



Citation for published version:

Lester, C, Yates, CA & Baker, RE 2018, 'Robustly simulating biochemical reaction kinetics using multi-level Monte Carlo approaches', *Journal of Computational Physics*, vol. 375, pp. 1401-1423.
<https://doi.org/10.1016/j.jcp.2018.06.045>

DOI:

[10.1016/j.jcp.2018.06.045](https://doi.org/10.1016/j.jcp.2018.06.045)

Publication date:

2018

Document Version

Peer reviewed version

[Link to publication](#)

Publisher Rights

CC BY-NC-ND

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Robustly simulating biochemical reaction kinetics using multi-level Monte Carlo approaches

Christopher Lester[†], Christian A. Yates^{*}, Ruth E. Baker[†]

[†]*Mathematical Institute, Woodstock Road, Oxford, OX2 6GG, UK* and ^{*}*Department for Mathematical Sciences, Claverton Down, Bath, BA2 7AY, UK.*

Abstract

In this work, we consider the problem of estimating summary statistics to characterise biochemical reaction networks of interest. Such networks are often described using the framework of the Chemical Master Equation (CME). For physically-realistic models, the CME is widely considered to be analytically intractable. A variety of Monte Carlo algorithms have therefore been developed to explore the dynamics of such networks empirically. Amongst them is the multi-level method, which uses estimates from multiple ensembles of sample paths of different accuracies to estimate a summary statistic of interest. In this work, we develop the multi-level method in two directions: (1) to increase the robustness, reliability and performance of the multi-level method, we suggest an improved variance reduction method for generating the sample paths of each ensemble; and (2) to improve computational performance, we provide a different mechanism for choosing which ensembles should be included in the multi-level algorithm.

Keywords: Biochemical reaction networks, Stochastic simulation, Multi-level Monte Carlo approaches, Variance reduction.

1. Introduction

Stochastic approaches are commonly used to model a wide variety of biological processes [1, 2, 3]. This work is concerned with individual-based models that represent the behaviour of single particles or organisms within a biological system of interest. We use discrete reaction events to describe the interactions between individual particles; the time evolution of the system is described by the Chemical Master Equation (CME) [4]. Except for a small number of special cases, the CME is both analytically and numerically intractable [5], and Monte Carlo simulation is often used to explore the model dynamics instead. Typically, a number of sample paths are generated, and used to estimate relevant summary statistics.

The multi-level method was described for stochastic differential equations by Giles [6], and in 2012, Anderson and Higham [7] developed an approach suitable for biochemical reaction networks. The multi-level method produces point estimators in a cost-effective manner by combining a

hierarchy of tau-leap estimators of differing accuracy in a telescoping sum. When compared with traditional simulation methods, the multi-level method has been shown to be capable of reducing the CPU time taken to estimate summary statistics by orders of magnitude [7]. A wide body of subsequent research has focussed on generalising the multi-level method used by Anderson and Higham [8, 9, 10]. A number of general performance improvements [11] have also been described, but current, state-of-the-art implementations of the multi-level method still rely heavily on user input and configuration. Additionally, a number of unusual implementation issues remain unresolved [11, 9].

This work has two aims. Firstly, to improve the efficiency and reliability of the multi-level method, we change the procedure by which the terms in the aforementioned telescoping sum are computed. Secondly, we provide further improvements in computational performance by including R-leap estimators, instead of tau-leap estimators, in the telescoping sum.

This work is arranged as follows: in Section 2 we review background material, so that the multi-level method can be described in Section 3. In Sections 4 and 5 we develop a new and reliable implementation of the multi-level method (the first aim of our work). Our new simulation method is compared numerically with a traditional approach in Section 6. A new multi-level framework is then set out and tested numerically in Section 7 (our second aim). Conclusions are then drawn in Section 8.

2. Stochastic biochemical networks

We study a biochemical network comprising N species, S_1, \dots, S_N , that may interact through M reaction channels, R_1, \dots, R_M . In this work, we will assume that the system of interest is well-stirred, and we describe the dynamics of the biochemical network using the CME [5]. At time t , the population, or copy number, of species S_i is denoted by $X_i(t)$, and the state vector, $\mathbf{X}(t)$, is given by

$$\mathbf{X}(t) := [X_1(t), \dots, X_N(t)]^T. \quad (1)$$

We associate two quantities with each reaction channel R_j . The first is the stoichiometric or state-change vector,

$$\boldsymbol{\nu}_j := [\nu_{1j}, \dots, \nu_{Nj}]^T, \quad (2)$$

where ν_{ij} is the change in the copy number of S_i caused by reaction R_j taking place. The second quantity is the propensity function, $p_j(\mathbf{X}(t))$. For infinitesimally small dt , the rate $p_j(\mathbf{X}(t))$ is defined as follows:

$$p_j(\mathbf{X}(t))dt := \mathbb{P}[R_j \text{ occurs in } [t, t + dt)].$$

Since the system is well-stirred, the reaction activity can, for example, be modelled with mass action kinetics. Where mass action kinetics are used, the propensity function of reaction R_j , p_j , is proportional to the number of possible combinations of reactant particles in the system [5].

2.1. The random time change representation

In order to analyse and understand the CME framework mathematically, the random time change representation (RTCR) is now set out. The RTCR is a path-wise representation of the *same* stochastic process as the CME, and was first described by Kurtz [12]. The RTCR describes the dynamics of our chosen biochemical network by using a set of inhomogeneous Poisson processes. We first describe the inhomogeneous Poisson process framework, and then provide the RTCR.

Consider a (homogeneous) Poisson process of fixed rate λ , labelled as \mathcal{Y}^λ . Then further suppose we have a Poisson process of unit rate, \mathcal{Y}^1 . As Poisson processes count the number of ‘arrivals’ over time, they can be compared by considering the distribution of the number of arrivals by some time t . If $\mathcal{Y}^\lambda(0, t)$ and $\mathcal{Y}^1(0, t)$ represent the number of arrivals over the time interval $(0, t]$ in each of our Poisson processes, then there is an equality in distribution, that is $\mathcal{Y}^\lambda(0, t) \sim \mathcal{Y}^1(0, \lambda \cdot t)$. It is therefore possible to re-scale time to transform a unit rate Poisson process to one of arbitrary (but known) rate. We will now allow λ to change its value. The value of λ can change with time, or, alternatively, depending on the arrivals of the Poisson process at earlier times. The inhomogeneous Poisson process formalism allows the rate parameter, λ , to change dynamically.

The number of times reaction R_j (for $j = 1, \dots, M$) takes place (‘fires’) over the time interval $(0, T]$ is given by a inhomogeneous Poisson counting process

$$\mathcal{Y}_j \left(0, \int_0^T p_j(\mathbf{X}(t)) dt \right),$$

where \mathcal{Y}_j is a unit-rate Poisson process, and $\mathcal{Y}_j(\alpha, \beta)$ is defined as

$$\mathcal{Y}_j(\alpha, \beta) := \# \text{ of arrivals in } (\alpha, \beta]. \quad (3)$$

Every time reaction R_j occurs, the state vector (see Equation (1)), is updated by adding the appropriate stoichiometric vector (see Equation (2)) to it. Therefore, by considering all possible reactions over the time interval $(0, T]$, we can determine the state vector at time T as

$$\mathbf{X}(T) = \mathbf{X}(0) + \sum_{j=1}^M \mathcal{Y}_j \left(0, \int_0^T p_j(\mathbf{X}(t)) dt \right) \cdot \nu_j. \quad (4)$$

As highlighted in the introduction, except for understanding very simple reaction networks¹, Monte Carlo simulation is commonly used to understand the system’s features and properties [5]. As this work progresses, the RTCR will be used to explain and develop new Monte Carlo methods.

We start by discussing three methods for generating suitable sample paths: the direct, tau-leap and R-leap methods. Typically, these methods are used to generate an ensemble of \mathcal{N} sample paths, and the sample paths are then used to estimate summary statistics that characterise the system.

2.2. Gillespie Direct Method

A variety of methods can be used to generate the sample paths, but the Gillespie Direct Method (DM) is the most widely-used algorithm [5, 13, 14]. The DM is a serial algorithm in the sense that sample paths are generated by simulating individual reactions in the order in which they take place. A key benefit of the DM is that it can be implemented in a straightforward manner, as shown in Algorithm 1. However, because every reaction that takes place in a sample path is simulated individually, the DM requires a relatively high level of computational resources.

Algorithm 1 The DM. This simulates a single sample path.

Require: initial conditions, $\mathbf{X}(0)$, and terminal time, T .

```

1: set  $\mathbf{X} \leftarrow \mathbf{X}(0)$  and set  $t \leftarrow 0$ 
2: loop
3:   for each  $R_j$ , calculate propensity values  $p_j$  and set  $p_0 \leftarrow \sum_{j=1}^M p_j$ 
4:   set  $\Delta \leftarrow \text{Exp}(p_0)$ 
5:   if  $t + \Delta > T$  then
6:     break
7:   end if
8:   choose reaction  $R_k$  to fire next:  $R_k$  fires with probability  $p_k/p_0$ 
9:   set  $\mathbf{X} \leftarrow \mathbf{X} + \boldsymbol{\nu}_k$ , and set  $t \leftarrow t + \Delta$ 
10: end loop

```

The summary statistics estimated using Monte Carlo simulation contain a statistical error. This arises as we have generated only a subset of the possible sample paths: with a different ensemble of realisations, the estimate, $\hat{\mathcal{Q}}$, will be slightly different, and therefore estimates have an inherent uncertainty to them. More precisely, if the sample variance of the chosen summary

¹For example, where the propensity functions are either constant or linear functions of the reactant species, analytic solutions are available.

statistic is \mathcal{V} , and \mathcal{N} sample values have been used to estimate the summary statistic, then the *estimator variance* is given by \mathcal{V}/\mathcal{N} . The estimator variance can be used to construct a confidence interval [11].

Approximate simulation methods can be implemented to reduce the computational resources required for Monte Carlo simulation. In short, the dynamics of a chosen biochemical reaction network are simplified so that fewer computational resources are required to simulate each sample path. Examples of such approximation schemes include the widely-used tau-leap [15], and R-leap [16] methods. We discuss both the tau- and R-leap methods in turn.

2.3. The tau-leap method

The tau-leap method generates approximate sample paths efficiently by taking time-steps of length τ through time, and firing multiple reactions during each time-step. We will work with a fixed choice of time-step, τ . During each time-step, the key assumption we make is that the reaction propensities are constant. It can then be shown that the number of times that reaction channel R_j fires during each time-step is a Poisson *random variate*, with known parameter [15]. By repeatedly advancing by a time-step of length τ through time², a sample path can be generated. An implementation of the tau-leap method is provided in Algorithm 2: to emphasise that the dynamics of a tau-leap sample path are different to the dynamics of the process \mathbf{X} , in our algorithm we denote the population of a tau-leap sample path as \mathbf{Z} .

Algorithm 2 The tau-leap method. This simulates a sample path using fixed time-step τ .

Require: initial conditions, $\mathbf{Z}(0)$, time-step τ , and terminal time, T .

```

1: set  $\mathbf{Z} \leftarrow \mathbf{Z}(0)$  and set  $t \leftarrow 0$ 
2: while  $t < T$  do
3:   for each  $R_j$  do
4:     calculate propensity value  $p_j(\mathbf{Z})$ 
5:     generate  $K_j \sim \mathcal{P}(p_j(\mathbf{Z}) \cdot \tau)$ 
6:   end for
7:   set  $\mathbf{Z} \leftarrow \mathbf{Z} + \sum_{j=1}^M K_j \cdot \boldsymbol{\nu}_j$ 
8:   set  $t \leftarrow t + \tau$ 
9: end while

```

Under reasonably general circumstances, it can be shown that the tau-leap method can generate sample paths more quickly than the DM [5]. However, if we choose to estimate a summary

²Here, and throughout the rest of the manuscript, we will assume that τ divides T exactly.

statistic with an ensemble of sample paths generated using the tau-leap method, then the resultant estimate will be biased. The bias is a consequence of using approximate reaction propensities in Algorithm 2; the bias typically scales as $\mathcal{O}(\tau)$ [17]. Roughly speaking, the total CPU time scales with the number of time-steps, i.e. $\mathcal{O}(1/\tau)$. Thus, τ must be chosen to balance the competing demands of speed and accuracy. Having explained how one might implement the tau-leap method, we turn to discussing its mathematical representation.

