

METHODOLOGY ARTICLE

Open Access

Efficient parametric analysis of the chemical master equation through model order reduction

Steffen Waldherr^{1*} and Bernard Haasdonk²

Abstract

Background: Stochastic biochemical reaction networks are commonly modelled by the chemical master equation, and can be simulated as first order linear differential equations through a finite state projection. Due to the very high state space dimension of these equations, numerical simulations are computationally expensive. This is a particular problem for analysis tasks requiring repeated simulations for different parameter values. Such tasks are computationally expensive to the point of infeasibility with the chemical master equation.

Results: In this article, we apply parametric model order reduction techniques in order to construct accurate low-dimensional parametric models of the chemical master equation. These surrogate models can be used in various parametric analysis task such as identifiability analysis, parameter estimation, or sensitivity analysis. As biological examples, we consider two models for gene regulation networks, a bistable switch and a network displaying stochastic oscillations.

Conclusions: The results show that the parametric model reduction yields efficient models of stochastic biochemical reaction networks, and that these models can be useful for systems biology applications involving parametric analysis problems such as parameter exploration, optimization, estimation or sensitivity analysis.

Keywords: Stochastic biochemical network, Model reduction, Reduced basis, Genetic regulatory network, Computational efficiency, Parameter estimation

Background

The chemical master equation (CME) is the most basic mathematical description of stochastic biomolecular reaction networks [1,2]. The CME is a generally infinite-dimensional linear differential equation. It characterizes the temporal development of the probabilities that the network is in any of its possible configurations, where the different configurations are characterized by the molecular copy numbers of the network's chemical species.

Due to its infinite dimension, the CME is usually not directly solvable, not even with numerical methods. A recent breakthrough in the numerical treatment of the CME was the establishment of the finite state projection (FSP) method by Munsky and Khammash [3]. They

showed that it is possible to compute a good approximation to the real solution by projecting the CME to a suitable finite subdomain of the network's state space, and solving the resulting finite-dimensional linear differential equation on that domain. Nevertheless, the FSP approach still yields very high-dimensional models which are computationally expensive to simulate, even for small biochemical networks. The efficient simulation of the CME is an area of active research, and recently other simulation methods have been developed that can also be used for larger networks [4,5].

Despite this progress, the direct simulation of the CME remains a computational bottleneck for common model analysis tasks in systems biology. It is especially problematic for tasks which require the repeated simulation of the model using different parameter values, for example identifiability analysis, parameter estimation, or model sensitivity analysis. Thereby, while a single or a few evaluations of a CME model with the FSP or other approaches may

*Correspondence: waldherr@ist.uni-stuttgart.de

¹ Institute for Systems Theory and Automatic Control, University of Stuttgart, Pfaffenwaldring 9, Stuttgart, Germany
Full list of author information is available at the end of the article

still be computationally feasible, the necessity of many repeated simulations will quickly render higher-level analysis tasks infeasible.

Mathematical methods that approximate the behaviour of a high-dimensional original model through a low-dimensional reduced model are a common way to deal with complex models. Especially for linear differential equations, model order reduction is a well established field and several methods to compute reduced order models are available [6]. Note that the step of generating a reduced model is usually computationally more expensive than a single or even a few simulations of the original high-dimensional model. But the simulation of the resulting reduced models is frequently orders of magnitude faster than the solution of the original model. So, model reduction is worth the effort if many repeated simulations are to be expected. Unfortunately, for analysis tasks which require the repeated model simulation with different parameters, classical model reduction methods are not helpful. With these methods, the reduced model depends on specific parameter values in the original model, and the reduction needs to be redone for different parameter values. Thus, for the mentioned analysis tasks, the model reduction process would have to be repeated for each new parameter value, and no gain in computational efficiency would typically be possible. While classical model reduction techniques have been applied to the CME in the past [7], they are not so suitable for parametric analysis tasks.

Fortunately, model reduction methods where parameters from the original model are retained as adjustable parameters also in the reduced model are now being developed. These methods allow to compute a reduced model which uses the same parameters as the original model, and where the reduced model can directly be simulated with any choice of parameter values [8-11].

The purpose of this paper is to introduce the application of these parametric model reduction methods to finite-state approximations of the chemical master equation, and to show possible usage scenarios of such an approach. The structure is as follows. In the following section, we introduce some background and notation concerning the modelling of chemical reaction networks and parametric model order reduction. We also show how the parametric model order reduction methods can in fact be applied to the CME. Afterwards, we apply the reduction technique on two reaction network models and corresponding parametric analysis tasks.

Methods

We start with some preparatory background on the chemical master equation (CME) and parametric model order reduction. This serves in particular to fix the notation used throughout the remainder of the article. Then the

application of parametric model order reduction to the CME is introduced.

The chemical master equation

The structure of a biochemical reaction network is characterized completely by the list of involved species, denoted as X_1, X_2, \dots, X_n , and the list of reactions, denoted as

$$\sum_{i=1}^n \sigma_{ij} X_i \rightarrow \sum_{i=1}^n \varphi_{ij} X_i, \quad j = 1, \dots, m, \quad (1)$$

where m is the number of reactions in the network, and the factors $\sigma_{ij} \in \mathbb{N}_0$ and $\varphi_{ij} \in \mathbb{N}_0$ are the stoichiometric coefficients of the reactant and product species, respectively [12]. The net change in the amount of species i occurring through reaction j is given by

$$N_{ij} = \varphi_{ij} - \sigma_{ij}. \quad (2)$$

Reversible reactions can always be written in the form (1) by splitting the forward and reverse path into two separate irreversible reactions.

For a stochastic network model, the variables of interest are the probabilities that the network is in any of the possible states which are characterized by the molecular copy numbers of the individual species X_1, X_2, \dots, X_n . We denote the molecular copy number of X_i by $[X_i] \in \mathbb{N}_0$. Then, the state variables of the stochastic model are given by the real numbers

$$p(t, x_1, x_2, \dots, x_n) = \text{Prob}([X_1] = x_1, [X_2] = x_2, \dots, [X_n] = x_n \text{ at time } t), \quad (3)$$

for $x_i \in \mathbb{N}_0, i = 1, \dots, n$. As a short-hand notation for (3), we write $p(t, x)$, with $x \in \mathbb{N}_0^n$.

