ARTICLE IN PRESS

Journal of Computational Physics xxx (2013) xxx–xxx

Contents lists available at SciVerse ScienceDirect

Journal of Computational Physics

journal homepage: www.elsevier.com/locate/jcp

Multiscale stochastic simulations of chemical reactions with regulated scale separation

Petros Koumoutsakos *, Justin Feigelman

Chair of Computational Science, Clausiusstrasse 33, ETH Zurich, CH-8092, Switzerland

ARTICLE INFO

Article history:
Available online xxxx

Keywords:
Stochastic simulation algorithm
Stiff systems
Heterogeneous multiscale methods
Flow averaging integrators
Cutoff phenomenon

ABSTRACT

We present a coupling of multiscale frameworks with accelerated stochastic simulation algorithms for systems of chemical reactions with disparate propensities. The algorithms regulate the propensities of the fast and slow reactions of the system, using alternating micro and macro sub-steps simulated with accelerated algorithms such as τ and R-leaping. The proposed algorithms are shown to provide significant speedups in simulations of stiff systems of chemical reactions with a trade-off in accuracy as controlled by a regulating parameter. More importantly, the error of the methods exhibits a cutoff phenomenon that allows for optimal parameter choices. Numerical experiments demonstrate that hybrid algorithms involving accelerated stochastic simulations can be, in certain cases, more accurate while faster, than their corresponding stochastic simulation algorithm counterparts.

1. Introduction

Stochastic simulation algorithms (SSA) have become indispensable in modeling the evolution of chemical reactions as they pertain to systems biology [1] in particular when the number of molecules involved in these reactions is small to moderate [2,3]. These reactions are governed by the Chemical Master Equation (CME) [4] and while SSA reduces the computational complexity of the CME it remains prohibitively expensive for certain systems as it implies the computation of every reaction event. In recent years a number of algorithms have been developed to accelerate SSA. The most prominent is τ-leaping [5], which introduces an *a priori* specified time step compatible with the *leap condition*, and a corresponding Poisson process to compute the firings across reaction channels. The R-leaping algorithm of [6,7], specifies instead the total number of firings across all reaction channels and the time step is derived from a Gamma distribution while the firings over the reaction channels are sampled from correlated Binomial distributions. Both algorithms operate under the assumption that propensities remain approximately constant during each step while R-leaping gains further computational speedup by exploiting the stiffness of systems with large propensity disparities.

Overcoming stiffness in chemical reaction networks is the subject of much current research in the area of systems biology. Stiff systems of reactions are common in important biological models, for example in a gene-regulatory pathway that involves very fast protein binding kinetics and relatively slow mRNA transcript production. A number of techniques have been developed in order to reduce the computational cost of evolving a stiff system by distinguishing between fast and slow reactions. These algorithms include the Slow-Scale method [8], the nested stochastic simulation algorithm [9], the Seamless Method [10], the Hybrid Multiscale Monte Carlo [11], and others [12]. In a related context generalized frameworks such as the Heterogeneous Multiscale Modeling (see [13] and references therein) and Equation Free [14] originally developed for handling stiffness in systems of ordinary differential equations and coupling solvers operating at different time scales,

* Corresponding author. Tel.: +41 44 632 5258; fax: +41 44 632 1703.
E-mail address: petros@ethz.ch (P. Koumoutsakos).

0021-9991/$ - see front matter © 2012 Elsevier Inc. All rights reserved.
http://dx.doi.org/10.1016/j.jcp.2012.11.030

have introduced the concept of solving fast and slow reactions separately in properly-adjusted alternating steps with some method-specific coupling between scales. These methods can be naturally extended to hybrid algorithms that couple multiscale simulations frameworks, such as the HMM, with stochastic simulations algorithms. A key concept in the hybrid algorithms is the identification and differential processing of the fast and slow reactions of the system as determined dynamically by the reaction propensities during the evolution of the system. In this context, E et al. [9,13] have used drift averaging techniques [15] to accelerate stochastic simulations while Bayati et al. extended the related flow averaging principle, called *FLow AVeraging integratORS* (FLAVORS) [16] to SSA (FLAVOR-SSA) [17].

In this work, we couple multiscale simulation frameworks with the τ and R-leaping accelerated stochastic simulation techniques. To the best of our knowledge this coupling is presented here for the first time.The coupling relies on a regulation parameter to scale the propensities of the fast and slow reactions and varying the number of microsteps inside each macrostep of the simulation. The proposed algorithms are a direct extension of FLAVOR-SSA [17] and have similarities to the slow-scale stochastic simulation algorithm [8], the (hybrid) MSMC algorithm [18,11], the seamless [10] and nested stochastic simulation [9] algorithms, the equation free SSA [19] and others. The fundamental difference from previous methods is the rescaling of the fast reaction propensities during the macroscopic simulation steps, which no previous methods except for FLAVOR-SSA have included, and the novel coupling of multi scale frameworks with leaping algorithms. We remark that the results of the simulations exhibit a cutoff phenomenon for the regulating parameter of the fast reactions, similar to what has been observed in FLAVOR-SSA [17], implying an optimal trade-off of accuracy and speedup.

The paper is organized as follows: In Section 2 we describe the accelerated multi scale algorithms, followed by results (Section 3) that are further discussed in Section 4 and closing with a summary and an outline of future work.

