A provably convergent control closure scheme for the Method of Moments of the Chemical Master Equation

Vincent Wagner[1,a], **Robin Strässer**[1,b], **Frank Allgöwer**[b], **and Nicole Radde**[a,✉]

[1]Joint first authors
[a]Institute for Stochastics and Applications, University of Stuttgart, Pfaffenwaldring 9, 70569, Germany
[b]Institute for Systems Theory and Automatic Control, University of Stuttgart, Pfaffenwaldring 9, 70569, Germany

In this article, we introduce a novel moment closure scheme based on concepts from Model Predictive Control (MPC) to accurately describe the time evolution of the statistical moments of the solution of the Chemical Master Equation (CME). The Method of Moments, a set of ordinary differential equations frequently used to consider the first n_{m} moments, is generally not closed since lower-order moments depend on higher-order moments. To overcome this limitation, we interpret the moment equations as a nonlinear dynamical system, where the first n_{m} moments serve as states and the closing moments serve as control input. We demonstrate the efficacy of our approach using two example systems and show that it outperforms existing closure schemes. For polynomial systems, which encompass all mass-action systems, we provide probability bounds for the error between true and estimated moment trajectories. We achieve this by combining convergence properties of a priori moment estimates from stochastic simulations with guarantees for nonlinear reference tracking MPC. Our proposed method offers an effective solution to accurately predict the time evolution of moments of the CME, which has wide-ranging implications for many fields, including biology, chemistry, and engineering.

Chemical Master Equation | Method of Moments | Moment Closure | Stochastic Simulation Algorithm | Control Theory | Model Predictive Control

Correspondence: *nicole.radde@simtech.uni-stuttgart.de*

DRAFT

Introduction

The dynamics of well-stirred chemical reaction systems are described by the Chemical Master Equation (CME). The CME is a set of linear differential equations for the time-dependent probability distribution over the set of all possible configurations of the system. Since this set is usually large or even infinite except for very small systems, solving the CME directly is often not feasible. One remedy for this problem is to consider only the first n_{m} moments of this probability distribution, thereby reducing the number of equations from the number of configurations to the number of considered moments.

Two well-known methods to estimate the statistical moments of the CME solution are the Stochastic Simulation Algorithm (SSA) and the Method of Moments (MoM). The SSA (1, 2) generates sample paths from the underlying Markov process, which can be used to approximate either the full distribution or its moments via Monte Carlo integration. SSA-based moment approximation has several advantages. First, it inherently produces physically meaningful moment estimates, for example, means and variances are non-negative. Second, empirical moment estimates are consistent and thus converge in probability to the true moments in the large sample limit. On the other hand, approximating the CME solution or its moments only based on SSA samples is often tedious, as the computational load of generating each sample grows with the number of simulated molecules. Moreover, a huge number of samples is needed for an accurate estimate if for example some species have a high probability to become zero, if the system is stiff, or if it includes rare events.

The MoM is an infinite set of coupled Ordinary Differential Equations (ODEs) which describes the evolution of the statistical moments of the CME solution. Mathematically, the equation for each moment is obtained by multiplying the CME with a moment-specific test function and summation over all configurations, as described in (3). The solution of this moment ODE system can be obtained with state-of-the-art solvers and the needed computational power is nearly independent of the number of simulated molecules. However, one challenging aspect of the MoM ODEs is their interdependence structure. For systems with quadratic propensities, a moment of given order depends on the next higher-order moment, leading to an infinite dependency hierarchy, and the set of equations for the first n_{m} moments is not closed. It is worth mentioning that closing the MoM equations for the first n_{m} moments with the true higher-order moments, would they be known hypothetically, results in exact moment estimates. Plenty of solution approaches have been proposed to solve this closure problem. We will briefly outline some of them and refer the interested reader to (4) for an in-depth review of the general moment closure problem.

The most trivial closure technique simply neglects the influence of all higher-order moments on the considered lower-order moments and is thus effectively a truncation of the ODE system. We, therefore, refer to it as Truncation Closure (TC). TC is trivially implementable without any additional manual or computational effort.

Beyond this naïve approach, a popular class of methods assumes that the CME solution is similar to a specific family of distributions and uses the higher-order moments of these respective distributions for closure (see, e.g., (5–10)). Other

approaches express higher-order moments as potentially nonlinear functions of the considered lower-order moments (see, e.g., (11–15)). A third interesting class reconstructs the underlying distribution at one or several time steps, e.g., by exploiting the maximum entropy principle, and uses the reconstruction to infer the closure moments (16–18).

Many of the proposed closure schemes work very well in practice. For special cases, there are even theoretical results justifying specific closure schemes (19–21). However, since there is generally no relation between the trajectories of the closed moment equations and the CME, both are decoupled. Hence, moment closure methods suffer from undesired artifacts such as unphysical moments and potentially large deviations from the true moment trajectories, which constitutes the major drawback of closure methods. Even worse, except for systems with only reactions of up to order one, for which the moment equations are naturally closed, we do not have criteria at hand which allow for an a priori judgement of the quality of a moment closure solution based on the set of chemical reactions, the initial species amounts and number n_{m} of considered moments. Therefore, a closure scheme that generally and provably closes the MoM ODEs in accordance with the true CME solution would be highly desirable since it solves this major problem.

In this article, our aim is to address this decoupling issue by providing a closure that guarantees that the distance between the moment closure solutions and the true moments is bounded and moments are physically plausible. To achieve this ambitious goal, we employ a few SSA samples that are responsible for coupling the closed moment ODEs to the true CME solution. This idea is related to the work in (22) that uses a few SSA trajectories to roughly estimate the closure moments and subsequently uses Kalman filtering to eradicate the noise. While their proposed closure method relies on *sampled* closing moments that are typically affected by non-negligible noise, we propose to obtain the closing moments as a control input. To this end, we interpret the closed moment equations as a nonlinear system, in which the moments up to order n_{m} serve as states and the closing moments serve as a control input. For this *Control Closure (CC) scheme*, we can build on a variety of controller design methods for nonlinear systems from the literature (23). Particularly for the class of polynomial dynamics, which contains for example all mass-action systems, a popular approach for nonlinear control relies on sum-of-squares optimization, which reformulates non-negativity conditions as semi-definite programs (24, 25). However, sum-of-squares optimization is generally computationally challenging. Instead, we rely on Model Predictive Control (MPC) (26, 27) and, in particular, use concepts of nonlinear reference tracking MPC as presented in (28, 29). For CC, we characterize SSA-based estimates and confidence intervals of the moments by exploiting concentration inequalities of random variables. We then combine these with the MoM by formulating an optimal control problem in which the resulting optimal control input serves as the proposed moment closure. We validate our approach on two example systems, for which we obtain very good performance compared to TC and stand-alone SSA estimation. In particular, the higher-order moment estimates obtained by our CC scheme are indeed close to the true moments. Moreover, for polynomial systems, we provide guarantees that the moment trajectories inferred from CC satisfy the moment closure equations and lie in the given confidence intervals while staying close to the sampled moment estimates.

Notation: For a symmetric matrix A, we write $A \succ 0$ ($A \succeq 0$) if A is positive (semi-)definite. We write $\|x\|_A$ for the weighted norm $x^\top A x$ of a vector x with $A \succ 0$. By $\mathbb{N}_0$ and $\mathbb{R}_{\geq 0}$ we denote the natural numbers including zero and all non-negative real numbers, respectively. We use the vectorial index $\alpha \in \mathbb{N}_0^n$ with $|\alpha| = \alpha_1 + ... + \alpha_n$ to define for a vector $x \in \mathbb{R}^n$ the monomial $x^\alpha = x_1^{\alpha_1} \cdots x_n^{\alpha_n}$. We denote the class of functions $\alpha : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$, which are continuous, strictly increasing, and satisfy $\alpha(0) = 0$ by $\mathcal{K}$. We write $\mathcal{K}_\infty$ for the set of functions $\alpha \in \mathcal{K}$ satisfying $\lim_{k\to\infty} \alpha(k) = \infty$. By $\mathcal{L}$ we denote the class of functions $\beta : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$, which are continuous and decreasing with $\lim_{k\to\infty} \beta(k) = 0$.

SSA trajectories as control reference for moment closure

In this section, we present a CC scheme to close the moment equations of the CME. Subsection 2.1 first introduces an interpretation of the closed moment equations from the perspective of control theory. Second, we derive probability bounds for SSA-based moment estimation in Subsection 2.2, which define a confidence region that contains the true moments with a certain user-defined probability. Subsequently, we use these SSA-based estimates and probability bounds in Subsection 2.3 to formulate an optimization problem whose solution is the best possible estimate of the moments according to a pre-defined optimization criterion, while satisfying the closed moment equations, staying close to the SSA-based moment estimates, and lying inside the confidence region. For this optimization problem, we then prove the convergence of the CC to the true CME moments.

