

Hybrid Numerical Solution of the Chemical Master Equation

Thomas A. Henzinger

Institute of Science and Technology, Austria

Linar Mikeev

Saarland University, Germany

Maria Mateescu

Ecole Polytechnique Fédérale de Lausanne, Switzerland

Verena Wolf

Saarland University, Germany

ABSTRACT

We present a numerical approximation technique for the analysis of continuous-time Markov chains that describe networks of biochemical reactions and play an important role in the stochastic modeling of biological systems. Our approach is based on the construction of a stochastic hybrid model in which certain discrete random variables of the original Markov chain are approximated by continuous deterministic variables. We compute the solution of the stochastic hybrid model using a numerical algorithm that discretizes time and in each step performs a mutual update of the transient probability distribution of the discrete stochastic variables and the values of the continuous deterministic variables. We implemented the algorithm and we demonstrate its usefulness and efficiency on several case studies from systems biology.

Keywords

Markov process, biochemical reaction network, chemical master equation, stochastic hybrid model

1. INTRODUCTION

A common dynamical model in systems biology is a system of ordinary differential equations (ODEs) that describes the time evolution of the concentrations of certain proteins in a biological compartment. This macroscopic model is based on the theory of chemical kinetics and assumes that the concentrations of chemical species in a well-stirred system change deterministically and continuously in time. It provides an appropriate description of a chemically reacting system as long as the numbers of molecules of the chemical species are large. However, in living cells the chemical populations can be low (e.g. a single DNA molecule, tens or a few hundreds of RNA or protein molecules). In this case the underlying assumptions of the ODE approach are violated and a more detailed model is necessary, which takes into account the inherently discrete and stochastic nature of chemical reactions [24, 30, 8, 27, 34]. The theory of stochastic chemical kinetics provides an appropriate description by means of

a discrete-state Markov process, that is, a continuous-time Markov chain (CTMC) that represents the chemical populations as random variables [9, 10]. If n is the number of different types of molecules, then we describe the state of the system at a certain time instant by an n -dimensional random vector whose i -th entry represents the number of molecules of type i . In the thermodynamic limit (when the number of molecules and the volume of the system approach infinity) the Markov model and the macroscopic ODE description are equal [21]. Therefore, the ODE approach can be used to approximate the CTMC only if all populations are large.

The evolution of the CTMC is given by a system of linear ordinary differential equations, known as the *chemical master equation* (CME). A single equation in the CME describes the time derivative of the probability of a certain state at all times $t \geq 0$. Thus, the solution of the CME is the probability distribution over all states of the CTMC at a particular time t , that is, the transient state probabilities at time t . The solution of the CME can then be used to derive measures of interest such as the distribution of switching delays [23], the distribution of the time of DNA replication initiation at different origins [26], or the distribution of gene expression products [35]. Moreover, many parameter estimation methods require the computation of the posterior distribution because means and variances do not provide enough information to calibrate parameters [15].

The more detailed description of chemical reactions using a CTMC comes at a price of significantly increased computational complexity because the underlying state space is usually very large or even infinite. Therefore, Monte Carlo simulation is in widespread use, because it allows to generate random trajectories of the model while requiring only little memory. Estimates of the measures of interest can be derived once the number of trajectories is large enough to achieve the desired statistical accuracy. However, the main drawback of simulative solution techniques is that a large number of trajectories is necessary to obtain reliable results. For instance, in order to halve the confidence interval of an estimate, four times more trajectories have to be generated. Consequently, often stochastic simulation is only feasible with a very low level of confidence in the accuracy of the results.

Recently, efficient numerical algorithms have been developed to compute an approximation of the CME [16, 25, 3, 5, 6, 13,

32, 7, 18]. Many of them are based on the idea of restricting the analysis of the model during a certain time interval to a subset of states that have “significant” probability. While some of these methods rely on an a priori estimation of the geometric bounds of the significant subset [16, 25, 3], others are based on a conversion to discrete time and they decide dynamically which states to consider at a certain time step [5, 6, 32].

If the system under consideration contains large populations, then the numerical algorithms mentioned above perform poorly. The reason is that the random variables that represent large populations have a large variance. Thus, a large number of states have a significant probability, which renders the numerical approximation of the distribution computationally expensive or infeasible.

In this paper we use a stochastic hybrid approach to efficiently approximate the solution of systems containing both small and large populations. More precisely, we maintain the discrete stochastic representation for small populations, but at the same time we exploit the small relative variance of large populations and represent them by continuous deterministic variables. Since population sizes change over time we decide dynamically (“on-the-fly”) whether we represent a population by a continuous deterministic variable or keep the discrete stochastic representation. Our criterion for changing from a discrete to a continuous treatment of a variable and vice versa is based on a population threshold.

For the solution of the stochastic hybrid model, we propose a numerical approximation method that discretizes time and performs a mutual update of the distributions of the discrete stochastic variables and the values of the continuous deterministic variables. Hence, we compute the solution of a CME with a reduced dimension as well as the solution of a system of (non-linear) ordinary differential equations. The former describes the distribution of the discrete stochastic variables and the latter the values of the continuous deterministic variables, and the two descriptions depend on each other. Assume, for instance, that a system has two chemical species. The two population sizes at time t are represented by the random variables $X(t)$ and $Y(t)$, where $X(t)$ is large and $Y(t)$ is small. Then, we consider for $Y(t)$ all events $Y(t) = y$ that have significant probability, i.e., $Pr(Y(t) = y)$ is greater than a certain threshold. For $X(t)$ we consider the conditional expectations $E[X(t) | Y(t) = y]$ and assume that they change continuously and deterministically in time. We iterate over small time steps $h > 0$ and, given the distribution for $Y(t)$ and the values $E[X(t) | Y(t) = y]$, we compute the distribution of $Y(t+h)$ and the values $E[X(t+h) | Y(t+h) = y]$. Again, we restrict our computation to those values of y that have significant probability.

To demonstrate the effectiveness of our approach, we have implemented the algorithm and applied it successfully to several examples from systems biology. Our most complex example has 6 different chemical species and 10 reactions. We compare our results with our earlier purely discrete stochastic approach and with the purely continuous deterministic approach in terms of running times and accuracy.

Related Work. Different hybrid approaches have been proposed in the literature [12, 31, 29]. As opposed to our approach, they focus on Monte Carlo simulation and consider the problem of multiple time scales. They do not use deterministic variables but try to reduce the computational complexity of generating a trajectory of the model by approximating the number of reactions during a certain time step. The closest work to ours is the hybrid approach proposed by Hellander and Lötstedt [14]. They approximate large populations by normally distributed random variables with a small variance and use Monte Carlo simulation to statistically estimate the probability distribution of the remaining populations with small sizes. They consider a single ODE to approximate the expected sizes of the large populations. As opposed to that, here we consider a set of ODEs to approximate the expected sizes of the large populations *conditioned* on the small populations. This allows us to track the dependencies between the different populations more accurately. Moreover, instead of a statistical estimation of probabilities, we provide a direct numerical method to solve the stochastic hybrid model. The direct numerical method that we use for the computation of the probability distributions of the stochastic variables has shown to be superior to Monte Carlo simulation [5]. Another difference is that the method in [14] does not allow a dynamic switching between stochastic and deterministic treatment of variables.

Finally, our approach is related to the stochastic hybrid models considered in [4, 2] and to fluid stochastic Petri nets [17]. These approaches differ from our approach in that they use probability distributions for the different values a continuous variable can take. In our setting, at a fixed point in time we only consider the conditional expectations of the continuous variables, which is based on the assumption that the respective populations are large and their relative variance is small. This allows us to provide an efficient numerical approximation algorithm that can be applied to systems with large state spaces. The stochastic hybrid models in [4, 2, 17] cannot be solved numerically except in the case of small state spaces.

2. DISCRETE-STATE STOCHASTIC MODEL

According to Gillespie’s theory of stochastic chemical kinetics, a well-stirred mixture of n molecular species in a volume with fixed size and fixed temperature can be represented as a continuous-time Markov chain $(\mathbf{X}(t), t \geq 0)$ [9]. The random vector $\mathbf{X}(t) = (X_1(t), \dots, X_n(t))$ describes the chemical populations at time t , i.e., $X_i(t)$ is the number of molecules of type $i \in \{1, \dots, n\}$. Thus, the state space of \mathbf{X} is $\mathbb{Z}_+^n = \{0, 1, \dots\}^n$. The state changes of \mathbf{X} are triggered by the occurrences of chemical reactions, which come in m different types. For $j \in \{1, \dots, m\}$ let $\mathbf{u}_j \in \mathbb{Z}^n$ be the *change vector* of the j -th reaction type, that is, $\mathbf{u}_j = \mathbf{u}_j^- + \mathbf{u}_j^+$ where \mathbf{u}_j^- contains only non-positive entries that specify how many molecules of each species are consumed (*reactants*) if an instance of the reaction occurs and vector \mathbf{u}_j^+ contains only non-negative entries that specify how many molecules of each species are produced (*products*). Thus, if $\mathbf{X}(t) = \mathbf{x}$ for some $\mathbf{x} \in \mathbb{Z}_+^n$ with $\mathbf{x} + \mathbf{u}_j^-$ being non-negative, then $\mathbf{X}(t + dt) = \mathbf{x} + \mathbf{u}_j$ is the state of the system after the occurrence of the j -th reaction within the infinitesimal time




Figure 1: Illustration of the exclusive switch in Ex. 1 (picture is adapted from [22]). The stochastic hybrid model with only three discrete stochastic states and two differential equations per state.

interval $[t, t + dt)$.

