ADAPTIVE DISCRETE GALERKIN METHODS APPLIED TO THE CHEMICAL MASTER EQUATION[*]

P. DEUFLHARD[†], W. HUISINGA[‡], T. JAHNKE[§], AND M. WULKOW[¶]

Abstract. In systems biology, the stochastic description of biochemical reaction kinetics is increasingly being employed to model gene regulatory networks and signalling pathways. Mathematically speaking, such models require the numerical solution of the underlying evolution equation, known as the chemical master equation (CME). Up to now, the CME has primarily been treated by Monte-Carlo techniques, the most prominent of which is the stochastic simulation algorithm (Gillespie 1976). The paper presents an alternative, which focuses on the discrete partial differential equation (PDE) structure of the CME. This allows to adopt ideas from adaptive discrete Galerkin methods as first suggested in (Deuflhard, Wulkow 1989) for polyreaction kinetics and independently developed in (Engblom 2006). Among the two different options of discretizing the CME as a discrete PDE, Engblom had chosen the method of lines approach (first space, then time), whereas we strongly advocate to use the Rothe method (first time, then space) for clear theoretical and algorithmic reasons. Numerical findings at two rather challenging problems illustrate the promising features of the proposed method and, at the same time, indicate lines of necessary further improvement of the method worked out here.

Key words. adaptive discrete Galerkin methods, adaptive Rothe method, discrete Chebyshev polynomials, stochastic reaction kinetics, chemical master equation

AMS subject classifications. 41A10, 60H35, 92C45, 65L60

Introduction. Stochastic effects are of crucial importance in the understanding of gene regulatory networks and signalling cascades [15, 37, 32, 33]. Based on this insight, there is a growing interest in stochastic modelling approaches to chemical reaction kinetics. The design of efficient numerical techniques and reliable approximations schemes for the solution of the chemical master equation (CME) is now an active field of research. In almost all approaches, the CME is solved by Monte-Carlo techniques, generating a statistically large ensemble of realizations of the associated continuous–time / discrete–state–space Markov jump process. The most prominent of these approaches is the stochastic simulation algorithm due to Gillespie in [20, 21]. This algorithm, however, requires an update of the system at each time when one of the reaction channels fires, thus causing enormous computational cost in the case of highly reactive systems. In order to gain relevant information (such as, e.g., expectations), a sufficiently large number of realizations has to be computed, which multiplies the computational work. Therefore, research has mainly concentrated on how to improve the efficiency of this kind of algorithm (cf. [19, 23, 6, 36, 1, 25, 7, 38]).

As an alternative to Monte-Carlo approaches, there is a growing interest in directly solving the CME, i.e., approximating the probability distribution as a function of time and space [35, 34, 5, 31, 16, 28]. After truncation of the state space to a sufficiently large finite domain, the original CME is reduced to a finite dimensional linear system of ordinary differential equations (ODEs). Since the state space is typically

still too huge for standard approaches to be applicable, efficient numerical techniques to approximate the matrix exponential, like Krylov-based or sparse grid techniques, are used.

The present paper deals with an alternative idea based on a change of perspective. In fact, from a mathematical point of view, the CME may be understood as a countable system of ordinary differential equations (CODE) or, equivalently, as a discrete partial differential equation (PDE). For its numerical treatment, the main challenge is the fact that each single state of the state space corresponds to one degree of freedom in the CME. Even in case of a small system containing, e.g., only three molecular species and at most 10^2 copy numbers per species, a total number of $(10^2)^3 = 10^6$ coupled ordinary differential equations (ODEs) has to be solved.

This situation is comparable to the one in polyreaction kinetics. There, too, single molecules (monomers) are linked together to long chains whose dynamical mathematical description gives rise to huge numbers of ODEs, say up to $10^4 - 10^6$ in realistic examples, sometimes even not known in advance. Two of the present authors have a long standing expertise in modelling and efficient simulation of polymer kinetics [13, 42, 45, 43, 27, 14, 44]. This has resulted in comprehensive numerical solvers and user-friendly software tools (e.g. Predici, Parsival), which serve as basis and often as reference for modelling applications in more than hundred publications in this field [8, 44]; for review articles see, e.g., [46, 29]. However, as a distinguishing feature, polymerization systems possess a natural coordinate, the polymer chain length. For this reason, the techniques used in polyreaction kinetics will need some thorough reconsideration in the present context. A first step in this direction has been performed recently by Engblom [16]. He has worked out a method of lines (MOL) approach, i.e. first state space approximation, then time discretization, so that one is left with a system of ordinary differential equations (ODEs). In what follows, we propose an approach in the frame of a Rothe method (ROM), i.e. first time discretization, then state space approximation, which leads to a sequence of stationary problems and thus conveniently allows for adaptively chosen state subspaces. This so-called adaptive Rothe method has been introduced by Bornemann [4, 2] in 1990 for scalar parabolic equations and extended by Lang to parabolic systems of reaction-diffusion type up to challenging real life problems [30]. Already in 1992, Wulkow had transferred the approach to polyreaction kinetics in his thesis [42, 43]. The ROM will be combined with a Galerkin h-p-method in higher dimensions which already has been successfully applied to similar problems with continuous property coordinates [26].

The article is organized as follows. In Section 1, we first set the scene and introduce the chemical master equation focusing on its discrete PDE structure. Moreover, we recall the two possible discretization options for this discrete PDE, the MOL and the ROM, and discuss their suitability to allow for state space adaptivity. On this basis, in Section 2, we work out details of discrete Galerkin methods for the CME. Adaptive Galerkin methods in one dimension are recalled and extended via a tensor product approach to multiple dimensions. The arising reduced linear systems require the efficient numerical evaluation of matrix elements and right-hand sides. In this context, Gauss-Christoffel summation plays an important role. Finally, in Section 3, we apply our approach to two challenging model problems exhibiting typical features of stochastic reaction kinetics of variable degree.

1. Preliminary Considerations. In this section, we first formulate the chemical master equation (CME) and reveal its structure as a discrete partial differential equation (PDE). From this perspective, we then discuss two options concerning the

order of time discretization and state space approximation in view of adaptivity.

1.1. Chemical Master Equation as Discrete PDE. Consider a well-mixed system of volume V with d chemical species $S_1, \ldots, S_d$ involved in M reactions $R_1, \ldots, R_M$. The state of the system is characterized by a vector

$$X(t) = (X_1(t), \ldots, X_d(t)) \in \mathbb{N}^d$$

with entries $X_i(t)$ denoting the numbers of molecules of S_i. (Here and below, $\mathbb{N}$ is the set of nonnegative integers.) In the stochastic modelling approach treated here, the *discrete* vector $X(t)$ is understood to be *random*. For a brief introduction of the stochastic modelling approach in the context of chemical reaction kinetics see, e.g., [20, 22], for a general introduction to Markov processes see, e.g., [41]. The reaction probability for each reaction R_j is specified by the *propensity function* $\alpha_j = \alpha_j(X(t), t)$, which typically equals the product of a rate constant c_j and the number of possible combinations of reactant molecules involved in reaction R_j. The most frequently arising reaction types are listed in Table 1.1.

chemical reaction			α_j
S_k	$\rightarrow$	S_i	$c_j X_k(t)$
$S_k + S_l$	$\rightarrow$	S_i	$c_j X_k(t) X_l(t)$
$S_k + S_l$	$\rightarrow$	S_i	$c_j X_k(t)(X_k(t) - 1)/2$

TABLE 1.1
Stochastic propensity functions

Once a reaction R_j takes place, the number of molecules for each species changes according to the stoichiometric vector $\nu_j \in \mathbb{N}^d$, i.e., $X(t) \rightarrow X(t) + \nu_j$. Note that the first two propensity terms above have the same form as in the continuous case, while the third one clearly reveals the discrete nature of the process.

The time evolution of the *probability distribution function* (PDF)

$$p(t, x) = \mathrm{P}\big[X_1(t) = x_1, \ldots, X_d(t) = x_n\big], \quad x \in \mathbb{N}^d$$

of the random vector $X(t)$ is given by the *chemical master equation* (CME)

$$\partial_t p(t, x) = \sum_{m=1}^{M} \big[\alpha_m(x - \nu_m) p(t, x - \nu_m) - \alpha_m(x) p(t, x)\big], \qquad (1.1)$$

see, e.g., [20, 22, 17, 41]. If the argument $(x - \nu_m)$ contains negative entries, the corresponding term has to be omitted because the PDF $p(t, \cdot)$ is only defined on $\mathbb{N}^d$. Hence, for convenience we set

$$\alpha_m(x) = 0 \text{ and } p(t, x) = 0 \quad \text{for all} \quad x \notin \mathbb{N}^d.$$

Upon defining the infinitesimal generator

$$\big(\mathcal{A}p(t, \cdot)\big)(x) = \sum_{m=1}^{M} \big[\alpha_m(x - \nu_m) p(t, x - \nu_m) - \alpha_m(x) p(t, x)\big], \qquad (1.2)$$

4

the CME with initial condition $p(0) = \phi$ can be rewritten in operator form as the *abstract Cauchy problem*

$$\partial_t p(t) = \mathcal{A}p(t), \qquad p(0) = \phi \qquad (1.3)$$

where $p(t) = p(t, \cdot)$. We note that (1.1) or (1.3) are *linear* equations, which will be useful for the construction of algorithms.