2.4. Representing the tau-leap method using Poisson processes

In this section, we use the RTCR to motivate and represent the tau-leap method. We recall that the state \mathbf{X} of our biochemical reaction network evolves according to Equation (4), which we now restate:

$$\mathbf{X}(T) = \mathbf{X}(0) + \sum_{j=1}^M \mathcal{Y}_j \left(0, \int_0^T p_j(\mathbf{X}(t)) dt \right) \cdot \boldsymbol{\nu}_j.$$

We continue to represent the state of a tau-leap process at time t as $\mathbf{Z}(t)$. Following from Algorithm 2, we suppose that $[0, T]$ is divided into K equal time-steps of length τ . The tau-leap assumption is that the propensities can change only at fixed times $t = k \cdot \tau$ (for $k = 1, 2, \dots$). Therefore, when we implement the tau-leap method we use the following approximation:

$$\int_0^T p_j(\mathbf{Z}(t)) dt \approx \sum_{k=0}^{K-1} p_j(\mathbf{Z}(\tau \cdot k)) \cdot \tau. \quad (5)$$

If we insert assumption (5) into Equation (4), then the evolution of the state of the tau-leap process, $\mathbf{Z}(t)$, is described by

$$\mathbf{Z}(T) = \mathbf{Z}(0) + \sum_{j=1}^M \mathcal{Y}_j \left(0, \sum_{k=0}^{K-1} p_j(\mathbf{Z}(\tau \cdot k)) \cdot \tau \right) \cdot \boldsymbol{\nu}_j, \quad (6)$$

where the \mathcal{Y}_j ($j = 1, \dots, M$) are unit-rate Poisson processes. We now rearrange Equation (6) so that it is easier to work with. First we set³

$$P_{k,j} = \sum_{k'=0}^k p_j(\mathbf{Z}(\tau \cdot k')) \cdot \tau, \quad (7)$$

with the special case of $P_{-1,j} = 0$. Then, we re-arrange Equation (6) to give

$$\mathbf{Z}(T) = \mathbf{Z}(0) + \sum_{k=0}^{K-1} \sum_{j=1}^M \mathcal{Y}_j(P_{k-1,j}, P_{k,j}) \cdot \boldsymbol{\nu}_j. \quad (8)$$

The tau-leap method can be seen as a method for iterating over k : for each reaction R_j , at each step we calculate the number of events in the Poisson process \mathcal{Y}_j between positions (or internal

³Note that j indexes the reaction, R_j ; and k indexes the time.

times) $P_{k-1,j}$ and $P_{k,j}$ (given by $\mathcal{Y}_j(P_{k-1,j}, P_{k,j})$). Equation (8) can therefore be re-arranged into an update formula. For $k = 1, \dots, K$,

$$\mathbf{Z}(k \cdot \tau) = \mathbf{Z}((k-1) \cdot \tau) + \sum_{j=1}^M \mathcal{Y}_j(P_{k-1,j}, P_{k,j}) \cdot \boldsymbol{\nu}_j. \quad (9)$$

2.5. The R-leap method

The R-leap method is an example of a simulation algorithm that generates approximate sample paths by firing a user-specified number of reactions during each step. The R-leap method was described by Auger et al. [16]. The R-leap method differs from the tau-leap method in that each ‘tau-leap’ covers a fixed time interval (with the number of reaction events determined by simulation), whereas the time interval covered by each ‘R-leap’ must be determined by simulating a random number (though the number of reaction events is fixed).

The algorithm proceeds as follows. The user specifies how many reactions will fire during each simulation step: this quantity is labelled as \mathcal{K} . The algorithm starts at time $t = 0$, and the state vector, \mathbf{Z} , is used to describe the molecular populations. At each step, the time for the \mathcal{K} reactions to take place, Δ , is generated. Then, the precise combination of reactions that take place (e.g. three R_1 reactions, four R_2 reactions, and so on) is chosen. Finally, the state vector, \mathbf{Z} , and time, t , are updated. These steps are repeated as required. We now discuss the details of the algorithm.

In the Gillespie DM (see Algorithm 1), the waiting time until the next reaction is given by an exponential variate with rate $p_0 = \sum_{j=1}^M p_j$. Consequently, given the fixed propensity values, the total waiting time for \mathcal{K} reactions to fire in the R-leap method algorithm, Δ , is given by the sum of \mathcal{K} exponential variates⁴. This sum is Gamma distributed, i.e. $\Delta \sim \Gamma(\mathcal{K}, 1/p_0)$, with shape parameter \mathcal{K} and scale parameter $1/p_0$.

Having calculated the distribution of the time period over which the next \mathcal{K} reaction events occur, we must decide on the specific combination of reaction types that take place. This can be achieved with a conditional binomial method that proceeds as follows. For each of the \mathcal{K} reaction events that we will simulate, the probability that it is a reaction of type R_j is given by p_j/p_0 . We start by considering R_1 . A binomial random number, $\mathcal{B}(\mathcal{K}, p_1/p_0)$ is simulated, using \mathcal{K} trials, each with probability p_1/p_0 of success. This provides the number of reactions of type R_1 ; we label this quantity as K_1 . There are therefore $\mathcal{K} - K_1$ reactions that still need to be assigned to a type, and we know that these reactions are *not* R_1 reactions. Thus, we generate a binomial random number, $K_2 \sim \mathcal{B}(\mathcal{K} - K_1, p_2/(p_0 - p_1))$ to decide how many reactions of type R_2 have fired. In

⁴As the propensities are only updated every \mathcal{K} reactions, p_0 is fixed for this many reactions.

this case, there are $\mathcal{K} - K_1$ trials, each with conditional probability $p_2/(p_0 - p_1)$ of success. The number of reactions of type R_3 is given by $K_3 \sim \mathcal{B}(\mathcal{K} - K_1 - K_2, p_3/(p_0 - p_1 - p_2))$. This process repeated until all \mathcal{K} reactions have been assigned a type. To ensure that the final step is carried out properly, so that the sample path terminates at time $t = T$, we require the following rule:

Rule 1. Conditioned on the value of Δ , where Δ is the waiting time for \mathcal{K} events to fire, the number of events that fire after traversing Δ' units of time, where $\Delta' < \Delta$, will be given by

$$\mathcal{K}' = \mathcal{B}(\mathcal{K} - 1, \Delta'/\Delta). \quad (10)$$

This follows as $\mathcal{K} - 1$ events take place strictly inside the interval of length Δ , with these events uniformly distributed throughout the interval.

A pseudo-code implementation of the R-leap method is provided in Algorithm 3. Note that, when \mathcal{K} decreases, the accuracy of the R-leap method increases, and when $\mathcal{K} = 1$, the DM is recovered. The computational cost scales as $\mathcal{O}(\mathcal{K}^{-1})$.

Algorithm 3 The R-leap method. This simulates a single sample path using fixed jump size, \mathcal{K} .

Require: initial conditions, $\mathbf{Z}(0)$, jump size, \mathcal{K} , and terminal time, T .

```

1: set  $\mathbf{Z} \leftarrow \mathbf{Z}(0)$  and set  $t \leftarrow 0$ 
2: while  $t < T$  do
3:   for each  $R_j$ , calculate propensity value  $p_j(\mathbf{Z})$ , and set  $p_0 \leftarrow \sum_{j=1}^M p_j$ 
4:   generate  $\Delta \sim \Gamma(\mathcal{K}, 1/p_0)$ 
5:   if  $t + \Delta > T$  then
6:     generate  $\mathcal{K} \sim \mathcal{B}(\mathcal{K} - 1, (T - t)/\Delta)$ , and set  $t \leftarrow T$ 
7:   else
8:     set  $t \leftarrow t + \Delta$ 
9:   end if
10:  for  $j = 1, \dots, M$  do
11:    generate  $K_j \sim \mathcal{B}(\mathcal{K}, p_j / \sum_{j'=j}^M p_{j'})$  and set  $\mathcal{K} \leftarrow \mathcal{K} - K_j$ 
12:  end for
13:  set  $\mathbf{Z} \leftarrow \mathbf{Z} + \sum_{j=1}^M K_j \cdot \boldsymbol{\nu}_j$ 
14: end while

```

2.6. Outlook

In the next section, we will introduce the multi-level method. Our investigation of the multi-level method is motivated by the dramatic computational savings it leads to: compared with traditional simulation methods, the multi-level method can potentially reduce the CPU time required to estimate a summary statistic by orders of magnitude [11].

3. Multi-level Monte Carlo

In 2008, Giles [6] described and implemented the multi-level framework; an efficient Monte Carlo scheme for stochastic differential equations. The multi-level method was subsequently extended to efficiently estimate summary statistics of discrete-state, continuous time Markov chains by Anderson and Higham [7]. The important contribution by Anderson and Higham [7] provides one way in which the multi-level technique can be used to model discrete-state systems, but there are many alternative ways in which the multi-level technique can be implemented. In this work, we will consider some of these.

In an effort to increase computational efficiency, the multi-level method divides the work done in calculating a summary statistic of interest into parts, known as *levels*. By way of example, the statistic we are interested in might be the average population of species S_i at time T . We write the chosen summary statistic as

$$\mathcal{Q} = \mathbb{E}[f(\mathbf{X})]. \quad (11)$$

Our aim is to estimate \mathcal{Q} to within a given statistical accuracy as efficiently as possible. For each level $\ell = 0, 1, \dots, L, L + 1$ (where L will be specified later), we will define an estimator \mathcal{Q}_ℓ such that

$$\mathcal{Q} = \sum_{\ell=0}^{L+1} \mathcal{Q}_\ell. \quad (12)$$

Each value of ℓ represents a level. We will independently estimate each \mathcal{Q}_ℓ , and then sum our estimates up, so that we have an estimated value of \mathcal{Q} . As outlined in Section 1, a confidence interval can be constructed around our estimated value of \mathcal{Q} . The size of the confidence interval depends on the estimator variance, with the estimator variance of \mathcal{Q} equal to the sum of the estimator variances of the estimates for \mathcal{Q}_ℓ , (for $\ell = 0, \dots, L + 1$). The key savings provided by the multi-level method arise because each of the estimates of $\mathcal{Q}_1, \dots, \mathcal{Q}_{L+1}$ can be calculated using a variance reduction technique. For each level of accuracy, a variance reduction technique cleverly lowers the number of sample paths required to estimate a statistic of interest, which means that, with a given level of accuracy, Monte Carlo simulation can be completed quickly. We explain the principle of variance reduction as follows:

Variance reduction. Suppose we wish to estimate a summary statistic of a stochastic process, ϕ , given by $\mathbb{E}[f(\phi)]$, where $f(\cdot)$ is a suitable function. Further, we suppose that estimating $\mathbb{E}[f(\phi)]$ is computationally intensive, because $f(\phi)$ has a relatively high sample variance, $\mathcal{V}_{f(\phi)}$. If we estimate $\mathbb{E}[f(\phi)]$ using \mathcal{N} sample paths, then the estimator variance is given by $\mathcal{V}_{f(\phi)}/\mathcal{N}$.

If there is a different stochastic process, ψ , and function, $g(\cdot)$, such that

$$\mathbb{E}[f(\phi)] = \mathbb{E}[g(\psi)], \quad (13)$$

then we can instead estimate $\mathbb{E}[g(\psi)]$, and any resultant estimate is also an estimate for $\mathbb{E}[f(\phi)]$. The bias of the estimate remains unchanged. However, the estimator variance of the estimate is now given by $\mathcal{V}_{g(\psi)}/\mathcal{N}$, where $\mathcal{V}_{g(\psi)}$ is the variance of $g(\psi)$. Thus, if $\mathcal{V}_{g(\psi)} < \mathcal{V}_{f(\phi)}$, fewer sample paths are required to achieve a given confidence in the estimator when $\mathbb{E}[g(\psi)]$, and not $\mathbb{E}[f(\phi)]$, is estimated. Therefore, the overall simulation time can be reduced by estimating $\mathbb{E}[g(\psi)]$ instead of $\mathbb{E}[f(\phi)]$.

Returning to the multi-level method: if the sample paths used to estimate \mathcal{Q}_ℓ have variance \mathcal{V}_ℓ , and there are \mathcal{N}_ℓ such sample paths, then estimator variance of \mathcal{Q}_ℓ is given by $\mathcal{V}_\ell/\mathcal{N}_\ell$. When we define \mathcal{Q}_ℓ below, we use the principle of variance reduction to choose \mathcal{Q}_ℓ so that \mathcal{V}_ℓ is small. This means that \mathcal{N}_ℓ does not need to be big, and the sample paths can be completed quickly.

