Available online at www.sciencedirect.com
ScienceDirect
Journal of Computational Physics 227 (2007) 100–122
JOURNAL OF COMPUTATIONAL PHYSICS
www.elsevier.com/locate/jcp

Hybrid method for the chemical master equation

Andreas Hellander, Per Lötstedt *

Division of Scientific Computing, Department of Information Technology, Uppsala University, SE-75105 Uppsala, Sweden

Received 24 August 2006; received in revised form 31 May 2007; accepted 20 July 2007
Available online 7 August 2007

Abstract

The chemical master equation is solved by a hybrid method coupling a macroscopic, deterministic description with a mesoscopic, stochastic model. The molecular species are divided into one subset where the expected values of the number of molecules are computed and one subset with species with a stochastic variation in the number of molecules. The macroscopic equations resemble the reaction rate equations and the probability distribution for the stochastic variables satisfy a master equation. The probability distribution is obtained by the Stochastic Simulation Algorithm due to Gillespie. The equations are coupled via a summation over the mesoscale variables. This summation is approximated by Quasi-Monte Carlo methods. The error in the approximations is analyzed. The hybrid method is applied to three chemical systems from molecular cell biology.

MSC: 65M20; 65M60

Keywords: Master equation; Reaction rate equations; Stochastic simulation algorithm; Monte Carlo method; Quasi-Monte Carlo method; Stochastic chemical kinetics

1. Introduction

In a well stirred chemical system, the chemical reactions are often modeled by the reaction rate equations. These equations form a system of nonlinear, coupled ordinary differential equations (ODEs). Such a macroscopic model provides a good description of the time evolution for the concentrations of the chemical species of the system in many cases e.g. when the number of molecules of each kind is large and in the absence of critical phenomena. On the other hand, the species with low copy number are not well described by a deterministic and macroscopic model since they are subject to random fluctuations which cannot be neglected and in many cases have a great impact on the behavior of the system. In a biological cell, the underlying assumptions for the reaction rate equations are often violated [1,9,31,34,37,40,42,43]. At least some species are usually present in low copy numbers. For example, mRNA usually exists in one or a few copies, while transcription factors may be present in the range from ten to hundreds of molecules. Yet other components could be active

* Corresponding author. Tel.: +46 18 4712972; fax: +46 18 523049.
E-mail addresses: andreas.hellander@it.uu.se (A. Hellander), perl@it.uu.se (P. Lötstedt).

0021-9991/$ - see front matter © 2007 Elsevier Inc. All rights reserved.
doi:10.1016/j.jcp.2007.07.020

in large numbers and approach macroscopic values. Thus, a realistic model must take the inherent randomness into account and therefore need to be of stochastic nature. A disadvantage with stochastic models is the increase in computational complexity compared to the reaction rate equations. Another source of computational difficulties is the different scales both in time and in reaction rates [4,8,11,17,20,25,35,36].

One way to model coupled chemical reactions stochastically at a mesoscopic level is to the use the Stochastic Simulation Algorithm (SSA) proposed by Gillespie [16]. This Monte Carlo algorithm yields a correct realization of the process, but the computing time required to approximate the probability distribution of the species in the system is often dictated by the reactions involving the molecules with the largest copy numbers or the fastest reaction rates. They may well be the components where the stochastic description is the least important. The convergence rate is also slow for this method and it can be computationally cumbersome to obtain detailed information of the probability distributions when the number of different reacting molecules is large.

The underlying stochastic process is often assumed to be memory lacking or *Markovian*. Then the time evolution of the probability distribution is described by a difference–differential equation, the chemical master equation [15,24]. One molecular species corresponds to one spatial dimension in the equation. Analytical solutions of this equation are known only for very simple chemical systems and numerical solution is necessary for realistic systems. Direct numerical solution of the master equation suffers from the curse of dimensionality as the computational work and storage requirements grow exponentially with the number of dimensions or reacting species. Consequently, this often limits the size of the models to four or maybe five dimensions.

Different ways to mitigate or avoid the exponential growth have been proposed, either by approximations of the master equation [41] or dimension reduction by introducing assumptions about the behavior of different components [43]. In the first case, the master equation is approximated by the Fokker–Planck equation, a partial differential equation derived from a truncated Taylor expansion of the master equation [15,24]. The discretized Fokker–Planck equation can be solved with fewer variables compared to the master equation, but this approach is still limited by an unfavorable rise in computational time with increasing number of species. For certain low-dimensional problems though, solution of the Fokker–Planck is orders of magnitude faster than the SSA [41]. The second approach relies on some prior knowledge of the system in order to reduce the dimension of the problem. While this can result in a considerable reduction of the complexity, a profound knowledge of the biological system is required to introduce some simplifying assumptions. For moderate numbers of reacting species, model reduction with Krylov spaces [2,30] and sparse grid methods [21] are alternatives.

With SSA [16] the work grows linearly with the number of species but for systems with different time scales the method is slow owing to the explicit time stepping. By allowing more reactions to take place in a time step or assuming approximate relations between the species, longer time steps are possible. Either the species are split into a slow and a fast set [11,36] or the reactions are partitioned in this way [7,8,20,38]. In [4,5], both reactions and species are partitioned. Then a quasi-steady-state assumption [36] or a partial equilibrium assumption [4,39] is made allowing a simplified treatment of the fast variables.

There are efficient numerical methods to solve ODEs and the fast components are modeled by deterministic equations and other components are treated with a stochastic model in [20,25]. The assumption in common in these papers is that there is a separation in the *time* scales.

A partitioning of the chemical compounds into a subset of variables that can be treated as normally distributed with a small variance and a subset of variables that need a stochastic treatment is suggested in [28]. The scales of the *variances* of the stochastic variables are assumed here to be separated into two subsets. Equations are derived for the expected values of the first subset assuming that the quotient between the standard deviation and the expected value is small. The equations can be solved given the probability density function (PDF) of the stochastic variables. This PDF satisfies a Fokker–Planck equation and is solved by a finite volume scheme in [13] but the dimension of the stochastic problem is then restricted to, say, five. Bounds are derived on the approximations due to the small variances.

Here, we apply SSA to the stochastic part and compute the coupling to the deterministic part using Monte Carlo and Quasi-Monte Carlo summation [3,18]. The equations for the expected values are integrated in time by an unconditionally stable, implicit method. Hence, if the species in the fast reactions coincide with the species with small variance and are treated macroscopically, then the small time steps in SSA are avoided as in [20,25]. On the other hand, if some of the fast variables have large variances and are in the stochastic regime,

then their trajectories can be simulated with the more efficient versions of SSA e.g. in [4,8,39]. Our method can be regarded either as a means to introduce stochasticity in some components in the reaction rate equations or as a way of improving the efficiency of SSA by reducing the number of species in its system state vector. In this way, the quality of large macroscopic models can be improved by allowing a subset of the variables to be stochastic or large stochastic models can be reduced to permit shorter simulation times.

The partitioning of the variables can be based on one of the following possibilities:

1. Species with large copy numbers L have a standard deviation of order $L^{1/2}$ and their quotient $L^{-1/2}$ is small for large L [24, p. 248]. These species are candidates for a macroscopic treatment. An example where L is large for some of the species is found e.g. in [7].
2. Estimation of the expected values and the covariances by solving the equations for them [10,24] or from SSA with a few trajectories will tell which variables are amenable to macroscopic approximation. The equations for the first and second moments are a system of nonlinear ordinary differential equations and the work grows as a polynomial of low order in the number of species. The cost of simulation of a small number of trajectories with SSA is negligible compared to simulation of the system with good accuracy.
3. Biological insight will provide hints how to partition the species. An example of a reduction of a system using the quasi-steady-state assumption is found in [43].

Our method differs from previously published hybrid methods by assuming a separation in the variance of the species, evaluation of the propensities for the macroscopic scale without ruining the complexity of the algorithm, providing error estimates depending on parameters in the method, adaptive control of the errors due to the time discretization, and by computing the PDFs with good numerical accuracy.

In the next section, the nonlinear system of differential equations for the expected values and the master equation for the PDF of the stochastic variables are derived. The solution algorithm for the coupled system is described in Section 3. The errors in the solution are discussed and the computational work is estimated in Section 4. Three systems in molecular biology are simulated in Section 5 using our algorithm. In one example, the behavior of the macroscopic model is sensitive to a parameter but addition of stochastic variables removes this sensitivity. The splitting is based on criterion 3 above. In another example, the full SSA simulation is compared to a mixed macroscopic–mesoscopic model where the mesoscpoic variables have estimated large variances compared to the expected values as in criterion 2. The difference in the probability distribution of critical components is small while the savings in computational work are significant. Finally, conclusions are drawn in the last section.

The notation in the paper is as follows. The ith element of a vector $\mathbf{v}$ is denoted by v_i. If $v_i \geqslant 0$ for all i, then we write $\mathbf{v} \geqslant 0$. The ℓ_p-norm of $\mathbf{v}$ of length N is $\|\mathbf{v}\|_p = \left(\sum_{i=1}^{N} |v_i|^p\right)^{1/p}$. The set of integer numbers is written $\mathbb{Z}$ and $\mathbb{Z}_+$ denotes the non-negative integer numbers. In the same manner, $\mathbb{R}$ denotes the real numbers and $\mathbb{R}_+$ is the non-negative real numbers.

2. The system of equations

Assume that we have a chemical system with N active molecular species X_i, $i = 1, \ldots, N$, and that x_i denotes the number of molecules of substrate X_i. The system has a state vector $\mathbf{x} \in \mathbb{Z}_+^N$, and a reaction r in the system is a transition from a state $\mathbf{x}_r$ to $\mathbf{x}$ so that $\mathbf{x}_r = \mathbf{x} + \mathbf{n}_r$ with $\mathbf{n}_r \in \mathbb{Z}^N$. Only a few components of $\mathbf{n}_r$ are non-zero. The probability of the reaction to occur per unit time is the non-negative propensity $w_r(\mathbf{x}_r, t)$. The change in the state vector by a reaction r can now be written

$$\mathbf{x}_r \xrightarrow{w_r(\mathbf{x}_r,t)} \mathbf{x}, \mathbf{n}_r = \mathbf{x}_r - \mathbf{x}. \tag{1}$$

The PDF $p(\mathbf{x}, t)$ for the system to be in the state $\mathbf{x}$ at time t satisfies the chemical master equation [15,24]. With a splitting of $\mathbf{n}_r$ into two parts so that

$$\mathbf{n}_r = \mathbf{n}_r^+ + \mathbf{n}_r^-, \quad n_{ri}^+ = \max(n_{ri}, 0), \quad n_{ri}^- = \min(n_{ri}, 0),$$

the master equation for R reactions is

$$\frac{\partial p(\mathbf{x},t)}{\partial t} = \sum_{\substack{r=1 \\ \mathbf{x}+\mathbf{n}_r^- \geqslant 0}}^{R} w_r(\mathbf{x}+\mathbf{n}_r,t)p(\mathbf{x}+\mathbf{n}_r,t) - \sum_{\substack{r=1 \\ \mathbf{x}-\mathbf{n}_r^+ \geqslant 0}}^{R} w_r(\mathbf{x},t)p(\mathbf{x},t). \tag{2}$$

It follows from [14] that the total probability $\sum_{\mathbf{x}\in\mathbb{Z}_+^N} p(\mathbf{x},t)$ is constant in time.

In order to reduce the computational complexity of solving (2), $\mathbf{x}$ is split in [28] into two parts $\mathbf{x}^{\mathrm{T}} \to (\mathbf{x}^{\mathrm{T}}, \mathbf{y}^{\mathrm{T}})$ with $\mathbf{x} \in \mathbb{Z}_+^m$, $\mathbf{y} \in \mathbb{R}^n$, and $N = m + n$. In the same manner, the transition vector $\mathbf{n}_r$ is split $\mathbf{n}_r^{\mathrm{T}} \to (\mathbf{m}_r^{\mathrm{T}}, \mathbf{n}_r^{\mathrm{T}})$ for reaction r. The dimensions of $\mathbf{m}_r$ and $\mathbf{n}_r$ are $\mathbf{m}_r \in \mathbb{Z}^m, \mathbf{n}_r \in \mathbb{Z}^n$. The corresponding stochastic variables are $X_i, i = 1,\ldots,m$, and $Y_i, i = 1,\ldots,n$. The assumption is that the stochastic variables Y_i are mutually independent, independent of X_i, and normally distributed with a small variance. Then the PDF of the full system is written

$$p(\mathbf{x},\mathbf{y},t) = \gamma_n p_0(\mathbf{x},t) \exp\left(-\sum_{j=1}^{n} \frac{(y_j - \phi_j(t))^2}{2\sigma_j^2}\right), \tag{3}$$

with the normalizing constant $\gamma_n = (2\pi)^{-n/2}\prod_{j=1}^n \sigma_j^{-1}$. Equations will be derived for p_0 and $\boldsymbol{\phi} \in \mathbb{R}^n$ such that p approximately fulfills (2).