The transitions from one state to another are determined by chemical reactions according to (1). The changes in the molecule numbers are described by the stoichiometric reaction vectors

$$v_j = (N_{1j} \ N_{2j} \ \dots \ N_{nj})^T \in \mathbb{Z}^n. \quad (4)$$

To avoid needlessly complicated cases, we assume $v_j \neq v_k$ for $j \neq k$.

The probabilities of the network being in any of the possible states x evolve over time, and their evolution is governed by the chemical master equation (CME) as derived by [1]. From a given molecular state x , one can compute the propensity v_j that reaction j takes place according to the law of mass action as

$$v_j(x, \theta) = \theta_j \prod_{i=1}^n \binom{x_i}{\sigma_{ij}}, \quad (5)$$

where $\theta = (\theta_j)_{j=1}^m$ is the vector of reaction rate constants, which are model parameters depending on the physical properties of the molecules involved in the reactions. The

propensities are related to the probability that reaction j will occur in a short time interval of length dt when the system is in state x :

$$\text{Prob}(\text{reaction } j \text{ occurs in } [t, t + dt] | [X] = x) = v_j(x, \theta) dt + o(dt). \quad (6)$$

Taking the possible transitions and the corresponding reaction propensities together yields the chemical master equation (CME), a linear differential equation where the variables are the probabilities that the system is in each of the possible molecular states x :

$$\frac{d}{dt} p(t, x) = \sum_{j=1}^m (v_j(x - v_j, \theta) p(t, x - v_j) - v_j(x, \theta) p(t, x)), \quad (7)$$

for $x \in \mathbb{N}_0^n$. The CME (7) is subject to an initial condition $p(t_0, x) = p_0(x)$ for $x \in \mathbb{N}_0^n$.

Despite being linear, the CME is hard to solve numerically. This is due to the problem that the state space is for most systems infinite-dimensional, since all possible states $x \in \mathbb{N}_0^n$ of the reaction network (1) must in general be considered. Instead of directly solving the CME (7), a number of alternative approaches to study the stochastic dynamics of biochemical reaction networks have been suggested. The most common approach is to generate a simulated realization of the stochastic process described by the reaction network (1), using for example the Gillespie algorithm [13]. In this approach, the probabilities $p(t, x)$ for the possible system states are obtained from many simulated realizations. However, since this requires a large number of realizations, it is computationally expensive.

As a more direct approach, Munsky and Khammash [3] have proposed the finite state projection (FSP), where the CME is solved on a finite subset of the state space. Here, this subset is denoted by Ω , and is defined as

$$\Omega = \{x^{(i)} \mid i = 1, \dots, d\} \subset \mathbb{N}_0^n, \quad (8)$$

where the $x^{(i)}$ are the system states for which the probabilities are computed in the projected model. The underlying assumption is that the probabilities for other states will be very low on the time scale of interest—otherwise the FSP may not yield good approximations to the solution of the CME. In particular we assume the time interval of interest to be given by $[0, T]$ for final time $T > 0$. The probabilities for the states $x^{(i)}$ in Ω are written in the vector $P(t)$ approximating $p(x, t)$ at the finite number of states Ω :

$$P(t) = (P_i(t))_{i=1, \dots, d} \approx \left(p(t, x^{(i)}) \right)_{i=1, \dots, d} \in [0, 1]^d. \quad (9)$$

The equation to be solved with the FSP approximation is

$$\begin{aligned} \frac{d}{dt} P(t) &= A(\theta) P(t) \quad \text{for } t \in (0, T) \\ P(0) &= P_0, \end{aligned} \quad (10)$$

where $A(\theta) \in \mathbb{R}^{d \times d}$ is the matrix of state transition propensities, and $P_0 = (p_0(x^{(i)}))_{i=1, \dots, d}$ is a vector of initial probabilities for the states in Ω . The elements of the matrix $A(\theta)$ are computed as

$$\begin{aligned} A_{ii}(\theta) &= - \sum_{r=1}^m v_r(x^{(i)}, \theta) \\ A_{ij}(\theta) &= \begin{cases} v_r(x^{(j)}, \theta) \text{ if } x^{(j)} = x^{(i)} + v_r, r = 1, \dots, m \\ 0 \text{ otherwise.} \end{cases} \end{aligned} \quad (11)$$

We will frequently omit the parameter dependence of the solution (and other parametric quantities). Hence the solution $P(t)$, as abbreviation of $P(t, \theta)$, of (10) is an approximation to the solution $p(t, x)$ of the original CME on the domain Ω . Munsky and Khammash [3] have also derived an upper bound on the error between the solution $P(t)$ computed via the FSP, and the solution of the original CME $p(t, x)$ on Ω .

Here, we consider in addition an output vector $y \in \mathbb{R}^p$ defined by

$$y(t) = CP(t), \quad (12)$$

with $C \in \mathbb{R}^{p \times d}$. Examples for relevant outputs are the probability that the molecular copy numbers are in a certain domain $\bar{\Omega} \subset \Omega$, which is achieved by the row vector output matrix C defined by $C_i = 1$ if $x^{(i)} \in \bar{\Omega}$, otherwise $C_i = 0$, with $p=1$, or the expected molecular copy numbers, given by

$$y_e(t) = \sum_{i=1}^d x^{(i)} P_i(t), \quad (13)$$

i.e. $C = (x^{(1)}, \dots, x^{(d)})$ with $p=n$.

The basic motivation for the model reduction presented here is that we are interested in parametric analysis of the model, where the model (10) has to be solved many times with different values for the parameters θ . Due to the typical high dimensions of the matrix $A(\theta)$, already a single simulation is computationally expensive, and analysis tasks requiring many repeated simulations are often computationally infeasible. Thus, the primary goal is to derive a reduced model which is rapidly solvable and provides an approximation $\hat{y}(t)$ to the output $y(t)$, potentially without any consideration of the original state vector $P(t)$.

Order reduction of parametric models

Model order reduction of parametric problems is a very active research field in systems theory, engineering and

applied mathematics. We refer to [8,10,11] and references therein for more information on the topic.

Here, we apply the reduction technique for parametric problems presented in [9] adopted to our notation. It is based on two biorthogonal global projection matrices $V, W \in \mathbb{R}^{d \times r}$ with $r \ll d$ and $W^T V = Id$, where r is the dimension of the reduced model. The matrix V is assumed to span a space that approximates the system state variation for all parameters and times. The construction of such matrices will be detailed in the next subsection.

