

AN EFFICIENT FINITE DIFFERENCE METHOD FOR PARAMETER SENSITIVITIES OF CONTINUOUS TIME MARKOV CHAINS*

DAVID F. ANDERSON[†]

Abstract. We present an efficient finite difference method for the computation of parameter sensitivities that is applicable to a wide class of continuous time Markov chain models. The estimator for the method is constructed by coupling the perturbed and nominal processes in a natural manner, and the analysis proceeds by utilizing a martingale representation for the coupled processes. The variance of the resulting estimator is shown to be an order of magnitude lower due to the coupling. We conclude that the proposed method produces an estimator with a lower variance than other methods, including the use of common random numbers, in most situations. Often the variance reduction is substantial. The method is no harder to implement than any standard continuous time Markov chain algorithm, such as “Gillespie’s algorithm.” The motivating class of models, and the source of our examples, are the stochastic chemical kinetic models commonly used in the biosciences, though other natural application areas include population processes and queuing networks.

Key words. finite differences, variance reduction, parameter sensitivities, next reaction method, Gillespie, random time change, continuous time Markov chain, martingale

AMS subject classifications. Primary, 60H35, 65C99; Secondary, 92C40

DOI. 10.1137/110849079

1. Introduction. We develop a new finite difference method for the computation of parameter sensitivities that is applicable to a wide class of continuous time Markov chain models. For $k \in \{1, \dots, M\}$, let $\zeta_k \in \mathbb{R}^d$ denote the possible transition directions for a continuous time Markov chain, and let $\lambda_k : \mathbb{R}^d \rightarrow \mathbb{R}$ denote the respective intensity, or propensity functions.¹ The random time change representation of Kurtz for the model is

$$(1) \quad X(t) = X(0) + \sum_{k=1}^M Y_k \left(\int_0^t \lambda_k(X(s)) ds \right) \zeta_k,$$

where the Y_k are independent, unit-rate Poisson processes. See, for example, [13], [7, Chapter 6] or the recent survey [5]. The infinitesimal generator for the model (1) is the operator \mathcal{A} satisfying

$$(\mathcal{A}f)(x) = \sum_k \lambda_k(x)(f(x + \zeta_k) - f(x)),$$

where $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is chosen from a sufficiently large class of functions. Without loss of generality, we assume throughout that the state space of the process, \mathcal{S} , is a subset of \mathbb{Z}^d .

Consider a family of models (1) indexed by a set of parameters, which we denote by the vector θ . Even when there are good theoretical reasons for believing the

*Received by the editors September 26, 2011; accepted for publication (in revised form) June 6, 2012; published electronically September 4, 2012. This work was supported by NSF-DMS-1009275. <http://www.siam.org/journals/sinum/50-5/84907.html>

[†]Department of Mathematics, University of Wisconsin, Madison, WI 53706 (anderson@math.wisc.edu).

¹In the language of probability theory, the functions λ_k are nearly universally termed intensity functions, whereas in the biosciences they are nearly universally termed propensity functions.

model is a reasonable description of some phenomenon, usually the parameters are not known precisely and have to be estimated experimentally. Depending on the setup and the parameters in question, it may be difficult to obtain good estimates. Thus, it is important to analyze how sensitive features of interest in the model are to variation in the parameters. For ease of exposition we take θ to be a scalar, though note that it is trivial to extend all of the ideas of the paper to the setting of $\theta \in \mathbb{R}^\ell$ for some $\ell > 0$.

We let $f : \mathcal{S} \rightarrow \mathbb{R}$ be a function of the state of the system that gives a measurement of interest. For example, f could be the abundance of one of the components at a particular time. Define $J(\theta) \stackrel{\text{def}}{=} \mathbb{E}f(X^\theta(t))$, where the θ dependence is being made explicit. The problem of interest is to efficiently approximate $J'(\theta)$.

There are a number of methods that can be used for the computation of such parameter sensitivities in this setting, including finite differences, likelihood ratios and Girsanov transformations, and infinitesimal perturbation analysis, each with its own benefits and drawbacks; see, for example, [6, 11, 17, 18]. We focus on finite difference methods, which, due to their simplicity, are the most popular choice. Specifically, in this paper a new finite difference method is introduced that is easy to implement, is analytically tractable, and typically produces an estimate with a given tolerance with substantially lower computational complexity than that obtained using the other methods currently known to the author.

While continuous time Markov chain models of the general form (1) are used ubiquitously in both industry and the sciences to model natural phenomena ranging from population processes to queueing networks, we feel the method developed here will be most useful in the study of stochastic models of biochemical reaction networks. We will therefore choose the language of biochemistry throughout and also choose this area as the setting for our examples.

A biochemical reaction network is a chemical system involving multiple reactions and chemical species. If the abundances of the constituent molecules of a reaction network are sufficiently high, then their concentrations are typically modeled by a coupled set of ordinary differential equations. If, however, the abundances are low, then the standard deterministic models do not provide a good representation of the behavior of the system and stochastic models are used. The simplest stochastic models of such networks [12, 16] treat the system as a continuous time Markov chain with the state, X , being the number of molecules of each species and with reactions modeled as possible transitions of the chain. More explicitly, if the k th reaction happens at time t , then the system is updated by the *reaction vector* ζ_k ,

$$X(t) = X(t-) + \zeta_k.$$

Letting $\lambda_k : \mathbb{R}^d \rightarrow \mathbb{R}$ denote the intensity, or propensity, of the k th reaction, we see that this stochastic model satisfies (1).

As will be pointed out in the following sections, the strategy being proposed here is in some ways similar to the common reaction path (CRP) method proposed in [18], which is also a quite capable estimator for finite differences. However, there are important differences. First, the actual coupling, and hence simulation, of the relevant processes is different. The coupling proposed here tends to provide an estimator with a lower variance, especially when the process is considered for moderate to large times. Second, the coupling proposed here lends itself to analysis more readily than that used in [18], as the centered counting processes used in our coupling are martingales with respect to the natural filtration of the process [2]. Third, the method

being proposed here is as easy to implement as the usual Gillespie algorithm or next reaction method. The strategy employed in [18], on the other hand, requires being quite careful with the seeds of the pseudorandom number generators used since one independent seed is required per reaction channel per sample path generated. Finally, the coupling proposed here essentially converts the problem of generating two paths of a continuous time Markov chain into a problem of generating one path of a different continuous time Markov chain with an enlarged state space. Therefore, all analytical and computational techniques developed for the study of continuous time Markov chains, of which there are many, will be employable on this larger system and therefore applicable to the problem of computing sensitivities.

The most common finite difference coupling used today for the approximation of sensitivities is probably an implementation of Gillespie’s algorithm plus using common random numbers (CRN). We will also discuss this coupling and give the relevant stochastic representation of it. We will conclude that the coupling proposed here will produce a lower variance estimator than CRN for the same reasons that it produces a lower variance estimator than the CRP scheme of [18].

The main goals of this paper are to introduce the new method and to provide the mathematical analysis of the expected squared difference between the two relevant processes, though some relevant examples will also be provided. In section 2, we formally introduce our mathematical model of interest, including all technical assumptions. In section 3, we develop our new finite difference estimator and provide sharp analytical bounds. We also discuss the long time behavior of the introduced estimator and compare it with both CRP and CRN. We conclude that the proposed method will be quite superior for moderate and large time scales. In section 4, we provide examples demonstrating our main results.

2. The formal setup. We consider the family of models

$$(2) \quad X^\theta(t) = X^\theta(0) + \sum_{k=1}^M Y_k \left(\int_0^t \lambda_k^\theta(X^\theta(s)) ds \right) \zeta_k,$$

where the Y_k are independent, unit-rate Poisson processes, the vector θ represents a given choice of parameters that we are making explicit in the notation, and all other notation is as before. The assumption that there are a finite number of possible jump directions ζ_k can almost certainly be weakened. However, this assumption makes the analysis significantly cleaner, and all the motivating models (such as those arising from biochemistry) naturally satisfy such a condition. We define $F^\theta : \mathbb{Z}^d \rightarrow \mathbb{R}^d$ by

$$F^\theta(z) \stackrel{\text{def}}{=} \sum_k \lambda_k^\theta(z) \zeta_k.$$

We make the following running assumptions throughout the remainder of the paper. The first is that the intensity functions are uniformly (in θ) globally Lipschitz. The second is that the intensity functions scale at least linearly with perturbations to θ .

Assumption 1. We suppose that there is a $K_1 > 0$ for which

$$|\lambda_k^\theta(x) - \lambda_k^\theta(y)| + |F^\theta(x) - F^\theta(y)| \leq K_1|x - y|$$

for all k, θ of interest and all $x, y \in \mathcal{S}$.