The marginal PDF of p in (3) satisfies

$$p_0(\mathbf{x},t) = \int p(\mathbf{x},\mathbf{y},t)\,\mathrm{d}\mathbf{y},$$

where the domain of the integral is $\mathbb{R}^n$. The scaling of p_0 is such that the total probability satisfies

$$\sum_{\mathbf{x}\in\mathbb{Z}_+^m} \int p(\mathbf{x},\mathbf{y},t) = \sum_{\mathbf{x}\in\mathbb{Z}_+^m} p_0(\mathbf{x},t) = 1. \tag{4}$$

The expected value of Y_k is

$$E[Y_k] = \sum_{\mathbf{x}\in\mathbb{Z}_+^m} \int y_k p(\mathbf{x},\mathbf{y},t)\,\mathrm{d}\mathbf{y} = \phi_k(t).$$

A differential–difference equation for p_0 is derived in [28] assuming that $\sigma_j \leqslant \sigma$ and that σ is small. Ignoring terms proportional to σ^2 for a vanishing σ, the equation is for R reactions

$$\frac{\partial p_0(\mathbf{x},t)}{\partial t} = \sum_{\substack{r=1 \\ \mathbf{x}+\mathbf{m}_r^- \geqslant 0}}^{R} w_r(\mathbf{x}+\mathbf{m}_r,\boldsymbol{\phi}(t),t)p_0(\mathbf{x}+\mathbf{m}_r,t) - \sum_{\substack{r=1 \\ \mathbf{x}-\mathbf{m}_r^+ \geqslant 0}}^{R} w_r(\mathbf{x},\boldsymbol{\phi}(t),t)p_0(\mathbf{x},t). \tag{5}$$

This is a master equation for p_0 with propensities depending on the expected values $\boldsymbol{\phi}$. It follows from (5) that the total marginal probability $\sum_{\mathbf{x}\in\mathbb{Z}_+^m} p_0(\mathbf{x},t)$ is constant (cf.(2)) as presupposed in (4).

The differential equation for $\phi_j, j = 1,\ldots,n$, is for a single reaction, $R = 1$,

$$\begin{aligned}
\frac{\mathrm{d}\phi_j}{\mathrm{d}t} &= \sum_{\mathbf{x}\in\mathbb{Z}_+^m} \int y_j \frac{\partial p}{\partial t}\mathrm{d}\mathbf{y} \\
&= \sum_{\mathbf{x}+\mathbf{m}_r^- \geqslant 0} w_r(\mathbf{x}+\mathbf{m}_r,\boldsymbol{\phi}(t),t)p_0(\mathbf{x}+\mathbf{m}_r,t)(\phi_j - n_{rj}) - \sum_{\mathbf{x}-\mathbf{m}_r^+ \geqslant 0} w_r(\mathbf{x},\boldsymbol{\phi}(t),t)p_0(\mathbf{x},t)\phi_j \\
&= \phi_j \sum_{\mathbf{x}\in\mathbb{Z}_+^m} \frac{\partial p_0(\mathbf{x},t)}{\partial t} - n_{rj} \sum_{\mathbf{x}+\mathbf{m}_r^- \geqslant 0} w_r(\mathbf{x}+\mathbf{m}_r,\boldsymbol{\phi}(t),t)p_0(\mathbf{x}+\mathbf{m}_r,t) \\
&= -n_{rj} \sum_{\mathbf{x}+\mathbf{m}_r^- \geqslant 0} w_r(\mathbf{x}+\mathbf{m}_r,\boldsymbol{\phi}(t),t)p_0(\mathbf{x}+\mathbf{m}_r,t),
\end{aligned} \tag{6}$$

by (2), (5), the conservation of the total marginal probability, and after ignoring small terms of $\mathcal{O}(\sigma^2)$, see [28]. The equation is simplified by removing the shift $\mathbf{m}_r$ so that for R reactions

$$\frac{\mathrm{d}\phi_j}{\mathrm{d}t} = -\sum_{r=1}^{R} n_{rj} \sum_{\mathbf{x}\in\mathbb{Z}_+^m} w_r(\mathbf{x}, \boldsymbol{\phi}(t), t) p_0(\mathbf{x}, t), \quad j = 1, \ldots, n. \tag{7}$$

This system of differential–summation equations is equal to the reaction rate equations for the chemical system when the distribution of all species is assumed to be normal with small variances.

Suppose that the mesoscopic $\mathbf{X}$ participates in only the first ρ reactions. Then $\mathbf{m}_r = 0$, $r = \rho + 1, \ldots, R$, and the right hand side in (5) is reduced to

$$\sum_{\substack{r=1 \\ \mathbf{x}+\mathbf{m}_r^- \geqslant 0}}^{\rho} w_r(\mathbf{x} + \mathbf{m}_r, \boldsymbol{\phi}(t), t) p_0(\mathbf{x} + \mathbf{m}_r, t) - \sum_{\substack{r=1 \\ \mathbf{x}-\mathbf{m}_r^+ \geqslant 0}}^{\rho} w_r(\mathbf{x}, \boldsymbol{\phi}(t), t) p_0(\mathbf{x}, t). \tag{8}$$

If $\rho = 0$, then $\partial p_0/\partial t = 0$ and $p_0(\mathbf{x}, t) = p_0(\mathbf{x}, 0)$. The summation over $\mathbf{x}$ in (7) to obtain new propensities should then be performed in a preparatory stage before the simulation starts.

Suppose that w_r depends on $\mathbf{x}$ only in the κ first reactions. Then the right hand side in (7) is

$$\begin{aligned} &-\sum_{r=1}^{\kappa} n_{rj} \sum_{\mathbf{x}\in\mathbb{Z}_+^m} w_r(\mathbf{x}, \boldsymbol{\phi}(t), t) p_0(\mathbf{x}, t) - \sum_{r=\kappa+1}^{R} n_{rj} \sum_{\mathbf{x}\in\mathbb{Z}_+^m} w_r(\boldsymbol{\phi}(t), t) p_0(\mathbf{x}, t) \\ &\quad = -\sum_{r=1}^{\kappa} n_{rj} \sum_{\mathbf{x}\in\mathbb{Z}_+^m} w_r(\mathbf{x}, \boldsymbol{\phi}(t), t) p_0(\mathbf{x}, t) - \sum_{r=\kappa+1}^{R} n_{rj} w_r(\boldsymbol{\phi}(t), t). \end{aligned} \tag{9}$$

If $\kappa = 0$, then (7) are the reaction rate equations.

Two equations have been derived for p_0 in (5) and for ϕ_j in (7). They are coupled via the expected values of the propensity of each reaction in (7) and the expected values of the species in the propensities in (5). The dimension of the stochastic problem has been reduced from N in (2) to m in (5). The cost for this reduction is the increased number of dependent variables from 1 in (2) to $n + 1$ in (5) and (7) but the work grows at most as a polynomial of low order in n solving (5) and (7) compared to an exponential growth in $n = N$ when is (2) solved.

3. Solution algorithm

The two systems of Eqs. (5) and (7) are solved numerically by a hybrid method where p_0 in (5) is determined by SSA and $\boldsymbol{\phi}$ in (7) is computed by a deterministic time stepping method.

3.1. Solution of the master equation

The SSA is applied to the chemical system defined by the master equation for p_0 in (5) with the right hand side (8). Trajectories of the system are simulated by updating the state vectors after each reaction in [16]. Then the probability for the system to be in a state $\mathbf{x}$ at time t^v is approximated by

$$p_0(\mathbf{x}, t^v) \approx \frac{1}{M} \sum_{j=1}^{M} \Psi_j, \qquad \Psi_j = \begin{cases} 1, & \mathbf{x}_j = \mathbf{x} \text{ at } t^v, \\ 0, & \text{otherwise}, \end{cases} \tag{10}$$

where M is the number of trajectories. The error in p_0 due to the finite M is analyzed in Section 4.1.

The time evolution of the state is simulated by SSA in the following way. Let w^v be the sum of the propensities at t^v so that

$$w^v = \sum_{r=1}^{\rho} w_r(\mathbf{x}^v, \boldsymbol{\phi}(t^v), t^v), \tag{11}$$

cf. (8). In the direct method in [16], the next reaction after t^v occurs at $t^v + \Delta\tau^v$, where the time increment $\Delta\tau^v$ is exponentially distributed with mean $1/w^v$. If the partitioning of the species into $\mathbf{x}$ and $\mathbf{y}$ based on the variances also is such that $\mathbf{m}_r = 0$ for the largest propensities w_r, i.e. w_r is small for $r \leqslant \rho$, then the expected value of $\Delta\tau^v$ is

longer than the time step without partitioning, making SSA more efficient. The reaction number μ with $1 \leqslant \mu \leqslant R$ is chosen with probability $w_\mu(\mathbf{x}^v, \boldsymbol{\phi}(t^v), t^v)/w^v$. With the stoichiometric matrix S defined by

$$S = (\mathbf{m}_1, \mathbf{m}_2, \ldots, \mathbf{m}_\rho)$$

and the unit vector $\mathbf{e}_j \in \mathbb{R}^\rho$ with 1 in the jth position, the algorithm can be written as a stochastic process for each trajectory

$$\mathbf{X}^{v+1} = \mathbf{X}^v - S\mathbf{e}_\mu, \quad t^{v+1} = t^v + \Delta\tau^v. \tag{12}$$

The M state vectors $\mathbf{X}_j^v$ are stored in a trajectory matrix $\mathbf{T}^v$ such that $\mathbf{T}_{j\cdot}^v = (\mathbf{X}_j^v)^{\mathrm{T}}, j = 1, \ldots, M$. The evaluation of p_0 in (10) for a given $\mathbf{x}$ is simplified if the rows of the matrix are sorted such that $\mathbf{T}_{j1}^v$ increases for increasing j. Then each interval in j with a constant $\mathbf{T}_{j,l}^v$ is sorted in ascending values of $\mathbf{T}_{j,l+1}^v$ for $l = 1, \ldots, M-1$. Identical rows are removed and the frequency of the row is saved. Then the row in $\mathbf{T}^v$ corresponding to a particular $\mathbf{x}_k$ is easily found by binary search.

