Available online at www.sciencedirect.com
ScienceDirect
Chemical Physics Letters 451 (2008) 136–140
CHEMICAL PHYSICS LETTERS
www.elsevier.com/locate/cplett

Accelerated stochastic and hybrid methods for spatial simulations of reaction–diffusion systems

Diego Rossinelli, Basil Bayati, Petros Koumoutsakos *

Computational Science and Engineering Laboratory, ETH Zurich, CH-8092, Switzerland

Received 19 July 2007; in final form 13 November 2007
Available online 24 November 2007

Abstract

Spatial distributions characterize the evolution of reaction–diffusion models of several physical, chemical, and biological systems. We present two novel algorithms for the efficient simulation of these models: Spatial τ-Leaping ($S\tau$-Leaping), employing a unified acceleration of the stochastic simulation of reaction and diffusion, and Hybrid τ-Leaping ($H\tau$-Leaping), combining a deterministic diffusion approximation with a τ-Leaping acceleration of the stochastic reactions. The algorithms are validated by solving Fisher's equation and used to explore the role of the number of particles in pattern formation. The results indicate that the present algorithms have a nearly constant time complexity with respect to the number of events (reaction and diffusion), unlike the exact stochastic simulation algorithm which scales linearly.

1. Reaction–diffusion: stochastic and deterministic models

Reaction–diffusion models are used to describe processes ranging from pattern formation in nature [1] and epidemics [2] to cancer induced angiogenesis [3]. These models are usually formulated either in terms of deterministic rate equations or by using stochastic descriptions of the underlying molecular processes. The stochastic description provides detailed information about the dynamics of the reaction–diffusion process, albeit at a significant computational cost over deterministic simulations.

The stochastic simulation algorithm (SSA) [4,5] has been used extensively in biochemical modeling ([6,7] and references therein) of reactions that assume a homogeneous spatial distribution of the species involved. A number of algorithms [8–10] have been presented for the acceleration of the SSA for homogeneous systems. In recent years the SSA has been extended to simulations involving spatially inhomogeneous molecular distributions undergoing diffusion and reaction processes [11–13]. The algorithm presented in [11,12] scales almost linearly with the number of events, but requires them to be scheduled thus prohibiting parallel execution. In [13] the computational time is reduced by splitting the reaction–diffusion phenomena into two distinct diffusion and reaction phases. This splitting may introduce numerical artifacts for systems close to a microscopic level as the reaction and diffusion processes happen concurrently, in particular for systems that involve too few particles to be insensitive to this kind of splitting. Recent works have examined the qualitative behavior of stochastic systems and have provided extensions for the deterministic systems to include leading order corrections for molecular noise [14,15], hence losing some of the descriptive benefits of a completely stochastic simulation but with the advantage of a relative reduction in computational cost.

A number of issues remain open in spatial SSA, such as the modeling of the diffusion rates in complex geometries, algorithms of increased computational efficiency and accuracy, and the enforcement of the homogeneity assumption [7].

In this Letter, we present two algorithms for the accelerated simulation of spatial reaction–diffusion processes: an accelerated spatial stochastic algorithm ($S\tau$-Leaping) employing a unified τ-Leaping procedure for the stochastic

* Corresponding author. Fax: +41 1 632 17 03.
E-mail address: petros@inf.ethz.ch (P. Koumoutsakos).

0009-2614/$ - see front matter © 2007 Elsevier B.V. All rights reserved.
doi:10.1016/j.cplett.2007.11.055

simulation of both processes, and a hybrid method ($H\tau$-Leaping), combining a deterministic description for diffusion with a τ-Leaping acceleration of the stochastic reactions. Both of the algorithms are validated in simulations of Fisher's equation [2]. In addition, we explore the role of the number of particles in the pattern forming Gray–Scott equations [16].