System theoretic interpretation of the closed moment equations. The CME describes for each configuration $z \in Z(\Omega)$ the time-dependent probability $p(z,t)$ as

$$\frac{\partial p(z,t)}{\partial t} = \sum_{j=1}^{n_r} a_j(z-\nu_j)p(z-\nu_j,t) - \sum_{j=1}^{n_r} a_j(z)p(z,t),$$

with reaction propensities $a_j(z)$, configuration change vectors ν_j, and sums over all feasible reactions. Moment equations are obtained by multiplication with a moment-specific test function $T(z) : Z(\Omega) \to \mathbb{R}$ and summing over all configurations (3), i.e.,

$$\sum_{z\in Z(\Omega)} T(z)\frac{\partial p(z,t)}{\partial t} = \sum_{j=1}^{n_r} \mathbb{E}_Z\left[(T(z+\nu_j) - T(z))a_j(z)\right].$$

For example, for $T(z) = z_i$ and $T(z) = (z_k - \mathbb{E}_Z(z_k))(z_l - \mathbb{E}_Z(z_l))$, respectively, we get for the expectation $\mathbb{E}_{Z_i}(t)$ and

the covariance $\mathbb{C}_{Z_k,Z_l}(t)$ between Z_k and Z_l

$$\frac{\mathrm{d}\mathbb{E}_{Z_i}(t)}{\mathrm{d}t} = \sum_{j=1}^{n_r} \nu_j^i \mathbb{E}_Z(a_j(z)) \tag{1a}$$

$$\frac{\mathrm{d}\mathbb{C}_{Z_k,Z_l}(t)}{\mathrm{d}t} = \sum_{j=1}^{n_r} \{\nu_j^k \mathbb{E}_Z\left[(z_l - \mathbb{E}_{Z_l}(t))a_j(z)\right] \tag{1b}$$
$$+ \nu_j^l \mathbb{E}_Z\left[(z_k - \mathbb{E}_Z(z_k))a_j(z)\right]\} + \sum_{j=1}^{n_r} \nu_j^k \nu_j^l \mathbb{E}_Z(a_j(z)),$$

with ν_j^k denoting the k-th component of the vector ν_j (3).
The general moment equations Eq. (1) can be interpreted as a controlled system of the form

$$\dot{x}(t) = f(x(t), u(t)), \quad x(0) = x_0, t \geq 0, \tag{2}$$

where $x \in \mathbb{R}^n$ contains the first n_m moments that appear as states and $u \in \mathbb{R}^m$ defines the control input and corresponds to the higher-order moments that are needed to close the system. Although the proposed CC scheme is in principle valid for arbitrary dimensions and moment orders, for ease of notation we restrict ourselves to systems with only one chemical species ($n_\mathrm{c} = 1$) for which we are interested in the expected value and the variance ($n_\mathrm{m} = 2$), i.e.,

$$\mathbb{E}_Z(t) = \sum_{k \in Z(\Omega)} k\, p(Z = k, t)$$
$$\mathbb{V}_Z(t) = \sum_{k \in Z(\Omega)} (k - \mathbb{E}_Z(t))^2\, p(Z = k, t).$$

In this case, a second order Taylor series expansion of $a_j(z)$ about $z = \mathbb{E}_Z(t)$ in Eq. (1) leads to (see Appendix A)

$$\frac{\mathrm{d}\mathbb{E}_Z(t)}{\mathrm{d}t} \approx \sum_{j=1}^{n_r} \nu_j \left[a_j(\mathbb{E}_Z(t)) + \frac{\mathrm{d}^2 a_j(\mathbb{E}_Z(t))}{\mathrm{d}z^2} \frac{\mathbb{V}_Z(t)}{2}\right] \tag{3a}$$

$$\frac{\mathrm{d}\mathbb{V}_Z(t)}{\mathrm{d}t} \approx \sum_{j=1}^{n_r} \Bigg(\nu_j^2 a_j(\mathbb{E}_Z(t)) \tag{3b}$$
$$+ \mathbb{V}_Z(t)\left[2\nu_j \frac{\mathrm{d}a_j(\mathbb{E}_Z(t))}{\mathrm{d}z} + \frac{\nu_j^2}{2} \frac{\mathrm{d}^2 a_j(\mathbb{E}_Z(t))}{\mathrm{d}z^2}\right]$$
$$+ \mathbb{S}_Z(t)\nu_j \frac{\mathrm{d}^2 a_j(\mathbb{E}_Z(t))}{\mathrm{d}z^2}\Bigg).$$

In this representation, the skewness

$$\mathbb{S}_Z(t) = \sum_{k \in Z(\Omega)} (k - \mathbb{E}_Z(t))^3\, p(Z = k, t)$$

is the closure moment.
For mass-action systems with reactions of at most quadratic propensities, the MoM ODEs are at most quadratic in the moments, i.e., there are no products of more than two moments. Moreover, the moments of order $(n_\mathrm{m} + 1)$ affect only the n_m-th moment and occur only linearly in the corresponding ODE (3). In this case, the Taylor expansion Eq. (3) is exact and can be used for further specification of the dynamical system Eq. (2) to

$$\dot{x}(t) = f(x(t)) + Bu(t), \tag{4}$$

where $n = \sum_{k=1}^{n_\mathrm{m}} n_\mathrm{c}^k$, $m = n_\mathrm{c}^{n_\mathrm{m}+1}$, $B = \begin{bmatrix} 0 & \tilde{B}^\top \end{bmatrix}^\top \in \mathbb{R}^{n \times m}$ with $\tilde{B} \in \mathbb{R}^{(n_\mathrm{c}^{n_\mathrm{m}}) \times m}$, and $f : \mathbb{R}^n \to \mathbb{R}^n$ is a quadratic drift term, i.e.,

$$f(x) = \begin{bmatrix} f_1(x) & \cdots & f_n(x) \end{bmatrix}^\top$$

with

$$f_j(x) = \sum_{\substack{\alpha \in \mathbb{N}_0^n \\ |\alpha| \leq 2}} s_\alpha x^\alpha,\ s_\alpha \in \mathbb{R}\ \forall j = 1, \ldots, n.$$

For $n_\mathrm{c} = 1$ and $n_\mathrm{m} = 2$, the dynamical systems Eq. (2) and Eq. (4) are of dimension $n = 2$, the closing moment u is of dimension $m = 1$, and the matrix B in Eq. (4) reduces to a vector, i.e., $B = \begin{bmatrix} 0 & \cdots & 0 & \tilde{B} \end{bmatrix}^\top$ with $\tilde{B} \in \mathbb{R}$.
With this simplification, we finally have a two-dimensional controlled system for the expectation value $\mathbb{E}_Z(t)$ and the variance $\mathbb{V}_Z(t)$, which reads

$$\dot{x}_1(t) = f_1(x_1(t), x_2(t)) \tag{5a}$$
$$\dot{x}_2(t) = f_2(x_1(t), x_2(t)) + \tilde{B}u(t), \tag{5b}$$

where $f_{1,2}(x_1, x_2)$ are polynomial functions in their arguments.

Probability bounds for SSA-based moment estimation. Empirical moment estimates $\hat{\mathbb{E}}_Z^{(N)}(t)$ from SSA sample paths $z^{(i)}(t)$, $i = 1, \ldots, N$, are consistent and converge for each time instance t^* in probability to the true moment $\mathbb{E}_Z(t^*)$ for large sample sizes, i.e.,

$$\operatorname*{plim}_{N \to \infty} \left(\hat{\mathbb{E}}_Z^{(N)}(t^*) = \frac{1}{N} \sum_{i=1}^{N} z^{(i)}(t)\right) = \mathbb{E}_Z(t^*).$$

In an attempt to quantify this convergence for a finite number of SSA trajectories N, we are interested in a probabilistic bound for $\left|\hat{\mathbb{E}}_Z^{(N)}(t^*) - \mathbb{E}_Z(t^*)\right|$, or in other words, confidence intervals of the mean $\mathbb{E}_Z(t^*)$. Estimating these intervals is one of the most fundamental questions of frequentist statistics and consequently, the amount of related literature is vast. Ideally, we would like to present a probabilistic bound of the form

$$P\left(\left|\hat{\mathbb{E}}_Z^{(N)}(t^*) - \mathbb{E}_Z(t^*)\right| \geq a\right) \leq \mathcal{B}\left(a, N, \{z^{(i)}(t^*)\}_{i=1,\ldots,N}\right),$$

where the bounding function $\mathcal{B}$ only depends on the desired width $a > 0$ of the confidence interval, the number of drawn samples N, and the samples themselves. However, it seems to be impossible to obtain such a purely sample-based bound $\mathcal{B}$. This subsection, therefore, outlines the most prominent approaches together with their respective additional assumptions and refers the interested reader to relevant sources for more information. We would like to stress that the proposed

CC scheme does not rely on any particular bound. Thus, every applicant of our technique is able to choose bounds according to each individual situation.
One of the most prominent representatives of this class of inequalities is the Chebyshev-Inequality (CI)

$$P\left(\left|\hat{\mathbb{E}}_Z^{(N)}(t^*)-\mathbb{E}_Z(t^*)\right|\geq a\right)\leq \min\left(1,\frac{\mathbb{V}_Z(t^*)}{Na^2}\right)\quad \forall a\in\mathbb{R}, \tag{6}$$

or closely related variants thereof (30). However, its bound depends on the unknown *true* variance $\mathbb{V}_Z(t^*)$.
In a similar fashion, the Central Limit Theorem (CLT) can be combined with the Berry-Esseen theorem (see (31, 32)) to show

$$\begin{aligned} &P\left(\left|\hat{\mathbb{E}}_Z^{(N)}(t^*)-\mathbb{E}_Z(t^*)\right|\geq a\right) \\ &\leq \min\left(1,1-\operatorname{erf}\left(\sqrt{\frac{Na^2}{2\mathbb{V}_Z(t^*)}}\right)+\frac{2c_{\mathrm{be}}\rho}{\sqrt{N\mathbb{V}_Z(t^*)^3}}\right), \end{aligned} \tag{7}$$

which is detailed in Appendix B. Here, c_{be} is the Berry-Esseen constant that is known to be upper bounded by 0.4748 (33), erf denotes the error function, and ρ is the *true* third centralized absolute moment

$$\rho=\sum_{k\in Z(\Omega)}\left|k-\mathbb{E}_Z(t^*)\right|^3\,P_{Z=k}(t)\,.$$

Again, knowledge of the *true* variance $\mathbb{V}_Z(t^*)$ and the just-defined moment ρ is necessary to obtain the desired bound.
In contrast, Hoeffding's Inequality (HI) (34) provides a bound independent of higher-order moments but instead requires bounded states $0\leq Z^{(i)}(t)\leq b$, $b\in\mathbb{R}$. It reads

$$P\left(\left|\hat{\mathbb{E}}_Z^{(N)}(t^*)-\mathbb{E}_Z(t^*)\right|\geq a\right)\begin{cases}\leq \min\left(1,2e^{-\frac{2a^2N}{b^2}}\right) & a\leq b,\\ =0 & a>b.\end{cases} \tag{8}$$