Assumption 2. We suppose there is a $K_2 > 0$ so that for all k and all $\epsilon < 1$,

$$\sup_{x \in \mathcal{S}} [|\lambda_k^{\theta+\epsilon}(x) - \lambda_k^\theta(x)| + |F^{\theta+\epsilon}(x) - F^\theta(x)|] \leq K_2\epsilon.$$

Assumption 1 can almost certainly be weakened to a local Lipschitz condition, in which case analytical methods similar to those found in [9] and/or [15] can be applied. Proving our main results in such generality, while possible and certainly worth doing in future work, will be significantly messier, and we feel the main points of the analysis will be lost. Note that Assumption 1 automatically holds if \mathcal{S} is a bounded set. In the chemical setting, mild assumptions on the intensity functions ensure that the nonnegative orthant is forward invariant. Therefore, \mathcal{S} is bounded if there is a vector $\xi \in \mathbb{Z}_{>0}^d$ for which $\xi \cdot \zeta_k \leq 0$ for all k . For example, such a ξ exists if mass is conserved. Assumption 1 also holds if the intensity functions are simply set to zero outside of a compact subset of \mathbb{Z}^d , which has the same effect as analyzing the standard model up until a stopping time τ , defined to be the time the processes leave a given compact set (essentially using a localization argument).

Another relevant situation in which Assumption 1 holds is when the equivalent *scaled* models are analyzed. While we point the reader to [2, 3, 4, 10] for a thorough description of this model in the biosciences, we will briefly discuss it here. For some parameter of the system, N , we let X^N denote the process with i th component $X_i^N = X_i/N^{\alpha_i}$, where the $\alpha_i \geq 0$ are chosen so that $X_i^N = O(1)$. Under mild assumptions on the intensity functions λ_k , it can be shown that X^N satisfies

$$(3) \quad X^N(t) = X^N(0) + \sum_k Y_k \left(\int_0^t N^{\beta_k + \nu_k \cdot \alpha} \lambda_k(X^N(s)) ds \right) \zeta_k^N,$$

where $\zeta_{k,i}^N = \zeta_{k,i}/N^{\alpha_i}$, and β_k is chosen so that $\lambda_k(X^N(\cdot)) = O(1)$. This scaled model is $O(1)$ and therefore more readily satisfies Assumption 1, where now it is understood that the state space is

$$\mathcal{S}^N \stackrel{\text{def}}{=} \{z \in \mathbb{R}^d \mid z_i = x_i N^{-\alpha_i}, x \in \mathcal{S}\}.$$

In many applications, it is more relevant to compute expectations and parameter sensitivities of the scaled model (3) than of the unscaled version (1), though we do not revisit this point in the current paper.

2.1. The basic problem and the benefits of variance reduction. We let $f : \mathbb{Z}_{\geq 0}^d \rightarrow \mathbb{R}$ be a function of the state of the system which gives a measurement of interest and define

$$J(\theta) \stackrel{\text{def}}{=} \mathbb{E}f(X^\theta(t)).$$

The problem of interest is to efficiently estimate $J'(\theta)$, where we recall that we are making the simplifying assumption that θ is one-dimensional.

To estimate $J'(\theta)$ the centered finite difference is often used:

$$(4) \quad J'(\theta) \approx \frac{\mathbb{E}f(X^{\theta+\epsilon/2}(t)) - \mathbb{E}f(X^{\theta-\epsilon/2}(t))}{\epsilon},$$

as its bias is $O(\epsilon^2)$ [6]. That is,

$$J'(\theta) = \frac{\mathbb{E}f(X^{\theta+\epsilon/2}(t)) - \mathbb{E}f(X^{\theta-\epsilon/2}(t))}{\epsilon} + O(\epsilon^2).$$

This should be compared with the forward difference, which has a bias of $O(\epsilon)$:

$$J'(\theta) = \frac{\mathbb{E}f(X^{\theta+\epsilon}(t)) - \mathbb{E}f(X^\theta(t))}{\epsilon} + O(\epsilon).$$

The estimator for (4) using centered finite differences is

$$(5) \quad D_R(\epsilon) = \frac{1}{R} \sum_{i=1}^R d_{[i]}(\epsilon),$$

with

$$(6) \quad d_{[i]}(\epsilon) = \frac{f(X_{[i]}^{\theta+\epsilon/2}(t)) - f(X_{[i]}^{\theta-\epsilon/2}(t))}{\epsilon},$$

where $X_{[i]}^\theta$ represents the i th path generated with parameter choice θ , and R is the number of paths generated. If $X_{[i]}^{\theta+\epsilon/2}(t)$ and $X_{[i]}^{\theta-\epsilon/2}(t)$ are computed independently, the variance of $d_{[i]}(\epsilon)$ is $O(\epsilon^{-2})$, and, hence, the variance of $D_R(\epsilon)$ is $O(R^{-1}\epsilon^{-2})$. Note that

$$\begin{aligned} \mathbb{E}(D_R(\epsilon) - J'(\theta))^2 &= \text{Var}(D_R(\epsilon)) + (\mathbb{E}D_R(\epsilon) - J'(\theta))^2 \\ &= O(R^{-1}\epsilon^{-2}) + O(\epsilon^4), \end{aligned}$$

which, for a given R , is minimized when $\epsilon = O(R^{-1/6})$, at a value that is $O(R^{-2/3})$. Therefore, the optimal convergence rate to the exact value, in the sense of confidence intervals, is $O(R^{-1/3})$ [6].

Many computations are performed with a target variance (which yields a target size of the confidence interval). Denoting the target variance by V^* , we see that the number of paths required is then approximated by the solution to

$$\text{Var}\left(\frac{1}{R} \sum_{i=1}^R d_{[i]}(\epsilon)\right) = \frac{1}{R} \text{Var}(d(\epsilon)) = V^* \implies R = \frac{1}{V^*} \text{Var}(d(\epsilon)).$$

Thus, decreasing the variance of $d(\epsilon)$ lowers the computational complexity (total number of computations) required to solve the problem. The basic idea of coupling, in the context of this paper, is to lower the variance of $d(\epsilon)$ by simulating $X^{\theta+\epsilon/2}$ and $X^{\theta-\epsilon/2}$ simultaneously so that the two processes are highly correlated or “coupled.” That is, instead of generating paths independently, we want to generate a *pair* of paths $(X^{\theta+\epsilon/2}, X^{\theta-\epsilon/2})$ so that for appropriate choices of f , the variance of $f(X^{\theta+\epsilon/2}) - f(X^{\theta-\epsilon/2})$ is reduced. The basic idea of any such coupling is to reuse, or share, some portion of the driving “noise” in the generation of each process. As already alluded to in the introduction, one such finite difference method that achieved a substantial reduction in variance due to coupling can be found in [18], which we discuss in more detail in later sections.

In section 3.1, we will develop a new coupling technique so that the variance of $d_{[i]}(\epsilon)$ in (6) is $O(\epsilon^{-1})$, a full order of magnitude lower (in ϵ) than when the paths were generated independently. This will lead to a finite difference method with an optimal convergence rate, in the sense of the above paragraph, of $O(R^{-2/5})$, achieved when $\epsilon = O(R^{-1/5})$. More importantly, however, the variance of the estimator (5) will be $O(R^{-1}\epsilon^{-1})$, which should be compared with a variance of $O(R^{-1}\epsilon^{-2})$ when independent paths are used. Thus, the number of paths (and computational complexity) required to solve a given problem will be reduced by an order of ϵ .

3. Coupled finite differences. In section 3.1, we discuss how to couple the requisite processes for the coupled finite difference method being proposed here. In section 3.2, we provide sharp bounds on the variance of the estimator.

3.1. Coupling the processes. Whether using the forward or centered difference, the main problem is to intelligently produce two paths generated from systems whose parameters differ by an order of ϵ . A good coupling should satisfy three things:

- (i) it should minimize the variance of the difference (6);
- (ii) it should be easy to simulate; and
- (iii) it should be analytically tractable.

We will show that the coupling (7) below satisfies each of these requirements; however, we begin by motivating the coupling by two simpler problems that capture the core idea.