1.1. Stochastic modeling of reaction–diffusion processes

Reaction–diffusion phenomena can be represented by stochastic models, where particles in a domain move via Brownian motion and are subject to molecular collisions. In the present spatial simulations, the domain is decomposed into independent cells such that a reactant molecule can only react with other reactants in its cell while diffusion events are modeled as unimolecular transitions to neighboring cells.

We consider a total of N species and a domain that is discretized into a set of uniform cells, $\mathbf{C}$, subject to the same set of reactions, $\mathbf{R}$. We denote by $a_r(\mathbf{u}^c), r \in \mathbf{R}, c \in \mathbf{C}$, the propensity of the reaction r in the cell c and $\boldsymbol{v}_r^c = (v_{1r}, \ldots, v_{Nr})$, the corresponding stoichiometric vector. The set of diffusion transitions is $\mathbf{D}$, and $\boldsymbol{v}_d^{(i,c)}$ is the stoichiometric vector of the diffusion transition $d \in \mathbf{D}$ for the species i in the cell c. We can write the reaction–diffusion process in a unified framework in terms of generic transitions:

$$\sum_{i=1}^{N} \alpha_i^j A_i^j \to \sum_{i=1}^{N} \beta_i^j B_i^j, \quad j = 1, \ldots, M, \tag{1}$$

where j is the index of the transition, M is the number transitions, A_i is the species undergoing a transition, B_i is the species in the resulting transition, and α_i and β_i are the stoichiometric values. As an example, the reaction transitions for the pattern forming Gray–Scott [16] model are expressed as

$$U_0^{x,y,z} + 2U_1^{x,y,z} \to 3U_1^{x,y,z}, \tag{2}$$

$$U_1^{x,y,z} \to U_2^{x,y,z}. \tag{3}$$

The diffusion process can be represented by transitions to neighboring cells:

$$U_i^{x,y,z} \xrightarrow{\frac{d_i}{dl^2}} U_i^{x-1,y,z} \quad U_i^{x,y,z} \xrightarrow{\frac{d_i}{dl^2}} U_i^{x+1,y,z} \tag{4}$$

$$U_i^{x,y,z} \xrightarrow{\frac{d_i}{dl^2}} U_i^{x,y-1,z} \quad U_i^{x,y,z} \xrightarrow{\frac{d_i}{dl^2}} U_i^{x,y+1,z} \tag{5}$$

$$U_i^{x,y,z} \xrightarrow{\frac{d_i}{dl^2}} U_i^{x,y,z-1} \quad U_i^{x,y,z} \xrightarrow{\frac{d_i}{dl^2}} U_i^{x,y,z+1}, \tag{6}$$

where $U_i^{x,y,z}$ denotes the molecular species i inside the cell indexed as (x, y, z), d_i is the diffusion coefficient for species i and dl represents the cell sizes in all of the dimensions.

1.2. Deterministic modeling of reaction–diffusion processes

The stochastic model, presented in Section 1.1, can be represented by a deterministic reaction–diffusion model under the assumption of an infinite number of particles in the system. In the deterministic model, we evolve the concentration of substances, $u_i = u_i(\mathbf{x}, t)$, $i \in \{1, \ldots, N\}$, according to partial differential equations of the form

$$\frac{\partial u_i}{\partial t} = d_i \Delta u_i + f^{(i)}(\mathbf{u}), \tag{7}$$

where $f^{(i)}$ denotes the rate of change in concentrations due to the reactions, and d_i is the diffusion coefficient of substance i. This equation can be solved numerically, using techniques such as finite difference or particle strength exchange methods [17].

