

A stochastic framework for the design of transient and steady state behavior of biochemical reaction networks

Ania A. Baetica¹, Ye Yuan², Jorge Gonçalves^{2,3}, and Richard M. Murray¹

Abstract—Stochasticity plays an essential role in biochemical systems. Stochastic behaviors of bimodality, excitability, and fluctuations have been observed in biochemical reaction networks at low molecular numbers. Stochastic dynamics can be captured by modeling the system using a forward Kolmogorov equation known in the biochemical literature as the chemical master equation. The chemical master equation describes the time evolution of the probability distributions of the molecule species. We develop a stochastic framework for the design of these time evolving probability distributions that includes specifying their uni-/multi-modality, their first moments, and their rate of convergence to the stationary distribution. By solving the corresponding optimizations programs, we determine the reaction rates of the biochemical systems that satisfy our design specifications. We then apply the design framework to examples of biochemical reaction networks to illustrate its strengths and limitations.

I. INTRODUCTION AND MOTIVATION

Biological behavior is commonly described using deterministic, nonlinear, continuous-time models [1]. For these models, multiple frameworks have been proposed in [2], [3] for the design of chemical reaction network behaviors. However, the deterministic description of chemical reaction network kinetics is not appropriate if the chemical species are at low molecular numbers or if stochastic fluctuations are important in the time evolution of the system [4]. As such, chemical reaction network kinetics inside living cells are better captured by discrete stochastic models since reactant molecules are often at low copy numbers and subject to random motion [5]. Experimental evidence in [4], [6] highlights stochastic effects in living cells by showing copy-number fluctuations in genetically identical cells and distinct cell fate decisions in populations of clonal cells.

In order to capture the observed discrete stochastic behavior, the chemical reactions in the network can be modeled as a Markov jump process [5]. Every state of this process is a vector of the concentration of species in the reaction network at a fixed time. The state vector evolves in time with dynamics given by a forward Kolmogorov equation, known in the biochemical literature as the chemical master equation (CME). The distribution of states evolves in time according to an infinite-dimensional ODE specified by the CME. The coefficients in the ODE are determined by rate constants and by the stoichiometry and propensity functions of the chemical reaction network. Analytical solutions to the CME are only available for specific examples of chemical reaction networks (e.g. monomolecular reaction networks [7]). Most commonly, no analytical solutions are known and Monte Carlo-based techniques are used to approximate the solutions [8]. One possible method is to truncate the

infinite-dimensional ODE by a finite state projection (FSP) and to obtain a finite-dimensional ODE approximation with bounded error [9].

We propose a stochastic framework for the design of the time evolving distributions of states, irrespective of knowledge of an analytical solution to the CME. We are able to capture design features of the chemical species' distributions such as their uni-/multi-modality, their first moment and shape, and their rate of convergence to a stationary distribution. These design features could not be captured in a deterministic framework; even the first moment of the distributions might be altered by stochastic effects [10].

The design features we chose were inspired by unanswered questions in the design of genetic regulatory circuits. Our insight comes from the problem of designing a simple genetic toggle switch [11]. The toggle switch has both unimodal and bimodal transients, as well as a wide range of gene expression levels in the cell population. The phenotypic heterogeneity of the cell population is poorly understood and not typically designed for. It would help control this heterogeneity to specify the modality of the transient distributions: uni-/multi-modal, the genes' expression levels, and the switching time. We formulate these design specifications mathematically using [12] as a guideline and we discuss how they result in remarkably different behaviors in the cell population in Section II C.

Even after selecting design features that are relevant to the design of biochemical reaction networks, the stochastic design problem is challenging to formulate mathematically. Our main challenges are that the exponential operator in the solution to the truncated CME has a dearth of exploitable mathematical properties [13] and a prohibitive computational cost. The exponential is not separable, which prevents us from leveraging a problem formulation in terms of relative entropy optimization as in [14]. We also considered its tensor projection as in [15], [16], [17], but the orthogonal bases that we projected on were depleted of biological meaning; it was unclear how to combine orthogonal basis polynomials in the space of projection such that design features of uni-/multi-modality of distributions were expressed. Such a formulation would create overly elaborate problems that lose track of biological implementation. To avoid these issues, we simply consider the Taylor approximation to the exponential operator and compute bounds on the error of this approximation in Section II D.

In Section III, we implement the design problem formulation for two examples of biochemical reaction networks: protein production-degradation and the Schlögl model [12].