We consider the problem of trying to understand the difference between $Z_1(t)$ and $Z_2(t)$, where Z_1, Z_2 are Poisson processes with rates 13.1 and 13, respectively. We let Y_1 and Y_2 be independent unit-rate Poisson processes and set

$$\begin{aligned} Z_1(t) &= Y_1(13t) + Y_2(0.1t), \\ Z_2(t) &= Y_1(13t), \end{aligned}$$

where we use the additivity property of Poisson processes. The important point to note is that both processes Z_1 and Z_2 are using the process $Y_1(13t)$ to generate simultaneous jumps. The process Z_1 then uses the auxiliary process $Y_2(0.1t)$ to jump the extra times that Z_2 does not. The processes Z_1, Z_2 will jump together the vast majority of times and in this way will be very tightly coupled.² The coupling above also already hints at the main points of the mathematical analysis that will be carried out in section 3.2 as

$$Z_1(t) - Z_2(t) = Y_2(0.1t),$$

and so

$$\begin{aligned} \mathbb{E}|Z_1(t) - Z_2(t)| &= \mathbb{E}Y_2(0.1t) = 0.1t, \\ \mathbb{E}(Z_1(t) - Z_2(t))^2 &= \mathbb{E}Y_2(0.1t)^2 = 0.1t + 0.01t^2. \end{aligned}$$

More generally, if Z_1 and Z_2 are nonhomogeneous Poisson processes with intensities $f(t)$ and $g(t)$, respectively, then we could let Y_1, Y_2 , and Y_3 be independent, unit-rate Poisson processes and define

$$\begin{aligned} Z_1(t) &= Y_1 \left(\int_0^t f(s) \wedge g(s) ds \right) + Y_2 \left(\int_0^t f(s) - (f(s) \wedge g(s)) ds \right), \\ Z_2(t) &= Y_1 \left(\int_0^t f(s) \wedge g(s) ds \right) + Y_3 \left(\int_0^t g(s) - (f(s) \wedge g(s)) ds \right), \end{aligned}$$

where we are using that, for example,

$$Y_1 \left(\int_0^t f(s) \wedge g(s) ds \right) + Y_2 \left(\int_0^t f(s) - (f(s) \wedge g(s)) ds \right) \stackrel{D}{=} Y \left(\int_0^t f(s) ds \right),$$

²In this case, the long-run percentage of jumps that are shared can be quantified precisely as $13/(13 + 0.1) \approx 0.99923$.

where Y is a unit-rate Poisson process and we define $a \wedge b \stackrel{\text{def}}{=} \min\{a, b\}$. Thus, we are coupling the processes by splitting up the intensity functions into two pieces, one shared

$$f(s) \wedge g(s),$$

and the other not, and then using the same noise, Y_1 , on the shared portion.

We return to the problem of coupling the main processes of interest to us. For ease of notation, we will couple the processes $X^{\theta+\epsilon}$ and X^θ as opposed to $X^{\theta+\epsilon/2}$ and $X^{\theta-\epsilon/2}$, with the understanding that generating the centered difference is performed in the obvious manner. We generate our coupled processes $(X^{\theta+\epsilon}, X^\theta)$ via

$$\begin{aligned} (7) \quad X^{\theta+\epsilon}(t) &= X^{\theta+\epsilon}(0) + \sum_k Y_{k,1} \left(\int_0^t \lambda_k^{\theta+\epsilon}(X^{\theta+\epsilon}(s)) \wedge \lambda_k^\theta(X^\theta(s)) ds \right) \zeta_k \\ &\quad + \sum_k Y_{k,2} \left(\int_0^t \lambda_k^{\theta+\epsilon}(X^{\theta+\epsilon}(s)) - \lambda_k^{\theta+\epsilon}(X^{\theta+\epsilon}(s)) \wedge \lambda_k^\theta(X^\theta(s)) ds \right) \zeta_k, \\ X^\theta(t) &= X^\theta(0) + \sum_k Y_{k,1} \left(\int_0^t \lambda_k^{\theta+\epsilon}(X^{\theta+\epsilon}(s)) \wedge \lambda_k^\theta(X^\theta(s)) ds \right) \zeta_k \\ &\quad + \sum_k Y_{k,3} \left(\int_0^t \lambda_k^\theta(X^\theta(s)) - \lambda_k^{\theta+\epsilon}(X^{\theta+\epsilon}(s)) \wedge \lambda_k^\theta(X^\theta(s)) ds \right) \zeta_k, \end{aligned}$$

where the $Y_{k,i}$ are unit-rate Poisson processes and all other notation is as before. Thus, and just as in the example pertaining to the nonhomogeneous Poisson processes above, the effect of the intensity function $\lambda_k^{\theta+\epsilon}$ on the process $X^{\theta+\epsilon}$ has been split into two pieces: one of size $\lambda_k^{\theta+\epsilon}(X^{\theta+\epsilon}(s)) \wedge \lambda_k^\theta(X^\theta(s)) ds$, and one of size

$$\lambda_k^{\theta+\epsilon}(X^{\theta+\epsilon}(s)) - \lambda_k^{\theta+\epsilon}(X^{\theta+\epsilon}(s)) \wedge \lambda_k^\theta(X^\theta(s)) ds.$$

Further, since the two processes $X^{\theta+\epsilon}$ and X^θ share the contribution of each of the terms with intensity

$$\lambda_k^{\theta+\epsilon}(X^{\theta+\epsilon}(s)) \wedge \lambda_k^\theta(X^\theta(s)) ds,$$

we expect them to be highly correlated. It is important to note that the marginal processes have the same distributions as the respective processes generated via (2). This fact can be seen by noting that (7) is a continuous time Markov chain and that the transition rates of the marginal processes are identical to those of (2) with the corresponding rate constants. Note also that the coupling (7) is essentially the same as in the toy problems above where we coupled Z_1 and Z_2 .

A coupling similar to (7) first appeared in [14]. More recently, it was used in [2] to study the strong error of different approximation methods in the discrete stochastic case and in [3] to generate paths so as to apply multilevel Monte Carlo techniques in the continuous time Markov chain setting. The application of the coupling (7) towards the problem of parametric sensitivity analysis is the main contribution of this paper.

As discussed at the end of section 1, the process $(X^{\theta+\epsilon}, X^\theta)$ satisfying (7) is a continuous time Markov chain with state space $\mathbb{Z}^d \times \mathbb{Z}^d$. Therefore, all analytical and

computational techniques developed for the study of continuous time Markov chains will be applicable to this system and, hence, to the problem of computing sensitivities.

Before proceeding with the analysis, we give the algorithm for generating a path $(X^{\theta+\epsilon}, X^\theta)$ via (7). We note that the method below is the next reaction method applied to (7) [1, 8]. See [1] for a thorough explanation of how the next reaction method is equivalent to simulating representations of the forms considered here. Below, we will denote a uniform $[0, 1]$ random variable by $\text{rand}(0, 1)$, and we remind the reader that if $U \sim \text{rand}(0, 1)$, then $\ln(1/U)$ is an exponential random variable with a parameter of one. All random variables generated are assumed to be independent of each other and all previous random variables. It is assumed that the processes start with the same initial condition, though this can be weakened in the obvious manner. Finally, we note that it is also possible to simulate the continuous time Markov chain (7) by the obvious adaption of Gillespie's direct, or optimized direct, algorithm. While we do not formally provide that algorithm here, it will be problem specific as to which implementation (Gillespie versus next reaction method) is more efficient.

ALGORITHM 1 (simulation of the representation (7)). **Initialize.** Set $X^{\theta+\epsilon} = X^\theta = x$ and $t = 0$. For each k and $i \in \{1, 2, 3\}$, set

- $P_{k,i} = \ln(1/u_{k,i})$, where $u_{k,i}$ is $\text{rand}(0, 1)$, and
- $T_{k,i} = 0$.

Repeat the following steps:

- (i) For each k , set
 - $A_{k,1} = \lambda_k^{\theta+\epsilon}(X^{\theta+\epsilon}) \wedge \lambda_k^\theta(X^\theta)$,
 - $A_{k,2} = \lambda_k^{\theta+\epsilon}(X^{\theta+\epsilon}) - A_{k,1}$, and
 - $A_{k,3} = \lambda_k^\theta(X^\theta) - A_{k,1}$.
- (ii) For each k and $i \in \{1, 2, 3\}$, set

$$\Delta t_{k,i} = \begin{cases} (P_{k,i} - T_{k,i})/A_{k,i} & \text{if } A_{k,i} > 0, \\ \infty & \text{if } A_{k,i} = 0. \end{cases}$$

- (iii) Set $\Delta = \min_{k,i} \{\Delta t_{k,i}\}$, and let $\mu \equiv \{k, i\}$ be the indices where the minimum is achieved.
- (iv) Set $t = t + \Delta$.
- (v) Update state vectors according to reaction ζ_μ (where the minimum occurred in step (iii)):

$$(X^{\theta+\epsilon}, X^\theta) = \begin{cases} (X^{\theta+\epsilon}, X^\theta) + (\zeta_k, \zeta_k) & \text{if } i = 1, \\ (X^{\theta+\epsilon}, X^\theta) + (\zeta_k, 0) & \text{if } i = 2, \\ (X^{\theta+\epsilon}, X^\theta) + (0, \zeta_k) & \text{if } i = 3. \end{cases}$$

- (vi) For each k and $i \in \{1, 2, 3\}$, set $T_{k,i} = T_{k,i} + A_{k,i} \times \Delta$.
- (vii) Set $P_\mu = P_\mu + \ln(1/u)$, where u is $\text{rand}(0, 1)$.
- (viii) Return to step (i) or quit.