3.2. Solution of the differential–summation equation

The time derivative of the expected value in (7) is approximated by an implicit backward differentiation formula of order 2 (BDF2) with variable time steps [19]. The increment between t^k and t^{k+1} is denoted by Δt^k. Then the new value of $\boldsymbol{\phi}^{k+1}$ at t^{k+1} is computed by

$$\begin{aligned}
&\alpha_0^k \boldsymbol{\phi}^{k+1} = \Delta t^k \mathbf{F}(\boldsymbol{\phi}^{k+1}, t^{k+1}) - \alpha_1^k \boldsymbol{\phi}^k - \alpha_2^k \boldsymbol{\phi}^{k-1},\\
&\boldsymbol{\omega}(\mathbf{x}, \boldsymbol{\phi}, t) = -\sum_{r=1}^{R} \mathbf{n}_r w_r(\mathbf{x}, \boldsymbol{\phi}, t),\\
&F_j(\boldsymbol{\phi}, t^{k+1}) = -\sum_{\mathbf{x}^{k+1} \in \mathbb{Z}_+^m} \omega_j(\mathbf{x}^{k+1}, \boldsymbol{\phi}, t^{k+1}) p_0(\mathbf{x}^{k+1}, t^{k+1}), \quad j = 1, \ldots, n,\\
&\alpha_0^k = (1 + 2\theta^k)/(1 + \theta^k), \quad \alpha_1^n = -(1 + \theta^k), \quad \alpha_2^k = (\theta^k)^2/(1 + \theta^k),\\
&\theta^k = \Delta t^k / \Delta t^{k-1},
\end{aligned} \tag{13}$$

see [19,29]. A predicted value $\hat{\boldsymbol{\phi}}^{k+1}$ is computed with an explicit scheme using

$$\begin{aligned}
&\hat{\alpha}_0^k \hat{\boldsymbol{\phi}}^{k+1} = \Delta t^k \mathbf{F}(\boldsymbol{\phi}^n, t^n) - \hat{\alpha}_1^k \boldsymbol{\phi}^n - \hat{\alpha}_2^k \boldsymbol{\phi}^{k-1},\\
&\hat{\alpha_0}^k = 1/(1 + \theta^k), \quad \hat{\alpha_1}^n = \theta^k - 1, \quad \hat{\alpha_2}^k = -(\theta^k)^2/(1 + \theta^k),
\end{aligned} \tag{14}$$

see [29]. The local temporal discretization error τ^{k+1} in the difference approximation is proportional to $(\Delta t^k)^2$. It is estimated by subtracting $\hat{\boldsymbol{\phi}}^{k+1}$ from $\boldsymbol{\phi}^{k+1}$ and dividing by Δt^k. The next time step Δt^{k+1} is chosen after comparing τ^{k+1} with an error tolerance so that τ^{n+2} is expected to satisfy the tolerance in the next step, see Section 4.4.

The system of nonlinear equations in (13) satisfied by $\boldsymbol{\phi}^{k+1}$ is solved by Newton iterations. The initial guess of $\boldsymbol{\phi}^{k+1}$ in the iterations is $\hat{\boldsymbol{\phi}}^{k+1}$. The elements of the Jacobian $\mathbf{J}$ in Newton's method are then computed as

$$J_{ij} = \left(\sum_{\mathbf{x} \in \mathbb{Z}_+^m} \omega_i(\mathbf{x}, \boldsymbol{\phi} + \mathbf{e}_j \Delta\phi_j, t) p_0(\mathbf{x}, t) - \sum_{\mathbf{x} \in \mathbb{Z}_+^m} \omega_i(\mathbf{x}, \boldsymbol{\phi}, t) p_0(\mathbf{x}, t) \right) \Big/ \Delta\phi_j, \quad i, j = 1, \ldots, n, \tag{15}$$

i.e. using first order forward differences. The second summation over $\mathbf{x}$ needs to be evaluated only once with this approximation of the derivatives. The number of participating species in every reaction is small implying that $\mathbf{J}$ is rather sparse. This fact or if κ is small in (9) makes the summation over $\mathbf{x}$ necessary only for a small fraction of $\mathbf{J}$. To take advantage of this, the system specification includes a list of the entries of the Jacobian that needs to be computed. It also states if the element is dependent of the stochastic subset. If not, it can be evaluated faster. The Jacobian is stored in factorized form and is recomputed only if the convergence of the iterations is too slow.

The stability of the integration in (13) is insensitive to the stiffness of $\boldsymbol{\omega}$ with respect to $\boldsymbol{\phi}$. If possible, species involved in reactions with propensities potentially contributing to the stiffness in the chemical system should be treated deterministically in $\boldsymbol{\phi}$.

The PDF p_0 in (13) is calculated from (10) where all the trajectories have been advanced by (12) so that $t^{\nu+1} \leqslant t^{k+1}$ but in the next SSA-step $t^{\nu+2} > t^{k+1}$. The time steps in (12) and (13) are then synchronized so that $t^{\nu+1} = t^{k+1}$. All the M state vectors are constant in the interval $[t^{\nu+1}, t^{\nu+2})$. The value of $\boldsymbol{\phi}^\nu$ in (11) at t^ν, $t^k < t^\nu \leqslant t^{k+1}$ is taken to be constant during the SSA steps and equals the value computed at t^k. Extrapolation of $\boldsymbol{\phi}$ to t^ν by

$$\boldsymbol{\phi}(t^\nu) = \boldsymbol{\phi}(t^k) + (t^\nu - t^k)(\boldsymbol{\phi}(t^k) - \boldsymbol{\phi}(t^{k-1}))/(t^k - t^{k-1})$$

would improve the accuracy in the evaluation of w_r in (5).

The time integration algorithm is summarized below.

1. Initialization: $k = 0$, $\nu = 0$, $t^0 = 0$, generate M state vectors $\mathbf{x}_j$ with a distribution given by the PDF $p_0(\mathbf{x}, 0)$, choose $t^{k=1} = \Delta t^0$ and $\phi_j(0)$, $j = 1, \ldots, n$.
2. Advance the solution of the master Eq. (5) from the previous step t^k with $\boldsymbol{\phi}(t^k)$ until $t^{\nu+1} \leqslant t^{k+1}$ but $t^{\nu+2} > t^{k+1}$.
3. Solve (13) for the expected values $\boldsymbol{\phi}^{k+1}$ given $p_0(\mathbf{x}, t^{\nu+1})$ determined by the M state vectors and (10) at $t^{\nu+1}$.
4. Compute a new time step Δt^{k+1}, $t^{k+2} = t^{k+1} + \Delta t^{k+1}$, $k := k + 1$, goto 2.

3.3. Approximation of the summation

The sums over $\mathbf{x}$ in (13) and (15) can be expensive to compute with exponential growth of the work in the dimension of $\mathbf{x}$ even if the summation is restricted to a finite but still multidimensional domain in $\mathbb{Z}_+^m$ where $p_0 \neq 0$. An alternative to summation over $\mathbb{Z}_+^m$ is to approximate the sums by a Monte Carlo (MC) or a Quasi-Monte Carlo (QMC) approach.

Let $\Omega = \{\mathbf{x} \mid 0 \leqslant x_i \leqslant x_{\max}\} \subset \mathbb{Z}_+^m$ for some $x_{\max} > 0$ and assume that $p_0 = 0$ outside Ω. Then in (13)

$$\sum_{\mathbf{x}\in\Omega} \boldsymbol{\omega}(\mathbf{x}, \boldsymbol{\phi}, t) p_0(\mathbf{x}, t) = E[\boldsymbol{\omega}(\mathbf{X}_0, \boldsymbol{\phi}, t)] \approx \frac{\zeta}{K} \sum_{k=1}^{K} \boldsymbol{\omega}(\mathbf{x}_k, \boldsymbol{\phi}, t) p_0(\mathbf{x}_k, t), \tag{16}$$

where $\mathbf{X}_0$ has the PDF p_0. Let $x_{j,\min}$ and $x_{j,\max} \leqslant x_{\max}$ be such that $p_0 = 0$ if $x_j < x_{j,\min}$ or $x_j > x_{j,\max}$ for every j. Then the factor in (16) is

$$\zeta = \prod_{j=1}^{\rho} ((x_{j,\max} + 1) - x_{j,\min}).$$

The K quadrature points $\mathbf{x}_k$ are chosen by an MC or QMC method.

In a standard MC method, the sum in (16) is evaluated by generating pseudorandom vectors $\mathbf{x}_k$ with x_{ki} having uniform distribution in $[x_{i,\min}, x_{i,\min} + 1, \ldots, x_{i,\max}]$. The method converges slowly with the rate $K^{-1/2}$ but independently of m and the regularity of $\boldsymbol{\omega} p_0$ [3,18].

An alternative is to use QMC methods where the sequence of quasi-random $\mathbf{x}_k$ is generated deterministically. For x_{ki}, a quasi-random number ξ in $[0, 1)$ is first generated according to some rule. Then it is scaled so that x_{ki} is the integer part of $x_{j,\min} + ((x_{j,\max} + 1) - x_{j,\min})\xi$. The convergence rate for smooth functions is now of $\mathcal{O}(K^{-1}(\log K)^m)$, which is an improvement compared to MC at least for a moderate m [3]. The summation errors in MC and QMC are estimated in Sections 4.2 and 4.3.

The sum (16) can be computed by an acceptance–rejection method [3,18] where a uniformly distributed trial point $\mathbf{x}_k$ is accepted in the sum with probability proportional to p_0. Then the summands can be written $v_k \omega_j(\mathbf{x}_k, \boldsymbol{\phi}, t)$ where the weights v_k are one or zero depending on if the point is accepted or rejected. The convergence rate of QMC depends on the smoothness of the summand and it may deteriorate because of the lack of regularity in $v_k \omega_j(\mathbf{x}_k, \boldsymbol{\phi}, t)$. By introducing a linear approximation of v_k close to the switch from 1 to 0, better convergence is achieved in [3,32]. The convergence of the MC and QMC methods are compared using Faure sequences [12] for QMC in a numerical example in Section 5.

4. Analysis of the algorithm

The three major sources of approximation errors are the calculation of the PDF in Section 3.1, the evaluation of the sums with MC or QMC in Section 3.3, and the time integration of the differential–summation equation in Section 3.2. These errors are analyzed and computable estimates of them are proposed in this section and the growth of the computational work is estimated when the number of trajectories M and the number of quadrature points K increase.

4.1. Error in the marginal PDF

The stochastic variable Ψ_j at t^ν in (10) is 1 with probability $p_0(\mathbf{x}, t^\nu)$ in trajectory j of the chemical system. Then according to the law of the large numbers [3,18], p_0 is approximated at $\mathbf{x} = \mathbf{x}_k$ by

$$\bar{p}^\nu_{0k} = \frac{1}{M}\sum_{j=1}^{M} \Psi_j.$$

In order to estimate the variance of the PDF at t^ν and $\mathbf{x}_k$ the original trajectory matrix is subdivided in J parts each one of size $\widetilde{M}$ and with an approximation $\tilde{p}^{\nu,j}_{0k}$. The variance σ^2_{Mk} in $\tilde{p}^{\nu,j}_{0k}$ is estimated by the sample variance

$$s^2_{Mk} = \frac{1}{J-1}\sum_{j=1}^{J}\left(\tilde{p}^{\nu,j}_{0k} - \bar{p}^\nu_{0k}\right)^2, \tag{17}$$

where the pooled estimate is given by $\bar{p}^\nu_{0k} = \frac{1}{J}\sum_{j=1}^{J}\tilde{p}^{\nu,j}_{0k}$. The error ϵ_{Mk} in $\bar{p}^\nu_{0k}$ is normally distributed by the central limit theorem, $\mathcal{N}(0, \sigma^2_{Mk}/J)$. If σ^2_Ψ is the variance of Ψ_j, then σ^2_{Mk} is $\sigma^2_\Psi/\widetilde{M}$ and the variance of $\bar{p}^\nu_{0k}$ is $\sigma^2_\Psi/(J\widetilde{M}) = \sigma^2_\Psi/M$. An approximate 95% confidence interval for ϵ_{Mk} is given by the Student t-distribution [18]

$$|\epsilon_{Mk}| \leqslant 2s_{Mk}/\sqrt{J}. \tag{18}$$

4.2. Error in the expected values

Suppose that the PDF in (16) is known by its exact value. The sum is the expected value of $\boldsymbol{\omega}$ at time t with the parameter $\boldsymbol{\phi}$. In an MC method, the sum is computed at t^ν by determining trial vectors $\mathbf{x}_k$ with uniformly distributed components. Let the sum of the ith component at t^ν be evaluated Q times using K trial vectors each time

$$S_{qi} = \frac{\zeta}{K}\sum_{k=1}^{K}\omega_i(\mathbf{x}^\nu_{K(q-1)+k}, \boldsymbol{\phi}^\nu, t^\nu)p_0(\mathbf{x}^\nu_{K(q-1)+k}, t^\nu), \quad q = 1, \ldots, Q. \tag{19}$$

Usually, Q is much smaller than K. The pooled estimate of $E[\boldsymbol{\omega}(\mathbf{X}, \boldsymbol{\phi}, t)]$ is

$$S_i = \frac{1}{Q}\sum_{q=1}^{Q} S_{qi}. \tag{20}$$

It approximates the expected value with the error ϵ_{Ki}. The sample variance s^2_i of S_{qi} is

$$s^2_i = \frac{1}{Q-1}\sum_{q=1}^{Q}(S_{qi} - S_i)^2. \tag{21}$$

It follows again from the central limit theorem that the error in S_i with probability 0.95 is bounded by

$$|\epsilon_{Ki}| \leqslant 2s_i/\sqrt{Q}. \tag{22}$$

The error ϵ_{Ki} is normally distributed, $\mathcal{N}(0, \sigma^2_{Ki}/Q)$, where σ^2_{Ki} is the variance of S_{qi}. According to the central limit theorem, σ_{Ki} decays as $K^{-1/2}$ when K increases. Consequently, the bound on ϵ_{Ki} is proportional to $1/\sqrt{KQ}$.

