independent, a natural choice is $\mathbb{E}[a(Z(t))]$. If we know $\mathbb{E}[a(Z(t))]$ analytically, we can construct a left endpoint approximation $\bar{\lambda}_1(t) = \mathbb{E}[a(Z(t_j))]$ for $t \in [t_j, t_{j+1})$. If $\mathbb{E}[a(Z(t))]$ is not analytically tractable, we could use some reasonable approximation of it. Examples include using the solution of the corresponding (deterministic) reaction rate equations as an approximation or, alternatively, estimate $\mathbb{E}[a(Z(t))]$ via an independent forward Monte Carlo simulation.

The above idea doesn’t take the observation y into account and we note that the variability in the Poisson weight W is dependent on y . Thus an alternative choice for the piecewise constant Poisson intensity $\bar{\lambda}_2$ is to improve the probability of the observed value y . So we impose the condition that the expectation of the process (as predicted by the intensity $\bar{\lambda}_2$ lands precisely at the desired state: $v'' \int_0^T \bar{\lambda}_2(t) dt = y - y_0$, while keeping $\bar{\lambda}_2$ close to the average intensity $\bar{\lambda}_1$. To be specific, we solve the following constrained optimization problem to get intensity $\bar{\lambda}_2$:

$$\bar{\lambda}_2 = \operatorname{argmin}_{\lambda \in \mathbb{R}_+^{m \times N}} \|\lambda - \bar{\lambda}_1\|_F \quad \text{s.t. } v''r = \Delta y, \text{ where } r_i = \sum_{j=1}^N \lambda_{i,j} \Delta t. \quad (28)$$

Since we need the inhomogeneous Poisson intensity to be strictly positive, in order to avoid zero intensity, if the solution to (28) does not return strictly positive λ_{ij} we use the optimization problem (29) instead. The set S in (29) is defined as follows. Let $\underline{\lambda}_i = \min\{a_i(z) \mid a_i(z) > 0, z \in \mathbb{Z}_+^n\}$ for $i = 1, \dots, m$. Set $S = \{\lambda \in \mathbb{R}_+^{m \times N} \mid \lambda_{i,j} \geq \underline{\lambda}_i, \text{ for } j = 1, \dots, N\}$.

$$\bar{\lambda}_2 = \operatorname{argmin}_{\lambda \in S} \|\lambda - \bar{\lambda}_1\|_F \quad \text{s.t. } v''r = \Delta y, \text{ where } r_i = \sum_{j=1}^N \lambda_{i,j} \Delta t. \quad (29)$$

Here $\|\cdot\|_F$ stands for the Frobenius norm. In our numerical examples, we use MATLAB program `fmincon` to solve this optimization problem.

The difference between the two choices $\bar{\lambda}_1$ and $\bar{\lambda}_2$ are illustrated by Examples 2, 3 and 1. The naive approach is also provided for comparison. In all examples,

Method	Observation	TVE (95% interval)	ESF (W)	ESF (\tilde{L})	ESF	Runtime
Naive	$y = 4$	0.1188, [0.1104, 0.1272]	N/A	N/A	0.25	0.05
Targeting ($\bar{\lambda}_1$)	(common observation)	0.0722, [0.0673, 0.0771]	0.95	0.56	0.66	0.3
Targeting ($\bar{\lambda}_2$)	$P(Y(T) = 4) = 0.245$	0.0760, [0.0702, 0.0818]	0.95	0.51	0.61	0.3
CP (approx)		0.0834, [0.0783, 0.0886]	N/A	N/A	0.52	31.3
Naive	$y = 7$	0.3646, [0.3393, 0.3900]	N/A	N/A	0.03	0.05
Targeting ($\bar{\lambda}_1$)	(rare observation)	0.0940, [0.0880, 0.1001]	0.76	0.37	0.38	0.3
Targeting ($\bar{\lambda}_2$)	$P(Y(T) = 7) = 0.027$	0.0962, [0.0897, 0.1028]	0.98	0.21	0.37	0.3
CP (approx)		0.0907, [0.0851, 0.0962]	N/A	N/A	0.43	33.3

Table 1: Example 2. Reversible isomerization reaction $S_1 \longleftrightarrow S_2$. We took $z_0 = (10, 0)$, $c = (1, 1.5)$, $T = 1$, $t = 0.7$, $\Delta t = 0.1$, while taking a common observation as well as a rare observation. We used $N_r = 100$ trials, while each trial used filter sample size $N_s = 1000$.

we choose a “common observation” for y , meaning a value of y close to the mode of marginal distribution for $Y(T)$ and a “rare observation” for y , meaning a value of y that has a significantly smaller probability compared to the common observation.