2. Accelerated τ/R-leaping algorithms as micro/macrosteppers in multiscale frameworks

Stochastic simulation algorithms [20,21] are Kinetic Monte Carlo Methods realizing reaction events governed by the chemical master equation. The algorithms simulate all events in a system of N reaction channels ($R_k,\ k = 1,\ldots,N$) each characterized by a propensity a_k. In SSA the reaction rates a_k are defined by

$$a_k(\mathbf{X}) := c_k \Omega \prod_{j=1}^{J} \left\{ \frac{X_j(X_j-1)\cdots(X_j-r_j+1)}{\Omega^{r_j}} \right\}, \tag{1}$$

where c_k is the reaction rate, Ω is the system volume, X_j is the number of molecules of species j, r_j (a constant) is the number of reactants of the species j. The discrete-state of the system is denoted by $\mathbf{x} \in \mathbb{N}^M$, where N is the number of species in the system. The SSA may be expressed as follows:

0. Initialize the time $t = t_0$ and the system's state $\mathbf{X} = \mathbf{X}_0$.
1. For the system in state $\mathbf{x}$ at time t, evaluate all the partial reaction channel propensities $a_j(\mathbf{X})$ and their total $a_0(\mathbf{X}) := \sum_{j=1}^{N} a_j(\mathbf{X})$.
2. Generate time step τ and identify reaction channel j. For τ sample an exponential random variable with parameter $a_0(\mathbf{X})$ and for j sample a discrete random variable on $\{1,\ldots,N\}$ with $\mathbb{P}[j = k] = \frac{a_k(\mathbf{X})}{a_0(\mathbf{X})}$.
3. Perform a reaction by replacing $t \to t + \tau$ and $X \to X + \boldsymbol{\nu}_j$ where $\boldsymbol{\nu}_j$ is the stoichiometric vector that denotes the change induced by reaction j.
4. Record $(\mathbf{X}, t)$. Return to step 1, or end.

The SSA generates random variables $\mathbf{X}(t)$ that sample from the probability distribution $\mathbf{P}(\mathbf{X}_0, t)$ which is the solution of the master equation for the chemical species at time t. The accuracy of the SSA realizations are however coming at the expense of its computational cost.

In order to accelerate SSA a number of algorithms have been proposed. The first algorithm introduced in this context is the τ-leaping [5], which under the assumption of unchanged propensities for the reaction channels (the *leap condition*), samples M independent Poisson distribution to compute the firings across reaction channels. An alternative approach is R-leaping [6], where the number of reactions across all channels is specified and N correlated Binomials are sampled to define the number of reactions in each channel. The accuracy of both techniques hinges on the validity of the assumption of unchanged propensities, while we note that R-leaping can be advantageous in systems with disparate reaction rates. This advantage stems from the fact that the algorithm is independent from the order of reactions and by sorting the reactions according to their propensity the number of specified reactions can be satisfied by sampling only a small subset of reactions. We refer the reader to the original articles of [5,6] as well to works that describe algorithms that ensure the validity of the *leap condition* [8,22].

One of the key concepts of the Multiscale Modeling frameworks for systems of differential equations, such as HMM [13] and Equation Free [14] is the alternation between steps that handle the fast and slow reactions of the system. In this context, the system of N reactions R is separated into fast reactions $R^f := \{R_i^f\}_{i=1}^{N^f}$ and slow reactions $R^s := \{R_i^s\}_{i=1}^{N^s}$, $N = N^f + N^s$, $R = R^f \cup R^s$ with corresponding reaction propensities $a^f := \{a_i^f\}_{i=1}^{N^f}$ and $a^s := \{a_i^s\}_{i=1}^{N^s}$ defined so that the scale separation $\frac{1}{\epsilon} := \frac{\min(a^f)}{\max(a^s)} \geqslant \frac{\alpha}{\beta}, \forall \alpha, \beta \in a^f \cup a^s$ is maximal.

Please cite this article in press as: P. Koumoutsakos, J. Feigelman, Multiscale stochastic simulations of chemical reactions with regulated scale separation, J. Comput. Phys. (2013), http://dx.doi.org/10.1016/j.jcp.2012.11.030

The algorithm alternates between microscopic and macroscopic stages and the microscopic stage is further divided into $M \geqslant 1$ substeps. Two regulation parameters are introduced: during the microscopic step fast reactions are rescaled by a factor $\xi \in (0,1)$ while during the microscopic stage, slow reactions are controlled by a factor $\chi \in \{0,1\}$. During each substep, either one τ-leaping or R-leaping (or SSA) update is performed using the rescaled propensities, and the system molecularities and time is updated accordingly. The state of the system following each microscopic step is recorded into a temporary buffer, and a weighted average of the buffered states is used to compute the reaction propensities for the subsequent macroscopic stage. The hybrid multiscale algorithm is summarized as follows:

0. Initialize the time $t = T_0$ and the state of the macroscopic system $x(T_0)$.
1. Initialize the microscopic system with the current value of the macroscopic system: $x_0 \leftarrow x(t)$.
2. For $k = 1, \ldots, N$:
   (a) Compute the reaction propensities of the current microstate $a(x_k) = a^f(x_k) \cup \chi a^s(x_k)$, and the rescaled total reaction propensity $\tilde{a}_0 = \sum_{a \in a^f} a(x_k) + \chi \sum_{a \in a^s} a(x_k), \chi \in \{0,1\}$
   (b) Compute Δx and $\tau(x_k)$ depending on the method used (SSA, τ,R-leaping)
   (c) Fire the reaction(s), and update the microscopic system: $x_{k+1} \leftarrow x_k + \Delta x, t \leftarrow t + \tau$
3. Compute the averaged species values: $\bar{x} = \sum_{i=1}^{M} K_i x_i$, with $\sum_{i=1}^{M} K_i = 1$
4. Compute the reaction propensities for the macroscopic system using the averaged species values: $a_j(\bar{x})$
5. Rescale the fast reactions by a factor ξ such that $a(x(t)) = \xi a^f \cup a^s$ and the total reaction propensity $\tilde{a}_0 = \xi \sum_{a \in a^f} a(x_k) + \sum_{a \in a^s} a(x_k)$, $\xi \in (0,1)$
6. Compute Δx and $\Delta t(x(t))$ depending on the method used (SSA, τ,R-leaping)
7. Fire the reaction(s), and update the macroscopic system: $x(t + \Delta t) \leftarrow x(t) + \Delta x$, $t \leftarrow t + \Delta t$
8. Repeat from 1 until final time is reached

The methods are differentiated by the variables ξ and χ which regulate the rescaling in the different time scales, and a vector of weights K. Negative molecularities may be produced by the leaping algorithms and are managed using the τ selection algorithm of Cao et al. [23], or by choosing a small value for ϵ and θ for τ and R-leaping [6], respectively.

Three special instances of the general algorithm described above are designated as HMM-S, FLAVOR-S, and Extended-FLAVOR-S (EF-S) where S is replaced by τ, R or SSA depending on the particular implementation. The parameters for each method as implemented in this study, based on parameters reported in the referenced publications are summarized in Table 1. It is important to note, in particular for the HMM methods, that a number of versions have been reported [24,13,10] that have studied variants of the parameters described in Table 1. We refer the reader to a recent review [25] and references therein, for a thorough discussion. Here for HMM we consider the version presented in [13,10].

3. Results

The methods were applied to test-systems of differing size, time-dependence and magnitude of scale separation, and number of steady states. Performance in each case was evaluated in terms of the root mean square of the error in the average molecularity of a sample species relative to the same species in SSA.

Rapid isomerization and reaction system (RIR). This systems consists of two species S_1 and S_2, and four reactions with a pair of rapid reversible isomerization reactions, and slow production and decay of S_2 [17]:

$$R_1 : S_1 \xrightarrow{2000} S_2$$
$$R_2 : S_2 \xrightarrow{2000} S_1$$
$$R_3 : \emptyset \xrightarrow{100} S_2$$
$$R_4 : S_2 \xrightarrow{1} \emptyset$$

Under the initial conditions of 50 molecules for S_1 and S_2, the propensities of reactions R_1 and R_2 are 1000 times as large as the propensity of R_3 giving an initial scale separation of $1/\epsilon = 1000$, with large scale separation persisting throughout the

Table 1
Summary of parameters for three versions of multiscale algorithms, as implemented in this study.

	M	ξ	χ	K_i
HMM [13]	$\geqslant 1$	0	0	1/M
EF [14]	>1	(0,1)	1	$\delta_{i,M}$
FLAVOR [16]	1	(0,1)	1	1

Please cite this article in press as: P. Koumoutsakos, J. Feigelman, Multiscale stochastic simulations of chemical reactions with regulated scale separation, J. Comput. Phys. (2013), http://dx.doi.org/10.1016/j.jcp.2012.11.030

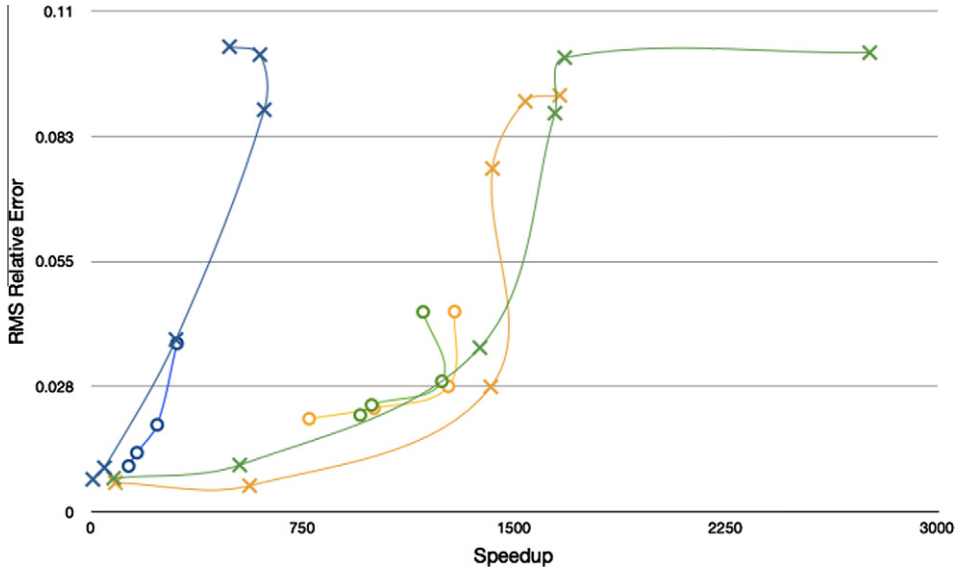

Fig. 1. RIR Performance: Speedup versus error in the RIR system for FLAVOR (X), and HMM (O) with $\xi = 10^{-j}, j = 1, \ldots, 6$, and $M \in \{3, 6, 9, 12\}$. SSA is shown in blue, τ-leaping in green, R-leaping in orange. Large values of ξ and M are towards the lower-left, and small values are towards the top-right. The markers denote simulated points and the lines are second-order interpolations. The optimal choice is $\xi \approx 10^{-2}$ where error is negligible. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