As rigorously derived by Gillespie [10], each reaction type has an associated *propensity function*, denoted by $\alpha_1, \dots, \alpha_m$, which is such that $\alpha_j(\mathbf{x}) \cdot dt$ is the probability that, given $\mathbf{X}(t) = \mathbf{x}$, one instance of the j -th reaction occurs within $[t, t + dt)$. The value $\alpha_j(\mathbf{x})$ is proportional to the number of distinct reactant combinations in state \mathbf{x} . More precisely, if $\mathbf{x} = (x_1, \dots, x_n)$ is a state for which $\mathbf{x} + \mathbf{u}_j^-$ is nonnegative then

$$\alpha_j(\mathbf{x}) = \begin{cases} c_j & \text{if } \mathbf{u}_j^- = (0, \dots, 0), \\ c_j \cdot x_i & \text{if } \mathbf{u}_j^- = -\mathbf{e}_i, \\ c_j \cdot x_i \cdot x_\ell & \text{if } \mathbf{u}_j^- = -\mathbf{e}_i - \mathbf{e}_\ell, \\ c_j \cdot \binom{x_i}{2} = c_j \cdot \frac{x_i \cdot (x_i - 1)}{2} & \text{if } \mathbf{u}_j^- = -2 \cdot \mathbf{e}_i, \end{cases} \quad (1)$$

where $i \neq \ell$, $c_j > 0$ is a constant, and \mathbf{e}_i is the vector with the i -th entry 1 and all other entries 0. We set $\alpha_j(\mathbf{x}) = 0$ whenever the vector $\mathbf{x} + \mathbf{u}_j^-$ contains negative entries, that is, when not enough reactant molecules are available. The constant c_j refers to the probability that a randomly selected pair of reactants collides and undergoes the j -th chemical reaction. Thus, if N is the volume (in liters) times Avogadro's number, then c_j

- scales inversely with N in the case of two reactants,
- is independent of N in the case of a single reactant,
- is proportional to N in the case of no reactants.

Since reactions of higher order (requiring more than two reactants) are usually the result of several successive lower order reactions, we do not consider the case of more than two reactants.

EXAMPLE 1. We consider a gene regulatory network, called the *exclusive switch* [22]. It consists of two genes with a common promotor region. Each of the two gene products P_1 and P_2 inhibits the expression of the other product if a molecule is bound to the promotor region. More precisely, if the promotor region is free, molecules of both types P_1 and P_2 are produced. If a molecule of type P_1 (P_2) is bound to the promotor region, only molecules of type P_1 (P_2) are produced, respectively. We illustrate the system in Fig. 1. The system has five chemical species of which two have an infinite range, namely P_1 and P_2 . If $\mathbf{x} = (x_1, \dots, x_5)$ is the current state, then the first two entries represent the populations of P_1 and P_2 , respectively. The entry x_3 denotes the number of unbound DNA molecules which is either zero or one. The entry x_4 (x_5) is one of a molecule of type P_1 (P_2) is bound to the promotor region and zero otherwise.

The chemical reactions are as follows. Let $j \in \{1, 2\}$.

- We describe production of P_j by $\text{DNA} \rightarrow \text{DNA} + P_j$. Thus, $\mathbf{u}_j = \mathbf{e}_j$ and $\alpha_j(\mathbf{x}) = c_j \cdot x_3$.
- We describe degradation of P_j by $P_j \rightarrow \emptyset$ with $\mathbf{u}_{j+2} = -\mathbf{e}_j$ and $\alpha_{j+2}(\mathbf{x}) = c_{j+2} \cdot x_j$.
- We model the binding of P_j to the promotor by $\text{DNA} + P_j \rightarrow \text{DNA}.P_j$ with $\mathbf{u}_{j+4} = -\mathbf{e}_j - \mathbf{e}_3 + \mathbf{e}_{j+3}$ and $\alpha_{j+4}(\mathbf{x}) = c_{j+4} \cdot x_j \cdot x_3$.
- For unbinding of P_j we use $\text{DNA}.P_j \rightarrow \text{DNA} + P_j$ with $\mathbf{u}_{j+6}(\mathbf{x}) = \mathbf{e}_j + \mathbf{e}_3 - \mathbf{e}_{j+3}$ and $\alpha_{j+6}(\mathbf{x}) = c_{j+6} \cdot x_{j+3}$.
- Finally, we have production of P_j if a molecule of type P_j is bound to the promotor, i.e., $\text{DNA}.P_j \rightarrow \text{DNA}.P_j + P_j$ with $\mathbf{u}_{j+8}(\mathbf{x}) = \mathbf{e}_j$ and $\alpha_{j+8}(\mathbf{x}) = c_{j+8} \cdot x_{j+3}$.

Depending on the chosen parameters, the probability distribution of the exclusive switch is bistable, i.e. most of the probability mass concentrates on two distinct regions in the state space. In particular, if binding to the promotor is likely, then these two regions correspond to the two configurations where either the production of P_1 or the production of P_2 is inhibited.

The Chemical Master Equation. For $\mathbf{x} \in \mathbb{Z}_+^n$ and $t \geq 0$, let $p(\mathbf{x}, t)$ denote the probability that the current population vector is \mathbf{x} , i.e., $p(\mathbf{x}, t) = \text{Pr}(\mathbf{X}(t) = \mathbf{x})$. Let $\mathbf{p}(t)$ be the row vector with entries $p(\mathbf{x}, t)$. Given $\mathbf{u}_1^-, \dots, \mathbf{u}_m^-, \mathbf{u}_1^+, \dots, \mathbf{u}_m^+$, $\alpha_1, \dots, \alpha_m$, and some initial distribution $\mathbf{p}(0)$, the Markov chain \mathbf{X} is uniquely specified if the propensity functions are of the form in Eq. (1). The evolution of \mathbf{X} is given by the chemical master equation (CME), which equates the change $\frac{d}{dt}p(\mathbf{x}, t)$ of the probability in state \mathbf{x} and the sum over all reactions of the ‘‘inflow’’ $\alpha_j(\mathbf{x} - \mathbf{u}_j^-) \cdot p(\mathbf{x} - \mathbf{u}_j^-, t)$ and ‘‘outflow’’ $\alpha_j(\mathbf{x}) \cdot p(\mathbf{x}, t)$ of probability [20]. Thus,

$$\frac{d}{dt}p(\mathbf{x}, t) = \sum_{j=1}^m (\alpha_j(\mathbf{x} - \mathbf{u}_j^-) \cdot p(\mathbf{x} - \mathbf{u}_j^-, t) - \alpha_j(\mathbf{x}) \cdot p(\mathbf{x}, t)). \quad (2)$$

Since the CME is linear it can be written as $\frac{d}{dt}\mathbf{p}(t) = \mathbf{p}(t) \cdot Q$, where Q is the generator matrix of \mathbf{X} with $Q(\mathbf{x}, \mathbf{x} + \mathbf{u}_j^+) = \alpha_j(\mathbf{x})$ and $Q(\mathbf{x}, \mathbf{x}) = -\sum_{j=1}^m \alpha_j(\mathbf{x})$. If Q is bounded, then Eq (2) has the general solution

$$\mathbf{p}(t) = \mathbf{p}(0) \cdot e^{Qt}, \quad (3)$$

where the matrix exponential is defined as $e^{Qt} = \sum_{i=0}^{\infty} \frac{(Qt)^i}{i!}$. If the state space is infinite, then we can only compute approximations of $\mathbf{p}(t)$ and even if Q is finite, the size of the matrix Q is often large because it grows exponentially with the number of state variables. Moreover, even if Q is sparse, as it usually is because the number of reaction types is small compared to the number of states, standard numerical solution techniques for systems of first-order linear equations of the form of Eq. (2), such as uniformization [19], approximations in the Krylov subspace [28], or numerical integration [33], are infeasible. The reason is that the number of nonzero entries in Q often exceeds the available memory capacity for systems of realistic size. If the populations of all species remain small (at most a few hundreds) then the solution of the CME can be efficiently approximated using projection methods [16, 25, 3] or fast uniformization methods [5, 6, 32]. The idea of these methods is to avoid an exhaustive state space exploration and, depending on a certain time interval, restrict the analysis of the system to a

subset of states.