In Example 2, $S_1 \longleftrightarrow S_2$, we picked $I_1(T)$ (corresponding to $S_1 \longrightarrow S_2$) as free reaction set $I_2 = I_1 - \Delta y$. we first took a smaller system (small total copy number): $Z(0) = (10, 0)$. We chose parameters $c = (1, 1.5)$ and also $t = 0.7$, $T = 1$, $\Delta t = 0.1$. The results are shown in Table 1. We also tried with a larger system with initial state $Z(0) = (100, 100)$. The results are shown Table 2. In Example 3, we took $X(0) = (20, 20, 20)$, $c = (0.5, 1, 0.1, 1)$, $t = 1$, $T = 1$. The results are shown in Table 3.

We first observe that in all cases the naive method performed poorly (both mean TVE and ESF) compared to the targeting methods. See Table 1, Table 2 and Table 3. The intensity choices $\bar{\lambda}_1$ and $\bar{\lambda}_2$ have similar performance in the common observation scenarios. In the rare observation scenarios, $\bar{\lambda}_2$ tends to have better performance. We also note that the optimization often gives a better effective sample fraction (ESF) for the Poisson weights, especially when the observation is rare. We also observe that by a large $\bar{\lambda}_1$ has similar or better ESF for the Girsanov weights when compared to $\bar{\lambda}_2$. This is expected since the choice of $\bar{\lambda}_1$ is solely focused on reducing the variance of the Girsanov weight. It is also notable that $\bar{\lambda}_2$ has better ESF for the Poisson weights especially when rare observations are concerned. This is not surprising since the choice of $\bar{\lambda}_2$ is based on a steering towards the observation y to improve the ESF of the Poisson weights. The overall

Method	Observation	TVE (95% interval)	ESF (W)	ESF (\tilde{L})	ESF	Runtime (seconds)
Naive	$y = 80$	0.1868, [0.1826, 0.1910]	N/A	N/A	0.06	19.50
Targeting ($\bar{\lambda}_1$)	(common observation)	0.0665, [0.0646, 0.0684]	0.86	0.51	0.43	21.10
Targeting ($\bar{\lambda}_2$)	($p = 0.0575$)	0.0626, [0.0607, 0.0646]	0.88	0.50	0.44	21.93
Naive	$y = 98$	0.6764, [0.6595, 0.6933]	N/A	N/A	0.004	19.03
Targeting ($\bar{\lambda}_1$)	(rare observation)	0.0928, [0.0902, 0.0954]	0.62	0.35	0.21	21.30
Targeting ($\bar{\lambda}_2$)	($p = 0.0043$)	0.0755, [0.0733, 0.0776]	0.87	0.35	0.31	21.95

Table 2: Example 2. Reversible isomerization reaction $S_1 \longleftrightarrow S_2$. We took $z_0 = (100, 100)$, $c = (1, 1.5)$, $T = 2$, $t = 0.7$, $\Delta t = 0.25$, while taking a common observation $y_T = 80$ as well as a rare observation $y_T = 98$. We used $N_r = 100$ trials, while each trial used filter sample size $N_s = 10,000$.

Method	Observation	TVE (95% interval)	ESF (W)	ESF (\tilde{L})	ESF	Runtime (seconds)
Naive	common observation	0.2146, [0.2117, 0.2174]	N/A	N/A	0.16	0.331
Targeting, $\bar{\lambda}_1$	$y = 24$	0.1695, [0.1669, 0.1722]	0.83	0.20	0.21	0.773
Targeting, $\bar{\lambda}_2$		0.1699, [0.1673, 0.1726]	0.84	0.20	0.21	0.770
Naive	rare observation	0.4347, [0.4286, 0.4408]	N/A	N/A	0.04	0.326
Targeting, $\bar{\lambda}_1$	$y = 20$	0.2184, [0.2152, 0.2217]	0.70	0.12	0.14	0.791
Targeting, $\bar{\lambda}_2$		0.2103, [0.2071, 0.2135]	0.84	0.14	0.15	0.794

Table 3: Example 3. $S_1 \longleftrightarrow S_2, S_1 + S_2 \longleftrightarrow S_3$. We took $X(0) = (20, 20, 20)$, $c = (0.5, 1, 0.1, 1)$, $t = T = 1$, $\Delta t = 0.1$. We used $N_r = 1,000$ trials, filter sample size $N_s = 1,000$.

ESF is also better for $\bar{\lambda}_2$ in the case of rare observations (with one exception), assuring that the constrained optimization is having the desired effect.