Of course, the estimators $\mathcal{Q}_0, \dots, \mathcal{Q}_{L+1}$ can be produced in many different ways. In this section, we will use the tau-leap method to construct the aforementioned estimators. We now describe the calculations performed on each of levels $\ell = 0, \dots, L + 1$.

The base level. The first level, $\ell = 0$, is known as the base level. We will generate sample paths using the tau-leap method, with constant time-step τ_0 . The simulation cost of each sample path is denoted \mathcal{C}_0 . We will choose a large value for τ_0 , so that \mathcal{C}_0 is relatively small. We label the complete trajectory of each of the \mathcal{N}_0 sample paths we generate as $\mathbf{Z}_0^{(r)}$, where $r = 1, \dots, \mathcal{N}_0$. Thus, we estimate \mathcal{Q}_0 as

$$\mathcal{Q}_0 \approx \frac{1}{\mathcal{N}_0} \sum_{r=1}^{\mathcal{N}_0} f\left(\mathbf{Z}_0^{(r)}\right).$$

We expect that our estimate, \mathcal{Q}_0 , will be highly biased due to the large choice of τ_0 .

The intermediate correction levels. Subsequent, intermediate levels, indexed as $\ell = 1, \dots, L$, are known as the correction levels. We consider pairs of sample paths, $(\mathbf{Z}_\ell, \mathbf{Z}_{\ell-1})$, where \mathbf{Z}_ℓ is simulated using the tau-leap method with a time-step τ_ℓ , and $\mathbf{Z}_{\ell-1}$ is also simulated using the tau-leap method, but with a time-step $\tau_{\ell-1}$. We follow Anderson and Higham [7] in choosing an integer refinement factor, \mathcal{M} , so that $\tau_\ell = \tau_{\ell-1}/\mathcal{M}$ (i.e. τ_ℓ is \mathcal{M} times smaller than $\tau_{\ell-1}$). Then, our estimator for level ℓ is

$$\mathcal{Q}_\ell \approx \frac{1}{\mathcal{N}_\ell} \sum_{r=1}^{\mathcal{N}_\ell} \left[f\left(\mathbf{Z}_\ell^{(r)}\right) - f\left(\mathbf{Z}_{\ell-1}^{(r)}\right) \right].$$

The CPU cost of generating each pair of sample paths is denoted as \mathcal{C}_ℓ . The stochastic simulation method that is used to estimate each sample value will be carried out with a variance reduction technique that ensures that the variance, \mathcal{V}_ℓ , is small. In total there are L intermediate levels that reduce the estimator bias.

The final level. Optionally, the final level, $L + 1$, removes all remaining bias⁵. We compare pairs of sample paths, $(\mathbf{X}, \mathbf{Z}_L)$, one of which is unbiased, and the other a tau-leap method sample path with time-step τ_L . The estimator on this level is equal to

$$\mathcal{Q}_{L+1} \approx \frac{1}{\mathcal{N}_{L+1}} \sum_{r=1}^{\mathcal{N}_{L+1}} \left[f(\mathbf{X}^{(r)}) - f(\mathbf{Z}_L^{(r)}) \right].$$

The CPU time required for each individual pair of sample paths is given by \mathcal{C}_{L+1} .

3.1. Configuring the multi-level method

We assess the efficiency of the multi-level method by determining the total CPU time required to estimate a summary statistic of interest. As each sample point on level ℓ takes \mathcal{C}_ℓ units of CPU time to generate, the total CPU time is given by $\sum_{\ell=0}^{L+1} \mathcal{C}_\ell \mathcal{N}_\ell$. Each level is independently simulated; therefore, the total estimator variance, $\widehat{\mathcal{V}}$, is given by

$$\widehat{\mathcal{V}} = \sum_{\ell=0}^{L+1} \mathcal{V}_\ell / \mathcal{N}_\ell. \quad (14)$$

If we wish to constrain the overall estimator to within a statistical error of ε , and the values of \mathcal{V}_ℓ are known, then it can be shown that the total computational cost is minimized by choosing \mathcal{N}_ℓ (for $\ell = 0, 1, \dots, L + 1$) to be [11]

$$\mathcal{N}_\ell = \left\{ \frac{1}{\varepsilon} \sum_{m=0}^{L+1} \sqrt{\mathcal{V}_m \cdot \mathcal{C}_m} \right\} \sqrt{\frac{\mathcal{V}_\ell}{\mathcal{C}_\ell}}. \quad (15)$$

The values of \mathcal{V}_ℓ are, however, usually unknown, and must therefore be estimated. Typically, a small number of trial sample paths are generated, and \mathcal{V}_ℓ is thus estimated; we will discuss the merits of this approach in Section 6.

If Equation (15) is used to specify \mathcal{N}_ℓ , then the total simulation cost is given by

$$\mathcal{C} = \frac{1}{\varepsilon} \left\{ \sum_{\ell=0}^{L+1} \sqrt{\mathcal{V}_\ell \cdot \mathcal{C}_\ell} \right\}^2. \quad (16)$$

Thus, it is to our advantage to ensure that the sample variance on each level, \mathcal{V}_ℓ , is as small as possible. We outline the variance reduction technique employed by Anderson and Higham [7] to reduce the variance of the correction levels (i.e. where $\ell = 1, \dots, L$), and mention a second variance reduction technique for the final level (i.e. $\ell = L + 1$).

We first explain how a variance reduction technique benefits the correction level estimators, \mathcal{Q}_ℓ , for $\ell = 1, \dots, L$. To generate the r -th sample value of \mathcal{Q}_ℓ , $\left[f(\mathbf{Z}_\ell^{(r)}) - f(\mathbf{Z}_{\ell-1}^{(r)}) \right]$, we need to

⁵If a low-bias (as opposed to an unbiased) estimate can be tolerated, then this step can be skipped, and the overall CPU time is consequently reduced.

generate two sample paths using the tau-leap method, but with different time-steps. As we are constructing a Monte Carlo estimator, we require each of the sample values to be independent of the other sample values. The key point to note is that for each r , there is no need for $f(\mathbf{Z}_\ell^{(r)})$ and $f(\mathbf{Z}_{\ell-1}^{(r)})$ to be independent of one another. Recall that

$$\begin{aligned} \mathcal{V} \left[f(\mathbf{Z}_\ell^{(r)}) - f(\mathbf{Z}_{\ell-1}^{(r)}) \right] &= \mathcal{V} \left[f(\mathbf{Z}_\ell^{(r)}) \right] + \mathcal{V} \left[f(\mathbf{Z}_{\ell-1}^{(r)}) \right] \\ &\quad - 2 \text{Cov} \left[f(\mathbf{Z}_\ell^{(r)}), f(\mathbf{Z}_{\ell-1}^{(r)}) \right]. \end{aligned}$$

We note it is therefore in our interests for $\mathbf{Z}_\ell^{(r)}$ and $\mathbf{Z}_{\ell-1}^{(r)}$ to exhibit a strong, positive correlation. This, in turn, should give rise to a lower estimator variance, and the total CPU time given by Equation (16) is reduced.

The final level, $\ell = L + 1$, involves coupling an unbiased sample path, $\mathbf{X}^{(r)}$, to a biased sample path, $\mathbf{Z}_L^{(r)}$. The estimator variance of samples of \mathcal{Q}_{L+1} can be similarly reduced. Reference [11] contains additional information on how to do this.

In Sections 4 and 5, we will discuss two distinct variance reduction approaches. The key idea is to use the same random input, as far as possible, for both \mathbf{Z}_ℓ and $\mathbf{Z}_{\ell-1}$. We explain the Split Propensity Method (SPM) [7] in Section 4, the previous gold-standard for multi-level Monte Carlo. The Common Process Method (CPM) is explained in Section 5: this has been previously used to develop low-variance estimators for parameter sensitivity analysis [18], but, until now, has yet to be used for multi-level Monte Carlo simulation. Having outlined the multi-level method, we will compare the two different implementations. We assess each implementation for robustness and improvements in numerical performance.

4. Split propensity method

Here we present the method developed by Anderson and Higham [7] that we call the SPM. The SPM involves the *coupling* of two tau-leap processes that are required to provide sample paths \mathbf{Z}_ℓ and $\mathbf{Z}_{\ell-1}$ (we discuss \mathcal{Q}_{L+1} separately below). For brevity, we will refer to a pair of sample paths as comprising a coarse path (referring to $\mathbf{Z}_{\ell-1}$, with $\tau = \tau_C = \tau_0/\mathcal{M}^{\ell-1}$) and a fine path (referring to \mathbf{Z}_ℓ , with $\tau = \tau_F = \tau_0/\mathcal{M}^\ell$). The SPM is premised as follows: we aim to keep the r -th sample paths of the approximate processes with time-steps τ_C and τ_F as similar to each other as possible over the time period of interest. The SPM is implemented by creating ‘virtual reaction channels’. Each reaction, R_j , is split into three virtual channels, labelled as R_j^1 , R_j^2 and R_j^3 , defined such that:

- R_j^1 : reaction R_j fires in both the coarse and fine sample paths;

- R_j^2 : reaction R_j fires only in the coarse sample path;
- R_j^3 : reaction R_j fires only in the fine sample path.

If R_j has propensity p_j^C in the coarse sample path, and p_j^F in the fine sample path, then the propensities of R_j^k , for $k \in \{1, 2, 3\}$, are given by

$$p_j^1 = \min \{p_j^F, p_j^C\}, \quad p_j^2 = p_j^C - p_j^1, \quad p_j^3 = p_j^F - p_j^1. \quad (17)$$

Note that at least one of p_j^2 and p_j^3 are zero. Then, if the state vectors \mathbf{Z}_ℓ and $\mathbf{Z}_{\ell-1}$ are similar, then we expect that, for each j , $p_j^C \sim p_j^F$, so that $p_j^1 \gg p_j^2, p_j^3$. In this case, the majority of reactions will occur through reaction channels of the form R_j^1 , so the state vectors \mathbf{Z}_ℓ and $\mathbf{Z}_{\ell-1}$ ought to remain somewhat similar after each time-step. Full pseudo-code is presented in Algorithm 4.

Algorithm 4 The SPM-coupled tau-leap method. This simulates a pair of sample paths.

Require: initial conditions, $\mathbf{Z}(0)$, time-steps τ_C, τ_F , and terminal time, T .

```

1: set  $\mathbf{Z}^C \leftarrow \mathbf{Z}(0)$ ,  $\mathbf{Z}^F \leftarrow \mathbf{Z}(0)$  and  $t \leftarrow 0$ . Set constant  $\mathcal{M} \leftarrow \tau_C/\tau_F$ 
2: for each  $t \in \{0, \tau_{\ell-1}, 2 \cdot \tau_{\ell-1}, \dots, T - \tau_{\ell-1}\}$  do
3:   for each  $R_j$ , calculate propensity value  $p_j^C(\mathbf{Z}^C)$ 
4:   for each  $s \in \{t, t + \tau_\ell, \dots, t + (\mathcal{M} - 1) \cdot \tau_\ell\}$  do
5:     for each  $R_j$ , calculate propensity value  $p_j^F(\mathbf{Z}^F)$ 
6:     for each  $R_j$ , calculate virtual propensities  $p_j^1, p_j^2$  and  $p_j^3$ 
7:     for each  $R_j$  and for each  $k \in \{1, 2, 3\}$  do
8:       generate  $K_{jk} \sim \mathcal{P}(p_j^k \cdot \tau_\ell)$ 
9:     end for
10:    set  $\mathbf{Z}^C \leftarrow \mathbf{Z}^C + \sum_{j=1}^M (K_{j1} + K_{j2}) \cdot \boldsymbol{\nu}_j$ 
11:    set  $\mathbf{Z}^F \leftarrow \mathbf{Z}^F + \sum_{j=1}^M (K_{j1} + K_{j3}) \cdot \boldsymbol{\nu}_j$ 
12:   end for
13: end for

```

4.1. Estimating \mathcal{Q}_{L+1}

The final estimator, \mathcal{Q}_{L+1} , couples a tau-leap process \mathbf{Z}_L (with time-step given by τ_L), with an exact process, \mathbf{X} . The SPM procedure demonstrated above is followed to produce three virtual channels. The resultant system is simulated with a suitable algorithm; pseudo-code is provided in Algorithm 5.