The gain of computational efficiency in repeated simulations comes from a separation of the simulation task into a computationally expensive “offline” phase and a computationally cheap “online” phase. In the offline phase, suitable projection matrices V and W are computed without fixing specific parameter values. With the projection matrices, a reduced model with the same free parameters as the original model is computed. In the online phase, the reduced model is simulated with the actually chosen parameter values, which is typically several orders of magnitude faster than the simulation of the original model. For analysis tasks with repeated simulations, only the online phase has to be repeated for different choices of the parameter values, yielding an overall gain in computational efficiency.

Decomposition in parametric and non-parametric part

The reduction technique assumes a separable parameter dependence of the full system matrices and the initial condition. This means, we assume that there exist a suitable small constant $Q_A \in \mathbb{N}$, parameter independent components $A^{[q]} \in \mathbb{R}^{d \times d}$ and parameter dependent scalar coefficient functions $\vartheta_A^{[q]}(\theta)$ for $q = 1, \dots, Q_A$ such that

$$A(\theta) = \sum_{q=1}^{Q_A} \vartheta_A^{[q]}(\theta) A^{[q]} \quad (14)$$

and similarly for the system matrix C and initial condition P_0 . We assume that $\theta \in \mathcal{P}$ stems from some domain $\mathcal{P} \subset \mathbb{R}^m$ of admissible parameters. In the next step, the reduced component matrices and initial conditions are determined by

$$A_r^{[q]} := W^T A^{[q]} V, \quad C_r^{[q]} := C^{[q]} V, \quad P_{0r}^{[q]} := W^T P_0^{[q]}. \quad (15)$$

for $q = 1, \dots, Q_A$. The resulting quantities $A_r^{[q]}$, $C_r^{[q]}$, and $P_{0r}^{[q]}$ are r -dimensional vectors or matrices and independent of the high dimension d . The basis computation and the computation of these reduced system components is performed once and parameter-independently in the offline-phase. Then, in the online-phase, for any new

parameter θ the reduced system matrices and the initial condition are assembled by

$$A_r(\theta) = \sum_{q=1}^{Q_A} \vartheta^{[q]}(\theta) A_r^{[q]} \quad (16)$$

and similarly for $P_{r0}(\theta)$ and $C_r(\theta)$. The low dimensional reduced system that remains to be solved is

$$\begin{aligned} \frac{d}{dt} P_r(t) &= A_r(\theta) P_r(t) \quad \text{for } t \in (0, T) \\ P_r(0) &= P_{r0}(\theta) \\ \hat{y}(t) &= C_r(\theta) P_r(t). \end{aligned} \quad (17)$$

From the reduced state $P_r(t)$, an approximate state for the full system can be reconstructed at any desired time by $\hat{P}(t) = V P_r(t)$. Also the difference between the approximated output $\hat{y}(t)$ and the output $y(t)$ of the original model can be bounded by so called error estimators. A-posteriori error bounds for the reduced systems as considered here are given in [9].

Basis generation

Different methods for the computation of the projection bases V and W exist. In systems theory, methods like balanced truncation, Hankel-norm approximation or moment matching are applied, that approximate the input-output behaviour of a linear time-invariant system [6]. The resulting reduced models can be applied for varying input signals. Extensions to parametric problems exist, e.g. [8,11]. As we do not have varying inputs in the problem studied here, we consider snapshot-based approaches to be more suitable. This means, the projection bases are constructed by solution snapshots, i.e. special solutions computed for selected parameter values.

The generation of projection matrices V and W must be done in such a way, that they are globally well approximating the system states over the parameter and time domain. A possible way to achieve this is the POD-Greedy algorithm, which has been introduced in [14] and is meanwhile a standard procedure in reduced basis methods [15,16]. The algorithm makes use of a repeated proper orthogonal decomposition (POD) of trajectories $P : [0, T] \rightarrow \mathbb{R}^d$, which for our purposes can be defined as

$$POD(P) := \arg \min_{v \in \mathbb{R}^d, \|v\|=1} \int_0^T \|P(t) - (v^T P(t))v\|^2 dt. \quad (18)$$

Intuitively, $POD(P) \in \mathbb{R}^d$ is a state space vector representing the single dominant mode that minimizes the squared mean projection error. Computationally, this minimization task is solved by a reformulation as a suitable eigenvalue problem. Consider the correlation matrix $C = \int_0^T P(t)P(t)^T dt$. Then, $v^* = POD(P) \in \mathbb{R}^d$ is an

eigenvector corresponding to the largest eigenvalue λ_{max} of C , i.e., $Cv^* = \lambda_{max}v^*$. For additional theoretical and computational details on POD we refer to [17,18]. We further require a finite subset of parameters $\mathcal{P}_{train} \subset \mathcal{P}$, that are used in the basis generation process. As error indicator $\Delta(\theta, V)$ we use the projection error of the full system trajectory on the reduced space spanned by the orthonormal columns of V , i.e.

$$\Delta(\theta, V) := \int_0^T \|P(t, \theta) - VV^T P(t, \theta)\|^2 dt. \quad (19)$$

The POD-Greedy procedure which is given in the pseudo-code below, starts with an arbitrary orthonormal initial basis $V_{N_0} \in \mathbb{R}^{d \times N_0}$ and performs an incremental basis extension. The algorithm repeatedly identifies the currently worst resolved parameter (a), orthogonalizes the corresponding full trajectory with the current reduced space (b), computes a POD of the error trajectory (c), and inserts the dominant mode into the basis (d).

function $V = \text{POD-Greedy}(\mathcal{P}_{train}, V_{N_0}, \varepsilon_{tol})$

1. $N := N_0$
2. while $\varepsilon_N := \max_{\theta \in \mathcal{P}_{train}} \Delta(\theta, V_N) > \varepsilon_{tol}$
 - (a) $\theta^* := \arg \max_{\theta \in \mathcal{P}_{train}} \Delta(\theta, V_N)$
 - (b) $E(t) := P(t, \theta^*) - V_N(V_N^T P(t, \theta^*))$
 - (c) $v_{N+1} := \text{POD}(E)$
 - (d) $V_{N+1} := [V_N, v_{N+1}]$
 - (e) $N := N + 1$

3. end while

Note that the algorithm is implemented such that the simulation of the full model, yielding $P(t, \theta)$ in (19), is only performed once for each θ in the training set \mathcal{P}_{train} .