Finally, in Example 1, we took $x_0 = 1000$, $c = 4$, $T = 0.5$, $t = 0.2$, and time mesh $\Delta t = 0.02$. See Table 4. In this example, the free dimension $m_1 = 0$, while the determined dimension $m_2 = 1$, so the Poisson weight will always be the same for different samples making ESF equal to 1 for the Poisson weights. This explains why $\bar{\lambda}_1$ does better than $\bar{\lambda}_2$ because only the Girsanov weights matter. In terms of computational time (the Tables 1, 2, 3 and 4 do not include the computation time for the intensities $\bar{\lambda}_1$ or $\bar{\lambda}_2$ which don't grow with filter size N_s and are relatively small), the naive method is only modestly faster. In Example 2 with small molecular copy numbers (Table 1), the naive method is significantly faster because the total number of reaction events is small and hence the overhead in the binomial random numbers generated per subintervals of length Δt are nonnegligible for the targeting methods.

Method	Observation	Total Variation Error (95% interval)	Effective Sample Fraction (ESF)	Runtime (seconds)
Naive	$x_T = 368$ (common observation)	1.1353, [1.1077, 1.1630]	0.027	2.3
Targeting ($\bar{\lambda}_1$)		0.2037, [0.1995, 0.2080]	0.919	2.6
Targeting ($\bar{\lambda}_2$)		0.2072, [0.2029, 0.2116]	0.900	2.7
CP (approx)		0.3485, [0.3414, 0.3556]	0.322	2.4
CP (exact)		0.1957, [0.1920, 0.1994]	1.000	2.4
Naive	$x_T = 404$ (rare observation)	1.9036, [1.8934, 1.9139]	0.002	2.4
Targeting ($\bar{\lambda}_1$)		0.1979, [0.1942, 0.2016]	0.923	2.6
Targeting ($\bar{\lambda}_2$)		0.2297, [0.2248, 0.2347]	0.654	2.4
CP (approx)		0.3351, [0.3275, 0.3427]	0.324	2.1
CP (exact)		0.1909, [0.1872, 0.1947]	1.000	2.3

Table 4: Example 1. We took $x_0 = 1000$, $c = 2$, $T = 0.5$, $t = 0.2$, $\Delta t = 0.02$. We used $N_r = 100$ trials, while each trial has filter sample size $N_s = 1000$. We note that in our experiment for the rare observation case, the naive method has only 80% of the trials provided meaningful results, and the rest of 20% runs of the naive filter failed when none of the landed at the target state, resulting a zero denominator issue.

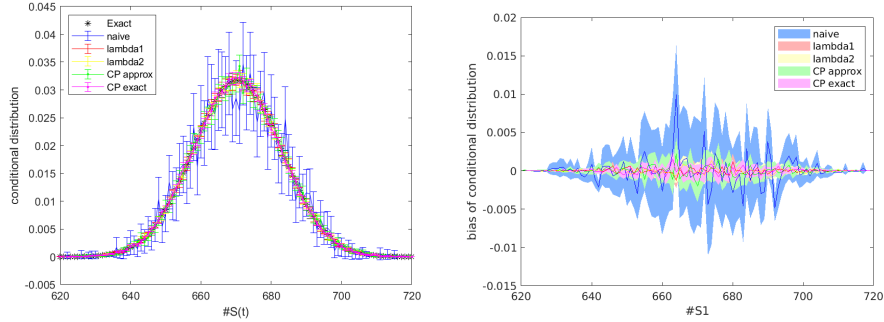


Figure 1: Example 1. Pure death example. Estimated conditional distribution compared with the actual conditional distribution. A common observation. $x_0 = 1000$, $c = 2$, $T = 0.5$, $t = 0.2$. The estimate came with 95% confident interval, which was estimated based on $N_r = 100$ trials, while each trial used filter sample size $N_s = 1000$.

6.2 Choice of free reaction channels

When we solve linear equations $y = Y_0 + v''I(T)$ to ensure $Y(T) = y$, we need to choose a set of reactions to be free, which determines the total reaction count of the rest of the reactions. There are different options. For instance, in Example 2, $S_1 \longleftrightarrow S_2$, to solve $y - Y_0 = I_1(T) - I_2(T)$, we could choose I_1 to be the free reaction, then set $I_2(T) = I_1(T) - y + Y_0$. Alternatively, if we choose I_2 to be free, then $I_1(T) = I_2(T) - y - Y_0$. We show the results in Table 5 and we see that there is a noticeable difference in the performances (we use $\hat{\lambda}_2$) of these two choices. This is a topic that may be worth exploring in the future. In the case of Example 2, this can be explained by some analysis which approximates Poissons by Gaussians. In the interest of space, we do not show it here.


