A Spectral Methods-Based Solution of the Chemical Master Equation for Gene Regulatory Networks

Michael Nip
Center for Control, Dynamical Systems, and Computation
University of California, Santa Barbara
California, USA 93106
mdnip@engineering.ucsb.edu

João P. Hespanha
Center for Control, Dynamical Systems, and Computation
University of California, Santa Barbara
California, USA 93106
hespanha@ece.ucsb.edu

Mustafa Khammash
Department of Biosystems Science and Engineering
ETH-Zürich
4058 Basel, Switzerland
mustafa.khammash@bsse.ethz.ch

***Abstract*— We present a new method to approximate the time evolution of the probability density function (PDF) for molecule counts in gene regulatory networks, modeled by the Chemical Master Equation (CME). A key feature of our method is that molecular states can be aggregated to reduce the computational burden without the need for assumptions like time-scales separation. We observe that the CME is amenable to the use of spectral methods adapted from partial differential equations and our method results from expanding the solution using carefully selected basis functions. The method is illustrated in the context of an example taken from the field of systems biology.**

I. Introduction

The evolution of a spatially homogeneous mixture of reacting particles in thermal equilibrium is often modeled using a stochastic formulation in which the reaction system is represented by a jump Markov process where each state represents the population count of each of the constituent species [1]. In this framework, the evolution of the probability density of the system's chemical populations is governed by the Chemical Master Equation (CME). In most cases the CME cannot be solved explicitly and various Monte Carlo simulation techniques have been used to find approximations of the probability densities by producing either detailed or approximate realizations of each process [1][2][3][4]. However, for many systems, biologically important events may occur rarely necessitating the generation of a prohibitively large set of realizations to obtain sufficiently precise statistics.

Others have attempted to solve the CME directly for the evolution of the probability densities. Many methods such as the finite state projection [5] and the sliding window abstraction [6] are based on truncating the state space to a finite subset containing the majority of the probability mass. These approaches have the advantage of providing explicit error bounds on their approximations of the densities. Unfortunately, to guarantee that such an approximation has a low error, it is often necessary to include a large number of states in the truncation, rendering many systems computationally intractable through such methods.

Various approximate methods have been developed that trade accurate density information for computational tractability, often replacing the discrete state-space description with a continuous one. These include Van Kampen's Linear Noise Approximation (LNA) [7], Moment Closure methods [8][9], and Chemical Langevin Equation (CLE) treatments [10][11]. These methods tend to give an accurate description of the dynamics when the population counts of all species remain large, but can perform poorly even when a single species exhibits low molecular counts. This is a significant limitation when one needs to model the (boolean) activation state of genes that necessarily have low molecular counts.

The classes of methods described in the two paragraphs above are in some sense complementary and there has been much recent work attempting to combine the best features of these, leading to the so called hybrid approaches. Many are based on exploiting a time-scale separation to partition the system into subsets of fast and slow reactions. These methods are then based on coupling an approximate method such as τ-leaping [12] or the Chemical Langevin Equation [13][14] for the fast species with an efficient variant of the Gillespie algorithm for the slow species to produce a new Monte Carlo algorithm. Other methods are based on partitioning the chemical species into a subset with large average molecule count and a subset with low molecule count and making an ODE approximation for the dynamics of the high count species [15][16]. These methods assume that the variance of each high count species is small compared to its mean; an assumption that, as we shall see, is not always valid.

Recently, spectral methods adapted from the numerical PDE literature have been employed to calculate approximate solutions of the CME. These expand the probability distribution as a linear combination of spectral basis functions and then use a Galerkin projection to map the dynamics onto a lower dimensional subspace spanned by a finite subset of the basis functions. In [17], the CME is expanded in terms of the discrete Charlier polynomials to obtain a global approximation of the solution. Discrete Chebychev polynomials have been used in [18] along with an appropriate scaling and translation to construct a global basis for the distribution from bases for disjoint rectangular patches. In regions of the state space where the fine details of the distribution are unnecessary, the basis functions can be sparse, but where

fine detail is needed, one should use a rich set of functions. Both in [17], and [18], the projection of the dynamics used was dictated by the choice of basis functions: each basis was orthogonal under a particular inner product and the authors chose to use the orthogonal projection consistent with that inner product.