The error in the QMC method is estimated by scrambling the sequences as in [22,33]. The digits of a quasi-random number in a base are permuted randomly to obtain the scrambled number. The advantage with scrambling is that the rapid convergence of QMC is preserved while allowing for error estimates as in MC. The sum S_{qi} is computed as in (19). The error estimate is given by (22). Since $\sigma_{Ki} \sim K^{-1}$, the right hand side in (22) decreases as $1/(K\sqrt{Q})$ for increasing K and Q and a moderate m.

4.3. The total summation error

With $\boldsymbol{\omega}_k^v = \boldsymbol{\omega}(\mathbf{x}_k^v, \boldsymbol{\phi}^v, t^v)$, we derive the total error in the summation from (16)

$$E[\boldsymbol{\omega}(\mathbf{X}_0, \boldsymbol{\phi}, t^v)] = \frac{\zeta}{K}\sum_{k=1}^{K} \boldsymbol{\omega}_k^v p_0(\mathbf{x}_k^v, t^v) + \boldsymbol{\epsilon}_K = \frac{\zeta}{K}\sum_{k=1}^{K} \boldsymbol{\omega}_k^v \bar{p}_{0k}^v + \frac{\zeta}{K}\sum_{k=1}^{K} \boldsymbol{\omega}_k^v (p_0(\mathbf{x}_k^v, t^v) - \bar{p}_{0k}^v) + \boldsymbol{\epsilon}_K. \tag{23}$$

The second sum on the right hand side is a sum of independent and normally distributed variables $\epsilon_{Mk} = p_0(\mathbf{x}_k^v, t^v) - \bar{p}_{0k}^v$. The difference ϵ_i^v between the expected value and the sum $\mathbf{S}$ is in the ith component

$$\epsilon_i^v = E[\omega_i(\mathbf{X}_0, \boldsymbol{\phi}, t^v)] - S_i = \epsilon_{Ki} + \frac{\zeta}{K}\sum_{k=1}^{K} \omega_{ki}^v \epsilon_{Mk}. \tag{24}$$

It follows from Section 4.1 that the sum Σ_i in (24) has the distribution

$$\Sigma_i \sim \mathcal{N}\left(0, \left(\frac{\zeta}{K}\right)^2 \sum_{k=1}^{K} (\omega_{ki}^v)^2 \frac{\sigma_{Mk}^2}{J}\right).$$

If $\sigma_{Mk} \leqslant \sigma_M$ for all k, then the variance of Σ_i is bounded by

$$\mathrm{Var}(\Sigma_i) \leqslant \frac{\zeta}{K}\frac{\sigma_M^2}{J}\frac{\zeta}{K}\sum_{k=1}^{K} \left(\omega_{ki}^v\right)^2. \tag{25}$$

By (24), Section 4.2, (25), and the bound $\sigma_{Ki} \leqslant \sigma_K, i = 1, \ldots, n$, the total summation error ϵ_i^v is normally distributed with zero mean value and a variance that is bounded by

$$\mathrm{Var}(\epsilon_i^v) \leqslant \frac{\sigma_K^2}{Q} + \frac{\zeta}{K}\frac{\sigma_M^2}{J}\frac{\zeta}{K}\sum_{k=1}^{K} (\omega_{ki}^v)^2, \quad i = 1, \ldots, n. \tag{26}$$

The variance σ_K^2 and the upper bound σ_M^2 in (26) are estimated by the sample variances in (17) and (21). Since σ_K^2 decays with increasing K and σ_M^2 with increasing M, $\mathrm{Var}(\epsilon_i^v)$ can be made as small as we wish by choosing K and M sufficiently large.

4.4. The integration error

The evaluation of $\mathbf{F}$ in (13) is perturbed by the error $\boldsymbol{\epsilon}^{k+1}$ caused by the summation. It is equal to $\boldsymbol{\epsilon}^v$ in (24), where t^v is the time for the most advanced trajectory in time in the SSA but still $t^v \leqslant t^{k+1}$. The discretization error in the approximation of the time derivative is denoted by $\boldsymbol{\tau}^{k+1}$. Following [19] and [29], the local error in the corrector solution $\boldsymbol{\phi}^{k+1}$ of (13) is

$$\begin{aligned} \boldsymbol{\phi}(t^{k+1}) - \boldsymbol{\phi}^{k+1} &= \alpha_0^{-1}\Delta t^k(\boldsymbol{\tau}_c^{k+1} - \boldsymbol{\epsilon}^{k+1}), \\ \boldsymbol{\tau}_c^{k+1} &= -\frac{1+\theta^k}{6\theta^k}(\Delta t^k)^2 \boldsymbol{\phi}''' + \mathcal{O}((\Delta t)^3), \end{aligned} \tag{27}$$

where $\Delta t = \max(\Delta t^k, \Delta t^{k-1})$. The local error in the predictor $\hat{\boldsymbol{\phi}}^{k+1}$ from (14) is

$$\begin{aligned} \boldsymbol{\phi}(t^{k+1}) - \hat{\boldsymbol{\phi}}^{k+1} &= \hat{\alpha}_0^{-1}\Delta t^k(\boldsymbol{\tau}_p^{k+1} - \boldsymbol{\epsilon}^k), \\ \boldsymbol{\tau}_p^{k+1} &= \frac{1}{6\theta^k}(\Delta t^k)^2 \boldsymbol{\phi}''' + \mathcal{O}((\Delta t)^3). \end{aligned} \tag{28}$$

The error equation approximately satisfied by the global error $\mathbf{e}$ in the corrector is

$$\frac{\mathrm{d}\mathbf{e}}{\mathrm{d}t} - \frac{\partial \mathbf{F}}{\partial \boldsymbol{\phi}}\mathbf{e} = \boldsymbol{\tau}_c - \boldsymbol{\epsilon}. \tag{29}$$

The aim is to control the driving right hand side in (29) by taking Δt^k sufficiently small and K and M sufficiently large.

The leading term in the temporal discretization error $\boldsymbol{\tau}_c^{k+1}$ is computed by first combining (27) and (28) and then obtaining an estimate of $\boldsymbol{\phi}'''(t^{k+1})$. This estimate is inserted into $\boldsymbol{\tau}_c^{k+1}$ in (27) to arrive at an expression for the jth entry

$$\tau_{cj}^{k+1} - \epsilon_j^{k+1} = -\frac{1+2\theta^k}{2+3\theta^k}\left(\frac{\phi_j^{k+1} - \hat{\phi}_j^{k+1}}{\Delta t^k} + (1+\theta^k)\epsilon_j^k + \epsilon_j^{k+1}\right) + \mathcal{O}(\Delta t^3). \tag{30}$$

Let the variance of ϵ_j^l be denoted by σ_ϵ^2. Since ϵ_j^l is normally distributed with mean 0, $\mathcal{N}(0,\sigma_\epsilon^2)$, (see Section 4.3) and independent of $\epsilon_j^\ell, \ell \neq l$, the leading term in (30) is

$$\mathcal{N}\left(-\frac{1+2\theta^k}{2+3\theta^k}\,\frac{\phi_j^{k+1} - \hat{\phi}_j^{k+1}}{\Delta t^k}, \left(\frac{1+2\theta^k}{2+3\theta^k}\right)^2((1+\theta^k)^2+1)\sigma_\epsilon^2\right). \tag{31}$$

The first term in (30) depends on Δt^2 but the second and third terms are independent of Δt. Suppose that we want the error in the right hand side of (29) to be less than η_j for all components j with a 95% confidence interval. Choose a fraction γ, $0<\gamma<1$, a time step Δt^k, K, Q, M, and J such that

$$\begin{aligned}
&\left|\frac{\phi_j^{k+1} - \hat{\phi}_j^{k+1}}{\Delta t^k}\right| \leqslant \frac{2+3\theta^k}{1+2\theta^k}\gamma\eta_j,\\
&\sigma_\epsilon \leqslant \left(\frac{\sigma_K^2}{Q} + \frac{\zeta}{K}\frac{\sigma_M^2}{J}\frac{\zeta}{K}\sum_{k=1}^{K}(\omega_{kj}^v)^2\right)^{1/2} \leqslant \frac{2+3\theta^k}{1.96(1+2\theta^k)\sqrt{1+(1+\theta^k)^2}}(1-\gamma)\eta_j
\end{aligned} \tag{32}$$

for all j using (26). Then by (30) we have

$$|\tau_{cj}^{k+1} - \epsilon_j^{k+1}| \leqslant \eta_j \tag{33}$$

with probability 0.95. The variance σ_{Ki}^2 and the bound σ_M are estimated by (21) and (17). The parameter γ determines the fraction of the bound on $\tau_{cj}^{k+1} - \epsilon_j^{k+1}$ to be satisfied by the first, deterministic term in the error. The tolerance parameter η_j is a function of the relative and absolute error tolerances

$$\eta_j = \max\left(RelTol * \left|\frac{\mathrm{d}\phi_j^{k+1}}{\mathrm{d}t}\right|, AbsTol\right). \tag{34}$$

With constant time steps, $\theta^k = 1$, the upper bounds in (32) are simplified to

$$\left|\frac{\phi_j^{k+1} - \hat{\phi}_j^{k+1}}{\Delta t^k}\right| \leqslant 1.67\gamma\eta_j, \quad \sigma_\epsilon \leqslant 0.38(1-\gamma)\eta_j.$$

4.5. Computational work

Here we consider the complexity of the hybrid method, and discuss when use of the hybrid method will be advantageous compared to SSA.

The work of SSA grows linearly with the number of reactions R for one trajectory. Suppose that the work in SSA of the full system to determine M trajectories is $c_{\mathrm{SSA}}R\,M$ and that the corresponding work for the reduced system in the hybrid algorithm is $c_{\mathrm{SSA}}^{\mathrm{hyb}}\rho M$. Usually, $c_{\mathrm{SSA}}^{\mathrm{hyb}} \leqslant c_{\mathrm{SSA}}$ because some stiffness in the full system is removed by the partitioning.

Assuming $c_{\mathrm{SSA}} = c_{\mathrm{SSA}}^{\mathrm{hyb}}$, a speedup of order R/ρ of the SSA part is obtained when the system is reduced. Often, reactions involving the variables taken in the deterministic subset are relatively fast due to large copy numbers, and then $c_{\mathrm{SSA}}^{\mathrm{hyb}}$ will be much smaller than c_{SSA} and the potential speedup is large. The maximal possible gain given a hybrid splitting is

$$g_{\max} = Rc_{\mathrm{SSA}}/\rho c_{\mathrm{SSA}}^{\mathrm{hyb}} \tag{35}$$

and is obtained if the overhead due to the deterministic part of the solver is negligible compared to the time spent in SSA in the hybrid system. This maximal speedup can be estimated from one simulation of the full system and one with the hybrid solver using few trajectories.