6.3 Two-stage targeting algorithm

We introduce the *two-stage method* as another attempt to address the issue of growing variance of the weights. For estimating the system at time t ($0 < t < T$), especially when t is close to T , we could split the interval $[0, T]$ into two subintervals $[0, t_0]$ and $[t_0, T]$ where $0 \leq t_0 \leq t \leq T$ and run the (unconditional) forward simulation in time interval $[0, t_0]$ and then apply the targeting algorithm in the interval $[t_0, T]$. This way, we reduce the length of time interval where the weights grow disparately. We also explore the choice of t_0 for splitting the interval.

Method	Observation	TVE (95% interval)	ESF (W)	ESF (\tilde{L})	ESF
Targeting ($\bar{\lambda}_2$), I_1 free	common observation	0.1014, [0.0942, 0.1086]	0.77	0.22	0.31
Targeting ($\bar{\lambda}_2$), I_2 free	$P(Y(T) = 4) = 0.245$	0.0978, [0.0887, 0.1070]	0.95	0.32	0.39
Targeting ($\bar{\lambda}_2$), I_1 free	rare observation	0.0940, [0.0880, 0.1001]	0.76	0.38	0.38
Targeting ($\bar{\lambda}_2$), I_2 free	$P(Y(T) = 7) = 0.027$	0.0921, [0.0864, 0.0977]	0.98	0.24	0.37

Table 5: Example 2. Reversible Reaction $S_1 \longleftrightarrow S_2$. We took $z_0 = (10, 0)$, $c = (1, 1.5)$, $T = 1$, $t = 0.7$, $\Delta t = 0.1$, while taking a common observation as well as a rare observation. We used $N_r = 100$ trials, while each trial used filter sample size $N_s = 1000$.

We use Example 3



where species S_3 is observed, to illustrate the two-stage algorithm and compare it with just using a single stage.

We took $X(0) = (20, 20, 20)$, $c = (0.5, 1, 0.1, 1)$, $t = T = 1$, and applied the naive method and several variants of targeting algorithm to estimate the conditional distribution of $X(t) = X(T)$. In this example, we picked first three reactions as “free reactions” and the last one is determined by $I_4 = I_3 - \Delta y$.

For the single stage, i.e. targeting over the interval $[0, T]$, we took mesh size $\Delta t = 0.1$. Given the nonlinearity of the system, we used the deterministic system to determine $\lambda(t)$. That is, we solved the reaction rate equations (MATLAB ODE solver)

$$\begin{aligned} \frac{dz_1}{dt} &= -c_1 z_1 + c_2 z_1 - c_3 z_1 z_2 + c_4 z_3, & \frac{dz_2}{dt} &= c_1 z_1 - c_2 z_1 - c_3 z_1 z_2 + c_4 z_3, \\ \frac{dz_3}{dt} &= c_3 z_1 z_2 - c_4 z_3, \end{aligned}$$

to obtain (piecewise constant) $\bar{\lambda}_1$ and then $\bar{\lambda}_2$ as described earlier.

For the two-stage targeting algorithm, since the targeting is only applied in the second stage $[t_0, T]$ which is intended to be a small interval, we simply used a constant value for the intensity. We explored two possibilities for the choice of this constant vector intensity. One option we took was to use the sample mean of the propensity from the end of the first stage: $\bar{\lambda} = \frac{1}{N_s} \sum a(\tilde{Z}^{(i)}(t_0))$. The other option is to assign each individual sample i , a propensity that takes account of the state $\tilde{Z}^{(i)}(t_0)$. We chose $\lambda^{(i)} \in \mathbb{R}_+^4$ to be $\lambda^{(i)} = \max\{a(\tilde{Z}^{(i)}(t_0), a(1, 1, 1))\}$. The reason

we introduce $a(1, 1, 1)$ is that $a(\tilde{Z}^{(i)}(t_1))$ could possibly contain zero components, while our method requires strictly positive intensities, so we enforce a small lower bound corresponding to the situation where there is one molecule for each species present. Our choice $a(1, 1, 1)$ is usually a small quantity compared to $a(\tilde{Z}^{(i)}(t_0))$ while sufficiently safeguarding the strict positivity of intensity $\lambda(t)$. Numerical results are summarized in Table 6 for a common observation and Table 7 for a rare observation.

One-stage targeting algorithm estimates the conditional distribution more accurately than the naive method. For the two-stage targeting algorithm, there is an optimal choice of t_0 to split the forward evolution stage and the targeting stage. When the length of the second interval $[t_0, T]$ is large, the growing disparity of the Girsanov weights could be a problem. However, if t_0 is very close to the final time T , we may be faced with very disparate Poisson weights W . With the optimal choice of t_0 , the two-stage targeting appears to perform better than the single stage targeting. On the other hand, within the two-stage algorithm, it appears that the using a common intensity $\bar{\lambda}$ and individual propensities tend to give similar performance.

