Journal of Computational Physics 266 (2014) 89–100

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

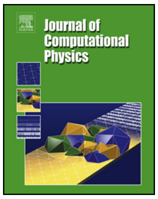

Local error estimates for adaptive simulation of the reaction–diffusion master equation via operator splitting

Andreas Hellander [a,*], Michael J. Lawson [b], Brian Drawert [b], Linda Petzold [b]

[a] *Department of Information Technology, Uppsala University, Box 337, SE-75105, Uppsala, Sweden*
[b] *Department of Computer Science, University of California Santa Barbara, Santa Barbara, CA 93106-5070, USA*

ARTICLE INFO

Article history:
Received 17 May 2013
Received in revised form 24 January 2014
Accepted 1 February 2014
Available online 20 February 2014

Keywords:
Reaction–diffusion master equation
Operator splitting
Local error
Adaptivity
Hybrid methods

ABSTRACT

The efficiency of exact simulation methods for the reaction–diffusion master equation (RDME) is severely limited by the large number of diffusion events if the mesh is fine or if diffusion constants are large. Furthermore, inherent properties of exact kinetic-Monte Carlo simulation methods limit the efficiency of parallel implementations. Several approximate and hybrid methods have appeared that enable more efficient simulation of the RDME. A common feature to most of them is that they rely on splitting the system into its reaction and diffusion parts and updating them sequentially over a discrete timestep. This use of operator splitting enables more efficient simulation but it comes at the price of a temporal discretization error that depends on the size of the timestep. So far, existing methods have not attempted to estimate or control this error in a systematic manner. This makes the solvers hard to use for practitioners since they must guess an appropriate timestep. It also makes the solvers potentially less efficient than if the timesteps were adapted to control the error. Here, we derive estimates of the local error and propose a strategy to adaptively select the timestep when the RDME is simulated via a first order operator splitting. While the strategy is general and applicable to a wide range of approximate and hybrid methods, we exemplify it here by extending a previously published approximate method, the diffusive finite-state projection (DFSP) method, to incorporate temporal adaptivity.

1. Introduction

To understand biological systems on the cellular level, it is often essential to account for the impact of noise due to small molecule count. For example, it has been demonstrated that stochasticity can have a profound effect on gene regulatory systems [23,8]. Spatial distribution of molecules in a cell can result in locally small populations of key chemical species, such that noise drives essential behavior, as in the case of symmetry breaking across many eukaryotic cell types [30]. Spatial stochastic modeling has already begun to yield new insights in systems such as spatiotemporal oscillators [10,27,26], MAPK signaling [28], self-organization of proteins into clusters [7] and polarization of proteins on the cell membrane in yeast [1]. Several modeling frameworks are used to model spatial stochastic system, the two most commonly used in

---

[*] Corresponding author.
E-mail address: andreas.hellander@it.uu.se (A. Hellander).

http://dx.doi.org/10.1016/j.jcp.2014.02.004

systems biology being continuous space Brownian Dynamics (BD) methods exemplified by GFRD [31] and the mesoscopic reaction–diffusion master equation (RDME), the latter being the focus of this paper.

In the traditional RDME, space is subdivided into subvolumes that can individually be treated as well-mixed. Reactions within a subvolume are expressed in the form of the chemical master equation (CME) [13] and realizations of the process can be generated using Gillespie's stochastic simulation algorithm (SSA) [12]. Molecules can move freely between neighboring voxels via diffusive jumps, which are modeled as linear jump events in a Markov process. Optimized exact simulation methods such as the next subvolume method (NSM) [7] can be used to generate statistically correct realizations of the RDME. As with all exact methods applied to RDME models, the NSM suffers from a potentially high computational cost due to having to explicitly simulate each diffusion event. The number of diffusive transfers between voxels grows rapidly as the mesh resolution is made finer, and as a result the majority of computation time tends to be spent on sampling diffusion events. Additionally, these methods are also inherently serial, which has thwarted attempts to increase efficiency via parallelization.