Fast Solution of the Discrete Stochastic Model. Here, we present a method similar to our previous work [6] that efficiently approximates the solution of the CME if the chemical populations remain small. We use it in Section 4 to solve the discrete part of the stochastic hybrid model.

The algorithm, called fast RK4, is based on the numerical integration of Eq. (2) using an explicit fourth-order Runge-Kutta method. The main idea is to integrate only those differential equations in Eq. (2) that correspond to states with “significant probability”. This reduces the computational effort significantly since in each iteration step only a comparatively small subset of states is considered. We dynamically decide which states to drop/add based on a fixed probability threshold $\delta > 0$. Due to the regular structure of the Markov model the approximation error of the algorithm remains small since probability mass is usually concentrated at certain parts of the state space. The farther away a state is from such a “significant set” the smaller is its probability. Thus, the total error of the approximation remains small. Unless otherwise specified, in our experiments we fix δ to 10^{-14} . This value that has been shown to lead to accurate approximations [6].

The standard explicit fourth-order Runge-Kutta method applied to Eq. (2) yields the iteration step [33]

$$\mathbf{p}(t+h) = \mathbf{p}(t) + h \cdot (\mathbf{k}_1 + 2 \cdot \mathbf{k}_2 + 2 \cdot \mathbf{k}_3 + \mathbf{k}_4)/6, \quad (4)$$

where $h > 0$ is the time step of the method and the vectors $\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3, \mathbf{k}_4$ are given by

$$\begin{aligned} \mathbf{k}_1 &= \mathbf{p}(t) \cdot Q, & \mathbf{k}_3 &= (\mathbf{p}(t) + h \cdot \frac{\mathbf{k}_2}{2}) \cdot Q, \\ \mathbf{k}_2 &= (\mathbf{p}(t) + h \cdot \frac{\mathbf{k}_1}{2}) \cdot Q, & \mathbf{k}_4 &= (\mathbf{p}(t) + h \cdot \mathbf{k}_3) \cdot Q. \end{aligned} \quad (5)$$

Note that the entries $k_1(\mathbf{x}), \dots, k_4(\mathbf{x})$ of state \mathbf{x} in the vectors $\mathbf{k}_1, \dots, \mathbf{k}_4$ are given by

$$\begin{aligned} k_1(\mathbf{x}) &= \sum_{j=1}^m (\alpha_j(\mathbf{x} - \mathbf{u}_j) \cdot p(\mathbf{x} - \mathbf{u}_j, t) - \alpha_j(\mathbf{x}) \cdot p(\mathbf{x}, t)), \\ k_{i+1}(\mathbf{x}) &= \sum_{j=1}^m (\alpha_j(\mathbf{x} - \mathbf{u}_j) \cdot (p(\mathbf{x} - \mathbf{u}_j, t) + h \cdot k_i(\mathbf{x} - \mathbf{u}_j)/2) \\ &\quad - \alpha_j(\mathbf{x}) \cdot (p(\mathbf{x}, t) + h \cdot k_i(\mathbf{x})/2)) \text{ for } i \in \{1, 2\}, \\ k_4(\mathbf{x}) &= \sum_{j=1}^m (\alpha_j(\mathbf{x} - \mathbf{u}_j) \cdot (p(\mathbf{x} - \mathbf{u}_j, t) + h \cdot k_3(\mathbf{x} - \mathbf{u}_j)) \\ &\quad - \alpha_j(\mathbf{x}) \cdot (p(\mathbf{x}, t) + h \cdot k_3(\mathbf{x}))). \end{aligned} \quad (6)$$

In order to avoid the explicit construction of Q and in order to work with a dynamic set Sig of significant states that changes in each step, we use for a state \mathbf{x} a data structure with the following components:

- a field $\mathbf{x}.prob$ for the current probability of \mathbf{x} ,
- fields $\mathbf{x}.k_1, \dots, \mathbf{x}.k_4$ for the four terms in the equation of state \mathbf{x} in the system of Eq. (5),
- for all j with $\mathbf{x} + \mathbf{u}_j \geq \mathbf{0}$ a pointer to the successor state $\mathbf{x} + \mathbf{u}_j$ as well as the rate $\alpha_j(\mathbf{x})$.

We start at time $t = 0$ and initialize the set Sig as the set of all states that have initially a probability greater than δ , i.e. $Sig := \{\mathbf{x} \mid p(\mathbf{x}, 0) > \delta\}$. We perform a step of the iteration in Eq. (4) by traversing the set Sig five times. In the first four rounds we compute $\mathbf{k}_1, \dots, \mathbf{k}_4$ and in the final round

```

1 choose step size  $h$ ;
2 for  $i = 1, 2, 3, 4$  do //traverse  $Sig$  four times
3 //decide which fields from state data structure
4 //are needed for  $k_i$ 
5 switch  $i$ 
6   case  $i = 1$ :  $coeff := 1$ ;  $field := prob$ ;
7   case  $i \in \{2, 3\}$ :  $coeff := h/2$ ;  $field := k_{i-1}$ ;
8   case  $i = 4$ :  $coeff := h$ ;  $field := k_{i-1}$ ;
9  $\mathbf{x}.k_i := \mathbf{x}.k_i$ ;
10 for all  $\mathbf{x} \in Sig$  do
11   for  $j = 1, \dots, m$  with  $\mathbf{x} + \mathbf{u}_j \geq \mathbf{0}$  do
12      $\mathbf{x}.k_i := \mathbf{x}.k_i - coeff \cdot \mathbf{x}.field \cdot \alpha_j(\mathbf{x})$ ;
13     if  $\mathbf{x} + \mathbf{u}_j \notin Sig$  then
14        $Sig := Sig \cup \{\mathbf{x} + \mathbf{u}_j\}$ ;
15      $(\mathbf{x} + \mathbf{u}_j).k_i := (\mathbf{x} + \mathbf{u}_j).k_i + coeff \cdot \mathbf{x}.field \cdot \alpha_j(\mathbf{x})$ ;
16 for all  $\mathbf{x} \in Sig$  do
17    $\mathbf{x}.prob := \mathbf{x}.prob + h \cdot (\mathbf{x}.k_1 + 2 \cdot \mathbf{x}.k_2 + 2 \cdot \mathbf{x}.k_3 + \mathbf{x}.k_4)/6$ ;
18    $\mathbf{x}.k_1 := 0$ ;  $\mathbf{x}.k_2 := 0$ ;  $\mathbf{x}.k_3 := 0$ ;  $\mathbf{x}.k_4 := 0$ ;
19   if  $\mathbf{x}.prob < \delta$  then
20      $Sig := Sig \setminus \{\mathbf{x}\}$ ;

```

Table 1: A single iteration step of the fast RK4 algorithm, which approximates the solution of the CME.

we accumulate the summands. While processing state \mathbf{x} in round i , $i < 5$, for each reaction j , we transfer probability mass from state \mathbf{x} to its successor $\mathbf{x} + \mathbf{u}_j$, by subtracting a term from $k_i(\mathbf{x})$ (see Eq. (6)) and adding the same term to $k_i(\mathbf{x} + \mathbf{u}_j)$. A single iteration step is illustrated in pseudocode in Table 1. In line 20, we ensure that Sig does not contain states with a probability less than δ . We choose the step size h in line 1 as suggested in [33]. In line 2-15 we compute the values $k_1(\mathbf{x}), \dots, k_4(\mathbf{x})$ for all $\mathbf{x} \in Sig$ (see Eq. (5)). The fifth round starts in line 16 and in line 17 the approximation of the probability $p(\mathbf{x}, t+h)$ is calculated. Note that the fields $\mathbf{x}.k_1, \dots, \mathbf{x}.k_4$ are initialized with zero.

Clearly, for the solution of the CME the same ideas as above can be used for many other numerical integration methods. Here, we focus on the explicit RK4 method and do not consider more advanced numerical integration methods to keep our presentation simple. The focus of this paper is not on particular numerical methods to solve differential equations but rather on general strategies for the approximate solution of the stochastic models that we consider. Moreover, we do not use uniformization methods as in previous work since uniformization is inefficient for very small time horizons. But small time steps are necessary for the solution of the hybrid model in order to take into account the dependencies between the stochastic and the deterministic variables.

3. DERIVATION OF THE DETERMINISTIC LIMIT

The numerical approximation presented in the previous section works well as long as only the main part of the probability mass is concentrated on a small subset of the state space. If the system contains large populations then the probability mass distributes on a very large number of states whereas the information content is rather low since we dis-

tinguish, for instance, the cases of having $X_i(t) = 10000$, $X_i(t) = 10001$, etc. In such cases no direct numerical approximation of the CME is possible without resorting to Monte Carlo techniques or discarding the discreteness of the state space. If all populations are large the solution of \mathbf{X} can be accurately approximated by considering the deterministic limit of \mathbf{X} . Here, we shortly recall the basic steps for the derivation of the deterministic limit. For a detailed discussion, we refer to Kurtz [21].