In mathematical terms, the CME may be understood either as a countable system of ODEs (often abbreviated as CODEs) or, equivalently, as a discrete partial differential equation (PDE), wherein the continuous derivatives are replaced by discrete differences. The theory of CODEs strongly indicates that CODEs are structurally quite different from ODE systems, even of high dimension, but quite close to PDEs [9, 10].

Upon adopting the discrete PDE point of view, we are naturally led to consider adaptive discrete Galerkin methods as first introduced in [13, 45] and further pursued in [42, 43]. Surveys on the development of the topic since then can be found in [44, 14]. However, the numerical challenge in solving the CME are its many degrees of freedom, even for reaction systems with a relatively small number of species. In Section 3.2 we give an example of a reaction system with just two chemical species whose total number of molecules is confined to 10^4, resulting in an CME with 10^8 degrees of freedom. Therefore, modifications of the original ideas of discrete Galerkin methods will be necessary in the present context.

1.2. Integrating the CME. Assume that the exact CME solution p is contained in some Hilbert space $\mathcal{H}$ with inner product $\langle \cdot, \cdot \rangle$. In view of the theory [43] we aim at a space with bounded statistical moments and $\mathcal{H} \hookrightarrow l^2$ (see also Section 2.2). Then $p \in \mathcal{H}$ can be represented in terms of basis functions $\{q_k\}$, which span $\mathcal{H}$, so that

$$p(t, x) = \sum_{k=0}^{\infty} a_k(t) q_k(x). \qquad (1.4)$$

Upon truncating this expansion, we naturally arrive at some approximation $p_r(t) \in \mathcal{H}_r \subset \mathcal{H}$ to $p(t)$ called the *Galerkin approximation*

$$p_r(t, x) = \sum_{k=0}^{r} a_k^{[r]}(t) q_k(x). \qquad (1.5)$$

In special cases, one can even guarantee that $a_k^{[r]} = a_k$, i.e. that the coefficients are independent of the truncation index. For simplicity, the superscript $[r]$ will be omitted. Crucial details to be set in this approximation frame are:
- choice of basis functions $\{q_k\}$,
- choice of truncation index r,
- computation of expansion coefficients $\{a_k^{[r]}(t)\}$ via the discrete PDE (1.3).

In this section, we will only discuss the last two points. The first point is left to the subsequent Section 2.

Method of lines (MOL): first space, then time. This approach is the most popular one to tackle PDEs. After choosing a finite dimensional subspace $\mathcal{H}_r \subset \mathcal{H}$, we are left with a *finite dimensional ODE*. Given a fixed truncation subspace $\mathcal{H}_r$, we insert

the corresponding representation (1.5) into the CME (1.1) and take the inner product with a basis function in $\mathcal{H}_r$:

$$\left\langle \sum_{k=0}^{r} q_k \partial_t a_k(t),\ q_l \right\rangle = \left\langle \sum_{k=0}^{r} (\mathcal{A} q_k) a_k(t),\ q_l \right\rangle.$$

This leads us to evolution equations for the expansion coefficients $a(t) = (a_k(t))$:

$$\Gamma_r \frac{\mathrm{d}}{\mathrm{d}t}\, a(t) = A\, a(t)$$

in terms of the "mass" matrix $\Gamma_r = (\langle q_k, q_l \rangle) \in \mathbb{R}^{r \times r}$ and the "stiffness" matrix $A = (\langle \mathcal{A} q_k, q_l \rangle) \in \mathbb{R}^{r \times r}$ representing the Galerkin approximation of the generator $\mathcal{A}$. For further details see, e.g., [13].

In chemical reaction kinetics, the resulting ODEs are typically stiff and must therefore be solved by some implicit numerical integration scheme. For example, the implicit Euler method for step size τ leads to a linear system of the kind

$$(\Gamma_r - \tau A)\Delta\eta_0 = \tau A\, \eta_0, \tag{1.6}$$
$$\eta_1 = \eta_0 + \Delta\eta_0,$$

where the vectors $\eta_i \in \mathbb{R}^r$ contain the Galerkin coefficients, and $\Delta\eta_0 = \eta_1 - \eta_0$ is the corresponding difference. Note that the given initial values η_0 correspond to some $p_r(t) \in \mathcal{H}_r$.

Rothe method (ROM): first time, then space. This approach has been introduced and theoretically analyzed in [4, 2] and immediately transferred to polyreaction kinetics in [42, 43]. After discretization in time first, we are left with a sequence of *stationary problems*.

For linear problems like the CME considered here, Bornemann [3] had even designed a special time integration scheme. This scheme incorporates an easily accessible *temporal error estimate*. Starting from $u_0 = p(0)$, a second order *temporal* approximation u_n of the CME solution $p(t_n)$ at $t_n = t_0 + n\tau$ ($n = 1, 2, ...$) is obtained recursively by first performing one step of the implicit Euler method

$$(\mathcal{I} - \tau\mathcal{A})\Delta u_n^{(1)} = \tau\mathcal{A}\, u_n, \tag{1.7}$$
$$u_{n+1}^{(1)} = u_n + \Delta u_n^{(1)},$$

followed by the second order correction step

$$(\mathcal{I} - \tau\mathcal{A})\Delta u_n^{(2)} = -\frac{\tau}{2}\mathcal{A}\, \Delta u_n^{(1)}, \tag{1.8}$$
$$u_{n+1} = u_{n+1}^{(1)} + \Delta u_n^{(2)}.$$

In contrast to the MOL approach, the temporal approximations are still defined on the state space $\mathbb{N}^d$. As add-on, we may gain the temporal error estimate

$$\mathrm{eps}_T = \|\Delta u_1\|, \tag{1.9}$$

defined in some suitable norm. A new time step τ_{new} is then proposed on the basis of an old time step τ by the usual formula

$$\tau_{\mathrm{new}} = \sqrt{\rho \frac{\mathrm{TOL}}{\mathrm{eps}_T}}\, \tau, \tag{1.10}$$

wherein $\rho < 1$ is an additional safety factor. All arising norms of state space terms can be approximated to prescribed accuracy within the Galerkin setting. Suppose we are given some Galerkin subspace $\mathcal{H}_r$; then we may compute an associated spatial error estimate (details to be worked out in the Section 2 below, see eq. (2.4))

$$\|p_r - p\| \le \text{TOL}_p. \tag{1.11}$$

Here TOL_p is an imposed spatial error tolerance. Following an argument by Bornemann in [3], this error tolerance is linked to the user-specified tolerance TOL by

$$\text{TOL}_p = \frac{1}{32}\text{TOL} .$$

State space adaptivity. By construction, the subspace $\mathcal{H}_r$ has to be kept unchanged over the integration step in the MOL approach. A change of the subspace after an integration step is possible, but a rather subtle task requiring a careful control of interpolation errors. Moreover, earlier computations in polyreaction kinetics have confirmed the undesirable experience that, already after one integration step, the approximation tends to "leave" the preassigned subspace $\mathcal{H}_r$ with the consequence of either extremely small time steps or oscillatory numerical artifacts. For this reason, a "moving weight function" concept has been introduced in [13] which results in adaptive *time dependent* basis functions $\{q_k(t)\}$. A comparable approach has recently been introduced by Engblom [16] without knowing about the earlier more general results in [13]. However, on the basis of numerical experience and theoretical investigations in [42], the MOL approach as a whole has turned out to be not sufficiently robust for complex polyreaction kinetics. There are several reasons for this. The change of the approximating subspace is prohibitively complicated in the MOL approach. Furthermore, most of the operators cannot be preprocessed analytically, but require numerical summation. In a MOL this introduces non-smooth right-hand sides leading to very small and erratic time steps in sophisticated ODE solvers. In fact, the whole development of new methods for countable systems got stuck by this obstacles before the ROM had been introduced [44].

As already mentioned above, the ROM permits an easy adaptation of Galerkin subspaces for the solution of the arising stationary value problems (1.7) or (1.8). Moreover, this method is also clearly preferable for theoretical reasons that are elaborated in [4, 42]. On this basis, we suggest this approach also for the present context of the CME.

2. Realization of Discrete Galerkin Methods. In this section, we work out details on the choice of basis functions $\{q_k\}$. Given an inner product associated with the Hilbert spaces $\mathcal{H}_r \subset \mathcal{H}$, which may be weighted or unweighted, we choose as a basis an orthogonal systems of polynomials of discrete variables. For this purpose, we partly recall material already given in [13, 42] and modify it for use in the present CME context. In the next subsection we will introduce the 1d case, leaving the multiple dimensional case to the subsequent Section 2.3.