Note that at most two of $A_{k,1}, A_{k,2}, A_{k,3}$ will be nonzero at each step. Further, it will often be that $A_{k,1} \gg \max\{A_{k,2}, A_{k,3}\}$ and the processes will move together the vast majority of the time (which is, of course, the whole point of such a coupling), showing that the cost of generating the path $(X^{\theta+\epsilon}, X^\theta)$ will be less than the cost of generating two paths via the representation (2). This fact is observed in the data collected on the numerical examples in section 4.

The common reaction path method. We now revisit the point that the strategy being proposed here is similar to that proposed in [18], where instead of the coupling (7) the authors used what is equivalent to

$$(8) \quad \begin{aligned} X^{\theta+\epsilon}(t) &= X^{\theta+\epsilon}(0) + \sum_k Y_k \left(\int_0^t \lambda_k^{\theta+\epsilon}(X^{\theta+\epsilon}(s)) ds \right) \zeta_k, \\ X^\theta(t) &= X^\theta(0) + \sum_k Y_k \left(\int_0^t \lambda_k^\theta(X^\theta(s)) ds \right) \zeta_k, \end{aligned}$$

where the Y_k are independent unit-rate Poisson processes and all other notation is as before. The key point is that they are using the same Poisson processes for the generation of each path. As stated in section 1, the estimator built with paths generated via (8) is quite capable in many circumstances. The main differences between their method and the one being proposed in this paper via (7) are the following:

1. The processes generated via (7) are generally coupled tighter than those generated via (8), resulting in a lower variance for the estimator, sometimes substantially so. However, sometimes the coupling (8) produces a lower variance estimator than (7) when the terminal time T is small. These facts will be demonstrated via examples in section 4 and discussed more below.
2. The model (7) is more amenable to analysis, as the centered counting processes of (7) are martingales with respect to the natural filtration [2], which is not the case for (8).
3. Implementation of (7) is simpler than that of (8), as (7) does not require the generation of many independent seeds for the pseudorandom number generator. In fact, simulation of (7) is no more challenging than simulating any continuous time Markov chain.
4. The coupling (7) makes the problem of computing the difference between two paths into one of computing a single path of a different continuous time Markov chain with an enlarged state space.

The following example is chosen to highlight the advantages of the coupling (7) over those of (8).

Example 1. Consider the simple model in which an mRNA molecule is created and degraded,



which is equivalent to an $M/M/\infty$ queue with arrival rate θ and service rate 0.1. Here we are using the common convention of putting the rate constant of a reaction next to the corresponding reaction vector. We suppose that we want to understand the sensitivity of the expected number of mRNA molecules with respect to the parameter $\theta \approx 2$. We consider how the different representations (7) and (8) “should” behave on this model, whose representation via (2) is

$$(10) \quad X^\theta(t) = X(0) + Y_1(\theta t) - Y_2 \left(\int_0^t 0.1 X^\theta(s) ds \right).$$

For $\epsilon > 0$, let $(X^{\theta+\epsilon}, X^\theta)$ satisfy (7), which for this example is

$$\begin{aligned} X^{\theta+\epsilon}(t) &= X^{\theta+\epsilon}(0) + Y_{1,1}(\theta t) + Y_{1,2}(\epsilon t) \\ &\quad - Y_{2,1} \left(\int_0^t 0.1 X^\theta(s) ds \right) - Y_{2,2} \left(\int_0^t 0.1 (X^{\theta+\epsilon}(s) - X^\theta(s)) ds \right), \\ X^\theta(t) &= X^\theta(0) + Y_{1,1}(\theta t) - Y_{2,1} \left(\int_0^t 0.1 X^\theta(s) ds \right), \end{aligned}$$

where we have used that with this coupling $X^{\theta+\epsilon}(t) \geq X^\theta(t)$ for all $t \geq 0$ if $\epsilon \geq 0$. Therefore, assuming that $X^{\theta+\epsilon}(0) = X^\theta(0)$, we have

$$X^{\theta+\epsilon}(t) - X^\theta(t) = Y_{1,2}(\epsilon t) - Y_{2,2} \left(\int_0^t 0.1 (X^{\theta+\epsilon}(s) - X^\theta(s)) ds \right).$$

Setting $Z^{\theta,\epsilon} = X^{\theta+\epsilon} - X^\theta$, we see that $Z^{\theta,\epsilon}$ itself can be viewed as the solution to (10), though with zero initial condition and input rate ϵ . The mean and variance can be solved as functions of time and satisfy

$$(11) \quad \mathbb{E}Z^{\theta,\epsilon}(t) = \mathbb{E}(X^{\theta+\epsilon}(t) - X^\theta(t)) = \frac{\epsilon}{0.1}(1 - e^{-0.1t}),$$

$$(12) \quad \text{Var}(Z^{\theta,\epsilon}(t)) = \text{Var}(X^{\theta+\epsilon}(t) - X^\theta(t)) = \frac{\epsilon}{0.1}(1 - e^{-0.1t}).$$

On the other hand, if $(X^{\theta+\epsilon}, X^\theta)$ satisfy the coupling (8), then

$$\begin{aligned} X^{\theta+\epsilon}(t) &= X^{\theta+\epsilon}(0) + Y_1(\theta t + \epsilon t) - Y_2 \left(\int_0^t 0.1 X^{\theta+\epsilon}(s) ds \right), \\ X^\theta(t) &= X^\theta(0) + Y_1(\theta t) - Y_2 \left(\int_0^t 0.1 X^\theta(s) ds \right), \end{aligned}$$

and

$$(13) \quad X^{\theta+\epsilon}(t) - X^\theta(t) = Y_1(\theta t + \epsilon t) - Y_1(\theta t) - \left[Y_2 \left(\int_0^t 0.1 X^{\theta+\epsilon}(s) ds \right) - Y_2 \left(\int_0^t 0.1 X^\theta(s) ds \right) \right].$$

In this case, we still have that $\mathbb{E}[X^{k+\epsilon}(t) - X^k(t)]$ satisfies the right-hand side of (11). However, the variance cannot be calculated with such ease as for (12). We note, however, that for large t , we will have $\theta t + \epsilon t \gg \theta t$, and therefore anticipate

$$\int_0^t 0.1 X^{\theta+\epsilon}(s) ds \gg \int_0^t 0.1 X^\theta(s) ds,$$

implying that the two processes $X^{\theta+\epsilon}$ and X^θ should decouple and behave independently. This is demonstrated in a numerical example in section 4 where we show that the variance of the difference (13) converges to 40, which is the same as if $X^{\theta+\epsilon}$ and X^θ were generated independently, and is *substantially* larger than the bound given in (12) for small ϵ .

As will be discussed immediately below, we also expect to see this “decoupling” when Gillespie’s algorithm is implemented with common random numbers (CRN). This also will be demonstrated by examples in section 4.

The above example gives a heuristic as to why the coupling (7) will tend to give a lower variance than (8). For processes generated by (7), whenever $X^{\theta+\epsilon}(t) \approx X^\theta(t)$ during the course of the simulation, the processes have recoupled, regardless of the the history of the process up to that time. On the other hand, if $X^{\theta+\epsilon}$ and X^θ are generated via (8), then $X^{\theta+\epsilon} \approx X^\theta$ need not imply

$$\int_0^t \lambda_k^{\theta+\epsilon}(X^{\theta+\epsilon}(s))ds \approx \int_0^t \lambda_k^{\theta+\epsilon}(X^\theta(s))ds,$$

and so the two processes could be exploring completely different portions of the Poisson processes. Thus, even when $X^{\theta+\epsilon}(t) = X^\theta(t)$ in the course of the simulation of (8), the integrated intensities will not be equal, and so the processes are no longer coupled as tightly as they were at time zero. As time increases, this problem could get worse and the processes can decouple completely (as happens in the example above).