When we use a first order Taylor approximation of the exponential operator, the design problems reduce to solving a linear program and a semi-definite program for the pre-specified convergence rate to the stationary distribution [18]. There exist very efficient, scalable computational tools, such as CVX, to solve these problems [19], [20]. However, the error of the approximation may be large and we suggest using polynomial optimization methods as an alternative. Our ability to obtain a solution also depends on the number of design features we specify and on the number of molecule counts for each species. We ultimately believe that we can find solutions for biochemical reaction networks with several species.

Our paper is organized as follows: In Section II, we set up the design problem and evaluate the error in the approximation of the exponential operator. In Section III, we implement and solve design problems for two classic examples of biochemical reaction networks. Section IV contains discussion of the applicability and limitations of our stochastic design framework, as well as an outline for future work.

II. DESIGN PROBLEM SETUP

A. Notation

Let $n \geq 1$, n integer. Let $P \in [0, 1]^n$ be the n -dimensional probability vector set. For $p = (p_1, \dots, p_n) \in P$, it must be that $p_i \geq 0$ and $\sum_{i=1}^n p_i = 1$.

B. Background on stochastic chemical kinetics

We start by describing a chemically reacting network that contains N distinct species $\{S_1, \dots, S_N\}$. The dynamical state of the system at time $t \geq 0$ is described by the state vector $x(t) = (x_1(t), \dots, x_N(t))$, where $x_i(t)$ is the integer population of species S_i at time t for all $1 \leq i \leq N$. There are M distinct monomolecular or bimolecular reactions $\{R_1, \dots, R_M\}$ that can change the system's state, according to the propensity function associated with each chemical reaction.

The CME describes how stochastically reacting chemical species behave in a well-stirred solution at thermal equilibrium in a fixed, finite volume [8]. The chemical kinetics of the N reacting molecular species are modeled as a discrete-state, continuous-time Markov process on the distribution state vector $p(x, t)$, which denotes the probability that the system will be in state x at time t . The CME gives the time evolution law for $p(x, t)$ as

$$\frac{\partial p}{\partial t}(x, t) = \sum_{j=1}^M (a_j(x - \xi_j)p(x - \xi_j, t) - a_j(x)p(x, t)) \quad (1),$$

where ξ_j is the j^{th} column of the stoichiometry matrix and a_j is the j^{th} propensity function associated with the chemical reaction network.

This equation is also referred to as the forward Kolmogorov equation for a jump Markov process. More compactly, the CME is a linear, infinite-dimensional ODE

$$\frac{dp}{dt}(x, t) = H(c)p(x, t) \quad (2),$$

where $c = (c_1, \dots, c_M)$ are the rate reaction parameters of the M chemical reactions.

Using the standard truncation given by the finite state projection algorithm in [9], we consider only a finite number S of states in each species in the chemical reaction network. Then $H(c)$ is finite-dimensional and we represent it affinely as

$$H(c) = \sum_{j=1}^M c_j H_j \quad (3).$$

Hence, equation (2) is equivalent to

$$\frac{dp}{dt}(x, t) = \sum_{j=1}^M c_j H_j p(x, t) \quad (4).$$

The matrices H_j are sparse, S -dimensional, and correspond uniquely to reaction R_j .

The solution to equation (4) is given by

$$p(x, t) = e^{\sum_{j=1}^M c_j H_j t} p(x, 0) \quad (5).$$

C. Problem formulation

Our formulation of a stochastic design framework for biochemical reaction networks is a two-part contribution: (1) we analytically describe the desired transient and stationary behavior using our design features and (2) we find a solution for the design problem under these constraints.

The design features we chose as constraints for the probability distribution vector are:

- (i) uni-/multi-modality
- (ii) fixed first moment
- (iii) rate of convergence to the stationary distribution

Our inclusion of design feature (i): the uni-/multi-modality of distributions is motivated by experimental evidence showing the presence of multi-modal (bimodal) transients in genetic switching in the λ phage, the lactose operon, and in cellular signal transduction pathways in mammalian cells [21]. The Gardner et al. [11] toggle switch is probably the first synthetic gene regulatory circuit to display multi-modality of the transient probability distributions. An illustration of the genetic toggle switch behavior is presented in Fig. 1.