Method	TVE (95% interval)	ESF (W)	ESF (\bar{L})	ESF	Runtime (seconds)
Naive	0.2146, [0.2117, 0.2174]	N/A	N/A	0.16	0.331
One stage $\lambda_1, \Delta t = 0.1$	0.1695, [0.1669, 0.1722]	0.83	0.20	0.21	0.773
One stage $\lambda_2, \Delta t = 0.1$	0.1699, [0.1673, 0.1726]	0.84	0.20	0.21	0.770
Two stage, $\bar{\lambda}, t_0 = 0.5$	0.1375, [0.1346, 0.1404]	0.80	0.36	0.35	0.55
Two stage, $\bar{\lambda}, t_0 = 0.8$	0.1133, [0.1117, 0.1150]	0.74	0.66	0.54	0.57
Two stage, $\bar{\lambda}, t_0 = 0.9$	0.1113, [0.1098, 0.1128]	0.74	0.81	0.63	0.50
Two stage, $\bar{\lambda}, t_0 = 0.99$	0.1852, [0.1826, 0.1878]	0.39	0.94	0.39	0.49
Two stage, $\bar{\lambda}, t_0 = 0.999$	0.2555, [0.2519, 0.2590]	0.29	0.96	0.29	0.54
Two stage, $\lambda^{(i)}, t_0 = 0.5$	0.1616, [0.1576, 0.1655]	0.79	0.31	0.28	0.54
Two stage, $\lambda^{(i)}, t_0 = 0.8$	0.1045, [0.1031, 0.1059]	0.86	0.69	0.72	0.52
Two stage, $\lambda^{(i)}, t_0 = 0.9$	0.1025, [0.1010, 0.1039]	0.82	0.85	0.78	0.50
Two stage, $\lambda^{(i)}, t_0 = 0.99$	0.1826, [0.1802, 0.1851]	0.40	0.98	0.40	0.50
Two stage, $\lambda^{(i)}, t_0 = 0.999$	0.2559, [0.2525, 0.2594]	0.29	0.99	0.29	0.55

Table 6: Example 3. A common observation. We took $X(0) = (20, 20, 20)$, $c = (0.5, 1, 0.1, 1)$, $t = T = 1$, $y_T = 24$. We used $\bar{\lambda}$ to denote common intensity and $\lambda^{(i)}$ for individual intensities. We used $N_r = 1,000$ trials, filter sample size $N_s = 1,000$.

Method	TVE (95% interval)	ESF (W)	ESF (\tilde{L})	ESF	Runtime (seconds)
Naive	0.4347, [0.4286, 0.4408]	N/A	N/A	0.04	0.326
One stage $\lambda_1, \Delta t = 0.1$	0.2184, [0.2152, 0.2217]	0.70	0.12	0.14	0.791
One stage $\lambda_2, \Delta t = 0.1$	0.2103, [0.2071, 0.2135]	0.84	0.14	0.15	0.794
Two stage, $\bar{\lambda}, t_0 = 0.5$	0.1677, [0.1653, 0.1701]	0.64	0.24	0.25	0.59
Two stage, $\bar{\lambda}, t_0 = 0.8$	0.1627, [0.1601, 0.1654]	0.50	0.46	0.31	0.50
Two stage, $\bar{\lambda}, t_0 = 0.9$	0.1885, [0.1860, 0.1909]	0.34	0.66	0.25	0.56
Two stage, $\bar{\lambda}, t_0 = 0.99$	0.3963, [0.3911, 0.4015]	0.08	0.92	0.07	0.53
Two stage, $\bar{\lambda}, t_0 = 0.999$	0.5545, [0.5472, 0.5617]	0.05	0.96	0.05	0.54
Two stage, $\lambda^{(i)}, t_0 = 0.5$	0.2115, [0.2070, 0.2160]	0.64	0.19	0.19	0.58
Two stage, $\lambda^{(i)}, t_0 = 0.8$	0.1641, [0.1619, 0.1663]	0.52	0.49	0.32	0.51
Two stage, $\lambda^{(i)}, t_0 = 0.9$	0.1890, [0.1866, 0.1914]	0.33	0.68	0.26	0.50
Two stage, $\lambda^{(i)}, t_0 = 0.99$	0.3940, [0.3887, 0.3992]	0.07	0.89	0.07	0.53
Two stage, $\lambda^{(i)}, t_0 = 0.999$	0.5662, [0.5585, 0.5739]	0.05	0.92	0.05	0.56

Table 7: Example 3. A rare observation. We took $X(0) = (20, 20, 20)$, $c = (0.5, 1, 0.1, 1)$, $t = T = 1$, $y_T = 20$. We used $\bar{\lambda}$ to denote common intensity and $\lambda^{(i)}$ for individual intensities. We used $N_r = 1,000$ trials, filter sample size $N_s = 1,000$.

