

On the numerical solution of the chemical master equation with sums of rank one tensors

Markus Hegland¹Jochen Garcke²

(Received 20 January 2011; revised 4 August 2011)

Abstract

We show that sums of rank one tensors (or separable functions) representing the so-called Candecomp/Parafac or *CP-decomposition* is used effectively to solve the chemical master equations as in many cases the effective tensor rank of the probability distribution only grows slowly with time. Both theoretical bounds and computational experiments are presented which support this claim. The proposed numerical algorithm is thought to provide an effective tool for the computational study of stochastic biochemical systems involving large numbers of different chemical species.

Contents

1 Introduction

C629

<http://anziamj.austms.org.au/ojs/index.php/ANZIAMJ/article/view/3895> gives this article, © Austral. Mathematical Soc. 2011. Published August 10, 2011. ISSN 1446-8735. (Print two pages per sheet of paper.) Copies of this article must not be made otherwise available on the internet; instead link directly to this URL for this article.

<i>1 Introduction</i>	C629
2 Low rank approximations of $p(t)$	C631
3 The algorithm	C635
4 Computational experiments	C638
References	C642

1 Introduction

A well stirred chemical system is characterised by a vector $X(t)$ of copy numbers of its relevant chemical species. According to Kurtz and collaborators [2], such a system is modelled as a stochastic process $X(t)$ satisfying

$$X(t) = X(0) + \sum_{i=1}^k Y_i \left(\int_0^t \lambda_i(X(s)) ds \right) z_i. \quad (1)$$

Here $X(0)$ is the initial condition, $Y_i(\nu)$ are independent standard Poisson processes with parameter ν , $\lambda_i(x) \in \mathbb{R}$ is the propensity and $z_i \in \mathbb{Z}^d$ the stoichiometric vector of reaction i for $i = 1, \dots, k$. One can use this equation to generate simulation algorithms like the Gillespie method [3] and even the kinetic rate equations as by the application of the law of large numbers one gets $Y_i(\nu) \rightarrow \nu$ for $\nu \rightarrow \infty$. Section 2 uses a related consequence of the law of large numbers.

Here we consider the conditional distribution $p(t; x) = \Pr[X(t) = x \mid X(0)]$ of the stochastic process $X(t)$ defined by equation (1). If $x \in \mathbb{Z}^d$ is interpreted as an index then $p(t; x)$ is the x th entry of a tensor $p(t)$. Such a conditional distribution is used, for example, to compute the likelihood of data $X(t_1), \dots, X(t_N)$ in a maximum likelihood method for the determination of parameters of the propensities λ_i , see for example the article by Tian et al. [8].

Van Kampen [9] showed that $\mathbf{p}(\mathbf{t})$ satisfies the *chemical master equations* (CME)

$$\frac{d\mathbf{p}(\mathbf{t})}{dt} = \mathbf{A}\mathbf{p}(\mathbf{t}), \quad (2)$$

where

$$\mathbf{A}\mathbf{p} = \sum_{i=1}^k (\mathcal{S}_{z_i} - \mathbf{I}) \lambda_i \otimes \mathbf{p}, \quad (3)$$

and where $\mathcal{S}_z \mathbf{u}(\mathbf{x}) = \mathbf{u}(\mathbf{x} + \mathbf{z})$ denotes the shift along \mathbf{z} and $\lambda_i \otimes \mathbf{p}$ is the Hadamard or point-wise product.