Multi-modality is a purely stochastic behavior that cannot be reproduced or accounted for by deterministic modeling. Gardner et al. themselves give an incomplete explanation on why it appears in the genetic toggle switch: "the stochastic nature of gene expression causes variability in the location of the switching threshold and thus blurs the [deterministic] bifurcation point" [11]. Currently, multi-modality in gene regulatory circuits is not well understood and there are no analytical tools to control it based on the CME. We hope that our mathematical formulation of this design feature will shed light onto how to design for uni-/multi-modal distributions.

The design problem's mathematical formulation is to find reaction rate vector $c = (c_1, \dots, c_M)$ such that the probability distribution vector $p(x, t)$ is constrained according to our choice of design features for time point values $t \in T = \{t_1, \dots, t_k\}$, where $k \geq 1$ is the number of time points.




Fig. 1. The transient distributions for the toggle switch are bimodal, while the initial and stationary distributions are unimodal. The bimodal transient can be visualized as the cell population phenotype with roughly equal proportions of red and green fluorescent protein expression. The unimodal initial distribution is pictured on the left and the unimodal stationary distribution on the right. Figure partially reproduced from Portle et al. [22]

Find $c = (c_1, \dots, c_M)$ such that:

$$f_0 p_0 \leq \mu_0, f p^* \leq \mu_f \quad (6)$$

$$f_i e^{\sum_{j=1}^M c_j H_j t_i} p_0 \leq \mu_i \quad (7)$$

$$(H(c) - p^* 1_M)^T (H(c) - p^* 1_M) \leq \mu^2 \mathbb{I}_{M \times M} \quad (8)$$

$$H(c) p^* = p^* \quad (9)$$

$$H(c) = \sum_{j=1}^M c_j H_j \quad (10)$$

$$p_0 \in X_0, p^* \in X_f, p_{t_i} \in X_i \quad (11)$$

$$X_0, X_i, X_f \subseteq P, \forall 1 \leq i \leq k \quad (12)$$

where p_0 and p^* are the initial and stationary distributions, respectively; f_0, f_i, f are pre-selected projection operators that induce uni- or multi-modality of distributions; X_0, X_i, X_f are pre-selected subsets of P ; μ, μ_i, μ_f are the tightness of the bounds, for all $1 \leq i \leq k$.

The inequalities in equations (6) and (7) impose design features (i) and (ii) at time points $\{t_1, \dots, t_k\}$ under appropriate choices of operators. For example, an operator that imposes unimodality and first moment equal to value m can be chosen to be the function $g : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}, g(x) = (x - m)^2$ [12]. In Section III of our paper, we give more examples of projection operator choices.

Remark 1: The inequality in equation (8) reduces to a semi-definite program (SDP) by using the Schur complement formulation. According to [18], the bound μ can be used to tune the largest singular value norm of matrix $H(c)$. Thus, μ controls the rate of convergence to the stationary distribution through the solution of the SDP.

Remark 2: We clarify that design features (i) and (ii) apply to the marginal probability distributions of biochemical reactants in networks with more than just one species, $N > 1$. In order to marginalize the probability distributions, we

multiply the operators f, f_0 and $f_i, 1 \leq i \leq k$, by the appropriate marginalization matrices of size $M \times M^{N-1}$.

Finding a solution to the design problem is equivalent to checking the feasibility of a corresponding reachability problem. We let Y_0, Y_i , and Y_f be the subsets of P where inequalities in equations (6) and (7) hold respectively, for $1 \leq i \leq k$. Our problem is to find a reaction vector c such that there is a feasible probability distribution trajectory from set $X_0 \cap Y_0$ to set $X_f \cap Y_f$ that passes through the sets $X_i \cap Y_i$ for $1 \leq i \leq k$ and approaches $X_f \cap Y_f$ at the pre-determined rate μ . Finding a feasible solution to this reachability problem is equivalent to solving the design problem set up in equations (6)-(12). See a graphical representation in Fig. 2.




Fig. 2. Another way of thinking about our formulation is in the form of a reachability problem: $p(t)$ is a solution if it goes through sets $X_0 \cap Y_0, X_1 \cap Y_1, \dots, X_k \cap Y_k$ and approaches $X_f \cap Y_f$ at a pre-determined rate μ . The sets are drawn in rectangular shapes for illustrative purposes.