5. Common process method

In this section, we present the CPM. As mentioned, this method has been previously used to produce low variance estimates for parametric sensitivities [18]. As with the SPM, the CPM involves the *coupling* of two tau-leap processes that are required to provide sample paths for \mathbf{Z}_ℓ and $\mathbf{Z}_{\ell-1}$, so that we can efficiently estimate \mathcal{Q}_ℓ , where $\ell = 1, \dots, L$. As before, we will refer to a pair of sample paths as comprising a coarse (referring to $\mathbf{Z}_{\ell-1}$) and a fine (referring to \mathbf{Z}_ℓ) path.

5.1. Estimating \mathcal{Q}_ℓ , where $\ell = 1, \dots, L + 1$

We now discuss the CPM, and use the RTCR provided by Equation (8) to represent the processes \mathbf{Z}_ℓ (with time-step τ_ℓ) and $\mathbf{Z}_{\ell-1}$ (with time-step $\tau_{\ell-1}$):

$$\begin{aligned}\mathbf{Z}_\ell(T) &= \mathbf{Z}_\ell(0) + \sum_{k=0}^K \sum_{j=1}^M \mathcal{Y}_j(P_{\ell,k-1,j}, P_{\ell,k,j}) \cdot \boldsymbol{\nu}_j; \\ \mathbf{Z}_{\ell-1}(T) &= \mathbf{Z}_{\ell-1}(0) + \sum_{k=0}^{K'} \sum_{j=1}^M \mathcal{Y}_j(P_{\ell-1,k-1,j}, P_{\ell-1,k,j}) \cdot \boldsymbol{\nu}_j,\end{aligned}$$

where

$$P_{\ell,k,j} = \sum_{k'=0}^k p_j(\mathbf{Z}_\ell(\tau_\ell \cdot k')) \cdot \tau_\ell, \quad (18)$$

and the \mathcal{Y}_j ($j = 1, \dots, M$) are unit-rate Poisson processes. The CPM method can produce a low variance estimate for \mathcal{Q}_ℓ by using the same set of M Poisson processes, i.e. $\mathcal{Y}_1, \dots, \mathcal{Y}_M$, for both sample paths \mathbf{Z}_ℓ and $\mathbf{Z}_{\ell-1}$.

This CPM scheme can be implemented by essentially running the tau-leap algorithm twice. We explain the procedure in detail in the following paragraphs, but the method can be summarised as follows:

- during the first phase, we simulate a sample path \mathbf{Z}_ℓ with time-step τ_ℓ . In doing so, the total number of times each Poisson process, \mathcal{Y}_j (for $j = 1, \dots, M$), has fired over each time-step is stored into memory (see Section 5.1.1);
- during the second phase, we simulate a sample path $\mathbf{Z}_{\ell-1}$, making use of the Poisson processes stored in memory and using interpolation (as phases one and two will use distinct time-steps) as required (see Section 5.1.2).

A sample path for the process \mathbf{Z}_ℓ will be simulated according to the pseudo-code provided in Algorithm 6; a sample path for process $\mathbf{Z}_{\ell-1}$ will be simulated according to Algorithm 7.

Algorithm 5 This simulates a pair of SPM-coupled sample paths: an exact path and a path with the tau-leap method using time-step τ .

Require: initial conditions, $\mathbf{X}(0)$, time-step, τ , and terminal time, T .

```

1: set  $\mathbf{X} \leftarrow \mathbf{X}(0)$ ,  $\mathbf{Z} \leftarrow \mathbf{X}(0)$ ,  $t \leftarrow 0$ , and  $t^* \leftarrow \tau$ 
2: for each  $R_j$ , calculate propensity values  $p_j^X(\mathbf{X})$  and  $p_j^Z(\mathbf{Z})$ 
3: loop
4:   for each  $R_j$ , set  $p_j^* = \max\{p_j^X, p_j^Z\}$ , and set  $p_0^* = \sum_{j=1}^M p_j^*$ 
5:   set  $\Delta \leftarrow \text{Exp}(p_0^*)$ 
6:   if  $t + \Delta > T$  then
7:     break
8:   else if  $t + \Delta > t^*$  then
9:     set  $t \leftarrow t^*$ , and  $t^* \leftarrow t^* + \tau$ 
10:    for each  $R_j$ , calculate propensity values  $p_j^Z(\mathbf{Z})$ 
11:    else
12:      choose index  $k$ , where  $k$  has probability  $p_k^*/p_0^*$  of being chosen
13:      with probability  $p_k^X/p_k^*$ , set  $\mathbf{X} \leftarrow \mathbf{X} + \boldsymbol{\nu}_k$ 
14:      with probability  $p_k^Z/p_k^*$ , set  $\mathbf{Z} \leftarrow \mathbf{Z} + \boldsymbol{\nu}_k$ 
15:      set  $t \leftarrow t + \Delta$ 
16:      for each  $R_j$ , calculate propensity values  $p_j^X(\mathbf{X})$ 
17:    end if
18:  end loop

```

5.1.1. Phase one: simulating \mathbf{Z}_ℓ

We first explain how to simulate a sample path of the tau-leap process \mathbf{Z}_ℓ . We use the RTCR as a starting point. At time $t = 0$, the population is equal to the initial condition, $\mathbf{Z}(0)$. The populations at later times, $t = k \cdot \tau_\ell$ (for $k = 1, 2, \dots$), are given by $\mathbf{Z}(k \cdot \tau_\ell)$. By setting $k = 1, 2, \dots$, the value of $\mathbf{Z}(k \cdot \tau_\ell)$ can be recursively calculated according to Equation (6), that is

$$\mathbf{Z}_\ell(k \cdot \tau_\ell) = \mathbf{Z}_\ell((k-1) \cdot \tau_\ell) + \sum_{j=1}^M \mathcal{Y}_j(P_{\ell, k-1, j}, P_{\ell, k, j}) \cdot \boldsymbol{\nu}_j.$$

Recall that the quantity $\mathcal{Y}_j(P_{\ell, k-1, j}, P_{\ell, k, j})$ represents the number of arrivals of the unit-rate Poisson process \mathcal{Y}_j over the interval $(P_{\ell, k-1, j}, P_{\ell, k, j}]$. The value of $\mathcal{Y}_j(P_{\ell, k-1, j}, P_{\ell, k, j})$ is Poisson distributed with parameter $P_{\ell, k, j} - P_{\ell, k-1, j}$. Note that, in accordance with Equation (7),

$$P_{\ell, k, j} - P_{\ell, k-1, j} = p_j(\mathbf{Z}_\ell(k \cdot \tau_\ell)) \cdot \tau_\ell.$$

The number of times reaction R_j occurs over the time span $((k-1) \cdot \tau_\ell, k \cdot \tau_\ell]$ is given by the Poisson random variate⁶, $\mathcal{P}(p_j(\mathbf{Z}_\ell(k \cdot \tau_\ell)) \cdot \tau_\ell)$.

The tau-leap method provided in Algorithm 6 is therefore implemented. At each step of the algorithm ($k = 1, 2, \dots$), we will store the tuple $\langle p_j(\mathbf{Z}_\ell(k \cdot \tau_\ell)) \cdot \tau, K_j \rangle$, where K_j describes the number of times reaction R_j occurs over the time span. The tuple is stored in an ordered list \mathcal{F}_j .

The ordered list \mathcal{F}_j therefore stores details of the arrivals of the unit-rate Poisson process \mathcal{Y}_j , and is indexed by $j = 1, \dots$.

Example. Suppose that $\mathcal{F}_j = \{\langle 2.1, 3 \rangle, \langle 4.0, 7 \rangle, \langle 1.7, 3 \rangle, \dots\}$. This means that we know that over the interval (of the unit-rate process) $(0.0, 2.1]$, there were three arrivals in the unit-rate Poisson process \mathcal{Y}_j . We also know that there were seven arrivals over the interval $(2.1, 6.1]$. Thus, over the interval $(0.0, 6.2]$, a total of ten arrivals were observed.

Therefore, the ordered lists, \mathcal{F}_j (for $j = 1, \dots, M$), provide *partial information* about the arrivals of the unit-rate Poisson process \mathcal{Y}_j . Each ordered list, \mathcal{F}_j , contains information about the total number of arrivals between specific positions, but does *not* contain information about the precise time at which each arrival is observed. Additional details will therefore be generated when they are needed for the coarse path. In the next section, we show how to use these ordered lists to generate a sample path of the process $\mathbf{Z}_{\ell-1}$ according to the CPM.

Algorithm 6 Phase one of the CPM-coupled tau-leap method. This simulates a single fine sample path, and records partial details of the Poisson processes associated with each reaction channel.

Require: initial conditions, $\mathbf{Z}(0)$, time-step, τ , and terminal time, T .

- 1: set $\mathbf{Z} \leftarrow \mathbf{Z}(0)$ and set $t \leftarrow 0$
 - 2: **while** $t < T$ **do**
 - 3: **for** each R_j **do**
 - 4: calculate propensity value $p_j(\mathbf{Z})$
 - 5: generate $K_j \sim \mathcal{P}(p_j(\mathbf{Z}) \cdot \tau)$
 - 6: add the tuple $\langle p_j(\mathbf{Z}) \cdot \tau, K_j \rangle$ to the end of ordered list \mathcal{F}_j
 - 7: **end for**
 - 8: set $\mathbf{Z} \leftarrow \mathbf{Z} + \sum_{j=1}^M K_j \cdot \boldsymbol{\nu}_j$ and set $t \leftarrow t + \tau$
 - 9: **end while**
-

⁶Note that, as expected, this quantity is consistent with Algorithm 2.

5.1.2. Phase two: simulating $\mathbf{Z}_{\ell-1}$

The CPM method uses the same unit-rate Poisson process, \mathcal{Y}_j , to fire the R_j reactions in each of \mathbf{Z}_ℓ and $\mathbf{Z}_{\ell-1}$. Therefore, in this section we describe how to generate a sample path of process $\mathbf{Z}_{\ell-1}$ using the information stored in the ordered lists, \mathcal{F}_j (for $j = 1, \dots, M$). As summarised above, the ordered list, \mathcal{F}_j , contains only an outline of the Poisson process, \mathcal{Y}_j (for $j = 1, \dots, M$), and further details of these Poisson processes need to be filled in as required. The population of $\mathbf{Z}_{\ell-1}$ is determined by setting $k = 1, 2, \dots$, and recursively calculating

$$\mathbf{Z}_{\ell-1}(k \cdot \tau_{\ell-1}) = \mathbf{Z}_{\ell-1}((k-1) \cdot \tau_{\ell-1}) + \sum_{j=1}^M \mathcal{Y}_j(P_{\ell-1, k-1, j}, P_{\ell-1, k, j}) \cdot \boldsymbol{\nu}_j.$$

From the ordered list \mathcal{F}_j we can directly read off the number of arrivals of the Poisson process \mathcal{Y}_j at positions $P_{\ell-1, 0, j}, P_{\ell-1, 1, j}, P_{\ell-1, 2, j}, \dots$ (for $j = 1, \dots, M$). At other positions, interpolation will be required. We use the following two rules to interpolate the Poisson process:

Rule 2. If there are K arrivals over the interval (α, γ) then, for β between α and γ , the number of arrivals over the interval (α, β) is binomially distributed⁷

$$\mathcal{B}\left(K, \frac{\beta - \alpha}{\gamma - \alpha}\right). \quad (19)$$

Rule 3 (a caveat). If there are K_1 arrivals over the interval (α, γ) , and K_2 arrivals over the interval (γ, β) , then interpolation using Rule 2 must be individually performed on the intervals (α, γ) and (γ, β) . Rule 2 cannot be directly applied to the entire interval (α, β) with $K_1 + K_2$ arrivals.

Rules 2 and 3 provide all the tools needed for Monte Carlo simulation. We first illustrate interpolation of a Poisson process with an example; pseudo-code is then provided in Algorithm 7.

Example. We return to our earlier example of an ordered list, $\mathcal{F}_j = \{\langle 2.1, 3 \rangle, \langle 4.0, 7 \rangle, \langle 1.7, 3 \rangle, \dots\}$. Suppose we wish to determine $\mathcal{Y}_j(0.0, 5.0)$, the number of arrivals in the unit-rate Poisson process by time 5.0. We first apply Rule 3: there are three arrivals over the time span $(0.0, 2.1]$, and we need to determine how many further arrivals are observed over the time span $(2.1, 5.0]$. We apply Rule 2 as follows: over the interval $(2.1, 6.1]$ there are seven arrivals, and we want to know how many of these arrivals occur inside the sub-interval $(2.1, 5.0)$. A binomial variate, $\mathcal{B}(7, 2.9/4)$ is generated to interpolate the Poisson process.