Let T_{SSA} be the total time for SSA and T_{hyb} the total time in the hybrid algorithm. If the fraction spent in SSA in the hybrid algorithm is ϑ, the speedup is

$$\frac{T_{\mathrm{SSA}}}{T_{\mathrm{hyb}}} = \vartheta g_{\max}. \tag{36}$$

If the splitting results in a large potential speedup, then SSA will not in general be the dominating part of the hybrid algorithm. The overhead associated with the solution of the deterministic equations (7) is estimated by considering the other major contributions to the algorithm. Assuming that a total of QK QMC vectors are used in the evaluation of the sum, p_0 has to be calculated at QK points. For the evaluation of p_0, the trajectory matrix $\mathbf{T}$ is first sorted according to rows using merge sort [6, p. 28] with a work of $\mathcal{O}(mM\log M)$. Obviously, many rows of $\mathbf{T}$ will consist of equal vectors for M reasonably large. They are removed in the sorted list with a work of $\mathcal{O}(mM)$, giving a modified trajectory matrix $\hat{\mathbf{T}}$ with $\widehat{M} \leqslant M$ rows. One quadrature point is then found in $\hat{\mathbf{T}}$ using binary search with work proportional to $m\log\widehat{M}$. In total, the work of evaluating the sum is $\mathcal{O}(mM\log M) + \mathcal{O}(mQK\log\widehat{M})$. The evaluation of the full Jacobian $\mathbf{J}$ and $\mathbf{F}$ would need work of $\mathcal{O}(n^2QK)$ and $\mathcal{O}(nQK)$, respectively, and the solution of the system of linear equations in the Newton iterations costs $\mathcal{O}(n^2)$ operations with a factorized $\mathbf{J}$. Thanks to the sparsity of the Jacobian the work is typically much lower, and for larger models an iterative method could be used. The dominant terms in the computational time in one time step of the hybrid algorithm are

$$c_{\mathrm{SSA}}^{\mathrm{hyb}}\rho M + c_1 mM\log M + c_2 mQK\log\widehat{M}. \tag{37}$$

In merge sort, we have $c_1 = 2$. With N_s deterministic time steps, ϑ in (36) can be estimated as

$$\vartheta \approx \frac{1}{1 + N_s\left(\frac{2m\log M}{c_{\mathrm{SSA}}^{\mathrm{hyb}}\rho} + \frac{c_2 mQK\log\widehat{M}}{c_{\mathrm{SSA}}^{\mathrm{hyb}}\rho M}\right)} \tag{38}$$

For a given system, QK is chosen to achieve an acceptable summation error. For large values of QK and small values of M, the term related to QMC summation will be large, and ϑ small, i.e. speedups far from $g_{\max}$ will be obtained. When M is increased with constant QK, this term will decrease and ϑ increase until the contribution from the summation term is small. The term associated with the sorting will increase slowly as $\log M$, and eventually ϑ will decrease due to this term. With a large value of $c_{\mathrm{SSA}}^{\mathrm{hyb}}$ it will be possible to keep ϑ close to 1. On the other hand, the maximal speedup $g_{\max}$ will not be large. We will see a scenario when this is the case in Section 5.2, and one example when $g_{\max}$ is large but ϑ small, still resulting in a significant speedup in Section 5.3. Obviously, the fewer timesteps N_s taken in the macroscopic solver, the higher value of ϑ, and the greater the speedup. Fewer time steps are usually taken with an adaptive algorithm. As a last observation, note that without summation with an MC or a QMC method, the work would grow exponentially with m.

5. Numerical results

The hybrid algorithm is applied to the simulation of three different chemical systems. The first system is small with two metabolites and two enzymes and serves as an illustration of some of the limitations of the hybrid method, but also develops guidelines for how to choose some of the parameters in the algorithm. The second system models a molecular clock where the reaction rate equations fail to produce oscillations for a certain parameter value. Introducing a few stochastic variables motivated by the simplification in [43]

makes the model much more robust to changes in the parameter. This is an example where addition of stochastic noise improves the original model based on the reaction rate equations. Here we show that the hybrid method is able to capture this behavior. The third system is a mitogen-activated protein kinase (MAPK) signaling cascade with 22 molecular species. Only three of these species are treated stochastically. The results from the simulations with the full system and SSA and the reduced system and our algorithm are similar but the computational work to simulate the reduced system with good accuracy is an order of magnitude lower.

The most time consuming parts of the hybrid solver, i.e. SSA, evaluation of p_0, QMC summation, and computation of the Jacobian $\mathbf{J}$, are implemented in C and wrapped as mex-files in a MATLAB environment. Different systems are defined by one function that performs SSA on the reduced system, one function that returns the right hand side of the differential–summation equations and one function that defines the sparsity pattern of the Jacobian matrix. To facilitate the generation of different hybrid splittings, a graphical tool has been developed in Python that reads models in the SBML format [23] and after manual selection of the stochastic variables automatically generates the necessary C files given a desired partitioning. The reaction rate equations are integrated by MATLAB's `ode15s`. The code for generation of the quasi-sequence is written as a mex-file calling a Fortran subroutine from [22]. This part is executed once initially and contributes little to the total CPU time if the number of time steps is sufficiently large. Some parts of the algorithm can be run in parallel using openMP, to take advantage of the multicore capabilities of most modern computers. However, all the experiments here have been run with a serial code on an Intel Core Duo MacBook with 2.0 GHz processors, 2MB L2 cache and 2 GB RAM.

5.1. Metabolites controlled by enzymes

The first system is a simple generic model with two metabolites A and B and two enzymes E_A and E_B as in [28]. The hybrid method does not outperform SSA for this system, as we will see, but the system is small and the behavior of different parts of the algorithm is easily investigated. It will provide insight into some of the limitations of the algorithm, but also give guidelines for how to choose some of the parameters.

The production of A and B is regulated by the enzymes. The reactions are

$$\emptyset \xrightarrow{\frac{k_a e_A}{1+\frac{a}{K_i}}} A \qquad \emptyset \xrightarrow{\frac{k_b e_B}{1+\frac{b}{K_i}}} B$$

$$A + B \xrightarrow{k_2 ab} \emptyset$$

$$A \xrightarrow{\mu a} \emptyset \qquad B \xrightarrow{\mu b} \emptyset$$

$$\emptyset \xrightarrow{\frac{k_{e_A}}{1+\frac{a}{K_r}}} E_A \emptyset \qquad \xrightarrow{\frac{k_{e_B}}{1+\frac{b}{K_r}}} E_B$$

$$E_A \xrightarrow{\mu e_A} \emptyset \qquad E_B \xrightarrow{\mu e_B} \emptyset$$

The reaction constants above are given in Table 1.

There are no fast time scales for the reactions and simulation with SSA is easily conducted for this system, so we would not expect the hybrid algorithm to result in a considerable speedup in this case. The first and second central moments of the distribution, μ and σ^2, are computed by simulation with SSA. Table 2 shows the maximal value of $\sigma(t)/\mu(t)$ for the different species during 1000 s and Table 3 the correlation coefficients at $t = 1000$ s using 10^5 trajectories.

The standard deviation relative to the mean value is neither very large nor very small.

Table 1
The parameters for the metabolite–enzyme model

k_a	k_b	k_2	K_i	μ	k_{eA}	k_{eB}	K_r
0.3	0.3	0.001	60	0.002	0.02	0.02	30

Table 2
The maximal value of $\sigma(t)/\mu(t)$ during the first 1000 s of simulation

A	B	E_A	E_B
0.52	0.52	0.35	0.35

Table 3
The correlation coefficients at $t = 1000$ s

	A	B	E_A	E_B
A	1	−0.75	0.59	−0.34
B	−0.75	1	−0.34	0.59
E_A	0.59	−0.34	1	0.22
E_B	−0.34	0.59	0.22	1

The system is partitioned so that the metabolites A and B are treated as stochastic variables and the enzymes E_A and E_B are assumed to be normally distributed with a small variance. In this case, the number of stochastic variables decreases from $N = 4$ in the full model to $m = 2$ and $n = 2$, the number of reactions is reduced from nine to $\rho = 5$ in (8). The number of reactions with propensities depending on $\mathbf{x}$, κ, is seven in (9).

The marginal probabilities for the metabolites and the enzymes calculated from $M = 10^5$ trajectories of SSA are plotted in Fig. 1. The system is close to the steady state at $t = 1000$ s. The fluctuations of the species A and B in (10) appear to be larger than for the enzymes in the figure, in agreement with the values in Table 2. Fig. 2 shows the system simulated with the hybrid solver to the final time $t = 1000$ s. The time steps are chosen adaptively with a relative tolerance of 0.01 and an absolute tolerance 10^{-4}. The maximal time step $\Delta t_{\max}$ is 5. The summation is performed by using $Q = 10$ scrambled sequences of $K = 2^{15}$ points. They are generated and stored initially and then used to form estimates of the sum and the error as in (20) and (22). The maximal value of the estimated leading terms of the error in (30) and σ_ϵ in (32) are shown in Fig. 3. For this system, the error tolerance with $\gamma = 0.5$ was easily met, and the time step reached its maximal value after approximatively 55 s and remained fixed for the rest of the simulation.

Clearly, the hybrid solver in Fig. 2 does not capture the characteristic shape of the marginal distribution of A and B sufficiently well. At $t = 1000$ s, the correlation coefficient between the metabolites was −0.34. These values are to be compared to Tables 2 and 3. We expect the hybrid algorithm to give inaccurate results in this case, since the simplifying assumptions in the derivation of the underlying equations are violated. First of all, the relative standard deviations of the species in the deterministic subset are not very small. Moreover, the correlation between the species in the stochastic and deterministic subsets is significant. A good splitting

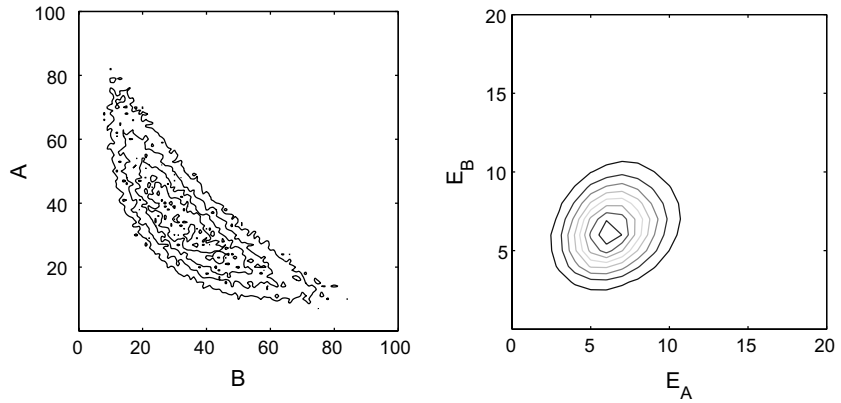

Fig. 1. Isolines of the marginal probability density at $t = 1000$ for A and B (left) and E_A and E_B (right) computed by SSA with $M = 10^5$.

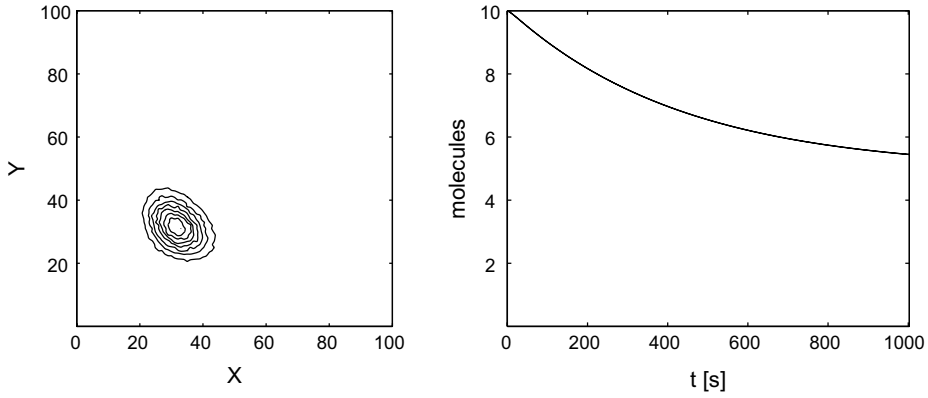

Fig. 2. The probability density for A and B at $t = 1000$ (left) and the expected values for E_A and E_B (right) computed with the hybrid solver.