2. Computational methods

2.1. Spatial τ-Leaping

2.1.1. Choosing the maximal time step for reaction transitions

Cao et al. in [18] have provided a computationally efficient method for calculating the time step for the τ-Leaping method without the need for evaluating derivatives. We follow [18] by creating a bound for the molecular population in each cell:

$$\tau^{\text{reaction}} = \min_{c \in \mathbf{C}} \{\tau_c^{\text{reaction}}\}, \tag{8}$$

and for each cell we have

$$\tau_c^{\text{reaction}} = \min_{i \in \mathbf{I}} \left\{ \frac{\max\{\epsilon u_i^c / g_i, 1\}}{|\hat{\mu}_{i,c}^{\text{reaction}}(\mathbf{u})|}, \frac{\max\{\epsilon u_i^c / g_i, 1\}}{(\hat{\sigma}_{i,c}^{\text{reaction}}(\mathbf{u}))^2} \right\}, \tag{9}$$

where we have let ϵ be a control parameter such that $0 < \epsilon \ll 1$, g_i is the *highest order of reaction*, $\mathbf{I}$ is the set of different species and $\hat{\mu}_{i,c}^{\text{reaction}}(\mathbf{u})$ and $(\hat{\sigma}_{i,c}^{\text{reaction}}(\mathbf{u}))^2$ are given as follows:

$$\hat{\mu}_{i,c}^{\text{reaction}}(\mathbf{u}) = \sum_{r \in \mathbf{R}} v_{ir}^c a_r(\mathbf{u}^c), \tag{10}$$

$$(\hat{\sigma}_{i,c}^{\text{reaction}}(\mathbf{u}))^2 = \sum_{r \in \mathbf{R}} (v_{ir}^c)^2 a_r(\mathbf{u}^c). \tag{11}$$

2.1.2. Choosing the maximal time step for diffusion transitions

We can use the simple structure of the diffusion transitions in order to accelerate the computation of $\tau^{\text{diffusion}}$

$$\tau^{\text{diffusion}} = \min_{c \in \mathbf{C}} \{\tau_c^{\text{diffusion}}\}, \tag{12}$$

$$\tau_c^{\text{diffusion}} = \min_{i \in \mathbf{I}} \left\{ \frac{\max\{\epsilon u_i^c, 1\}}{|\hat{\mu}_{i,c}^{\text{diffusion}}(\mathbf{u})|}, \frac{\max\{\epsilon u_i^c, 1\}}{(\hat{\sigma}_{i,c}^{\text{diffusion}}(\mathbf{u}))^2} \right\}. \tag{13}$$

The denominators can be computed as

$$\hat{\mu}_{i,c}^{\text{diffusion}}(\mathbf{u}) = \frac{1}{dl^2} \sum_{c' \in N(c)} u_i^{c'} - u_i^c, \tag{14}$$

$$(\hat{\sigma}_{i,c}^{\text{diffusion}}(\mathbf{u}))^2 = \frac{1}{dl^2} \sum_{c' \in N(c)} u_i^{c'} + u_i^c, \tag{15}$$

where $N(c)$ denotes the set of neighboring cells of c. Since Eq. (15) will always be greater than Eq. (14), the formula for $\tau_c^{\text{diffusion}}$ is simplified to:

$$\tau_c^{\text{diffusion}} = \min_{i\in\mathbf{I}} \left\{ \frac{\max\{\epsilon u_i^c, 1\}}{\left(\hat{\sigma}_{i,c}^{\text{diffusion}}(\mathbf{u})\right)^2} \right\}. \tag{16}$$

2.1.3. Applying the transitions

The time step, τ, is chosen as the minimum of the two time steps,

$$\tau = \min\{\tau^{\text{reaction}}, \tau^{\text{diffusion}}\}. \tag{17}$$

We perform the transitions on the entire solution, $\mathbf{u} = \{\mathbf{u}^c\}_{c\in\mathbf{C}}$, according to the following formula:

$$\mathbf{u}(t+\tau) = \mathbf{u}(t) + \sum_{c\in\mathbf{C}} \sum_{r\in\mathbf{R}} \boldsymbol{\nu}_r^c \mathscr{P}(a_r(\mathbf{u}^c), \tau) + \sum_{c\in\mathbf{C}} \sum_{i\in I} \sum_{d\in\mathbf{D}} \boldsymbol{\nu}_d^{(i,c)} \mathscr{P}\left(\frac{d_i u_i^c}{dl^2}, \tau\right), \tag{18}$$

where $\mathscr{P}(.)$ is a sample from a Poisson distribution.