In this paper we consider a class of models of gene regulatory networks that explicitly include the full binding dynamics of transcription factors. Here it is assumed that the binding dynamics occur on a time scale that is not fast enough to be averaged out by the rest of the system and are therefore important to the overall function of the network. We demonstrate that spectral methods are particularly well suited to solving CMEs of this type. In addition, we state the general mathematical framework applying spectral methods to CME models that is more general than Galerkin projection. Other examples of spectral methods that have been applied in the PDE literature are collocation, the method of moments, and least squares, all of which fit within our framework [19]. We also propose a novel heuristic for calculating stationary distributions based on this framework.

The remainder of this paper is organized as follows. The next section introduces models of gene regulatory networks that include transcription factor binding dynamics. Section 3 introduces the mathematical framework for general spectral methods for these models and introduces the method for approximating the stationary distribution. Section 4 applies these methods to a simple example system from systems biology.

II. Gene Regulatory Networks

We assume that the state of our system of chemical reactions is expressed by a pair (x, g) where g takes values in a finite set $\mathcal{Q} = \{1, \ldots, Q\}$ and x is a vector of integers that can potentially be unbounded. We shall refer to the g component of the state as the *low count species* and to x as the *high count species*. In a generic regulatory circuit, the subcomponent g could correspond to the configuration of occupied/vacant binding sites for a transcription factor, whereas x could be a vector with molecule counts of mRNA, protein, and/or transcription factors.

For a stochastic chemical reaction network with R reaction channels, for the kth reaction, let $\omega_k(x, g)$ denote the reaction propensity, and η_k^x and η_k^g the components of the stoichiometric vector affecting the high count and low count species, respectively, when channel k fires. The Chemical Master Equation describing the time evolution of the probability density function is given by:

$$\begin{aligned}\dot{p}(x,g;t) \;=\; & \\ & - \;\; p(x,g;t)\sum_{\mu=1}^{R} p(x,g;t)\omega_k(x,g) \\ & + \;\; \sum_{\mu=1}^{R} p(x-\eta_k^x, g-\eta_k^g;t)\omega_k(x-\eta_k^x, g-\eta_k^g)\end{aligned} \tag{1}$$

where each $p(x, g; t)$ is nonnegative, less than 1, and the sum of all $p(x, g; t)$ over all possible values of x and q is equal to 1. In general, this describes an infinite dimensional coupled linear system of differential equations since for each pair (x, g), equation (1) specifies one differential equation.

The state space on which the process evolves is $\mathbb{Z}_{\geq 0}^{n_x} \times \mathcal{Q}$ and, for each fixed time t, $p(x, g; t)$ is an ℓ^1 function over the index-set $\mathbb{Z}_{\geq 0}^{n} \times \mathcal{Q}$.

III. Spectral Projection

Let $\{\psi_\alpha(x, g) : \alpha \in A\}$ be a basis for $\ell^1(\mathbb{Z}_{\geq 0}^{n} \times \mathcal{Q})$ where A is some infinite (but countable) index set. Expanded in this basis, the chemical master equation reads:

$$\begin{aligned}&\sum_{\alpha\in A} \dot{c}_\alpha(t)\psi_\alpha(x,g) \\ &= -\sum_{\mu=1}^{R}\sum_{\alpha\in A} c_\alpha(t)\psi_\alpha(x,g)\omega_k(x,g) \\ &+ \sum_{\mu=1}^{R}\sum_{\alpha\in A} c_\alpha(t)\psi_\alpha(x-\eta_k^x, g-\eta_k^g)\omega_k(x-\eta_k^x, g-\eta_k^g)\end{aligned} \tag{2}$$

