Journal of Computational Physics 346 (2017) 1–13

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

Coupling sample paths to the thermodynamic limit in Monte Carlo estimators with applications to gene expression

Ethan Levien *, Paul C. Bressloff

Department of Mathematics, University of Utah, Salt Lake City, UT 84112, USA

ARTICLE INFO

Article history:
Received 8 September 2016
Received in revised form 27 May 2017
Accepted 31 May 2017
Available online 13 June 2017

Keywords:
Chemical reaction networks
Monte Carlo
Variance reduction
Piecewise deterministic Markov process
Gillespie algorithm

ABSTRACT

Many biochemical systems appearing in applications have a multiscale structure so that they converge to piecewise deterministic Markov processes in a thermodynamic limit. The statistics of the piecewise deterministic process can be obtained much more efficiently than those of the exact process. We explore the possibility of coupling sample paths of the exact model to the piecewise deterministic process in order to reduce the variance of their difference. We then apply this coupling to reduce the computational complexity of a Monte Carlo estimator. Motivated by the rigorous results in [1], we show how this method can be applied to realistic biological models with nontrivial scalings.

1. Introduction

Large stochastic biochemical reaction networks are a popular modeling framework for investigating cellular processes [2], but the complexity and population sizes involved in realistic models pose major computational challenges. However, when there is a separation of scales, such models lend themselves to a number of model reduction techniques that are useful for course grained analysis. One example occurs when there is a separation in species abundances [3,4]. If some subset of chemical species in a reaction network are extremely abundant, then reaction channels involving those species will generally occur much faster than reactions involving less abundant species. Another examples occurs when the model parameters vary over many orders of magnitude. For example, even if a species is in a very low abundance, it is possible that the reaction rates are such that a certain reaction involving this species occurs on a different timescale than other reactions involving the same species. One approach to analyzing the qualitative properties of such a multiscale model involves rescaling the system and taking a thermodynamic limit to obtain a piecewise deterministic Markov process (PDMP). A number of recent studies have provided rigorous errors bounds for this type of reduction [5,3,6,7]. While the PDMP yields useful information about stochastic effects of the rare species, quantitative information about the stochastic fluctuations of the abundant species is lost. On the other hand, in many systems, particularly those with feedback between the rare and abundant chemical species, there is an interest in quantifying the stochastic effects due to these fluctuations [8]. A common method for resolving these fluctuations is the diffusion approximation. While the diffusion approximation is often thought to be computationally advantageous, recent work on classically scaled population models has shown that this method yields only moderate computational gains [9]. Moreover, the error between the PDMP and the exact model is fixed. However, it is sometimes desirable to control this quantity, especially when the separation of scales is only moderate.

* Corresponding author.
E-mail addresses: levien@math.utah.edu (E. Levien), bressloff@math.utah.edu (P.C. Bressloff).

http://dx.doi.org/10.1016/j.jcp.2017.05.050
0021-9991/

An alternative to multiscale reduction techniques is to develop methods for accelerating stochastic simulation algorithms such as the Gillespie algorithm [10–12]. For example, there have been numerous studies of the method of τ-leaping in an effort to accelerate simulations of continuous time Markov chains [13,14]. More recently, multi-level methods that couple τ-leaping approximations at different resolutions have been used to reduce variances in Monte Carlo estimators [15,16,9]. Variance reduction techniques that utilize probabilistic couplings have also appeared earlier in the context of stochastic differential equations (SDEs) and Markov Chain Monte Carlo methods [17,18]. While there has been some work that leverages multiscale reduction techniques for Monte Carlo estimators [19], to our knowledge the idea of using these techniques directly as a variance reduction tool has not been studied.