While this boundedness is a strong assumption, the number of system molecules is often bounded by conservation laws anyways. In case the system can indeed grow infinitely, a configuration state truncation can provide such an upper bound b, which is, however, not uniquely defined. Both named possibilities will be demonstrated in one of our example systems, respectively.
By defining an upper bound $\alpha\in(0,1)$ for the probability in one of the inequalities Eq. (6), Eq. (7), and Eq. (8), each inequality can be resolved for the parameter a. In case of the HI, this leads to

$$a(\alpha)=b\sqrt{\frac{\ln 2-\ln\alpha}{2N}}$$

exemplary, and the true moment lies with probability $1-\alpha$ in the symmetric interval $[\hat{\mathbb{E}}_Z^{(N)}(t)-a(\alpha),\hat{\mathbb{E}}_Z^{(N)}(t)+a(\alpha)]$ around the SSA-estimate, which serves as $(1-\alpha)$ confidence interval. In addition, physical constraints such as the non-negativity of expectation and variance might be used to tighten these intervals even further if desired. We denote the union of confidence intervals for all t^* in a given time interval $\mathcal{T}:=[0,T_{\mathrm{end}}]$ with $T_{\mathrm{end}}>0$ as confidence region, which describes a symmetric tube around $\hat{\mathbb{E}}_Z^{(N)}(t)$ for $t\in\mathcal{T}$. For a bounded state space, also the variance could be constrained. While we restricted this subsection to bounds for $\mathbb{E}_Z(t^*)$, many results can be generalized to higher-order moments and similarly included in the CC scheme presented later. For these and many other related results, we can (upon many others) recommend the textbook of Bercu et al. (35) and the articles of Savage (30), Nagaev (36), and Slagle (37).

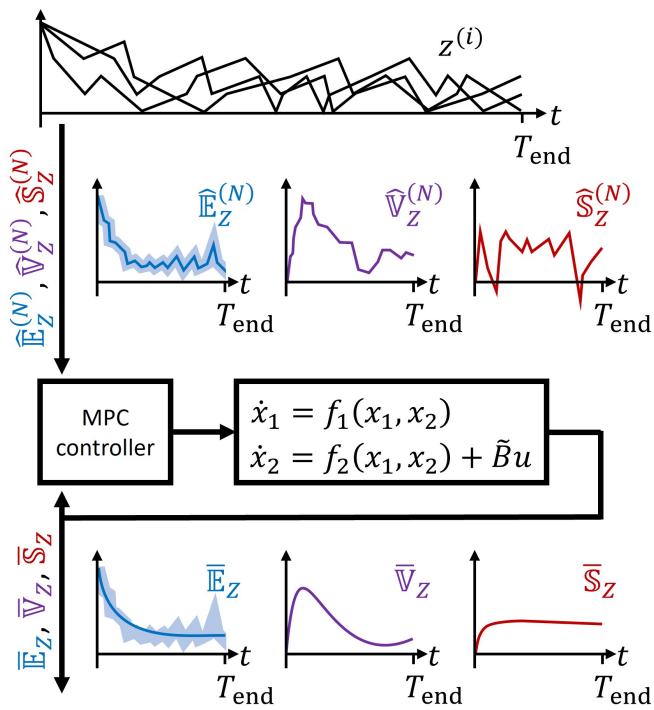

Fig. 1. *Control closure scheme for one chemical species and two considered moments.* According to the dynamics in Eq. (5), the state $x\in\mathbb{R}^2$ consists of expectation and variance, and the skewness, which constitutes the input $u(t)$, appears in f_2 as a linear term with factor $\tilde{B}$. The controller takes SSA-based estimations $\hat{\mathbb{E}}_Z^{(N)}(t)$ and $\hat{\mathbb{V}}_Z^{(N)}(t)$ and, if applicable, respective confidence regions as inputs and uses an MPC algorithm that returns moment estimates $\bar{x}=\left(\bar{\mathbb{E}}_Z,\bar{\mathbb{V}}_Z\right)$, and $\bar{u}=\bar{\mathbb{S}}_Z$. The time indication is omitted for brevity.

Optimal CC scheme for CME moment estimation. Based on the system theoretic representation Eq. (4) of the moment equations as a dynamical system, we formulate a CC scheme that relies on a priori SSA stand-alone moment estimates and derived confidence regions and uses MPC concepts to find an optimal input $\bar{u}$. In particular, we formulate an optimization problem to find a control input $\bar{u}$ which guarantees that the MoM with CC leads to estimated moments that minimally deviate from the SSA-based moment estimation $\hat{\mathbb{E}}_Z^{(N)}(t)$ with respect to a specified cost function and are guaranteed to lie within the pre-defined confidence region.
A graphical overview of the proposed CC is depicted in Fig. 1. Given SSA trajectories $z^{(i)}(t)$ over the finite time interval $\mathcal{T}$, we obtain estimates of the first $(n_{\mathrm{m}}+1)$ moments. We collect the first n_{m} moments in $\hat{x}(t)\in\mathbb{R}^n$ and all moments of order $(n_{\mathrm{m}}+1)$ in $\hat{u}(t)\in\mathbb{R}^m$ for $t\in\mathcal{T}$. Where applicable, we then define the constraint sets $\mathcal{X}(t)\subseteq\mathbb{R}^n$ and

$\mathcal{U}(t) \subseteq \mathbb{R}^m$ as the SSA-based confidence regions in Subsection 2.2. We want to ensure $(x(t), u(t)) \in \mathcal{X}(t) \times \mathcal{U}(t)$ for all $t \in \mathcal{T}$. Both $\hat{x}$ and $\hat{u}$ as well as $\mathcal{X}(t)$ and $\mathcal{U}(t)$ influence the MPC controller, which is explained in the following and returns moment estimates $\bar{x}(t)$ and $\bar{u}(t)$ that lie within $\mathcal{X}(t)$ and $\mathcal{U}(t)$.

As we seek the optimal moment estimation on the interval $\mathcal{T}$, we need to define an optimality criterion first. To this end, we define the open-loop cost function

$$J(x(0), u(\cdot)) = \int_0^{T_{\text{end}}} \ell(t, x(t), u(t)) \mathrm{d}t \tag{9}$$

with the stage cost

$$\begin{aligned} \ell(t, x(t), u(t)) = \|x(t) - \hat{x}(t)\|_Q + \|u(t) - \hat{u}(t)\|_R \\ + \|u(t) - u(t-\delta)\|_S \end{aligned} \tag{10}$$

for some user-defined positive definite $Q \in \mathbb{R}^{n \times n}$, positive semi-definite $R, S \in \mathbb{R}^{m \times m}$, and sampling period $\delta > 0$. These matrices can be chosen, e.g., to weigh some moments more than others. If there is no estimate $\hat{u}(t)$ of the closing moments, we can just choose $R = 0$ or $\hat{u}(t) = 0$. The last term in Eq. (10) is used to penalize the rate of change of the resulting input u. The matrix S can be seen as a tuning parameter: For a larger weight S, the difference in u at two consecutive time points $t - \delta$ and t is penalized more, and, hence, the optimization yields a smoother trajectory for the control input.

Definition 1: For weights $Q \succ 0$, $R \succeq 0$, and $S \succeq 0$ and an initial condition $x(0) = x_0$, the optimal reachable trajectory $(x^*(t), u^*(t))$ on the interval $\mathcal{T}$ is the minimizer of

$$\min_{u(\cdot)} J(x(0), u(\cdot)) \tag{11a}$$

$$\begin{aligned} \text{s.t. } \dot{x}(t) &= f(x(t)) + Bu(t), \\ &(x(t), u(t)) \in \mathcal{X}(t) \times \mathcal{U}(t), \ \forall t \in \mathcal{T}, \\ &x(0) = x_0. \end{aligned} \tag{11b}$$

This definition characterizes the best possible moment estimation of the CME according to the given SSA samples and the cost function in Eq. (9) while additionally satisfying the closed moment equations and staying inside the confidence region. Hence, the solution $u^*(t)$ is the optimal CC for the moment equations on $\mathcal{T}$.

Guaranteed convergent CC using a short-horizon optimization. For an increasing size of the interval $\mathcal{T}$, the nonlinear optimization problem in Eq. (11) becomes numerically challenging. To overcome this issue, we present in the following a short-horizon optimization based on MPC which is guaranteed to converge to the optimal reachable trajectory $(x^*(t), u^*(t))$ without explicitly calculating it. In particular, we use an MPC scheme based on the results in (29), where the authors propose a reference tracking MPC scheme without terminal ingredients for *unreachable* reference trajectories of discrete-time nonlinear systems based on techniques from economic MPC (38, 39). This MPC scheme ensures practical stability of the *unknown* optimal reachable reference $(x^*(t), u^*(t))$. Based on this, the results were translated to the case of continuous-time nonlinear systems in (28), which makes it suitable for our setting.

To obtain an optimal estimate of the CME moments which is close to the SSA-based estimation, we want to track the optimal reachable trajectory Eq. (11) without explicitly computing $(x^*(t), u^*(t))$. To this end, we define the short-horizon cost function

$$J_T(x(t), t, u(\cdot)) = \int_t^{t+T} \ell(t, x(t), u(t)) \mathrm{d}t$$

and propose the following: At each time $t \geq 0$, given the current state $x(t)$ and an admissible reference $(\hat{x}(t), \hat{u}(t)) \in \mathcal{X}(t) \times \mathcal{U}(t)$, we solve a short-horizon MPC scheme characterized by the optimization problem

$$\min_{\bar{u}(\cdot;t)} J_T(x(t), t, \bar{u}(\cdot;t)) \tag{12a}$$

$$\begin{aligned} \text{s.t. } \dot{\bar{x}}(\tau;t) &= f(\bar{x}(\tau;t)) + B\bar{u}(\tau;t), \\ &(\bar{x}(\tau;t), \bar{u}(\tau;t)) \in \mathcal{X}(t) \times \mathcal{U}(t), \\ &\bar{x}(t;t) = x(t), \\ &\forall \tau \in [t, t+T] \end{aligned} \tag{12b}$$

with the prediction horizon T satisfying $\delta \leq T \leq T_{\text{end}}$. The optimal trajectory resulting from Eq. (12) is denoted by $(\bar{x}^*(\cdot;t), \bar{u}^*(\cdot;t))$, while the control input u is chosen as $\bar{u}(\tau) = \bar{u}^*(\tau;t)$ for $\tau \in [t, t+\delta)$. The MPC problem Eq. (12) is solved repeatedly after each sampling period δ. The closed-loop dynamics using the MPC controller read then

$$\dot{\bar{x}}(t) = f(\bar{x}(t)) + B\bar{u}(t). \tag{13}$$

Due to the finite time horizon T, the solution to Eq. (12) does in general not exactly follow the optimal trajectory $(x^*(t), u^*(t))$. While the application of our proposed CC based on Eq. (12) is generally possible and requires only the definition of some hyperparameters, we can even establish, under mild regularity conditions, recursive feasibility of the MPC problem Eq. (12).