Typically, $\mathbf{p}(\mathbf{t})$ is a \mathbf{d} -dimensional tensor with $M^{\mathbf{d}}$ entries where M is the size of state space in one dimension. This exponential dependency of the complexity on the dimension is the curse of dimension. However, there are tensors which can be represented in a sparse format and therefore do not suffer under this curse [6]. In particular, if

$$\mathbf{p}(\mathbf{t}; \mathbf{x}) = \prod_{i=1}^{\mathbf{d}} \mathbf{p}_i(\mathbf{t}; x_i)$$

for some \mathbf{p}_i where x_i is the i th component of \mathbf{x} , then one only requires the storage of $\mathbf{d}M$ numbers, namely M for each \mathbf{p}_i . This is a rank one tensor. More generally, a tensor of rank r is of the form

$$\mathbf{p}(\mathbf{t}; \mathbf{x}) = \sum_{j=1}^r s_j \prod_{i=1}^{\mathbf{d}} \mathbf{p}_{j,i}(\mathbf{t}; x_i). \quad (4)$$

The storage requirement here is $r\mathbf{d}M$ and there is no curse for low r . In the following we provide some evidence that such low rank r approximations of the solution of the master equation are effective.

A related approach to the one discussed here was suggested by Jahnke and Huisinga [5]. Their approach is based on the Dirac–Frenkel–McLachlan variational principle while the algorithm suggested here is based on more general

approximation with low rank tensors. We also provide some explanation why the chemical master equation admits low rank approximations. Alternatives to tensor product approximations were suggested by Munsky and Khammash [7] using the finite state projection method and by Hegland et al. [4] using (dimension) adaptive sparse grids. While especially the adaptive sparse grids have been shown to work well even for very high dimensions, the theory of these earlier approaches for high dimensions is less clear. As the chemical master equations have a natural low rank structure we think that the tensor product approaches are more promising for some very high dimensional problems, and in particular for modelling of reaction cascades.

Section 2 discusses when and why the solution of the chemical master equation can be approximated by low rank tensors. In particular, the law of large numbers is used to demonstrate the accuracy of (local) rank one approximations. Section 3 introduces the new algorithm which is based on a combination of a standard initial value problem solver and low rank compression. To show the potential of the approach, Section 4 provides some computational results and show that the new method works well at least for two dimensional problems where it appears to have a spectral-like complexity.

2 Low rank approximations of $p(t)$

When solving the CME for the conditional probability $p(t; x) = \Pr[X(t) = x \mid X(0) = x_0]$ one often observes that there exist ϵ -approximations for which the rank grows initially ($p(x, 0)$ has rank one), then has some phases where the rank is relatively stable and others where the rank decreases. We show under which conditions these situations arise. For this we consider ϵ -approximations obtained by approximating the propensities λ_i of the CME. The effect of such a modification on the solution is bounded by the following lemma.

Lemma 1. Let $\mathbf{p}(t)$ and $\tilde{\mathbf{p}}(t)$ solve the chemical master equations

$$\frac{d\mathbf{p}}{dt} = \sum_{i=1}^k (\mathbf{S}_{z_i} - \mathbf{I}) \lambda_i \circledast \mathbf{p}, \quad (5)$$

$$\frac{d\tilde{\mathbf{p}}}{dt} = \sum_{i=1}^k (\mathbf{S}_{z_i} - \mathbf{I}) \tilde{\lambda}_i \circledast \tilde{\mathbf{p}}, \quad (6)$$

for $\mathbf{p}(0) = \tilde{\mathbf{p}}(0)$. If $\mathbf{X}(t)$ is the stochastic process with distribution $\mathbf{p}(t)$ then

$$\|\mathbf{p}(t) - \tilde{\mathbf{p}}(t)\|_1 \leq 2 \sum_{i=1}^k \int_0^t \mathbb{E} \left[|\lambda_i(\mathbf{X}(s)) - \tilde{\lambda}_i(\mathbf{X}(s))| \right] ds.$$