In this paper we explore the idea of coupling reduced models to exact models as a variance reduction tool for Monte Carlo estimators. We develop a new efficient Monte Carlo estimator for multiscale chemical reaction networks that are sufficiently near a thermodynamic limit. The key insight is that, since only a small fraction of the degrees of freedom of the PDMP are stochastic, one can efficiently compute statistics of the process. On the other hand, if one wants to resolve demographic noise in the full model it is necessary to perform a large number of Monte Carlo simulations. By coupling the full stochastic model to the PDMP one can reduce the variance by a factor inversely related to the system size, and hence a smaller number of simulations need to be performed to achieve a given error tolerance. For practical applications the desired error tolerance of the Monte Carlo estimator scales with this factor. Hence, the coupled Monte Carlo estimator has the potential to speed up computations by a fractional power of the error tolerance. Our results extend the idea of variance reduction developed in [15,16], and provide a new computational application of the theory developed in previous work on PDMP approximations, or partial thermodynamic limits [5,3,6].

The paper is organized as follows. In section 2 we introduce some background material related to mathematical modeling chemical process. In section 3 we derive an approximation by taking a thermodynamic limit of a multi-scale system. We also discuss our approach to simulating this process. Our approach to variance reduction is introduced in section 4, and in section 4.1 we present an algorithm for coupling the exact process to the multi-scale approximation. Finally, we apply our method to models of gene expression in section 5 where it is shown that significant computational gains can be made in comparison to crude Monte Carlo methods.

2. Background

2.1. Stochastic chemical reaction networks in the classical setting

We consider a system involving d chemical species, denoted $\mathcal{X} = \{\mathcal{X}_i\}_{i\in\mathcal{I}}$ with $\mathcal{I} = \{1, 2, \ldots, d\}$. The species interactions are prescribed by p reaction channels, denoted $\mathcal{R} = \{\mathcal{R}_j\}_{j\in\mathcal{J}}$ with $\mathcal{J} = \{1, 2, \ldots, p\}$. Let x_i be the number of $\mathcal{X}_i$ and set $x = (x_1, \ldots, x_d)$. Then the j-th reaction takes the form

$$\mathcal{R}_j : \sum_{i=1}^{d} K_{j,i}^{\text{in}} \mathcal{X}_i \longrightarrow \sum_{i=1}^{d} K_{j,i}^{\text{out}} \mathcal{X}_i,$$

where $K_{j,i}^{\text{in}}, K_{j,i}^{\text{out}}$ are known as *stochiometric coefficients*. When such a reaction occurs the state x is changed according to

$$x_i \to x_i + K_{j,i}, \quad K_{j,i} = K_{j,i}^{\text{out}} - K_{j,i}^{\text{in}}.$$

More complicated multi-step reactions can always be decomposed into these fundamental single-step reactions with appropriate stochiometric coefficients. In practice, most reactions involve collisions between pairs of molecules, so that $\sum_i K_{j,i}^{\text{in}} = 1$ or 2. In the so-called *classical setting*, the abundances of each species are assumed to be the same order of magnitude. Therefore, if all the abundances are large then we can describe the evolution of the system by a set of deterministic kinetic equations involving the scaled variables $z_i = x_i/S$. Here S is a dimensionless quantity representing the system size, which in gene networks is typically taken to be the characteristic number of proteins. Alternatively, it could represent some volume scale factor. For a set of p reactions, the kinetic equations take the form (strictly speaking in the thermodynamic limit $S \to \infty$)

$$\frac{dz_i}{dt} = \sum_{j=1}^{p} K_{j,i} \left[\kappa_j \prod_{l=1}^{d} z_l^{K_{j,l}^{\text{in}}} \right] \equiv \sum_{j=1}^{p} K_{j,i} \bar{\alpha}_j(z), \tag{2.1}$$

where κ_j is a constant that depends on the probability that a collision of the relevant molecules actually leads to a reaction. The product term is motivated by the idea that in a well-mixed container there is a spatially uniform distribution of each type of molecule, and the probability of a collision depends on the probability that each of the reactants is in the same local region of space. Ignoring any statistical correlations, the latter is given by the product of the individual components. The S independent functions $\bar{\alpha}_j$ are known as transition intensities or *propensities*. These intensities are the leading order term of the scaled transition intensities for a finite population ($S < \infty$), which can be derived combinatorially. These are given by