5.2. Estimating \mathcal{Q}_{L+1}

The final estimator, \mathcal{Q}_{L+1} , couples a tau-leap process \mathbf{Z}_L (with time-step given by τ_L), with an exact process, \mathbf{X} . The CPM is implemented as outlined in the following paragraph.

⁷This follows as the K arrivals are uniformly distributed over (α, γ) .

Algorithm 7 Phase two of the CPM-coupled tau-leap method. This simulates a single coarse sample path from the Poisson processes stored during phase one.

Require: initial conditions, $\mathbf{Z}(0)$, time-step, τ , ordered lists, \mathcal{F}_j (with $j = 1, \dots, M$), and terminal time, T .

```

1: set  $\mathbf{Z} \leftarrow \mathbf{Z}(0)$  and set  $t \leftarrow 0$ 
2: while  $t < T$  do
3:   for each  $R_j$  do
4:     calculate propensity value  $p_j(\mathbf{Z})$ 
5:     set  $P \leftarrow 0, K \leftarrow 0$ 
6:     while  $P < p_j(\mathbf{Z}) \cdot \tau$  do
7:       if  $\mathcal{F}_j = \emptyset$  then
8:         break
9:       else
10:        read and then delete  $(P', K')$  from the front of  $\mathcal{F}_j$ 
11:        set  $P \leftarrow P + P'$  and  $K \leftarrow K + K'$ 
12:       end if
13:     end while
14:     if  $P > p_j(\mathbf{Z}) \cdot \tau$  then
15:       generate  $K' \sim \mathcal{B}(K', (P - p_j(\mathbf{Z}) \cdot \tau) / P')$ 
16:       set  $K_j \leftarrow K - K'$ 
17:       add the tuple  $\langle P - p_j(\mathbf{Z}) \cdot \tau, K' \rangle$  to the front of ordered list  $\mathcal{F}_j$ 
18:     else if  $P < p_j(\mathbf{Z}) \cdot \tau$  then
19:       generate  $K' \sim \mathcal{P}(p_j(\mathbf{Z}) \cdot \tau - P)$ 
20:       set  $K_j \leftarrow K + K'$ 
21:     else
22:       set  $K_j \leftarrow K$ 
23:     end if
24:   end for
25:   set  $\mathbf{Z} \leftarrow \mathbf{Z} + \sum_{j=1}^M K_j \cdot \boldsymbol{\nu}_j$  and  $t \leftarrow t + \tau$ 
26: end while

```

To simulate a sample path of the exact process, \mathbf{X} , every reaction must be individually simulated. Therefore, we will need to determine every required arrival time of each Poisson process, \mathcal{Y}_j . The arrival times of each Poisson process are then saved into memory; the same set of Poisson processes is used to generate a sample path for \mathbf{Z}_ℓ , the tau-leap process with time-step

Algorithm 8 Phase one of the CPM-coupled DM and tau-leap method. This simulates a single, exact sample path, and records the Poisson processes associated with each reaction channel.

Require: initial conditions, $\mathbf{X}(0)$, and terminal time, T .

```

1: set  $\mathbf{X} \leftarrow \mathbf{X}(0)$ , and set  $t \leftarrow 0$ 
2: for each  $R_j$ , set  $p_j \leftarrow 0$ , generate  $T_j \leftarrow \text{Exp}(1)$ , and store  $T_j$  as the first element of  $\mathcal{F}_j$ 
3: loop
4:   for each  $R_j$ , calculate propensity values  $p_j(\mathbf{X})$  and calculate  $\Delta_j$  as
           
$$\Delta_j = \frac{T_j - p_j}{p_j}$$

5:   set  $\Delta \leftarrow \min_j \Delta_j$ , and  $k \leftarrow \text{argmin}_j \Delta_j$ 
6:   if  $t + \Delta > T$  then
7:     break
8:   end if
9:   set  $\mathbf{X}(t + \Delta) \leftarrow \mathbf{X}(t) + \boldsymbol{\nu}_j$ , set  $t \leftarrow t + \Delta$ , and for each  $R_j$ , set  $p_j \leftarrow p_j + p_j \cdot \Delta$ 
10:  generate  $u \sim \text{Exp}(1)$ , then set  $T_k \leftarrow T_k + u$  and append  $u$  to end of  $\mathcal{F}_k$ 
11: end loop

```

τ_L . A range of algorithms can be satisfactorily implemented to simulate the sample paths. We will adapt the Modified Next Reaction Method (MNRM), as described by Anderson [19]. The CPM proceeds in two phases, that can be outlined as follows:

- during the first phase, a sample path for process \mathbf{X} is generated using the MNRM. The waiting times of Poisson process \mathcal{Y}_j are recorded in an ordered list⁸, \mathcal{F}_j (Algorithm 8);
- during the second phase, a sample path for process \mathbf{Z} is generated using the MNRM. The waiting times for Poisson process \mathcal{Y}_j are determined using ordered list \mathcal{F}_j (Algorithm 9).

6. Numerical experiments

In this section, we consider three representative case studies. In each case, we will test our two tau-leap based multi-level coupling methods: the SPM and the CPM. We will also compare the SPM and CPM implementations of the multi-level method with the regular DM.

Algorithmic performance will depend on the chosen algorithm parameters: L , \mathcal{M} and τ_0 . A simple search procedure can easily discard clearly inefficient algorithm parameter choices, leav-

⁸For $\ell = L + 1$, the ordered list \mathcal{F}_j contains only numbers that correspond to the individual arrival times; but where $\ell = 1, \dots, L$, the list contains tuples.

Algorithm 9 Phase two of the CPM-coupled DM and tau-leap method. This simulates an approximate sample path from the Poisson processes stored during phase one.

Require: initial conditions, $\mathbf{Z}(0)$, terminal time, T , and ordered lists, \mathcal{F}_j ($j = 1, \dots, M$).

```

1: set  $\mathbf{Z} \leftarrow \mathbf{Z}(0)$ ,  $t \leftarrow 0$ ,  $t^* \leftarrow \tau$ 
2: for each  $R_j$ , set  $p_j \leftarrow 0$ 
3: for each  $R_j$ , set  $T_j$  to be the first element of list  $\mathcal{F}_j$ , then delete the first element of  $\mathcal{F}_j$ 
4: for each  $R_j$ , calculate propensity values  $p_j(\mathbf{Z})$ 
5: loop
6:   for each  $R_j$ , calculate  $\Delta_j$  as
           
$$\Delta_j = \frac{T_j - p_j}{p_j} \tag{20}$$

7:   set  $\Delta \leftarrow \min_j \Delta_j$ , and  $k \leftarrow \operatorname{argmin}_j \Delta_j$ 
8:   if  $t + \Delta > T$  then
9:     break
10:  else if  $t + \Delta > t^*$  then
11:    set  $t \leftarrow t^*$ , and  $t^* \leftarrow t^* + \tau$ 
12:    for each  $R_j$ , set  $p_j \leftarrow p_j + p_j \cdot (t^* - t)$ , then recalculate propensity  $p_j$ 
13:  else
14:    set  $\mathbf{Z}(t + \Delta) \leftarrow \mathbf{Z}(t) + \boldsymbol{\nu}_k$ , set  $t \leftarrow t + \Delta$ , and for each  $R_j$ , set  $p_j \leftarrow p_j + p_j \cdot \Delta$ 
15:    if  $\mathcal{F}_k \neq \emptyset$  then
16:      let  $u$  be the first element of  $\mathcal{F}_k$ : set  $T_k \leftarrow T_k + u$ , and delete the first element of  $\mathcal{F}_k$ 
17:    else
18:      generate  $u \sim \operatorname{Exp}(1)$ , then set  $T_k \leftarrow T_k + u$ 
19:    end if
20:  end if
21: end loop

```


ing us with a range of possible values for L , \mathcal{M} and τ_0 that should be investigated in greater detail. By using the algorithm parameters earmarked for further investigation, we compare the computational performance of the SPM and CPM variance reduction methods on the multi-level method. We will therefore compare the relative performance of our algorithms using a wide range of choices of L , \mathcal{M} and τ_0 that might be implemented by an end-user of the algorithm.

Our sample paths have been generated by using C++, according to the 2011 standard (C++11). Where possible, the C++ Standard Template Library has been used, and all calculations have been performed with double precision. Sample paths were generated by using an AMD FX-4350 CPU,

with an advertised clock-speed of 4.2 GHz.

6.1. Case study 1: a gene regulatory network model

The first case study is a model of gene expression that has been used by Anderson and Higham [7] to test the multi-level method. In this model, mRNA is produced, and this leads to the production of proteins. The proteins bind to form dimers, and the mRNA and protein molecules may also decay. The biochemical reaction network is stated as:



We write the numbers of mRNA, protein and dimer molecules at time t , respectively, as $\mathbf{X}(t) = [X_1(t), X_2(t), X_3(t)]^T$ and consequently the initial condition can be expressed as $\mathbf{X}(0) = [0, 0, 0]^T$. The parameters are as supplied, and are dimensionless. We will estimate the dimer population at terminal time $T = 1$: this is given by $\mathbb{E}[X_3(1)]$. As a reference, this calculation can be undertaken with the DM. It takes a little over two hours (7223 seconds) to generate the 4.8×10^6 sample paths required for the estimate of $\mathbb{E}[X_3(1)] = 3714.3 \pm 1.0$. Here ‘ ± 1.0 ’ indicates a 95% confidence interval of semi-length 1.0.

We now use the multi-level method to estimate $\mathbb{E}[X_3(1)]$ to the same level of accuracy. We test both the SPM and CPM implementations of the multi-level method. In Figure 1, we show the average CPU time taken by the respective multi-level method implementations. The values of L , \mathcal{M} and τ_0 are varied; each average is computed over 100 test runs of the complete multi-level algorithm. Our fastest CPM configuration takes, on average, 120.9 seconds to estimate the dimer population (using $\mathcal{M} = 2$, $\tau_0 = 1/8$ and $L = 7$), and our most efficient SPM implementation takes, on average, 156.5 seconds to estimate the same quantity (using $\mathcal{M} = 3$, $\tau_0 = 1/9$ and $L = 5$). For System (21), our most efficient SPM-implementation of the multi-level method therefore takes approximately 29% longer to run than the comparable CPM implementation. In fact, for each test case that we considered, when compared with the CPM, the SPM implementation requires more CPU time to estimate the required summary statistic.

Whilst the CPM implementation can reduce the average duration of multi-level Monte Carlo simulation, another key benefit is that its run-time is far more predictable than the run-time of the SPM implementation. The black lines in Figure 1 indicate the range occupied by the 10-th to 90-th percentiles of the total CPU time. It is clear that the CPU times of the CPM implementation are very tightly clustered around the mean CPU time, whilst, for the SPM, the CPU times display a higher variance. Moreover, the lower run-time of the CPM means that, even

if a sub-optimal set of algorithm parameters is used, then the effect on the CPU time required is somewhat limited.