Proof: Let \mathbf{A} and $\tilde{\mathbf{A}}$ be the operators of the equations (5) and (6), respectively such that

$$\mathbf{A}\mathbf{u} = \sum_{i=1}^k (\mathbf{S}_{z_i} - \mathbf{I}) \lambda_i \circledast \mathbf{u} \quad \text{and} \quad \tilde{\mathbf{A}}\mathbf{u} = \sum_{i=1}^k (\mathbf{S}_{z_i} - \mathbf{I}) \tilde{\lambda}_i \circledast \mathbf{u},$$

for $\mathbf{u} \in \ell_1(\mathbb{Z}^d)$. Let in the following $\|\mathbf{u}\|_1 = \sum_{x \in \mathbb{Z}^d} |\mathbf{u}(x)|$ be the $\ell_1(\mathbb{Z}^d)$ norm. A standard variation of the constant argument gives

$$\mathbf{p}(t) - \tilde{\mathbf{p}}(t) = \int_0^t e^{\tilde{\mathbf{A}}(t-s)} (\mathbf{A} - \tilde{\mathbf{A}}) \mathbf{p}(s) ds.$$

Using the triangle inequality and the ℓ_1 -contractivity of $e^{\tilde{\mathbf{A}}t}$ one gets the bound

$$\|\mathbf{p}(t) - \tilde{\mathbf{p}}(t)\|_1 \leq \int_0^t \|(\mathbf{A} - \tilde{\mathbf{A}})\mathbf{p}(s)\|_1 ds.$$

Another application of the triangle inequality combined with the contractivity of the shifts \mathbf{S}_{z_i} then shows

$$\|\mathbf{p}(t) - \tilde{\mathbf{p}}(t)\|_1 \leq 2 \sum_{i=1}^k \int_0^t \|\lambda_i \circledast \mathbf{p}(s) - \tilde{\lambda}_i \circledast \mathbf{p}(s)\|_1 ds.$$

The bound follows from the definitions of the ℓ_1 norm and expectation. ♠

If \mathbf{A} has operator rank one, that is $\mathbf{A} = \mathbf{A}_1 \otimes \cdots \otimes \mathbf{A}_d$, then the rank of the solution $\mathbf{p}(t)$ of the corresponding CME does not change over time. Lemma 1 then provides a condition under which a constant rank approximation of $\mathbf{p}(t)$ exists—when \mathbf{A} applied to $\mathbf{p}(t)$ is well approximated by a rank one operator.

If the propensities λ_i have tensor rank one, then the rank of $\lambda_i \circledast \mathbf{p}$ is the same as the rank of \mathbf{p} . As in this case the rank of $\mathbf{S}_{z_i} \lambda_i$ is also one, the rank of $\mathbf{A}\mathbf{p}$ (which has $2k$ terms) is $2k$ times the rank of \mathbf{p} . For sufficiently small time steps Δt one approximates $\mathbf{p}(t + \Delta t)$ using the Euler method as

$$\mathbf{p}(t + \Delta t) = \mathbf{p}(t) + \Delta t \mathbf{A} \mathbf{p}(t).$$

The rank of $\mathbf{p}(t + \Delta t)$ is then at most $2k + 1$ times the rank of $\mathbf{p}(t)$. Consequently we expect to find that ϵ -approximations of $\mathbf{p}(t)$ have a tensor rank which grows over time. While the discussion suggests exponential growth in time the growth is slower in practice.

In order to understand what may happen in practice, we consider a simple example with two reactions $\mathcal{S}_1 \rightarrow \mathcal{S}'_1$ and $\mathcal{S}_2 \rightarrow \mathcal{S}'_2$. The state is defined by the counts of \mathcal{S}_1 and \mathcal{S}_2 . Assume that the propensities are $\lambda_1(\mathbf{x}) = \kappa_1 x_1$ and $\lambda_2(\mathbf{x}) = \kappa_2 x_1 x_2 / \mathbf{n}_1$ where \mathbf{n}_1 is the initial count of \mathcal{S}_1 . Choose $\tilde{\lambda}_1 = \lambda_1$ and $\tilde{\lambda}_2(\mathbf{x}) = \lambda_2(\mathbb{E}[\mathbf{X}_1(t)], x_2) = \kappa_2 e^{-\kappa_1 t} x_2$. The identity $\mathbb{E}[\mathbf{X}_1(t)] = e^{-\kappa_1 t} \mathbf{n}_1$ is obtained from the explicit determination of the marginal distribution of $\mathbf{X}_1(t)$. This distribution satisfies a master equation with propensity $\lambda_1(x_1) = \kappa_1 x_1$ and is thus binomial. The approximating system has two statistically independent reactions where the second propensity depends on time. The approximating tensor $\tilde{\mathbf{p}}(t)$ thus has rank one.