Definition 2: The MPC problem Eq. (12) is called *feasible* at time t if there exists at least one $\bar{u}(\cdot;t)$ such that the constraints are satisfied.

Further, we show that the resulting closed-loop system practically tracks the optimal reachable reference $(x^*(t), u^*(t))$ based on results in (28). Moreover, the proposed CC recovers the true moments in the limit of infinitely many sample trajectories, which follows from the consistency of the empirical moments estimates. Before stating our main theoretical result, we make the following assumption.

Assumption 3: System Eq. (4) representing the MoM is uniformly suboptimally operated off the trajectory x^* according to (39, Def. 12).

If the minimizer $(x^*(t), u^*(t))$ is unique, then Assumption 3 is not restrictive due to the positive definite weight matrix Q using some turnpike properties (compare (29, Sec. IV-D)). The following result establishes the theoretical properties of the proposed CC scheme.

Algorithm 1 Control closure scheme for MoM

Input:

Initial condition $z(0)$, N SSA trajectories $z^{(i)}(t)$ on the interval $\mathcal{T} = [0, T_{\text{end}}]$, sampling interval $\delta > 0$, prediction horizon $T \leq T_{\text{end}}$, weight matrices $Q \succ 0$, $R \succeq 0$, $S \succeq 0$.

Output:

Optimal CME moment estimation $(\bar{x}(t), \bar{u}(t))$ for $t \in \mathcal{T}$.

SSA-based estimation:

1: Use $z^{(i)}(t)$ to calculate a bias-free estimation of all moments up to order $(n_{\text{m}} + 1)$. Collect these in $\hat{\mathbb{E}}_Z^{(N)}(t), \hat{\mathbb{V}}_Z^{(N)}(t), \hat{\mathbb{S}}_Z^{(N)}(t), \ldots$

2: Set the MPC internal variables by allocating the first n_{m} moments in $\hat{x}(t)$ and the moments of order $(n_{\text{m}} + 1)$ in $\hat{u}(t)$ for $t \in \mathcal{T}$.

3: If applicable, define the confidence regions described in Section 2.2 as constraint sets $\mathcal{X}(t)$ and $\mathcal{U}(t)$.

Short-horizon optimization:

4: Initialize the estimate $\bar{x}(0) = x(0)$.

5: **for** $k = 0, ..., (T_{\text{end}} - T)/\delta$ **do**

6: Update current sampling time $t_k = \delta k$.

7: Initialize current estimation $\bar{x}(t_k; t_k) = \bar{x}(t_k)$.

8: Solve Eq. (12) to obtain the optimal prediction $(\bar{x}^*(\tau; t_k), \bar{u}^*(\tau; t_k))$ for $\tau \in [t_k, t_k + T]$.

9: Apply the CC input $\bar{u}(t) = \bar{u}^*(t; t_k)$ for $t \in [t_k, t_k + \delta)$.

10: Obtain the optimal moments $\bar{x}(t)$ for $t \in [t_k, t_k + \delta)$ via Eq. (13).

11: **end for**

Theorem 1: Suppose Assumption 3 holds. Then, there exists a $\tilde{T}_0$ such that for all $T > \tilde{T}_0$ the CC $\bar{u}(t)$ based on the MPC optimization Eq. (12) leads to a moment estimation that converges to a neighborhood of the *optimal* MoM-based estimates Eq. (11) of the CME and is guaranteed to satisfy physical properties and the given confidence region of the moments. Moreover, the CC recovers the *true* moments for infinitely many sample trajectories, i.e., $N \to \infty$.

Proof: See Appendix C for the proof and the precise definition of the achieved neighborhood to which the system converges. ■

The proposed optimal estimation of the CME moments based on the presented CC scheme is summarized in Algorithm 1. Note that the lower bound $\tilde{T}_0$ in Theorem 1 is, in general, unavailable since it depends on existent but typically unknown functions. Hence, the theorem shows the existence of a sufficiently long prediction horizon without stating the horizon explicitly. Nevertheless, for a fixed sampling interval δ, the closed-loop system stays around the optimal trajectory $(x^*(t), u^*(t))$, which is as close as possible to the unreachable reference $(\hat{x}(t), \hat{u}(t))$, given that a large enough prediction horizon is used. We stress further that when $\delta \to 0$, T should go to infinity (28). Possible future work could use ideas from (40), which presents a nonlinear MPC scheme without terminal ingredients for tracking possibly unreachable setpoints using artificial steady-states (41). The therein proposed results use less restrictive assumptions and might be exploited to indeed derive computable bounds on the necessary prediction horizon, which is beyond the scope of this work.

Simulation examples

In this section, we validate the performance of our CC scheme on two example systems. Both considered systems converge to a steady-state distribution on their respective configuration-transition graph, as discussed in (18). Algorithm 1 was implemented in Matlab using MPCTools (42) with its interface to the nonlinear optimization software CasADi (43).

System 1 The first exemplary system consists of a constant production rate and quadratic degradation and is therefore one of the most simple systems involving quadratic propensities. The reactions read

$$\emptyset \xrightarrow{\mathrm{k}_1} \mathrm{A} \qquad 2\,\mathrm{A} \xrightarrow{\mathrm{k}_2} \emptyset$$

and the MoM ODEs are given by

$$\begin{aligned}\frac{\mathrm{d}\mathbb{E}_Z(t)}{\mathrm{d}t} &= k_1 + 2k_2\mathbb{E}_Z(t) - 2k_2\mathbb{E}_Z^2(t) - 2k_2\mathbb{V}_Z(t),\\ \frac{\mathrm{d}\mathbb{V}_Z(t)}{\mathrm{d}t} &= k_1 - 4k_2\mathbb{E}_Z(t) + 4k_2\mathbb{E}_Z^2(t)\\ &\quad - 8k_2\mathbb{E}_Z(t)\mathbb{V}_Z(t) + 8k_2\mathbb{V}_Z(t) - 4k_2\mathbb{S}_Z(t).\end{aligned}$$

In principle, there is no limit for the number of molecules in System 1 and the space of all possible system configurations $Z(\Omega)$ is, therefore, $\mathbb{N}_0$. In order to still solve the CME numerically, we truncate $Z(\Omega)$ to 31 system configurations, leading to $Z(\Omega) = \{0, ..., 30\}$. While this seems to be a very drastic measure, the probability mass not captured by the first 31 configurations is indeed negligible in our setting, which justifies this approximation.

A comparison of the performance of different approaches for moment trajectory estimation for System 1 is shown in Fig. 2 for expectation value (top), variance (top second), and skewness (top third). The CME lines represent the true, numeric moment solutions and serve as references. SSA-based moment estimates $\hat{x}(t)$ and $\hat{u}(t)$ are shown for $N = 10$ and $N = 100$ SSA sample paths. The results of TC fails here completely.

The result of our CC scheme utilize the average over $N = 10$ SSA sample paths as a reference in the objective function. Confidence regions

$$\mathcal{X}_1(t) = [\max(0, \hat{x}_1(t) - a(\alpha = 0.95)), \hat{x}_1(t) + a(\alpha = 0.95)]$$

were set for the expectation value by using HI and CI bounds. Please note that the CI bounds require knowledge of the true variance $\mathbb{V}_Z$ and central absolute third-order moment ρ to be valid. As those quantities are typically not available, the CI bounds depicted in Fig. 2, 3 are calculated using SSA-based estimators (e.g. $\hat{\mathbb{V}}_Z^{(10)}$) of the respective moments and therefore only approximate the true CI bounds.

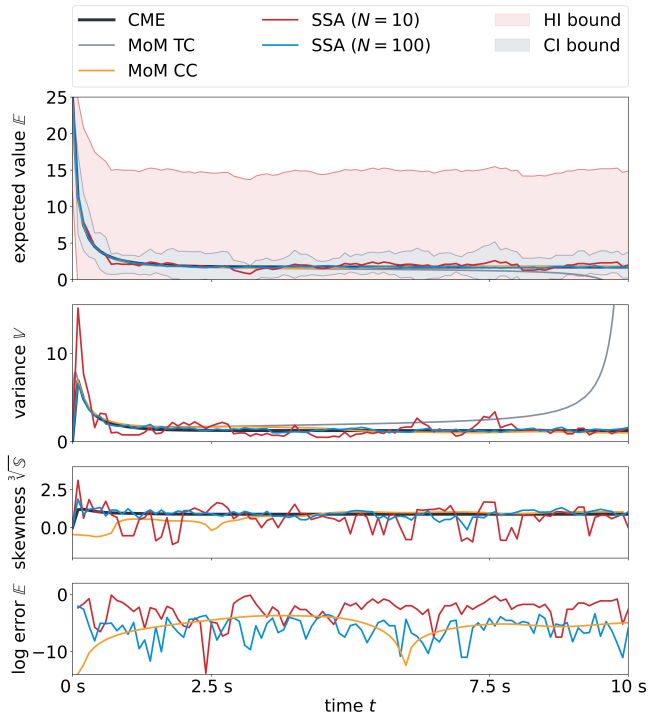

Fig. 2. Comparison of moment estimation approaches for System 1. Top. Shown are moment trajectories estimated via SSA samples ($N = 10$ and $N = 100$) and via TC and CC using the average of $N = 10$ SSA trajectories as reference input. The CME reference solution was obtained by numerical integration. Shaded regions are confidence bounds for $N = 10$ obtained with $\alpha = 0.95$. Middle and bottom. Errors for expectation and skewness. Results were obtained with initial condition $z(0) = 25$ molecules, which translates to $\mathbb{E}_Z(0) = 25$ and $\mathbb{V}_Z(0) = 0$ for the moment equations. Reaction rates were set to $[k_1, k_2] = [1.25\,\mathrm{s}^{-1}, 0.25\,\mathrm{s}^{-1}]$ and the system was simulated over a virtual time of $T_{\mathrm{end}} = 10\,\mathrm{s}$. For CC, we used the parameters $Q = \mathrm{diag}\{1000, 30\}$, $R = 1$, $S = 100$, and $\delta = 0.1\,\mathrm{s}$.