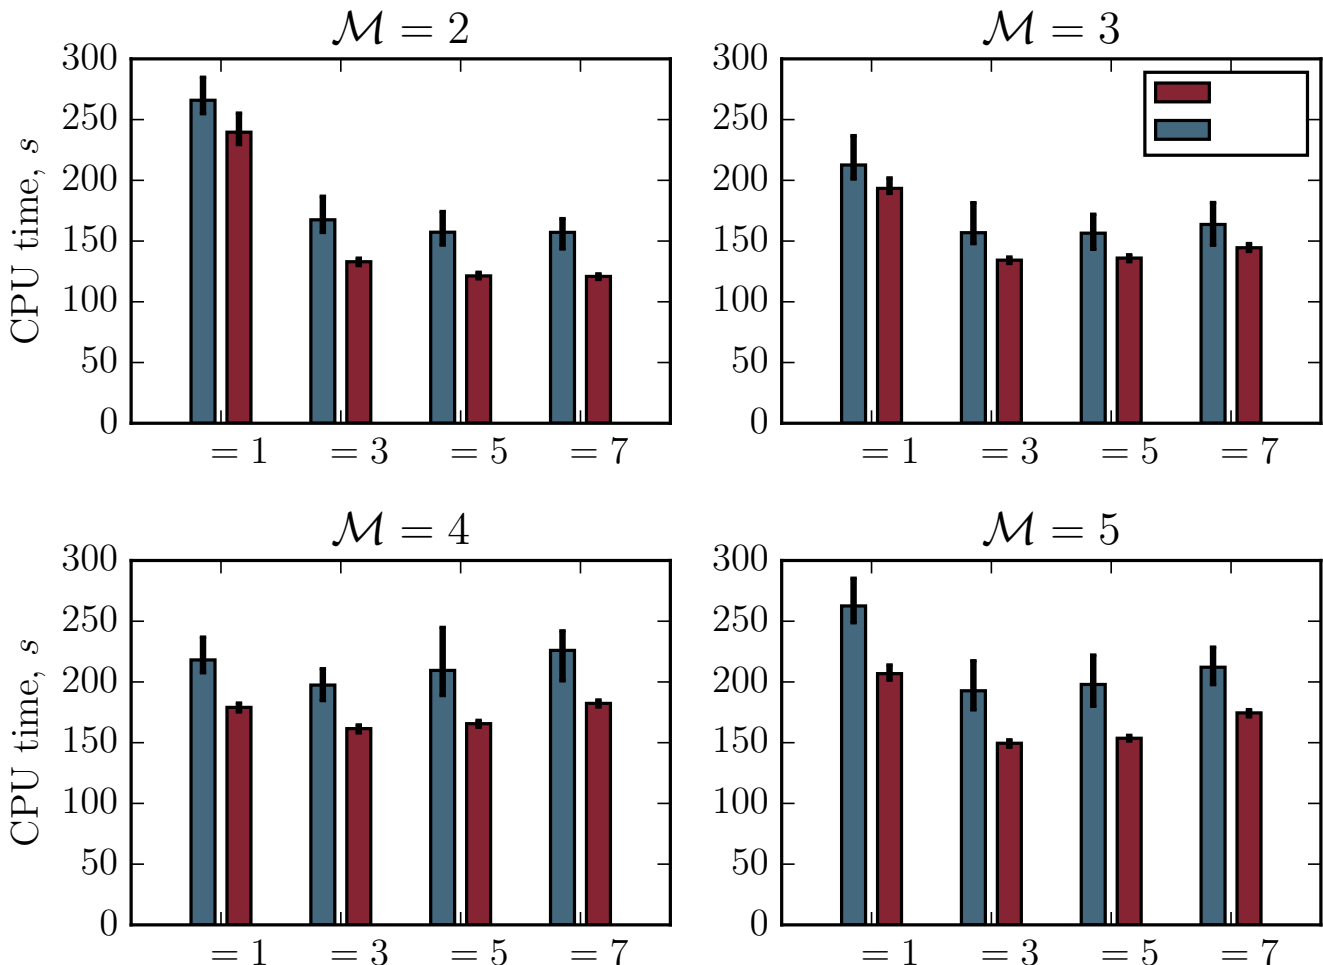


Figure 1: The average CPU time required by the multi-level method to estimate $\mathbb{E}[X_3(1)]$ for System (21). We vary \mathcal{M} and L ; the estimator is unbiased. The black bars indicate the range occupied by the 10-th to 90-th percentiles of the data. The values of τ_0 are: $\mathcal{M} = 2 \Rightarrow \tau_0 = 1/8$; $\mathcal{M} = 3 \Rightarrow \tau_0 = 1/9$; $\mathcal{M} = 4 \Rightarrow \tau_0 = 1/16$ and $\mathcal{M} = 5 \Rightarrow \tau_0 = 1/5$.

To compare the SPM with the CPM in more detail, we will need to concentrate on specific values of τ_0 , \mathcal{M} and L . We will not hand-pick algorithm parameters that give the CPM any potential advantage over the SPM. We take $\tau_0 = 1/9$, $\mathcal{M} = 3$ and $L = 5$. With the aforementioned algorithm parameters, the CPM multi-level algorithm takes an average of 136.0 seconds to run (15.1 seconds slower than the fastest CPM configuration), whilst the SPM method runs in an average of 156.5 seconds (the fastest SPM configuration). In this case, the SPM takes approximately 15% longer to run than the CPM method.

In Figure 2, we show the empirical mean, sample variance, kurtosis and CPU time of each level estimator, \mathcal{Q}_ℓ , when each of the SPM and CPM are used. The same base level ($\ell = 0$)




Figure 2: The expected mean, sample variance, normalised kurtosis and CPU time for each level ℓ in our multi-level simulator for System (21). We have taken $\tau_0 = 1/9$, $\mathcal{M} = 3$ and $L = 5$, and used 10^5 sample paths on each level. Note that the mean values for the SPM and CPM shown in the upper-left diagram overlap.

is used for the SPM and CPM. The variances, \mathcal{V}_ℓ , are substantially lower when the CPM, and not the SPM, is used. However, the CPU time taken to generate each sample, \mathcal{C}_ℓ , is typically higher when the CPM is used. For the higher levels, the use of the SPM results in a very high kurtosis (this effect has previously been described as representing a ‘catastrophic decoupling’ – see reference [11]). A high kurtosis means that more of the sample variance can be attributed to infrequent but substantial deviations from the mean (as compared with frequent but modestly

sized deviations from the mean).

The effect of using the CPM implementation on the overall CPU time is outlined as follows. Subject to choosing \mathcal{N}_ℓ according to Equation (14), the total CPU time is given by Equation (16), i.e.

$$\frac{1}{\varepsilon^2} \left\{ \sum_{\ell=0}^{L(+1)} \sqrt{\mathcal{C}_\ell \cdot \mathcal{V}_\ell} \right\}^2 .$$

When the CPM is implemented, different values of \mathcal{V}_ℓ and \mathcal{C}_ℓ are inserted into Equation (16), and the result is that, with this case study, the total CPU time is reduced.

We now explore the variation in CPU times between successive iterations of the multi-level method. Estimated values of \mathcal{V}_ℓ and \mathcal{C}_ℓ are used to populate Equation (14), and variations in these estimates cause variations in the CPU time expended by the multi-level method. There are two competing approaches for estimating \mathcal{V}_ℓ and \mathcal{C}_ℓ :

- the ‘one-step calibration’ approach can be used. A small number of initial sample paths are generated for each level (for example, 10^2 or 10^3 sample paths), and then \mathcal{V}_ℓ and \mathcal{C}_ℓ are estimated. Equation (14) is then evaluated, and the requisite number of sample paths is generated. This is the approach used by Anderson and Higham [7];
- a repeated update, ‘dynamic calibration’ approach can also be used. In this case, the estimated values of \mathcal{V}_ℓ are repeatedly updated as the algorithm progresses, and the required number of simulations on each level is refined. Where appropriate, additional sample paths can be generated. A detailed explanation is contained within our earlier work [11].

The CPU times shown in Figure 1 were generated with the ‘dynamic calibration’ procedure. We now evaluate the effect of using the CPM on the one-step and dynamic calibration procedures.

One-step calibration. We first generate $\mathcal{N} = 10^2$ sample paths from which to estimate \mathcal{V}_ℓ and \mathcal{C}_ℓ . In Figure 3 we compare the effect of using the one-step calibration procedure with the SPM and the CPM. For each of the SPM and CPM, the entire multi-level method is run, from start until finish, 1000 times. We show the absolute CPU time, and plot this against the resultant confidence interval semi-length (as before, we aim for a confidence interval of semi-length 1.0).

Our results indicate that the CPM implementation is far more likely to achieve the required estimator variance than the SPM implementation, and to do so with a broadly comparable CPU time. The mean CPU times for the SPM and CPM are 136.7 and 134.4 seconds, respectively. Regrettably, the required confidence interval semi-length is not necessarily achieved; and it transpires that the estimated sample variances (based on $\mathcal{N} = 10^2$ initial paths) are not sufficiently accurate, and therefore too few sample paths are generated. The issue is far more pronounced for

the SPM implementation as the required confidence interval is attained only 8% of the time. With the CPM, the proportion of runs that achieve the required confidence intervals rises substantially to 58%.

Dynamic calibration. The difficulties encountered with estimated sample variances, \mathcal{V}_ℓ , can be mitigated by using the ‘dynamic calibration’ procedure: after the sample paths for each level are completed, the variance estimates can be updated as appropriate, and Equation (14) used to recalculate the number of sample paths required for each level estimator. If the revised number of sample paths required for a given level is lower than the number of sample paths already generated, a ‘variance re-allocation’ procedure can be followed, so that the sundry sample paths are not wasted. We follow the procedure outlined in our earlier work [11] a total of 1000 times (we used $\mathcal{N} = 10^2$ initial paths to start the dynamic algorithm), we plot our results in Figure 4. The average CPU times of the SPM and CPM are now 158.0 seconds and 134.7 seconds, respectively⁹. The average CPU time required by the CPM is broadly similar to the CPU time required without variance reallocation (see Figure 3), but the CPU time for the SPM method increases significantly.

⁹The averages differ slightly from those presented in Figure 1: the sample size is different.





Figure 3: The full multi-level method is run 1000 times to estimate $\mathbb{E}[X_3(1)]$ for System (21) by following the ‘one-step calibration’ approach, and using both the SPM and CPM. The CPU times are plotted against the semi-length of the confidence intervals attained. The black crosses represent the mean values of the data. The target confidence interval size is shown with a dashed line. The CPM is clearly superior.

Therefore, in this case, the CPM remains superior to the SPM. A CPM implementation of the multi-level method results in a decreased CPU time, when compared with the SPM. The CPM is reliable, and can be used with sample variances estimated using a small number of preliminary sample paths.

6.2. Case study 2: Lotka-Volterra dynamics

In this second case study, we evaluate the performance of the multi-level method with a stochastic analogue of the Lotka-Volterra system [20]. The population dynamics of a predator, A , and its prey, B , are considered. The following reaction channels are defined:



Initially, the population of A , X_1 , and the population of B , X_2 , are both set to equal 1200. We estimate the population levels of System (22) at time $T = 3$.

System (22) clearly exhibits oscillatory dynamics, with the amplitude of the oscillations being highly unstable [21]. To our knowledge, the multi-level method has yet to be successfully applied to a system that exhibits such dynamics, even over a short time-interval. The average predator




Figure 4: The full multi-level method is run 1000 times to estimate $\mathbb{E}[X_3(1)]$ for System (21) by following the ‘dynamic calibration’ approach, and using both the SPM and CPM. The CPU times are plotted against the semi-length of the confidence intervals attained. The black crosses represent the mean values of the data. The target confidence interval size is shown with a dashed line. This approach is clearly preferable to that shown in Figure 3. The CPM still outperforms the SPM.

population, given by $\mathbb{E}[X_1(T)]$, can be estimated with the DM. This calculation takes a little over 25 minutes (1523 seconds), and, with 1.94×10^5 sample paths, we estimate $\mathbb{E}[X_1(T)] = 783.4 \pm 1.0$.

We now use both the SPM- and CPM-controlled multi-level method to estimate $\mathbb{E}[X_1(T)]$. In Figure 5 we show the average CPU times for a range of choices of the refinement factor, \mathcal{M} . In each case, we show the average CPU time achieved with the most efficient choice of L and τ_0 that we have found. The most efficient CPM algorithm we found requires 600.2 seconds of CPU time; this is 37% faster than the most efficient SPM algorithm, which requires 957.1 seconds. For each choice of \mathcal{M} , the most efficient SPM algorithm took substantially longer to run than the most efficient CPM method. Moreover, if we compare the most efficient sets of algorithm parameters for the CPM and SPM, with the DM, then the CPM method is 2.5 times faster than the DM, whereas the SPM by itself is 37% faster than the DM.





Figure 5: The average CPU time required by the multi-level method to estimate $\mathbb{E}[X_1(30)]$ for System (22). We vary \mathcal{M} and L ; the estimator is unbiased. The black bars indicate the range occupied by the 10-th to 90-th percentiles of the data. The values of τ_0 are: $\mathcal{M} = 2 \Rightarrow \tau_0 = T \cdot 2^{-12}$; $\mathcal{M} = 3 \Rightarrow \tau_0 = T \cdot 3^{-8}$; $\mathcal{M} = 4 \Rightarrow \tau_0 = T \cdot 2^{-12}$ and $\mathcal{M} = 5 \Rightarrow \tau_0 = T \cdot 2^{-5}$. In each case, taking $L = 1$ is optimal.