Proposition 2. *Let $\mathbf{p}(t)$ and $\tilde{\mathbf{p}}(t)$ be the probability distributions for the two systems defined in the previous paragraph. Then the error is bounded by*

$$\|\mathbf{p}(t) - \tilde{\mathbf{p}}(t)\|_1 \leq \sqrt{\frac{2}{\mathbf{n}_1 \kappa_1}} \kappa_2 (1 - e^{-\kappa_1 t}) \sqrt{\int_0^t \mathbb{E}[\mathbf{X}_2(s)^2] ds}$$

Proof: An application of Lemma 1 gives


$$\|\mathbf{p}(t) - \tilde{\mathbf{p}}(t)\|_1 \leq 2 \int_0^t \mathbb{E} \left[\left| \lambda_2(\mathbf{X}(s)) - \tilde{\lambda}_2(\mathbf{X}(s)) \right| \right] ds.$$

Inserting the values of the propensities one gets by Cauchy–Schwarz

$$\begin{aligned} \mathbb{E} \left[\left| \lambda_2(\mathbf{X}(s)) - \tilde{\lambda}_2(\mathbf{X}(s)) \right| \right] &\leq \frac{\kappa_2}{n_1} \sqrt{\text{var}(\mathbf{X}_1(s))} \sqrt{\mathbb{E}[\mathbf{X}_2(s)^2]} \\ &= \kappa_2 n_1^{-1/2} \sqrt{e^{-\kappa_1 s} (1 - e^{-\kappa_1 s})} \sqrt{\mathbb{E}[\mathbf{X}_2(s)^2]}. \end{aligned}$$

Integration over t and another application of Cauchy–Schwarz gives

$$\begin{aligned} &\int_0^t \sqrt{e^{-\kappa_1 s} (1 - e^{-\kappa_1 s})} \sqrt{\mathbb{E}[\mathbf{X}_2(s)^2]} ds \\ &\leq \sqrt{\int_0^t e^{-\kappa_1 s} (1 - e^{-\kappa_1 s}) ds} \sqrt{\int_0^t \mathbb{E}[\mathbf{X}_2(s)^2] ds}, \end{aligned}$$

from which the result follows. 

One can see that the error of the approximation $\tilde{\lambda}_2$ first increases and then decreases with time. Interestingly, the error also decreases with the initial count n_1 of \mathcal{S}_1 , this is a consequence of the law of large numbers.

When the error grows to a certain threshold, one partitions the domain into two parts and does the rank one approximation for each part separately. (An alternative to using domain partitioning would be to use a partition of unity.) This is how one gets higher rank approximations and one can use the proposition to get error bounds for all the components and thus for $\mathbf{p}(t)$.

To further illustrate the consequence of the law of large numbers consider the original system from above. This was solved for different parameters n_1 using the CMEPY software mentioned in Section 4 to get the exact solution. An approximation was obtained using the singular value decomposition (SVD)




FIGURE 1: Ranks of truncated SVD approximation.

and truncating all components with singular values $\sigma_i < 10^{-5}$. One can see that the ranks of the resulting approximation decrease with increasing n_1 . Thus the probability of larger systems are approximated by tensors with smaller ranks, see Figure 1.

3 The algorithm

The numerical solution of the initial value problem (2) uses a discretisation of time $t_0 = 0 < t_1 < \dots < t_K$. For example, one may choose $t_k = k\tau$ but more generally, the grid widths $\tau_k = t_k - t_{k-1}$ are not all equal. Denote a numerical approximation to $p(t_k)$ by $p_k \in \mathbb{R}^\Omega$ where Ω contains the support of $p(t_k)$. As the differential equation associated with the initial value problem (2) is stiff, one needs to use stable solvers. To illustrate the suggested approach we use very simple solvers.