time of interest. 10,000 samples were simulated for $\{\text{FLAVOR}, \text{HMM}\} \otimes \{\text{SSA}, \tau, \text{R}\}$ with $\xi = 10^{-j}$ for $j = 1, \ldots, 6$, and $M \in \{3, 6, 9, 12\}$.

The speedup and relative error versus SSA were computed as a function of ξ and M for FLAVOR and HMM respectively, shown parametrically in Fig. 1. The largest values of ξ, M are in the lower left near the origin, and the smallest values of ξ, M are near the top-right with the most speedup and error.

The cutoff phenomenon of FLAVOR described in [17] is identifiable in Fig. 1 by the strongly sigmoidal character of the plots for FLAVOR, particularly for FLAVOR-R and FLAVOR-τ. Also of note is the coincidence of HMM-SSA with FLAVOR-SSA and HMM-τ (R) with FLAVOR-τ (R). Each FLAVOR method shows the same qualitative trend as the corresponding HMM, but with increased domain and range. The speedup of HMM is likely hampered by the constraint that at least 1 micro-step per macrostep is performed, as per the definition of HMM. R-leaping and τ-leaping produce qualitatively similar speed-up/error profiles, likely because no benefit is achieved from ranking the propensities in the R-leaping scheme when there are so few reactions. For comparison, plain τ-leaping and R-leaping produce speedups of about 7× and 6× compared to SSA, respectively, with relative errors less than 1% (not shown).

3.1. Time adaptive system (TA)

The system consists of 4 species with 6 reactions, and has a scale separation that emerges over time due to the accumulation of molecules.

$$R1 : S_1 \xrightarrow{1} S_2$$
$$R2 : S_2 \xrightarrow{1} S_1$$
$$R3 : S_1 \xrightarrow{1} S_3$$
$$R4 : S_3 \xrightarrow{1} S_1$$
$$R5 : 2S_2 + S_3 \xrightarrow{4} 3S_4$$
$$R6 : 3S_4 \xrightarrow{12} 2S_2 + S_3$$

At the start of the simulation there are 100 molecules of S_1, and 3 molecules of S_2, S_3 and S_4. Reaction propensities are initially 100 for a_1, a_3, 3 for a_2, a_4, 72 for a_5, and 108 for a_6. Thus initially there is little scale separation, but as molecules of S_2 and S_3 accrue, they drive forward reaction R_5. The simulation software periodically samples the scale separation $1/\epsilon$, and switches to the multiscale algorithm when this exceeds 300.

Simulations were run for HMM, FLAVOR, and Extended-FLAVOR, with 2,000 samples for each method using SSA, τ-leaping and R-leaping, with $\xi = 10^{-j}, j = 1, \ldots, 6$ and $M \in \{3, 6, 9, 12\}$. The error control parameters were set to $\epsilon = 0.1$, and $\theta = 0.1$.

Please cite this article in press as: P. Koumoutsakos, J. Feigelman, Multiscale stochastic simulations of chemical reactions with regulated scale separation, J. Comput. Phys. (2013), http://dx.doi.org/10.1016/j.jcp.2012.11.030