The variance and skewness were not constrained in this setting. The two first estimated moments, expectation and variance, are smooth and closely follow the reference CME solution. Moreover, the expectation value clearly lies within $\mathcal{X}_1(t)$ for both bounds. The skewness deviates from the reference course for the first 2.5 seconds but approaches the true solution quite well afterward. This becomes especially obvious when considering the squared quadratic error between the true CME solution and its approximations in the bottom sub-graph. Not only is the error of the controlled MoM solution orders of magnitude smaller than the error of the underlying 10-sample SSA approximation, but it is also comparable in magnitude with the SSA approximation for $N = 100$ samples while being significantly smoother at the same time.

System 2 Our second example represents a reversible dimerization reaction

$$\mathrm{A}_2 \xrightarrow{\mathrm{k}_1} 2\,\mathrm{A} \qquad 2\,\mathrm{A} \xrightarrow{\mathrm{k}_2} \mathrm{A}_2.$$

The configuration space of System 2 is limited by a conservation law, $2\mathrm{A}_2(t) + \mathrm{A}(t) = c_{\mathrm{A}} = 2\mathrm{A}_2(0) + \mathrm{A}(0)$, which can be used to reduce the system to species A only. The MoM ODEs are given by

$$\begin{aligned}\frac{\mathrm{d}\mathbb{E}_Z(t)}{\mathrm{d}t} &= k_1 c_{\mathrm{A}} + (2k_2 - k_1)\mathbb{E}_Z(t) - 2k_2\mathbb{E}_Z^2(t) - 2k_2\mathbb{V}_Z(t),\\ \frac{\mathrm{d}\mathbb{V}_Z(t)}{\mathrm{d}t} &= 2k_1 c_{\mathrm{A}} - (2k_1 + 4k_2)\mathbb{E}_Z(t) + 4k_2\mathbb{E}_Z^2(t)\\ &\quad - 8k_2\mathbb{E}_Z(t)\mathbb{V}_Z(t) + (8k_2 - 2k_1)\mathbb{V}_Z(t) - 4k_2\mathbb{S}_Z(t).\end{aligned}$$

For this simulation example, we chose an initial number of 40 molecules of species A and $c_{\mathrm{A}} = 50$, which directly implies the existence of 5 A_2 molecules for $t = 0$ and $Z(\Omega) = \{0, ..., 50\}$. A comparison of results for different approaches is presented in Fig. 3 analogously to Fig. 2. First and foremost, courses of the true CME moments are extremely well captured by both moment closure approaches, and the expectation value trajectory again lies well within $\mathcal{X}_1(t)$. The reason for the similarity of results of both closure methods is that the control input skewness estimated by the CC scheme is quite close to zero during the whole simulation. In contrast, the SSA-based approximation is as shaky as for System 1.

Conclusion

Applying the MoM to approximate the CME's moments is a challenging task due to the interdependence structure of the considered moments. In this work, we interpreted the system of MoM ODEs in the framework of control theory to present a non-heuristic closure scheme that provably stays close to the true solution in a probabilistic sense. More specifically, our CC scheme draws a few SSA samples from the true CME solution and uses the corresponding sample-based moment estimates as a physically meaningful reference trajectory. The closure moments are then used as the system's control input in order to follow the SSA trajectories while still obeying the MoM dynamics. Indeed, the CC leads to physically meaningful results that closely approximate the true moments even when TC fails to do so.

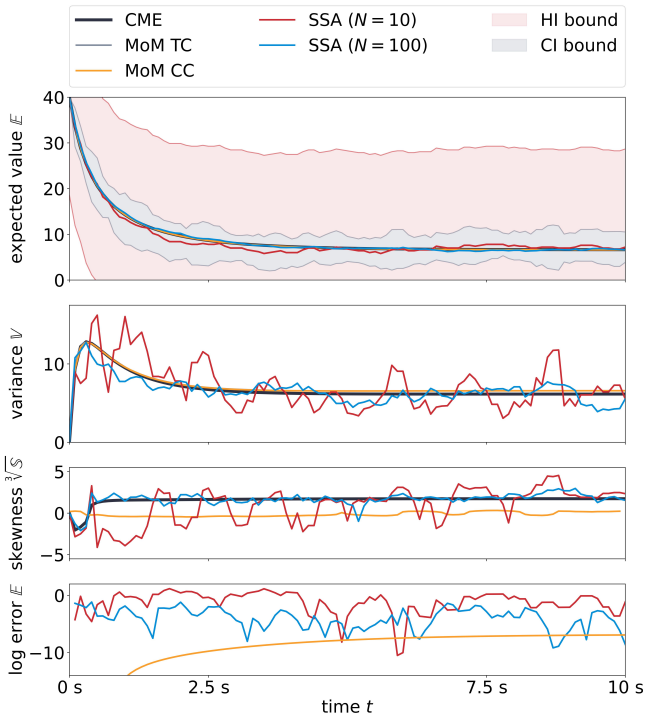

Fig. 3. Comparison of moment estimation approaches for System 2. Illustration of results similar to Fig. 2. Results were obtained with initial condition $z(0) = 40$ molecules and $c_{\mathrm{A}} = 50$ and reaction rates $[k_1, k_2] = [0.05\,\mathrm{s}^{-1}, 0.025\,\mathrm{s}^{-1}]$. Similar to Fig. 2, we used $Q = \mathrm{diag}\{1000, 30\}$, $R = 1$, $S = 100$, $\delta = 0.1\,\mathrm{s}$ for the CC.

Moreover, it outperforms purely sample-based approximation techniques as it smooths the sample averages according to the MoM dynamics. The MoM and hence basically all moment closure approaches require the considered moments of the CME solution to exist for the time interval of interest, which we have assumed (and numerically tested) in this work. However, it should be noted that this is a delicate issue since reaction systems with solutions whose moments are not defined can easily be constructed (44). These systems can show wildly stochastic behavior, and in particular, SSA-based moment estimates do not converge for increasing sample sizes. Moment-based methods terribly fail when applied to such systems without caution. Generally, this restriction has to be kept in mind when dealing with the MoM and related approaches. In this work, we have restricted ourselves to systems with one molecular species only, to keep the notation simple and to emphasize the main principle of the CC scheme itself rather than getting lost in complicated indexing. Our closure scheme can in principle also directly be applied to systems with more molecular species, and also considering more moments, as indicated with the more general dimensions at several places in the manuscript. However, a severe limitation is that the number of moments scales very badly with the number of molecular species in the system. For n_{c} chemical species, we have n_{c} expectation values, an $n_{\mathrm{c}} \times n_{\mathrm{c}}$ covariance matrix, n_{c}^3 skewness values, and the number grows likewise for higher moments. This is true for other moment approaches as well, but especially also affects the computation time in our case via repeated optimization. Finally, we note again that a particular strength of our method compared to other existing closure schemes is that we are able to provide probabilistic guarantees that the calculated moments stay in a neighborhood of the true moments and hence cannot arbitrarily diverge. Moreover, our scheme provides several tuning points to control this distance. First, the number of SSA samples correlates negatively with the size of the confidence intervals. Second, different weights can be given to different moments, and the smoothness of the input can be regulated via hyperparameters in the MPC objective function. Finally, computational effort can to some extent be regulated by appropriately choosing the horizon in the short-horizon MPC scheme. Overall, we believe that CC has great potential to be wildly applied to describe the moments of biochemical reaction systems.

Funding

Funded by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2075 – 390740016 and within grant AL 316/15-1 – 468094890. We acknowledge the support by the Stuttgart Center for Simulation Science (SimTech). V. Wagner and R. Strässer thank the Graduate Academy of the SC SimTech for its support.

Abbreviations

Control Closure (CC); Chebyshev-Inequality (CI); Central Limit Theorem (CLT); Chemical Master Equation (CME); Hoeffding's Inequality (HI); Method of Moments (MoM); Model Predictive Control (MPC); Ordinary Differential Equation (ODE); Stochastic Simulation Algorithm (SSA); Truncation Closure (TC)

Availability of data and materials

All code is available via https://fairdomhub.org/models/846.

Competing interests

The authors declare that they have no competing interests.

Authors' contributions

N.R. and F.A. supervised the project and acquired funding, V.W. and R.S. implemented the examples, all authors contributed to the theoretical results, R.S., V.W., and N.R. wrote the manuscript, all authors reviewed the manuscript.

Acknowledgments

This preprint was formatted using a LaTeX class by Ricardo Henriques that can be accessed here.