For concluding the basis generation, we set $W := V$. This satisfies the biorthogonality condition $W^T V = Id$, as V has orthonormal columns by construction. In practice the time-integrals in (18) are realized by a finite sampling of the time interval.

A theoretical underpinning for the POD-Greedy algorithm has recently been provided by the analysis of convergence rates [19]. This is based on the approximation-theoretical notion of the *Kolmogorov n -width* $d_N(\mathcal{F})$ of a given set $\mathcal{F} \subset \mathbb{R}^d$, which quantifies how well the set can be approximated by arbitrary N -dimensional linear subspaces of \mathbb{R}^d . The convergence statement for the case of exponential convergence then can be summarized as follows: If the set of solutions $\mathcal{F} := \{P(t, \theta) | t \in [0, T], \theta \in \mathcal{P}\} \subset \mathbb{R}^d$ is compact and has an exponentially decaying Kolmogorov n -width $d_N(\mathcal{F}) \leq Me^{-aN^\alpha}$ for some $M, a, \alpha > 0$ and all $N \in \mathbb{N}$, then the error sequence $(\varepsilon_N)_{N \in \mathbb{N}}$ generated by the POD-Greedy procedure (cf. the definition in Step 2. in the pseudo code) also decays with

an exponential rate, $\varepsilon_N \leq CMe^{-cN^\beta}$ with suitable constants $\beta, c, C > 0$ depending on M, a, α . Thus, if the set of solutions can be approximated by linear subspaces with an exponentially decaying error term, then the POD-Greedy algorithm will in fact find an approximation with an exponentially decaying error term, though possibly with suboptimal parameters in the error bound.

Extensions of the POD-Greedy algorithm exist, e.g. allowing more than one mode per extension step, performing adaptive parameter and time-interval partitioning, or enabling training-set adaptation [15,16,20].

Reduced models of the parametrized chemical master equation

In this section, we describe how to apply the reduction method for parametrized models presented in the previous section to FSP models for the chemical master equation.

As discussed in the previous section, the first step in the proposed reduction method is a decomposition of the d -dimensional system matrix $A(\theta)$ as in (14). Such a decomposition is possible for the case of mass action reaction propensities, as defined in (5), or generalized mass action, as recently suggested for the chemical master equation [21]. In this case, the length of the parameter vector θ is equal to the number of reactions m , and we decompose $A(\theta)$ into m terms as

$$A(\theta) = \theta_1 A^{[1]} + \dots + \theta_m A^{[m]}. \quad (20)$$

Hence, concerning the notation given before, we have $Q_A = m$ components $A^{[q]}$ and coefficient functions $\vartheta_A^{[q]}(\theta) = \theta_q$. Each matrix $A^{[q]}$ in this decomposition comes from just the transition propensities corresponding to reaction q , and is defined by

$$A_{ii}^{[q]} = - \prod_{k=1}^n (x_k^{(i)})^{\sigma_{kq}}$$

$$A_{ij}^{[q]} = \begin{cases} \prod_{k=1}^n (x_k^{(j)})^{\sigma_{kq}} & \text{if } x^{(j)} = x^{(i)} + v_q \\ 0 & \text{otherwise.} \end{cases} \quad (21)$$

More generally, such a decomposition is also possible if reaction rate propensities can be decomposed into the product of two terms, with the first term depending on parameters only, and the second term on molecule numbers only. This case is for example encountered when the temperature-dependance of the reaction rate constant is relevant, and the temperature T is a variable parameter in the Arrhenius equation $\theta = Ae^{-\frac{E_A}{RT}}$. Since the output matrix C and the initial condition P_0 are usually not depending on parameters in this framework, a decomposition of C and P_0 is not considered.

The situation is more difficult for reaction propensities involving for example rational terms with parameters in the denominator. The denominator parameters can not be included in the reduced order model by the decomposition outlined in (20) and (21). If variations in these parameters are however not relevant to the planned analysis, then they can be set to their nominal value, and the decomposition can directly be done as described above. Alternatively, approximation steps can be performed, such as Taylor series expansion or empirical interpolation [22], that generate an approximating parameter-separable expansion.

Results

In this section, we present the study of two example networks with the proposed model reduction method. With these examples, the applicability of the reduced modeling approach especially for analysis tasks requiring repeated simulations with different parameter values is illustrated. The first network is a bistable genetic toggle switch, where cells may switch randomly between two states, based on the model in [23]. For this network, the problem of parameter estimation with a reduced model is studied. The second network is a second-order genetic oscillator, based on [24], where we perform a sensitivity analysis over a wide parameter range.

Parameter estimation in a genetic toggle switch model

Network description

The genetic toggle switch considered here is an ovarian follicle switch model from [23]. It is a system of two genes which activate each other. The switch is modelled as a reaction network with two species X_1, X_2 , representing the gene products. The network reactions are specified in Table 1, and the network parameters in Table 2.

In [23], this network was shown to describe a bistable switch with two probability peaks, one close to $x^{(off)} = (0, 0)^T$ and the other close to $x^{(on)} = (V_1, V_2)^T$.

In the study [23], only the lower probability peak was of interest. Here, we are interested in the transition of the system from $x^{(off)}$ to $x^{(on)}$. Therefore, the system is truncated to a rectangle $\bar{\Omega} := \{0, \dots, 150\} \times \{0, \dots, 150\}$ such that $x^{(on)}, x^{(off)} \in \bar{\Omega}$, yielding a good approximation in the

Table 1 The follicle switch model

Reaction	Stoichiometry v_j	Propensity v_j
Production of X_1	$(1, 0)^T$	$u_1(k_1 + \frac{V_1 x_2^3}{M_1^2 + x_2^3})$
Degradation of X_1	$(-1, 0)^T$	$u_1 x_1$
Production of X_2	$(0, 1)^T$	$u_2(\frac{V_2 x_1^3}{M_2^2 + x_1^3})$
Degradation of X_2	$(0, -1)^T$	$u_2 x_2$

List of reactions and reaction propensity functions for the follicle switch model [23].