$\forall (x, g) \in \mathbb{Z}_{\geq 0}^{n} \times \mathcal{Q}$. Our goal is to develop an approximation to the equation (2) that accurately captures biologically meaningful quantities that can be computed from $p(x, g; t)$. Many such quantities can be obtained by computing functionals on the state space $\ell^1(\mathbb{Z}_{\geq 0}^{n} \times \mathcal{Q})$ that correspond to inner products. For example, the probability of the system occupying some subset S of the state space is a linear functional given by the inner product of the distribution with the indicator function $\mathbf{1}_S$ taking the value 1 only on S. The expected value of any function $f(x, g)$ also corresponds to the evaluation of a linear functional that can be expressed as the inner product of the distribution with a vector whose entries are the values of $f(x, g)$ at the different points in the state-space $\mathbb{Z}_{\geq 0}^{n} \times \mathcal{Q}$. In view of this, we take as given a collection $\{\phi_\beta : \beta \in B\}$ of linear functionals on ℓ^1 parameterized by an index set B, with the understanding that, while we may be willing to accept errors in the probability density $p(x, g; t)$, we want our approximation to the CME to accurately capture the evolution of the values of each functional ϕ_β along solutions to the CME, which is given by the following set of equations:

$$\begin{aligned}&\sum_{\alpha} \dot{c}_\alpha(t)\phi_\beta(\psi_\alpha(x,g)) \\ =\;& \sum_{\alpha}(-\sum_{\mu=1}^{R} c_\alpha(t)\phi_\beta(\psi_\alpha(x,g)\omega_k(x,g)) \\ +\;& \sum_{\mu=1}^{R} c_\alpha(t)\phi_\beta(\psi_\alpha(x-\eta_k^x, g-\eta_k^g)\omega_k(x-\eta_k^x, g-\eta_k^g)))\end{aligned} \tag{3}$$

with one differential equation for each ϕ_β. In order for the set of coefficients $\{c_\alpha(t)\}$ to solve equation (2) they must

necessarily solve the equation (3) for each functional ϕ_β, $\beta \in B$.

Our construction of the lower dimensional approximation to the CME relies on the premise that the distribution $p(x, g; t)$ lies close to a subspace of $\ell^1(Z^n_{\geq 0} \times \mathcal{Q})$ generated by a finite subset $\{\phi_\alpha : \alpha \in \hat{A}\}$, $\hat{A} \subset A$ of the original basis for $\ell^1(Z^n_{\geq 0} \times \mathcal{Q})$. To determine appropriate equations for the evolution of the coefficients $\{c_\alpha : \alpha \in \hat{A}\}$, we require that the equations (3) that express the evolution of the functionals $\{\phi_\beta : \beta \in B\}$ holds for every $\beta \in B$.

A. Projection onto the Reduced Basis

Choosing the truncated basis functions $\{\psi_\alpha : \alpha \in \hat{A}\}$ to form an orthogonal basis for an Hilbert space H and selecting the functionals $\{\phi_\beta : \beta \in B\}$ to be of the form $\phi_\beta(z) =< \psi_\beta, z >, \forall z \in H, \forall \beta \in A = B$, each equation (3) would directly give us the evolution of one of the coefficients. However, this is undesirable for two reasons: (1) it excludes the possibility to work with truncated basis that, while not orthogonal, have a better chance to approximate well $p(x, g; t)$ and (2) it typically prevents the use of functionals ϕ_β that provide biologically useful information. The price we pay, is that the system of equations (3) does not directly lead to one equation for each coefficient c_α. However, by stacking the coefficient $\{c_\alpha : \alpha \in A\}$ into a column vector $\mathbf{c}$, we can still write (3) as the following system of linear differential equations

$$\mathbf{Q}\dot{\mathbf{c}} = \tilde{\mathbf{A}}\mathbf{c} \tag{4}$$

where $\mathbf{Q}$ and $\tilde{\mathbf{A}}$ denote semi-infinite matrices with one row for each element of the finite set B and one column for each element of the infinite set A. Similarly, by stacking all the probabilities $p(x, g; t)$ in a infinite vector $\mathbf{p}(t)$ we can write

$$\mathbf{p} = \mathbf{B}\mathbf{c} \tag{5}$$

where $\mathbf{B}$ is the infinite matrix with one column for each element of the set $\mathbb{Z}^n_{\geq 0} \times \mathcal{Q}$ and one row for each element of A. The CME can be expressed as