2.2. Hybrid τ-Leaping

In order to further accelerate the spatial modeling of reaction–diffusion systems, we propose a hybrid scheme where the reactions are simulated stochastically while diffusion is simulated deterministically. This approximation is suitable since the diffusion process is typically two orders of magnitude faster than the reaction process [19]. We consider a system where the particles, $u_i = u_i(\mathbf{x}, t)$, evolve according to the following equation:

$$u_i(\mathbf{x}, t+\tau) = u_i(\mathbf{x}, t) + \mathscr{M}_1(d_i \Delta_{\text{d}} \mathscr{M}_2(u_i(\mathbf{x}, t))) + f_{\text{s}}^{(i)}(\mathbf{u}(\mathbf{x}, t)), \tag{19}$$

where $f_{\text{s}}^{(i)}$ represents the stochastically simulated reactions, Δ_{d} represents a deterministic diffusion operator, and $\mathscr{M}_1$ and $\mathscr{M}_2$ are mapping functions such that $\mathscr{M}_1 : \mathbb{R}_+^N \to \mathbb{N}^N$ and $\mathscr{M}_2 : \mathbb{N}^N \to \mathbb{R}_+^N$. The operators $\mathscr{M}_1$ and $\mathscr{M}_2$ convert function values between the discrete and continuum descriptions. The operator $\mathscr{M}_2$ is trivial as it converts from a discrete to a continuum model. However, care needs to be taken with $\mathscr{M}_1$ since we need to ensure both a concise mapping and also a conservation of mass within our system. The details of deriving $\mathscr{M}_1$ will be specified elsewhere.

The algorithm for the hybrid method involves choosing a value for τ and simulating the reactions in the volume. Then, using this τ, we simulate the diffusion process with second order finite differences. This procedure is performed iteratively using an Euler time integration method until the final time is reached.

3. Results

3.1. Validation: Fisher's equation

The proposed methods were validated by simulating Fisher's equation [2], a model for the spreading of an advantageous gene in a population inhabiting a one-dimensional space:

$$\frac{\partial p}{\partial t} = k\Delta p + mp(1-p). \tag{20}$$

Fisher's equation admits traveling wave solutions [20] with a speed of c, which is known analytically for certain initial conditions. We used a smoothed Heaviside step function [20], for the initial condition with $k = 1$ and $m = 1$, which has an analytical solution with a wavespeed of $c = \frac{5}{\sqrt{6}}$ per unit of time. In addition, we used zero-flux Neumann boundary conditions. This initial condition was distributed across the entire y and z-axes of our three-dimensional domain so that we should be able to obtain the same wavespeed in our simulation as in the one-dimensional case. We solved Fisher's equation with various resolutions in a domain of $[-14, 14] \times [-14, 14] \times [-14, 14]$ using the $S\tau$-Leaping and $H\tau$-Leaping methods with a varying number of particles per unit of concentration which we denote by P.

Fig. 1 presents the error of the propagating front with respect to the analytical solution of Fisher's equation. The plot on the left shows the error at a single time point ($t = 2$), and the plot on the right shows the error of the position of the front with respect to all time points. The results indicate that using $S\tau$-Leaping with a high number of particles is the more accurate method. We observe that this fact can be justified by the absence of any decoupling of reaction and diffusion processes. However, when using $H\tau$-Leaping, we can obtain a good approximation when the number of particles per concentration unit is above 1000.

In Fig. 2 we show the relative performance of both methods in simulating Fisher's equation by varying the number of cells and the number of particles. We observe that in both methods the number of cells is a more critical factor for CPU time as opposed to the number of particles. We note that the CPU time does not grow linearly with respect to the number of particles (and thus the number of events), which is the case with the exact spatial SSA algorithm.