Table 2 Parameters for the follicle switch model

k_1	V_1	M_1	u_1	V_2	M_2	u_2
4	75	25	$0.01 \frac{1}{\text{min}}$	75	25	$0.01 \frac{1}{\text{min}}$

Parameter values for the follicle switch model in Table 1.

finite state projection to the infinite-dimensional chemical master equation.

The next step is to apply the decomposition of the matrix $A(\theta)$ as described in the methods section. Note that $A(\theta)$ for the switch network contains rational terms with the parameters M_1 and M_2 . Considering these two parameters as fixed quantities, the truncated CME for the follicle switch can be written as

$$\dot{P}(t) = (k_1 A^{[1]} + V_1 A^{[2]} + u_1 A^{[3]} + V_2 A^{[4]} + u_2 A^{[5]})P(t), \quad (22)$$

where $A^{[i]}$, $i = 1, \dots, 5$ are of dimension $151^2 \times 151^2 = 22801 \times 22801$.

As initial condition we choose a probability distributed over some lower states

$$p(0, x) = \begin{cases} \frac{1}{210} & \text{for } x_1 + x_2 \leq 20 \\ 0 & \text{otherwise.} \end{cases} \quad (23)$$

For the parametric model reduction, we consider only variations in the parameters u_1 and u_2 . These influence both the steady state level of gene activity in the on-state as well as the switching kinetics and are thus of high biological significance in the model. Hence we set $\theta := (u_1, u_2)^T \in [0.005, 0.02]^2$ as the parametric domain \mathcal{P} . As final time we choose $T = 10^7$ which corresponds to a time range of approximately 19 years, i.e. about three times the half-life time of the off-state estimated in [23].

Some state plots from the simulation of the full model are shown in Figure 1. These snapshots clearly show the transition of the switch from the off-state with low values for x_1 and x_2 to the on state with high values. The parameter influence is mainly reflected in the speed of the transition: for the parameter vector $(u_1, u_2) = (0.005, 0.02)$ in the lower row, most of the probability is already arranged around the on-state at the end of the simulation time. In contrast, for the parameter vector $(u_1, u_2) = (0.05, 0.005)$ in the upper row, a significant portion of the probability is still located around the off-state at this time point. Also, the transition paths are different: in the first case, the values for x_2 are lower than the values for x_1 during the transition, while in the second case, this relation is reversed.

As typical simulation time for a single trajectory of the full system, we obtain 98.2 seconds on a IBM Lenovo 2.53 GHz Dual Core Laptop.




Figure 1 Illustration of solution snapshots of the switch model. Illustration of some solution snapshots $P(t)$ of the switch model (22) for parameter vector $(u_1, u_2) = (0.05, 0.005)$ (upper row) and $(u_1, u_2) = (0.005, 0.02)$ (lower row) at times $t=0, 2 \cdot 10^5, 5 \cdot 10^6,$ and $1 \cdot 10^7$ from left to right.

Basis generation

We generated a reduced basis with the POD-Greedy algorithm, where the training set was chosen as the vertices of a mesh with 9^2 logarithmically equidistant parameter values over the parameter domain \mathcal{P} . We set $\varepsilon_{tol} = 10^{-12}$ as target accuracy. We use the projection error as error measure, hence precompute the 81 trajectories for construction of the reduced basis. As initial basis we set $N_0 = 1$ and $V_{N_0} := P_0$ using the parameter independent initial condition.

The POD-Greedy algorithm produces a basis of 33 vectors and the overall computation of the reduced basis takes 7.9 hours, the dominating computation time being spent in the error evaluations and POD computations. Some of the resulting orthonormal basis vectors are illustrated in Figure 2. The error decay curve and the selected

parameters in the parameter domain are illustrated in Figure 3. We nicely observe an exponential error decay, which indicates a parametric smoothness of the solution manifold, cf. the convergence rate statement given before for the POD-Greedy algorithm. The selected parameters seem to be located at the boundary of the parameter domain, indicating that the model behaviour in between can well be interpolated from the model behaviours along the boundary of the parameter domain.

The final reduced model of dimension 33 can then be simulated in 0.135 seconds, corresponding to a computational speedup factor of more than 700.

Parameter estimation

We exemplify a possible application of the reduced order model in parameter estimation, where we assume that a





Figure 2 Basis vectors for the switch model. Illustration of the first eight basis vectors for the switch model generated by the POD-Greedy algorithm.



distorted output $y(t)$ as the expected values $E[x_1]$ is available from population-averaged measurements. The task is to estimate the parameter values u_1 and u_2 from such a noisy measurement.

The reference parameter is $\theta_{ref} = (u_1, u_2) = (0.01, 0.01)^T$, and, for the purpose of this example, the measured output is produced by simulating the original model with the reference parameter values and adding 5% relative random white noise $n(t)$ sampled from a standard normal distribution, $y_{meas}(t) := y(t, \theta_{ref})(1 + 0.05n(t))$. An illustration of the reference output $y(t, \theta_{ref})$ and the noisy signal $y_{meas}(t)$ is given in the left of Figure 4.

We want to recover the values of the parameters u_1 and u_2 based on fitting the reduced parametric model's output $\hat{y}(t, \theta)$ to the measured output $y_{meas}(t)$. As is commonly done in parameter estimation, we formulate a least squares cost function as