$$\dot{\mathbf{p}} = \mathbf{A}\mathbf{p} \tag{6}$$

for some infinite matrix with one row/column for each element of $Z^n_{\geq 0} \times Q$. Finally, we can also stack the values of all functionals $\{\phi_\beta \in B\}$ into a (finite) vector $\mathbf{d}$ with one entry for each element of B and write

$$\mathbf{d} = \mathbf{D}\mathbf{p}$$

where the semi-infinite matrix $\mathbf{D}$ has one row for each entry of the finite set B and one column for each element of the infinite set $\mathbb{Z}^n_{\geq 0} \times \mathcal{Q}$. Given a truncation of the corresponding basis elements, $\mathbf{B}$ can be partitioned into submatrices:

$$\mathbf{B} = \begin{bmatrix} \mathbf{B}_r & \mathbf{B}_\infty \end{bmatrix} \tag{7}$$

and $\mathbf{c}$ can be partitioned into subvectors:

$$\mathbf{c} = \begin{bmatrix} \mathbf{c}_r \\ \mathbf{c}_\infty \end{bmatrix} \tag{8}$$

Equation (4) can then be rewritten as

$$\mathbf{D}\mathbf{B}_r\dot{\mathbf{c}}_r + \mathbf{D}\mathbf{B}_\infty\dot{\mathbf{c}}_\infty = \mathbf{D}\mathbf{A}\mathbf{B}_r\mathbf{c}_r + \mathbf{D}\mathbf{A}\mathbf{B}_\infty\mathbf{c}_\infty$$

Truncation of the basis is equivalent to imposing that $\mathbf{c}_\infty = 0$ for all time resulting in the reduced system:

$$\mathbf{D}\mathbf{B}_r\dot{\mathbf{c}}_r = \mathbf{D}\mathbf{A}\mathbf{B}_r\mathbf{c}_r \tag{9}$$

We refer to equation (9) as the reduced dynamics. These equations can be viewed as a Method of Lines approximation of the dynamics. If $\mathbf{D}\mathbf{B}_r$ is invertible, then equation (9) can be multiplied on both sides by the inverse leading to a set of coupled ordinary differential equations which can be solved using numerical ODE integrators.

B. Approximating Stationary Distributions

For many applications it is desirable to obtain an estimate of the stationary distribution. For example, many hybrid algorithms assume that a subset of the dynamics is sufficiently fast and reaches stationarity to obtain a reduced model by averaging out the fast dynamics. Such methods require an estimate of the stationary distribution of the fast subset. However, Monte Carlos methods like SSA and the Finite State Projection are not particularly adequate for determining stationary distributions. SSA requires the generation of a large number of long sample paths, with the required time span of each simulation difficult to estimate *a priori.* In the Finite State Projection, the stationary distribution of the truncated process typically corresponds to all probability assigned to an artificially created absorbing state. In this section we propose a heuristic for finding an approximation of the stationary distribution using the spectral projection framework previously established.

Assuming that the system has a stationary distribution $\hat{\mathbf{p}}$ it must satisfy the equation:

$$0 = \mathbf{A}\hat{\mathbf{p}} \tag{10}$$

A candidate for the closest approximation of the stationary distribution in the reduced subspace is

$$\arg \inf_{\substack{||x||_1 = 1, \\ x \in R(\mathbf{B}_r)}} ||\mathbf{A}\mathbf{x}|| \tag{11}$$

which identifies the distribution in the reduced subspace that exhibits the slowest change. Since the subspace is assumed finite dimensional it is closed and the infimum is achieved. The infinite vector $\mathbf{A}\mathbf{x}$ is typically impossible to calculate, but it is possible to calculate the truncated vector $\mathbf{D}\mathbf{A}\mathbf{B}_r\mathbf{c}$ so we instead solve the related minimization problem:

$$\begin{aligned} &\textit{minimize} && ||\mathbf{D}\mathbf{A}\mathbf{B}_r\mathbf{c}_r|| \\ &\textit{subject to} && ||\mathbf{B}_r\mathbf{c}_r||_1 = 1 \end{aligned} \tag{12}$$

where the optimization variables are the coefficients $\mathbf{c}$ of the expansion. For any choice of norm in the objective function, the optimization is convex. When it is chosen to be the 2-norm, this is equivalent to finding the eigenvector

corresponding to the smallest eigenvalue of the positive semi-definite matrix

$$(\mathbf{DAB}_r)^*(\mathbf{DAB}_r) \tag{13}$$

and then normalizing with respect to the 1-norm.