3.2. Gray–Scott equations

In addition to Fisher's model, we also considered the Gray–Scott [16] model, which was described in Section 1.1 with Eqs. (2)–(6). The corresponding partial differential equations for this model are

$$\frac{\partial u}{\partial t} = d_u \Delta u - uv^2 + F(1-u), \tag{21}$$

$$\frac{\partial v}{\partial t} = d_v \Delta v + uv^2 - (F+\kappa)v, \tag{22}$$

where κ is a rate constant for the second chemical reaction, F is a dimensionless feed rate, d_u and d_v diffusion coefficients, $u = u(\mathbf{x}, t)$ represents the concentration of species U_0, and $v = v(\mathbf{x}, t)$ represents the concentration of species U_1.

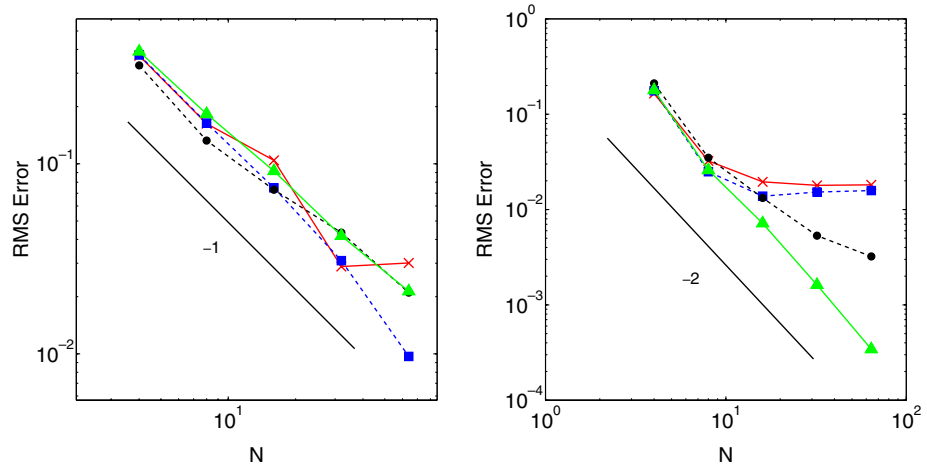

Fig. 1. Convergence for both methods, $S\tau$-Leaping and $H\tau$-Leaping, with different values for the number of particles per unit of concentration (P) and the number of cells per dimension (N). $S\tau$-Leaping is denoted by '- - ● - -' $P = 10^2$, '– ▲–' $P = 10^6$, and $H\tau$-Leaping '–×–' $P = 10^2$, '- - ■ - -' $P = 10^6$. Shown on the left, the RMS error at time $t = 2$. On the right, the RMS error in the position of the front with respect to time.