6.4 Comparison with using the conditional propensity

Here, we compare the targeting algorithm with the idea of using the conditional propensity function (when available) to generate \tilde{Z} as discussed in section 3.4 (also see [16]). We use Examples 1 and 2 since in those two examples, the conditional propensities can be computed as described in Section 5. However, as mentioned there, only in Example 1 it is possible to obtain samples from the non-exponential wait times per the conditional propensity. In Example 2, sampling from the non-exponential wait time is not practical. An alternative to sampling from the non-exponential wait time is to approximate the conditional propensity by holding its value constant at the most recent jump time, a strategy used in [16] which makes the wait times exponential.

In Example 1, as Table 4 shows, the exact time-dependent conditional propensity method (CP exact) has the best performance, as expected. However, our targeting method significantly outperformed the approximate conditional propensity method (CP approx) described above. In Example 2, as mentioned earlier, only the “CP approx” method is practical. For the small total copy number case, as shown in Table 1, “CP-approx” method slightly outperforms our targeting methods in the rare observation case while slightly underperforms our targeting methods in the common observation case. However, the computational time is much longer

than the targeting methods. Moreover, in Example 2, for the modestly large copy number case the “CP-approx” method is computationally prohibitive and hence is not shown in Table 2.

6.5 Resampling

Since the disparity of the weights will usually grow fast as time progresses, we could use the resampling technique to alleviate the issue [2, 25]. Example 2 with a larger final time $T = 10$ is used to illustrate this. We resampled after intervals of length Δs and for convenience we took $\Delta s = \Delta t = 0.25$, that is equal to the piecewise constant intervals of the intensity. See Table 8, where it is clear (from the expected TVE) that resampling improves the performance. However, the resampling procedure itself introduces extra randomness which could hurt the accuracy on the other hand. The frequency of resampling is an interesting question to explore as in particle filtering method.

Method	Observation	TVE (95% interval)	ESF (W)	ESF (\tilde{L})	ESF	Runtime (seconds)
Naive	$y_T = 80$	0.4178, [0.4058, 0.4299]	N/A	N/A	0.059	11.1
Targeting ($\bar{\lambda}_1$), without resampling	(common observation)	1.1456, [1.1126, 1.1786]	0.861	0.005	0.005	35.9
Targeting, resample ($\bar{\lambda}_1$), $\Delta s = 0.25$		0.1574, [0.1533, 0.1615]			0.948	39.2
Naive	$y_T = 98$	1.7246, [1.7000, 1.7492]	N/A	N/A	0.002	17.9
Targeting ($\bar{\lambda}_1$), without resampling	(rare observation)	1.3004, [1.2694, 1.3315]	0.804	0.004	0.003	27.5
Targeting, resample ($\bar{\lambda}_1$), $\Delta s = 0.25$		0.1830, [0.1779, 0.1881]			0.779	27.9

Table 8: Example 2. $S_1 \longleftrightarrow S_2$. We took $z_0 = (100, 100)$, $c = (1, 1.5)$, $T = 10$, $t = 9$, $\Delta t = 0.25$, while taking a common observation $y_T = 80$ as well as a rare observation $y_T = 98$. We used $N_r = 100$ trials, while each trial used filter sample size $N_s = 2,000$. We note that in our experiment for the rare observation case, the naive method has 98% of the trials successfully producing a meaningful estimate without encountering a zero denominator issue.

6.6 Increasing filter sample size

Finally, as a sanity check, we investigate if the expected TVE decreases with the number of particles N_s in the filter. Indeed, as seen in Table 9, the expected total variation error (TVE) decreases as filter sample size increases.

Method	Filter sample size	TVE (95% interval)	Runtime (seconds)
Naive	1,000	0.2171, [0.2078, 0.2265]	0.12
	2,000	0.1513, [0.1446, 0.1579]	0.33
	4,000	0.1063, [0.1015, 0.1110]	1.60
	8,000	0.0757, [0.0724, 0.0789]	2.92
One stage λ_1	1,000	0.1704, [0.1613, 0.1795]	0.29
	2,000	0.1177, [0.1115, 0.1240]	1.55
	4,000	0.0884, [0.0843, 0.0924]	4.28
	8,000	0.0597, [0.0570, 0.0625]	7.59

Table 9: Example 3. A common observation. We took $X(0) = (20, 20, 20)$, $c = (0.5, 1, 0.1, 1)$, $t = T = 1$, $\Delta t = 0.1$, $y_T = 24$. We used $N_r = 100$ trials, filter sample size are varied.