Our main challenge in finding a solution to the design problem is the exponential operator present in equation (7). The transition rate matrices $H_j, 1 \leq j \leq M$ do not typically commute, unless the associated chemical reaction is monomolecular, so the matrix exponential of the sum cannot be separated into a product of exponentials. A possible approach we considered was to use the tensor projection of the exponential operator as in [15], [16], [17], but it was overly elaborate to express the design features (i)-(iii) in terms of the orthogonal bases that we project onto. Our best approach has been to consider the Taylor approximation to the exponential operator and calculate the error of this approximation in Section II D. Following a Taylor approximation of order $l \geq 1$ of the exponential operator, the inequality in equation (7) is replaced by

$$f_i \sum_{v=0}^l \frac{1}{v!} \left(\sum_{j=1}^M c_j H_j t_i \right)^v p_0 \leq \mu_i, \forall 1 \leq i \leq k \quad (13).$$

Subsequently, the design problem formulation has linear constraints in equations (6) and (10), a semi-definite constraint in equation (8), and polynomial constraints in equations (9) and (13). The problem is polynomial of degree $l+1$ in variables c, p_0 , and p^* . In our implementation in Section III, we find it useful to assume knowledge of p_0 and p^* , acquired either

through experimental data or computer simulations. This reduces the degree of the polynomial problem to l , eliminates the inequality in equation (6), and makes the equality in equation (9) linear.

D. Error bound for the approximation of the exponential operator

Theorem 1: Let $A \in \mathbb{R}^{M \times M}$ be a transition rate matrix. Then the error bound for the approximation of the exponential operator $ge^{At}p_0$ by the truncated Taylor series $g \sum_{v=0}^l \frac{1}{v!} A^v t^v p_0$ of degree $l \in \mathbb{Z}_{\geq 1}$ is given by

$$gp_0 T_l(t) + \mathcal{O}(gp_0 T_l(\lambda_2 t)) \quad (14),$$

where $1 > \lambda_2 > \dots > \lambda_m$ are the eigenvalues of A without counting multiplicity. Here, T_l is the l^{th} degree Taylor polynomial, $T_l(t) = \sum_{i=l+1}^{\infty} \frac{1}{i!} t^i$, for all $t \geq 0$.

Proof: Let

$$\epsilon(t) = \sum_{i=l+1}^{\infty} \frac{1}{i!} g A^i p_0 t^i \quad (15)$$

be the residue following the truncation of the Taylor series.

We write transition matrix A in its Jordan form. Let $U \in \mathbb{R}^{M \times M}$ be an invertible matrix such that $A = UJU^{-1}$. Let the Jordan blocks be J_1, J_2, \dots, J_m , $m \geq 1$. The blocks correspond to eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_m$.

We separate each Jordan block $J_j = \lambda_j \mathbb{I}_j + N_j$, where \mathbb{I}_j is the identity matrix of size equal to that of block J_j and N_j is the corresponding nilpotent matrix. Then for each $j \geq 2$ and $i \geq l+1$,

$$U J_j^i U^{-1} = U (\lambda_j \mathbb{I}_j + N_j)^i U^{-1} \quad (16).$$

Since $|\lambda_j| < 1$ for any $j \geq 2$, then

$$J_j^i = \mathcal{O}(\lambda_j^i \mathbb{I}_j) \quad (17).$$

Given that $\lambda_2 > \dots > \lambda_m$, we obtain the final result

$$\epsilon(t) = gp_0 T_l(t) + \mathcal{O}(gp_0 T_l(\lambda_2 t)) \quad (18).$$

By applying Theorem 1 to the design problem, we calculate an error of

$$f_i p_0 T_l(t_i) + \mathcal{O}(f_i p_0 T_l(\lambda_2 t_i)) \quad (19)$$

for the approximation in equation (7) at each time step t_i , $1 \leq i \leq k$, where λ_2 is the second largest eigenvalue of matrix $H(c)$. This informs us to choose low norm reaction vector c , and to normalize f_i and p_0 in the implementation.

Remark 3: There is a clear trade-off between choosing a larger truncation order l with the effect of decreasing the approximation error and keeping the degree of the polynomial inequalities in the design problem low.

III. IMPLEMENTATION OF THE STOCHASTIC DESIGN FRAMEWORK

A. Protein production-degradation reaction network

We implement our design problem formulation on the gene regulatory network of protein production-degradation [1]. Here, protein production-degradation is modeled stochastically as a birth-death Markov process. The chemical reaction network has only two reactions



that represent the production and degradation of protein species A . The rates of the two reactions are c_1 and c_2 . The birth occurs according to a Poisson process with probability c_1 per unit time and the death occurs with probability per unit time proportional to $c_2 A(t)$.