C. Choice of Basis Functions

While many choices of bases in (2) are admissible, some will capture the dynamics better when truncating to a finite number of elements. [17] proposes the use of discrete Charlier polynomials, which are mutually orthogonal with respect to the Poisson distributions. [18] suggests the use of discrete Chebychev polynomials with a suitable scaling and translation are used to construct a tensor basis for the space, the elements of which are mutually orthogonal with respect to the standard inner product.

In this work we are especially interested in systems with state spaces of the form $\mathbb{Z}^n_{\geq 0} \times Q$, where the component of the state in the finite set Q typically represents the dynamics of transcription factor binding sites. A natural choice of basis for such system is obtained by taking the tensor product between vectors in a basis of $\ell^1(\mathbb{Z}^n_{\geq 0})$ with indicator functions on the set $\mathcal{Q}$. Specifically, we work with basis vectors for $\ell^1(\mathbb{Z}^n_{\geq 0} \times \mathcal{Q})$ of the form

$$\{\psi_j \otimes e_k\}_{j,k} \tag{14}$$

where ψ_j, $j \in \mathbb{Z}^n_{\geq 0}$ is an element of the basis of $l_1(\mathbb{Z}^n_{\geq 0})$ and e_k, $k \in \mathcal{Q}$, is the indicator function $e_k(g) = 1$ if $g = k$ and $e_k(g) = 0$ if $g \neq k$.

This choice of basis is motivated by the observation that by expressing the probability distribution $p(x, g; t) = P(X(t) = x, G(t) = g)$ as a linear combination of the vectors (14), we are implicitly expanding the distribution of $X(t)$ conditioned to $G(t)$ as a linear combination of the vectors ψ_j since:

$$\begin{aligned} & P(X(t) = x, G(t) = g) \\ = \;& \alpha(t) P(X(t) = x | G(t) = g) \\ = \;& \alpha(t) \sum_{j,k} c_{j,k} (\psi_j \otimes e_k)(x, g) \\ = \;& \alpha(t) \sum_{j} c_{j,g} \psi_\beta(x) \end{aligned} \tag{15}$$

where $\alpha(t) = P(G(t) = g)$. In view of this, the choice of the basis functions ψ_j should be dictated by the conditional distribution of $X(t)$ given $G(t)$.

In many systems, the binding configuration determines the qualitative dynamics of the mRNA, proteins, etc. For instance, the binding of a transcription factor to a promoter can drastically increase the transcription rate of the associated gene. In the next section, we describe a simple system where this is precisely the case and including the binding dynamics is crucial to obtaining an accurate model of the system.

IV. Case Study: Negative Feedback Circuit Exhibiting Bimodality

Consider the negative-feedback gene regulatory circuit modeled by

$$\begin{aligned} G_{on} &\xrightarrow{\frac{k_1 X}{TV}} G_{off} \\ G_{off} &\xrightarrow{\frac{k_2}{T}} G_{on} \\ G_{on} &\xrightarrow{k_3 V} G_{on} + X \\ \emptyset &\xrightarrow{k_4 V} X \\ X &\xrightarrow{k_5} \emptyset \end{aligned}$$

G acts as a promoter that randomly switches between the "on" and "off" states and the species X is produced at a high rate when G is "on" and at a lower basal rate when G is "off". The likelihood of G switching from "on" to "off" is influenced by the concentration of X: the higher the concentration, the more likely G will transition to "off". V is a rate parameter that reflects the volume in which the reactions take place, while T is a time-scale parameter that scales the toggle switch rate between the two states of the gene. Here, the rate of production of the species X scales linearly with the volume while the rate of degradation remains constant. This results in a steady state mean of X increasing linearly with increasing volume. Decreasing T decreases the mean waiting time between switching from "on" to "off" and vice versa.