7 Conclusion

In this paper we provided a new particle filtering method, *the targeting algorithm*, for statistical inference in stochastic reaction networks where exact partial state observations are available in discrete time snapshots. Our targeting algorithm is designed so that the simulated process always satisfies the observations. This is accomplished by splitting the reaction channels into “free” and “slaved” ones and making use of the fact that the partial state observations impose linear constraints among the reaction counts. Moreover, under the simulation probability measure, in between observations, the reaction count processes are taken to be inhomogeneous Poissons to facilitate easy interpolation in between observations.

We provided theoretical justification as well as numerical examples where the latter show superior performance compared to a prediction/correction approach. We also discussed and illustrated different choices within our targeting algorithm which include the selection of the inhomogeneous intensity for the proposals, the choice of free and slaved reactions and the two-stage targeting approach. The optimal choice among these options is the subject of future research. While our work focused on the case of exact observations (which also covers Poisson type noise as explained in Section 1), our method could be adapted to the case of additive Gaussian noise considered by other researchers [16, 9]. In fact, the analysis provided in Appendix A suggests that, at least in the case of small additive noise, our targeting approach should perform much better than a prediction/correction approach. This is the subject of future investigations.

Appendix A: Advantage of targeting over prediction/correction

Here we consider a probability space $(\Omega, \mathcal{F}, P_0)$ and an m -variate Poisson random variable $R(T) = (R'(T), R''(T))$ where $R'(T)$ is $\mathbb{Z}_+^{m_1}$ valued and $R''(T)$ is $\mathbb{Z}_+^{m_2}$ valued, with mean $M = (M', M'')$. Let C be a matrix and d be a vector of appropriate dimensions such that the event $\{R''(T) = CR'(T) + d\}$ makes sense. We wish to generate an identically distributed weighted sample $(K^{(i)}, W^{(i)})$ distributed like (K, W) so that the conditional distribution is given by

$$P_0\{R(T) = k \mid R''(T) = CR'(T) + d\} = \frac{\mathbb{E}_0[1_{\{k\}}(K)W]}{\mathbb{E}_0[W]}.$$

The targeting approach generates $K' \sim R'(T)$ and sets $K'' = CK' + d$ and sets

$$W = W_p = \sum_{k'} P_0\{K'' = Ck' + d\} 1_{\{k'\}}(K').$$

A more direct approach based on prediction/correction would be to generate $K \sim R(T)$ and set $W = W_i = 1_{\{K'' = CK' + d\}}$. We have used subscripts p and i to denote the Poisson weight and the indicator weight respectively. The effective sample fraction of a weight W (w.r.t. probability measure P_0) could be taken to be $\text{ESF} = \mathbb{E}_0^2[W]/E_0[W^2]$. We obtain the following formulae for ESF_p and ESF_i , the effective sample fractions of the Poisson weight and the indicator weight. First we note that

$$\mathbb{E}_0[W_p] = \mathbb{E}_0[W_i] = P_0\{R''(T) = CR'(T) + d\}, \quad (31)$$

so that both weights have the same mean which equals the probability of the event $\{R''(T) = CR'(T) + d\}$. On the other hand, we obtain

$$\mathbb{E}_0[W_i^2] = \mathbb{E}_0[W_i] = P_0\{R''(T) = CR'(T) + d\},$$

where as

$$\mathbb{E}_0[W_p^2] = \sum_{k'} P_0^2\{R''(T) = Ck' + d\} P_0\{R'(T) = k'\}.$$

Let us denote by

$$\bar{\rho} = \max_{k'} P_0\{R''(T) = Ck' + d\}.$$

Then, we see that $\mathbb{E}_0[W_p^2] \leq \bar{\rho} P_0\{R''(T) = CR'(T) + d\}$. Consequently, $\text{ESF}_p \geq \text{ESF}_i/\bar{\rho}$. As an upper bound of $\bar{\rho}$ we may use the maximum of the p.m.f. of $R''(T)$ which is an m_2 -variate Poisson. The maximum p.m.f. of a Poisson random

variable with mean M may be approximated by using the Stirling’s approximation (which holds here for large M) to obtain $(2\pi M)^{-1/2}$. Hence, we obtain the bound

$$\text{ESF}_p \geq \text{ESF}_i (2\pi)^{m_2/2} (M_{m_1+1} \cdots M_m)^{1/2},$$

which is valid when M_i are sufficiently large. It is clear that the larger the product $M_{m_1+1} \cdots M_m$, the more efficient the targeting method would be over the prediction/correction approach. As an example, if only a single species is observed, then there is only one constraint among the reaction counts and hence $m_2 = 1$. If the mean reaction count is modestly large, say 10, we expect an improvement in the effective sample fraction of $\sqrt{20\pi} \approx 7.9$. Clearly, the improvement gets better if the number of observed species increases.