CRN and Gillespie’s algorithm. The standard method of using CRN in the implementation of Gillespie’s algorithm will suffer the same defect as (8) in that for large times the coupled processes can decouple. To understand why, we need to give the correct representation for Gillespie’s algorithm (which is equivalent to simulating the embedded discrete time Markov chain). The following representation can be found in [5]. We define

$$\lambda_0(x) = \sum_k \lambda_k(x) \quad \text{and} \quad q_k(x) = \sum_{i=1}^k \lambda_i(x)/\lambda_0(x).$$

Let Y be a unit-rate Poisson process, and let $\{\xi_i\}$ be an independent and identically distributed sequence of uniform(0,1) random variables that are also independent of Y . Then let X satisfy

$$(14) \quad \begin{aligned} R_0(t) &= Y \left(\int_0^t \lambda_0(X(s))ds \right), \\ X(t) &= X(0) + \sum_k \zeta_k \int_0^t \mathbf{1}_{(q_{k-1}(X(s-)), q_k(X(s-))]}(\xi_{R_0(s-)})dR_0(s). \end{aligned}$$

The counting process R_0 is determining the jump times, which are seen to be exponential random variables with parameter $\lambda_0(X(s-))$. The uniform random variables are then used to select which reaction occurs, with the k th reaction being chosen with probability $\lambda_k(X(s-))/\lambda_0(X(s-))$. Simulation of (14) is called Gillespie’s algorithm (or simply simulating the embedded discrete time Markov chain). The standard CRN + Gillespie algorithm finite difference method, which is probably the most common coupling method used today in the context of finite differences, then consists of using the same Y and choice of $\{\xi_i\}$ for the construction of X^θ and $X^{\theta+\epsilon}$, and it will decouple for the same reasons as those of (8). This is demonstrated numerically in an example in section 4.

Naive couplings: A cautionary tale. At this point it may be tempting to try to couple the processes generated via (7) even tighter by using the same Poisson processes for each of $Y_{k,2}$ and $Y_{k,3}$. This would, in effect, be a highbred version of the common reaction path and coupled finite difference methods. That is, one may be

tempted to use

(15)

$$\begin{aligned} X^{\theta+\epsilon}(t) &= X^{\theta+\epsilon}(0) + \sum_k Y_{k,1} \left(\int_0^t \lambda_k^{\theta+\epsilon}(X^{\theta+\epsilon}(s)) \wedge \lambda_k^\theta(X^\theta(s)) ds \right) \zeta_k \\ &\quad + \sum_k Y_{k,2} \left(\int_0^t \lambda_k^{\theta+\epsilon}(X^{\theta+\epsilon}(s)) - \lambda_k^{\theta+\epsilon}(X^{\theta+\epsilon}(s)) \wedge \lambda_k^\theta(X^\theta(s)) ds \right) \zeta_k, \\ X^\theta(t) &= X^\theta(0) + \sum_k Y_{k,1} \left(\int_0^t \lambda_k^{\theta+\epsilon}(X^{\theta+\epsilon}(s)) \wedge \lambda_k^\theta(X^\theta(s)) ds \right) \zeta_k \\ &\quad + \sum_k Y_{k,2} \left(\int_0^t \lambda_k^\theta(X^\theta(s)) - \lambda_k^{\theta+\epsilon}(X^{\theta+\epsilon}(s)) \wedge \lambda_k^\theta(X^\theta(s)) ds \right) \zeta_k, \end{aligned}$$

where we are now using the same Poisson process for all of the auxiliary processes. In fact, this does not work: the marginal distributions of $(X^{\theta+\epsilon}, X^\theta)$ as generated by (15) are not the same as the original processes, and so this coupling should *not* be used. In fact, the marginal distributions can be so different that the coupled processes will converge to the wrong value as $\epsilon \rightarrow 0$. This fact is best demonstrated by an example.

Example 2. Consider the system arising from the single reaction



with $X_0 = 1$. The stochastic equation of the form (1) governing this system is

$$X^\theta(t) = 1 - Y \left(\theta \int_0^t X^\theta(s) ds \right).$$

Hence, the process can take only the values of one and zero and $\mathbb{E}X^\theta(1) = \exp\{-\theta\}$, which implies that $\frac{d}{d\theta}\mathbb{E}X^\theta(1) = -\exp\{-\theta\}$. In particular,

$$\left. \frac{d}{d\theta} \mathbb{E}X^\theta(1) \right|_{\theta=1} = -e^{-1}.$$

For the coupling (15) for this model we have

$$\epsilon^{-1} \mathbb{E}[X^{1+\epsilon/2} - X^{1-\epsilon/2}] = -e^{-1} \frac{2}{2+\epsilon} \left(e^{\epsilon/2} - e^{-\epsilon/2} \right) = -e^{-1} \epsilon + O(\epsilon^2),$$

which converges to zero as $\epsilon \rightarrow 0$. Perhaps the simplest way to compute the above expectation is to find the condition on the first jump times of the underlying Poisson processes that guarantee the random variable $X^{1+\epsilon/2} - X^{1-\epsilon/2}$ takes a value of negative one. This event has a probability of $O(\epsilon^2)$, implying the result. This computation is left to the interested reader.

3.2. Analytical results. The following theorem is the main analytical result of this paper and allows us to conclude that for any function f satisfying the assumptions of Theorem 3.1 and $(X^{\theta+\epsilon}, X^\theta)$ satisfying (7),

$$\text{Var} (f(X^{\theta+\epsilon}(t)) - f(X^\theta(t))) \leq C_{t,f,M} \epsilon$$

for some $C_{t,f,M} > 0$ depending upon t , f , and M (the number of reactions).

THEOREM 3.1. *Suppose $(X^{\theta+\epsilon}, X^\theta)$ satisfy (7) with our running assumptions (Assumptions 1 and 2). Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be a C^1 function with bounded first derivative on all $x \in \mathcal{S}$. Then, for any $T > 0$ there is a $C_{T,f,M} > 0$ for which*

$$\mathbb{E} \sup_{t \leq T} (f(X^{\theta+\epsilon}(t)) - f(X^\theta(t)))^2 \leq C_{T,f,M} \epsilon.$$

We provide two lemmas, giving the L^1 and L^2 bounds on the difference between $X^{\theta+\epsilon}$ and X^θ , before proving Theorem 3.1.

LEMMA 3.2. *Suppose $(X^{\theta+\epsilon}, X^\theta)$ satisfy (7) with our running assumptions (Assumptions 1 and 2). Then, for $T > 0$ there is a $C_{T,M} > 0$ for which*

$$\mathbb{E} \sup_{t \leq T} |X^{\theta+\epsilon}(t) - X^\theta(t)| \leq C_{T,M} \epsilon.$$

Proof. Let $T > 0$. For any $s \geq 0$,

$$(16) \quad X^{\theta+\epsilon}(s) - X^\theta(s) = M^{\theta,\epsilon}(s) + \int_0^s F^{\theta+\epsilon}(X^{\theta+\epsilon}(r)) - F^\theta(X^\theta(r)) dr,$$

where $M^{\theta,\epsilon}$ is a martingale with quadratic covariation

$$[M^{\theta,\epsilon}]_t = \sum_k \left(N_{k,2}^{\theta,\epsilon}(t) + N_{k,3}^{\theta,\epsilon}(t) \right) \zeta_k \zeta_k^T,$$

where

$$N_{k,2}^{\theta,\epsilon}(t) \stackrel{\text{def}}{=} Y_{k,2} \left(\int_0^t \lambda_k^{\theta+\epsilon}(X^{\theta+\epsilon}(s)) - \lambda_k^{\theta+\epsilon}(X^{\theta+\epsilon}(s)) \wedge \lambda_k^\theta(X^\theta(s)) ds \right),$$

$$N_{k,3}^{\theta,\epsilon}(t) \stackrel{\text{def}}{=} Y_{k,3} \left(\int_0^t \lambda_k^\theta(X^\theta(s)) - \lambda_k^{\theta+\epsilon}(X^{\theta+\epsilon}(s)) \wedge \lambda_k^\theta(X^\theta(s)) ds \right).$$

Therefore, for $s \leq t$,

$$\begin{aligned} |X^{\theta+\epsilon}(s) - X^\theta(s)| &\leq \sup_{r \leq t} |M^{\theta,\epsilon}(r)| + \int_0^s |F^{\theta+\epsilon}(X^{\theta+\epsilon}(r)) - F^\theta(X^{\theta+\epsilon}(r))| dr \\ &\quad + \int_0^s |F^\theta(X^{\theta+\epsilon}(r)) - F^\theta(X^\theta(r))| dr \\ &\leq \sup_{r \leq t} |M^{\theta,\epsilon}(r)| + K_2 \epsilon s + K_1 \int_0^s |X^{\theta+\epsilon}(r) - X^\theta(r)| dr \\ &\leq \sup_{r \leq t} |M^{\theta,\epsilon}(r)| + K_2 \epsilon t + K_1 \int_0^t \sup_{u \leq r} |X^{\theta+\epsilon}(u) - X^\theta(u)| dr, \end{aligned}$$

where K_1, K_2 are the constants of Assumptions 1 and 2, respectively. As the above inequality holds for all $s \leq t$, we have that

$$(17) \quad \sup_{s \leq t} |X^{\theta+\epsilon}(s) - X^\theta(s)| \leq \sup_{r \leq t} |M^{\theta,\epsilon}(r)| + K_2 \epsilon t + K_1 \int_0^t \sup_{u \leq r} |X^{\theta+\epsilon}(u) - X^\theta(u)| dr.$$