2.1. Global Techniques in 1d Revisited. In [13], a general setup has been defined in terms of a weight function $\psi(x\,;\rho)$ with $x \in \mathbb{N}$. The parameter ρ characterizes the weight function and can be time dependent, as suggested in the so-called "moving weight function" concept [13]. This weight function gives rise to the inner product

$$\langle u, v \rangle_\psi = \sum_{x \in \mathbb{N}^d} u(x)v(x)\psi(x\,;\rho). \tag{2.1}$$

inducing some weighted norm $\|u\|_\psi^2 = \langle u, u \rangle_\psi$ thus defining some discrete Hilbert space (sequence space)

$$\mathcal{H}_\psi = \{ u : \mathbb{N} \to \mathbb{R} : \|u\|_\psi < \infty \} \, . \qquad (2.2)$$

As a basis in this Hilbert space, orthogonal polynomials $\{q_k\}$ are defined such that

$$\langle q_k, q_l \rangle_\psi = \gamma_k \delta_{kl} \, , \quad \gamma_k > 0,$$

where δ_{kl} denotes the Kronecker symbol. Numerical experience exhibited that the more sophisticated choices of weight functions (cf. [13, 16, 45, 43, 16]), including the moving weight function concept, lead to global discrete Galerkin methods that are well-suited for monomodal solutions, but less appropriate to bimodal functions, which, however, do arise in important applications. For that reason, Wulkow [43] turned to localized unweighted polynomials, *discrete Chebychev polynomials* formally corresponding to the weight function

$$\psi(x) \equiv 1 \, . \qquad (2.3)$$

It is this latter setting which paved the way to the extreme success of discrete Galerkin methods in polymer chemistry. On this basis, we will suggest the same choice for the CME context—to be treated in the next section for more than one spatial dimension.

Spatial error estimation. Recall the type of global expansion (1.4) for the exact solution p and (1.5) for the Galerkin approximation p_r. Of course, we will choose the induced Hilbert space norm to get an error estimate $\|p_r - p\|_\psi$. If the coefficients $\{a_k\}$ are independent of the truncation index r, we obtain the exact expression

$$\|p_r - p\|_\psi = \left\| \sum_{k=r+1}^\infty a_k q_k(x; \rho) \right\|_\psi = \left(\sum_{k=r+1}^\infty |a_k|^2 \gamma_k \right)^{1/2} . \qquad (2.4)$$

Otherwise, which is the typical case, a straightforward modification can be applied, see [13]. Since, in this norm, the expansion coefficients tend to zero for $k \to \infty$, the spatial error estimate

$$\mathrm{eps}_r = \|p_{r+1} - p_r\|_\psi = |a_{r+1}| \sqrt{\gamma_{r+1}}$$

may serve the purpose. Clearly, this error estimate will work fine whenever the expansion coefficients decrease sufficiently fast. Otherwise, the expansion has to be continued a few steps beyond the step $r + 1$.

Algebraic Galerkin equations. At the core of discrete Galerkin methods, algebraic equations must be constructed. For simplicity, we choose the implicit Euler operator equation (1.7) for the first order approximation of $u_n^{(1)} = u_n + \Delta u_n^{(1)}$. Recall that the increment $\Delta u_n^{(1)}$ is defined on the entire state space $\mathbb{N}^d$. Denote by $\Delta \eta_n \in \mathcal{H}_r$ the Galerkin approximation of $\Delta u_n^{(1)}$, defined in terms of the coefficients $\Delta \eta_n = ((\Delta \eta_n)_1, \dots, (\Delta \eta_n)_r)$. Upon applying inner products with each orthogonal polynomial, we arrive at the formal relations

$$\left\langle \sum_{k=0}^r (q_k - \tau \mathcal{A} q_k)(\Delta \eta_n)_k, q_l \right\rangle_\psi = \langle \tau \mathcal{A} u_0, q_l \rangle_\psi , \qquad l = 0, \dots, r. \qquad (2.5)$$

Due to the orthogonality of the polynomial basis, this yields an algebraic system of the kind (1.6) already shown in the context of the MOL, i.e.

$$(\Gamma_r - \tau A)\Delta\eta_n = b \tag{2.6}$$

in terms of the $r \times r$-matrices

$$\Gamma_r = \text{diag}(\gamma_k) \text{ and } A = (\langle \mathcal{A}q_k, q_l \rangle_\psi)_{kl} \tag{2.7}$$

and with right-hand side coefficients $b = (b_1, \ldots, b_r)$. In a similar way we obtain the coefficients of the Galerkin approximation of the correction $\Delta u_n^{(2)}$. The solution of these linear equations is usually not difficult, since they are of low dimension r. The difficulty, however, lies in the computation of the elements of Γ, A, and b, which requires the approximation of the inner products that are defined over $\mathbb{N}$.

Gauss-Christoffel summation. In order to compute the inner products $\langle \cdot, \cdot \rangle_\psi$ we have to approximate infinite sums. In [11, Section 9.7], an adaptive discrete multigrid algorithm (code SUMMATOR), which has been developed by Wulkow, has been presented to approximate large sums efficiently. In the course of further improvement of the polyreaction algorithms, discrete Gauss-Christoffel methods have been worked out that are intimately linked to the structure of the weighted inner products. For the trivial weight $\psi \equiv 1$, which we suggest to apply in the CME context, this leads to a discrete Gauss–Legendre quadrature, i.e., to a high-order summation technique. On the basis of the theory for Gauss-Christoffel quadrature, the nodes and weights can easily be computed also in the discrete case. Given a truncation index r, a triangular eigenvalue problem must be solved, for details see, e.g., [11]. It has been shown in [43] that the "aliasing error" introduced by the Gauss-Christoffel summation does not affect the quality of a Galerkin approximation of order r, if only at least $r + 1$ nodes are used.

2.2. Localization and the Galerkin h-p-method in 1d. In practice, the global approach outlined above has proven to be too restrictive, in particular when dealing with multi-modal distributions. Therefore we abandon the global approximation strategy and introduce a localization principle based on decomposing the state space into intervals. Furthermore, we choose the weight function to be the constant function $\psi \equiv 1$. The Hilbert space corresponding to $\psi \equiv 1$ is the space of all square summable sequences $\mathcal{H}_\psi = l^2 = \{u : \mathbb{N} \to \mathbb{R} : \sum_{x \in \mathbb{N}} |u(x)|^2 < \infty\}$. Square summable distributions $p(t) \in l^2$ do in general not possess bounded statistical moments, which is a desirable property. In order to guarantee bounded statistical moments, we consider the family of Hilbert spaces $\mathcal{H}_{\psi(\cdot,\rho)}$ with weight function $\psi(x, \rho) = (1 - \rho)\rho^x$ for $x \in \mathbb{N}$. It can be shown that the infinitesimal generators corresponding to the CME are typically discrete shift operators that have unique solutions within this family of Hilbert spaces. Basically such operators are Lipschitz-continuous within the scale, but not within one fixed $\mathcal{H}_{\psi(\cdot,\rho)}$. Now, if we require $p(0) \in \mathcal{H}_{\psi(\cdot,\rho)}$, we can ensure that $p(t)$ has bounded statistical moments for $t > 0$.

Consider a partition of the state space $\mathbb{N}$ given by

$$\mathbb{N} = \left[L^{(1)}, U^{(1)}\right] \cup \cdots \cup \left[L^{(M)}, U^{(M)}\right] \cup \left[L^{(M+1)}, \infty\right) \tag{2.8}$$

with finite discrete intervals $I^{(m)} = [L^{(m)}, U^{(m)}] = \{L^{(m)}, L^{(m)} + 1, \ldots, U^{(m)}\}$ and the semi-infinite interval $I_\infty = [L^{(M+1)}, \infty) = \{L^{(M+1)}, \ldots\}$. We assume that $L^{(i)} \leq$

$U^{(i)} < L^{(j)} \leq U^{(j)}$ for $i < j$, and set $L^{(1)} = 0$ and $L^{(m+1)} = U^{(m)} + 1$. In the sequel, we will describe how to approximate $p(t)$ locally on each interval $I^{(m)}$. Define $p_r^{(m)}(t)$ to be the Galerkin approximation of order r to $p(t)$ on $I^{(m)}$. In general, r will depend on the interval $I^{(m)}$.

We first consider the semi-infinite interval $I^{(\infty)}$. Choosing $p_0^{(\infty)}(t) \equiv 0$, the resulting approximation error on $I^{(\infty)}$ is

$$|p(t) - p_0^{(\infty)}(t)|_{I^{(\infty)}} = \sum_{x \in I^{(\infty)}} |p(t,x)|^2. \qquad (2.9)$$

By assumption, the distribution $p(t)$ possesses bounded zero, first and second statistical moments μ_0, μ_1 and μ_2, respectively. Consequently, the above approximation error will get arbitrarily small, if $x_{max} = L^{(M+1)} - 1$ is chosen large enough. It has been shown in [44] that

$$x_{max} = \frac{\mu_1}{\mu_0} + \kappa \cdot \sqrt{\frac{\mu_2}{\mu_0} - \left(\frac{\mu_1}{\mu_0}\right)^2} \qquad (2.10)$$

is a reasonable choice, where κ is a safety factor with a typical value of $\kappa = 10$. The expression (2.10) can be derived by applying the Chebyshev inequality of statistics to the probability that $p(t,x) > \varepsilon$ for $\varepsilon > 0$ and $x > x_{\max}$. This effectively accounts to restricting the semi-infinite approximation problem to a finite approximation problem on $[0, x_{max}]$. It is important to remark that x_{max} depends on the time evolution of the system and is not known *a priori*. The adaptation of x_{max} from time step to time step will be discussed in the next section (see *Adaptation of state space truncation*).