We constrain the transient distribution to be unimodal and of mean 100 using operator $f(x) = (x - 100)^2$ and we assume that the stationary distribution is pre-determined by a Gaussian distribution with the same mean. The initial probability distribution is a Dirac delta function of height 1. Our simulation results give reaction rates $c_1 = 3.9894$ and $c_2 = 0.0397$. H_1 and H_2 are the same as in [12]. The number of states in the FSP truncation is $S = 201$ and the convergence rate to the stationary distribution is $\mu = 0.1$.

The results can be seen in Fig. 3 and Fig. 4. The approximation error is $\mathcal{O}(10^{-9})$.




Fig. 3. Time-evolution of the unimodal transient distributions.





Fig. 4. We plot the pre-specified stationary distribution and the operator $f(x) = (x - 100)^2$ that imposes the uni-modality of the transients.

Remark 4: We want to clarify that the solution to the optimization problem is not unique. The reaction rates c_1 and

c_2 can take other values and they can certainly be adjusted by tuning the bounds $\mu_0, \mu_i, \mu_f, \forall 1 \leq i \leq k$.

B. Schlögl chemical reaction network

The Schlögl chemical reaction network [23] exhibits bistability in the deterministic model and bimodality in the CME-based model. The set of reactions is as follows:



Here, concentrations of A and B are kept constant (buffered) and

$$a_1(X) = k_1 A \frac{1}{2} X(X-1) \quad (23)$$

$$a_2(X) = k_2 \frac{1}{6} X(X-1)(X-2) \quad (24)$$

$$a_3(X) = k_3 B \quad (25)$$

$$a_4(X) = k_4 X \quad (26)$$

are the propensity functions. We return to our previous notation by setting $c_1 = k_1 A$, $c_2 = k_2$, $c_3 = k_3 B$, and $c_4 = k_4$. See Gunawan et al. [24] for an in depth discussion of the chemical reaction network and [12] for the CME expression. The analysis of the deterministic model of the reaction network informs us that there is a bifurcation into two possible steady states with values $s_1 = 84.79$ and $s_2 = 569.9$. We construct our operators centered around these values.

Using operator $f_{unimodal}(x) = (x - s_1)^2$, we are able to impose an unimodal constraint on the transient distributions for rate reaction values $c_1 = 1.0710 \times 10^{-5}$, $c_2 = 21.9939 \times 10^{-15}$, $c_3 = 0.3668$, $c_4 = 0.0049$. We expect the coefficients to span many orders of magnitude [12]. We chose small rate reaction values in order to prevent the exponential operator from blowing up. The convergence rate is $\mu = 0.001$.

See our results in Fig. 5.




Fig. 5. We plot the time evolution of the uni-modal transients and compare it to the stationary distribution. Not all transients are displayed.

Then, we impose a bimodal transient constraint as in [12] using operator

$$f_{bimodal}(x) = \begin{cases} \min((x - s_1)^2, 14920) & \text{if } x \geq 328 \\ \min((x - s_2)^2, 14920) & \text{otherwise.} \end{cases}$$

and, simultaneously, a unimodal stationary constraint $f^*(x) = (x - s_1)^2$.

Our results are presented in Fig. 6.




Fig. 6. We plot the time evolution of the distributions. In part a, the initial distributions is pictured. We move through the transients in parts b-e. Part f has the stationary distribution. Not all transients are displayed.

We start from an initial distribution p_0 consisting of two Dirac delta functions with different weights and we move through a bimodal transient towards the unimodal steady state distribution p^* . It is possible to find a solution to the problem irrespective of the placement and the heights of the Dirac delta functions. We show this in Fig. 7 with a second unimodal stationary distribution choice $f^*(x) = (x - s_2)^2$. It is also possible to define an initial distribution p_0 with Gaussian distributions replacing of Dirac delta functions and also to replace the piece-wise function with a sum of Gaussian distributions centered at s_1 and s_2 . In all these cases, we are able to obtain solutions to the design problem.




Fig. 7. We plot the time evolution of the distributions. In part a, the initial distributions is pictured. We move through the transients in parts b-e. Part f has the stationary distribution. Not all transients are displayed.