Appendix A: Method of Moment ODEs

In this article, we consider chemical systems consisting of one chemical species. We are interested in the expected value $\mathbb{E}_Z(t)$ and variance $\mathbb{V}_Z(t)$ of the time-dependent number of

system molecules $Z(t)$. To derive the MoM equations, we start with Proposition 2.5 in (3). For the special case of one chemical species, the equations read

$$\frac{\mathrm{d}\mathbb{E}_Z(t)}{\mathrm{d}t} = \sum_{j=1}^{n_r} \nu_j \mathbb{E}_Z(a_j(Z(t))), \tag{14a}$$

$$\frac{\mathrm{d}\mathbb{V}_Z(t)}{\mathrm{d}t} = \sum_{j=1}^{n_r} 2\nu_j \mathbb{E}_Z\left((Z-\mathbb{E}_Z(t))a_j(Z(t))\right) + \sum_{j=1}^{n_r} \nu_j^2 \mathbb{E}_Z(a_j(Z(t))). \tag{14b}$$

A crucial next step is a reasonable approximation of $\mathbb{E}_Z(a_j(Z(t)))$. To this end, we assume that the system's propensities are perfectly approximated by a second-order Taylor approximation. This assumption is fulfilled for all mass-action kinetics, as their propensities are polynomials of order two or less. More specifically, our assumptions imply

$$\begin{aligned} a_j(Z(t)) &= a_j(\mathbb{E}_Z(t)) \\ &+ \frac{\mathrm{d}a_j(\mathbb{E}_Z(t))}{\mathrm{d}Z(t)}(Z(t)-\mathbb{E}_Z(t)) \\ &+ \frac{\mathrm{d}a_j(\mathbb{E}_Z(t))}{\mathrm{d}Z^2(t)}\frac{(Z(t)-\mathbb{E}_Z(t))^2}{2}. \end{aligned}$$

Applying the expected value to both sides of the equation results in the desired characterization

$$\begin{aligned} \mathbb{E}_Z(a_j(Z(t))) &= a_j(\mathbb{E}_Z(t)) \\ &+ \frac{\mathrm{d}a_j(\mathbb{E}_Z(t))}{\mathrm{d}Z(t)}\mathbb{E}_Z(Z(t)-\mathbb{E}_Z(t)) \\ &+ \frac{\mathrm{d}a_j(\mathbb{E}_Z(t))}{\mathrm{d}Z^2(t)}\frac{\mathbb{E}_Z\left((Z(t)-\mathbb{E}_Z(t))^2\right)}{2} \\ &= a_j(\mathbb{E}_Z(t)) + \frac{\mathrm{d}a_j(\mathbb{E}_Z(t))}{\mathrm{d}Z^2(t)}\frac{\mathbb{V}_Z(t)}{2}. \end{aligned} \tag{15}$$

By inserting Eq. (15) into the original MoM ODEs Eq. (14), we obtain the final MoM equations

$$\begin{aligned} \frac{\mathrm{d}\mathbb{E}_Z(t)}{\mathrm{d}t} &= \sum_{j=1}^{n_r} \nu_j \left[a_j(\mathbb{E}_Z(t)) + \frac{\mathrm{d}^2 a_j(\mathbb{E}_Z(t))}{\mathrm{d}z^2}\frac{\mathbb{V}_Z(t)}{2}\right], \\ \frac{\mathrm{d}\mathbb{V}_Z(t)}{\mathrm{d}t} &= \sum_{j=1}^{n_r}\Bigg(\nu_j^2 a_j(\mathbb{E}_Z(t)) \\ &+ \mathbb{V}_Z(t)\left[2\nu_j\frac{\mathrm{d}a_j(\mathbb{E}_Z(t))}{\mathrm{d}z} + \frac{\nu_j^2}{2}\frac{\mathrm{d}^2 a_j(\mathbb{E}_Z(t))}{\mathrm{d}z^2}\right] \\ &+ \mathbb{S}_Z(t)\nu_j\frac{\mathrm{d}^2 a_j(\mathbb{E}_Z(t))}{\mathrm{d}z^2}\Bigg). \end{aligned}$$

Appendix B: Deriving a probabilistic upper bound for SSA-based moment estimation from the CLT

The theorem of Berry-Esseen (31, 32) is an improved version of the CLT. We use it to derive a probabilistic bound for the expectation estimation error of a sample average.

Theorem 2 (Berry-Esseen) For $a \in \mathbb{R}$ and N samples, consider the cumulative density F_N of the shifted and scaled sample average $\hat{\mathbb{E}}_Z^{(N)}$

$$F_N(a) := P\left(\frac{\hat{\mathbb{E}}_Z^{(N)} - \mathbb{E}_Z}{\sqrt{\frac{\mathbb{V}_Z}{N}}} \le a\right) = P\left(\hat{\mathbb{E}}_Z^{(N)} - \mathbb{E}_Z \le a\sqrt{\frac{\mathbb{V}_Z}{N}}\right). \tag{16}$$

Then, we conclude

$$|F_N(a) - \Phi(a)| \le \frac{c_{\mathrm{be}}\rho}{\sqrt{N\mathbb{V}_Z^3}}, \tag{17}$$

where Φ is the standard-normal cumulative distribution.
We derive the probabilistic bound by first reformulating the quantity of interest and using the definition of F_N. Subsequently, Theorem 2 is employed to achieve an upper bound. After this, terms are rearranged and we introduce the error function that is related to Φ by the equality

$$\Phi(a) = \frac{1}{2}\left(1 + \mathrm{erf}\left(\frac{a}{\sqrt{2}}\right)\right). \tag{18}$$

Altogether, the derivation reads as

$$\begin{aligned} &P\left(\left|\hat{\mathbb{E}}_Z^{(N)} - \mathbb{E}_Z\right| \ge a\right) \\ &= P\left(\hat{\mathbb{E}}_Z^{(N)} - \mathbb{E}_Z \ge a\right) + P\left(\hat{\mathbb{E}}_Z^{(N)} - \mathbb{E}_Z \le -a\right) \\ &= 1 - P\left(\hat{\mathbb{E}}_Z^{(N)} - \mathbb{E}_Z \le a\right) + P\left(\hat{\mathbb{E}}_Z^{(N)} - \mathbb{E}_Z \le -a\right) \\ &\overset{\text{Eq. (16)}}{=} 1 - F_N\left(a\sqrt{\frac{N}{\mathbb{V}_Z}}\right) + F_N\left(-a\sqrt{\frac{N}{\mathbb{V}_Z}}\right) \\ &\overset{\text{Eq. (17)}}{\le} 1 - \Phi\left(a\sqrt{\frac{N}{\mathbb{V}_Z}}\right) + \Phi\left(-a\sqrt{\frac{N}{\mathbb{V}_Z}}\right) + \frac{2c_{\mathrm{be}}\rho}{\sqrt{N\mathbb{V}_Z^3}} \\ &= 2\Phi\left(-a\sqrt{\frac{N}{\mathbb{V}_Z}}\right) + \frac{2c_{\mathrm{be}}\rho}{\sqrt{N\mathbb{V}_Z^3}} \\ &\overset{\text{Eq. (18)}}{=} 1 + \mathrm{erf}\left(-a\sqrt{\frac{N}{2\mathbb{V}_Z}}\right) + \frac{2c_{\mathrm{be}}\rho}{\sqrt{n\mathbb{V}_Z^3}} \\ &= 1 - \mathrm{erf}\left(\sqrt{\frac{Na^2}{2\mathbb{V}_Z}}\right) + \frac{2c_{\mathrm{be}}\rho}{\sqrt{N\mathbb{V}_Z^3}}. \end{aligned}$$

Appendix C - Theoretical properties

In this section, we first provide the theoretical background on MPC which is necessary to establish the theoretical guarantees of our CC scheme, and second, we establish the proof of Theorem 1. More precisely, we present theoretical guarantees for the class of polynomial systems, including all mass-action systems, leading to probability bounds for the error between true and estimated moment trajectories. In particular, we present results on closed-loop (practical) stability of the closed-loop Eq. (13) based on (28). To this end, we recall the notion of practical stability.

Lemma 3 ((45, Thm. 2.4),(27, Thm. 2.23)) Let V be a time-varying practical Lyapunov function with

$$\alpha_1(\|e(t)\|) \leq V(e(t),t) \leq \alpha_2(\|e(t)\|),$$
$$V(e(t+\delta),t+\delta) - V(e(t),t) \leq -\alpha_3(\|e(t)\|) + \theta$$

for all $V(e(t),t) \leq V_{\max}$, $V_{\max} > \eta := \alpha_2(\alpha_3^{-1}(\theta)) + \theta$ and $\alpha_1, \alpha_2, \alpha_3 \in \mathcal{K}, \theta > 0$. Then, for all initial conditions $V(e(0),0) \leq V_{max}$, the origin $e = 0$ is uniformly practically asymptotically stable, i.e., the system uniformly converges to the set $\{(e,t) | V(e,t) \leq \eta\}$.

In the following, we establish practical stability of system Eq. (4) and show that the MPC problem in Eq. (12) is recursively feasible. To this end, we need the following local incremental stabilizability assumption (28, 29).

Assumption 4: There exist a continuous control law $\kappa : \mathcal{X} \times \mathcal{X} \times \mathcal{U} \to \mathbb{R}^m$, a δ-Lyapunov function $V_\delta : \mathcal{X} \times \mathcal{X} \times \mathcal{U} \to \mathbb{R}_{\geq 0}$, that is continuous in the first argument and satisfies $V_\delta(\xi,\xi,v) = 0$ for all $(\xi,v) \in \mathcal{X} \times \mathcal{U}$, and positive parameters $c_{\delta,l}, c_{\delta,u}, \delta_{\text{loc}}, k_{\max}, \rho$, such that for all initial conditions $(x(0),\xi(0),v(0)) \in \mathcal{X} \times \mathcal{X} \times \mathcal{U}$ with $V_\delta(x(0),\xi(0),v(0)) \leq \delta_{\text{loc}}$ hold

$$c_{\delta,l}\|x-\xi\|^2 \leq V_\delta(x,\xi,v) \leq c_{\delta,u}\|x-\xi\|^2,$$
$$\|\kappa(x,\xi,v) - v\| \leq k_{\max}\|x-\xi\|,$$
$$V_\delta(x(t),\xi(t),v(t)) \leq e^{-\rho(t-s)} V_\delta(x(s),\xi(s),v(s)),$$
$$0 \leq s \leq t,$$
$$(x,\xi,v) \in \mathcal{X} \times \mathcal{X} \times \mathcal{U},$$

with

$$\dot{x} = f(x) + B\kappa(x,\xi,v), \qquad \dot{\xi} = f(\xi) + Bv.$$

This assumption basically requires that the controller $\kappa(x,\xi,v)$ leads to perfect tracking of (ξ,v) when x is sufficiently close to ξ. The constant $k_{\max}$ serves as a local Lipschitz bound on the nonlinear control law κ.