$$J(\theta) = \int_0^T (y_{meas}(t) - \hat{y}(t, \theta))^2 dt, \quad (24)$$

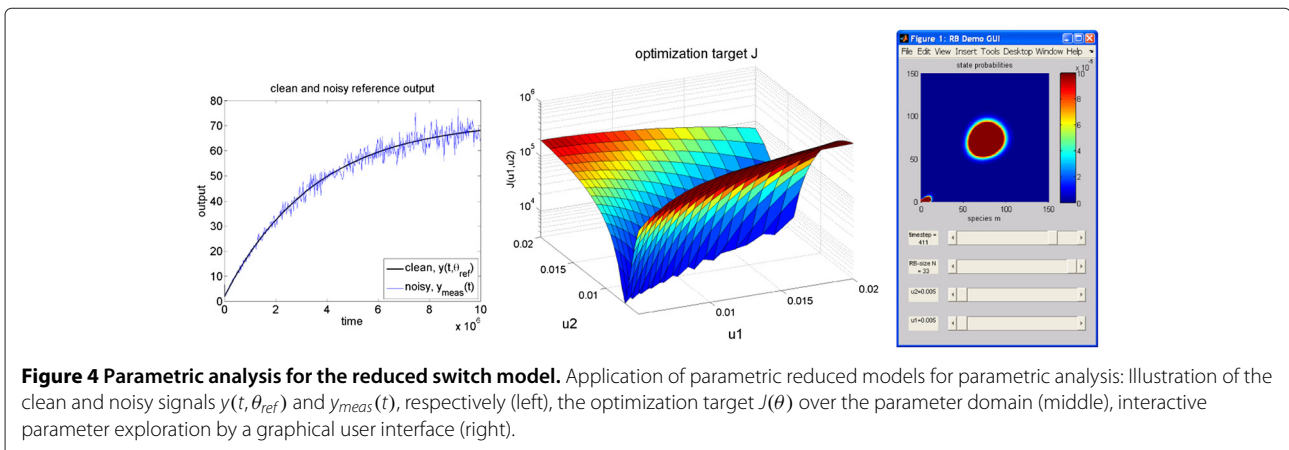
and estimate the parameters by

$$\theta_{est} = \arg \min_{\theta \in \mathcal{P}} J(\theta). \quad (25)$$

In such an optimization problem, typically many forward simulations are required for adjusting \hat{y} to the measurement. This is a particular beneficial scenario for reduced order models, as these simulations can be computed rapidly.

In order to gain a deeper insight into the optimization problem (25), we plot the values of the error functional $J(\theta)$ over the parameter domain (middle of Figure 4). Using the reduced model, the computation of the required $21^2 = 441$ trajectories is realized in less than one minute. This would be a significant computational effort when using a non-reduced model.

From the cost function plot, we observe a narrow area of parameters which seem to produce a similar output as the reference parameter θ_{ref} . This shows that the two model parameters are not simultaneously identifiable from the considered output, and indicates that there may exist a



functional dependence between the parameters u_1 and u_2 such that the model yields similar outputs $y(t)$.

Assuming a functional dependence of u_1 and u_2 we now consider the 1-dimensional optimization problem along the line $u_2 = u_{2,ref} = 0.01$. We would like to recover u_1 from the optimization problem. The corresponding value of the cost function is $J(\theta_{ref}) = 3330.68$, indicating a significant contribution of the noise. This restricted optimization problem is well conditioned and the optimization with a standard active set algorithm by MATLAB's command `fmincon` yields the estimated parameter $\theta_{est} := (u_{1,est}, 0.01)$ with $u_{1,est} = 0.0100204$, using 27 evaluations of the cost function. This accounts to a relative error in the u_1 value of 0.204%, hence excellent recovery. We refrain from plotting the recovered output $\hat{y}(t, \theta_{est})$ as it is visually indistinguishable from the output in the left of Figure 4. Interestingly, the optimization target value $J(\theta_{est}) = 3329.56$ implies $J(\theta_{est}) < J(\theta_{ref})$, which may stem from a slight approximation error in the reduced model or from the effects of the measurement noise.

The right plot in Figure 4 illustrates another application of reduced parametric models: We incorporated the model in an interactive graphical user interface in *RBmatlab*, a matlab package for model order reduction, available for download at www.morepas.org. This allows interactive parameter variations and instantaneous simulation response.

Sensitivity analysis in a stochastic oscillator

Network description

The second case study is built on a genetic oscillator model showing stochastic resonance, which was presented in [24]. The oscillator is based on a negative feedback loop between two genes with one gene having positive autoregulation. The oscillator is modelled as a reaction network with two species X_1, X_2 , representing the gene products. The network reactions are specified in Table 3, with parameters as in Table 4. In the original model in [24], the dynamics were described as stochastic differential equation for the amounts of X_1 and X_2 , coming from a Langevin approximation to the stochastic dynamics [12]. For the framework used in this paper, the dynamics have to be described directly by the underlying

Table 3 The oscillator model

Reaction	Stoichiometry v_j	Propensity v_j
Production of X_1	$(1, 0)^T$	$\frac{k_1 s^2}{k_2 s + x_2}$
Degradation of X_1	$(-1, 0)^T$	$k_3 x_1$
Production of X_2	$(0, 1)^T$	$k_4 s + \frac{k_5 x_2^2 x_1}{k_6 s^2 + x_2^2}$
Degradation of X_2	$(0, -1)^T$	$k_7 x_2$

List of reactions and reaction propensity functions for the oscillator model adopted from [24].

Table 4 Parameters for the oscillator model

k_1	k_2	k_3	k_4	k_5	k_6	k_7	s
$15 \frac{1}{s}$	0.2	$1 \frac{1}{s}$	$10 \frac{1}{s}$	$100 \frac{1}{s}$	6.5	$100 \frac{1}{s}$	10

Parameter values for the oscillator model from Table 3.

CME. To achieve this, we introduce the parameter s which maps the dimensionless state variables from [24] to actual molecule numbers as required for the CME. Thus, s is also a measure for the network's noise level: the higher s , the larger the molecule number that is considered, and the smaller the noise level will be.

The network model in Table 3 shows oscillations only in a stochastic description. The deterministic model has a unique asymptotically stable equilibrium point, but in a stochastic model, fluctuations may push the molecular numbers beyond a certain threshold, inducing a dynamical response along a slow manifold, which corresponds to one oscillatory period [24]. Depending on the noise level, such responses will be initiated more or less often, corresponding to a more or less regular oscillatory pattern.

The system is truncated to the rectangle $\tilde{\Omega} := \{0, \dots, 300\} \times \{0, \dots, 300\}$, which contains the relevant system states for the parameter ranges of interest.

Similarly as in the switch example, the reaction propensity expressions contain rational terms in the parameters s, k_2 , and k_6 . These three cannot be decomposed directly, so we do the decomposition described in the methods section for the other five parameters only. With this decomposition, the truncated CME for the genetic oscillator can be written as

$$\dot{P}(t) = \left(k_1 A^{[1]} + k_3 A^{[2]} + k_4 A^{[3]} + k_5 A^{[4]} + k_7 A^{[5]} \right) P(t), \quad (26)$$

where $A^{[i]}, i = 1, \dots, 5$ are of dimension $301^2 \times 301^2 = 90601 \times 90601$. The initial condition for (26) is chosen as a uniform distribution over the rectangle $\{0, \dots, 50\} \times \{0, \dots, 50\}$:

$$p(0, x) = \begin{cases} \frac{1}{51^2} & \text{for } x_1 \leq 50, x_2 \leq 50 \\ 0 & \text{otherwise.} \end{cases} \quad (27)$$

The time scale of interest for the model in (26) is for $0 \leq t \leq T = 6$. At the end of the interval, the probability distribution seems to approach a steady state.