It remains to characterize the Galerkin approximations on the finite intervals $I^{(m)}$ for $m = 1, \ldots, M$. Consider the orthogonal basis $\{T_k^{(m)} : k = 0, \ldots, (U^{(m)} - L^{(m)})\}$ of discrete Chebyshev polynomials satisfying

$$\left\langle T_k^{(m)}, T_j^{(m)} \right\rangle_{I^{(m)}} = \sum_{x \in I^{(m)}} T_k^{(m)}(x) \, T_j^{(m)}(x) = \gamma_j^{(m)} \delta_{k,j}.$$

Then, we can represent $p(t)$ on $I^{(m)}$ in terms of the Chebyshev polynomials

$$p(t,x) = \sum_{k=0}^{U^{(m)}-L^{(m)}} a_k^{(m)} \, T_k^{(m)}(x).$$

For simplicity, we ignore the time-dependence of the coefficients and the polynomials and write $a_k^{(m)}$ and $T_k^{(m)}(x)$ instead of $a_k^{(m)}(t)$ and $T_k^{(m)}(t,x)$. Again the Galerkin approximation $p_r^{(m)}(t)$ to $p(t)$ on $I^{(m)}$ is defined by some suitable truncation of the above expansion

$$p_r^{(m)}(t,x) = \sum_{k=0}^{r} a_k^{(m)} \, T_k^{(m)}(x)$$

with polynomial order $r \leq U^{(m)} - L^{(m)}$, which in general will depend on I_m. The resulting error may be estimated by

$$\mathrm{eps}_r^{(m)} = \|p_{r+1}^{(m)}(t) - p_r^{(m)}(t)\|_{I^{(m)}} = \left| a_{r+1}^{(m)} \right| \sqrt{\gamma_{r+1}^{(m)}}.$$

In brief, we summarize the resulting decomposition and polynomial approximations by the element-order pattern

$$\Delta_1 = \big\{(I_1, r_1), (I_2, r_2), \ldots, (I_M, r_M)\big\}.$$

We expect that for an efficient approximation $r_m \ll U^{(m)} - L^{(m)}$ on most of the elements. The whole approach using local refinement *and* higher order approximations is called *Galerkin h-p-method*. For rather smooth distributions, we can expect that only a few intervals with moderate polynomial orders are sufficient. However, for distributions with steep flanks or for multi-modal distributions, the approach allows an automated local adaptation to avoid bad and inefficient approximation. This, in particular, is the reason for abandoning the global approach based on a weighted Galerkin ansatz space and leads to efficient approximations even in higher dimensions. Moreover, in an early, more theoretically paper about h-p-methods for PDEs [24] it has been demonstrated on the basis of examples with known analytical solutions that exponential convergence can be obtained by such an approach. This could be verified for a wide range of practical applications for the discrete case.

2.3. A Tensor Product Approach to Multiple Dimensions. After these preparations we can extend the approach to higher dimensions.

Domains in d dimensions. Analogous to the $1d$ case, we restrict the semi-infinite domain $\mathbb{N}^d$ to a finite domain

$$D = [0, x_{max,1}] \otimes \cdots \otimes [0, x_{max,d}],$$

where $x_{max,i}$ denotes the upper bound along the ith dimension. Next we decompose D into disjoint rectangles $\{D^{(m)} : m = 1, \ldots M\}$ such that

$$D = \bigcup_{m=1}^{M} D^{(m)} \quad \text{with} \quad D^{(m)} = [L^{(m,1)}, U^{(m,1)}] \otimes \cdots \otimes [L^{(m,d)}, U^{(m,d)}].$$

This tensor ansatz allows us to directly make use of the one-dimensional h-p-method. Given the multi-index $k = (k_1, \ldots, k_d)$, we define

$$T_k^{(m)}(x) = T_{k_1}^{(m,1)}(x_1) \cdots T_{k_d}^{(m,d)}(x_d)$$

as the product of the discrete Chebyshev polynomials $T_{k_i}^{(m,i)}$ on the intervals $[L^{(m,i)}, U^{(m,i)}]$ in the ith dimension. We have to take care that the boundaries of the rectangles appear only once (see Figure 3.1 for illustration). A Galerkin approximation on a rectangle $D^{(m)}$ is then given by

$$p_r^{(m)}(t, x) = \sum_{k=1}^{r^{(m)}} a_k^{(m)} T_k^{(m)}(x)$$

with multi-index $r^{(m)} = (r_1^{(m)}, \ldots, r_d^{(m)}) \in N^d$. This is a highly non-uniform and flexible structure, since for each axis on each element an independent polynomial system is used. The resulting element-order pattern is given by

$$\Delta_d = \{(D^{(1)}, r^{(1)}), \ldots, (D^{(M)}, r^{(M)})\}.$$

It should be noted that the evaluation of the approximation for a given $x \in D$ can be efficiently performed by expressing $p_r^{(m)}(t)$ as the result of a sequence of $1d$-Galerkin approximations (which in turn can be computed by fast summation based on the three-term recurrence of the polynomials [11]).

Refinement strategy. Our aim is to find local approximations $p_r^{(m)}(t)$ on each domain $D^{(m)}$ such that the overall approximation

$$p_r(t,x) = \sum_{m=1}^{M} p_r^{(m)}(t,x) \cdot 1_{D^{(m)}}(x)$$

with characteristic function

$$1_{D^{(m)}}(x) = \left\{ \begin{array}{ll} 1 & \text{if } x \in D^{(m)} \\ 0 & \text{if } x \notin D^{(m)} \end{array} \right.$$

satisfies

$$\text{eps}_D = \|p_r - p\|_D = \left(\sum_{m=1}^{M} \left\| p(t) - p_r^{(m)}(t) \right\|_{D^{(m)}}^2 \right)^{1/2} \leq \text{TOL}_X. \qquad (2.11)$$

Since the general structure of the tensor product approach has now become apparent, we restrict the following considerations to two dimensions for sake of clarity.

An element-order pattern Δ_2 has to be generated by a multi-level algorithm using certain refinement strategies. Based on an initial pattern it has to be decided step-by-step how to change elements or polynomial orders. To do so, *directional* error estimates are necessary, which can be used to predict the approximation errors after increase of the polynomial order or refinement of (some of) the intervals (for details, see [12]). In view of an efficient algorithmic realization, the element-order-pattern is chosen and changed such that the amount of work necessary to compute the final approximation is as small as possible. Therefore only rectangles are changed with a local error larger than some threshold. The threshold is computed based on all error predictions such that the errors within the decomposition are equilibrated. However, one cannot set up a local tolerance for single domains *a priori*, since the number of domains on the final level is not known beforehand. Additionally all refinement steps are chosen in view of the obtained (and necessary) gain of accuracy per work.

From time step to time step of the ROM, the grid has to be coarsened in some sense, since otherwise moving peaks or changing shapes would lead to a monotone increase of expansion coefficients. For the coarsening the error estimates and the now available average error per element are used. All domains having a local error lower than a certain percentage of the average error will be changed by reducing the polynomial order by one. If a minimal polynomial order is reached, such elements $D^{(m)}$ will be merged with neighbors $D^{(n)}$ with $L^{(m,i)} = L^{(n,i)} \wedge U^{(m,i)} = U^{(n,i)}$ for at least one dimension i. This strategy is grid-conservative and flexible at the same time. Additionally, the number of levels required to fulfill the stationary tolerance of the next time step is kept small.

Adaptation of state space truncation. In the one-dimensional case an update of $x_{\max}$ can easily be realized by adding or deleting intervals at the right boundary. This is more complicated for the higher-dimensional case, since here all directions are affected simultaneously. Therefore all coordinate directions are stretched or compressed by the factor $f_i = (y_{\max,i} + 1)/(x_{\max,i} + 1)$, where $y_{\max,i}$ denotes the new estimate. At the same time care has to be taken that all resulting $L^{(m,i)}$ and $U^{(m,i)}$ are still defined as natural numbers. Finally, when compressing the whole grid Δ_d, it has to be ensured that no degenerated domains $D^{(m)}$ arise.

Efficient evaluation of Galerkin matrix entries. In order to illustrate how the setup of the Galerkin equations works in $2d$, we consider the special operator

$$(\mathcal{A}p)(x_1, x_2) = \alpha(x_1 - 1, x_2 + 1)p(t, x_1 - 1, x_2 + 1) \qquad (2.12)$$

as it appears as a typical term in the CME (see Sec. 3) with propensity a and distribution $p(t)$.

For the treatment within the Galerkin method, the operator $\mathcal{A}$ has to be applied to all basis functions $T_k^{(n)} = T_k^{(n,1)} T_l^{(n,2)}$ of all ansatz elements $D^{(n)} = I^{(n,1)} \otimes I^{(n,2)}$, and to be tested with all such pairs from test elements $D^{(m)} = I^{(m,1)} \otimes I^{(m,2)}$. Thus we have to compute double sums $S = S(i, j, k, l; n, m)$ of the form

$$S = \sum_{x_1 \in I^{(n,1)}} \sum_{x_2 \in I^{(n,2)}} T_i^{(m,1)}(x_1) T_j^{(m,2)}(x_2) \alpha(x_1 - 1, x_2 + 1) T_k^{(n,1)}(x_1 - 1) T_l^{(n,2)}(x_2 + 1)$$

$$(2.13)$$

for all i, j, k, l and $1 \le m, n \le M$. Here, the complexity is reduced by the fact that only neighboring elements can lead to non-zero S. The intersection of the edges of neighboring elements will be denoted by

$$I_S^1 = [\max(L^{(m,1)}, L^{(n,1)} - 1), \min(U^{(m,1)}, U^{(n,1)} - 1)]$$
$$I_S^2 = [\max(L^{(m,2)}, L^{(n,2)} + 1), \min(U^{(m,2)}, U^{(n,2)} + 1)].$$

Note that the I_S^i may consist of a single point, which will contribute to the summation.

