

System Theory Applications in Biology: From Stochastic Chemical Kinetics to Deterministic Model Invalidation

Mustafa Khammash and Brian Munsky

Abstract—Ideas from System Theory lie behind many of the new powerful methods being developed in the burgeoning field of Systems Biology. In this paper, we show two examples of this: one in the area of stochastic chemical kinetics, and the other in biological model invalidation. Stochastic chemical kinetics has gained a lot of attention in the last few years. In order to capture certain important dynamics in the subcellular environment, it is necessary to model molecular interactions at the gene level as discrete stochastic events. The dynamics of such processes is typically described by probability distributions, which evolve according to the set of linear ordinary differential equations known as the chemical master equation (CME). Until recently, it has been believed that the CME could not be solved analytically except in the most trivial of problems, and the CME has been analyzed almost exclusively with Monte Carlo (MC) algorithms. However, concepts from linear systems theory have enabled the Finite State Projection (FSP) approach and have significantly enhanced our ability to solve the CME without resorting to MC simulations. In this paper we review the FSP approach as well as a variety of systems theory based modifications to the FSP algorithm that dramatically improve the computational efficiency of the algorithm and expand the class of solvable problems. Notions such as observability, controllability and minimal realizations enable large reductions in the order of models and increase efficiency with little to no loss in accuracy. Model reduction techniques based upon linear perturbation theory allow for the systematic projection of multiple time scale dynamics onto a slowly varying manifold. Our second example shows the application of systems ideas in the area of biological model invalidation. As a specific case study, we use a dynamic model of the bacterial heat-shock response to demonstrate the approach. Using recent sum-of-squares techniques we show that the heat-shock model, when stripped from a certain protein-protein interaction that implements a certain feedback loop, cannot account for the input-output data regardless of the parameter choice for the model. In essence, such a deficient model is invalidated. Such conclusions are essential for pointing out the likelihood of missing components or interactions, thereby guiding new biological experiments.

I. INTRODUCTION

With new experimental tools, contemporary molecular biologists are discovering complex gene regulatory networks, which control the expression of diverse biological traits. As more data becomes available, these networks become so complicated that it is becoming increasingly difficult to

Mustafa Khammash is with the Department of Mechanical Engineering, University of California, Santa Barbara, CA 93106-5070 khammash@engr.ucsb.edu

Brian Munsky is with the Department of Mechanical Engineering, University of California, Santa Barbara, CA 93106-5070 brianem@engr.ucsb.edu

This material is based upon work supported by the National Science Foundation under Grant NSF-ITR CCF-0326576 and the Institute for Collaborative Biotechnologies through Grant DAAD19-03-D-0004 from the U.S. Army Research Office.

understand their function without the aid of clear quantitative models. With such models at hand, it becomes possible to capture the various interactions among the known players and to provide a holistic system level understanding of the underlying biology. The new understanding that emerges also provides insight into how one may alter these networks and affect their function, making it possible to design and construct new biological parts and systems, or to re-design existing, natural biological systems for useful purposes. The latter is the main aim of the new field of synthetic biology.

In modeling biological phenomena at the cellular level, one might be initially tempted to propose deterministic models to describe the concentrations of key proteins and other molecules within the cell as a function of time. This is feasible so long as there are huge numbers of each reactant species, such that fluctuations in these concentrations are inconsequential. At the sub cellular level, however, processes can, and often do, depend upon individual molecular interactions. A single transcription factor binding to a single gene may initiate production of a key protein and eventually result in a cascade of events that affects the whole cell. Here random fluctuations have great impact on the cell, one that cannot be captured with a deterministic model. Instead, these processes must be modeled on the mesoscopic scale using discrete and stochastic models. This is the subject of stochastic chemical kinetics, an area that has received a lot of attention recently and one where notions from system theory are beginning to have a big impact. In the first part of this paper, we describe a new and very promising direction in the field of stochastic chemical kinetics that relies on finite state projections and ideas from system theory to provide probability densities of important biological states within prespecified errors. In the second part of this paper, we show the application of systems ideas in the area of biological model invalidation. As a specific example, we apply recent sum-of-squares techniques to show that a dynamic model of the bacterial heat-shock response must include a certain protein-protein interaction. Without the feedback loop enacted by that reaction, there exists no set of model parameters for which the system will exhibit the observed input-output data. This type of model invalidation is essential for the discovery of missing components or interactions and can guide new biological experiments.

In the next section, we give an overview of a new approach for stochastic computations that arise in stochastic chemical kinetics. In subsection II-A we briefly review the mesoscopic model for chemical kinetics. In subsection II-B we introduce the Finite State Projection (FSP) method, which implements

a new approach for getting approximate solutions of the so called chemical master equation. In subsection II-C, we review a few model reduction techniques which are easily applied to the FSP algorithm. These include several ideas from System Theory including concepts of controllability, observability, and time-scale separation. In section III we give an overview of how model invalidation for biological models can be performed using sum of squares techniques. Subsection III-A gives the theoretical background for the use of SOS methods in model invalidation, while subsection III-B demonstrates how these advanced methods can be used to invalidate a deficient biological model of the bacterial heat shock response.

II. A SYSTEM THEORETIC APPROACH TO STOCHASTIC CHEMICAL KINETICS

A. The Mesoscopic Model for Chemical Kinetics

In a mesoscopic model of an N species chemically reacting process, the configuration of the system is defined by the populations of the N species, $\mathbf{x} := [\xi_1 \ \xi_2 \ \dots \ \xi_N]^T$. Under the assumption that the chemical system is well mixed, the system behaves as a continuous time, discrete space Markov process. We are interested in computing the probabilities of the various possible system configurations as functions of time. For the configurations $\{\mathbf{x}_1, \mathbf{x}_2, \dots\}$, we define the corresponding probability density state vector as $\mathbf{P}(t) := [p_1(t), p_2(t), \dots]^T$. The evolution of $\mathbf{P}(t)$ over time can be described by the possibly infinite dimensional set of linear ordinary differential equations known as the Chemical Master Equation (CME) [1]: $\dot{\mathbf{P}} = \mathbf{A}\mathbf{P}$. Until recently, the CME had not been directly solved except for the most trivial problems, and analyses are often conducted using Monte Carlo (MC) algorithms.