This system can be made to exhibit a bimodal stationary distribution for species X, which is difficult to capture by methods that rely on a continuous approximation of the dynamics. In the regime of slowly switching gene, the peaks of the distributions of X conditioned on G will be offset from one another as the system will slowly switch between two subsystems that are effectively isolated from each other. The distribution of X for the isolated subsystems is Poisson with parameter $\frac{k_4 V}{k_5}$ for the subsystem conditioned to "off" and with parameter $\frac{(k_3+k_4)V}{k_5}$ for the subsystem conditioned to "on". For k_3 large enough compared to k_4, the marginal distribution of X will then have two distinct peaks corresponding to the peak of each conditional. Note that in this parameter regime, it is essential to include the binding/unbinding dynamics of the promoter in the Markov model. Assuming an equilibrium model for the state of the binding site through an averaging argument destroys the bimodality of the full stationary distribution.

We encoded the state space of the process as the ordered pair (x, g) with x counting the number of molecules of X present and $g \in \{0, 1\}$ with $g = 0$ corresponding to the "off" state of the gene and $g = 1$ corresponding to the "on" state. For this example, we used the tensor basis:

$$\{e^{-\frac{(x-x_j)^2}{2\sigma^2}} \otimes e_k\}_{j,k} \tag{16}$$

where the first term of the product are elements of a basis for $\ell^1(\mathbb{Z}_{\geq 0})$ and e_k, $k \in \{0, 1\}$ is the indicator function of the event $g = k$. The different basis vectors for $\ell^1(\mathbb{Z}_{\geq 0})$ are

obtained using Gaussian distributions with means equal to $5, 10, 15, \ldots$ all with the same standard deviation equal to 15. The use of Gaussian distributions for the basis functions is motivated by the observation that the state distribution typically resemble Gaussians and therefore one hopes that a low-order representation can be achieved with such basis functions. For the linear functionals we used the evaluation mappings:

$$\phi_{(x,g)}(f) = f(x,g) \tag{17}$$

for any $f \in \ell^1(\mathbb{Z}^n_{geq0} \times \mathcal{Q})$ and with $g \in \{0,1\}$ and $x \in \{0,5,10,\ldots\}$ corresponding to the peak of each Gaussian. That is, for each pair $(x,g) \in B$, $B = \{(x,g) : x \in \{0,5,10,\ldots\}, g \in \{0,1\}\}$, $\phi_{(x,g)}$ is the linear functional that evaluates the distribution at (x,g). This corresponds to being able to identify whether or not the binding site is occupied (since we included both e_0 and e_1) and being able to count the number of transcription factors but with limited resolution.

Equations (16) and (17) can be translated into the language of semi-infinite matrices in the following way. Assuming the lexicographic ordering of the state space and using the corresponding matrix $\mathbf{A}$, the matrices encoding the choice of basis functions and linear functionals have the block diagonal form:

$$\mathbf{B}_r = \begin{bmatrix} \mathbf{B}_r^0 & \mathbf{0} \\ \mathbf{0} & \mathbf{B}_r^1 \end{bmatrix}, \quad \mathbf{D} = \begin{bmatrix} \mathbf{D}^0 & \mathbf{0} \\ \mathbf{0} & \mathbf{D}^1 \end{bmatrix} \tag{18}$$

with the columns of $\mathbf{B}_r^0$ and $\mathbf{B}_r^1$ representing the basis for the distribution conditioned on G_{off} and G_{on}, respectively, and with $\mathbf{D}^0$ and $\mathbf{D}^1$ representing measurements of the amount of X present in each of the gene states:

$$\mathbf{B}_r^0 = \mathbf{B}_r^1 = \begin{bmatrix} | & | & \\ e^{-\frac{(x-0)^2}{2\sigma^2}} & e^{-\frac{(x-5)^2}{2\sigma^2}} & \ldots \\ | & | & \end{bmatrix} \tag{19}$$

$$\mathbf{D}^0 = \mathbf{D}^1 = \begin{bmatrix} e_0 \\ e_5 \\ \vdots \end{bmatrix} \tag{20}$$

The reduced model is given by equation (III-A). $\mathbf{DB}_r$ in this case happens to be invertible so we can directly apply conventional numerical ODE solvers.