Some state plots are given in Figure 5. One observes a significant effect of the parameter k_4 on the amplitude of the oscillations. The simulation time for the detailed model was in average 7.3 minutes on a Dell desktop computer with 3.2 GHz dual-core Intel 4 processor and 1 GB RAM, without including the computation time for the construction of the state transition matrix $A(\theta)$.




Figure 5 Illustration of solution snapshots of the oscillator model. Illustration of some solution snapshots $P(t)$ of the oscillator CME model for parameter values $k_4 = 15$ (upper row) and $k_4 = 30$ (lower row) at times $t = 0, 0.2, 0.6, 6.0$ from left to right.

Basis generation

For the basis generation, the parameter k_4 was assumed to vary within the interval $[10, 100]$. A reduced basis with the POD-Greedy algorithm was computed from a training set of 30 logarithmically equidistant parameters over the parameter domain (Figure 6). As in the switch example, the target accuracy was chosen as $\varepsilon_{tol} = 10^{-12}$, and the initial basis was chosen from the initial condition $V_1 := P_0$.

The POD-Greedy algorithm produces a basis of 109 vectors, with an overall computation time of 16.5 hours on the hardware as in the previous subsection. The first




Figure 6 Parametric analysis results for the oscillator model. Sensitivity analysis of oscillation amplitude over a parameter k_4 interval. Blue line shows oscillatory amplitude over the parameter k_4 predicted from the reduced model. Red dots are validation results from a simulation of the original model. Triangles on the parameter axis indicate parameter values which were used in the construction of the reduced basis.

20 basis vectors are shown in Figure 7. It is apparent that several of the basis vectors are directly included in order to reproduce the different amplitudes of oscillations that will occur under variations of the parameter k_4 . The error decay curve is shown in Figure 8, displaying an exponential error decay as also observed for the switch example.

With the reduced basis $V \in \mathbb{R}^{90601 \times 109}$, we can construct a reduced parametric model for the CME of the oscillator as

$$\begin{aligned} \dot{P}_r(t) &= (k_4 A_r^{[3]} + A_r^{[0]}) P_r(t) \\ P_r(0) &= V^T P(0), \end{aligned} \quad (28)$$

with $A_r^{[3]} = V^T A^{[3]} V \in \mathbb{R}^{109 \times 109}$ and $A_r^{[0]} = V^T (k_1 A^{[1]} + k_3 A^{[2]} + k_5 A^{[4]} + k_7 A^{[5]}) V \in \mathbb{R}^{109 \times 109}$. Note that since only k_4 has been varied in the reduction process, the other parameters are no longer present as parameters in the reduced model, but just take their nominal values. While the same basis V could be used to construct another reduced model where all parameters are retained, it is unlikely that this other model will be a good approximation of the original one for varying values of the other parameters.

Sensitivity analysis of the oscillation amplitude

As an application of the reduced order parametric model obtained in the previous section, we study the variations of oscillatory amplitude over a parameter range. Specifically, we consider 200 equally spaced values for the parameter k_4 in the interval $[12, 40]$ and compute the probability that the amount of X_2 is larger than 100:

$$\text{Prob}(x_2 > 100) = \sum_{x: x_2 > 100} p(T, x), \quad (29)$$




Figure 7 Basis vectors for the oscillator model. First 20 basis vectors for the oscillator model.

with $T = 6$ the final time of the simulation. The results are shown in Figure 6 and show a clear decay of oscillatory amplitude for increasing values of k_4 . Due to the significant time savings from the reduced model, this sensitivity curve can be computed with a high resolution.

To evaluate the quality of the reduced model, we also computed the probability (29) using the original model (26) at two points within the considered interval for the

parameter k_4 . As shown in Figure 6, the results from the original model are in perfect agreement with the predictions from the reduced model at these points. Since the points at which the original model was evaluated in this experiment were not part of the training set (shown as triangles on the parameter axis in Figure 6), this shows that it is in fact possible to extrapolate the reduced model to parameter values that were not used to construct the basis.




Figure 8 Results of the POD-Greedy algorithm for the oscillator model. Error decay curve for the oscillator model.

Conclusions

In this paper, we have introduced the application of parametric model reduction methods to finite-state approximations of the chemical master equation. We have also presented two case studies where these methods are applied to CME models of different networks in order to make parametric analysis tasks computationally efficient. By this, it has become clear that parametric model reduction methods are a very useful tool for the analysis of stochastic biochemical reaction network described by the CME.

Especially analysis tasks where many repeated simulations of a network with different parameter values are required can profit significantly from parametric model reduction. This includes for example sensitivity analysis or parameter optimization tasks such as identifiability analysis or estimation. Moreover, the significant speedup of

the simulation for the reduced model allows an interactive exploration of the network's dynamics within the parameter space within a suitable graphical user interface.

This contribution is just a first step in the application of parametric model reduction methods to the CME. One particularly important aspect that we have not discussed here is the computation of error estimates for certifying that the simulation output of the reduced model is within some tolerance of the corresponding simulation output of the original model. To maintain computational efficiency, the error estimation should be done without actually simulating the original model. Error estimation methods have been developed for parametric model reduction of generic models [9], but tighter estimates could likely be obtained by taking into account the special structure of the CME models. Recent work for example refined the previous generic error bounds for stable models [25].

Competing interests

Both authors declare that they have no competing interests.

Authors contributions

SW and BH conceived of the study, performed the study, and wrote the manuscript. Both authors read and approved the final manuscript.

Acknowledgements

We thank Wolfgang Halter for programming support in the oscillator case study. The authors would like to thank the German Research Foundation (DFG) for financial support of the project within the Cluster of Excellence in Simulation Technology at the University of Stuttgart. BH also acknowledges the Baden-Württemberg Stiftung gGmbH for funding. This work was also supported by the German Research Foundation (DFG) within the funding programme Open Access Publishing.