The most widely used MC algorithm for mesoscopic chemical kinetics is the Stochastic Simulation Algorithm (SSA) [2]. In this algorithm, one uses a series of random numbers to simulate each individual reaction event in the chemical process. At each time, t , the process has a specific configuration, $\mathbf{x}(t) = \mathbf{x}_i$, and can change through at most M possible reactions which cause the process to transition from \mathbf{x}_i to $\mathbf{x}_j = \mathbf{x}_i + \nu_\mu$, where ν_μ is the stoichiometric vector of the μ^{th} reaction. Each μ^{th} reaction has an infinitesimal probability of occurring in the interval $(t, t + dt]$, and this probability is given by $a_\mu(\mathbf{x}_i, t)dt$, where $a_\mu(\mathbf{x}_i, t)$ is known as the *propensity function*. To use the SSA, one must assume that a_μ does not depend upon time, so that the time until the next reaction, τ , becomes an exponentially distributed random variable with a mean equal to $a_0(\mathbf{x}_i) = \left(\sum_{\mu=1}^M a_\mu(\mathbf{x}_i)\right)^{-1}$. In the MC scheme, one can choose τ from this exponential distribution, and one can choose which of the M reactions occurs according to the distribution defined by $\left\{\frac{a_1(\mathbf{x}_i)}{a_0(\mathbf{x}_i)}, \dots, \frac{a_M(\mathbf{x}_i)}{a_0(\mathbf{x}_i)}\right\}$. One then updates the time to $t + \tau$ and the configuration to $\mathbf{x}_i + \nu_\mu$ where μ is the index of the chosen reaction. The process is continued until the final time of interest. While this basic MC algorithm produces detailed realizations for the dynamics of

the mesoscopic model, the computational cost increases with the number of reactions in each process. In these cases, one may give up some of the accuracy of the SSA for faster approximate MC schemes such as time-leaping methods and system-partitioning methods.

Suppose that one can make the assumption that many reaction events will occur in a period of time without significantly changing the reaction propensity functions. This assumption is known as the *τ leap assumption* and is the basis of the τ leaping approximations to the SSA [3]. Rather than simulate every individual reaction event, one specifies a time interval and chooses the number of times each reaction channel fires during that interval according to a simple Poisson distribution. This approximate method can be much faster than the original SSA and performs particularly well for systems in which all molecular species are present in large numbers. However, there are two major obstacles to this approach, which occur when there are very small chemical populations in the system. First, if too many critical reactions are included in a single leap, some molecular populations may become negative and yield a meaningless result. More recent versions of τ -leaping, including binomial τ leaping [4] and adaptive and implicit τ leaping [5], [6] are more robust than their predecessors in this regard. However, the second obstacle to τ leaping methods comes from the fact that their accuracy remains severely compromised when very small populations of interacting molecules result in fast, dramatic changes in propensity functions.

The second approximation to the SSA involves splitting the system into slow and fast parts. The dynamics of the fast part are approximated, and the remaining slow dynamics are simulated with a reduced MC scheme [7], [8], [9]. For such a scheme to work, it is necessary that there is a sufficient gap between the between fast and slow reactions.

B. The Finite State Projection (FSP) algorithm

We recently proposed a new analytical approach to solving the CME: the Finite State Projection (FSP) algorithm [10]. Because the CME is a linear system whose solution is constrained to be non-negative with a sum of one, we showed that any finite state projection of the CME provides an analytical approximation to the CME within guaranteed error bounds. The FSP algorithm provides a means of systematically choosing a projection of the CME, which satisfies any prespecified accuracy requirement. Furthermore, the FSP algorithm is based upon linear systems theory, and it is ripe for the further application of modern controls theory as we will review in this paper.

In the chemical master equation, $\dot{\mathbf{P}}(t) = \mathbf{A}(t)\mathbf{P}(t)$, the generator matrix $\mathbf{A}(t)$ is comprised of the propensity functions for transitions from one configuration to another and is defined by the reactions and the enumeration of the configuration space. As is the case for all generator matrices, diagonal elements of $\mathbf{A}(t)$ are non-positive; off-diagonal elements are non-negative; and columns sum to zero.

If the configuration space were finite, then it would be straightforward to solve the CME using either a matrix

exponential for a constant \mathbf{A} or an ODE solver for a time-varying $\mathbf{A}(t)$. However, the dimension of $\mathbf{P}(t)$ may be extremely large or infinite, and the general problem may not be solved so easily. In this case, a projection may be made to achieve an arbitrarily accurate approximation. In order to show how this projection works, we use index sets of the form $J = \{j_1, j_2, j_3, \dots\}$. For any vector \mathbf{v} , we let \mathbf{v}_J denote the subvector of \mathbf{v} chosen according to J , and for any matrix \mathbf{A} , we let \mathbf{A}_{IJ} denote the submatrix of \mathbf{A} whose rows and columns have been chosen according to I and J , respectively. For example, if $I = \{3, 1\}$ and $J = \{3, 2\}$, then:

$$\begin{bmatrix} a & b & c \\ d & e & f \\ g & h & k \end{bmatrix}_J = \begin{bmatrix} k & h \\ c & b \end{bmatrix}.$$

With this notation, we can restate the theorems from [10] as follows, where we have extended Theorem 2.2 to include the case of time-varying $\mathbf{A}(t)$.