Figure 1 compares the time evolution of the marginal distribution for X using the spectral approximation method proposed here with the one obtained from Monte Carlo simulations using 10^4 SSA sample paths. Both distributions assumed that the system is initialized with $(x,g) = (0,0)$ The spectral approximation produces qualitatively similar dynamics to the Monte Carlo with good agreement of the centering of the peaks of each mode of the distribution. Note, however, that the values of the peaks for the approximation are smaller than those produced by Monte Carlo and that the approximation distribution decays much more slowly. This lack of fast decay is a consequence of using Gaussians in the choice of basis since the decay rate of the approximation is bound below by the decay rate of any one of the Gaussians used.

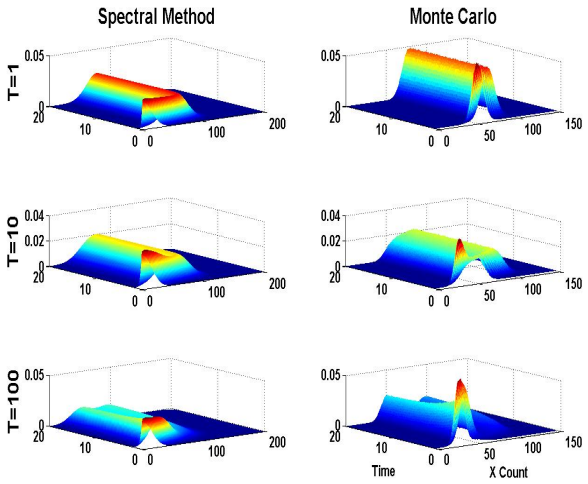

Fig. 1. Comparison of the time evolutions of the marginal PDF of species X generated by the spectral approximation (left column) with 10^5 Monte Carlo simulations (right column). In the case of the slow binding/unbinding dynamics ($T = 100$), the spectral method successfully captures the bimodality of the solution.

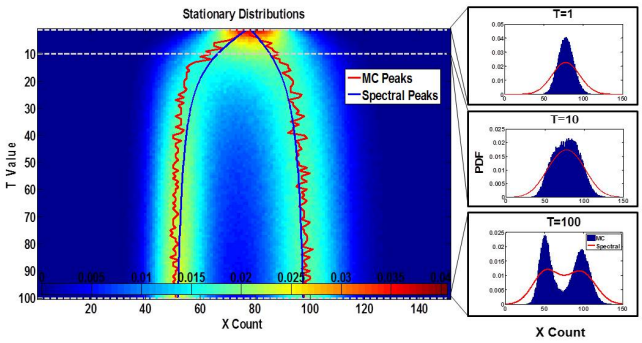

Fig. 2. Comparison of estimates of the stationary marginal PDF of species X generated by the spectral approximation with 10^5 Monte Carlo simulations. The pseudo-color plot shows PDFs estimated by Monte Carlo over a range of binding/unbinding speeds. The red curves show the peaks of the distribution estimated by Monte Carlo, while the blue curves show the peaks estimated by the spectral method. At right, comparisons of the stationary distribution predicted by the spectral method with Monte Carlo estimates. The spectral method successfully captures the bimodality of the solution in the slow switching regime.

Figure 2 compares the approximate stationary distribution generated by the heuristic with SSA sample paths of simulated time length of 100. The approximate solution was found by selecting the 2-norm for the objective function and performing the corresponding eigenvalue/eigenvector calculation. The spectral approximation is qualitatively similar to the Monte Carlo estimate. Note here that this solution suffers from the same inaccuracies as the transient solution: peak values are smaller, and the distribution does not decay as rapidly.

V. Conclusions/Future Research

In this paper we considered a class of models of gene regulatory networks that explicitly include the full binding dynamics of transcription factors. Here we assumed that the binding dynamics occur on a time scale that is not fast enough to be averaged out by the rest of the system and are therefore important to the overall function of the network. We demonstrated that general spectral methods have characteristics well suited to solving CMEs of this type assuming a good choice of basis functions can be found. We also introduced a novel method for approximating the stationary distribution based on the spectral framework.