To speed up simulation of the RDME many methods rely on operator splitting. By splitting the operators, most often with a Lie–Trotter scheme [29], the reaction and diffusion steps can be solved independently. While diffusion carries the bulk of the computational burden in exact solvers, it is often possible to take advantage of the structure and linear nature of the discretized diffusion equation to speed up this step in an operator-split solver. Examples of approximate methods that have been proposed to speed up the simulation of the RDME by reducing the cost of the diffusive step include methods based on tau-leaping [25,21], the multinomial simulation algorithm [19], spatially adaptive hybrid methods [11] and the diffusive finite state projection method (DFSP) [6]. While splitting in itself does not resolve the issue of the inherent stiffness of the diffusion operator, the continued introduction of methods for simulating the RDME via operator splitting highlights the potential of this approach. Another recent use of operator splitting in the RDME context is the use of Lie–Trotter splitting to simulate fractional diffusion [2]. Yet another advantage of splitting is that it converts a largely serial problem, which is known to be difficult to parallelize in an efficient manner, into a naturally parallelizable one. For the existing approximate and hybrid methods for the RDME, splitting the physics (reaction and diffusion) is necessary. For parallel implementations, another possibility is to split the computational grid into blocks, as proposed by Arampatzis et al. [14], where the error introduced by operator splitting at block boundaries was analyzed. Our analysis here is different since it applies to the case of splitting the reaction and diffusion operators.

Splitting the operators introduces an error that depends on the size of the splitting time step, however previous algorithms that rely on operator splitting have not attempted an *a priori* error estimator. Without such an estimator these methods have no way to automatically control the splitting error. This limits their usefulness for practitioners. From an efficiency point of view, not knowing and controlling the error might lead to the use of unnecessary small timesteps at the price of slower simulations. In this work we seek to address theses issues by presenting estimators of the local error in probability, mean and variance for a first order Lie–Trotter splitting of the RDME. The estimators allow control of the splitting error for spatial stochastic simulation and enables any method based on operator splitting to be implemented adaptively.

This paper is organized as follows: in Section 2 we introduce the RDME. In Section 3 we outline how to simulate the RDME using operator splitting. We derive our estimator for the local error in the PDF, mean and variance, and demonstrate the accuracy of the estimator for an example problem. Finally, in Section 4 we present how the local error estimates can be used to extend an approximate method for the RDME to incorporate temporal adaptivity.

2. Spatial stochastic simulation using the RDME

Given a system with N_s chemical species X_s reacting in a physical domain Ω, discretize Ω into N_v non-overlapping voxels $\mathcal{V}_i$, with volume $|\mathcal{V}_i|$, and let the state of the system be described by the $(N_v \times N_s)$ state matrix $\mathbf{x}$, where the element x_{is} is the copy number of species X_s in voxel $\mathcal{V}_i$. Let $\mathbf{x}_{i,\cdot}$ denote the i-th row of $\mathbf{x}$ and $\mathbf{x}_{\cdot,s}$ the s-th column. The reaction network consists of N_r chemical reactions $r = 1 \dots N_r$. The $(1 \times N_s)$ stoichiometry vector $\mathbf{n}_{ir}$ describes the change in the state, $\tilde{\mathbf{x}}_{i,\cdot} = \mathbf{x}_{i,\cdot} + \mathbf{n}_{ir}$, when reaction r occurs in voxel i and the propensity function for that reaction is $a_{ir}(\mathbf{x}_{i,\cdot})$.

Diffusion of species X_s along the edge (2D) or face (3D) connecting voxels $\mathcal{V}_i$ and $\mathcal{V}_j$ is modeled as a linear jump event, or diffusive transfer,

$$X_{is} \xrightarrow{d_{ijs}} X_{js}, \tag{1}$$

with propensity function $\mu(\mathbf{x}) = d_{ijs} x_{is}$. The change in state due to the diffusive transfer is described by the $(N_v \times 1)$ stoichiometry vector ν_{ijs}, such that the new state is $\tilde{\mathbf{x}}_{\cdot,s} = \mathbf{x}_{\cdot,s} + \nu_{ijs}$. All entries of ν_{ijk} are zero except $\nu_{ijs}(i) = -1$ and $\nu_{ijs}(j) = 1$. In the case of reactions only, the probability density function $p(\mathbf{x}, t) \equiv p(\mathbf{x}, t | \mathbf{x}_0, 0)$ obeys the master equation

$$\frac{\mathrm{d}}{\mathrm{dt}} p(\mathbf{x}, t) = \mathcal{M} p(\mathbf{x}, t) \equiv \sum_{i=1}^{N_v} \sum_{r=1}^{N_r} a_{ir}(\mathbf{x}_{i\cdot} - \mathbf{n}_{ir}) p(\mathbf{x}_{1\cdot}, \dots, \mathbf{x}_{i\cdot} - \mathbf{n}_{ir}, \dots, \mathbf{x}_{N\cdot}, t) - \sum_{i=1}^{N_v} \sum_{r=1}^{N_r} a_{ir}(\mathbf{x}_{i\cdot}) p(\mathbf{x}, t). \tag{2}$$

For the case of one subvolume, $N_v = 1$, (2) reduces to the CME for a well stirred system. For a system with only diffusion, the master equation takes the form