Theorem 2.1¹ If all off-diagonal elements of $\mathbf{A} \in \mathbb{R}^{n \times n}$ are non-negative, then for any index set J ,

$$[\exp(\mathbf{A})]_J \geq \exp([\mathbf{A}]_J) \geq \mathbf{0}.$$

Theorem 2.2 Consider any distribution which evolves according to the linear ODE: $\dot{\mathbf{P}}(t) = \mathbf{A}(t)\mathbf{P}(t)$. Let $\Phi_J(t_2, t_1)$ be the state transition matrix from time t_1 to time t_2 of the J -indexed finite state projection system $\dot{\mathbf{P}}_{FSP} = \mathbf{A}_J \mathbf{P}_{FSP}$.

If for $\varepsilon > 0$, and $t_f \geq 0$,

$$\mathbf{1}^T \Phi_J(t_f, 0) \mathbf{P}_J(0) \geq 1 - \varepsilon, \quad (1)$$

then

$$\Phi_J(t_f, 0) \mathbf{P}_J(0) \leq \mathbf{P}_J(t_f), \text{ and} \quad (2)$$

$$\|\Phi_J(t_f, 0) \mathbf{P}_J(0) - \mathbf{P}_J(t_f)\|_1 \leq \varepsilon \quad (3)$$

Proof: We begin by proving (2). Let J' denote the complement of J . The evolution of the probability distribution on the set J is governed by:

$$\dot{\mathbf{P}}_J(t) = \mathbf{A}_J(t) \mathbf{P}_J(t) + \mathbf{A}_{JJ'}(t) \mathbf{P}_{J'}(t), \quad (4)$$

where the submatrix $\mathbf{A}_{JJ'}(t)$ is nonnegative for any generator $\mathbf{A}(t)$. The solution for (4) is

$$\mathbf{P}_J(t_f) = \Phi(t_f, 0) \mathbf{P}_J(0) + \int_0^{t_f} \Phi(t_f, \tau) \mathbf{A}_{JJ'}(\tau) \mathbf{P}_{J'}(\tau) d\tau$$

Since $\mathbf{A}_{JJ'}(t)$, $\mathbf{P}_{J'}(t)$, and $\Phi(t, \tau)$ are all nonnegative for $t \geq \tau \geq 0$, we obtain the inequality in (2).

Since all probability distributions are non-negative and sum to one we are assured that $\|\mathbf{P}_J(t_f)\|_1 \leq 1$ and the condition (1) becomes:

$$\|\Phi(t_f, 0) \mathbf{P}_J(0)\|_1 \geq \|\mathbf{P}_J(t_f)\|_1 - \varepsilon, \quad (5)$$

Finally, applying (2) and rearranging terms yields (3) and completes the proof. ■

In the original form of [10], we assumed that \mathbf{A} did not vary in time, and the result is the same except that

$\Phi(t_2, t_1) = \exp(\mathbf{A}(t_2 - t_1))$. Using this FSP theorem, we can use an algorithmic approach to add and remove states in the finite projection until we obtain an error, ε , that is less than a prespecified bound.

C. Speeding up the FSP

In its most basic form, the FSP method is not feasible for every problem. The number of ODEs required for the finite projection to meet a given accuracy requirement may be far too large. However, there are many additional tools available from systems theory that can help us to meet this challenge. In the time invariant case, Krylov subspace methods for sparse systems could effectively enable the computation of the matrix exponential times a vector [11]. Furthermore, several readily available tools from modern control theory facilitate lower order approximations of larger systems and promise significant reductions in computational cost [12]. State aggregation, and multi-scale partitioning also provide enormous benefits over the original FSP [13], [14], [15]. Some of the underlying concepts of these reductions are summarized below in the context of time-invariant systems.

1) Observability, controllability, and minimal realizations:

Because the FSP approach formulates the CME as a finite dimensional problem, it opens the analysis to linear systems theory based model reductions as we explored previously in [12]. We can pose the initial value CME problem as an equivalent impulse response problem: $\dot{\mathbf{P}}(t) = \mathbf{A}\mathbf{P}(t) + \mathbf{b}\delta(t)$, where $\mathbf{b} = \mathbf{P}(0)$.

Suppose that we wish only to compute the output $\mathbf{y}(t) = \mathbf{C}\mathbf{P}(t)$. For example, \mathbf{y} may include statistical information such as means or variances, or could correspond to the probability of certain important portions of the configuration set. The resulting problem now takes on a familiar form:

$$\begin{aligned} \dot{\mathbf{P}}(t) &= \mathbf{A}\mathbf{P}(t) + \mathbf{b}\delta(t); \\ \mathbf{y}(t) &= \mathbf{C}\mathbf{P}(t). \end{aligned} \quad (6)$$

For systems constrained to a finite configuration set, or for systems that have been projected onto a finite configuration set using the FSP, this standard representation is open to a host of computational tools. In [12] we show that one can use concepts very similar to observability to reduce the system once it has been put into this form. With additional, more powerful, model reductions it may be possible to further enhance the applicability of the FSP.

2) *Multiple time scale partitioning:* Many biological models have certain reactions that occur much more frequently than others. These frequent reactions take up the vast majority of the total computational effort. In MC algorithms like the SSA [2], the majority of the simulated reactions correspond to those with large propensities. In the case of the CME or its finite state projection, this time scale disparity results in numerical stiffness. There has been significant progress in developing approximate MC algorithms to deal with these concerns [7], [8]. In these, the fast dynamics are essentially averaged, and the slow dynamics are simulated assuming instantaneous thermal equilibrium. The FSP algorithm is also amenable to time-partitioning approximation

¹For proof and additional details, see [10].

schemes that speed up computation at a small cost to the accuracy [13], [14]. In the original works, the time scale separation is carried out using perturbation theory. In this paper we take a linear systems theory approach.