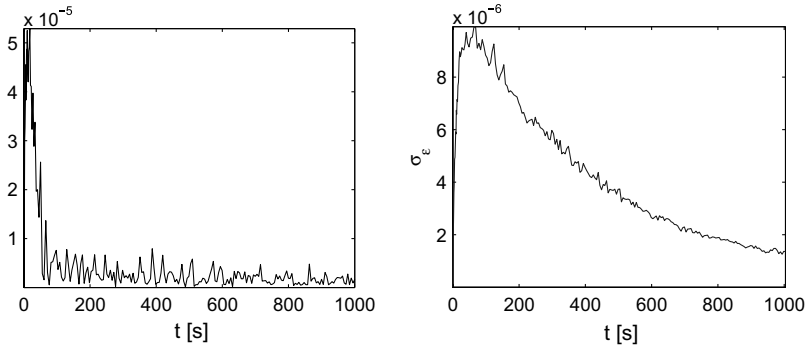

Fig. 3. The maximal local time discretization errors (left) and the bound on σ_ϵ (right) with $K = 2^{15}$, $M = 10^5$.

ideally keeps highly correlated species within the stochastic subset. However, the distribution obtained with this solver compares very well with the results in [28] where a Fokker–Planck equation approximates the master equation.

Different summation techniques for (16) are evaluated in Fig. 4 with the PDF based on $M = 10^5$ realizations at $t = 1000$. The slope of the curves has been determined by a least squares fit to the data. The sums

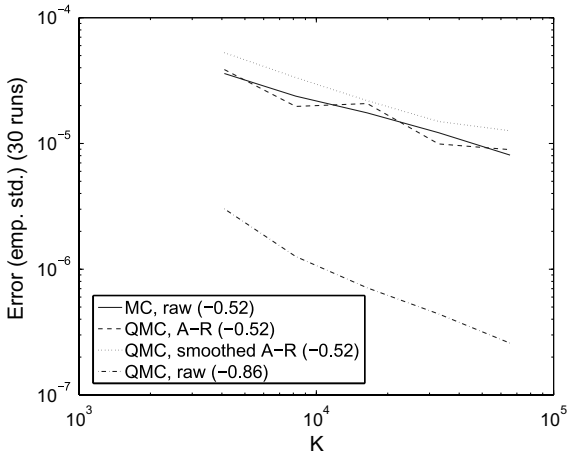

Fig. 4. The sample standard deviation for different MC and QMC strategies. The number in parenthesis is the inclination of the curve.

$\mathbf{S}_q$, $q = 1, \ldots, Q$, are computed for $Q = 30$. Then the sample or empirical standard deviation is determined componentwise as in (21).

In Fig. 4, the first component of the sample standard deviation is displayed. Raw MC summation with pseudorandom numbers is compared to raw QMC summation using numbers from the Faure sequence generated with the algorithm in [22]. The acceptance–rejection method (A–R) [3,18] and a smoothing alternative [3,32] are tested. The discontinuity in the A–R method is replaced by a linear function between 0 and 1 in [32]. One important consideration is the fact that when for example the A–R method is used to generate numbers from the distribution, fewer evaluation points will be accepted and used than in the raw MC method. This means that K in (16) will be different for different methods. The errors reported in Fig. 4 refer to the error obtained with a fixed number of trial points K because we are interested in which method gives the smallest error per generated random number to account for the cost in generating the numbers and most of all in evaluating the PDF. The best convergence rate with the lowest error is achieved by standard QMC. This is the preferred method in the sequel.

From (32) and Section 4.2 some conclusions can be drawn concerning how to choose K and Q. The bound of σ_ϵ^2 in (32) consists of two terms, one corresponding to σ_K^2 and the other to σ_M^2. Fig. 5 shows these terms separately as computed during the simulation. In this case, the second term dominates.

The state space is quite small in this example. An upper bound on the second term $\sigma_{\epsilon 2}^2$ in the error is

$$\sigma_{\epsilon 2}^2 = \frac{\zeta}{K}\frac{\sigma_M^2}{J}\frac{\zeta}{K}\sum_{k=1}^{K}(\omega_{kj}^v)^2 \leqslant \frac{\zeta^2\sigma_M^2}{JK}\max_{\mathbf{x}\in\Omega}\omega_i(\mathbf{x},\phi,t)^2. \tag{39}$$

If the reactions are fast, and the macroscopic equations are stiff, then ω_i can be large. To ensure that the $\sigma_{\epsilon 2}^2$ term stays small it is advantageous to choose large values of K and M. The error in the QMC approximation of the sum decreases as $1/(K\sqrt{Q})$. For a constant KQ the smallest error is obtained for $Q = 1$. The price when $Q = 1$ is that no simple error estimation is available. The results of the other two simulations in Sections 5.2 and 5.3 are consistent with these conclusions.

The Matlab tool 'Profiler' has been invoked to determine the time spent in different parts of the algorithm. First the hybrid solver with full error estimation is compared to SSA. The system was simulated to $t = 1000$ s with fixed time step $\Delta t = 1$ s and ten QMC sequences of $2^{15} \approx 33 \times 10^3$ points. The number of trajectories was $M = 10^5$. In this setting, simulation with SSA took 296 s, while the hybrid algorithm required 1360 s simulation time, i.e. SSA was about 4.6 times faster. This is explained by the small size of the system and no fast scales are separated. Table 4 shows the performance when the system is simulated to $t = 1000$ s with adaptive

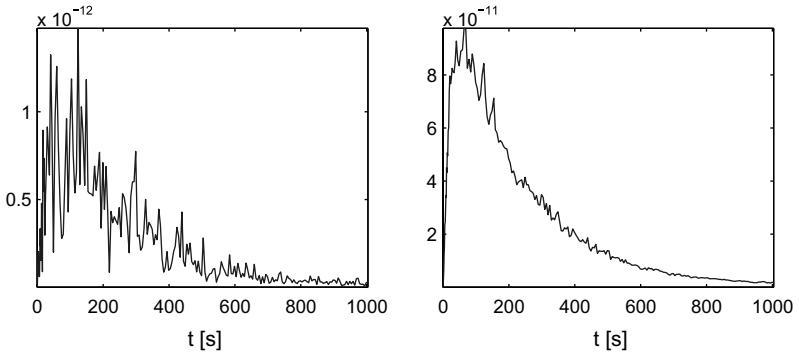

Fig. 5. The first term involving σ_K^2 (left) and the second term involving σ_M^2 (right) in (32) with $K = 2^{15}$, $M = 10^5$.

Table 4
Time spent in SSA in the hybrid method and a comparison with full SSA

Number of trajectories	10^5	10^6
Time spent in SSA [%]	71.0	78.2
Total CPU time [s]	184	1676
$T_{\mathrm{SSA}}/T_{\mathrm{hyb}}$	0.95	1.04

time step selection with a maximal time step of $\Delta t = 5$ s, using two QMC sequences of $K = 2^{16} \approx 65 \times 10^3$ points and no summation error estimation.

In this case, the hybrid algorithm is as fast as SSA for the full system.

5.2. Circadian rhythm model

The oscillator in [1,43] is a model for circadian rhythms. This kind of control system assures that periodic oscillations of certain molecular species appear in order to establish a circadian rhythm in the organism.

The model has nine variables. Two genes, D_a and D_r and their corresponding mRNA, M_a and M_r, are controlled by an activator and a repressor A and R, synthesized from the respective mRNA. The activator and repressor can associate and form a complex C, in which the activator A is degraded. The variables D_a' and D_r' are the genes D_a and D_r with a bound activator. In the model it is assumed that there is only one gene coding for the repressor and the activator. Thus $D_a + D_a' = 1$, and the same holds true for the repressor gene. In the numerical simulations, the problem is scaled by setting the initial conditions for the genes to $D_a = D_r = 0.2$. The 18 reactions for the nine molecular species are

$$\left.\begin{array}{l} D_a' \overset{\theta_a D_a'}{\longrightarrow} D_a \\ D_a + A \overset{\gamma_a D_a A}{\longrightarrow} D_a' \\ D_r' \overset{\theta_r D_r'}{\longrightarrow} D_r \\ D_r' + A \overset{\gamma_r D_r A}{\longrightarrow} D_r' \end{array}\right\} \quad \left.\begin{array}{l} \emptyset \overset{\alpha_a' D_a'}{\longrightarrow} M_a \\ \emptyset \overset{\alpha_a D_a}{\longrightarrow} M_a \\ M_a \overset{\delta_{ma} M_a}{\longrightarrow} \emptyset \end{array}\right\} \quad \left.\begin{array}{l} \emptyset \overset{\beta_a M_a}{\longrightarrow} A \\ \emptyset \overset{\theta_a D_a'}{\longrightarrow} A \\ \emptyset \overset{\theta_r D_r'}{\longrightarrow} A \\ A \overset{\delta_a A}{\longrightarrow} \emptyset \\ A + R \overset{\gamma_c AR}{\longrightarrow} C \end{array}\right\} \quad \left.\begin{array}{l} \emptyset \overset{\alpha_r' D_r'}{\longrightarrow} M_r \\ \emptyset \overset{\alpha_r D_r}{\longrightarrow} M_r \\ M_r \overset{\delta_{mr} M_r}{\longrightarrow} \emptyset \end{array}\right\} \quad \left.\begin{array}{l} \emptyset \overset{\beta_r M_r}{\longrightarrow} R \\ R \overset{\delta_r R}{\longrightarrow} \emptyset \\ C \overset{\delta_a C}{\longrightarrow} R \end{array}\right\} \tag{40}$$

The reaction constants are found in Table 5. For these parameters the system exhibits a limit cycle but if the parameter δ_r is sufficiently small then the macroscopic reaction rate equations quickly reach a stable fixed point and the oscillations stop, see Fig. 6. It is shown by Vilar et al in [43] that a mesoscopic description of the system continues to produce reliable oscillations. The stochastic noise is obviously sufficient to perturb

Table 5
Parameters for the Vilar oscillator

α_A	α_a'	α_r	α_r'	β_a	β_r	δ_{ma}	δ_{mr}
50.0	500.0	0.01	50	50.0	5.0	10.0	0.5
δ_a	δ_r	γ_a	γ_r	γ_c	Θ_a	Θ_r	
1.0	0.2	1.0	1.0	2.0	50.0	100.0	

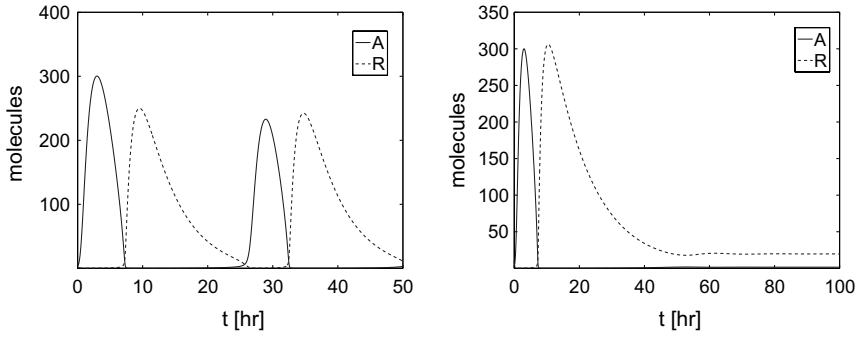

Fig. 6. Time evolution of the repressor R computed with the reaction rate equations. The parameter δ_r is 0.2 (left) and 0.08 (right).

the trajectories away from the fixed point to initiate new cycles. To be useful, the hybrid algorithm should capture this behavior qualitatively. The parameter δ_r has two different values 0.2 and 0.08 in the numerical experiments.

The partitioning is inspired by the discussion in [43]. In this type of control system, we expect a high correlation between many of the species. Even if the species A and R have large copy numbers most of the time, at critical time intervals they drop near zero. Here, the variables are partitioned into two subsets: A and R are treated stochastically, corresponding to $\mathbf{X}$ in Section 2 with $m = 2$, while the other variables are treated deterministically, corresponding to $\mathbf{Y}$ with $n = 7$, with $g_{\max} = 2.0$. There are ten stochastic reactions ($\rho = 10$ in (8)).

The initial conditions for the deterministic variables are 0.2 for the genes, 10 for C and 0 for the other species. The stochastic variables are initiated as normal distributions centered around 10. The hybrid system was solved with adaptive time step selection with relative tolerance 0.05 and absolute tolerance 10^{-3}. The probability density p_0 is approximated with $M = 10^6$ trajectories in (10) and $K = 2^{20} \approx 1.05 \times 10^6$ quasi-random points in (16) are used in each sequence in the integration algorithm and $Q = 5$.