The inner product S defined in eq. (2.13) is approximated by Gauss-Christoffel summation $\tilde{S}$ based on the nodes $x_1^{(1)}, \ldots, x_1^{(n_1)}$ and $x_2^{(1)}, \ldots, x_2^{(n_2)}$ and corresponding weights $v_1, \ldots, v_{n_1}$ and $w_1, \ldots, w_{n_2}$, respectively. The sums can be evaluated in two steps. First, for all Gauss-Christoffel nodes of the first direction, the second sum is computed for all j and l:

$$\tilde{S}_2(x_1^{(\gamma)}) = \sum_{\beta=1}^{n_2} v_\beta \, \alpha(x_1^{(\gamma)}, x_2^{(\beta)}) T_j^{(m,2)}(x_2^{(\beta)}) T_l^{(n,2)}(x_2^{(\beta)} + 1).$$

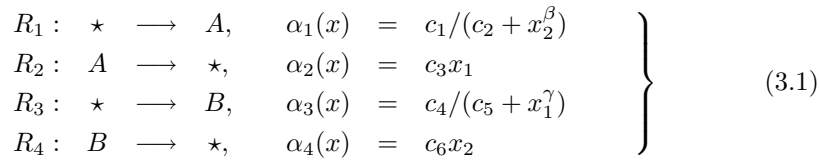Then $\tilde{S}$ can be expressed in terms of the intermediate approximations $\tilde{S}_2$:

$$\tilde{S} = \sum_{\gamma=1}^{n_1} w_\gamma T_i^{(m,1)}(x_1^{(\gamma)}) T_k^{(n,1)}(x_1^{(\gamma)} - 1) \tilde{S}_2(x_1^{(\gamma)}).$$

In many cases, the propensity factorizes according to $\alpha(x_1, x_2) = \alpha_1(x_1) a_2(x_2)$. Then, $\tilde{S}_2(x_1^{(\gamma)})$ is independent of $x_1^{(\gamma)}$ and $\tilde{S}$ is just the product of two 1d Galerkin sums for shift operators. This approximation is exact, if the α_i are polynomials and the number of nodes is chosen accordingly. The structure of $\tilde{S}$ shows that starting from a non-uniform, n-dimensional grid the final treatment of operators can be reduced to a one-dimensional evaluation on single discrete intervals. This strategy is crucial for the whole tensor product approach, because it allows to treat all terms of the equations presented in Section 3 in a structured way.

A summary of the entire algorithm in its essential parts is given in the Appendix. We remark that the main computational effort of the algorithms consists in the treatment of the single operators (cf. eq. (2.12)) leading to a quadratic matrix of dimension number of degrees of freedom. The practical experience with this algorithmic approach in different fields shows that this number of degrees of freedom does not increase exponentially with the dimension due to the local adaptation scheme and only moderately with the required tolerance.

3. Numerical Examples. In this section, we give two numerical examples to illustrate the adaptive discrete Galerkin approach described above. The first one, a genetic toggle switch, has been selected for comparison purposes, since it has already been used as a test problem in [39] and [16]. The second one is an artificial test problem constructed to generate special difficulties for any kind of tensor product approaches, since there the dynamics asymptotically collapses to the diagonal between two sources.

3.1. Bistable Toggle Switch. As a first application, we consider a model for the simulation of a genetic toggle switch in E. coli; cf. [18]. The model consists of two competing repressors, A and B, transcribed by two constitutive promoters. Each of the two repressors can inhibit the production of the competing repressor by binding to the corresponding genetic control sequences of the promoter. The system includes the following reactions:

$$
\left.
\begin{array}{llllll}
R_1: & \star & \longrightarrow & A, & \alpha_1(x) & = & c_1/(c_2 + x_2^\beta) \\
R_2: & A & \longrightarrow & \star, & \alpha_2(x) & = & c_3 x_1 \\
R_3: & \star & \longrightarrow & B, & \alpha_3(x) & = & c_4/(c_5 + x_1^\gamma) \\
R_4: & B & \longrightarrow & \star, & \alpha_4(x) & = & c_6 x_2
\end{array}
\right\}
\tag{3.1}
$$

Here and below, the symbol $\star$ denotes a species which is present in such a large quantity that its molecule number can be considered as constant. The positive constants c_1, c_2 and c_3, c_4 determine the maximal rate of synthesis of the repressors A and B, respectively. The degradation rates of the repressors are denoted by c_5 and c_6, while coefficients $\beta > 1$ and $\gamma > 1$ specify the cooperativity of the two promoters.

The reaction R_1 is constructed in such a way that the corresponding propensity $\alpha_1(x)$ is small whenever x_2 is large. Hence, the transcription of A is inhibited if many copies of B are present. Conversely, a large copy number of A inhibits the production of new B since $x_1 \gg 1$ implies $\alpha_3(x) \ll 1$. These two scenarios (large number of A and small number of B, or large number of B and small number of A) correspond to the two stable steady states of the traditional reaction rate equations

$$
\begin{aligned}
\dot{y}_1 &= \frac{c_1}{c_2 + y_2^\beta} - c_3 y_1 \\
\dot{y}_2 &= \frac{c_4}{c_5 + y_1^\gamma} - c_6 y_2
\end{aligned}
\tag{3.2}
$$

(cf. [18]). The solution of (3.2), however, does not provide an appropriate description since a single trajectory can only converge to one of the steady states, whereas in the real biological system flipping between stable states due to chemical or thermal induction is possible (cf. [18]). This switching behavior can only be reproduced by a stochastic description including fluctuations which can induce transitions from one steady state to the other one.

Due to the bistability of the toggle switch, the solution of the corresponding CME

$$
\frac{\partial}{\partial t} p(t, x_1, x_2) = \frac{c_1}{c_2 + x_2^\beta} \big(p(t, x_1 - 1, x_2) - p(t, x_1, x_2) \big)
\tag{3.3}
$$

$$
+ c_3(x_1 + 1) p(t, x_1 + 1, x_2) - c_3 x_1 p(t, x_1, x_2)
\tag{3.4}
$$

$$
+ \frac{c_4}{(c_5 + x_1^\gamma)} \big(p(t, x_1, x_2 - 1) - p(t, x_1, x_2) \big)
\tag{3.5}
$$

$$
+ c_6(x_2 + 1) p(t, x_1, x_2 + 1) - c_6 x_2 p(t, x_1, x_2)
\tag{3.6}
$$

is a bimodal PDF. Bimodality is of prime importance in many biological applications since it indicates the presence of two different scenarios or macro-states. Hence it is crucial that a numerical method constructed for solving the CME captures the bimodality correctly.

In [39], the parametrisation

$$
\begin{array}{llllllll}
c_1 & = & c_4 & = & 3 \cdot 10^3 s^{-1}, & c_2 & = & c_5 & = & 1.1 \cdot 10^4, \\
c_3 & = & c_6 & = & 0.001 s^{-1}, & \beta & = & \gamma & = & 2
\end{array}
\tag{3.7}
$$

has been used, and both the stationary distribution and the time-dependent solution in the interval $[0, 10^4]$ have been computed on the domain $[0, 399] \times [0, 399]$. Instead of solving the full discrete CME, however, the authors presented an approximation of the Fokker-Planck equation, a PDE known to be, in some sense, the continuous counterpart of the CME (see, e.g., [17]). In the case of the genetic toggle switch, this may yield a reasonable approximation, but in situations where the discrete nature of the reaction system is crucial (cf. [28, 40]), replacing the CME by the Fokker-Planck equation causes a large modelling error in addition to the numerical approximation error. Hence, the approach proposed in [39] can actually not be considered as a viable method for the CME, but rather as a method for the Fokker-Planck equation.

The genetic toggle switch has been investigated again in [16]. There, a computation based on the full CME has been presented, but the computational domain was decreased in [16] to approximately $[0, 200] \times [0, 200]$ by a rescaling of the parameters. Note that, by construction, our adaptive discrete Galerkin method does not require any such downscaling.

The following simulations have been performed with a special $2d$-version of the program package PREDICI®. In our first example, the full CME of (3.1) with parameters (3.7) was solved up to a predefined tolerance of 0.03. As an initial distribution, the product Gaussian

$$
\frac{1}{2\pi\sigma_1\sigma_2} \exp\left(-\frac{(x_1 - \mu_1)^2}{2\sigma_1^2}\right) \exp\left(-\frac{(x_2 - \mu_2)^2}{2\sigma_2^2}\right)
\tag{3.8}
$$

has been evaluated and normalized on the discrete state space for $\mu_1 = \mu_2 = 133$ and $\sigma_1 = \sigma_2 = \sqrt{133}$. Figure 3.1 shows the PDF at $t = 6 \cdot 10^4$. By this time, the PDF has almost converged to the stationary distribution, and a comparison of Figure 3.1 with the right panel of Figure 3 in [39] shows indeed a very good agreement. The bimodality is clearly visible and indicates the relevant biological information, namely the bistability of the toggle switch. Along with the contour lines, the h-p-grid is depicted; the numbers indicate the maximal order of the polynomial basis. As desired, the discrete Galerkin method refines either the subdivision of space or the polynomial order or both in regions where the solution changes significantly, and the grid shows a nearly perfect symmetry. The number of degrees of freedom of this h-p-grid is about 1600.