In the configuration space, some subsets of configuration points are connected by fast reactions. Through a simple permutation of the CME, these points can be clustered together to form m fast groups. In turn, these clusters are connected to one another by slower reactions. After some permutation, the N dimensional finite state projection of the CME can be written:

$$\dot{\mathbf{P}} = \mathbf{A}\mathbf{P} = (\mathbf{H} + \epsilon\mathbf{G})\mathbf{P}, \quad (7)$$

where \mathbf{H} is a block diagonal matrix, $\mathbf{H} = \text{diag}\{\mathbf{H}_1, \mathbf{H}_2, \dots, \mathbf{H}_m\}$, and each \mathbf{H}_i is the generator matrix for the i^{th} fast cluster, and $\epsilon\mathbf{G}$ is the generator matrix of the reactions that take the system from one cluster to another. Each \mathbf{H}_i has a single eigenvalue equal to zero, and its corresponding left and right eigenvectors are \mathbf{u}_i and \mathbf{v}_i , respectively. We define the following matrices:

$$\mathbf{U} = \begin{bmatrix} \mathbf{u}_1 & 0 & \dots \\ 0 & \mathbf{u}_2 & \dots \\ \vdots & \vdots & \ddots \end{bmatrix}, \text{ and } \mathbf{V} = \begin{bmatrix} \mathbf{v}_1 & 0 & \dots \\ 0 & \mathbf{v}_2 & \dots \\ \vdots & \vdots & \ddots \end{bmatrix}.$$

Let $\mathbf{S} = [\mathbf{V} \quad \mathbf{R}]$ be a square matrix in which the columns of \mathbf{R} are the remaining $N - m$ right eigenvectors of \mathbf{H} . The inverse of \mathbf{S} is given by $\mathbf{S}^{-1} = [\mathbf{U}^T \quad \mathbf{L}^T]^T$ such that we have the following similarity transformation for \mathbf{H} :

$$\mathbf{S}^{-1}\mathbf{H}\mathbf{S} = \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Lambda} \end{bmatrix}, \quad \mathbf{\Lambda} = \text{diag}(\lambda_{m+1}, \dots, \lambda_N).$$

where the first m diagonal elements correspond to the zero eigenvalues of the \mathbf{H}_i blocks. Note that both \mathbf{S} and \mathbf{S}^{-1} have block structures similar to \mathbf{H} . The columns of \mathbf{S} and the rows of \mathbf{S}^{-1} are eigenvectors of the smaller \mathbf{H}_i blocks and are relatively inexpensive to compute.

We assume that non-zero eigenvalues of \mathbf{H} are ordered so that $0 > \text{Re}\{\lambda_{m+1}\} \geq \text{Re}\{\lambda_{m+2}\}, \dots \geq \text{Re}\{\lambda_N\}$. Applying the coordinate transformation $[\mathbf{y}_1^T(t) \quad \mathbf{y}_2^T(t)]^T = \mathbf{S}^{-1}\mathbf{P}(t)$, (7) becomes:

$$\begin{bmatrix} \dot{\mathbf{y}}_1(t) \\ \dot{\mathbf{y}}_2(t) \end{bmatrix} = \begin{bmatrix} \epsilon\mathbf{UGV} & \epsilon\mathbf{UGR} \\ \epsilon\mathbf{LGV} & \mathbf{Q} \end{bmatrix} \begin{bmatrix} \mathbf{y}_1(t) \\ \mathbf{y}_2(t) \end{bmatrix}, \quad (8)$$

where we have defined the matrix $\mathbf{Q} = \mathbf{\Lambda} + \epsilon\mathbf{LGR}$.

There are two important observations to make regarding this transformed system. First, the matrix \mathbf{UGV} is itself a generator for a Markov process in that it satisfies the two sufficient conditions: (i) its columns sum to zero, and (ii) its off-diagonal elements are non-negative. To show that, note that $\mathbf{1}^T\mathbf{U} = \mathbf{1}^T$ and therefore $\mathbf{1}^T\mathbf{UG} = \mathbf{1}^T\mathbf{G} = \mathbf{0}$. Furthermore, $[\mathbf{UGV}]_{ij} = \mathbf{u}_i\mathbf{G}_{J_i J_j}\mathbf{v}_j$, where \mathbf{u}_i and \mathbf{v}_j are non-negative for any (i, j) and the submatrix $\mathbf{G}_{J_i J_j}$ is non-negative for any $i \neq j$. Hence the off-diagonal elements of \mathbf{UGV} are indeed non-negative. The second observation that one can make is that for $\epsilon \ll |\text{Re}\{\lambda_{m+1}\}|$, linear perturbation theory assures us that the matrix \mathbf{Q} is Hurwitz,

and its eigenvalues are close to $\{\lambda_{m+1}, \lambda_{m+2}, \dots, \lambda_N\}$. In particular if we will let $\tilde{\lambda}$ denote the real part of the least stable eigenvalue of \mathbf{Q} , we know that $\tilde{\lambda} \approx \text{Re}\{\lambda_{m+1}\}$.

With these observations in mind, we now examine the forced dynamics: $\dot{\mathbf{y}}_2(t) = \epsilon\mathbf{LGV}\mathbf{y}_1(t) + \mathbf{Q}\mathbf{y}_2(t)$, which has a solution comprised of a zero-state and a zero-input response: $\mathbf{y}_2(t) = \mathbf{y}_2^{zs}(t) + \mathbf{y}_2^{zi}(t)$. Because \mathbf{Q} is Hurwitz, with eigenvalues all having real parts less than or equal to $\tilde{\lambda}$, the zero-input response, $\mathbf{y}_2^{zi}(t)$, is bounded by an exponentially decaying expression. Therefore, there exists a constant K_1 such that $|\mathbf{y}_2^{zi}(t)|_1 \leq K_1 \exp(\tilde{\lambda}t)$, for all $t \geq 0$.