6.3. Case study 3: a logistic growth model

The third case study is of a stochastic logistic growth model that comprises one species, and the following two reaction channels:



Initially, the population of A is given as¹⁰ $X(0) = 50$. We will simulate System (23) until a terminal time $T = 3$, and estimate the mean population at that time, $\mathbb{E}[X(T)]$.

¹⁰As there is only one species, we will suppress the subscript and work with $X(t)$ instead of $X_1(t)$.

The DM estimates $\mathbb{E}[X(T)] = 999.6 \pm 1.0$ using 3800 sample paths in 16.2 seconds. We have run the multi-level algorithm with a wide variety of algorithm parameters (i.e. L , \mathcal{M} , and τ_0) on this system. Our efforts are briefly summarised in Figure 6, which is constructed as follows: for each choice of \mathcal{M} , we find the optimal L and τ_0 for each of the SPM and CPM, and present our results.

We find that the multi-level method can be effectively implemented with the SPM, but the CPM approach is less efficient for this example. The SPM can estimate $\mathbb{E}[X(3)]$ to a 95% confidence interval of semi-length 1.0 within 2.1 seconds, which means that the DM requires 7.7 times as long to perform the calculation to the same level of accuracy. Our best result for the CPM is less efficient: our implementation requires 9.8 seconds of CPU time to estimate $\mathbb{E}[X(3)]$ to within the required statistical accuracy. Whilst the CPM substantially is slower than the SPM, the CPM is still faster than the DM. We provide a brief discussion of these results in Section 8.




Figure 6: The average CPU time required by the multi-level method to estimate $\mathbb{E}[X(3)]$ for System (23). We vary \mathcal{M} and L ; the estimator is unbiased. The black bars indicate the range occupied by the 10-th to 90-th percentiles of the data. The values of τ_0 and L are as follows: for the SPM, $\mathcal{M} = 2 \Rightarrow \tau_0 = T \cdot 2^{-6}$ and $L = 3$; $\mathcal{M} = 3 \Rightarrow \tau_0 = T \cdot 3^{-4}$ and $L = 1$; $\mathcal{M} = 4 \Rightarrow \tau_0 = T \cdot 2^{-4}$ and $L = 1$; and $\mathcal{M} = 5 \Rightarrow \tau_0 = T \cdot 5^{-3}$ and $L = 1$. For the CPM, $\mathcal{M} = 2 \Rightarrow \tau_0 = T \cdot 2^{-9}$ and $L = 1$; $\mathcal{M} = 3 \Rightarrow \tau_0 = T \cdot 3^{-6}$ and $L = 1$; $\mathcal{M} = 4 \Rightarrow \tau_0 = T \cdot 2^{-9}$ and $L = 1$; and $\mathcal{M} = 5 \Rightarrow \tau_0 = T \cdot 5^{-4}$ and $L = 1$.

7. Multi-level Monte Carlo with the R-leap method

In this section, we present a new implementation of the multi-level method. With a view to improving computational performance, we will use the R-leap method to construct an efficient multi-level algorithm. Once we have described our method, we demonstrate the performance of our algorithm with an example.

7.1. Variance reduction with the R-leap method

Our description of the R-leap multi-level method will mimic the description of the tau-leap multi-level method described in Section 3. We write the summary statistic of interest, \mathcal{Q} , as the following telescoping sum:

$$\mathcal{Q} = \sum_{\ell=0}^{L+1} \mathcal{Q}_\ell.$$

The values of \mathcal{Q}_ℓ , for $\ell = 0, \dots, L+1$, are now determined as follows:

- on the *base level*, where $\ell = 0$, we generate sample paths using the R-leap method, with \mathcal{K}_0 reactions simulated at each step. We use a large value of \mathcal{K}_0 , so that we can quickly produce the sample paths required to estimate \mathcal{Q}_0 . The r -th such sample paths is labelled as $\mathbf{Z}_0^{(r)}$, and we use the scalar $Z_0^{(r)}$ to represent the point statistic, $Z_0^{(r)} := f(\mathbf{Z}_0^{(r)})$. Accordingly, the estimator for level 0 is given by

$$\mathcal{Q}_0 := \mathbb{E}\left[Z_0^{(r)}\right] \approx \frac{1}{\mathcal{N}_0} \sum_{r=1}^{\mathcal{N}_0} Z_0^{(r)};$$

- the *correction levels* ($\ell = 1, \dots, L$) require the generation of pairs of sample paths, $[\mathbf{Z}_{\ell-1}, \mathbf{Z}_\ell]$. The ‘fine’ sample path \mathbf{Z}_ℓ is generated using the R-leap method, where \mathcal{K}_ℓ reactions are fired during each step. The ‘coarse’ sample path $\mathbf{Z}_{\ell-1}$ is also generated with the R-leap method, but with $\mathcal{K}_{\ell-1}$ reactions at each step. As with the tau-leap multi-level method, we choose a refinement factor, \mathcal{M} , so that $\mathcal{K}_\ell = \mathcal{K}_{\ell-1}/\mathcal{M}$. The estimator for level ℓ is

$$\mathcal{Q}_\ell := \mathbb{E}\left[Z_\ell^{(r)} - Z_{\ell-1}^{(r)}\right] \approx \frac{1}{\mathcal{N}_\ell} \sum_{r=1}^{\mathcal{N}_\ell} \left[Z_\ell^{(r)} - Z_{\ell-1}^{(r)}\right];$$

- finally, and optionally, the *final correction level*, $L+1$, removes all remaining bias:

$$\mathcal{Q}_{L+1} := \mathbb{E}\left[X^{(r)} - Z_L^{(r)}\right] \approx \frac{1}{\mathcal{N}_{L+1}} \sum_{r=1}^{\mathcal{N}_{L+1}} \left[X^{(r)} - Z_L^{(r)}\right].$$

The sample paths for \mathcal{Q}_0 are performed with the regular R-leap method; pseudo-code is provided in Algorithm 3. As before, we will use an algorithm that *couples* sample paths, so that $\mathcal{Q}_1, \dots, \mathcal{Q}_{L+1}$ can be estimated with a low variance. Note that, unlike the tau-leap method implementation, a special algorithm for the ‘final estimator’ (i.e. \mathcal{Q}_{L+1}) is not required. If we set $\mathcal{K}_{L+1} = 1$, then \mathcal{Q}_{L+1} can be estimated with exactly the same method as $\mathcal{Q}_1, \mathcal{Q}_2, \dots, \mathcal{Q}_L$. As before, the multi-level method will only reduce computational costs if we can efficiently estimate $\mathcal{Q}_1, \dots, \mathcal{Q}_{L+1}$: we now discuss techniques for doing so.

When we implement the R-leap method within the multi-level scheme, we use a coupling method that combines elements of the CPM and the SPM. We start by referring to Section 2.5, which states that at each step of the R-leap algorithm, two quantities are stochastically generated:

1. the time-period covered by that step;
2. the precise combination of reactions that fire during that time-period.

We will let sample paths \mathbf{Z}_ℓ (which we will call the fine path) and $\mathbf{Z}_{\ell-1}$ (which we will call the coarse path) advance by different time-periods at each step (point 1 above). A variance reduction technique is used to choose the time-periods that each sample path traverses. Then, to achieve maximal variance reduction, we will also ensure that, as far as possible, the same reactions fire in each sample path (point 2 above).

In order to simultaneously generate a pair of sample paths, $[\mathbf{Z}_\ell, \mathbf{Z}_{\ell-1}]$, at each step of the coupled simulation algorithm a total of $\mathcal{K}_\ell (= \min\{\mathcal{K}_\ell, \mathcal{K}_{\ell-1}\})$ reaction events will take place in *each* sample path. In particular:

- a Gamma variate, $\Delta = \Gamma(\mathcal{K}_\ell, 1)$, is generated. The time-period spanned by this step in each of the coarse and fine sample paths is then determined by a rescaling argument. In distribution, $\Gamma(\mathcal{K}, \theta) \sim \theta \cdot \Gamma(\mathcal{K}, 1)$. Therefore, the time-period for fine path is given by $\Delta / \sum_{j=1}^M p_j^F$, and for the coarse system by $\Delta / \sum_{j=1}^M p_j^C$;
- the precise combination of reactions is chosen as follows. For each of the \mathcal{K}_ℓ reactions that take place in *each* sample path, the probability that it is a R_j reaction is given by $p_j^C / \sum_{j'=1}^M p_{j'}^C$ (for the coarse path), and $p_j^F / \sum_{j'=1}^M p_{j'}^F$ (for the fine path). Our coupling method must therefore fire reactions with these probabilities. For each R_j (for $j = 1, \dots, M$), we define the probabilities b_j^m , where $m \in \{1, 2, 3\}$ as:

$$b_j^1 = \min \left\{ \frac{p_j^C}{\sum_{j'=1}^M p_{j'}^C}, \frac{p_j^F}{\sum_{j'=1}^M p_{j'}^F} \right\}; \quad (24)$$

$$b_j^2 = \frac{p_j^C}{\sum_{j'=1}^M p_{j'}^C} - b_j^1; \quad b_j^3 = \frac{p_j^F}{\sum_{j'=1}^M p_{j'}^F} - b_j^1.$$

We interpret each probability, b_j^m , where $m \in \{1, 2, 3\}$ by noting¹¹:

- b_j^1 represents the probability reaction R_j takes place in both the coarse and the fine paths;
- b_j^2 represents the probability reaction R_j takes place in only the coarse path;
- b_j^3 represents the probability reaction R_j takes place in only the fine path;

As with the regular R-leap method, a conditional binomial method is used to choose the precise combination of reactions that take place. A total of \mathcal{K}_ℓ reaction events must take

¹¹Note that, for each j , at least one of b_j^2 and b_j^3 will be zero.

place in each of the coarse and the fine paths. We would, if possible, like the same reactions to fire in each sample path. Thus, we first consider the ‘common’ reactions – i.e. those that occur with probability b_j^1 (for $j = 1, \dots, M$). Once the ‘common’ reactions have been determined, then we will determine the reactions specific to either the coarse or the fine path. Note that, as we are coupling sample paths, we expect that the propensities of the fine and coarse paths are similar for each R_j , i.e. $p_j^F \sim p_j^C$, so that $p_j^F / \sum_{j'=1}^M p_{j'}^F \sim p_j^C / \sum_{j'=1}^M p_{j'}^C$. Following Equation (24), we conclude that $b_j^1 \gg b_j^2, b_j^3$. Therefore, most events will be in the form of a ‘common’ reaction, and so will take place in both the coarse and the fine sample paths.

- the propensity values are updated as appropriate. The propensities associated with the fine path are updated at every step of the algorithm (i.e. after the required \mathcal{K}_ℓ reaction events have taken place in the fine path), whilst the propensities of the coarse path are updated every \mathcal{M} steps of the algorithm (i.e. after $\mathcal{M} \cdot \mathcal{K}_\ell = \mathcal{K}_{\ell-1}$ reaction events have taken place in the coarse path).

As mentioned, once the common reactions have been completed, the remaining reactions events must be performed. The R-leap coupling method is presented as pseudo-code in Algorithm 10. We now proceed to present numerical results in Section 7.2. The computational performance of our new method is compared with the tau-leap multi-level method and traditional simulation methods.

7.2. Numerical experimentation

We return to Case study 3 presented in Section 6.3. The reaction channels are



with initial conditions as described in Section 6.3. As before, we estimate $\mathcal{Q} = \mathbb{E}[X(3)]$; for reference, we note that the DM estimates \mathcal{Q} to be 999.6 ± 1.0 within 16.2 seconds. The R-leap multi-level method has been implemented with a range of choices of \mathcal{M} (the refinement factor) and L (that controls the number of levels). We seek an unbiased estimate, which means that $\mathcal{K}_0 = \mathcal{M}^{L+1}$, $\mathcal{K}_1 = \mathcal{M}^L$, \dots , $\mathcal{K}_L = \mathcal{M}$, and $\mathcal{K}_{L+1} = 1$. Our results are collated in Figure 7. The most efficient R-leap multi-level implementation estimates \mathcal{Q} using 3.6 seconds of CPU time. The DM method therefore takes approximately 4.5 times longer than our R-leap multi-level method.

Our results demonstrate that the R-leap multi-level method is a feasible alternative to the tau-leap multi-level method. If we consider the numerical performance of Case study 3 when the tau-leap multi-level method is used, it is clear that for this example, R-leap is superior to a CPM

Algorithm 10 The coupled R-leap method. This simulates a pair of sample paths.