Figure 3.2 shows the time-evolution of the degrees of freedom for the tolerances 0.1, 0.03, 0.01. It can be nicely seen how the adaptive strategies refine the approximation according to the requirements of the changing distribution. It can also be recognized that the number of variables increases more or less linearly with the required tolerance - a sign for nearly exponential convergence rate of the discrete h-p-method, even if a wider range of tolerances would be necessary to proof this here (see remark in Section 2.2). It should be pointed out that even when the desired accuracy is rather high, only about 3000 degrees of freedom have to be handled, whereas a naïve
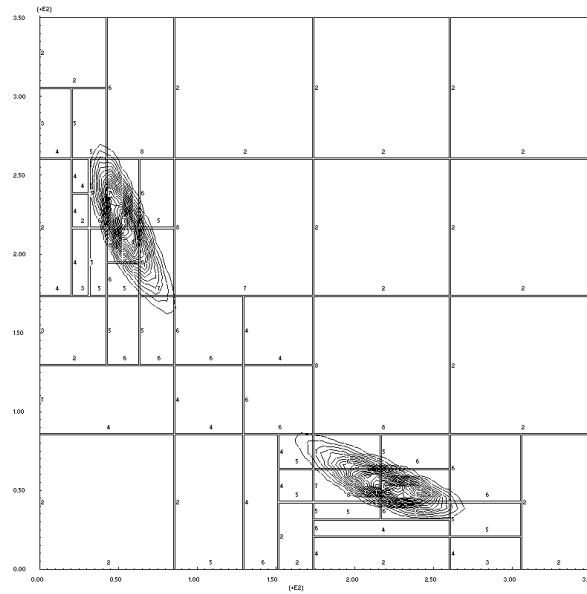
FIG. 3.1. *Stationary distribution of the genetic toggle switch (3.1), (3.7) obtained from a dynamic simulation on the time-interval* $[0\,,\,6\cdot 10^4]$. *The boxes and numbers indicate the h-p-grid and the maximal order of the polynomial basis.*

computation of the PDF on the considered state space $0,\ldots,349\times 0,\ldots,349$ would require the solution of $350^2 = 122500$ ODEs. This corresponds to a reduction of more than 97.5%.
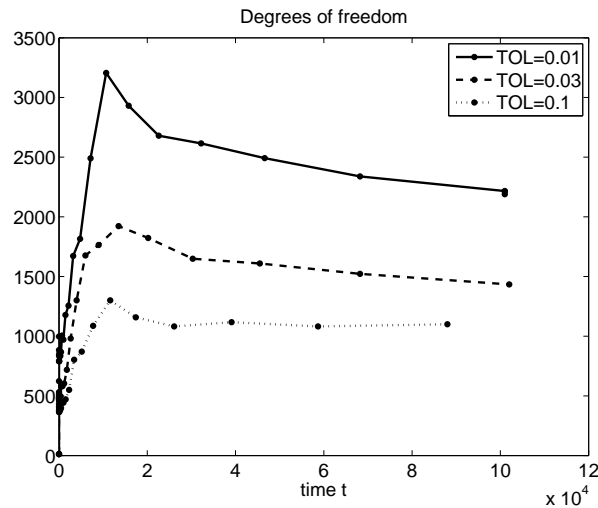

FIG. 3.2. *Total number of degrees of freedom as a function of time for the tolerances 0.1, 0.03, 0.01.*

Next, in order to study an asymmetric and more challenging situation, we chose $\mu_1 = 100$, μ_2, $\sigma_1 = 10$, and $\sigma_2 = 3$ in (3.8). The initial domain has been chosen to $[0, 150] \times [0, 20]$. In the course of the integration this range is extended adaptively. The

panels of Figure 3.3 show the time evolution of the PDF. First, the main part of the probability mass is attracted by the lower steady state since it is closer to the center of the initial PDF than the upper one. At $t = 16000$ the distribution is stretched a bit, and the additional grid refinements around $(50, 200)$ seem to be dispensable.
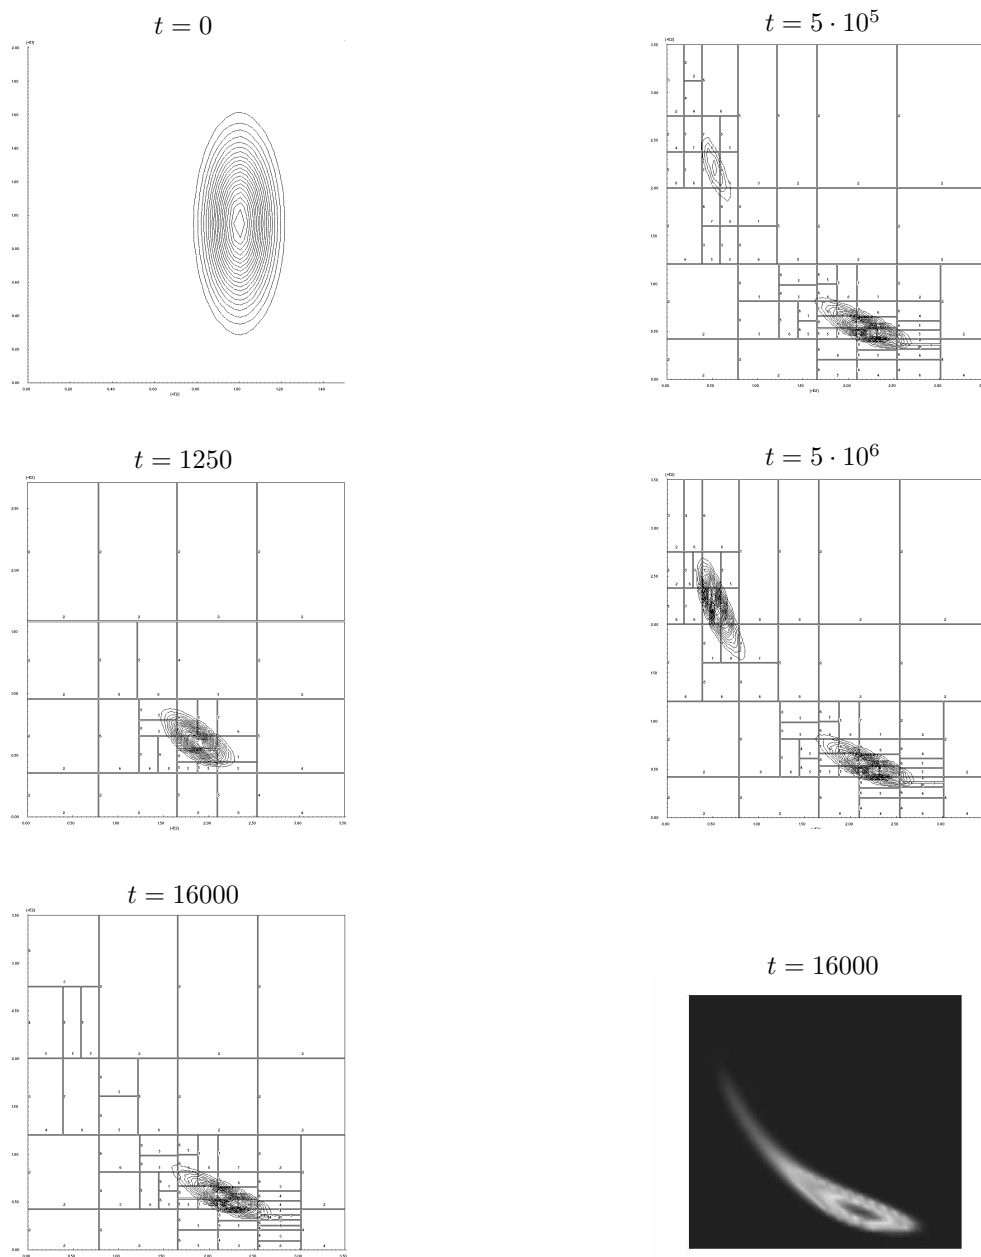

FIG. 3.3. *"Movie" of the PDF of the genetic toggle switch (3.1), (3.7). The panels shows the solution at different times. In the lower right panel the solution at $t = 16000$ is depicted in logarithmical greyscale.*