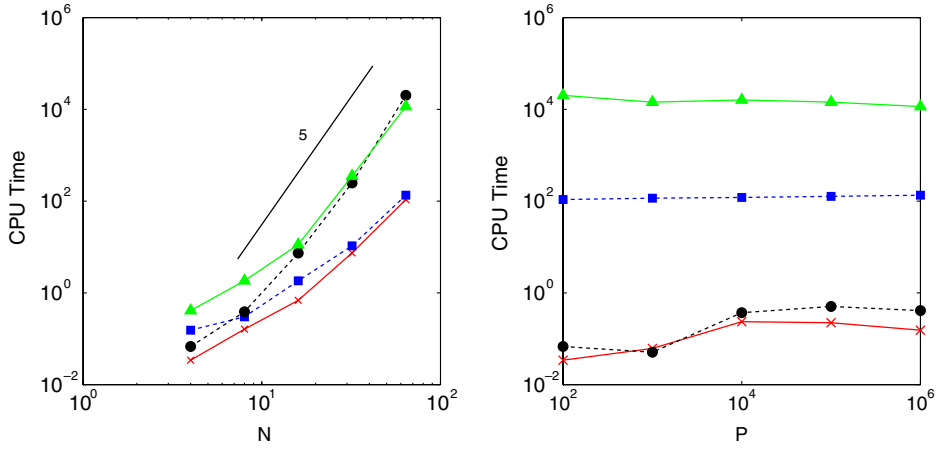

Fig. 2. CPU time in seconds for both methods, $S\tau$-Leaping and $H\tau$-Leaping, with different values for the number of particles per unit of concentration (P) and the number of cells per dimension (N). $S\tau$-Leaping is denoted by '- - ● - -', '– ▲ –', and $H\tau$-Leaping '–×–', '- - ■ - -'. Shown on the left, the CPU time with varying values of N, where $P = 10^2, 10^6, 10^2, 10^6$ for the four methods, respectively. On the right, the CPU time with respect to P, where $N = 4, 64, 4, 64$, respectively.

Eqs. (21) and (22) have been analyzed in their two-dimensional continuum form in [16] for a range of values for F and κ. We performed numerical simulations of the Gray–Scott equations in two and three-dimensions with periodic boundary conditions using deterministic, $H\tau$-Leaping (Section 2.2), and $S\tau$-Leaping approaches (Section 2.1) with varying levels of particles in order to determine whether we obtain qualitatively different patterns. Fig. 3 shows results for two-dimensional simulations of the Gray–Scott equations. We varied the number of particles in each cell while keeping $F = 0.04$, $\kappa = 0.06$, and integrated from $t = 0$ to $t = 1000$. There are notable differences in the solutions: as the number of particles increases, the stochastic simulations converge to the pattern observed by purely deterministic simulations of reactions and diffusion. Lastly, we solved the Gray–Scott equations in three-dimensions using a discretization of $128 \times 128 \times 128$ with

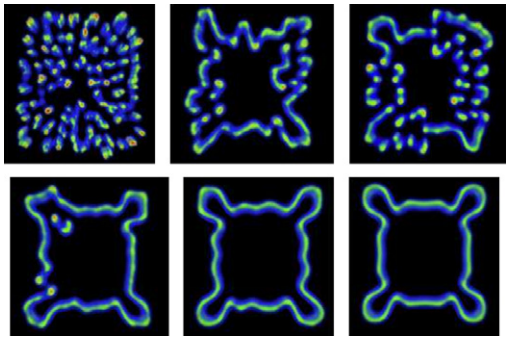

Fig. 3. Analysis of the role of the number of particles for the Gray–Scott equations solved with a 300×300 discretization with $F = 0.04$, $\kappa = 0.06$, $t = 1000$. From top left to bottom right the number of particles per unit of concentration is increased from 100, 1000, 1000, 5000, 10000, continuum, respectively. The methods used to solve the equations were the following (from top left to bottom right): $S\tau$-Leaping, $S\tau$-Leaping, $H\tau$-Leaping, $S\tau$-Leaping, $H\tau$-Leaping, deterministic.