By the definition of our transformation $|\mathbf{y}_1(t)|_1 = |\mathbf{UP}(t)|_1 = 1$, and $|\mathbf{LGV}\mathbf{y}_1(t)|_1$ is bounded. Since \mathbf{Q} is Hurwitz and the input is $O(\epsilon)$, we are guaranteed that the zero-state solution, $\mathbf{y}_2^{zs}(t)$ satisfies $|\mathbf{y}_2^{zs}(t)|_1 = O(\epsilon)$ for all $t \geq 0$. Combining the two solutions, we have the following bounds on $\mathbf{y}_2(t)$

$$|\mathbf{y}_2(t)|_1 \leq K_1 \exp(\tilde{\lambda}t) + O(\epsilon), \quad (9)$$

for all times $t \geq 0$.

The forced dynamics of $\mathbf{y}_1(t)$ given by the $\dot{\mathbf{y}}_1(t) = \epsilon\mathbf{UGV}\mathbf{y}_1(t) + \epsilon\mathbf{UGR}\mathbf{y}_2(t)$, has a solution at the chosen final time t_f :

$$\begin{aligned} \mathbf{y}_1(t_f) &= \exp(\epsilon\mathbf{UGV}t_f)\mathbf{y}_1(0) \\ &+ \epsilon \int_0^{t_f} \exp(\epsilon\mathbf{UGV}(t_f - \tau))\mathbf{UGR}\mathbf{y}_2(\tau) d\tau. \end{aligned}$$

Note that since \mathbf{UGV} is a infinitesimal generator of a Markov process, every column of $\exp(\mathbf{UGV}t)$ has a sum of exactly one for any $t \geq 0$, and $\|\exp(\epsilon\mathbf{UGV}(t - \tau))\|_1 = 1$ for all $\epsilon \geq 0$ and $t \geq \tau$. Therefore

$$|\mathbf{y}_1(t_f) - \exp(\epsilon\mathbf{UGV}t_f)\mathbf{y}_1(0)|_1 \leq \epsilon \int_0^{t_f} |\mathbf{UGR}\mathbf{y}_2(\tau)|_1 d\tau.$$

Combining this with (9) and defining the constant $K_2 = K_1 \|\mathbf{UGR}\|_1$, one obtains the following bound on the error of \mathbf{y}_1 at $t = t_f$:

$$\begin{aligned} &|\mathbf{y}_1(t_f) - \exp(\epsilon\mathbf{UGV}t_f)\mathbf{y}_1(0)|_1 \\ &\leq \epsilon \int_0^{t_f} K_2 \exp(\tilde{\lambda}\tau) + O(\epsilon) d\tau \\ &\leq \epsilon K_2 \frac{1}{|\tilde{\lambda}|} + t_f O(\epsilon^2). \end{aligned}$$

Therefore, for any fixed $t_f \geq 0$,

$$|\mathbf{y}_1(t_f) - \exp(\epsilon\mathbf{UGV}t_f)\mathbf{y}_1(0)|_1 = O(\epsilon). \quad (10)$$

Combining (9) and (10), we have the following bounds on our approximation error:

$$\left\| \begin{bmatrix} \mathbf{y}_1(t_f) \\ \mathbf{y}_2(t_f) \end{bmatrix} - \begin{bmatrix} \exp(\epsilon\mathbf{UGV}t_f)\mathbf{y}_1(0) \\ \mathbf{0} \end{bmatrix} \right\|_1 \leq K_1 \exp(\tilde{\lambda}t_f) + O(\epsilon).$$

Substituting the initial condition,

$$\begin{bmatrix} \mathbf{y}_1(0) \\ \mathbf{y}_2(0) \end{bmatrix} = \mathbf{S}^{-1}\mathbf{P}(0) = \begin{bmatrix} \mathbf{UP}(0) \\ \mathbf{LP}(0) \end{bmatrix},$$

and performing the reverse similarity transformation, $\mathbf{P}(t_f) = \mathbf{V}\mathbf{y}_1(t_f) + \mathbf{R}\mathbf{y}_2(t_f)$, yields:

$$\begin{aligned} & |\mathbf{P}(t_f) - \mathbf{V} \exp(\varepsilon \mathbf{U} \mathbf{G} \mathbf{V} t_f) \mathbf{U} \mathbf{P}(0)|_1 \\ & \leq K_1 \|\mathbf{R}\|_1 |\mathbf{L} \mathbf{P}(0)|_1 \exp(\tilde{\lambda} t_f) + O(\varepsilon). \end{aligned}$$

Thus, this reduced model differs from the full system by at most an exponentially decreasing transient term plus a term of order ε . However, the reduced system can have a much smaller dimension and its solution is much easier to compute.

III. MODEL INVALIDATION USING THE SUM OF SQUARES APPROACH

A. Theory

Model invalidation techniques provide a way to rigorously demonstrate the inability of a given model to capture observed system behaviors. The role of model invalidation techniques is to *invalidate* a model, by proving that some experimental data are inconsistent with the model, thus indicating that a refinement of the model is required.

In this section, we present a recently proposed methodology for invalidation of nonlinear models with noise and uncertain parameters [16]. The methodology relies on finding functions called *barrier certificates*. The existence of a barrier certificate generates a contradiction between model and some time-domain experimental data, in the sense that some level sets of this certificate act as barriers between possible model trajectories and data. The advantage of this approach is that barrier certificates can be computed using convex methods via the sum of squares decomposition and semidefinite programming, e.g., using the SOSTOOLS software [17], [18].

Given the system of ordinary differential equations

$$\dot{x}(t) = f(x(t), p, t), \quad (11)$$

where $x(t) \in \mathbb{R}^n$ is the vector of state variables, t is time, and $p \in \mathbb{R}^m$ is the parameter vector, assumed to take its value in a set $P \subset \mathbb{R}^m$. Suppose an experiment has been performed with the real system, and two measurements have been taken at time $t = 0$ and $t = T$. Suppose further that these measurements indicate that $x(0) \in X_0$ and $x(T) \in X_T$, where both X_0 and X_T are subsets of \mathbb{R}^n . Assuming that $x(t) \in X$ for all $t \in [0, T]$, where $X \subseteq \mathbb{R}^n$, the invalidation problem can be stated as follows:

Problem 1: Given the model (11), parameter set P , and trajectory information $\{X_0, X_T, X\}$, prove that for all possible parameter $p \in P$, the model (11) cannot produce a trajectory $x(t)$ such that $x(0) \in X_0$, $x(T) \in X_T$, and $x(t) \in X, \forall t \in [0, T]$. When such a proof is found, we say that the model (11) and parameter set P are invalidated by $\{X_0, X_T, X\}$.