However, the coloured contour plot (here plotted in logarithmical greyscale) indi-

cates that some of the probability mass already leaks to the other stationary state. As time evolves, more and more probability mass fluctuates very slowly from the lower to the upper steady state until finally an equilibrium between both states is reached in the stationary distribution. This simulation has been performed with tolerance 0.01 and a maximum of 2700 variables at intermediate states. The h-p-grid is nearly symmetric at the final stage. This example shows how the h-p-algorithm and the ROM automatically capture the dynamics and the structure of distributions. More pictures would lead too far off here, but, e.g., by setting $\mu = 10^{-4}$ or even $\mu = 10^{-5}$, intermediate distributions with maximal particle numbers up to several thousands occurred and could equally well be represented.
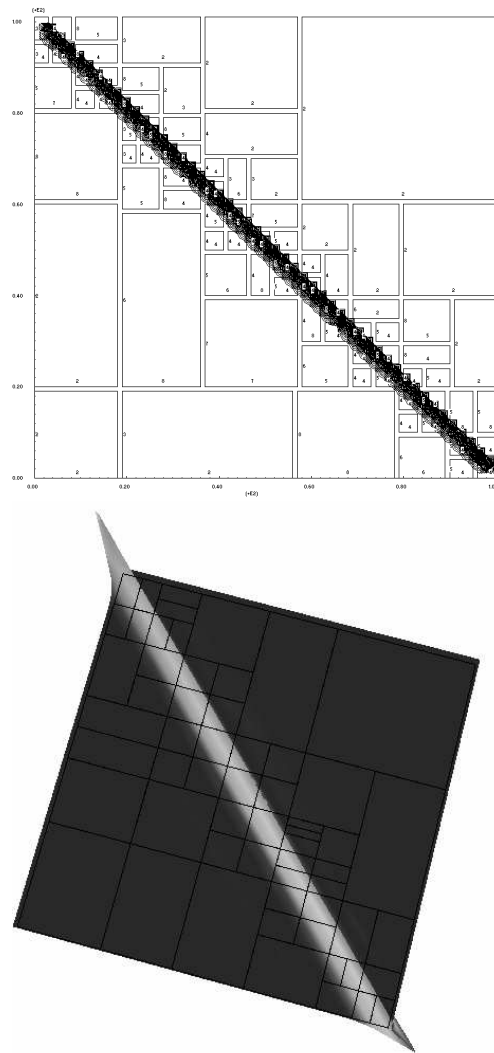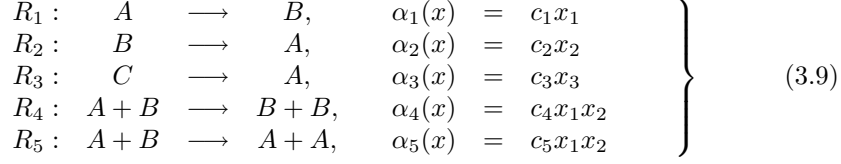

Fig. 3.4. *Top: Contour plot of the solution of the CME (3.10), (3.11) for $N = 100$ and* TOL $= 0.02$ *at $t = 5$. Bottom: Three dimensional visualization of the solution of the CME (3.10), (3.11) for $N = 10000$ and* TOL $= 0.05$ *at $t = 5$.*

3.2. A Challenging Test Problem. As a second example, we chose a comparatively simple model problem which, however, represents a real challenge for any numerical method based on a tensor product ansatz:

$$
\left.
\begin{array}{rcccrcl}
R_1: & A & \longrightarrow & B, & \alpha_1(x) & = & c_1 x_1 \\
R_2: & B & \longrightarrow & A, & \alpha_2(x) & = & c_2 x_2 \\
R_3: & C & \longrightarrow & A, & \alpha_3(x) & = & c_3 x_3 \\
R_4: & A+B & \longrightarrow & B+B, & \alpha_4(x) & = & c_4 x_1 x_2 \\
R_5: & A+B & \longrightarrow & A+A, & \alpha_5(x) & = & c_5 x_1 x_2
\end{array}
\right\}
\tag{3.9}
$$

Since there are three species A, B, C involved, the state space is actually three dimensional, but since the total number N of molecules remains invariant in all reactions, the molecule numbers of one of the species, say C, can be expressed in terms of the other ones:

$$
x_3 = N - x_1 - x_2.
$$

This actually reduces the three-dimensional problem to a two-dimensional one. The corresponding CME reads

$$
\frac{\partial}{\partial t} p(t, x_1, x_2) = c_1(x_1 + 1)p(t, x_1 + 1, x_2 - 1) - c_1 x_1 p(t, x_1, x_2) \tag{3.10}
$$
$$
+ c_2(x_2 + 1)p(t, x_1 - 1, x_2 + 1) - c_2 x_2 p(t, x_1, x_2)
$$
$$
+ c_3(N - x_1 - x_2 + 1)p(t, x_1 - 1, x_2) - c_3(N - x_1 - x_2)p(t, x_1, x_2)
$$
$$
+ c_4(x_1 + 1)(x_2 - 1)p(t, x_1 + 1, x_2 - 1) - c_4 x_1 x_2 p(t, x_1, x_2)
$$
$$
+ c_5(x_1 - 1)(x_2 + 1)p(t, x_1 - 1, x_2 + 1) - c_5 x_1 x_2 p(t, x_1, x_2).
$$

We suppose that at $t = 0$ there are N molecules of C and no molecules of A or B such that the initial distribution is given by

$$
p(0, x_1, x_2) = \left\{
\begin{array}{ll}
1 & \text{if } x_1 = x_2 = 0, \\
0 & \text{else.}
\end{array}
\right.
$$

It is easy to see that in the limit $t \longrightarrow 0$ all molecules of C disappear and that the exchange between A and B reaches an equilibrium. In the special case

$$
c_1 = \ldots = c_5 = 1, \tag{3.11}
$$

it can be shown that the stationary distribution is the uniform distribution on the line $\{(x_1, x_2) \in \mathbb{N}^2 \ : \ x_1 + x_2 = N\}$, i.e.

$$
\bar{p}(x_1, x_2) = \lim_{t \to \infty} p(t, x_1, x_2) = \left\{
\begin{array}{ll}
1/(N+1) & \text{if } x_1 + x_2 = N, \\
0 & \text{else.}
\end{array}
\right.
$$

This distribution is the worst case for any tensor product ansatz, because an exact representation of $\bar{p}$ requires $N + 1$ basis functions although $\bar{p}$ contains only $N + 1$ nonzero elements. Moreover, all nonzero elements have the same value such that any truncation produces the same error.

Figure 3.4 (upper panel) shows the result for N=100 at t=5 computed with TOL = 0.02 . At that time, the distribution is nearly reduced to the line $x_1 + x_2 = N$, and obviously this leads to difficulties regarding the domain decomposition into rectangles.

About 2500 variables are necessary to compute the solution. In comparison to the maximum number of about 10000, this is not a convincing reduction, and for smaller N, the reduction is even less efficient. However, if we do not try to approximate the final end of the process, but stop a bit earlier, the pay-off can be immense. Figure 3.4 (lower panel) shows the solution at t=4, but with N increased to $N = 10000$. To obtain a tolerance of TOL = 0.05 on the state space with 10000×10000 states, only about 2600 variables are required, which corresponds to a reduction of 0.999974%! A 3d-plot shot from the top shows how narrow even then the front already is.

Conclusion. We present an adaptive discrete Galerkin method for solving the chemical master equation (CME). In one state space dimension, such methods have had an impressive influence on the modelling of polyreaction kinetics. For the CME, however, the step towards multiple dimensions turns out to be crucial. The paper clearly shows that discrete Galerkin methods allow an efficient treatment of CME. In particular, examples as the one presented in the last figure could not be solved up to now - not even in two dimensions. However, the challenging second example makes it apparent that there is still a complexity barrier in view of more complicated cases and higher dimensions. Future research will therefore have to focus on how to overcome these difficulties. Nevertheless, the results presented here will lay a measure for any future improvements.

REFERENCES

[1] A. Alfonsi, E. Cancès, G. Turinici, B. Di Ventura, and W. Huisinga, *Adaptive simulation of hybrid stochastic and deterministic models for biochemical systems*, ESAIM Proceeding, 14 (2005), pp. 1–13.

[2] F. Bornemann, *An adaptive multilevel approach to parabolic equations: I. general theory and 1d implementation.*, IMPACT Comput Sci Eng, 2 (1990), pp. 279–317.

[3] ———, *An adaptive multilevel approach to parabolic equations: II. variable-order time discretization based on a multiplicative error correction.*, IMPACT Comput Sci Eng, 3 (1991), pp. 93–122.

[4] ———, *An adaptive multilevel approach to parabolic equations in two space dimensions*, PhD thesis, Freie Universität Berlin, 1991.

[5] K. Burrage, M. Hegland, S. MacNamara, and R.B. Sidje, *A Krylov-based finite state projection algorithm for solving the chemical master equation arising in the discrete modelling of biological systems*, in Markov Anniversary Meeting: An international conference to celebrate the 150th anniversary of the birth of A.A. Markov, A.N.Langville and W.J.Stewart, eds., Boson Books, 2006, pp. 21 – 38.

[6] K. Burrage and T. Tian, *Poisson Runge-Kutta methods for chemical reaction systems*, in Advances in Scientific Computing and Applications, Y. Lu W. Sun and T. Tang, eds., Science Press, Beijing/New York, 2004, pp. 82–96.

[7] K. Burrage, T. Tian, and P. Burrage, *A multi-scaled approach for simulating chemical reaction systems*, Progress in Biophysics and Molecular Biology, 85 (2004), pp. 217–234.