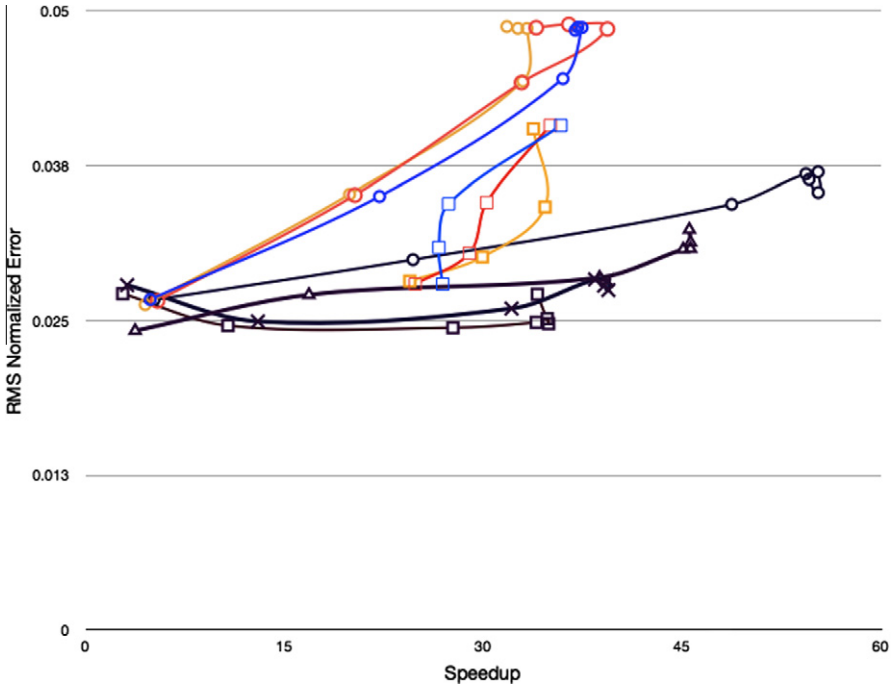

Fig. 2. Error versus speedup for the TA system. $\xi = 10^{-j}, j = 1, \ldots, 6$, $M \in \{3, 6, 9, 12\}$, with FLAVOR (∘), HMM (□). FLAVOR-SSA and HMM-SSA is shown in blue, τ in red, and R in orange. EF-SSA is shown in purple with M = 3 (∘), M = 6 (△), M = 9 (×), and M = 12 (□). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Fig. 2 summarizes the error and speedup relative to SSA for each algorithm, with error defined as the root mean square error in the molecularity of species 1 relative to SSA. EF shows an almost linear relationship between error and speedup, with error and speedup negatively correlated with ξ and M. EF provides the best overall performance, achieving speedups of about 50 with an error of about 4% for $M = 3$.

In this system, R-leaping and τ-leaping provide no speedup over SSA, which is likely due to the system's low molecularity. When there are so few molecules, no significant leap can be performed without changing the molecularities or propensities substantially; thus, these methods degenerate to a sort of inefficient SSA. Indeed the R-leaping without any accelerations is about 40% as fast as plain SSA, while modified τ-leaping [22] is marginally faster (not shown). Since R-leaping and τ-leaping were found to be inappropriate, no simulations were performed combining EF with the leaping algorithms.

3.2. Multiscale Schlögl Model

Samant et al. [11] introduced the multiscale Schlögl model, consisting of 2 species X and Y. Species X undergoes rapid reversible reactions and a slow reversible conversion to species Y:

$$R_1 : 2X \xrightarrow{144} 3X$$
$$R_2 : 3X \xrightarrow{6} 2X$$
$$R_3 : \emptyset \xrightarrow{5000} X$$
$$R_4 : X \xrightarrow{1276} \emptyset$$
$$R_5 : X \xrightarrow{0.0005} Y$$
$$R_6 : Y \xrightarrow{0.0001} X$$

Initial conditions are $X(0) = 10$, $Y(0) = 100$. When treated deterministically, the system has multiple steady states, $X^1_{\text{Eqm}} = 5$, $X^2_{\text{Eqm}} = 20$, $X^3_{\text{Eqm}} = 50$, with steady states 1 and 3 being stable, and 2 unstable. ODE solvers will typically arrive at one of the steady states, dependent on initial conditions, and fail to capture the system's inherent bistability. While the full SSA is able to capture the system's bistability, drift-averaging techniques such as Slow-Scale SSA [8] may not [11]. As shown in Fig. 3, FLAVOR-SSA is able to roughly reproduce the results of the full SSA, including the bistability of X.

The histogram of species X at time $t = 1000$ was compared for some methods with $\xi = 10^{-4}$ in each case, and $M = 10$ for EF and HMM (Fig. 4). 1000 samples were simulated for {FLAVOR, HMM, EF} $\otimes$ {SSA,τ,R} with $\xi = 10^{-j}, j \in \{4, 5, 6, 7\}$ and

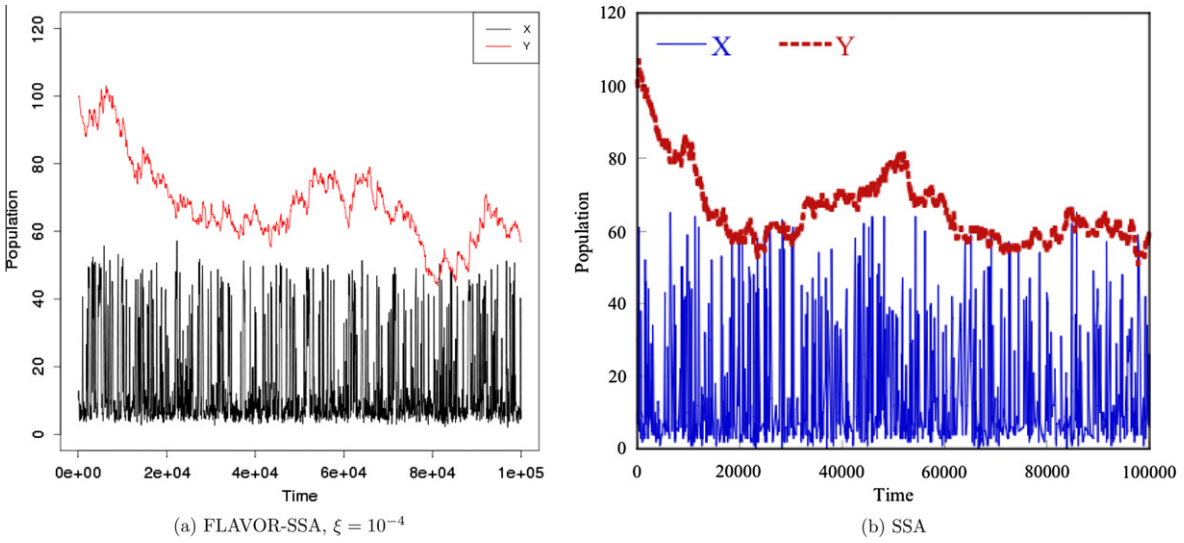

(a) FLAVOR-SSA, $\xi = 10^{-4}$ (b) SSA

Fig. 3. Trajectory of species X and Y in Schlögl system.