Sustained oscillations are obtained with the hybrid algorithm for $\delta_r = 0.2$ in Figs. 7 and 8. Fig. 8 also displays A and R from one trajectory determined by SSA for the full system. As can be seen, the solution with the hybrid method has the same qualitative behavior.

Fig. 9 shows the time steps taken by the hybrid solver and the error term σ_ϵ. Small time steps are needed in regions where the solution changes rapidly (cf. Fig. 7). Recomputing time steps after a failure to satisfy the error tolerance is necessary occasionally in this example. This incurs an extra cost when this oscillator is simulated with the hybrid algorithm. This is mostly due to the difficulty to satisfy the tolerance on σ_ϵ in (32). In fact, the tolerance was violated in a few time intervals, see Fig. 8 (right). This problem would likely be ameliorated by using one longer QMC sequence and more trajectories M, according to the discussion in

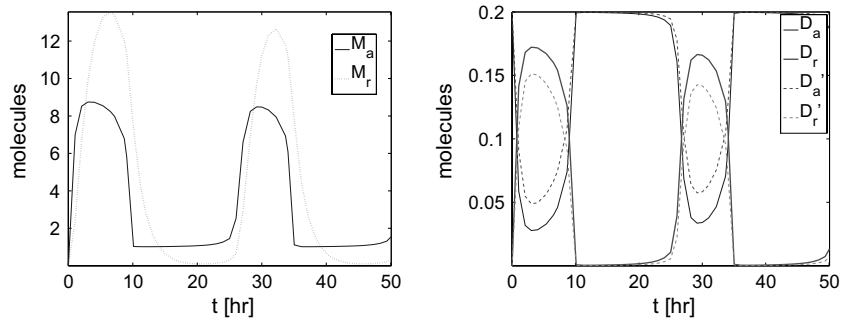

Fig. 7. Time evolution of the mRNA variables M_a and M_r (left) and the gene variables D_a, D'_a, D_r, and D'_r (right) computed with the hybrid method when $\delta_r = 0.2$.

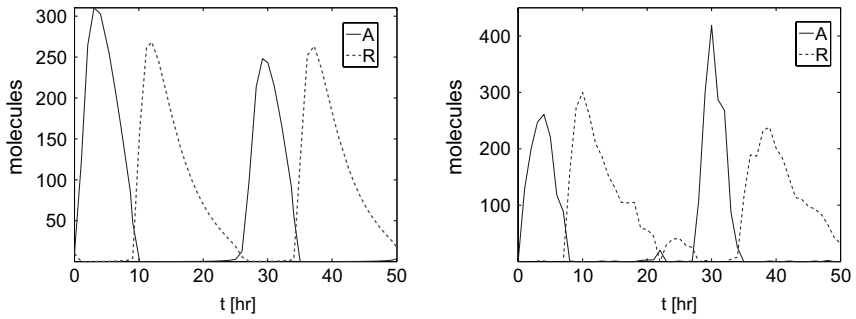

Fig. 8. Time evolution of the expected values of the activator A and the repressor R computed with the hybrid method (left) and one trajectory of SSA for the full system (right) when $\delta_r = 0.2$.

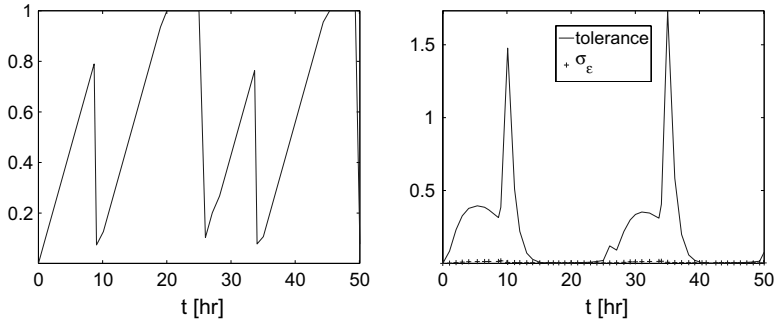

Fig. 9. The time steps taken by the solver (left), the error tolerance (34) and the error σ_ϵ in (32) (right).

the previous section. On the other hand, sometimes a small time discretization error is not the highest priority since the understanding of concepts is often the major goal rather than exact predictions.

The system was simulated with the hybrid solver when $\delta_r = 0.08$. Time integration was done adaptively with relative tolerance 0.05 and absolute tolerance 10^{-3}. The PDF was approximated with $M = 10^4$ trajectories and summation performed with one QMC sequence of $K = 2^{18} \approx 2.62 \times 10^5$ points. Fig. 10 shows the results from the simulation together with results from one trajectory of SSA for the full system. The solution with the hybrid algorithm exhibits oscillations as the fully stochastic model does in contrast to simulation with the reaction rate equations in Fig. 6.

The contribution of the SSA algorithm to the total time required to solve the system is studied in Table 6. Here, the system was simulated with a fixed time step $\Delta t = 0.5$ and one QMC sequence with $K = 2^{18}$. For this system, the stochastic simulations are more demanding than for the previous example in Section 5.1. The partitioning does not permit a substantial speedup and we expect the hybrid solver to approach $g_{\max} = 2$ for a large number of trajectories. Indeed, for $M = 10^6$ more than 95% of the CPU time of the hybrid method is due to SSA.

The summation error of the QMC quadrature determined by (22) for this problem is displayed in Fig. 11. The sum in (16) with p_0 at $t = 25$ s is computed using the Faure sequences for different number of trajectories M and different number of quadrature points K with $Q = 30$ in (21). The convergence rate and the negative slope increase with increasing number of trajectories. This is explained by the fact that a larger M implies a smoother p_0 and a smaller error in the quadrature.

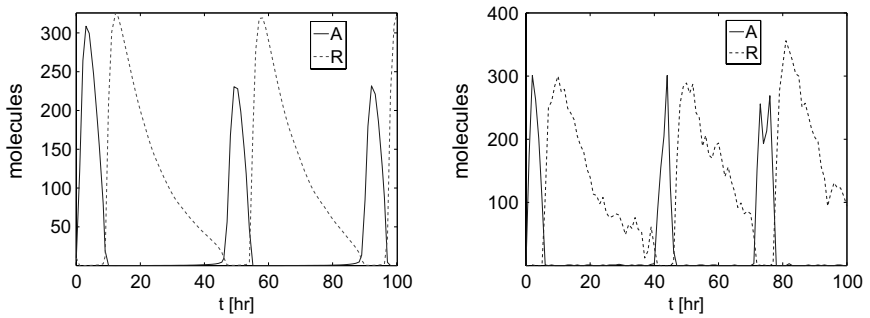

Fig. 10. Time evolution of the expected values of the activator A and the repressor R computed with the hybrid method (left) and one trajectory of SSA for the full system (right) when $\delta_r = 0.08$.

Table 6
Time spent in SSA when the system was simulated with a fixed time step $\Delta t = 0.5$ to the final time 50 s and $Q = 1$, $K = 2^{18}$

Number of trajectories	10^4	10^5	10^6
Time spent in SSA [%]	67.2	92.8	95.8
$T_{\mathrm{SSA}}/T_{\mathrm{hyb}}$	1.33	1.84	1.93

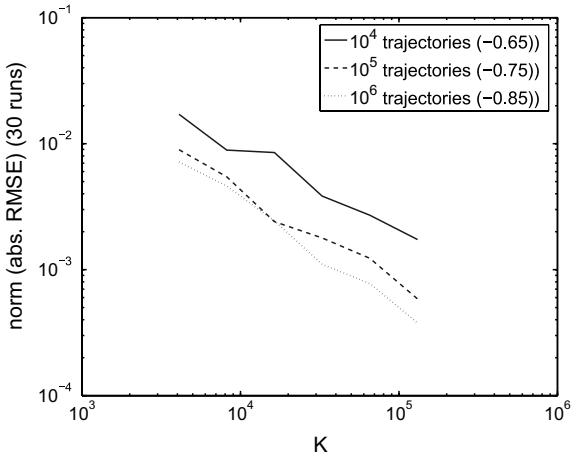

Fig. 11. The absolute summation error for different number of trajectories M and quadrature points K. The number in parenthesis is the inclination of the curve.

5.3. Signaling cascade

As a final example we will consider a model of a mitogen-activated protein kinase (MAPK) signaling cascade [26,27]. These receptor mediated signal transduction pathways are conserved regulatory systems, and the model consists of three sequentially acting kinases and their dephosphatases, phosphorylated forms of the kinases and the intermediate complexes formed in the reactions. The model has 22 species participating in 30 reactions. The output signal is doubly phosphorylated MAKP (MAPKpp), responding to the input signal kinase RAFK.

Suppose we are interested in the stochastic variation of the output signal MAPKpp, while the detailed stochastic information concerning the other components is less important. This is a scenario where the hybrid method will provide a significant speedup compared to the full SSA and more modeling accuracy than the reaction rate equations. No reactions are much faster than the other ones due to the propensities. The system is not too stiff for simulation with SSA, and a simulation with few trajectories gives an estimate of the variances of the species. The species, their initial values and the relative standard deviation at $t = 500$ s are found in Table 7. The computational cost of this calculation is very small compared to an accurate simulation of the system using 10^5–10^6 trajectories.

For modeling purposes, we define three categories based on the values of the relative standard deviation. High variance species are taken to be those with relative standard deviation greater than 0.5. Species with intermediate variance are those in the range 0.1–0.5 and low variance those with quotient smaller than 0.1.

To evaluate the performance of the hybrid method, the state space is divided so that the high variance species (MAPKpp, MAPKpMEKpp and MAPKppMAPKPH) are treated as stochastic variables, $\mathbf{X}$ and $m = 3$ in Section 2, and the other compounds as deterministic variables, $\mathbf{Y}$ and $n = 19$. In this way, 30 reactions for 22 species are reduced to six reactions with a stochastic component, $\rho = 6$ in (8), and 24 differential–summation equations, giving a maximal speedup $g_{\max} = 39$.

Initially the kinases and the dephosphatases are present in higher copy numbers than the others. Simulation with the reaction rate equations also shows that some other species relatively quickly reach higher levels. Reactions involving these species will typically be responsible for most of the time spent in SSA. Hence, they are candidates for macroscopic treatment. Obviously, simulations with higher initial conditions for these species benefit more from the hybrid algorithm.

As a benchmark we use the expected values and standard deviation of MAPKpp at $t = 500$ s, computed by SSA with 10^6 trajectories, yielding $\mu = 0.908$ and $\sigma = 0.983$, or with a 95% confidence interval $\mu = 0.908 \pm 0.002$. Simulation with the reaction rate equations gives $\mu = 0.86$.

The results of simulations with this splitting for increasing number of trajectories are collected in Table 8. One QMC sequence of 2^{14} points was used for the summation. The differential–summation equations were

Table 7
The species divided sorted by the quotient $\sigma(500)/\mu(500)$ based on 1000 trajectories of SSA

Species	Initial value	$\sigma(500)/\mu(500)$
RAFpRAFPH	0	0.02
MAPKPH	180	0.04
MEKPH	120	0.04
RAF	240	0.07
RAFK	60	0.08
MAPK	180	0.08
RAFp	0	0.11
MAPKMEKpp	0	0.16
MEK	120	0.19
MAPKpMAPKPH	0	0.20
RAFRAFK	0	0.21
MEKpMEKPH	0	0.21
MEKRAFp	0	0.21
RAFpRAFPH	0	0.22
MEKp	0	0.28
MAPKp	0	0.32
MEKpRAFp	0	0.34
MEKppMEKPH	0	0.35
MEKpp	0	0.48
MAPKppMAPKPH	0	0.70
MAPKpMEKpp	0	0.70
MAPKpp	0	1.08

Table 8
Speedup factors, standard deviation and expected value for MAPKpp computed with the hybrid algorithm

M	10^4	10^5	10^6
σ	0.94	0.94	0.94
μ	0.89 ± 0.02	0.875 ± 0.006	0.884 ± 0.002
T_{SSA}/T_{hyb}	6.6	13.03	11.83

The expected value μ is reported with a 95% confidence interval.

integrated adaptively with a relative tolerance 0.05 and an absolute tolerance 10^{-4} with a maximal time step of 0.5. The hybrid method produces results between the values computed with SSA for the full system and the deterministic mean value. Also, we obtain a large speedup. However, here the factor decreases for the higher values of M, see the discussion in Section 4.5.