However, when we impose a bimodal steady state distribution constraint, we are unable to find a satisfactory solution. This occurs because we implement equation (9) as the relaxation $\|H(c)p^* - p^*\| \leq \gamma$ for small γ . Hence, p^* is not forced to be an eigenvector of the transition rate matrix $H(c)$ and this does not ensure that there are no other eigenvectors corresponding to eigenvalues closer to 0. In our experience, the transient approaches p^* , but it ultimately decays to a stationary distribution corresponding to the eigenvector with the smallest eigenvalue. We choose not to implement equation (9) without the relaxation because the problem may be infeasible.

C. Reducing the error bound

If the approximation error bound is deemed too large, we can use a larger order approximation of the exponential operator to adjust it. In this case, the design problem becomes a polynomial optimization problem of order equal to that of the new Taylor approximation. Polynomial optimization problems (POPs) are computationally NP-hard [25]; but, in practice, solutions can usually be found for problems of small to moderate size [26], [27]. Using our formulation, we expect the polynomial optimization problems to be solvable for biochemical reaction networks with several species. Our ability to obtain a solution to the POP will also depend on the number of design features we specify and on the number of molecule counts allowed for each species.

IV. CONCLUSIONS AND FUTURE WORK

In this paper, we have developed and implemented a CME-based stochastic framework for the design of biochemical reaction networks. Our formulation of the stochastic design problem uses biologically meaningful design features for the setup of optimization problems. Their solutions are the rate reactions of the biochemical reaction networks. Our stochastic design framework might offer insight into what is even biologically possible to build; for example, we might want to know if it is possible to build a genetic switch with a uniform distribution transient. In particular, when designing gene regulatory circuits, it is challenging to predict what transient behavior might arise, how long the transient would last for, or if the stationary behavior will even follow our specifications. Using the design feature language we have developed, we can test for these questions. Future work will include applying our stochastic design framework to the class of genetic switches and testing out what is possible to build. When combined with forward simulation techniques, theoretical design work can be done by iterating between two, similarly to the design process followed in engineering problems.

The main limitation of our stochastic framework lies in the size of the problems we can solve accurately. A better approximation to the exponential operator might avoid the "curse of dimensionality", but none that we considered were viable. Hence, the polynomial optimization portion of the design problem formulation can only be solved for small to at most medium-sized problems. However, this might be sufficient to offer insight into the behavior of larger gene regulatory circuits, when combined with results in reducing multiscale stochastic models [28] or when using quasi-steady-state and quasi-equilibrium approximations [29]. In particular, we hope to use our framework to design multiscale genetic circuits with partial knowledge of rate reaction values.

ACKNOWLEDGMENTS

We would like to thank Venkat Chandrasekaran for discussion of optimization topics and for helping us formulate our problem. We would like to thank Thomas A. Catanach for the review of the manuscript. Research is supported

in part by the Air Force Office of Scientific Research, grant FA9550-14-1-0060. Ye Yuan and Jorge Gonçalves are supported by EPSRC.