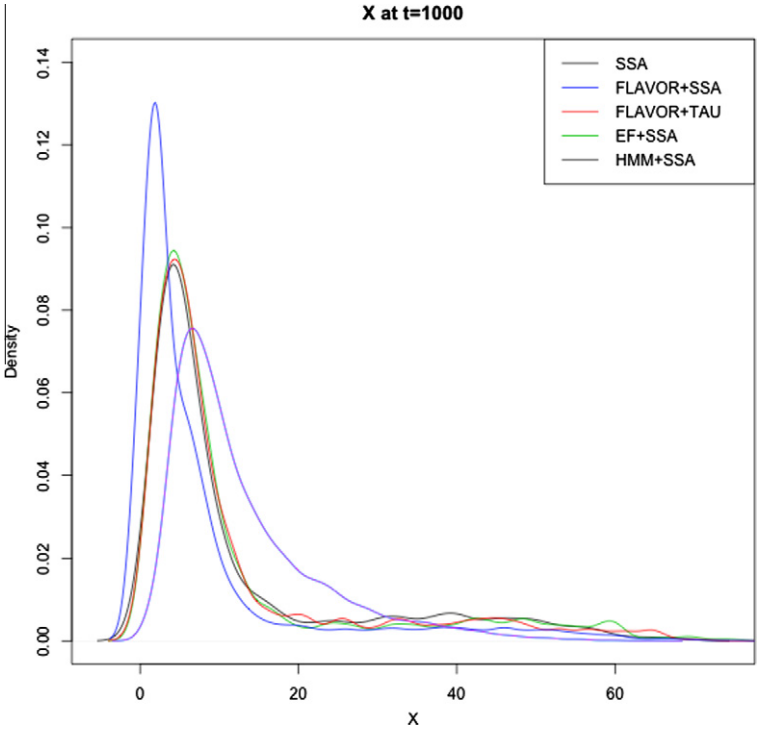

Fig. 4. Histogram for species X at time $t = 1000$ in the Schlögl system. $\xi = 10^{-4}$, $M = 10$.

$M \in \{1, 2, 5, 10, 15, 25, 50\}$. FLAVOR-$\tau$ and EF-SSA seem to reproduce SSA nearly exactly. However, FLAVOR-SSA (blue), and HMM-SSA (purple) are clearly shifted.

Analysis of the algorithms' performance reveals some interesting features in the (Fig. 5) errors reported as the L^2 norm of the histogram distances for species X. Firstly, there is little difference between τ-leaping and SSA for the EF-S and HMM-S simulations, as evidenced by the proximity of the solid red/blue circles, and the open red/blue symbols in Fig. 5. Due to the small molecularity of X, τ-leaping does not provide any speedup over SSA. Secondly, in some cases FLAVOR-τ is more accurate than FLAVOR-SSA, which is discussed below. Lastly, EF-SSA seems to provide the best solution, with low error and high speedup, and a nearly linear relationship error and speedup.

Please cite this article in press as: P. Koumoutsakos, J. Feigelman, Multiscale stochastic simulations of chemical reactions with regulated scale separation, J. Comput. Phys. (2013), http://dx.doi.org/10.1016/j.jcp.2012.11.030

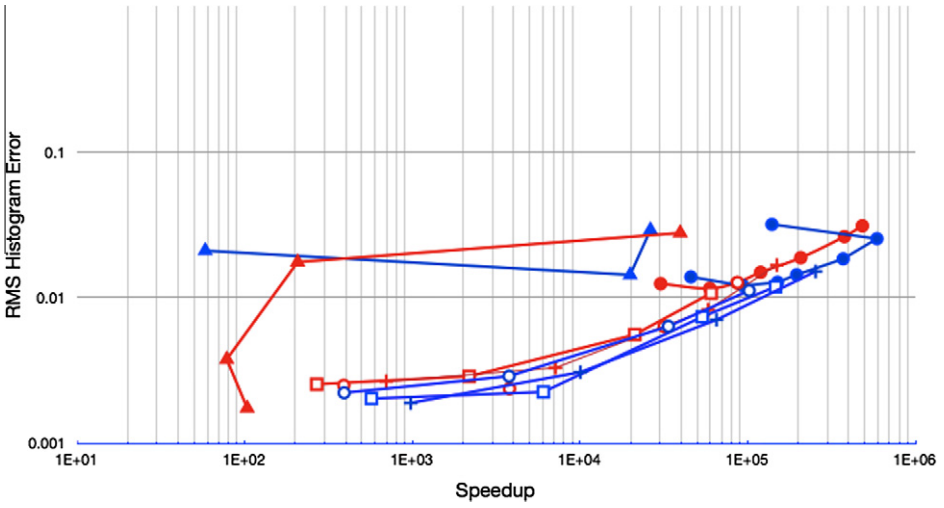

Fig. 5. Performance in the Schlögl system for FLAVOR-S (▲), HMM-S (●), and EF-S with $M = 5$ (+), $M = 10$ (○), and $M = 15$ (□). SSA is shown in blue, and τ in red. For FLAVOR-S and EF-S, $\xi = 10^{-j}, j \in \{4, 5, 6, 7\}$. For HMM-S, $M \in \{1, 2, 5, 10, 15, 25, 50\}$. Error is computed as the L^2 histogram error of species X relative to SSA. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

While FLAVOR-τ is more accurate than FLAVOR-SSA, it is not faster. For this system τ and R-leaping in general are not faster, as the number of molecules is small, e.g., 5 for X at equilibrium, and so leaping is generally infeasible. A possible explanation for the paradoxically *increased* accuracy of τ-leaping relative to FLAVOR-SSA is found in the discussion section.