By the Burkholder–Davis–Gundy inequality and the fact that $\sqrt{z} \leq z$ for all nonnegative integers, we have the existence of a $C_2 > 0$ for which

(18)

$$\begin{aligned} \mathbb{E} \sup_{r \leq t} |M^{\theta, \epsilon}(r)| &\leq C_2 \sum_k \int_0^t \mathbb{E} |\lambda_k^{\theta+\epsilon}(X^{\theta+\epsilon}(u)) - \lambda_k^\theta(X^\theta(u))| du \\ &\leq C_2 \sum_k \left[\int_0^t \mathbb{E} |\lambda_k^{\theta+\epsilon}(X^{\theta+\epsilon}(u)) - \lambda_k^\theta(X^{\theta+\epsilon}(u))| du \right. \\ &\quad \left. + \int_0^t \mathbb{E} |\lambda_k^\theta(X^{\theta+\epsilon}(u)) - \lambda_k^\theta(X^\theta(u))| du \right] \\ &\leq C_3 \epsilon t + C_4 \int_0^t \mathbb{E} |X^{\theta+\epsilon}(u) - X^\theta(u)| du \\ &\leq C_3 \epsilon t + C_4 \int_0^t \mathbb{E} \sup_{u \leq r} |X^{\theta+\epsilon}(u) - X^\theta(u)| dr, \end{aligned}$$

where C_3 and C_4 are constants independent of ϵ or T and depend linearly on M (the number of reactions). Taking expectations of (17), applying (18), and using Gronwall's inequality gives the desired result. \square

LEMMA 3.3. *Suppose $(X^{\theta+\epsilon}, X^\theta)$ satisfy (7) with our running assumptions (Assumptions 1 and 2). Then, for $T > 0$ there is a $C_{T,M} > 0$ for which*

$$\mathbb{E} \sup_{t \leq T} |X^{\theta+\epsilon}(t) - X^\theta(t)|^2 \leq C_{T,M} \epsilon.$$

Proof. Returning to (16) in the proof of Lemma 3.2, we have that

$$|X^{\theta+\epsilon}(s) - X^\theta(s)|^2 \leq 2|M^{\theta, \epsilon}(s)|^2 + 2T \int_0^s |F^{\theta+\epsilon}(X^{\theta+\epsilon}(r)) - F^\theta(X^\theta(r))|^2 dr.$$

The proof is now essentially the same as that of Lemma 3.2, though (18) in combination with Lemma 3.2 is used to bound the martingale term. \square

Note that the expected difference of the p th moment, for any $p \geq 1$, can be estimated in the same manner as above.

Example 3. To demonstrate the above lemmas, we consider the following simple example:



where the parameter of interest is the rate constant θ . For this example we have that $\lambda^\theta(x) \equiv \theta$, and so

$$\begin{aligned} X^{\theta+\epsilon}(t) &= Y_1(\theta t) + Y_2(\epsilon t), \\ X^\theta(t) &= Y_1(\theta t). \end{aligned}$$

Hence $X^{\theta+\epsilon}(t) - X^\theta(t) = Y_2(\epsilon t)$, and the statements of Lemmas 3.2 and 3.3 follow immediately. Further, the lemmas are shown to be sharp.

Proof of Theorem 3.1. By Taylor’s theorem combined with our assumption on f , we have that for some $C_f > 0$,

$$(f(X^{\theta+\epsilon}(t)) - f(X^\theta(t)))^2 \leq C_f |X^{\theta+\epsilon}(t) - X^\theta(t)|^2,$$

and the result follows by application of Lemma 3.3. \square

We return now to (5) and (6) and note that with $X^{\theta+\epsilon/2}, X^{\theta-\epsilon/2}$ generated via the coupling proposed here,

$$\begin{aligned} \text{Var}(d_{[i]}(\epsilon)) &= O(\epsilon^{-1}), \\ \text{Var}(D_R(\epsilon)) &= O(R^{-1}\epsilon^{-1}), \end{aligned}$$

which are an order of magnitude lower, in terms of ϵ , than the respective variances when the processes are generated independently. As discussed at the end of section 2.1, this fact leads to a decrease in the computational work (and simulation time) required to solve a given problem to a desired tolerance by a factor of ϵ , yielding potentially enormous savings.

4. Numerical examples. We compare our method with existing methods on three different models: a basic model for the production of mRNA and proteins, an $M/M/\infty$ queue, and a genetic toggle switch from [18]. Because the common reaction path (CRP) method of [18] tends to perform at least as well as the usual implementation of CRN with Gillespie’s algorithm, we choose to include only the CRP method in our comparison (except for one plot in Numerical Example 2 to demonstrate the decoupling alluded to at the end of section 3.1).

Numerical Example 1. Consider the model of gene transcription and translation



where a single gene is being translated into mRNA, which is then being transcribed into proteins. The final two reactions represent degradation of the mRNA and protein molecules, respectively. Assuming that there is a single gene copy, the stochastic equation (1) for this model is

$$(20) \quad \begin{aligned} X^\theta(t) = X^\theta(0) &+ Y_1(2t) \begin{pmatrix} 1 \\ 0 \end{pmatrix} + Y_2 \left(\int_0^t 10X_1^\theta(s)ds \right) \begin{pmatrix} 0 \\ 1 \end{pmatrix} \\ &+ Y_3 \left(\int_0^t \theta X_1^\theta(s)ds \right) \begin{pmatrix} -1 \\ 0 \end{pmatrix} \\ &+ Y_4 \left(\int_0^t X_2^\theta(s)ds \right) \begin{pmatrix} 0 \\ -1 \end{pmatrix}, \end{aligned}$$

where $X_1^\theta(t)$ and $X_2^\theta(t)$ give the number of mRNA and protein molecules at time t , respectively, and Y_1, Y_2 are independent unit-rate Poisson processes. Suppose the rate constant θ is of interest to us and we believe that $\theta \approx 1/4$. We would like to estimate the sensitivity of the mean number of protein molecules at time $T = 30$, say, with respect to the parameter $\theta \approx 1/4$. Here, it is a straightforward calculation to find that

$$\mathbb{E}X_2^\theta(30) \Big|_{\theta=1/4} \approx 79.941 \quad \text{and} \quad \frac{d}{d\theta} \mathbb{E}X_2^\theta(30) \Big|_{\theta=1/4} \approx -318.073$$

TABLE 1

95% confidence intervals and computational complexity for (i) crude Monte Carlo (CMC), (ii) the common reaction path (CRP) method of [18], and (iii) the coupled finite difference (CFD) method proposed here, applied to (19) in order to approximate $J(1/4)$ for different choices of R and ϵ . The exact value is $J(1/4) = -318.073$. Note that the bias of the centered finite difference is apparent when $\epsilon = 1/20$ and $R = 40,000$.

Method	R	$\epsilon = 1/20$	$\epsilon = 1/100$	# updates	CPU time
CMC	1,000	-276.2 ± 46.3	$-472.7.1 \pm 237.3$	$\approx 8.4 \times 10^6$	≈ 9.6 S
CRP	1,000	-323.1 ± 18.4	-321.0 ± 60.2	$\approx 8.4 \times 10^6$	≈ 10.1 S
CFD	1,000	-323.7 ± 8.7	-333.8 ± 28.0	$\approx 4.4 \times 10^6$	≈ 6.5 S
CMC	10,000	-324.8 ± 14.7	-305.4 ± 74.3	$\approx 8.4 \times 10^7$	≈ 98.8 S
CRP	10,000	-325.5 ± 5.8	-328.6 ± 18.6	$\approx 8.4 \times 10^7$	≈ 105.4 S
CFD	10,000	-320.0 ± 2.8	-316.6 ± 8.9	$\approx 4.4 \times 10^7$	≈ 64.9 S
CMC	40,000	-322.7 ± 7.5	-341.9 ± 37.3	$\approx 3.4 \times 10^8$	≈ 395.3 S
CRP	40,000	-319.6 ± 2.9	-310.6 ± 9.3	$\approx 3.4 \times 10^8$	≈ 411.5 S
CFD	40,000	-321.6 ± 1.4	-317.8 ± 4.4	$\approx 1.8 \times 10^8$	≈ 263.3 S

if the initial condition is $X^\theta(0) = [0, 0]^T$. Defining

$$J(\theta) \stackrel{\text{def}}{=} \frac{d}{d\theta} \mathbb{E} [X_2^\theta(30)],$$

our goal is to efficiently estimate $J(1/4)$, and we compare the following methods on this problem:

- (i) the usual crude Monte Carlo (CMC) estimator with independent samples,
- (ii) the CRP method of [18] using the coupled processes (8),
- (iii) the coupled finite difference (CFD) method being proposed in this paper using the coupling (7), and
- (iv) a Girsanov transformation method of Plyasunov and Arkin detailed in [17].