Require: initial conditions, $\mathbf{Z}(0)$, \mathcal{K} ($= \mathcal{K}_\ell$), \mathcal{M} ($= \mathcal{K}_{\ell-1}/\mathcal{K}_\ell$), and terminal time, T .

- 1: set $\mathbf{Z}^C \leftarrow \mathbf{Z}(0)$, $\mathbf{Z}^F \leftarrow \mathbf{Z}(0)$, $t^C \leftarrow 0$ and $t^F \leftarrow 0$
- 2: for each R_j , calculate propensity values $p_j^C(\mathbf{Z}^C)$ and $p_j^F(\mathbf{Z}^F)$
- 3: set flag \leftarrow **false**, iter $\leftarrow 0$
- 4: **while** flag is **false** **do**
- 5: set iter \leftarrow iter + 1
- 6: generate $\Delta \sim \Gamma(\mathcal{K}, 1)$
- 7: set $t^C \leftarrow t^C + \Delta / \sum_{j=1}^M p_j^C$ and $t^F \leftarrow t^F + \Delta / \sum_{j=1}^M p_j^F$
- 8: **if** $\max\{t^C, t^F\} > T$ **then**
- 9: set $\mathcal{K} \sim \mathcal{B}(\mathcal{K} - 1, \min\{(T - t^C) / (\Delta / \sum_{j=1}^M p_j^C), (T - t^F) / (\Delta / \sum_{j=1}^M p_j^F) - 1\})$
- 10: set $\max\{t^C, t^F\} \leftarrow T$
- 11: set flag \leftarrow **true**
- 12: **end if**
- 13: set $K^* \leftarrow \mathcal{K}$
- 14: for each R_j , calculate probabilities b_j^1 , b_j^2 and b_j^3 according to Equations (24)
- 15: **for** $j = 1, \dots, M$ **do**
- 16: generate $K_{j1} \sim \mathcal{B}(K^*, b_j^1 / (1 - \sum_{j'=0}^{j-1} b_{j'}^1))$ and set $K^* \leftarrow K^* - K_j$
- 17: **end for**
- 18: set $K^{*2} \leftarrow K^*$ and $K^{*3} \leftarrow K^*$
- 19: **for** $j = 1, \dots, M$ **do**
- 20: generate $K_{j2} \sim \mathcal{B}(K^{*2}, b_j^2 / \sum_{j'=j}^M b_{j'}^2)$ and set $K^{*2} \leftarrow K^{*2} - K_{j2}$
- 21: generate $K_{j3} \sim \mathcal{B}(K^{*3}, b_j^3 / \sum_{j'=j}^M b_{j'}^3)$ and set $K^{*3} \leftarrow K^{*3} - K_{j3}$
- 22: **end for**
- 23: set $\mathbf{Z}^C \leftarrow \mathbf{Z}^C + \sum_{j=1}^M (K_{j1} + K_{j2}) \cdot \boldsymbol{\nu}_j$
- 24: set $\mathbf{Z}^F \leftarrow \mathbf{Z}^F + \sum_{j=1}^M (K_{j1} + K_{j3}) \cdot \boldsymbol{\nu}_j$
- 25: for each R_j , calculate propensity values $p_j^F(\mathbf{Z}^F)$
- 26: **if** \mathcal{M} divides iter, for each R_j , calculate propensity values $p_j^C(\mathbf{Z}^C)$
- 27: **end while**
- 28: **if** $t^C < T$ **then** use Algorithm 3 to simulate \mathbf{Z}^C until time T . **end if**.
- 29: **if** $t^F < T$ **then** use Algorithm 3 to simulate \mathbf{Z}^F until time T . **end if**.

tau-leap method (which takes 9.8 seconds to estimate \mathcal{Q}), but inferior to a SPM implementation (which we recall takes 2.1 seconds to estimate \mathcal{Q}). The optimal method will depend on the




Figure 7: The average CPU time required by the R-leap multi-level method to estimate $\mathbb{E}[X(3)]$ for System (23). We vary M and L ; the estimator is unbiased. The black bars indicate the range occupied by the 10-th to 90-th percentiles of the data.

particular reaction network and summary statistics of interest.

8. Discussion

At the outset, we presented two related aims for this manuscript: firstly, to improve the robustness of the multi-level method; and, secondly, to improve the computational performance of the multi-level method. Through the use of the SPM and the CPM approaches, the multi-level method has the potential to dramatically reduce the CPU time required to carry out Monte Carlo simulation. Whilst the CPM method does not always outperform the SPM method (although it often does), we have demonstrated its effectiveness and reliability. The R-leap multi-level method has been demonstrated as a new and effective simulation technique. The computational performance of the multi-level method may still depend on the particular reaction network and the chosen summary statistic, but our new methods provide additional tools for accelerating stochastic simulation.

In Case study 1, we demonstrated that, for the gene regulatory reaction network, the CPM is nearly 30% more efficient than the SPM. Furthermore, we noted that the kurtoses of levels $\ell = 1, \dots, L + 1$ are substantially higher when the SPM is used, when compared with the CPM. As explained in Section 6, the high kurtoses make it difficult to estimate the sample variances of the SPM method. There are two possible consequences:

- the sample variance, \mathcal{V}_ℓ , is an under-estimate. The effect is that the required confidence interval semi-length is not faithfully attained;

- the sample variance \mathcal{V}_ℓ , is an over-estimate. In this case, too many sample paths are generated, and the algorithm takes substantially longer to run.

A more robust approach is therefore provided by the CPM. Over a wide range of algorithm parameters, the CPM is able to outperform the SPM.

In Case study 2, the CPM clearly outperformed the SPM. The Lotka-Volterra model, System (22), was a particularly challenging test case, and we follow Gillespie [21] in explaining why. Gillespie [21] argues that if an ODE modelling approach is followed, then, in the X_1 - X_2 plane, the solution trajectories of System (22) are closed orbits. Gillespie [21] then says that the addition of microscopic fluctuations (due to using a stochastic model) induces a ‘drunkard’s walk’ over the continuum of deterministic orbits, thereby resulting in unstable behaviour. This feature makes it difficult to ensure that pairs of sample paths are tightly coupled. In particular, if a pair of sample paths differs slightly in their state vectors, then the difference in state vectors continues to increase with the SPM method. The CPM does not suffer the same defect, and, as such, it is more efficient than the DM.

In Case study 3, a case study of a logistic growth model was presented. In this case, the SPM outperformed the CPM. This is due to the choice of estimator, $\mathcal{Q} = \mathbb{E}[X(T)]$. By the terminal time, $T = 3$, X is fluctuating rapidly about a steady state. As such, the state vectors at earlier times, $X(t)$, tend to have a very limited effect on \mathcal{Q} , and so the CPM method is of limited benefit. If, however, one is concerned with the time-averaged population of System (21), given by

$$\mathcal{Q} = \mathbb{E} \left[\frac{1}{T} \int_0^T X(t) dt \right],$$

then the CPM outperforms the SPM (details not shown).

8.1. Comparing tau-leap and R-leap multi-level methods

By treating the time traversed by each algorithm step of the R-leap method, and the particular combination of reactions that take place during that step, as two distinct and unrelated quantities that must be determined by Monte Carlo simulation, we have described a new variance reduction technique for multi-level simulation. The R-leap multi-level method was assessed by considering Case study 3. We demonstrated that the R-leap approach performed slightly better than the CPM-driven, tau-leap multi-level method, but that it did not perform as well as the SPM-driven method. In summary, the R-leap multi-level method has the potential to provide good computational performance.

The crucial difference between the tau-leap and R-leap multi-level methods is seen by comparing Equations (17) and (24). The tau-leap method couples two distinct sample paths by

considering the difference in *absolute propensity values* of corresponding reaction channels in the sample paths, whilst the R-leap method considers the difference in the corresponding *proportion of the total propensity* (or probability) attributable to each reaction channel. Variance reduction in the tau-leap and R-leap multi-level techniques arises in a different format, and the computational performance is therefore different.

8.2. Outlook

The CPM provides a natural framework for implementing the multi-level method. It is able to mitigate some of the difficulties associated with the previously-used SPM implementation. Future work will include categorising reaction networks and summary statistics in order to derive criteria to decide whether the CPM or SPM should be used for that particular problem. Hybrid approaches that combine the SPM and CPM can also be implemented [22]. The R-leap method has been successfully implemented, and future work will determine the problems it is most suited to handling. Ultimately, a refined multi-level method will dramatically reduce the computational burden of Monte Carlo simulation.

Bibliography

- [1] Hanna, J., Saha, K., Pando, B., Van Zon, J., Lengner, C. J., Creyghton, M. P., van Oudenaarden, A., and Jaenisch, R. Direct cell reprogramming is a stochastic process amenable to acceleration. *Nature*, **462**(7273):595–601, 2009.
- [2] Lande, R., Engen, S., and Saether, B.-E. *Stochastic Population Dynamics in Ecology and Conservation*. Oxford University Press, 2003.
- [3] Thattai, M. and Van Oudenaarden, A. Intrinsic noise in gene regulatory networks. *Proceedings of the National Academy of Sciences*, **98**(15):8614–8619, 2001.
- [4] Van Kampen, N. G. *Stochastic Processes in Physics and Chemistry*. Elsevier, 1992.
- [5] Gillespie, D. T., Hellander, A., and Petzold, L. R. Perspective: Stochastic algorithms for chemical kinetics. *Journal of Chemical Physics*, **138**(17):170901, 2013.
- [6] Giles, M. B. Multilevel Monte Carlo path simulation. *Operations Research*, **56**(3):607–617, 2008.
- [7] Anderson, D. and Higham, D. Multi-level Monte Carlo for continuous time Markov chains, with applications in biochemical kinetics. *SIAM Multiscale Modeling and Simulation*, **10**(1):146–179, 2012.
- [8] Lester, C., Yates, C. A., Giles, M. B., and Baker, R. E. An adaptive multi-level simulation algorithm for stochastic biological systems. *Journal of Chemical Physics*, **142**(2):024113, 2015.
- [9] Moraes, A., Tempone, R., and Vilanova, P. Multilevel hybrid Chernoff tau-leap. *BIT Numerical Mathematics*, **56**(1):189–239, 2016.
- [10] Hammouda, C. B., Moraes, A., and Tempone, R. Multilevel hybrid split-step implicit tau-leap. *Numerical Algorithms*, **74**(2):527–560, 2017.
- [11] Lester, C., Baker, R. E., Giles, M. B., and Yates, C. A. Extending the multi-level method for the simulation of stochastic biological systems. *Bulletin of Mathematical Biology*, **78**(8):16401677, 2016.
- [12] Kurtz, T. G. Representations of Markov processes as multiparameter time changes. *Annals of Probability*, pages 682–715, 1980.

- [13] Gillespie, D. T. Exact stochastic simulation of coupled chemical reactions. *Journal of Physical Chemistry*, **81**(25):2340–2361, 1977.
- [14] Gillespie, D. T. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *Journal of Computational Physics*, **22**(4):403–434, 1976.
- [15] Gillespie, D. Approximate accelerated stochastic simulation of chemically reacting systems. *Journal of Chemical Physics*, **115**(4):1716–1733, 2001.
- [16] Auger, A., Chatelain, P., and Koumoutsakos, P. R-leaping: accelerating the stochastic simulation algorithm by reaction leaps. *Journal of Chemical Physics*, **125**(8):084103, 2006.
- [17] Li, T. Analysis of explicit tau-leaping schemes for simulating chemically reacting systems. *SIAM Multiscale Modeling and Simulation*, **6**(2):417–436, 2007.
- [18] Rathinam, M., Sheppard, P. W., and Khammash, M. Efficient computation of parameter sensitivities of discrete stochastic chemical reaction networks. *Journal of Chemical Physics*, **132**(3):034103, 2010.
- [19] Anderson, D. A modified next reaction method for simulating chemical systems with time-dependent propensities and delays. *Journal of Chemical Physics*, **127**(21):214107, 2007.
- [20] Murray, J. *Mathematical Biology*, volume 2. Springer, 2002.
- [21] Gillespie, D. Exact stochastic simulation of coupled chemical reactions. *Journal of Physical Chemistry*, **81**(25):2340–2361, 1977.
- [22] Lester, C., Yates, C. A., and Baker, R. E. Efficient parameter sensitivity computation for spatially-extended reaction networks. *Journal of Chemical Physics*, **146**(4):044106, 2017.