References

- [1] David F Anderson. A modified next reaction method for simulating chemical systems with time dependent propensities and delays. *The Journal of chemical physics*, 127(21):214107, 2007.
- [2] Alan Bain and Dan Crisan. *Fundamentals of stochastic filtering*, volume 60. Springer Science & Business Media, 2008.
- [3] Patrick Billingsley. *Probability and measure*. John Wiley & Sons, 2017.
- [4] Pierre Brémaud. *Point processes and queues: martingale dynamics*. Springer-Verlag, 1981.
- [5] Silvia Calderazzo, Marco Brancaccio, and Bärbel Finkenstädt. Filtering and inference for stochastic oscillators with distributed delays. *Bioinformatics*, 35(8):1380–1387, 2019.
- [6] Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. Sequential monte carlo samplers. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 68(3):411–436, 2006.
- [7] Arnaud Doucet, Nando De Freitas, Neil James Gordon, et al. *Sequential Monte Carlo methods in practice*, volume 1. Springer, 2001.
- [8] Zhou Fang, Ankit Gupta, and Mustafa Khammash. Stochastic filters based on hybrid approximations of multiscale stochastic reaction networks. In

2020 59th IEEE Conference on Decision and Control (CDC), pages 4616–4621. IEEE, 2020.

- [9] Zhou Fang, Ankit Gupta, and Mustafa Khammash. Stochastic filtering for multiscale stochastic reaction networks based on hybrid approximations. *Journal of Computational Physics*, 467:111441, 2022.
- [10] Zhou Fang, Ankit Gupta, and Mustafa Khammash. Convergence of regularized particle filters for stochastic reaction networks. *SIAM Journal on Numerical Analysis*, 61(2):399–430, 2023.
- [11] Paul Fearnhead. Computational methods for complex stochastic systems: a review of some alternatives to mcmc. *Statistics and Computing*, 18:151–171, 2008.
- [12] Bert Fristedt, Naresh Jain, N Krylov, and Nikolai Vladimirovich Krylov. *Filtering and Prediction: A Primer*, volume 10. American Mathematical Soc., 2007.
- [13] Michael A Gibson and Jehoshua Bruck. Efficient exact stochastic simulation of chemical systems with many species and many channels. *The journal of physical chemistry A*, 104(9):1876–1889, 2000.
- [14] Daniel T Gillespie. Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.*, 81:2340–2361, 1977.
- [15] Daniel T Gillespie. Stochastic simulation of chemical kinetics. *Annu. Rev. Phys. Chem.*, 58:35–55, 2007.
- [16] Andrew Golightly and Chris Sherlock. Efficient sampling of conditioned markov jump processes. *Statistics and Computing*, 29:1149–1163, 2019.
- [17] Andrew Golightly and Darren J Wilkinson. Bayesian sequential inference for stochastic kinetic biochemical network models. *Journal of Computational Biology*, 13(3):838–851, 2006.
- [18] Andrew Golightly and Darren J Wilkinson. Bayesian parameter inference for stochastic biochemical network models using particle markov chain monte carlo. *Interface focus*, 1(6):807–820, 2011.

- [19] András Horváth and Daniele Manini. Parameter estimation of kinetic rates in stochastic reaction networks by the em method. In *2008 International Conference on BioMedical Engineering and Informatics*, volume 1, pages 713–717. IEEE, 2008.
- [20] Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. *ASME J. Basic Eng*, 1960.
- [21] Rudolph Emil Kalman and Richard S Bucy. New results in linear filtering and prediction theory. *ASME J. Basic Eng*, 1961.
- [22] Shuohao Liao, Tomáš Vejchodský, and Radek Erban. Tensor methods for parameter estimation and bifurcation analysis of stochastic reaction networks. *Journal of The Royal Society Interface*, 12(108):20150233, 2015.
- [23] Jun S Liu. Metropolized independent sampling with comparisons to rejection sampling and importance sampling. *Statistics and computing*, 6:113–119, 1996.
- [24] Muruhan Rathinam and Hana El Samad. Reversible-equivalent-monomolecular tau: A leaping method for “small number and stiff” stochastic chemical systems. *Journal of Computational Physics*, 224(2):897–923, 2007.
- [25] Muruhan Rathinam and Mingkai Yu. State and parameter estimation from exact partial state observation in stochastic reaction networks. *The Journal of Chemical Physics*, 154(3):034103, 2021.