Traditionally, model invalidation is achieved by running a very large number of simulations of (11) with the parameters p and initial conditions $x(0)$ randomly chosen from P and X_0 . If after such exhaustive simulations no trajectory is found that is consistent with the data, the model is declared invalidated. Clearly, such an invalidation, cannot be

considered conclusive, as it is impossible to test all parameter values p and all initial conditions $x(0)$.

As an alternative approach, we describe a method that is based on the creation of a certain function of state, parameters, and time, which is referred to as a barrier function, or a barrier certificate. A barrier certificate yields an *exact* proof of inconsistency between model and data by providing a barrier between feasible trajectories of the model starting at X_0 and the final measurement set X_T . One of the main advantages of this approach is that no simulations of the model are required. The barrier certificate approach is captured by the following:

Theorem 2 ([16]): Let the model (11) and the sets P, X_0, X_T, X be given, with $f(x, p, t)$ being continuous in x and t . Suppose that there exists a real-valued function $B(x, p, t)$ that is differentiable with respect to x and t , such that

$$\begin{aligned} B(x_T, p, T) - B(x_0, p, 0) &> 0 \\ \forall (x_T, x_0, p) &\in X_T \times X_0 \times P, \quad (12) \end{aligned}$$

$$\begin{aligned} \frac{\partial B}{\partial x}(x, p, t) f(x, p, t) + \frac{\partial B}{\partial t}(x, p, t) &\leq 0 \\ \forall (x, p, t) &\in X \times P \times [0, T]. \quad (13) \end{aligned}$$

Then the model (11) and its associated parameter set P are invalidated by $\{X_0, X_T, T\}$. The function $B(x, p, t)$ is referred to as a barrier certificate.

The proof of this result should be fairly intuitive. From the experimental data and (12), $B(x(t), p, T)$ is greater than $B(x(t), p, 0)$. Yet, (13) implies that $B(x(t), p, t)$ must not increase in time whenever $x(t)$ is a valid trajectory of the model (11). Hence, the model must be inconsistent with the experimental data.

The crux of the invalidation lies in the construction of barrier certificates. Fortunately, for models with polynomial vector fields and sets P, X_{T_i}, X, U_i described by polynomial equalities and inequalities, a new computationally tractable relaxation for constructing barrier certificates exists. This relaxation relies on the sum of squares decomposition and therefore can be achieved using convex programming. To see this, start with the model (11) with a polynomial $f(x, p, t)$ and parameter set P defined as follows:

$$P = \{p \in \mathbb{R}^m : g_{P,i}(p) \geq 0 \quad \forall i \in I_P\} \quad (14)$$

where the $g_{P,i}(p)$'s are polynomials in p , and I_P is an index set. Similarly, let the trajectory data be defined by

$$X_0 = \{x_0 \in \mathbb{R}^n : g_{0,i}(x_0) \geq 0 \quad \forall i \in I_0\}, \quad (15)$$

$$X_T = \{x_T \in \mathbb{R}^n : g_{T,i}(x_T) \geq 0 \quad \forall i \in I_T\}, \quad (16)$$

$$X = \{x \in \mathbb{R}^n : g_{X,i}(x) \geq 0 \quad \forall i \in I_X\}. \quad (17)$$

Then a barrier certificate can be computed by solving the convex optimization problem:

Program 3: Let the polynomial vector field $f(x, p, t)$ and the sets P, X_0, X_T, X in (14)–(17) be given. To invalidate the model (11), find a polynomial $B(x, p, t)$, a positive number ε , and sums of squares $M_{P,i}(x_0, x_T, p)$, $M_{0,i}(x_0, x_T, p)$,

$M_{T,i}(x_0, x_T, p)$, $N_{P,i}(x, p, t)$, $N_{X,i}(x, p, t)$, $N_t(x, p, t)$, such that the expressions

$$B(x_T, p, T) - B(x_0, p, 0) - \epsilon - \sum_{i \in I_P} M_{P,i}(x_0, x_T, p) g_{P,i}(p) - \sum_{i \in I_0} M_{0,i}(x_0, x_T, p) g_{0,i}(x_0) - \sum_{i \in I_T} M_{T,i}(x_0, x_T, p) g_{T,i}(x_T) \quad (18)$$

and

$$-\frac{\partial B}{\partial x}(x, p, t) f(x, p, t) - \frac{\partial B}{\partial t}(x, p, t) - \sum_{i \in I_P} N_{P,i}(x, p, t) g_{P,i}(p) - \sum_{i \in I_X} N_{X,i}(x, p, t) g_{X,i}(x) - N_t(x, p, t)(Tt - t^2) \quad (19)$$

are sums of squares.

Next, we demonstrate this approach by invalidating a biological model which lacks a key regulatory feedback loop.

B. Model Invalidation in The Heat Shock Response in *E. coli*

High temperatures cause cell proteins to unfold from their native three dimensional structures, resulting in loss of cellular function. Cells have evolved gene regulatory mechanisms to counter the effects of heat shock by expressing specific genes that encode heat shock proteins (hsps) whose role is to help the cell survive the consequence of the shock. In *E. coli*, the heat shock (HS) response is implemented through an intricate architecture of feedback loops centered around the σ - factor that regulates the transcription of the HS proteins under normal and stress conditions. The enzyme RNA polymerase (RNAP) bound to this regulatory sigma factor, σ^{32} , recognizes the HS gene promoters and transcribes specific HS genes. The HS genes encode predominantly molecular chaperones that are involved in refolding denatured proteins and proteases that function to degrade unfolded proteins.