The isolines for the marginal distribution of MAPKpp and MAPKpMEKpp using SSA for the full system and the hybrid solution are compared in Fig. 12, and the probability distribution of MAPKpp is found in

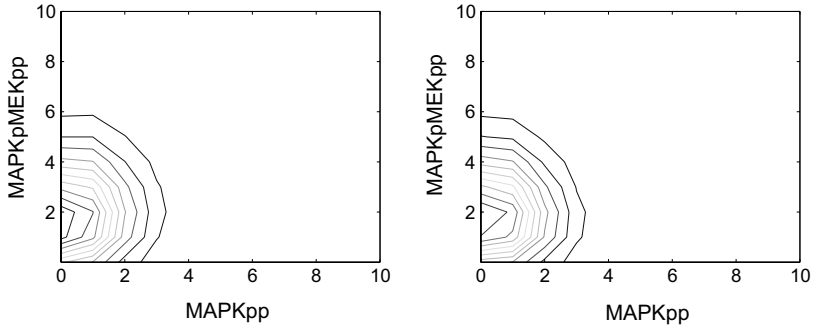

Fig. 12. Isolines of the marginal probability distribution for MAPKpp and MAPKpMEKpp with SSA (left) and the hybrid method (right).

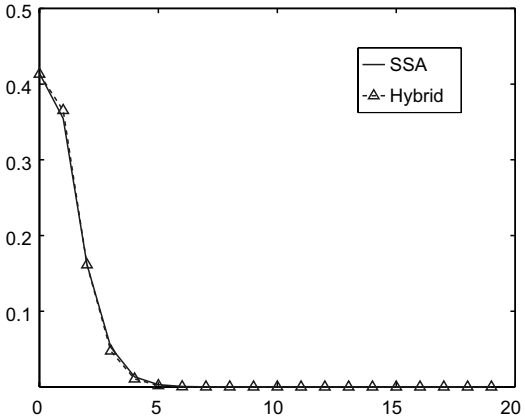

Fig. 13. Marginal probability distribution of MAPKpp computed with SSA and with the hybrid method.

Fig. 13. In both cases, the results were computed using $M = 10^6$ trajectories. It is evident in the figures that the hybrid solver is able to capture the stochastic properties of the activated MAPK species.

This system is still fairly small. For example, the same model extended to include scaffolds [27] has 89 variables participating in 300 chemical reactions. For a system of this size, if a good separation of the variables can be chosen, then the hybrid solver is superior compared to SSA.

6. Conclusions

We have shown that the proposed hybrid method is able to capture important features of the fully mesoscopic models, while keeping the number of stochastically treated variables at a manageable level.

For one of the model systems considered, the reactions are too few and the splitting is done in such a way that an improvement in execution time over the full SSA algorithm cannot be obtained. However, for systems where the partitioning results in a sufficient reduction of the number of reactions and/or an elimination of fast scales from SSA, a considerable speedup compared to a fully stochastic simulation is reached. Also, the hybrid method can be viewed upon as a way of improving the macroscopic model by introducing stochasticity in some components. With this viewpoint, the hybrid solver is more computationally demanding than the ODE models, but gives more realistic results at a comparatively low additional cost.

The major bottleneck in the time stepping scheme of the hybrid solver is, apart from the SSA simulation, the evaluation of the probability distribution function $p_0(\mathbf{x}, t)$. The number of evaluation points is determined by the performance of the summation algorithm. It is therefore crucial to choose a scheme that gives a small error with few evaluated quadrature points. Here, the sum is computed with a Quasi-Monte Carlo method.

The partitioning of the systems based on the variance of the species is often more intuitive than e.g. on fast and slow reactions. The first and second moments are estimated in the last example providing a splitting of the system in a low-dimensional stochastic part and many macroscopic variables. The computing time is reduced by an order of magnitude.

Acknowledgement

This work has been supported by the Swedish Foundation for Strategic Research.

References

[1] N. Barkai, S. Leibler, Circadian clocks limited by noise, Nature 403 (2000) 267–268.
[2] K. Burrage, M. Hegland, S. MacNamara, R.B. Sidje, A Krylov-based finite state projection algorithm for solving the chemical master equation arising in the discrete modeling of biological systems, in: A.N. Langville, W.J. Stewart (Eds.), Proceedings of the 150th Markov Anniversary Meeting, Boson Books, 2006, pp. 21–38.

[3] R.E. Caflisch, Monte Carlo and quasi-Monte Carlo methods, Acta Numer. (1998) 1–49.
[4] Y. Cao, D. Gillespie, L. Petzold, Multiscale stochastic simulation algorithm with stochastic partial equilibrium assumption for chemically reacting systems, J. Comput. Phys. 206 (2005) 395–411.
[5] Y. Cao, D.T. Gillespie, L.R. Petzold, The slow-scale stochastic simulation algorithm, J. Chem. Phys. 122 (2005) 014116.
[6] T.H. Cormen, C.E. Leiserson, R.L. Rivest, C. Stein, Introduction to Algorithms, second ed., The MIT Press, Cambridge, MA, 2001.
[7] W. E, D. Liu, E. Vanden-Eijnden, Nested stochastic simulation algorithm for chemical kinetic systems with disparate rates, J. Chem. Phys. 123 (2005) 194107.
[8] W. E, D. Liu, E. Vanden-Eijnden, Nested stochastic simulation algorithm for chemical kinetic systems with multiple time scales, J. Comput. Phys. 221 (2007) 158–180.
[9] J. Elf, J. Paulsson, O.G. Berg, M. Ehrenberg, Near-critical phenomena in intracellular metabolite pools, Biophys. J. 84 (2003) 154–170.
[10] S. Engblom, Computing the moments of high dimensional solutions of the master equation, Appl. Math. Comput. 180 (2006) 498–515.
[11] R. Erban, I.G. Kevrekidis, D. Adalsteinsson, T.C. Elston, Gene regulatory networks: a coarse-grained equation-free approach to multiscale computation, J. Chem. Phys. 124 (2006) 084106.
[12] H. Faure, Discrépance de suites associées à un système de numération (en dimension s), Acta Aritm. 41 (1982) 337–351.
[13] L. Ferm, P. Lötstedt, Numerical method for coupling the macro and meso scales in stochastic chemical kinetics, Technical Report 2006-001, Department of Information Technology, Uppsala University, Uppsala, Sweden, 2006. Available at <http://www.it.uu.se/research/publications/reports/2006-001/>.
[14] L. Ferm, P. Lötstedt, P. Sjöberg, Conservative solution of the Fokker–Planck equation for stochastic chemical reactions, BIT 46 (2006) S61–S83.
[15] C.W. Gardiner, Handbook of Stochastic Methods, second ed., Springer, Berlin, 2002.
[16] D.T. Gillespie, A general method for numerically simulating the stochastic time evolution of coupled chemical reactions, J. Comput. Phys. 22 (1976) 403–434.
[17] D.T. Gillespie, Approximate accelerated stochastic simulation of chemically reacting systems, J. Chem. Phys. 115 (2001) 1716–1733.
[18] P. Glasserman, Monte Carlo Methods in Financial Engineering, Springer, New York, 2004.
[19] E. Hairer, S.P. Nørsett, G. Wanner, Solving Ordinary Differential Equations I, Nonstiff Problems, second ed., Springer-Verlag, Berlin, 1993.
[20] E. Haseltine, J. Rawlings, Approximate simulation of coupled fast and slow reactions for stochastic chemical kinetics, J. Chem. Phys. 117 (2002) 6959–6969.
[21] M. Hegland, C. Burden, L. Santoso, S. MacNamara, H. Booth, A solver for the stochastic master equation applied to gene regulatory networks, J. Comput. Appl. Math. 205 (2007) 708–724.
[22] H.S. Hong, F.J. Hickernell, Algorithm 823: Implementing scrambled digital sequences, ACM Trans. Math. Softw. 29 (2003) 95–109.
[23] M. Hucka, A. Finney, A. Sauro, H.M. Bolouri, H. Doyle, J.C. Kitano, The systems biology markup language (SBML): a medium for representation and exchange of biochemical network models, Bioinformatics 19 (2003) 524–531.
[24] N.G. van Kampen, Stochastic Processes in Physics and Chemistry, North-Holland, Amsterdam, 1992.
[25] T.R. Kiehl, R.M. Mattheyses, M.K. Simmons, Hybrid simulation of cellular behaviour, Bioinformatics 20 (2004) 316–322.
[26] A. Levchenko, J. Bruck, P.W. Sternberg, Scaffold proteins may biphasically affect the levels of mitogen-activated protein kinase signaling and reduce its threshold properties, Proc. Natl. Acad. Sci. USA 97 (2000) 5818–5823.
[27] H. Lodish, A. Berk, P. Matsudaira, C.A. Kaiser, M. Krieger, M.P. Scott, S.L. Zipursky, J. Darnell, Molecular Cell Biology, fifth ed., Freeman, New York, 2004.
[28] P. Lötstedt, L. Ferm, Dimensional reduction of the Fokker–Planck equation for stochastic chemical reactions, Multiscale Meth. Simul. 5 (2006) 593–614.
[29] P. Lötstedt, S. Söderberg, A. Ramage, L. Hemmingsson-Fränden, Implicit solution of hyperbolic equations with space–time adaptivity, BIT 42 (2002) 134–158.
[30] S. Mac, K. Burrage, R.B. Sidje, Multiscale modeling of chemical kinetics via the master equation, Multiscale Meth. Simul., in press.
[31] H.H. McAdams, A. Arkin, Stochastic mechanisms in gene expression, Proc. Natl. Acad. Sci. USA 94 (1997) 814–819.
[32] B. Moskowitz, R.E. Caflisch, Smoothness and dimension reduction in Quasi-Monte Carlo Methods, Math. Comput. Modell. 23 (1996) 37–54.
[33] A.B. Owen, Monte Carlo variance of scrambled net quadrature, SIAM J. Numer. Anal. 34 (1997) 1884–1910.
[34] J. Paulsson, O.G. Berg, M. Ehrenberg, Stochastic focusing: fluctuation-enhanced sensitivity of intracellular regulation, Proc. Natl. Acad. Sci. USA 97 (2000) 7148–7153.
[35] J. Puchalka, A.M. Kierzek, Bridging the gap between stochastic and deterministic regimes in the kinetic simulations of the biochemical reaction networks, Biophys. J. 86 (2004) 1357–1372.
[36] C.V. Rao, A.P. Arkin, Stochastic chemical kinetics and the quasi-steady-state assumption: application to the Gillespie algorithm, J. Chem. Phys. 118 (2003) 4999–5010.
[37] C.V. Rao, D.M. Wolf, A.P. Arkin, Control exploitation and tolerance of intracellular noise, Nature 420 (2002) 231–237.
[38] H. Salis, Y. Kaznessis, Accurate hybrid stochastic simulation of a system of coupled chemical or biochemical reactions, J. Chem. Phys. 122 (2005) 054103.
[39] A. Samant, D.G. Vlachos, Overcoming stiffness in stochastic simulation stemming from partial equilibrium: a multiscale Monte Carlo algorithm, J. Chem. Phys. 123 (2005) 144114.
[40] L. Santoso, H.S. Booth, C.J. Burden, M. Hegland, A stochastic model of gene switches, ANZIAM J. 46 (2005) C530–C543.

[41] P. Sjöberg, P. Lötstedt, J. Elf, Fokker–Planck approximation of the master equation in molecular biology, Technical Report 2005-044, Department of Information Technology, Uppsala University, Uppsala, Sweden, Comput. Vis. Sci., in press, doi:10.1007/s00791-006-0045-6.
[42] M. Thattai, A. van Oudenaarden, Intrinsic noise in gene regulatory networks, Proc. Nat. Acad. Sci. USA 98 (2001) 8614–8619.
[43] J.M.G. Vilar, H.Y. Kueh, N. Barkai, S. Leibler, Mechanisms of noise-resistance in genetic oscillators, Proc. Nat. Acad. Sci. USA 99 (2002) 5988–5992.