REFERENCES

- [1] D. Del Vecchio, R. M. Murray, Biomolecular Feedback Systems (Book style), ch. 3-4, unpublished.
- [2] P. Gifani, Y. Yuan and J. Gonçalves, Reversed design approach. 2015.
- [3] I. Otero-Muras, J. R. Banga, Multicriteria global optimization for biocircuit design, BMC Systems Biology 2014; 8:113.
- [4] M. B. Elowitz, A. J. Levine, E. D. Siggia, P. S. Swain, Stochastic gene expression in a single cell, Science, 16 August 2002; 297(5584), pp. 1183-1186.
- [5] D. T. Gillespie, The Chemical Langevin Equation, Journal of Chemical Physics, 1 July 2000; 113(1), pp. 297-306.
- [6] H. H. McAdams, A. Arkin, Stochastic mechanisms in gene expression. Proc Natl Acad Sci U S A, 4 February 1997 Feb 4; 94(3), pp. 814-9.
- [7] T. Jahnke, W. Huisinga, Solving the chemical master equation for monomolecular reaction systems analytically, J Math Biol., January 2007; 54(1), pp. 1-26.
- [8] D. T. Gillespie, Stochastic simulation of chemical kinetics, Annu Rev Phys Chem, 2007; 58, pp. 35-55.
- [9] B. Munsky, M. Khammash, The finite state projection algorithm for the solution of the chemical master equation, Journal of Chemical Physics, 25 January 2006; 124, 044104.
- [10] J. Paulsson, O. G. Berg, M. Ehrenberg, Stochastic focusing: Fluctuation-enhanced sensitivity of intracellular regulation, Proc Natl Acad Sci U S A, 13 June 2000; 97, pp. 7148-7153.
- [11] T. S. Gardner, C. R. Cantor, J. J. Collins, Construction of a genetic toggle switch in *Escherichia coli*, Nature, 20 January 2000; 403, pp. 339-342.
- [12] N. C. Martins, J. M. Gonçalves, A linear programming approach to parameter fitting for the Master Equation, IEEE Transactions On Automatic Control, October 2009; 54(10), pp. 2451-2455.
- [13] C. Moler, C. V. Loan, Nineteen Dubious Ways to Compute the Exponential of a Matrix, Twenty-Five Years Later*, Siam Review, 3 February 2003; Vol. 45, No.1, pp. 349.
- [14] V. Chandrasekaran, P. Shah, Relative Entropy Relaxations for Signomial Optimization, arXiv:1409.7640 [math.OC], 26 September 2014.
- [15] M. Hegland, J. Garcke, On the numerical solution of the Chemical Master Equation with sums of rank one tensors, ANZIAM J., 4 August 2011; 52, pp. 628-643.
- [16] T. Jahnke, W. Huisinga, A dynamical low-rank approach to the Chemical Master Equation, Bulletin of Mathematical Biology, 21 August 2008; 70(8), pp. 2283-2302.
- [17] M. Nip, J. P. Hespanha, M. Khammash, A Spectral Methods-Based Solution of the Chemical Master Equation for Gene Regulatory Networks, 51st IEEE Conference on Decision and Control, December 2012; pp. 5354-5360.
- [18] L. Xiao, S. Boyd, Fast Linear Iterations for Distributed Averaging, System and Control Letters, 2004; 53, pp. 65-78.
- [19] CVX Research, Inc. CVX: Matlab software for disciplined convex programming, version 2.0. <http://cvxr.com/cvx>, April 2011.
- [20] M. Grant and S. Boyd. Graph implementations for nonsmooth convex programs, Recent Advances in Learning and Control (a tribute to M. Vidyasagar), V. Blondel, S. Boyd, and H. Kimura, editors, pages 95-110, Lecture Notes in Control and Information Sciences, Springer, 2008.
- [21] T. Tian, K. Burrage, Stochastic models for regulatory networks of the genetic toggle switch, PNAS., 21 March 2006; 103(22), pp. 83728377.
- [22] S. Portle, T. B. Causey, K. Wolf, G. N. Bennett, K. Sanb, N. Mantzaris, Cell population heterogeneity in expression of a gene-switching network with fluorescent markers of different half-lives, Journal of Biotechnology, Science, October 2006; 128(2007), pp. 362375.
- [23] D. T. Gillespie, Markov Processes: An Introduction for Physical Scientists, Academic Press, 1992, San Diego, CA.
- [24] R. Gunawan, Y. Cao, L. Petzold, F. J. Doyle III, Sensitivity Analysis of Discrete Stochastic Systems, Biophysical Journal, April 2005; 88(4), pp. 25302540.
- [25] P. A. Parrilo, Semidefinite programming relaxations for semialgebraic problems, Mathematical Programming, Springer-Verlag, 10 April 2003; 96(2), pp. 293-320.

- [26] H. Waki, S. Kim, M. Kojima, M. Muramatsu, Sums of squares and semidefinite programming relaxation for polynomial optimization problems with structured sparsity, *SIAM J. Optim.*, 2006; 17(1), pp. 218-242.
- [27] H. Waki, S. Kim, M. Kojima, M. Muramatsu, H. Sugimoto, Sparse-POP : a Sparse Semidefinite Programming Relaxation of Polynomial Optimization Problems, *ACM Transactions on Mathematical Software*, 2008; 35(2) 15.
- [28] K. Ball, T. G. Kurtz, L. Popovic, G. Rempala, Asymptotic Analysis of Multiscale Approximations to Reaction Networks, *The Annals of Applied Probability*, November 2006; 16(4), pp. 1925-1961.
- [29] B. Mélykúti, J. P. Hespanha, M. Khammash, Equilibrium Distributions of Simple Biochemical Reaction Systems for Time-Scale Separation in Stochastic Reaction Networks, *J. R. Soc. Interface*, 11 June 2014; 11:20140054.