As our goal is to stay close to the reference trajectory $(\hat{x}(t),\hat{u}(t))$, the MPC scheme should track the *unknown* reference $(x^*(t),u^*(t))$. To this end, the subsequent analysis assumes compactness of $\mathcal{X} \times \mathcal{U}$, which is satisfied for bounded state space $Z(\Omega)$. Since Subsection 2.2 only discusses probabilistic bounds of the expectation, but no bounds for higher-order moments, this requirement is generally violated for unconstrained moments. However, some system classes allow also the derivation of probabilistic bounds for higher-order moments and, alternatively, it is always possible to define a large enough interval around the SSA-based moment estimate which then contains the true moments with high probability and leads to the compactness of $\mathcal{X} \times \mathcal{U}$. Hence, using additionally the boundedness of the SSA-based estimate $(\hat{x}(t),\hat{u}(t))$ for the considered systems, the difference between the reference and the optimal reachable trajectory is bounded as well. We make the following assumption to guarantee constraint satisfaction.

Assumption 5: The optimal reachable reference trajectory $(x^*(t),u^*(t))$ from Definition 1 is such that $V_\delta(x(t),x^*(t),u^*(t)) \leq \delta_{\text{r}}$ implies

$$x(t) \in \mathcal{X}, \qquad \kappa(x(t),x^*(t),u^*(t)) \in \mathcal{U}$$

with V_δ, κ from Assumption 4 and some $\delta_{\text{r}} > 0$. There exists a function $\alpha_V \in \mathcal{K}$ such that $V_\delta(x(t),x^*(t),u^*(t)) \leq \delta_{\text{r}}$ and $V_\delta(y,x^*(t),u^*(t)) \leq \delta_{\text{r}}$ imply

$$\|V_\delta(x(t),x^*(t),u^*(t)) - V_\delta(y,x^*(t),u^*(t))\|$$
$$\leq \alpha_V(\|x(t)-x^*(t)\| + \|y-x^*(t)\|).$$

We denote the error with respect to the optimal reachable reference by $e_{\text{r}}(t) = x(t) - x^*(t)$ and the stage cost of the optimal reachable reference trajectory by $\ell^*(t) = \ell(t,x^*(t),u^*(t))$. Further, we make the following dissipativity assumption (28, 29, 46).

Assumption 6: There exists a time-varying storage function $\lambda : \mathbb{R} \times \mathcal{X} \to \mathbb{R}$ such that for all $(x(t),u(t)) \in \mathcal{X} \times \mathcal{U}$ the rotated stage cost L satisfies

$$L(t,x(t),u(t)) := \ell(t,x(t),u(t)) - \ell^*(t) - \dot{\lambda}(t,x(t)) \geq \alpha_l(\|e_{\text{r}}(t)\|)$$

for all $t \geq 0$, $\delta > 0$, with $\alpha_l \in \mathcal{K}_\infty$. Furthermore, for all $x(t) \in \mathcal{X}$ the storage function is uniformly bounded by

$$|\lambda(t,x(t))| \leq \gamma_\lambda(\|e_{\text{r}}(t)\|), \qquad \gamma_\lambda \in \mathcal{K}.$$

Moreover, we define the rotated cost function

$$\tilde{J}_T(x(t),t,u(\cdot;t)) = \int_t^{t+T} L(\tau,x(\tau;t),u(\tau;t))\mathrm{d}\tau$$

and the corresponding rotated MPC problem

$$\begin{aligned}\tilde{J}_T(x(t),t) := \min_{\tilde{u}(\cdot;t)} \; & \tilde{J}_T(x(t),t,\tilde{u}(\cdot;t)) \\ \text{s.t. } & \dot{\tilde{x}}(\tau;t) = f(\tilde{x}(\tau;t)) + B\tilde{u}(\tau;t)), \\ & (\tilde{x}(\tau;t),\tilde{u}(\tau;t)) \in \mathcal{X} \times \mathcal{U}, \\ & \tilde{x}(t;t) = x(t), \\ & \tau \in [t,t+T].\end{aligned}$$

We make the following assumption to establish a lower bound on the optimal cost.

Assumption 7: For any given positive constant $\delta \leq T$, there exists a function $\alpha_\delta \in \mathcal{K}$ satisfying that

$$\int_t^{t+T} \alpha_l(\|x(\tau;t) - x^*(\tau)\|)\mathrm{d}\tau \geq \alpha_\delta(\|e_{\text{r}}(t)\|)$$

for all feasible $x(\tau;t) \in \mathcal{X}, \tau \in [t,t+T]$, where $\alpha_{\delta_1}(r) \geq \alpha_{\delta_2}(r)$ if $\delta_1 \geq \delta_2$.

For practical tracking, we introduce the set $\mathbb{S}^\delta_{c_p} := \{(x,t) | \tilde{J}_T(x,t) \leq \alpha_\delta(c_p)\}$, with α_δ according to Assumption 7 and c_p defined by

$$\delta_p := \min\{\delta_{\text{loc}}, \delta_{\text{r}}\}, \qquad c_p := \sqrt{\frac{\delta_p}{c_{\delta,u}}}.$$

Theorem 4 ((28, Thm. 3.1)) Let Assumptions 4, 5, 6, and 7 be satisfied. Then, there exists a $\tilde{T}_0$ and a function $\tilde{\theta} \in \mathcal{L}$, such that for all $T > \tilde{T}_0$ and all initial conditions satisfying $(x(0),0) \in \mathbb{S}^\delta_{c_p}$, the MPC problem Eq. (12) is recursively feasible and the closed-loop system Eq. (13) satisfies

$$\alpha_\ell(\|e_{\text{r}}(t)\|) \leq \tilde{J}_T(x(t),t) \leq \alpha_u(\|e_{\text{r}}(t)\|)$$

$$\tilde{J}_T(x(t+\delta),t+\delta)-\tilde{J}_T(x(t),t)\leq-\alpha_\ell(\|e_{\mathrm{r}}(t)\|)+\tilde{\theta}(T-\delta)$$

for $t\geq 0$. The set $\mathbb{S}^{\delta}_{c_p}$ is positively invariant and the origin $e_{\mathrm{r}}=0$ is uniformly practically asymptotically stable under the closed-loop dynamics Eq. (13).

Proof: This result is taken from (28) and we omit the proof for conciseness. ■

The resulting closed-loop system practically tracks the *unknown* optimal reachable reference $(x^*(t),u^*(t))$. We note that the assumption on the initial condition in Theorem 4 can be relaxed by requiring a longer time-horizon $T\geq\tilde{T}_1>\tilde{T}_0$ (28, Thm. 3.2).

Remark 8: Note that the results in (28, 29) are only valid for *periodic* references. Even though in our case the reference is *non*periodic, we can still use the presented results to ensure closed-loop stability over any finite interval $\mathcal{T}$ (29, Rem. 8). More precisely, any reference $(\hat{x}(t),\hat{u}(t))$ on the interval $\mathcal{T}$ can be extended to a periodic reference $(\tilde{x}(t),\tilde{u}(t))$, where $(\tilde{x}(t),\tilde{u}(t))=(\hat{x}(t),\hat{u}(t))$ for all $t\in\mathcal{T}$ and the period length $p\geq T_{\mathrm{end}}$. Then, clearly, both reference trajectories with the same initial condition cannot be distinguished on the finite interval $\mathcal{T}$.

Now we can prove Theorem 1 establishing theoretical guarantees for our proposed CC scheme estimating the moments of the CME. *Proof:* [Proof of Theorem 1] We rely on the result in Theorem 4. To this end, we first need to show that the considered setting satisfies all assumptions of Theorem 4. Note that the nonlinear dynamics in Eq. (4) are twice differentiable. Further, the MoM dynamics Eq. (4) are such that the linearized system $(A(\xi),B)$ with $A(\xi)=\frac{\partial f}{\partial x}\big|_{\xi}$ is locally controllable around any point $(\xi,v)\in\mathcal{X}\times\mathcal{U}$. This results from the particular structure of the dynamics, where the moments of order $(k+1)$ enter only linearly in the dynamics of the moments of order k, i.e., $A_{j,j+1}(\xi)=A_{j,j+1}\in\mathbb{R}\setminus\{0\}$ for all $j=1,...,n-1$ and $\tilde{B}\neq 0$. Thus, for $n_{\mathrm{c}}=1$ and $n_{\mathrm{m}}=2$ we have

$$S_{(A(\xi),B)}\begin{bmatrix}B & A(\xi)B\end{bmatrix}=\begin{bmatrix}0 & A_{1,2}\tilde{B}\\ \tilde{B} & A_{2,2}(\xi)\tilde{B}\end{bmatrix}$$

with $S_{(A(\xi),B)}$ having full rank for all ξ. Hence, Assumption 4 is satisfied for system Eq. (4) (cf. (28, Rem. 2.1), (29, Prop. 1)). As discussed in (38, 47), Assumption 5 requires local continuity of the δ-Lyapunov function V_δ around the optimal reachable reference trajectory $(x^*(t),u^*(t))$ which is also ensured by local controllability of the system. Computing suitable storage functions for Assumption 6 is typically challenging for general nonlinear systems with arbitrary reference trajectories $(\hat{x}(t),\hat{u}(t))$. However, the existence of such a storage function can be shown assuming local controllability (48, Def. 3.6.4) and that the system is uniformly suboptimally operated off the trajectory x^* according to (39, Def. 12) (compare (29, Lm. 4)). Thus, Assumption 6 is satisfied for system Eq. (4) with Assumption 3. Moreover, as noted in (28), a lower bound satisfying Assumption 7 can be constructed for polynomial systems and, thus, in particular also for system Eq. (4), using convex optimization as proposed in (49). Hence, all the assumptions of Theorem 4 are satisfied.