Furthermore, one needs a tensor class which has two operations: addition of tensors and (point-wise, Hadamard) multiplication. The sum of a rank r_1 and a rank r_2 tensor is (at most) of rank $r = r_1 + r_2$, whereas the Hadamard product is a tensor of rank $r_1 r_2$.

We consider the *implicit Euler or Rothe method* where the approximation \mathbf{p}_k is defined as the solution of the linear system of equations

$$(\mathbf{I} - \tau_k \mathbf{A})\mathbf{p}_k = \mathbf{p}_{k-1}, \quad k = 1, \dots, K.$$

The discretisation τ_k is chosen such that the error of the Rothe method is at most some ϵ .

The determination of \mathbf{p}_k is done with a *gradient descent method*. This class of iterative methods defines a sequence $\mathbf{p}_k^{(n)}$ by

$$\mathbf{p}_k^{(n+1)} = \mathbf{p}_k^{(n)} - \omega_n [(\mathbf{I} - \tau_k \mathbf{A})\mathbf{p}_k^{(n)} - \mathbf{p}_{k-1}].$$

Iterate from $n = 0, \dots, N$ and choose N such that the error of the descent method is again bounded by ϵ . In the analysis one utilises the fact that the \mathbf{p}_{k-1} is also only known approximately. For example, the parameter ω_n is chosen according to the steepest descent method.

The main computational ingredient is the matrix vector product

$$\mathbf{A}\mathbf{p} = \sum_{i=1}^k (\mathcal{S}_{z_i} - \mathbf{I})\Lambda_i \mathbf{p}.$$

As Λ_i are diagonal we first consider the values of (the tensor) $(\lambda_i(\mathbf{x}))_{\mathbf{x}}$. In the simplest case which uses counts of species for the states and the law of mass action all these tensors are of rank one. Consequently the $\Lambda_i \mathbf{p}$ all have the same rank as \mathbf{p} , thus no rank reduction is required. In other cases a low rank representation of the tensor containing the values of the propensities is explicitly computed; here an approximation might be necessary to keep the rank low while controlling the error made. In two dimensions one uses

Algorithm 1: $A^{\text{CP}}(\mathbf{p})$: Operation of A on \mathbf{p}

```

1 propensities  $\lambda_i$  are given in CP-format (exact or approximated initially)
2 shift operators  $S_{z_i}$  are given in CP-format
3  $A\mathbf{p} = 0$ 
4 for  $i \in \{1, \dots, k\}$  do
5    $A_i\mathbf{p} = (S_{z_i}\lambda_i) \otimes (S_{z_i}\mathbf{p}) - \lambda_i \otimes \mathbf{p}$ 
6   possible rank reduction of  $A_i\mathbf{p}$   $\triangleright \text{cp}^\#(A_i\mathbf{p}) = 2 \cdot \text{cp}^\#(\lambda_i) \cdot \text{cp}^\#(\mathbf{p})$ 
7    $A\mathbf{p} = A\mathbf{p} + A_i\mathbf{p}$ 
8 end
9 possible rank reduction of  $A\mathbf{p}$   $\triangleright \text{cp}^\#(A\mathbf{p}) = \sum_{i=1}^k \text{cp}^\#(A_i\mathbf{p})$ 
10 return  $A\mathbf{p}$ 

```

the singular value decomposition (SVD), and for more dimensions one can use alternating least squares or other procedures [6]. The employed rank is chosen adaptively to get a desired accuracy. This is possible as the rank of the exact solution is finite and small. The value $S_{z_i}\mathbf{p}$ is obtained by shifting all the components representing \mathbf{p} by the value of component z_i . This shift operation is rank invariant.