At physiological temperatures, there is very little σ^{32} present and hence little transcription of the HS genes. When bacteria are exposed to high temperatures, σ^{32} first rapidly accumulates, allowing increased transcription of the HS genes and then declines to a new steady state level characteristic of the new growth temperature. There are two mechanisms by which σ^{32} levels are increased when the temperature is raised [19]. First, the translation rate of the rpoH mRNA (encoding σ^{32}) increases immediately, resulting in a fast 10-fold increase in the concentration of σ^{32} [20]. This mechanism implements what we refer to as the *feedforward control loop*. Second, during steady state growth, σ^{32} is rapidly degraded ($t_{1/2} = 1$ minute), but is stabilized for the first five minutes after temperature upshift, so that its concentration rapidly increases.

In vivo evidence is consistent with the following titration model for the HS response. The chaperone DnaK, and its cochaperone DnaJ are required for the rapid degradation of σ^{32} by the HS protease FtsH. Raising the temperature produces an increase in the cellular levels of unfolded proteins that then titrate DnaK/J away from σ^{32} , allowing it to bind

to RNA polymerase (resulting in increased transcription) and stabilizing it in the process. Together, increased translation and stabilization lead to a transient 15-20 fold increase in the amount of σ^{32} at the peak of the HS response. The accumulation of high levels of HS proteins leads to the efficient refolding of the denatured proteins thereby decreasing the pool of unfolded protein, freeing up DnaK/J to sequester this protein from RNA polymerase. This implements what is referred to as a *sequestration feedback loop*. Furthermore, this sequestration itself promotes the degradation of σ^{32} and results in feedback regulated degradation, mainly by the protease FtsH. We refer to this as the *FtsH degradation feedback loop*. The overall result is a decrease in the concentration of σ^{32} to a new steady state concentration that is dictated by the balance between the temperature-dependent translation of the rpoH mRNA and the level of σ^{32} activity modulated by the hsp chaperones and proteases acting in a negative feedback fashion.

1) *A Model for the HS Response:* In a previous work, we have developed a detailed deterministic mathematical model for the heat stress response in *E. coli* [21], [22], [23]. Using first order kinetics (law of mass-action). The model describes the molecular interactions described above, namely the synthesis of new proteins, and the association/dissociation activity of molecules. The dynamics described above were modeled using differential rate equations, generating a set of 31 Differential-Algebraic Equations (DAEs) of the form

$$\begin{aligned} \dot{X}(t) &= F(t; X; Y) \\ 0 &= G(t; X; Y) \end{aligned}$$

where X is a 11-dimensional vector whose elements are the differential variables and Y is a 20-dimensional vector whose elements are algebraic variables. This form is known as a *semi-explicit* DAE. The model possesses 27 kinetic rate parameters. Subsequently, a reduced order model was derived using insight into the system's architecture and separation principles in time and concentrations. As in the full model, this reduced model involves the dynamics of the basic building blocks of the HS response, namely the σ factor (S), the chaperones (D), and the protein folding mechanism. The model equations are as follows

$$\begin{aligned} \frac{dD_t}{dt} &= K_d S_f - \alpha_d D_t \\ \frac{dS_t}{dt} &= \eta(T) - \alpha_0 S_t - \alpha_s S : D \\ \frac{dU_f}{dt} &= K(T) P_{folded} - K_{fold} U : D \\ S : D &= K_s \cdot S_f \cdot D_f \\ U : D &= K_u \cdot U_f \cdot D_f \\ D_t &= D_f + U : D + S : D \\ S_t &= S_f + S : D \\ P_t &= P_{folded} + U_f + U : D \end{aligned} \quad (20)$$

where $U : D$ is the complex formed by the binding of the unfolded proteins U_f to D , $S : D$ is the complex formed by the binding of S to D , and P_t is the total number of proteins

Parameter	Value
K_d	3 min^{-1}
α_d	0.015 min^{-1}
$\eta(T)$	$10 \text{ molecule} \cdot \text{min}^{-1} @ T_1 \text{ \& } 60 @ T_2$
α_0	0.03 min^{-1}
α_s	3 min^{-1}
K_s	$0.05 \text{ molecule}^{-1}$
K_u	$0.0254 \text{ molecule}^{-1}$
$K(T)$	$40 \text{ min}^{-1} @ T_1 \text{ \& } 80 \text{ min}^{-1} @ T_2$
K_{fold}	6000 min^{-1}
P_t	$2 \times 10^6 \text{ molecules}$

TABLE I
PARAMETER VALUES FOR HEAT SHOCK MODEL

in the cell, considered here to be constant. The parameters used in this model are given in Table I.

We replace the algebraic constraints into the initial system (20), then use the facts that $S_t \ll D_t$ and that $U_f \gg 1$ in the wild type bacterial HS response and simplify the expression for S_f and D_f . Simple algebraic manipulations yield a compact description for the reduced order HS model:

$$\begin{aligned} \frac{dD_t}{dt} &= f_1(D_t, U_f, S_t) - \alpha_d D_t \\ \frac{dS_t}{dt} &= \eta(T) - \alpha_0 S_t - f_2(D_t, U_f, S_t) \\ \frac{dU_f}{dt} &= K(T)[P_t - U_f] - [K(T) + K_{fold}]D_t \end{aligned} \quad (21)$$

As in the original equations, the feedforward control is achieved by the temperature dependent function $\eta(T)$ in the ODE describing the dynamics of S_t . $f_1(D_t, U_f, S_t) = K_d \frac{S_t}{1 + \frac{K_s D_t}{K_u U_f}}$ and $f_2(D_t, U_f, S_t) = \alpha_s \frac{1 + K_u U_f}{1 + \frac{K_s D_t}{K_u U_f}} S_t$ describe the various feedback strategies implemented in the HS response. f_1 is the effect of the sequestration of S by D on D formation, while f_2 reflects the effect of the regulated degradation of S through the action of the sequestration itself. The dynamics of the third state U_f are much faster than those of S_t and D_t . Such stiffness is also strongly present in the full model and creates ill-conditioning and algorithms that don't exploit stiffness are almost certainly doomed to suffer from it. However, stiffness can also be exploited to robustly produce simplified models by singular perturbation, as was done in deriving the 3-state from the full model. By further setting $\frac{dU_f}{dt} = 0$ to obtain a quasi-steady state approximation, the third equation is then replaced by an algebraic one, and the result is again a differential-algebraic equation (DAE). The validity of this approximation has been verified by simulation which showed virtually no difference in the solution of the ODE as compared to that of the DAE.

