Journal of Computational Physics 328 (2017) 438–454

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

Information criteria for quantifying loss of reversibility in parallelized KMC

Konstantinos Gourgoulias *, Markos A. Katsoulakis, Luc Rey-Bellet

Department of Mathematics and Statistics, Lederle Graduate Research Tower, University of Massachusetts, 710 N. Pleasant Street, Amherst, MA, 01003-9305, United States

ARTICLE INFO

Article history:
Received 1 May 2016
Received in revised form 13 September 2016
Accepted 12 October 2016
Available online 19 October 2016

Keywords:
Parallel kinetic Monte Carlo
Operator splitting schemes
Long-time errors
Time-reversibility
Detailed balance
Entropy production
Information criteria

ABSTRACT

Parallel Kinetic Monte Carlo (KMC) is a potent tool to simulate stochastic particle systems efficiently. However, despite literature on quantifying domain decomposition errors of the particle system for this class of algorithms in the short and in the long time regime, no study yet explores and quantifies the loss of time-reversibility in Parallel KMC. Inspired by concepts from non-equilibrium statistical mechanics, we propose the entropy production per unit time, or entropy production rate, given in terms of an observable and a corresponding estimator, as a metric that quantifies the loss of reversibility. Typically, this is a quantity that cannot be computed explicitly for Parallel KMC, which is why we develop *a posteriori* estimators that have good scaling properties with respect to the size of the system. Through these estimators, we can connect the different parameters of the scheme, such as the communication time step of the parallelization, the choice of the domain decomposition, and the computational schedule, with its performance in controlling the loss of reversibility. From this point of view, the entropy production rate can be seen both as an information criterion to compare the reversibility of different parallel schemes and as a tool to diagnose reversibility issues with a particular scheme. As a demonstration, we use Sandia Lab's SPPARKS software to compare different parallelization schemes and different domain (lattice) decompositions.

1. Introduction

Kinetic Monte Carlo, also known as the n-fold way [1] or the Bortz–Kalos–Lebowitz algorithm [2], is a common tool among practitioners interested in simulating stochastic processes arising from chemical, biological, or agent-based models on lattices [3]. However, even sophisticated algorithms inevitably experience slowdown as the size of the system increases. In fact, it may even be the case that the system size prohibits the use of a serial simulation, for instance due to problems with storing the system in a single CPU's memory.

There is a substantial amount of work in addressing the efficiency issues when simulating larger time and length scales by using parallel algorithms for systems with either short-range [4–6] or long-range interactions [7,8]. Typically, those algorithms are based on domain decomposition of the lattice into sub-domains (see Fig. 1), and subsequently the simulation on each sub-domain according to a chosen computational schedule. One such algorithm is part of Sandia Labs' SPPARKS Monte Carlo code [9]. A new insight from [10] was that such parallel KMC algorithms that depend on short-range interactions

* Corresponding author.
E-mail addresses: gourgoul@math.umass.edu (K. Gourgoulias), markos@math.umass.edu (M.A. Katsoulakis), luc@math.umass.edu (L. Rey-Bellet).

http://dx.doi.org/10.1016/j.jcp.2016.10.031
0021-9991/

can be formulated as operator splitting schemes that approximate the exact process. This mathematical formulation allows both for performance and numerical error analysis of the schemes [11]. It was also leveraged in previous work [12] where, combined with information metrics, it allowed us to study the long time error behavior of the schemes.

In fact, the investigation of long-time errors for operator splitting schemes is of prime importance when using parallel KMC, as errors accumulate due to the domain decomposition procedure. This accumulation can affect the simulation dramatically at long times and make it uncertain for practitioners to sample from the correct stationary regime. Unfortunately, classical numerical analysis fails to quantify errors of splitting schemes, such as parallel KMC for long times, which in turn motivated our use of the relative entropy per unit time as a tool to study the performance of operator splitting schemes [12].