[8] Computing in Technology GmbH (CiT), *http://www.cit-wulkow.de*.

[9] K. Deimling, *Ordinary Differential Equations in Banach Spaces*, vol. 596 of Lecture Notes in Mathematics, Springer, Berlin/New York, 1977.

[10] ———, *Nonlinear Functional Analysis*, Springer, Berlin/New York, 1985.

[11] P. Deuflhard and A. Hohmann, *Numerical Analysis in Modern Scientific Computing: An Introduction*, vol. 43 of Texts in Applied Mathematics, Springer, 2003.

[12] P. Deuflhard, W. Huisinga, T. Jahnke, and M. Wulkow, *Adaptive discrete Galerkin methods applied to the chemical master equation*, Tech. Report 07-04, Zuse Institute Berlin, 2007.

[13] P. Deuflhard and M. Wulkow, *Computational treatment of polyreaction kinetics by orthogonal polynomials of a discrete variable*, IMPACT Comput Sci Eng, 1 (1989), pp. 269–301.

[14] ———, *Simulationsverfahren in der Polymerchemie*, in Mathematik in der Praxis, Springer International, 1995, pp. 117–136.

[15] M. B. Elowitz, E. D. Siggia, P. S. Swain, and A. J. Levine, *Stochastic gene expression in a single cell*, Science, 297 (2002), pp. 1183–1186.

[16] S. Engblom, *A discrete spectral method for the chemical master equation*, Submitted, (2006).

[17] C. W. Gardiner, *Handbook of Stochastic Methods*, Springer, Berlin, 2nd enlarged edition ed., 1985.

[18] T. S. Gardner, C. R. Cantor, and J. J. Collins, *Construction of a genetic toggle switch in Escherichia coli*, Nature, 403 (2000), pp. 339–342.

[19] M. A. Gibson and J. Bruck, *Efficient exact stochastic simulation of chemical systems with many species and many channels*, J. Phys. Chem. A, 104 (2000), pp. 1876–1889.

[20] D.T. Gillespie, *A general method for numerically simulating the stochastic time evolution of coupled chemical reactions*, J. Comput. Phys., 22 (1976), pp. 403–434.

[21] D. T. Gillespie, *Exact stochastic simulation of coupled chemical reactions*, J. Phys. Chem., 81 (1977), pp. 2340–2361.

[22] ———, *A rigorous derivation of the chemical master equation*, Physica A, 188 (1992), pp. 404–425.

[23] ———, *Approximate accelerated stochastic simulation of chemically reacting systems*, Journal of Chemical Physics, 115 (2001), pp. 1716–1733.

[24] W. Gui and I. Babushka.

[25] E. L. Haseltine and J. B. Rawlings, *Approximate simulation of coupled fast and slow reactions for stochastic chemical kinetics*, Journal of Chemical Physics, 117 (2002), pp. 6959–6969.

[26] H. Horn and M. Wulkow, *Simulation von Wachstum und Abtrag von Biomasse - Eine exemplarische Betrachtung für eine 2d-Modellierung*, Chemie Ingenieur Technik, 77 (2005).

[27] P. D. Iedema, M. Wulkow, and H. C. J. Hoefsloot, *Modelling molecular weight and degree of branching distribution of low-density polyethylene*, Macromolecules, 33 (2000), pp. 7173–7184.

[28] T. Jahnke and W. Huisinga, *A dynamical low-rank approach to the chemical master equation*, Submitted, (2007).

[29] Markus Kraft, *Modelling of particulate processes*, Kona, Powder and Particle, 23 (2005), pp. 18–35.

[30] J. Lang, *Adaptive Multilevel Solution of Nonlinear Parabolic PDE Systems*, vol. 16 of Lecture Notes in Computational Science and Engineering, Springer, 2001.

[31] S. MacNamara, K. Burrage, and R.B. Sidje, *Multiscale modeling of chemical kinetics via the master equation*, MMS, 6 (2008), pp. 1146–1168.

[32] H. H. McAdams and A. P. Arkin, *Stochastic mechanisms in gene expression*, PNAS, 94 (1997), pp. 814–819.

[33] ———, *It's a noisy business! Genetic regulation at the nanomolar scale*, Trends Genet, 15 (1999), pp. 65–69.

[34] Brian Munsky and Mustafa Khammash, *The finite state projection algorithm for the solution of the chemical master equation*, J. Chemical Physics, (2006).

[35] S. Peles, B. Munsky, and M. Khammash, *Reduction and solution of the chemical master equation using time-scale separation and finite state projection*, J Chem Phys, 125 (2006), pp. 204104–.

[36] C. V. Rao and A. P. Arkin, *Stochastic chemical kinetics and the quasi-steady-state assumption: Application to the Gillespie algorithm*, J. Chem. Phys., 118 (2003), pp. 4999–5010.

[37] J. M. Raser and E. K. O'Shea, *Control of stochasticity in eukaryotic gene expression*, Science, 304 (2004), pp. 1811–1814.

[38] H. Salis and Y. Kaznessis, *Accurate hybrid simulation of a system of coupled chemical or biochemical reactions*, J. Chem. Phys., 122 (2005).

[39] P. Sjöberg, P. Lötstedt, and J. Elf, *Fokker-Planck approximation of the master equation in molecular biology*, Comput. Vis. Sci., (2007).

[40] R. Srivastava, L. You, J. Summers, and J. Yin, *Stochastic vs. deterministic modeling of intracellular viral kinetics*, J. theor. Biol., 218 (2002), pp. 309–321.

[41] Daniel W. Stroock, *An Introduction to Markov Processes*, vol. 230 of Graduate Texts in Mathematics, Springer, 2005.

[42] M. Wulkow, *Numerical Treatment of countable systems of ordinary differential equations*, doctoral thesis, Freie Universität Berlin, 1990.

[43] ———, *Adaptive treatment of polyreactions in weighted sequence spaces*, IMPACT Comput Sci Eng, 4 (1992), pp. 153–193.

[44] ———, *The simulation of molecular weight distribution in polyreaction kinetics by discrete*

Galerkin methods, Macromol Theory Simul, 5 (1996), pp. 393–416.

[45] M. WULKOW AND P. DEUFLHARD, *Towards an efficient computational treatment of heterogeneous polymer reactions*, in Computational Ordinary Differential Equations, S. O. Fatunla, ed., University Press PLC, Ibadan, 1992, pp. 287–306.

[46] W. J. YOON, Y. S. KIM, I. S. KIM, AND K. Y. CHOI, *Recent advances in polymer reaction engineering: Modeling and control of polymer properties*, Korean J. Chem. Eng., 21 (2004), pp. 147–167.

Appendix.

Algorithmic flow. According to the different steps outlined in the preceding sections the algorithm can be summarized as follows. The most outer loop is the *Time Control.* Within the time control, the routine on *Galerkin approximation* is called. This, in turn, relies on the subroutines *Initial element-order-pattern* that coarsens the element-order-pattern from the previous time step and *Final element-order-pattern* that computes the necessary element-order-pattern to meet the spatial tolerance.

Time control. Given the current approximation u_n of $p(t_n)$ and the current time step τ.
1. Solve the stationary problem (1.7) with $\phi = u_n$ and spatial accuracy TOL_p using some Galerkin approximation as outlined below.
2. Solve system (1.8) using the same representation of the solution obtained in Step 1 and compute an estimate of the temporal error eps_T using eq. (1.9).
3. Compute the new step size according to eq. (1.10).
4. If the error eps_T is smaller than TOL, set $t = t + \tau$, $\phi = u_{n+1}$, $\tau = \tau_{\text{new}}$ and go to Step 1. Otherwise, reduce the step size τ and go to Step 2.

Galerkin approximation. Given the parameterization of the Galerkin ansatz space at the current time t, an approximation u_n of $p(t_n)$, the current time step τ and the spatial tolerance TOL_p.
1. Compute the parameterization of a new Galerkin ansatz space $\mathcal{H}_r$ at time $t + \tau$.
2. Solve the stationary problems (1.7) and (1.8) in the new Galerkin ansatz space with resulting spatial discretization error eps_D satisfying $\text{eps}_D \leq \text{TOL}_p$.

Initial element-order-pattern. Given the decomposition of the overall discretization domain of the previous step.
1. Eliminate all domains with corresponding low polynomial order, i.e., domains that are possibly not necessary.
2. Add such eliminated domains to the neighbored domain with lowest order.

Final element-order-pattern. Given an initial decomposition of the discretization domain $D^{(m)}$ and an error estimate $\text{eps}_D^{(m)}$ of the local expansion on each rectangle $D^{(m)}$.
1. Decide on the basis of local error prediction

$$\text{eps}_D^{(m)} = \left\| p(t) - p_r^{(m)}(t) \right\|_{D^{(m)}},$$

whether the rectangle $D^{(m)}$ has to be divided or the present polynomial order has to be increased.

2. Compute a threshold value to decide which domains have to be changed (divided or order increase).

3. Build up the Galerkin equations (2.5) that define the expansion coefficients on each domain $D^{(m)}$ for all local approximations $p_r^{(m)}$ and solve them.

4. Estimate the overall approximation error eps_D according to (2.11).

5. If $\text{eps}_D < \text{TOL}_p$ then stop, otherwise go to Step 1.