For all simulations, we assume an initial condition of zero mRNA and zero protein molecules. We will denote the number of sample paths used in the construction of the relevant estimators (5) via R and the perturbation in the centered finite difference via ϵ .

In Table 1, we provide the 95% confidence intervals computed using the CMC method with independent paths, the CRP method, and the CFD method for different choices of R , the number of paths simulated, and ϵ , the perturbation of θ . For each method and choice of R , we also provide (i) an approximate total number of steps (and random numbers used) over the course of the entire simulation, and (ii) an approximation of the CPU time required. These numbers, which quantify the computational work required from each method, are essentially independent of ϵ , and the numbers provided are the average of the actual values for the two different values of ϵ for a given method and choice of R . These numbers should be used as a reference for the computational complexity required by the different methods with the understanding that CPU time will depend greatly upon implementation (the author used MATLAB for all computations, which were performed on an Apple machine with a 2.2 GHz Intel i7 processor). In Table 2, we provide similar data for the Girsanov transformation method of [17].

While Tables 1 and 2 demonstrate that the method being proposed here can produce a more accurate estimate in less CPU time than the other methods, a more

TABLE 2

95% confidence intervals and computational complexity for the Girsanov transformation method of [17] applied to (19) in order to approximate $J(1/4)$. The exact value is $J(1/4) = -318.073$.

R	Approximation	# updates	CPU time
1,000	-441.5 ± 156.5	4.2×10^6	5.2 S
10,000	-324.2 ± 49.2	4.2×10^7	54.5 S
40,000	-327.8 ± 25.1	1.7×10^8	207.6 S

TABLE 3

Required R , # updates, and CPU time for each method to provide a 95% confidence of ± 6.0 . Each finite difference method used $\epsilon = 1/40$. The exact value is $J(1/4) = -318.073$.

Method	R	Approximation	# updates	CPU time
Girsanov	689,600	-312.1 ± 6.0	2.9×10^9	3,506.6 S
CMC	246,000	-319.3 ± 6.0	2.1×10^9	2,364.8 S
CRP	25,980	-316.7 ± 6.0	2.2×10^8	270.9 S
CFD	4,580	-319.9 ± 6.0	2.0×10^7	29.2 S

important statistic is the CPU time needed for each method to achieve a desired tolerance. Therefore, we applied each method until the 95% confidence interval was ± 6.0 . The finite difference methods were applied with a perturbation size of $\epsilon = 1/40$. Table 3 provides the number of updates needed by each method combined with the CPU time needed on our machine. We see that the CFD method was approximately 9 times more efficient than the CRP method, and vastly more efficient than both the Girsanov transformation method and the CMC method with independent samples.

Next, we simulated the system (20) 5,000 times using each of the different methods and plotted the variance of the estimators versus time up to $T = 60$; see Figure 1. The finite difference methods were computed with a perturbation of size $\epsilon = 1/40$. We note that the variance of each of the finite difference methods appears to converge, though the limiting value for the CFD method being proposed here converges to a value that is approximately 6.5 times lower than that of the CRP method and 52 times lower than the CMC method. Also note that the variance of the Girsanov transform method grows linearly in time, as expected, and is quite large for even moderate values of time, t .

Numerical Example 2. We revisit Example 1, which modeled an mRNA molecule being created and degraded (or an $M/M/\infty$ queue),

$$(21) \quad \emptyset \xrightleftharpoons[0.1]{\theta} M.$$

We suppose that we want to understand the sensitivity of the expected number of particles (or customers in the queuing setting) with respect to the parameter $\theta \approx 2$. In Figure 2 we provide a plot of the variances of the different estimators as functions of time. So as to demonstrate the different behaviors of the different estimators, the scales on both the time and variance axes are dramatically different for the different methods. For each of the methods, we chose $R = 1,000$, and we used $\epsilon = 1/100$ for the perturbation methods. Recall that in Example 1, we proved that the variance of the difference between $X^{\theta+\epsilon}$ and X^θ will converge to $\epsilon/0.1$ if they are coupled using (7).

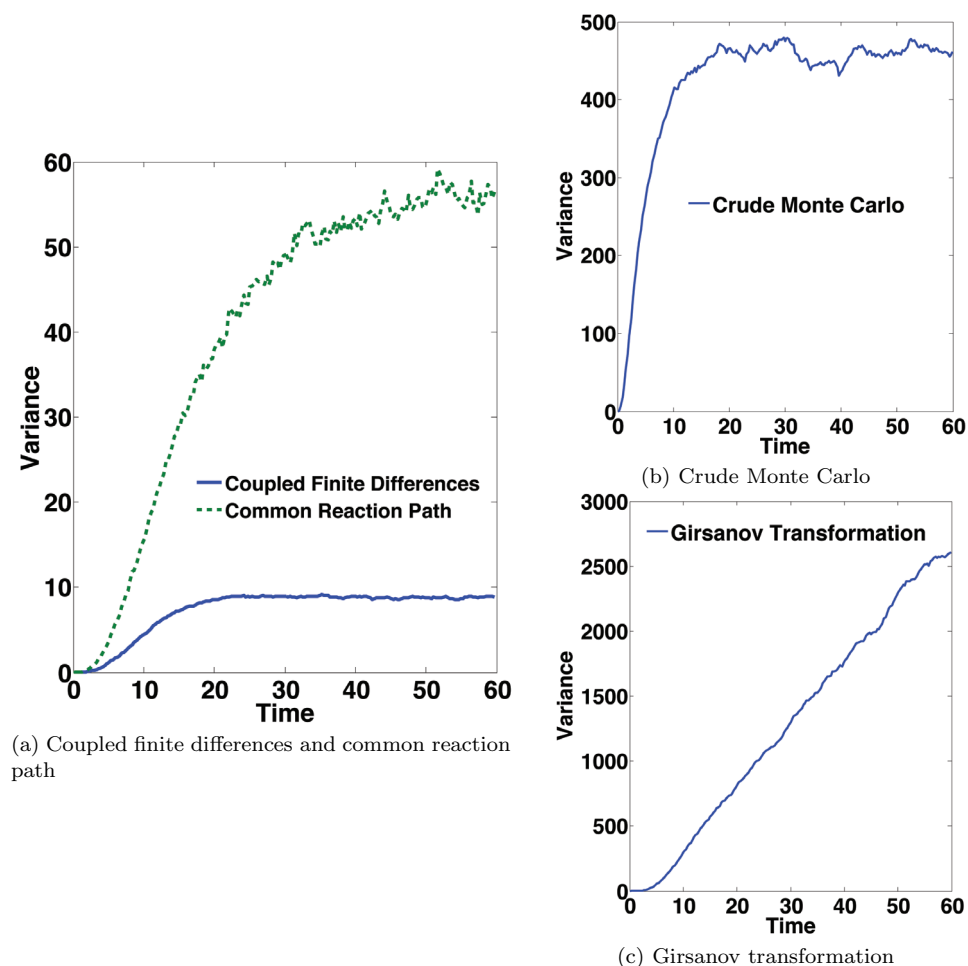


FIG. 1. Variance of the different estimators applied to (20). For each, $R = 5,000$ sample paths were used to construct the relevant estimators. For each of the finite difference methods (Figures 1(a) and 1(b)), a perturbation of $\epsilon = 1/40$ was used. Note that the scales on the variance axis are dramatically different for the different methods.

Therefore, the variance of the estimator (5) will converge to

$$\frac{\epsilon}{0.1} \frac{1}{\epsilon^2} \frac{1}{R} = \frac{1}{R} \frac{1}{0.1\epsilon}$$

as $t \rightarrow \infty$. In our case, $\epsilon = 1/100$ and $R = 1,000$, and the above value is equal to one. This predicted behavior is born out in Figure 2(e). Also in Example 1, we predicted (though did not prove) that after a long enough time the variance of both the CRP estimator and Gillespie's algorithm plus CRN should converge to the variance of the CMC estimator constructed with independent paths. In essence, we are predicting that the processes will decouple after a long enough time and behave independently. This behavior is demonstrated in Figures 2(a) and 2(c) (for CRP) and 2(f) (for Gillespie + CRN), though we note that the time for a full decoupling is quite large in this example. Also note that we plotted the variance of the estimator