Another aspect of long-time behavior, and the focus of this work, is on systems with time-reversible dynamics. That symmetry is often an integral part of the physical structure of the model, for example in the simulation of interacting diffusions or adsorption/desorption mechanisms. While in such cases the time-reversal symmetry is preserved under the serial KMC simulation (typically by enforcing the detailed balance condition), the time-discretization, domain decomposition, and breakdown of serial communication of the parallelized algorithm may lead to loss of detailed balance, and thus of reversibility. There exists some literature on constructing parallel algorithms that preserve the detailed balance (DB) condition [13]. In those algorithms, the scheme picks a schedule for sweeping over the lattice sub-domains, executes it by simulating each sub-domain forward in time for a fixed number of time steps according to the schedule, and then picks a new schedule. For the adjustment to the correct timescale, computation of an equilibrium autocorrelation function is also required. Although such schemes resemble the random Lie–Trotter splitting [11] and they can be numerically analyzed in a similar manner, we will not discuss them here, mainly due to the technical differences with schemes that employ a fixed computational schedule [10]. More specifically, our focus is on general partially asynchronous parallel algorithms, like the one in SPPARKS. For those, a user has to pick between different domain decompositions, time steps of the scheme, and a fixed schedule. These choices will impact the loss of reversibility of the scheme. Therefore, it makes sense to develop a theory that can connect the various parameters of the scheme with loss of reversibility.

Regarding the loss of reversibility of numerical schemes, in [14] the authors used the entropy production rate (EPR) as an information metric to quantify the loss of reversibility for the Euler–Maruyama and Milstein schemes for stochastic differential equations, as well as BBK schemes for Langevin dynamics. This idea was motivated by concepts in non-equilibrium statistical mechanics, originally developed to understand the long-time dynamics and the fluctuations in non-equilibrium steady states [15–19]. The authors in [14] used such non-equilibrium statistical mechanics methods as computable numerical tools to assess the loss of reversibility of numerical schemes for SDEs. More specifically, they computed the EPR with the Gallavotti–Cohen action functional [17] as an estimator for different numerical schemes. It was demonstrated that the scheme performance in controlling the loss of reversibility can vary greatly. In particular, the Euler–Maruyama scheme for SDEs with multiplicative noise can break reversibility in an unrecoverable manner regardless of the size of the time step [14, Theorem 3.7].

Our goal here is to apply a similar perspective for the study of splitting schemes in parallel KMC. However, in contrast with schemes for SDEs, for the class of systems that we can simulate in this manner the transition probabilities are either difficult to compute or not available at all. Because of this, a new approach is required, which is why we write the EPR as an asymptotic expansion in the scheme's time step by using the semigroup theory for Markov chains. We demonstrate that the coefficients of the expansion of the EPR depend on the transition rates of the model and can be estimated as ergodic averages by samples from the parallel algorithm. We also show that the required computations for the estimation of the coefficients scale with the size of the boundary between sub-domains on the lattice in a manner that depends on the scheme selected. Therefore, by appropriate normalization, we can calculate the entropy production rate per lattice site, i.e. independent of system size. As a result, we obtain an *a posteriori* expansion for the estimator of the EPR, which can be used as a diagnostic tool that can be calculated on a system of smaller size than the targeted one, and/or even ran with a simple serial implementation of the parallel algorithm.

This information-theoretical perspective is similar to the use of information metrics to assess discrepancy of models, algorithms, approximations, etc., to that applied to the study of long time errors for Parallel KMC [12], sensitivity analysis [20], and in studying loss of information due to coarse-graining in non-equilibrium systems [21].

The manuscript is organized as follows. In Section 2, we provide an introduction to Parallel Lattice KMC and the ideas behind its error analysis based on operator splitting. Section 3 is especially important, as we introduce the entropy production and entropy production per unit time. Those concepts will be the information-theoretical quantities used to study loss of reversibility for operator splitting schemes. Then, in Section 4, we discuss the estimation of the EPR, referring to specific examples and an implementation in SPPARKS. We use the estimates to compare two splitting schemes as well as discuss their loss of reversibility with respect to lattice decomposition and time step. Finally, in Section 5, we provide general results for the asymptotic behavior of EPR, and deduce from them estimators.

2. Background on parallel lattice KMC

Parallel Lattice KMC is an approximation to the exact, but serial, simulation algorithm. In implementations, it works by taking advantage of the spatial dependencies between the different events. For example, in a model with finite range interactions, the spins on two lattice sites can change with no error to the dynamics as long as the two are sufficiently far