We first define a set of functions β_j such that if N is large (recall that N is the volume times the Avogadro's number) then the propensity functions can be approximated as $\alpha_j(\mathbf{x}) \approx N \cdot \beta_j(\mathbf{z})$, where $\mathbf{z} = (z_1, \dots, z_n) = \mathbf{x} \cdot N^{-1}$ corresponds to the vector of concentrations of chemical species and belongs to \mathbb{R}^n . Recall the dependencies of c_j on the scaling factor N as described at the beginning of Section 2. For constants $k_j > 0$ that are *independent* of N ,

- $c_j = k_j \cdot N$ in the case of no reactants,
- $c_j = k_j$ in the case of a single reactant,
- $c_j = k_j/N$ in the case of two reactants.

From this, it follows that except for the case of bimolecular reactions, we can construct the functions β_j such that $\alpha_j(\mathbf{x}) = N \cdot \beta_j(\mathbf{z})$.

$$\beta_j(\mathbf{z}) = \frac{\alpha_j(\mathbf{x})}{N} = \begin{cases} \frac{c_j}{N} = k_j & \text{if } \mathbf{u}_j^- = (0, \dots, 0), \\ c_j \cdot \frac{x_i}{N} = k_j \cdot z_i & \text{if } \mathbf{u}_j^- = -\mathbf{e}_i, \\ c_j \cdot x_i \cdot \frac{x_\ell}{N} = k_j \cdot z_i \cdot z_\ell & \text{if } \mathbf{u}_j^- = -\mathbf{e}_i - \mathbf{e}_\ell, \end{cases}$$

where $i \neq \ell$. In the case of bimolecular reactants ($\mathbf{u}_j^- = -2 \cdot \mathbf{e}_i$), we use the approximation

$$\begin{aligned} N \cdot \beta_j(\mathbf{z}) &= k_j \cdot N \cdot z_i^2 = k_j \cdot x_i \cdot z_i = (\frac{1}{2} c_j N) \cdot x_i \cdot z_i \\ &= \frac{1}{2} c_j \cdot x_i^2 \approx \frac{1}{2} c_j \cdot x_i (x_i - 1) = \alpha_j(\mathbf{x}), \end{aligned}$$

which is accurate if x_i is large. In order to derive the deterministic limit for the vector $\mathbf{X}(t) = (X_1(t), \dots, X_n(t))$ that describes the chemical populations, we first write $\mathbf{X}(t)$ as

$$\mathbf{X}(t) = \mathbf{X}(0) + \sum_{j=1}^m \mathbf{u}_j \cdot C_j(t),$$

where $\mathbf{X}(0)$ is the initial population vector and $C_j(t)$ denotes the number of occurrences of the j -th reaction until time t . The process $C_j(t)$ is a counting process with intensity $\alpha_j(\mathbf{X}(t))$ and it can be regarded as a Poisson process whose time-dependent intensity changes according to the stochastic process $\mathbf{X}(t)$. Now, recall that a Poisson process $\tilde{Y}(t)$ with time-dependent intensity $\lambda(t)$ can be transformed into a Poisson process $Y(u)$ with constant intensity one, using the simple time transform $u = \int_0^t \lambda(s) ds$, that is, $Y(u) = Y(\int_0^t \lambda(s) ds) = \tilde{Y}(t)$. Similarly, we can describe $C_j(t)$ as a Poisson process with intensity one, i.e.,

$$C_j(t) = Y_j \left(\int_0^t \alpha_j(\mathbf{X}(s)) ds \right),$$

where Y_j are independent Poisson processes with intensity one. Hence, for $i \in \{1, \dots, n\}$

$$X_i(t) = X_i(0) + \sum_{j=1}^m u_{ji} \cdot Y_j \left(\int_0^t \alpha_j(\mathbf{X}(s)) ds \right), \quad (7)$$

where $\mathbf{u}_j = (u_{j1}, \dots, u_{jn})$. The next step is to define $\mathbf{Z}(t) = \mathbf{X}(t) \cdot N^{-1}$, that is, $\mathbf{Z}(t) = (Z_1(t), \dots, Z_n(t))$ contains the

concentrations of the chemical species in moles per liter at time t . Thus,

$$Z_i(t) = Z_i(0) + \sum_{j=1}^m u_{ji} \cdot N^{-1} \cdot Y_j \left(\int_0^t \alpha_j(\mathbf{X}(s)) ds \right), \quad (8)$$

and using the fact that $\alpha_j(\mathbf{x}) \approx N \cdot \beta_j(\mathbf{z})$ yields

$$Z_i(t) \approx Z_i(0) + \sum_{j=1}^m u_{ji} \cdot N^{-1} \cdot Y_j \left(N \cdot \int_0^t \beta_j(\mathbf{Z}(s)) ds \right). \quad (9)$$

By the law of large numbers, the unit Poisson process Y_j will approach $N \cdot u$ at time $N \cdot u$ for large $N \cdot u$. Thus, $Y_j(N \cdot u) \approx N \cdot u$ and hence,

$$Z_i(t) \approx Z_i(0) + \sum_{j=1}^m u_{ji} \cdot \int_0^t \beta_j(\mathbf{Z}(s)) ds. \quad (10)$$

The right-hand side of the above integral equation is the solution $\mathbf{z}(t)$ of the system of ODEs

$$\frac{d}{dt} \mathbf{z}(t) = \sum_{j=1}^m \mathbf{u}_j \cdot \beta_j(\mathbf{z}(t)). \quad (11)$$

As shown by Kurtz [21], in the large volume limit, where the volume and the number of molecules approach infinity (while the concentrations remain constant), $\mathbf{Z}(t) \rightarrow \mathbf{z}(t)$ in probability for finite times t . Note that the chemical concentrations $\mathbf{z}(t)$ evolve continuously and deterministically in time. This continuous deterministic approximation is reasonable if all species have a small relative variance and if they are only weakly correlated. The reason is that only in this case the assumption that $Z_i(t)$ is deterministic is appropriate. Note that for most models this is the case if the population of species i is large since this implies that $E[X_i(t)]$ is large whereas the occurrence of chemical reactions results only in a marginal relative change of the value of $X_i(t)$.

EXAMPLE 1 (CONT.). *The ODEs of the exclusive switch are given by*

$$\begin{aligned} \frac{d}{dt} z_1(t) &= k_1 \cdot z_3(t) - k_3 \cdot z_1(t) - k_5 \cdot z_1(t) \cdot z_3(t) \\ &\quad + k_7 \cdot z_4(t) + k_9 \cdot z_4(t) \\ \frac{d}{dt} z_2(t) &= k_2 \cdot z_3(t) - k_4 \cdot z_2(t) - k_6 \cdot z_2(t) \cdot z_3(t) \\ &\quad + k_8 \cdot z_5(t) + k_{10} \cdot z_5(t) \\ \frac{d}{dt} z_3(t) &= -k_5 \cdot z_1(t) \cdot z_3(t) - k_6 \cdot z_2(t) \cdot z_3(t) \\ &\quad + k_7 \cdot z_4(t) + k_8 \cdot z_5(t) \\ \frac{d}{dt} z_4(t) &= k_5 \cdot z_1(t) \cdot z_3(t) - k_7 \cdot z_4(t) \\ \frac{d}{dt} z_5(t) &= k_6 \cdot z_2(t) \cdot z_3(t) - k_8 \cdot z_5(t) \end{aligned}$$

where $z_1(t), z_2(t), z_3(t), z_4(t), z_5(t)$ denote the respective chemical concentrations. Moreover, $c_j = N^{-1} \cdot k_j$ for $j \in \{5, 6\}$ and $c_j = k_j$ for $j \notin \{5, 6\}$.

In [1], Ball et al. scale only a subset of the populations in order to approximate the behavior of the system if certain populations are large and others are small. Additionally, they take into account the different speeds of the chemical reactions. For a selected number of examples, they give analytical expressions for the distributions in the limit, i.e., when the scaling parameter approaches infinity. In the next section, we will construct a stochastic hybrid model that is equivalent to the one considered in [1] if we scale the continuous components and consider the deterministic limit.