It remains to show that the initial condition satisfies $(x(0),0)\in\mathbb{S}^{\delta}_{c_p}$. Since we assume in our setting exact knowledge of the initial condition, i.e., $x(0)=\hat{x}(0)=x^*(0)$ and $e_{\mathrm{r}}(0)=0$, we obtain $\tilde{J}_T(x(0),0)\leq\alpha_u(\|e_{\mathrm{r}}(0)\|)=0$ due to (28, Prop. 3.1). Hence, the initial condition $(x(0),0)$ of our setting is always contained in the set $\mathbb{S}^{\delta}_{c_p}$, and, thus, we can directly apply Theorem 4 to show the first result of Theorem 1 for the constraint set $\mathcal{X}\times\mathcal{U}$ consisting of the SSA-based confidence bounds combined with physical properties of the considered moments.

The second claim results from the fact that the SSA-based moment estimation $(\hat{x}(t),\hat{u}(t))$ converges to the true CME moments for $N\to\infty$. Since the MoM is exact if we know the true closure moments and, in the limit, the considered reference $(\hat{x}(t),\hat{u}(t))$ matches with the true moments and is thereby compatible with the dynamics Eq. (4), the MPC optimization also follows the true CME moments for $N\to\infty$. ■

Bibliography

1. Daniel T Gillespie. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *Journal of Computational Physics*, 22(4):403–434, 1976.
2. Daniel T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *Journal of Computational Chemistry*, 81(25):2340–2361, 1977.
3. Stefan Engblom. Computing the moments of high dimensional solutions of the master equation. *Applied Mathematics and Computation*, 180(2):498–515, 2006.
4. Christian Kuehn. Moment closure—a brief review. *Control of self-organizing nonlinear systems*, pages 253–271, 2016.
5. P. Whittle. On the use of the normal approximation in the treatment of stochastic processes. *Journal of the Royal Statistical Society: Series B (Methodological)*, 19(2):268–281, 1957.
6. I. Krishnarajah, A. Cook, G. Marion, and G. Gibson. Novel moment closure approximations in stochastic epidemics. *Bulletin of Mathematical Biology*, 67(4):855–873, 2005.
7. Isthrinayagy Krishnarajah, Glenn Marion, and Gavin Gibson. Novel bivariate moment-closure approximations. *Mathematical Biosciences*, 208(2):621–643, 2007.
8. Tiago Ramalho, Marco Selig, Ulrich Gerland, and Torsten A. Ensslin. Simulation of stochastic network dynamics via entropic matching. *Physical Review E*, 87(2):022719, 2013.
9. Eszter Lakatos, Angelique Ale, Paul D. W. Kirk, and Michael P. H. Stumpf. Multivariate moment closure techniques for stochastic kinetic models. *Journal of Chemical Physics*, 143 (9):094107, 2015.
10. Leo Bronstein and Heinz Koeppl. A variational approach to moment-closure approximations for the kinetics of biomolecular reaction networks. *Journal of Chemical Physics*, 148(1): 014105, 2018.
11. Abhyudai Singh and Joao Pedro Hespanha. Lognormal moment closures for biochemical reactions. In *Proc. 45th IEEE Conference on Decision and Control*, 2006.
12. Abhyudai Singh and João Pedro Hespanha. A derivative matching approach to moment closure for the stochastic logistic model. *Bulletin of Mathematical Biology*, 69(6):1909–1925, 2007.
13. Abhyudai Singh and João P. Hespanha. Approximate moment dynamics for chemically reacting systems. *IEEE Transactions on Automatic Control*, 56(2):414–418, 2011.
14. Mohammad Soltani, Cesar Augusto Vargas-Garcia, and Abhyudai Singh. Conditional moment closure schemes for studying stochastic dynamics of genetic circuits. *IEEE Transactions on Biomedical Circuits and Systems*, 9(4):518–526, 2015.
15. Mohammad Naghnaeian and Domitilla Del Vecchio. Robust moment closure method for the chemical master equation. In *IEEE Conference on Control Technology and Applications (CCTA)*, 2017.
16. Alexander Andreychenko, Linar Mikeev, and Verena Wolf. Model reconstruction for moment-based stochastic chemical kinetics. *ACM Transactions on Modeling and Computer Simulation*, 25(2):1–19, 2015.
17. Patrick Smadbeck and Yiannis N. Kaznessis. A closure scheme for chemical master equations. *Proc. of the National Academy of Sciences of the USA*, 110(35):14261–14265, 2013.
18. Vincent Wagner, Benjamin Castellaz, Marco Oesting, and Nicole Radde. Quasi-entropy closure: a fast and reliable approach to close the moment equations of the chemical master equation. *Bioinformatics*, 38(18):4352–4359, 2022.
19. Ramon Grima. A study of the accuracy of moment-closure approximations for stochastic chemical kinetics. *Journal of Chemical Physics*, 136(15):154105, 2012.
20. David Schnoerr, Guido Sanguinetti, and Ramon Grima. Validity conditions for moment closure approximations in stochastic chemical kinetics. *Journal of Chemical Physics*, 141(8): 084103, 2014.
21. David Schnoerr, Guido Sanguinetti, and Ramon Grima. Comparison of different moment-closure approximations for stochastic chemical kinetics. *Journal of Chemical Physics*, 143 (18):185101, 2015.
22. J. Ruess, A. Milias-Argeitis, S. Summers, and J. Lygeros. Moment estimation for chemically reacting systems by extended Kalman filtering. *Journal of Chemical Physics*, 135(16): 165102, 2011.

23. Hassan K. Khalil. *Nonlinear systems*. Prentice-Hall, Upper Saddle River, NJ, third edition, 2002.
24. Pablo A. Parrilo. *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*. Ph.D. dissertation, California Institute of Technology, 2000.
25. Man-Duen Choi, Tsit Lam, and B. Reznick. Sums of squares of real polynomials. *Proceedings of Symposia in Pure Mathematics*, 58, 1994.
26. James B. Rawlings, David Q. Mayne, and Moritz M. Diehl. *Model Predictive Control: Theory, Design, and Computation*. Nob Hill Publishing, Madison, WI, 2nd, paperback edition, 2020.
27. Lars Grüne and Jürgen Pannek. *Nonlinear model predictive control*. Springer, 2017.
28. Yushen Long and Lihua Xie. Unconstrained tracking MPC for continuous-time nonlinear systems. *Automatica*, 129:109680, 2021.
29. Johannes Köhler, Matthias A. Müller, and Frank Allgöwer. Nonlinear reference tracking: An economic model predictive control perspective. *IEEE Transactions on Automatic Control*, 64(1):254–269, 2018.
30. I Richard Savage. Probability inequalities of the Tchebycheff type. *Journal of Research of the National Bureau of Standards-B. Mathematics and Mathematical Physics B*, 65(3): 211–222, 1961.
31. Andrew C Berry. The accuracy of the gaussian approximation to the sum of independent variates. *Transactions of the American Mathematical Society*, 49(1):122–136, 1941.
32. Carl-Gustav Esseen. Fourier analysis of distribution functions. A mathematical study of the Laplace-Gaussian law. 1945.
33. I. G. Shevtsova. On the absolute constants in the Berry-Esseen-type inequalities. In *Doklady Mathematics*, volume 89, pages 378–381. Springer, 2014.
34. Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.
35. Bernard Bercu, Bernard Delyon, and Emmanuel Rio. *Concentration inequalities for sums and martingales*. Springer, 2015.
36. Sergey V Nagaev. Large deviations of sums of independent random variables. *The Annals of Probability*, pages 745–789, 1979.
37. CNP Slagle. One hundred probability and statistics inequalities. *arXiv:2102.07234*, 2021.
38. Lars Grüne. Economic receding horizon control without terminal constraints. *Automatica*, 49(3):725–734, 2013.
39. Matthias A. Müller, Lars Grüne, and Frank Allgöwer. On the role of dissipativity in economic model predictive control. *IFAC-PapersOnLine*, 48(23):110–116, 2015.
40. Raffaele Soloperto, Johannes Köhler, and Frank Allgöwer. A nonlinear MPC scheme for output tracking without terminal ingredients. *IEEE Transactions on Automatic Control*, 68 (4):2368–2375, 2023.
41. Daniel Limon, Antonio Ferramosca, Ignacio Alvarado, and Teodoro Alamo. Nonlinear MPC for tracking piece-wise constant reference signals. *IEEE Transactions on Automatic Control*, 63(11):3735–3750, 2018.
42. Michael J. Risbeck and James B. Rawlings. MPCTools: Nonlinear model predictive control tools for CasADi (Octave inferface), 2016.
43. Joel A. E. Andersson, Joris Gillis, Greg Horn, James B. Rawlings, and Moritz Diehl. CasADi – A software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*, 11(1):1–36, 2019.
44. Vincent Wagner and Nicole Radde. The impossible challenge of estimating non-existent moments of the chemical master equation. *bioRxiv:2023.02.08.527667*, 2023.
45. Lars Grüne and Marleen Stieler. Asymptotic stability and transient optimality of economic mpc without terminal conditions. *Journal of Process Control*, 24(8):1187–1196, 2014.
46. Andrea Alessandretti, A Pedro Aguiar, and Colin N. Jones. On convergence and performance certification of a continuous-time economic model predictive control scheme with time-varying performance index. *Automatica*, 68:305–313, 2016.
47. Matthias A Müller and Lars Grüne. Economic model predictive control without terminal constraints for optimal periodic behavior. *Automatica*, 70:128–139, 2016.
48. Eduardo D. Sontag. *Mathematical Control Theory: Deterministic Finite Dimensional Systems*, volume 6. Springer Science & Business Media, 1998.
49. Tanagorn Jennawasin, Michihiro Kawanishi, and Tatsuo Narikiyo. Performance bounds for optimal control of polynomial systems: A convex optimization approach. *SICE Journal of Control, Measurement, and System Integration*, 4(6):423–429, 2011.