4. Discussion

The work represents an extension of the FLAVORized SSA algorithm presented in [17]. The algorithms were applied to problems of different sizes and scale separations and depend critically on scale separation of the reaction propensities, as well as on the partial equilibrium of the fast reactions. When these two criteria are satisfied, multiscale algorithms can achieve several orders of magnitude speedup, with speedup and accuracy controllable functions of ξ and M.

Multiscale methods rely on the ability to partition the systems into fast and slow subsystems such that the fast subsystem converges in probability to its Quasi-Equilibrium distribution (QE PDF) on a time scale short in comparison to the waiting time of the next slow reaction. Different approaches have been taken for partitioning the system, including simply checking that the propensity exceeds some minimal value λ to be considered fast, as in Salis et al. [19], or using computational singular perturbation to estimate the relaxation times of the modes, and looking for a large gap in magnitudes, as in Samant et al. [11] The approach taken here is similar and seeks to maximize the separation between fast and slow group propensities. However this does not guarantee the convergence of the fast reactions over the microscopic stage. The examples in the literature focus primarily on systems with groups of (fast) reversible reaction pairs, assigning these to the fast reaction group. It is clear that these systems will always have a partial equilibrium. However, in more complicated systems it is not immediately obvious when the fast subsystem will converge. For example, in the scale separation algorithm failed in the case of the LacZ/LacY [26] system; a protein production reaction as identified as fast, but due to the lack of a partial-equilibrium on the microscopic time scale, the rescaling during the macroscopic stage produced inaccurate results. For such systems, it is necessary to manually identify the fast reactions that will converge over the microscopic time interval, which constitutes a limitation of the presented algorithms.

We note that the methods presented herein make no rigorous checks on the convergence of the fast subsystem to its QE PDF. Rather, they simulate the fast system a pre-specified number M times, with no guarantee of convergence. We note that M provides an additional degree of freedom for the speedup of the method and its value is at the moment largely determined by empirical estimates. However this does not deteriorate the computational efficiency of the method, as it is largely determined by the size of the macro scale time steps. The use of M sub steps is well suited to systems that are "near" to equilibrium, i.e. where the perturbation to the system during the macroscopic simulation step is due to few molecules. However, this clearly becomes inadequate as the perturbation is increased and a longer microscopic integration step is required before convergence is observed.

We remark that FLAVOR-SSA is in certain cases (see the Schlögl model) less accurate than FLAVOR-$\tau^{®}$. This is due to the fact that FLAVOR-SSA is not guaranteed to satisfy the time constraint (1.21) of [16]: $\tau \gg \epsilon \cdot \delta$, where τ is the microscopic time step of the fast subsystem, δ is the mesoscopic time step of the slow subsystem, and $1/\epsilon$ is the scale separation. If this inequality is not satisfied then the trajectories of the fast variables are not mixing with respect to their ergodic measures,

Please cite this article in press as: P. Koumoutsakos, J. Feigelman, Multiscale stochastic simulations of chemical reactions with regulated scale separation, J. Comput. Phys. (2013), http://dx.doi.org/10.1016/j.jcp.2012.11.030

meaning that the estimates of the slow-scale propensities at the beginning of the mesoscopic integration step will be inaccurate.

5. Conclusions and future work

We have presented a multiscale method for stochastic simulation of chemical reactions which can provide speedups of several orders of magnitude in systems with sufficiently large scale separation and molecularities. The proposed methods can be seen as a combination of multi scale frameworks, such as the HMM, FLAVOR or Equation Free, with accelerated stochastic simulation algorithms. We emphasise the observation that the error associated with the proposed methodologies is characterized by a cutoff phenomenon. This cutoff allows for the determination of an optimal trade-off between error and speedup. The extension of FLAVOR to Extended-FLAVOR provides additional control over the extent of convergence of the fast modes, affecting speedup and error. Ongoing work investigates the extensions of these methods with other variants of the HMM algorithm [25] as well as stochastic reaction–diffusion systems with variable spatial adaptivity [27]. We consider that an important line of work is to consider systems with multiple time-scales and investigate whether adaptively defined, multiple ξ, χ, M variables can be introduced to scale propensities on multiple time-scales.