4. STOCHASTIC HYBRID MODEL

A straightforward consequence of the CME is that the time derivative of the populations' expectations are given by

$$\frac{d}{dt}E[\mathbf{X}(t)] = \sum_{j=1}^m \mathbf{u}_j \cdot E[\alpha_j(\mathbf{X}(t))]. \quad (12)$$

If all reactions of the system involve at most one reactant, Eq. (12) can be simplified to

$$\frac{d}{dt}E[\mathbf{X}(t)] = \sum_{j=1}^m \mathbf{u}_j \cdot \alpha_j(E[\mathbf{X}(t)]). \quad (13)$$

because the propensity functions α_j are linear in \mathbf{x} . But in the case of bimolecular reactions, we have either $\alpha_j(\mathbf{x}) = c_j \cdot x_i \cdot x_\ell$ for some i, ℓ with $i \neq \ell$ or $\alpha_j(\mathbf{x}) = c_j \cdot x_i \cdot (x_i - 1)/2$ if the j -th reaction involves two reactants of type i . But this means that

$$E[\alpha_j(\mathbf{X}(t))] = c_j \cdot E[X_i(t) \cdot X_\ell(t)] \quad \text{or}$$

$$E[\alpha_j(\mathbf{X}(t))] = \frac{1}{2} \cdot c_j \cdot E[(X_i(t))^2] - E[X_i(t)],$$

respectively. In both cases new differential equations are necessary to describe the unknown values of $E[X_i(t) \cdot X_\ell(t)]$ and $E[(X_i(t))^2]$. This problem repeats and leads to an infinite system of ODEs. As shown in the sequel, we can, however, exploit Eq. (12) to derive a stochastic hybrid model.

Assume we have a system where certain species have a large population. In that case we approximate them with continuous deterministic variables. The remaining variables are kept discrete stochastic. This is done because it is usually infeasible or at least computationally very costly to solve a purely stochastic model with high populations since in the respective dimensions the number of significant states is large. Therefore, we propose to switch to a hybrid model where the stochastic part does not contain large populations. In this way we can guarantee an efficient approximation of the solution.

Formally, we split $\mathbf{X}(t)$ into small populations $\mathbf{V}(t)$ and large populations $\mathbf{W}(t)$, i.e. $\mathbf{X}(t) = (\mathbf{V}(t), \mathbf{W}(t))$. Let \tilde{n} be the dimension of $\mathbf{V}(t)$ and \hat{n} the dimension of $\mathbf{W}(t)$, i.e. $n = \tilde{n} + \hat{n}$. Moreover, let \tilde{D} and \hat{D} be the set of indices that correspond to the populations in \mathbf{V} and \mathbf{W} , respectively. Thus, $\tilde{D}, \hat{D} \subseteq \{1, \dots, n\}$ and $|\tilde{D}| = \tilde{n}$, $|\hat{D}| = \hat{n}$. We define $\tilde{\mathbf{u}}_j$ and $\hat{\mathbf{u}}_j$ as the components of \mathbf{u}_j that belong to \tilde{D} and \hat{D} , respectively. Under the condition that $\mathbf{V}(t) = \mathbf{v}$ and $\mathbf{W}(t) = \mathbf{w}$, we assume that for an infinitesimal time interval of length dt the evolution of \mathbf{W} is given by the stochastic differential equation

$$\mathbf{W}(t + dt) = \mathbf{W}(t) + \sum_{j=1}^m \hat{\mathbf{u}}_j \cdot \alpha_j(\mathbf{v}, \mathbf{w}) \cdot dt. \quad (14)$$

The evolution of \mathbf{V} remains unchanged, i.e.,

$$Pr(\mathbf{V}(t + dt) = \mathbf{v} + \tilde{\mathbf{u}}_j \mid \mathbf{V}(t) = \mathbf{v}, \mathbf{W}(t) = \mathbf{w}) = \alpha_j(\mathbf{v}, \mathbf{w}) \cdot dt$$

The density function $h(\mathbf{v}, \mathbf{w}, t)$ of the Markov process $\{(\mathbf{V}(t), \mathbf{W}(t)), t \geq 0\}$ can be derived in the same way as done by Horton et al. [17]. Here, for simplicity we consider only the case $\hat{n} = 1$ which means that $\mathbf{w} = w$ is a scalar. The generalization to higher dimensions is straightforward. If $w > 0$ then the following partial differential equation holds for h .

$$\begin{aligned} & \frac{\partial h(\mathbf{v}, w, t)}{\partial t} + \frac{\partial (h(\mathbf{v}, w, t) \cdot \sum_j \hat{\mathbf{u}}_j \cdot \alpha_j(\mathbf{v}, w))}{\partial w} \\ &= \sum_j \alpha_j(\mathbf{v} - \tilde{\mathbf{u}}_j, w) \cdot h(\mathbf{v} - \tilde{\mathbf{u}}_j, w, t) - \sum_j \alpha_j(\mathbf{v}, w) \cdot h(\mathbf{v}, w, t). \end{aligned}$$

If $w = 0$ then we have probability mass $g(\mathbf{v}, w, t)$ in state (\mathbf{v}, w) where

$$\begin{aligned} & \frac{\partial g(\mathbf{v}, w, t)}{\partial t} + h(\mathbf{v}, w, t) \cdot \sum_j \hat{\mathbf{u}}_j \cdot \alpha_j(\mathbf{v}, w) \\ &= \sum_j \alpha_j(\mathbf{v} - \tilde{\mathbf{u}}_j, w) \cdot g(\mathbf{v} - \tilde{\mathbf{u}}_j, w, t) - \sum_j \alpha_j(\mathbf{v}, w) \cdot g(\mathbf{v}, w, t). \end{aligned}$$

As explained in-depth by Horton et al., the above equations express that probability mass must be conserved, i.e. the change of probability mass in a "cell" with boundaries $(\mathbf{v}, w - dw)$ and $(\mathbf{v}, w + dw)$ equals the total mass of probability entering the cell minus the total mass leaving the cell.

In order to exploit the fact that the relative variance of \mathbf{W} is small, we suggest an approximative solution of the stochastic hybrid model given above. The main idea is not to compute the full density h and the mass function g but only the distribution of V as well as the *conditional expectations* $E[\mathbf{W}(t) = w \mid \mathbf{V}(t) = \mathbf{v}]$. Thus, in our numerical procedure the distribution of \mathbf{W} is approximated by the different values $E[\mathbf{W}(t) = \mathbf{w} \mid \mathbf{V}(t) = \mathbf{v}]$, $\mathbf{v} \in \mathbb{N}^{\tilde{n}}$ that are taken by $\mathbf{W}(t)$ with probability $Pr(\mathbf{V}(t) = \mathbf{v})$.

Assume that at time t we have the approximation $\mathbf{p}(t)$ of the discrete stochastic model as described in Section 2, that is, for all states \mathbf{x} that have a probability that is greater than δ we have $p(\mathbf{x}, t) > 0$ and for all other states \mathbf{x} we have $p(\mathbf{x}, t) = 0$. At time t the expectations of one or more populations reached a certain large population threshold. Thus, we switch to a hybrid model where the large populations (index set \tilde{D}) are represented as continuous deterministic variables $\mathbf{W}(t)$ while the small populations (index set \hat{D}) are represented by $\mathbf{V}(t)$. We first compute the vector of conditional expectations

$$\Psi_{\mathbf{v}}(t) := E[\mathbf{W}(t) = \mathbf{w} \mid \mathbf{V}(t) = \mathbf{v}] = \sum_{\mathbf{x}: \tilde{\mathbf{x}} = \mathbf{v}} \mathbf{w} \cdot p(\mathbf{x}, t).$$

Here, $\tilde{\mathbf{x}}$ is the subvector of \mathbf{x} that corresponds to \tilde{D} . We also compute the distribution $\mathbf{p}(t)$ of $\mathbf{V}(t)$ as

$$r(\mathbf{v}, t) := \sum_{\mathbf{x}: \tilde{\mathbf{x}} = \mathbf{v}} p(\mathbf{x}, t).$$

Now, we integrate the system for a small time interval of length $h > 0$. This is done in three steps as described below. We will write $\Psi_{\mathbf{v}}(t')$ for the approximation of $E[\mathbf{W}(t') = \mathbf{w} \mid \mathbf{V}(t') = \mathbf{v}]$. The i -th element of the \hat{n} -dimensional vector $\Psi_{\mathbf{v}}(t')$ is denoted by $\psi_{\mathbf{v}}(i, t')$. The value $r(\mathbf{v}, t')$ denotes the approximation of $Pr(\mathbf{V}(t) = \mathbf{v})$ where $t' \in [t, t + h)$. The vector $\mathbf{r}(t')$ contains the elements $r(\mathbf{v}, t')$.

(1) Update distribution. We first integrate $\mathbf{r}(t)$ for h time units according to a CME with dimension \tilde{n} to approximate the probabilities $Pr(\mathbf{V}(t + h) = \mathbf{v})$ by $r(\mathbf{v}, t + h)$, that is, $\mathbf{r}(t + h)$ is the solution of the system of ODEs

$$\begin{aligned} \frac{dr(\mathbf{v}, t')}{dt'} &= \sum_j \alpha_j(\mathbf{v} - \tilde{\mathbf{u}}_j, \Psi_{\mathbf{v} - \tilde{\mathbf{u}}_j}(t')) \cdot r(\mathbf{v} - \tilde{\mathbf{u}}_j, t') \\ &\quad - \sum_j \alpha_j(\mathbf{v}, \Psi_{\mathbf{v}}(t')) \cdot r(\mathbf{v}, t') \end{aligned}$$

with initial condition $\mathbf{r}(t)$. Note that this equation is as Eq. (2) except that the dimensions in \hat{D} are removed. Moreover, the population sizes \mathbf{w} are replaced by the conditional expectations $\Psi_{\mathbf{v}}(t')$.