The algorithm $A^{\text{CP}}(\mathbf{p})$ to compute $A\mathbf{p}$ in the CP-format is shown in Algorithm 1. The rank of $A\mathbf{p}$ after the application of the (approximated) tensors $(\lambda_i(\mathbf{x}))_x$ can be as large as

$$2 \sum_{i=1}^k \text{cp}^\#(\lambda_i) \cdot \text{cp}^\#(\mathbf{p}),$$

where $\text{cp}^\#$ equals r in the representation (4) of a tensor. Therefore a rank reduction algorithm often needs to be applied.

The full procedure is shown in Algorithm 2. The residual used in the gradient descent algorithm also needs to be stored as a low rank tensor. Due the multiplication with A the ranks of the residual might grow as well. Therefore a rank reduction algorithm needs to be applied which again needs to keep

Algorithm 2: Implicit Euler in CP-format

```

1 initial distribution  $\mathbf{p}_0$  is given in CP-format
2 time interval  $[0, \mathbf{t}_K]$  is divided into time steps  $\mathbf{t}_0 = 0 < \mathbf{t}_1 < \dots < \mathbf{t}_K$ 
3 for  $\mathbf{k} \in \{1, \dots, K\}$  do ▷ implicit Euler
4    $\tau_{\mathbf{k}} = \mathbf{t}_{\mathbf{k}} - \mathbf{t}_{\mathbf{k}-1}$ 
5    $\mathbf{p}^{\text{it}} = \mathbf{p}_{\mathbf{k}-1}$ 
6    $\text{res} = \delta_{\mathbf{t}} \cdot \mathbf{A}^{\text{CP}}(\mathbf{p}^{\text{it}})$ 
7   while not good enough do ▷ grad. desc.:  $(\mathbf{I} - \tau_{\mathbf{k}}\mathbf{A})\mathbf{p}_{\mathbf{k}} = \mathbf{p}_{\mathbf{k}-1}$ 
8      $\text{res}^{\text{op}} = \text{res} - \tau_{\mathbf{k}} \cdot \mathbf{A}^{\text{CP}}(\text{res})$ 
9      $\alpha = (\text{res}^{\text{op}}, \text{res}) / (\text{res}^{\text{op}}, \text{res}^{\text{op}})$ 
10     $\mathbf{p}^{\text{it}} = \mathbf{p}^{\text{it}} + \alpha \cdot \text{res}$ 
11    possible rank reduction of  $\mathbf{p}^{\text{it}}$  ▷  $\text{cp}^{\#}(\mathbf{p}^{\text{it}}) = \text{cp}^{\#}(\mathbf{p}^{\text{it}}) + \text{cp}^{\#}(\text{res})$ 
12     $\text{res} = \text{res} - \alpha \cdot \text{res}^{\text{op}}$ 
13    possible rank reduct. of  $\text{res}$  ▷  $\text{cp}^{\#}(\text{res}) = \text{cp}^{\#}(\text{res}) + \text{cp}^{\#}(\text{res}^{\text{op}})$ 
14  end
15   $\mathbf{p}_{\mathbf{k}-1} = \mathbf{p}^{\text{it}}$ ; possible rank reduction of  $\mathbf{p}_{\mathbf{k}-1}$ 
16 end
```

the error of the approximation under control.

4 Computational experiments

We report on some experiments with a two dimensional problem with propensities

$$\lambda_i = 60e^{-t} \left(\frac{0.5}{x_i + x_{1-i}} + \frac{0.5}{x_i + 1000} \right) x_i, \quad i = 1, 2. \quad (7)$$

This model was suggested to us by Kevin Burrage. It models competing T-cell clonotypes. The basis for this study is the CMEPY software [1]. For our purposes we added low rank approximations to CMEPY. The system converges towards a bistable distribution where one clonotype dominates in

each of the branches, see Figure 2. It has some diagonal structure which is often unfortunate for tensor based approaches.