Establishing error bound for the approximations proposed here is the topic of ongoing research.

VI. Acknowledgements

This material is based upon work supported by the Institute for Collaborative Biotechnologies through grant W911NF-09-D-0001 from the U.S. Army Research Office and by the National Science Foundation under Grants No. ECCS-0725485 and ECCS-0835847.

References

[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] Michael A. Gibson and Jehoshua Bruck. Efficient exact stochastic simulation of chemical systems with many species and many channels. *The Journal of Physical Chemistry A*, 104(9):1876–1889, 2000.

[3] Daniel T. Gillespie. Approximate accelerated stochastic simulation of chemically reacting systems. *The Journal of Chemical Physics*, 115(4):1716–1733, 2001.

[4] Yang Cao, Hong Li, and Linda Petzold. Efficient formulation of the stochastic simulation algorithm for chemically reacting systems. *The Journal of Chemical Physics*, 121(9):4059–4067, 2004.

[5] Brian Munsky and Mustafa Khammash. The finite state projection algorithm for the solution of the chemical master equation. *The Journal of Chemical Physics*, 124(4):044104, 2006.

[6] Thomas Henzinger, Maria Mateescu, and Verena Wolf. Sliding window abstraction for infinite markov chains. In Ahmed Bouajjani and Oded Maler, editors, *Computer Aided Verification*, volume 5643 of *Lecture Notes in Computer Science*, pages 337–352. Springer Berlin / Heidelberg, 2009. 10.1007/978-3-642-02658-4-27.

[7] NG Van Kampen. *Stochastic Processes in Physics and Chemistry*. North-Holland, Amsterdam and New York, 1992.

[8] João Pedro Hespanha and Abhyudai Singh. Stochastic models for chemically reacting systems using polynomial stochastic hybrid systems. *Int. J. on Robust Control,* Special Issue on Control at Small Scales: Issue 1, 15:669–689, Sep. 2005.

[9] Carlos A. Gomez-Uribe and George C. Verghese. Mass fluctuation kinetics: Capturing stochastic effects in systems of chemical reactions through coupled mean-variance computations. *The Journal of Chemical Physics*, 126(2):024109, 2007.

[10] Daniel T. Gillespie. The chemical langevin equation. *The Journal of Chemical Physics*, 113(1):297–306, 2000.

[11] Stewart N. Ethier and Thomas G. Kurtz. *Markov Processes: Characterization and Convergence*. Wiley-Interscience, New York, 2005.

[12] Jacek Puchalka and Andrzej M. Kierzek. Bridging the gap between stochastic and deterministic regimes in the kinetic simulations of the biochemical reaction networks. *Biophysical Journal*, 86(3):1357 – 1372, 2004.

[13] Eric L. Haseltine and James B. Rawlings. Approximate simulation of coupled fast and slow reactions for stochastic chemical kinetics. *The Journal of Chemical Physics*, 117(15):6959–6969, 2002.

[14] Howard Salis and Yiannis Kaznessis. Accurate hybrid stochastic simulation of a system of coupled chemical or biochemical reactions. *The Journal of Chemical Physics*, 122(5):054103, 2005.

[15] Andreas Hellander and Per Ltstedt. Hybrid method for the chemical master equation. *Journal of Computational Physics*, 227(1):100 – 122, 2007.

[16] Thomas A. Henzinger, Linar Mikeev, Maria Mateescu, and Verena Wolf. Hybrid numerical solution of the chemical master equation. In *Proceedings of the 8th International Conference on Computational Methods in Systems Biology*, CMSB '10, pages 55–65, New York, NY, USA, 2010. ACM.

[17] Stefan and Engblom. Spectral approximation of solutions to the chemical master equation. *Journal of Computational and Applied Mathematics*, 229(1):208 – 221, 2009.

[18] T. Jahnke M. Wulkow P. Deuflhard, W. Huisinga. Adaptive discrete gelerkin methods applied to the chemical master equation. *Sci. Comput.*, 30(6):2990–3011, 2008.

[19] John P. Boyd. *Chebyshev and Fourier Spectral Methods*. Dover Publications, Inc, 2001.