(2) Integrate. For each state \mathbf{v} with $r(\mathbf{v}, t) > \delta$, we compute an approximation $\Phi_{\mathbf{v}}(t + h)$ of the conditional expecta-

tion

$$E[\mathbf{W}(t+h) | \mathbf{V}(t') = \mathbf{v}, t' \in [t, t+h)],$$

that is, we assume that the system remains in state \mathbf{v} during $[t, t+h)$ and that the expected numbers of the large populations \mathbf{W} change deterministically and continuously in time. Thus, the \hat{n} -dimensional vector $\Phi_{\mathbf{v}}(t+h)$ is obtained by numerical integration of the ODE

$$\frac{d}{dt'} \Phi_{\mathbf{v}}(t') = \sum_{j=1}^m \hat{\mathbf{u}}_j \cdot \alpha_j(\mathbf{v}, \Phi_{\mathbf{v}}(t'))$$

with initial condition $\Phi_{\mathbf{v}}(t)$. The above ODEs are similar to Eq. (12) except that for $t' \in [t, t+h)$ the value $E[\alpha_j(\mathbf{X}(t'))]$ is approximated by $\alpha_j(\mathbf{v}, \Phi_{\mathbf{v}}(t'))$. For instance, if the j -th reaction is a bimolecular reaction that involves two populations with indices i, ℓ in \hat{D} then $E[\alpha_j(\mathbf{v}, \mathbf{W}(t')) | \mathbf{V}(t') = \mathbf{v}]$ is approximated by $c_j \cdot \phi_{\mathbf{v}}(i, t') \cdot \phi_{\mathbf{v}}(\ell, t')$ where the two last factors are the elements of the vector $\Phi_{\mathbf{v}}(t')$ corresponding to the i -th and ℓ -th population. Thus, in this case the correlations between the i -th and the ℓ -th populations are not taken into account which is reasonable if the two populations are large. Note that the correlations *are* taken into account when at least one population is represented as a discrete stochastic variable. If, for instance, $i \in \hat{D}$ and $\ell \in \hat{D}$, then we use the approximation $c_j \cdot v_i \cdot \phi_{\mathbf{v}}(\ell, t')$ where v_i is the entry in vector \mathbf{v} that represents the size of the i -th population.

(3) Distribute. In order to approximate $E[\mathbf{W}(t+h) | \mathbf{V}(t+h)]$ by $\Psi_{\mathbf{v}}(t+h)$ for all states \mathbf{v} , we have to replace the condition $\mathbf{V}(t') = \mathbf{v}, t' \in [t, t+h)$ by $\mathbf{V}(t+h) = \mathbf{v}$ in the conditional expectation $\Phi_{\mathbf{v}}(t+h)$ that was computed in step 2. This is done by “distributing” $\Phi_{\mathbf{v}}(t+h)$ according to the change in the distribution of $\mathbf{V}(t)$ as explained below. The idea is to take into account that \mathbf{V} enters state \mathbf{v} from \mathbf{v}' during the interval $[t, t+h)$. Assume that $[t, t+h)$ is an infinitesimal time interval and that $q(\mathbf{v}', \mathbf{v}, h)$, $\mathbf{v} \neq \mathbf{v}'$ is the probability to enter \mathbf{v} from \mathbf{v}' within $[t, t+h)$. Then

$$P(\mathbf{V}(t+h) = \mathbf{v}) = \sum_{\mathbf{v}' \neq \mathbf{v}} q(\mathbf{v}', \mathbf{v}, h) \cdot P(\mathbf{V}(t) = \mathbf{v}') + (1 - \sum_{\mathbf{v}' \neq \mathbf{v}} q(\mathbf{v}', \mathbf{v}, h)) \cdot P(\mathbf{V}(t) = \mathbf{v}). \quad (15)$$

Thus, we approximate $E[\mathbf{W}(t+h) | \mathbf{V}(t+h) = \mathbf{v}]$ as

$$\sum_{\mathbf{v}' \neq \mathbf{v}} \Phi_{\mathbf{v}'}(t+h) \cdot q(\mathbf{v}', \mathbf{v}, h) \cdot P(\mathbf{V}(t) = \mathbf{v}' | \mathbf{V}(t+h) = \mathbf{v}) \quad (16)$$

$$+ \Phi_{\mathbf{v}}(t+h) \cdot (1 - \sum_{\mathbf{v}' \neq \mathbf{v}} q(\mathbf{v}', \mathbf{v}, h)) \cdot P(\mathbf{V}(t) = \mathbf{v} | \mathbf{V}(t+h) = \mathbf{v}).$$

Obviously, we can make use of the current approximations $\mathbf{r}(t)$ and $\mathbf{r}(t+h)$ to compute the conditional probabilities $P(\mathbf{V}(t) = \mathbf{v}' | \mathbf{V}(t+h) = \mathbf{v})$. For a small time step h , $q(\mathbf{v}', \mathbf{v}, h) \approx h \cdot \alpha_j(\mathbf{v}', \Psi_{\mathbf{v}'}(t))$ if $\mathbf{v}' = \mathbf{v} + \hat{\mathbf{u}}_j$ and $q(\mathbf{v}', \mathbf{v}, h) \approx 0$ otherwise.

Using Eq. (16), we compute the approximation $\Psi_{\mathbf{v}}(t+h) \approx E[\mathbf{W}(t+h) | \mathbf{V}(t+h) = \mathbf{v}]$ as

$$\Psi_{\mathbf{v}}(t+h) = \sum_j \Phi_{\mathbf{v}-\hat{\mathbf{u}}_j}(t+h) \cdot \frac{p(\mathbf{v}-\hat{\mathbf{u}}_j, t)}{p(\mathbf{v}, t+h)} \cdot \alpha_j(\mathbf{v}-\hat{\mathbf{u}}_j, \Psi_{\mathbf{v}-\hat{\mathbf{u}}_j}(t)) \cdot h \quad (17)$$

$$+ \Phi_{\mathbf{v}}(t+h) \cdot \frac{p(\mathbf{v}, t)}{p(\mathbf{v}, t+h)} (1 - \sum_j \alpha_j(\mathbf{v}, \Psi_{\mathbf{v}}(t)) \cdot h).$$

Note that the first sum runs over all direct predecessors $\mathbf{v} - \hat{\mathbf{u}}_j$ of \mathbf{v} .

EXAMPLE 1 (CONT.). In the exclusive switch the expected number of molecules of type P_1 and/or P_2 may become high,




Figure 2: Probability distribution of the exclusive switch in Ex. 1.




Figure 3: The discrete stochastic part of the stochastic hybrid model of Ex. 1.

depending on the chosen parameters. If, for instance, $c_1 = c_2 = c_9 = c_{10} = 0.5$, $c_3 = c_4 = c_7 = c_8 = 0.005$, $c_5 = c_6 = 0.01$, and we start initially without any proteins, i.e. with probability one in state $\mathbf{y} = (0, 0, 1, 0, 0)$, then after 500 time units most of the probability mass is located around the states $\mathbf{x} = (92, 2, 0, 1, 0)$ and $\mathbf{x} = (2, 92, 0, 0, 1)$ (compare the plot in Fig. 2, left). Note that $x_3 = 0, x_4 = 1, x_5 = 0$ refers to the case that a molecule of type P_1 is bound to the promotor and $x_3 = x_4 = 0, x_5 = 1$ refers to the case that a molecule of type P_2 is bound to the promotor. Since for these parameters the system is symmetric, the expected populations of P_1 and P_2 are identical. Assume that at a certain time instant, both populations reach the threshold from which on we approximate them by continuous deterministic variables (we consider the unsymmetric case later, in Section 5). The remaining discrete model then becomes finite since only P_1 and P_2 have an infinite range in the original model ($\hat{n} = 2, \tilde{n} = 3$). More precisely, it contains only 3 states, namely the state \mathbf{v}_1 where the promotor is free ($x_3 = 1, x_4 = x_5 = 0$), the state \mathbf{v}_2 where P_1 is bound to the promotor ($x_3 = 0, x_4 = 1, x_5 = 0$), and the state \mathbf{v}_3 where P_2 is bound to the promotor ($x_3 = x_4 = 0, x_5 = 1$), see also Fig. 1. For $i \in \{1, 2\}$ let $W_i(t)$ be the population size of P_i . The differential equations which are used to approximate the conditional expectations $\Psi_{\mathbf{v}_j}(t+h)$, $j \in \{1, 2, 3\}$ are