References

[1] H. Kitano, Computational systems biology, Nature 420 (6912) (2002) 206–210.
[2] A. Arkin, J. Ross, H.H. McAdams, Stochastic kinetic analysis of developmental pathway bifurcation in phage lambda-infected *Escherichia coli* cells, Genetics 149 (4) (1988) 1633–1648.
[3] M. Barrio, K. Burrage, A. Leier, T.H. Tian, Oscillatory regulation of hes1: Discrete stochastic delay modelling and simulation, PLOS Comput. Biol. 2 (9) (2006) 1017–1030.
[4] D.T. Gillespie, A rigorous derivation of the chemical master equation, Physica A: Stat. Theoret. Phys. 188 (1–3) (1992) 404–425.
[5] D.T. Gillespie, Approximate accelerated stochastic simulation of chemically reacting systems, J. Chem. Phys. 115 (4) (2001) 1716–1733.
[6] A. Auger, P. Chatelain, P. Koumoutsakos, R-leaping: Accelerating the stochastic simulation algorithm by reaction leaps, J. Chem. Phys. 125 (8) (2006) 084103.
[7] E. Mjolsness, D. Orendorff, P. Chatelain, P. Koumoutsakos, An exact accelerated stochastic simulation algorithm, J. Chem. Phys. 130 (14) (2009) 144110.
[8] Y. Cao, D.T. Gillespie, L.R. Petzold, The slow–scale stochastic simulation algorithm, J. Chem. Phys. 122 (1) (2005) 014116.
[9] W.E.D. Liu, E. Vanden-Eijnden, Nested stochastic simulation algorithm for chemical kinetic systems with disparate rates, J. Chem. Phys. 123 (19) (2005).
[10] E. Weinan, Weiqing Ren, Eric Vanden-Eijnden, A general strategy for designing seamless multiscale methods, J. Chem. Phys. 228 (15) (2009) 5437–5453.
[11] A. Samant, B.A. Ogunnaike, D.G. Vlachos, A hybrid multiscale monte carlo algorithm (hymsmc) to cope with disparity in time scales and species populations in intracellular networks, BMC Bioinform. 8 (1) (2007) 175.
[12] C.V. Rao, A.P. Arkin, Stochastic chemical kinetics and the quasi-steady-state assumption: Application to the gillespie algorithm, J. Chem. Phys. 118 (11) (2003) 4999.
[13] W.E.D. Liu, B. Engquist, W. Li, X. Ren, E. Vanden-Eijnden, Heterogeneous multiscale methods: a review, Commun. Comput. Phys. 2 (3) (2007) 367–450, 06.
[14] I.G. Kevrekidis, C.W. Gear, G. Hummer, Equation-free: The computer-aided analysis of complex multiscale systems, AIChE J. 50 (7) (2004) 1346–1355.
[15] A.V. Skorokhod. Asymptotic methods in the theory of stochastic differential equations, volume 78 of Translations of Mathematical Monographs, American Mathematical Society, Providence, RI, 1989 [Translated from the Russian by H.H. McFaden].
[16] M. Tao, H. Owhadi, J.E. Marsden, Non-intrusive and structure preserving multiscale integration of stiff odes, sdes and hamiltonian systems with hidden slow dynamics via flow averaging, SIAM Multiscale Model. Simul. 8 (4) (2010) 1269–1324.
[17] B. Bayati, H. Owhadi, P. Koumoutsakos, A cutoff phenomenon in accelerated stochastic simulations of chemical kinetics via flow averaging (FLAVOR-SSA), J. Chem. Phys. 133 (24) (2010).
[18] A. Samant, D.G. Vlachos, Overcoming stiffness in stochastic simulation stemming from partial equilibrium: A multiscale monte carlo algorithm, J. Chem. Phys. 123 (14) (2005) 144114.
[19] H. Salis, Y.N. Kaznessis, An equation-free probabilistic steady-state approximation: Dynamic application to the stochastic simulation of biochemical reaction networks, J. Chem. Phys. 123 (21) (2005) 214106.
[20] A.B. Bortz, M.H. Kalos, J.L. Lebowitz, A new algorithm for monte carlo simulation of ising spin systems, J. Comput. Phys. (1975).
[21] D.T. Gillespie, Exact stochastic simulation of coupled chemical reactions, J. Phys. Chem. 81 (25) (1977) 2340–2361.
[22] Y. Cao, D. Gillespie, L. Petzold, Multiscale stochastic simulation algorithm w-, J. Comput. Phys. 206 (2) (2005) 395–411.
[23] Y. Cao, D.T. Gillespie, L.R. Petzold, Efficient step size selection for the tau-leaping simulation method, J. Chem. Phys. 124 (4) (2006) 044109.
[24] I. Fatkullin, E. Vanden-Eijnden, A computational strategy for multiscale systems with applications to Lorenz 96 model, J. Comput. Phys. 200 (2) (2004) 605–638.
[25] Assyr Abdulle, E. Weinan, Björn Engquist, Eric Vanden-Eijnden, The heterogeneous multiscale method, Acta Numer. 21 (2012) 1–87.
[26] K. Takahashi, A multi-algorithm, multi-timescale method for cell simulation, Bioinformatics 20 (4) (2004) 538–546.
[27] Basil Bayati, Philippe Chatelain, Petros Koumoutsakos, Adaptive mesh refinement for stochastic reaction–diffusion processes, J. Comput. Phys. 230 (1) (2011) 13–26.

Please cite this article in press as: P. Koumoutsakos, J. Feigelman, Multiscale stochastic simulations of chemical reactions with regulated scale separation, J. Comput. Phys. (2013), http://dx.doi.org/10.1016/j.jcp.2012.11.030