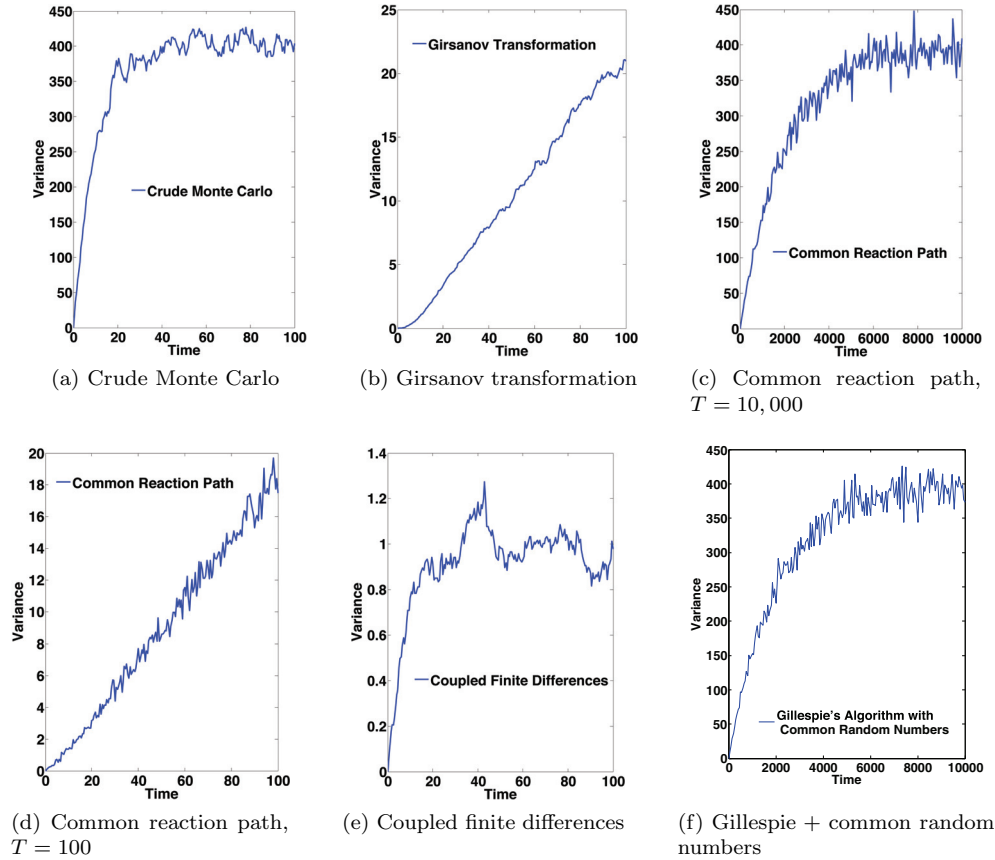


FIG. 2. Variance of the different estimators applied to (21). $R = 1,000$ sample paths were used to construct the relevant estimators. For each of the finite difference methods (plots 2(a), 2(c), 2(d), 2(e), and 2(f)), a perturbation of $\epsilon = 1/100$ was used. So as to demonstrate the different behaviors of the different estimators, the scales both on the time and variance axes are dramatically different for the different methods. Also, note that we plotted the variance of the estimator for the CRP method up to both times $T = 100$ and $T = 10,000$ so as to demonstrate the different behaviors exhibited. Note that the CRP and Gillespie + CRN methods appear equivalent for this example.

for the CRP method up to both times $T = 100$ and $T = 10,000$ so as to demonstrate the different behaviors exhibited. Finally, we point out that the full “decoupling” of the CRP method described here does not seem to take place in Example 1. The estimator built using the Girsanov transformation method exhibits a variance that grows linearly in time.

Next, we considered the sensitivity to the decay parameter at 0.1. That is, we considered



with $\theta = 0.1$. In Figure 3 we provide a plot of the variance of the CFD estimator versus the CRP estimator. Each plot was generated using 1,000 sample paths in which a perturbation of $\epsilon = 1/100$ was used. We again see the lower variance exhibited by the CFD estimator, though the difference is now less dramatic.




FIG. 3. Time plot of the variance of the CFD estimator versus the CRP estimator for the model (22) with decay rate perturbed. Each plot was generated using 1,000 sample paths. A perturbation of $\epsilon = 1/100$ was used.

Numerical Example 3. We consider a model for a genetic toggle switch found in [18]:



with intensity functions

$$\begin{aligned} \lambda_1(X(t)) &= \frac{\alpha_1}{1 + X_2(t)^\beta}, & \lambda_2(X(t)) &= X_1(t), \\ \lambda_3(X(t)) &= \frac{\alpha_2}{1 + X_1(t)^\gamma}, & \lambda_4(X(t)) &= X_2(t) \end{aligned}$$

and parameter choices

$$\alpha_1 = 50, \quad \alpha_2 = 16, \quad \beta = 2.5, \quad \gamma = 1.$$

We begin the process with initial condition $[0, 0]$ and consider the sensitivity of X_1 as a function of α_1 . In Figure 4 we provide a plot of the variance of the CFD estimator versus the CRP estimator as a function of time. Each plot was generated using 10,000 sample paths in which a perturbation of $\epsilon = 1/10$ was used. We see that the CFD method performs substantially better for times $t > 5$, whereas the CRP method performs better for shorter times, $t < 5$.




FIG. 4. Time plot of the variance of the CFD estimator versus the CRP estimator for the model (23). Each plot was generated using 10,000 sample paths. A perturbation of $\epsilon = 1/10$ was used. It is worth noting that here the CRP estimator outperforms the CFD estimator for times $t < 5$, though CFD still greatly outperforms CRP for larger times.

REFERENCES

- [1] D. F. ANDERSON, *A modified next reaction method for simulating chemical systems with time dependent propensities and delays*, J. Chem. Phys., 127 (2007), 214107.
- [2] D. F. ANDERSON, A. GANGULY, AND T. G. KURTZ, *Error analysis of tau-leap simulation methods*, Ann. Appl. Probab., 21 (2011), pp. 2226–2262.
- [3] D. F. ANDERSON AND D. J. HIGHAM, *Multilevel Monte Carlo for continuous time Markov chains, with applications in biochemical kinetics*, Multiscale Model. Simul., 10 (2012), pp. 146–179.
- [4] D. F. ANDERSON AND M. KOYAMA, *Weak error analysis of numerical methods for stochastic models of population processes*, Multiscale Model. Simul., to appear.
- [5] D. F. ANDERSON AND T. G. KURTZ, *Continuous time Markov chain models for chemical reaction networks*, in Design and Analysis of Biomolecular Circuits: Engineering Approaches to Systems and Synthetic Biology, H. Koepl, D. Densmore, G. Setti, and M. di Bernardo, eds., Springer, New York, 2011, pp. 3–42.
- [6] S. ASMUSSEN AND P. W. GLYNN, *Stochastic Simulation: Algorithms and Analysis*, Springer, New York, 2007.
- [7] S. N. ETHIER AND T. G. KURTZ, *Markov Processes: Characterization and Convergence*, John Wiley & Sons, New York, 1986.
- [8] M. GIBSON AND J. BRUCK, *Efficient exact stochastic simulation of chemical systems with many species and many channels*, J. Phys. Chem. A, 105 (2000), pp. 1876–1889.
- [9] M. HUTZENTHALER AND A. JENTZEN, *Convergence of the stochastic Euler scheme for locally Lipschitz coefficients*, Found. Comput. Math., 11 (2011), pp. 657–706.
- [10] H.-W. KANG AND T. G. KURTZ, *Separation of time-scales and model reduction for stochastic reaction networks*, Ann. Appl. Probab., to appear.
- [11] M. KOMOROWSKI, M. J. COSTA, D. A. RAND, AND M. P. H. STUMPF, *Sensitivity, robustness, and identifiability in stochastic chemical kinetics models*, Proc. Natl. Acad. Sci. USA, 108 (2011), pp. 8645–8650.
- [12] T. G. KURTZ, *The relationship between stochastic and deterministic models for chemical reactions*, J. Chem. Phys., 57 (1972), pp. 2976–2978.
- [13] T. G. KURTZ, *Approximation of Population Processes*, CBMS-NSF Regional Conf. Ser. in Appl. Math. 36, SIAM, Philadelphia, 1981.

- [14] T. G. KURTZ, *Representation and approximation of counting processes*, in *Advances in Filtering and Optimal Stochastic Control*, Lecture Notes in Control and Inform. Sci. 42, Springer, Berlin, 1982, pp. 177–191.
- [15] H. LAMBDA, J. C. MATTINGLY, AND A. M. STUART, *An adaptive Euler-Maruyama scheme for SDEs: Convergence and stability*, *IMA J. Numer. Anal.*, 27 (2007), pp. 479–506.
- [16] D. A. MCQUARRIE, *Stochastic approach to chemical kinetics*, *J. Appl. Probability*, 4 (1967), pp. 413–478.
- [17] S. PLYASUNOV AND A. P. ARKIN, *Efficient stochastic sensitivity analysis of discrete event systems*, *J. Comput. Phys.*, 221 (2007), pp. 724–738.
- [18] M. RATHINAM, P. W. SHEPPARD, AND M. KHAMMASH, *Efficient computation of parameter sensitivities of discrete stochastic chemical reaction networks*, *J. Chem. Phys.*, 132 (2010), 034103.