$$\begin{aligned} \frac{d\phi_{\mathbf{v}_1}(i, t')}{dt'} &= c_i - c_{2+i} \cdot \phi_{\mathbf{v}_1}(i, t') - c_{4+i} \cdot \phi_{\mathbf{v}_1}(i, t') \\ \frac{d\phi_{\mathbf{v}_2}(i, t')}{dt'} &= -c_{2+i} \cdot \phi_{\mathbf{v}_2}(i, t') + (c_7 + c_9) \cdot (2 - i) \\ \frac{d\phi_{\mathbf{v}_3}(i, t')}{dt'} &= -c_{2+i} \cdot \phi_{\mathbf{v}_3}(i, t') + (c_8 + c_{10}) \cdot (i - 1) \end{aligned}$$

where $\phi_{\mathbf{v}}(1, t')$ and $\phi_{\mathbf{v}}(2, t')$ are the elements of the vector $\Phi_{\mathbf{v}_1}(t')$. Note that each of the 3 states $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ has a system of two differential equations, one for P_1 and one for P_2 . The transition rates in the discrete stochastic part of the model are illustrated in Fig. 3. Thus, after solving the differential equations above to compute $\Phi_{\mathbf{v}_j}(t+h)$ we obtain the vector $\Psi_{\mathbf{v}_j}(t+h)$ of the two conditional expectations for P_1 and P_2 from distributing $\Phi_{\mathbf{v}_1}(t+h), \Phi_{\mathbf{v}_2}(t+h), \Phi_{\mathbf{v}_3}(t+h)$ among the 3 states as defined in Eq. (17). For the parameters used in Fig. 1, left, the conditional expectations of the

pset	purely stochastic			stochastic hybrid						purely determ.	
	ex. time	$ Sig $	error	pop. thres.	ex. time	$ Sig $	m1	m2	m3	ex. time	m1
1a	11h 46min	$8 \cdot 10^5$	$7 \cdot 10^{-5}$	50	15sec	$4 \cdot 10^2$	0.005	0.2	0.30	1sec	0.03
				100	1min 50sec	$3 \cdot 10^3$	0.004	0.2	0.30		
1b	7min 43sec	$5 \cdot 10^4$	$7 \cdot 10^{-7}$	50	1min 19sec	$6 \cdot 10^3$	0.01	0.19	0.30	1sec	0.03
				100	2min 50sec	$3 \cdot 10^4$	0.01	0.19	0.30		
2	4h 51min	$2 \cdot 10^5$	$4 \cdot 10^{-5}$	50	25sec	$4 \cdot 10^2$	0.06	0.08	0.09	1sec	0.45
				100	28sec	$6 \cdot 10^2$	0.06	0.07	0.09		
3	2min 21sec	$7 \cdot 10^5$	$6 \cdot 10^{-5}$	50	18sec	$6 \cdot 10^3$	0.02	0.08	0.16	1sec	0.05
				100	1min 41sec	$4 \cdot 10^4$	0.01	0.05	0.12		

Table 2: Results for the exclusive switch example.

states \mathbf{v}_2 and \mathbf{v}_3 accurately predict the two stable regions where most of the probability mass is located. The state \mathbf{v}_1 has small probability and its conditional expectation is located between the two stable regions. It is important to point out that, for this example, a purely deterministic solution cannot detect the bistability because the deterministic model has a single steady-state [22]. Finally, we remark that in this example the number of states in the reduced discrete model is very small. If, however, populations with an infinite range but small expectations are present, we use the truncation described in Section 2 to keep the number of states small.

If at time t a population, say the i -th population, is represented by its conditional expectations, it is possible to switch back to the original discrete stochastic treatment. This is done by adding an entry to the states \mathbf{v} for the i -th dimension. This entry then equals $\psi_{\mathbf{v}}(i, t)$. This means that at this point we assume that the conditional probability distribution has mass one for the value $\psi_{\mathbf{v}}(i, t)$. Note that here switching back and forth between discrete stochastic and continuous deterministic representations is based on a population threshold. Thus, if the expectation of a population oscillates we may switch back and forth in each period.

5. EXPERIMENTAL RESULTS

We implemented the numerical solution of the stochastic hybrid model described above in C++ as well as the fast solution of the discrete stochastic model described in Section 2. In our implementation we dynamically switch the representation of a random variable whenever it reaches a certain population threshold. We ran experiments with two different thresholds (50 and 100) on an Intel 2.5GHz Linux workstation with 8GB of RAM. In this section we present 3 examples to that we applied our algorithm, namely the exclusive switch, Goutsias’ model, and a predator-prey model. Our most complex example has 6 different chemical species and 10 reactions. We compare our results to a purely stochastic solution where switching is turned off as well as to a purely deterministic solution. For all experiments, we fixed the cutting threshold $\delta = 10^{-14}$ to truncate the infinite state space as explained in Sec. 2.

Exclusive Switch. We chose different parameters for the exclusive switch in order to test whether our hybrid approach works well if

- 1) the populations of P_1 and P_2 are large (a) or small (b),
- 2) the model is unsymmetric (e.g. P_1 is produced at a higher rate than P_2 and degrades at a slower rate than P_2),

- 3) the bistable form of the distribution is destroyed (i.e. promotor binding is less likely, unbinding is more likely).

The following table lists the parameter sets (psets):

pset	c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
1a	5	5	0.0005	0.0005	0.1	0.1	0.005	0.005	5	5
1b	0.5	0.5	0.0005	0.0005	0.1	0.1	0.005	0.005	0.5	0.5
2	5	0.5	0.0005	0.005	0.1	0.1	0.005	0.005	5	0.5
3	0.5	0.5	0.0005	0.0005	0.01	0.01	0.1	0.1	0.5	0.5

We chose a time horizon of $t = 500$ for all parameter sets. Note that in the case of pset 3 the probability distribution forms a thick line in the state space (compare the plot in Fig. 2, right). We list our results in Table 2 where the first column refers to the parameter set. Column 2 to 4 list the results of a purely stochastic solution (see Section 2) where “ex. time” refers to the execution time, $|Sig|$ to the average size of the set of significant states and “error” refers to the amount of probability mass lost due to the truncation with threshold δ , i.e. $1 - \sum_{\mathbf{x} \in Sig} p(\mathbf{x}, t)$. The columns 6-10 list the results of our stochastic hybrid approach and column 5 lists the population threshold used for switching in the representations in the stochastic hybrid model. Here, “m1”, “m2”, “m3” refer to the relative error of the first three moments of the joint probability distribution at the final time instant. For this, we compare the (approximate) solution of the hybrid model with the solution of the purely stochastic model. Since we have five species, we simply take the average relative error over all species. Note that even if a species is represented by its conditional expectations, we can approximate its first three moments by

$$E[\mathbf{W}(t)^i] \approx \sum_{\mathbf{v}} (\Psi_{\mathbf{v}}(t))^i \cdot r(\mathbf{v}, t)$$

where the i -th power of the vectors are taken component-wise. Finally, in the last two columns we list the results of a purely deterministic solution as explained in Section 3. The last column refers to the average relative error of the expected populations when we compare the purely deterministic solution to the purely stochastic solution. Note that the deterministic solution of the exclusive switch yields an accurate approximation of the first moment (except for pset 2) because of the symmetry of the model. It does, however, not reveal the bistability of the distribution. As opposed to that, the hybrid solution *does* show this important property. For pset 1 and 3, the conditional expectations of the 3 discrete states are such that two of them match exactly the two stable regions where most of the probability mass is located (see also Example 1 in Sec. 4). The remaining conditional

model	purely stochastic			stochastic hybrid						purely determ.	
	ex. time	$ Sig $	error	pop. thres.	ex. time	$ Sig $	m1	m2	m3	ex. time	m1
Goutsias	1h 16min	$1 \cdot 10^6$	$4 \cdot 10^{-7}$	50	8min 47sec	$1 \cdot 10^5$	0.001	0.07	0.13	1sec	0.95
				100	48min 57sec	$6 \cdot 10^3$	0.0001	0.0003	0.001		
p.-prey	6h 6min	$5 \cdot 10^5$	$1 \cdot 10^{-7}$	50	8min 56sec	$2 \cdot 10^4$	0.06	0.15	0.27	1sec	0.86
				100	1h 2min	$8 \cdot 10^4$	0.04	0.11	0.23		

Table 3: Results for Goutsias’ model and the predator-prey model.





Figure 4: Expected populations in Goutsias’ model.

expectation of the state where the promotor region is free has small probability and predicts a conditional expectation between the two stable regions. The execution time of the purely stochastic approach is high in the case of pset 1a, because the expected populations of P_1 and P_2 are high. This yields large sizes of Sig while we iterate over time. During the hybrid solution, we switch when the populations reach the threshold and the size of Sig drops to 3. Thus, the average number of significant states is much smaller. In the case of pset 1b, the expected populations are small and we use a deterministic representation for protein populations only during a short time interval (at the end of the time horizon). For pset 2, the accuracy of the purely deterministic solution is poor because the model is no longer symmetric. The accuracy of the hybrid solution on the other hand is independent of the symmetry of the model.