While smaller time steps seem to result in smaller ranks, the rank increases over the observation interval as seen in Figure 3(a). This tendency is also seen in Figure 3(b), where in addition one sees that the error is proportional to the threshold. The cost for a reduction of the error by a factor of ten is a rank increase by two. This means that by a linear increase in complexity one can reduce the error by a constant factor. This exponential error dependency is what is encountered for spectral methods and for radial basis function approximation and needs to be further investigated. (For grid based methods the error is usually proportional to a power of the cost.) Figure 3(c) shows that the largest observed rank of the residual in Algorithm 2 after thresholding is proportional to the step size, so smaller time steps allow a more compact representation. The final plot, Figure 3(d), provides an insight into the complexity of the solver as it displays the number of gradient descent steps required at every integration step. With smaller step sizes one requires fewer iterations but interestingly note that the actual number of iterations required is mostly very small. These last two observations together explain that the observed total run time decreases with smaller time steps, in contrast with what one would expect. But this observation has to be taken with a grain of salt since our current implementation is far from optimal.

Acknowledgements The work on CMEPY was supported by the ARC Centre in Bioinformatics. The work by the second author was supported through the MATHEON Project A18. We thank R. Fletcher-Costin for his work on CMEPY and S. Kamprath and M. Asher for their work on the tensor product approximation codes used in this project.







FIGURE 2: Solution p for the T-cell clonotypes model at time $t = 7.5$.




(a) solution rank over time.



(b) ranks versus error for different thresholds, each mark represents one time step.



(c) max. rank of update over time.



(d) iterations for gradient descent.

FIGURE 3: Some results for the T-cell clonotypes model.

References

- [1] CmePy Documentation. <http://fcostin.github.com/cmepy/>, 2010. C638
- [2] K. Ball, T. G Kurtz, L. Popovic, and G. Rempala. Asymptotic analysis of multiscale approximations to reaction networks. *Ann. Appl. Probab.*, 16(4):1925–1961, 2006. doi:10.1214/105051606000000420 C629
- [3] Daniel T. Gillespie. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *J. Comput. Phys.*, 22(4):403–434, 1976. doi:10.1016/0021-9991(76)90041-3 C629
- [4] Markus Hegland, Conrad Burden, Lucia Santoso, Shev MacNamara, and Hilary Booth. A solver for the stochastic master equation applied to gene regulatory networks. *J. Comput. Appl. Math.*, 205(2):708–724, 2007. doi:10.1016/j.cam.2006.02.053 C631
- [5] Tobias Jahnke and Wilhelm Huisinga. A dynamical low-rank approach to the chemical master equation. *Bull. Math. Biol.*, 70(8):2283–2302, 2008. doi:10.1007/s11538-008-9346-x C630
- [6] Tamara G. Kolda and Brett W. Bader. Tensor decompositions and applications. *SIAM Rev.*, 51(3):455–500, 2009. doi:10.1137/07070111X C630, C637
- [7] Brian Munsky and Mustafa Khammash. A multiple time interval finite state projection algorithm for the solution to the chemical master equation. *J. Comput. Phys.*, 226(1):818–835, 2007. doi:10.1016/j.jcp.2007.05.016 C631
- [8] T. Tian, S. Xu, J. Gao, and K. Burrage. Simulated maximum likelihood method for estimating kinetic rates in gene expression. *Bioinformatics*, 23(1):84–91, 2006. doi:10.1093/bioinformatics/btl552 C629

- [9] N. G. van Kampen. *Stochastic processes in physics and chemistry*, volume 888 of *Lecture Notes in Mathematics*. North-Holland Publishing Co., Amsterdam, 1981. C630

Author addresses

1. **Markus Hegland**, Mathematical Sciences Institute, The Australian National University, Australian Capital Territory 0200, AUSTRALIA.
<mailto:markus.hegland@anu.edu.au>
2. **Jochen Garcke**, Institut für Mathematik, Technische Universität Berlin, 10623 Berlin, GERMANY
<mailto:garcke@math.tu-berlin.de>