Author details

¹Institute for Systems Theory and Automatic Control, University of Stuttgart, Pfaffenwaldring 9, Stuttgart, Germany. ²Institute for Applied Analysis and Numerical Simulation, University of Stuttgart, Pfaffenwaldring 57, Stuttgart, Germany.

Received: 3 January 2012 Accepted: 18 May 2012

Published: 2 July 2012

References

1. Gillespie DT: **A rigorous derivation of the chemical master equation.** *Physica A: Statist Theor Phys* 1992, **188**(1-3):404–425. [http://www.sciencedirect.com/science/article/B6TVG-46FX396-7N/2/a0537c1efc0f5c330fa05b5e4ae61b98]
2. van Kampen NG: *Stochastic Processes in Physics and Chemistry.* Amsterdam, The Netherlands: North-Holland; 1981.
3. Munsky B, Khammash M: **The finite state projection algorithm for the solution of the chemical master equation.** *J Chem Phys* 2006, **124**(4):044104. [http://dx.doi.org/10.1063/1.2145882].
4. Jahnke T, Huisinga W: **A Dynamical Low-Rank Approach to the Chemical Master Equation.** *Bull Math Biol* 2008, **70**:2283–2302.
5. Hegland M, Hellander A, Lötstedt P: **Sparse grids and hybrid methods for the chemical master equation.** *BIT Numerical Mathematics* 2008, **48**:265–283.
6. Antoulas AC: *Approximation of Large-Scale Dynamical Systems.* Philadelphia, USA: SIAM; 2005.
7. Munsky B, Khammash M: **The Finite State Projection Approach for the Analysis of Stochastic Noise in Gene Networks.** *Automatic Control, IEEE Transactions on* 2008, **53**(Special Issue):201–214.
8. Baur U, Benner P: **Parametrische Modellreduktion mit dünnen Gittern.** In *GMA-Fachausschuss 1.30, Modellbildung, Identifizierung und*

- Simulation in der Automatisierungstechnik, Salzburg ISBN 978-3-9502451-3-4*; 2008:262–271.
9. Haasdonk B, Ohlberger M: **Efficient Reduced Models and A-Posteriori Error Estimation for Parametrized Dynamical Systems by Offline/Online Decomposition.** *MCMDs, Mathematical and Computer Modelling of Dynamical Systems* 2011, **17**(2):145–161.
 10. Daniel L, Siong O, Chay L, Lee K, White J: **Multi-parameter moment-matching model-reduction approach for generating geometrically parameterized interconnect performance models.** *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 2004, **23**(5):678–693.
 11. Moosmann C, Rudnyi E, Greiner A, Korvink J, Hornung M: **Parameter Preserving Model Order Reduction of a Flow Meter.** In *Technical Proceedings of Nanotech 2005*; 2005.
 12. Higham DJ: **Modeling and Simulating Chemical Reactions.** *SIAM Rev* 2008, **50**(2):347–368. [http://link.ajp.org/link/SIR/50/347/1].
 13. Gillespie DT: **Exact stochastic simulation of coupled chemical reactions.** *J Phys Chem* 1977, **81**(25):2340–2361. [http://pubs.acs.org/cgi-bin/abstract.cgi/jpchax/1977/81/i25/f-pdf/f_j100540a008.pdf]
 14. Haasdonk B, Ohlberger M: **Reduced Basis Method for Finite Volume Approximations of Parametrized Linear Evolution Equations.** *M2AN, Math Model Numer Anal* 2008, **42**(2):277–302.
 15. Eftang JL, Knezevic DJ, Patera AT: **An hp Certified Reduced Basis Method for Parametrized Parabolic Partial Differential Equations.** *MCMDs, Mathematical and Computer Modelling of Dynamical Systems* 2011, **17**(4):395–422.
 16. Knezevic D, Patera A: **A Certified Reduced Basis Method for the Fokker-Planck Equation of Dilute Polymeric Fluids: FENE Dumbbells in Extensional Flow.** *SIAM Journal of Scientific Computing* 2010, **32**(2):793–817.
 17. Volkwein S: **Model Reduction using Proper Orthogonal Decomposition.** 2011. [http://www.uni-graz.at/imawwww/volkwein/publist.html]. [Lecture Notes, University of Constance]
 18. Jolliffe IT: *Principal Component Analysis.* New York, USA: Springer-Verlag; 2002.
 19. Haasdonk B: **Convergence Rates of the POD-Greedy Method.** 2011. Simtech preprint 2011-23, University of Stuttgart, Germany.
 20. Haasdonk B, Dihlmann M, Ohlberger M: **A Training Set and Multiple Bases Generation Approach for Parametrized Model Reduction Based on Adaptive Grids in Parameter Space.** *MCMDs, Mathematical and Computer Modelling of Dynamical Systems* 2011, **17**(4):423–442.
 21. Wu J, Vidakovic B, Voit EO: **Constructing stochastic models from deterministic process equations by propensity adjustment.** *BMC Syst Biol* 2011, **5**:187. [http://dx.doi.org/10.1186/1752-0509-5-187]
 22. Barrault M, Maday Y, Nguyen N, Patera A: **An 'empirical interpolation' method: application to efficient reduced-basis discretization of partial differential equations.** *C R Math Acad Sci Paris Series I* 2004, **339**:667–672.
 23. Waldherr S, Wu J, Allgöwer F: **Bridging time scales in cellular decision making with a stochastic bistable switch.** *BMC Syst Biol* 2010, **4**:108. [http://www.biomedcentral.com/1752-0509/4/108]
 24. El-Samad H, Khammash M: **Coherence resonance: a mechanism for noise induced stable oscillations in gene regulatory networks.** In *Proc. of the 45th Conf. Dec. Contr. (CDC).* San Diego, USA; 2006:2382–2387.
 25. Hasenauer J, Löhning M, Khammash M, Allgöwer F: **Dynamical optimization using reduced order models: A method to guarantee performance** 2012. [Journal of Process Control, Online Publication before print]

doi:10.1186/1752-0509-6-81

Cite this article as: Waldherr and Haasdonk: Efficient parametric analysis of the chemical master equation through model order reduction. *BMC Systems Biology* 2012 **6**:81.