Goutsias’ Model. In [11], Goutsias defines a model for the transcription regulation of a repressor protein in bacteriophage λ . This protein is responsible for maintaining lysogeny of the λ virus in *E. coli*. The model involves 6 different species and the following 10 reactions.



We used the following parameters that differ from the original parameters used in [11] in that they increase the number of RNA molecules (because with the original parameters, all populations remain small).

c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	c_{10}
0.043	7e-4	71.5	3.9e-6	0.02	0.48	2e-4	9e-12	0.08	0.5

Table 3 shows the results for the Goutsias’ model where we use the same column labels as above. We always start initially with 10 molecules of RNA, M, and D, as well as 2 DNA molecules. We choose the time horizon as $t = 4$. Note that the hybrid solution as well as the purely deterministic

solution are feasible for much longer time horizons. The increase of the size of the set of significant states makes the purely stochastic solution infeasible for longer time horizons. As opposed to that the memory requirements of the hybrid solution remain tractable. In Fig. 4 we plot the means of two of the six species obtained from the purely stochastic (stoch), purely deterministic (determ), and the hybrid (hyb) solution. Note that a purely deterministic solution yields very poor accuracy (relative error of the means is 95%).

Predator Prey. We apply our algorithm to the predator prey model described in [9]. It involves two species A and B and the reactions are $A \rightarrow 2A$, $A + B \rightarrow 2B$, and $B \rightarrow \emptyset$. The model shows sustainable periodic oscillations until eventually the population of B reaches zero. We use this example to test the switching mechanism of our algorithm. We choose rate constants $c_1 = 1$, $c_2 = 0.03$, $c_3 = 1$ and start initially with 30 molecules of type A and 120 molecules of type B . For a population threshold of 50, we start with a stochastic representation of A and a deterministic representation of B . Then, around time 1.3 we switch to a purely stochastic representation since the expectation of B becomes less than 50. Around time $t = 6.1$ we switch the representation of A because $E[A(t)] > 50$, etc. We present our detailed results in Table 3. Similar to Goutsias’ model, the deterministic solution has a high relative error whereas the hybrid solution yields accurate results.

6. CONCLUSION

We presented a stochastic hybrid model for the analysis of networks of chemical reactions. This model is based on a dynamic switching between a discrete stochastic and a continuous deterministic representation of the chemical populations. Instead of solving the underlying partial differential equation, we propose a fast numerical procedure that exploits the fact that for large populations the conditional expectations give appropriate approximations. Our experimental results substantiate the usefulness of the method. As future work we plan to include a diffusion approximation for populations of intermediate size.

7. REFERENCES

- [1] K. Ball, T. G. Kurtz, L. Popovic, and G. Rempala. Asymptotic analysis of multiscale approximations to reaction networks. *The Annals of Applied Probability*, 16(4):1925–1961, 2006.
- [2] M. L. Bujorianu and J. Lygeros. Towards a general theory of stochastic hybrid systems. In *Stochastic Hybrid Systems*, volume 337 of *Lecture Notes in Control and Information Sciences*, pages 3–30. Springer, 2006.

- [3] K. Burrage, M. Hegland, F. Macnamara, and B. Sidje. A Krylov-based finite state projection algorithm for solving the chemical master equation arising in the discrete modelling of biological systems. In *Proceedings of the Markov 150th Anniversary Conference*, pages 21–38. Bosc Books, 2006.
- [4] C. Cassandras and J. Lygeros. Stochastic hybrid systems: Research issues and areas. In *Stochastic Hybrid Systems, (C.G. Cassandras, and J. Lygeros, Ed.s)*, pages 1–14. Taylor and Francis, 2006.
- [5] F. Didier, T. A. Henzinger, M. Mateescu, and V. Wolf. Approximation of event probabilities in noisy cellular processes. In *Proc. of CMSB*, volume 5688 of *LNBI*, page 173, 2009.
- [6] F. Didier, T. A. Henzinger, M. Mateescu, and V. Wolf. Fast adaptive uniformization of the chemical master equation. In *Proc. of HIBI*, pages 118–127. IEEE Computer Society, 2009.
- [7] S. Engblom. Spectral approximation of solutions to the chemical master equation. *Journal of Computational and Applied Mathematics*, 229(1):208 – 221, 2009.
- [8] N. Fedoroff and W. Fontana. Small numbers of big molecules. *Science*, 297:1129–1131, 2002.
- [9] D. T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.*, 81(25):2340–2361, 1977.
- [10] D. T. Gillespie. A rigorous derivation of the chemical master equation. *Physica A*, 188:404–425, 1992.
- [11] J. Goutsias. Quasiequilibrium approximation of fast reaction kinetics in stochastic biochemical systems. *J. Chem. Phys.*, 122(18):184102, 2005.
- [12] E. L. Haseltine and J. B. Rawlings. Approximate simulation of coupled fast and slow reactions for chemical kinetics. *J. Chem. Phys.*, 117:6959–6969, 2002.
- [13] M. Hegland, C. Burden, L. Santoso, S. Macnamara, and H. Booth. A solver for the stochastic master equation applied to gene regulatory networks. *J. Comput. Appl. Math.*, 205:708–724, 2007.
- [14] A. Hellander and P. Lötstedt. Hybrid method for the chemical master equation. *Journal of Computational Physics*, 227, 2007.
- [15] D. Henderson, R. Boys, C. Proctor, and D. Wilkinson. Linking systems biology models to data: a stochastic kinetic model of p53 oscillations. In *Handbook of Appl. Bayesian Analysis*. Oxford University Press, 2009.
- [16] T. Henzinger, M. Mateescu, and V. Wolf. Sliding window abstraction for infinite Markov chains. In *Proc. CAV*, volume 5643 of *LNCS*, pages 337–352. Springer, 2009.
- [17] G. Horton, V. G. Kulkarni, D. M. Nicol, and K. S. Trivedi. Fluid stochastic Petri nets: Theory, applications, and solution techniques. *European Journal of Operational Research*, 105(1):184–201, 1998.
- [18] T. Jahnke. An adaptive wavelet method for the chemical master equation. *SIAM J. Scientific Computing*, 31(6):4373–4394, 2010.
- [19] A. Jensen. Markoff chains as an aid in the study of Markoff processes. *Skandinavisk Aktuarietidskrift*, 36:87–91, 1953.
- [20] N. G. v. Kampen. *Stochastic Processes in Physics and Chemistry*. Elsevier, 3rd edition, 2007.
- [21] T. Kurtz. *Approximation of Population Processes*. Society for Industrial Mathematics, 1981.
- [22] A. Loinger, A. Lipshtat, N. Q. Balaban, and O. Biham. Stochastic simulations of genetic switch systems. *Phys. Rev. E*, 75(2):021904, 2007.
- [23] H. H. McAdams and A. Arkin. Stochastic mechanisms in gene expression. *PNAS, USA*, 94:814–819, 1997.
- [24] H. H. McAdams and A. Arkin. It’s a noisy business! *Trends in Genetics*, 15(2):65–69, 1999.
- [25] B. Munsky and M. Khammash. The finite state projection algorithm for the solution of the chemical master equation. *J. Chem. Phys.*, 124:044144, 2006.
- [26] P. Patel, B. Arcangioli, S. Baker, A. Bensimon, and N. Rhind. DNA replication origins fire stochastically in fission yeast. *Mol. Biol. Cell*, 17:308–316, 2006.
- [27] J. Paulsson. Summing up the noise in gene networks. *Nature*, 427(6973):415–418, 2004.
- [28] B. Philippe and R. Sidje. Transient solutions of Markov processes by Krylov subspaces. In *Proc. of the 2nd International Workshop on the Numerical Solution of Markov Chains*, pages 95–119. Kluwer Academic Publishers, 1995.
- [29] J. Puchalka and A. M. Kierzek. Bridging the gap between stochastic and deterministic regimes in the kinetic simulations of the biochemical reaction networks. *Biophysical Journal*, 86(3):1357 – 1372, 2004.
- [30] C. Rao, D. Wolf, and A. Arkin. Control, exploitation and tolerance of intracellular noise. *Nature*, 420(6912):231–237, 2002.
- [31] H. Salis and Y. Kaznessis. Accurate hybrid stochastic simulation of a system of coupled chemical or biochemical reactions. *J. Chem. Phys.*, 122, 2005.
- [32] R. Sidje, K. Burrage, and S. MacNamara. Inexact uniformization method for computing transient distributions of Markov chains. *SIAM J. Sci. Comput.*, 29(6):2562–2580, 2007.
- [33] W. J. Stewart. *Introduction to the Numerical Solution of Markov Chains*. Princeton University Press, 1995.
- [34] P. S. Swain, M. B. Elowitz, and E. D. Siggia. Intrinsic and extrinsic contributions to stochasticity in gene expression. *PNAS, USA*, 99(20):12795–12800, 2002.
- [35] A. Warmflash and A. Dinner. Signatures of combinatorial regulation in intrinsic biological noise. *PNAS*, 105(45):17262–17267, 2008.