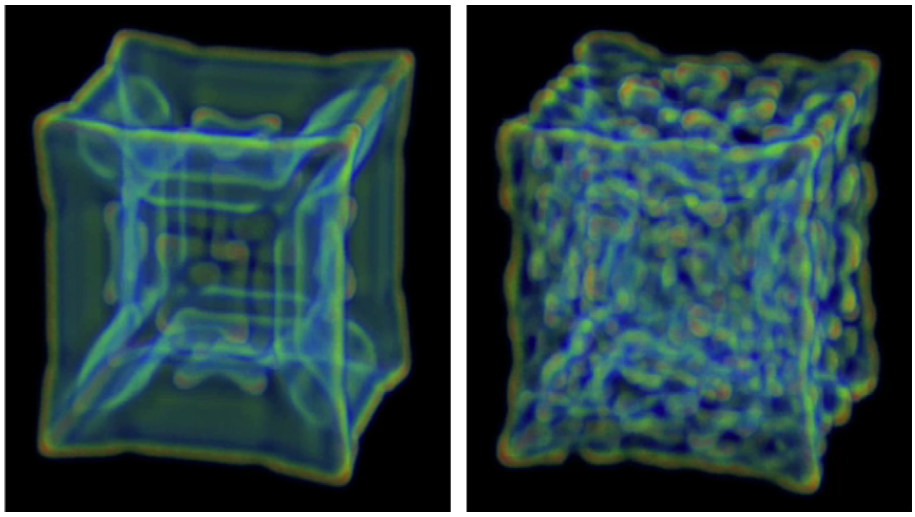

Fig. 4. Three-dimensional solutions of the Gray–Scott equations using (left) deterministic and $H\tau$-Leaping solvers (right) on a $128 \times 128 \times 128$ discretization with $F = 0.04,\ \kappa = 0.06,\ t = 1000$. The $H\tau$-Leaping solver was performed with 1000 particles per unit of concentration.

$F = 0.04, \kappa = 0.06$, and integrated from $t = 0$ to $t = 1000$ (Fig. 4). In three-dimensions, the noise from the low numbers of particles makes itself apparent and the solution notably differs from the deterministic solution.

4. Conclusion

We presented two novel numerical methods for the efficient simulation of reaction–diffusion processes as described by stochastic and hybrid models. In $S\tau$-Leaping, a unified τ-Leaping procedure was used for both the reaction and the diffusion processes whereas in $H\tau$-Leaping, diffusion was handled deterministically and the reactions stochastically. We validated the methods using the analytical solution of Fisher's equation and we investigated the role of the number of molecules in pattern forming Gray–Scott equations. The algorithms were shown to exhibit significant computational improvements over the exact spatial SSA.

Present work includes further accelerating the proposed algorithms using R-Leaping [10], the parallelization of the methods, and their application to biological models in order to verify the significance of spatially-dependent chemical reactions.

References

[1] E. Ben-Jacob, H. Levine, Nature 409 (6823) (2001) 985.
[2] R. Fisher, Ann. Eugenics 7 (1937) 355.
[3] V. Cristini, H.B. Frieboes, R. Gatenby, S. Caserta, M. Ferrari, J. Sinek, Clin. Cancer Res. 11 (19) (2005) 6772.
[4] D.T. Gillespie, J. Comput. Phys. 22 (4) (1976) 403.
[5] D. Gillespie, J. Phys. Chem. 81 (25) (1977) 2340.
[6] D.P. Tolle, N. Le Novere, Curr. Bioinform. 1 (3) (2006) 315.
[7] C. Lemerle, B. Di Ventura, L. Serrano, Febs. Lett. 579 (8) (2005) 1789.
[8] D. Gillespie, J. Chem. Phys. 115 (4) (2001) 1716.
[9] M. Gibson, J. Bruck, J. Phys. Chem. A 104 (9) (2000) 1876.
[10] A. Auger, P. Chatelain, P. Koumoutsakos, J. Chem. Phys. 125 (8) (2006) 084103.
[11] J. Elf, M. Ehrenberg, Syst. Biol. 1 (2) (2004) 230.
[12] J. Hattne, D. Fange, J. Elf, Bioinformatics 21 (12) (2005) 2923.
[13] J.V. Rodriguez, J.A. Kaandorp, M. Dobrzynski, J.G. Blom, Bioinformatics 22 (15) (2006) 1895.
[14] M. Scott, T. Hwa, B. Ingalls, Proc. Natl. Acad. Sci. USA 104 (18) (2007) 7402.
[15] A.B. Stundzia, C.J. Lumsden, J. Comput. Phys. 127 (1) (1996) 196.
[16] J. Pearson, Science 261 (5118) (1993) 189.
[17] P. Degond, S. Mas-Gallic, Math. Comput. 53 (188) (1989) 485.
[18] Y. Cao, D.T. Gillespie, L.R. Petzold, J. Chem. Phys. 124 (4) (2006) 044109.
[19] D. Bernstein, Phys. Rev. E 71 (4) (2005) 041103.
[20] J. Murray, Mathematical Biology: An Introduction, Springer-Verlag, 2002.