2) *Invalidation of the Model of Heat Shock Response in E. coli*: As we pointed out earlier, the creation of barrier certificates using SOSTOOLS can be used in model invalidation in biological modeling. The key ideas of this methodology can be illustrated in the context of the heat-shock example, where at least two feedback loops are involved in the regulation scheme. We will show rigorously that each loop adds its own important function to the overall system and that both

are necessary to explain the observed behavior of the heat shock system.

In previous work, we have used sensitivity analysis and confirmed that these feedback loops indeed increase the robustness to parametric uncertainty [21]. However, upon disabling the degradation (FtsH) feedback loop, one observes in simulation that the transient response to a temperature increase becomes considerably slower. Achieving a faster transient response in the absence of this (FtsH) feedback loop necessitates a substantial increase in the protein synthesis rate, and therefore, produces a larger number of chaperones. Therefore, it is reasonable to hypothesize that the (FtsH) feedback loop is essential for achieving a fast response to the heat disturbance.

To illustrate how we might actually prove such a hypothesis using the invalidation scheme in Section III-A, we perform an experiment with the system to obtain some data that will be used in the construction of a barrier certificate to invalidate a model. We shall assume that the "real system" is the model with the degradation (FtsH) loop (21) in place, and use that data to invalidate a 'deficient model' lacking this feedback. If we denote the state variables (D_t, S_t, U_f) by (x_1, x_2, x_3) , then the deficient model will just be $\dot{x} = f(x, p)$, where the vector field are defined by (21), without the degradation loop.

An "experiment" (we use a numerical simulation of the full system with both feedback loops in place) is performed, with the parameters fixed at the nominal values. We observe that the corresponding system trajectory satisfies $x(0) \in X_0$ and $x(25) \in X_T$, where

$$\begin{aligned} X_0 = \{ &(x_1, x_2, x_3) \in \mathbb{R}^3 : 0.9D_0 \leq x_1 \leq 1.5D_0 \\ &0.9S_0 \leq x_2 \leq 1.5S_0, 2.9U_0 \leq x_3 \leq 3.1U_0 \} \end{aligned} \quad (22)$$

$$\begin{aligned} X_T = \{ &(x_1, x_2, x_3) \in \mathbb{R}^3 : 1.5D_0 \leq x_1 \leq 2.5D_0, \\ &2S_0 \leq x_2 \leq 3S_0, 0.5U_0 \leq x_3 \leq 1.5U_0 \} \end{aligned} \quad (23)$$

with D_0 , S_0 , and U_0 denoting their steady state values at low temperature. Intervals are used in order to take into account the effects of measurement noise and uncertainty in initial conditions. We observe that from time $t = 0$ to $t = 25$, the state variables satisfy $x(t) \in X$, with

$$\begin{aligned} X = \{ &(x_1, x_2, x_3) \in \mathbb{R}^3 : 0.9D_0 \leq x_1 \leq 2.5D_0, \\ &0.9S_0 \leq x_2 \leq 8S_0, 0.2U_0 \leq x_3 \leq 4U_0 \}. \end{aligned} \quad (24)$$

which we add to our apriori data.

For the deficient model, we will focus on three parameters $p = (K_d, \alpha_0, \eta(T))$, and assume that the rest are fixed at the nominal values. Plausible ranges for these parameters define the parameter set P :

$$\begin{aligned} P = \{ &(K_d, \alpha_0, \eta(T)) \in \mathbb{R}^3 : 0.5\bar{K}_d \leq K_d \leq 5\bar{K}_d, \\ &0.5\bar{\alpha}_0 \leq \alpha_0 \leq 1.5\bar{\alpha}_0, 0.5\bar{\eta}(T) \leq \eta(T) \leq 1.5\bar{\eta}(T) \}, \end{aligned} \quad (25)$$

where \bar{K}_d , $\bar{\alpha}_0$, and $\bar{\eta}(T)$ denote their nominal values. We deliberately make the upper bound for K_d quite large, since one obvious way for obtaining a fast response is to increase

the number of chaperones, corresponding to increasing this parameter.

Using SOSTOOLS, we search for (and find) a barrier certificate for these model and data,

$$\begin{aligned}
B(x, t) = & 1.4233t - 4.4566x_2 - .0051737x_3 + .021541x_1tx_3 \\
& + .016599tx_1^2 + .020129tx_2^2 - .081828x_1t - .054129x_3t \\
& - .71041x_2t + .54734 \times 10^{-4}x_3tx_2 + .0022033x_1tx_2 \\
& + .0067639x_1 - .76233 \times 10^{-3}x_3x_2 + .0010018x_3x_1 \\
& + .0052243t^2x_1 + .02375t^2x_2 - .66663 \times 10^{-3}t^2x_3 \\
& - .68942 \times 10^{-3}t^2x_1^2 - .0024426t^2x_2^2 \\
& + .74858 \times 10^{-4}t^2x_3^2 + .0029471t^2 + .95341 \times 10^{-3}x_1^2 \\
& + .010339x_2^2 + .27438 \times 10^{-3}x_1t^2x_3 \\
& - .40973 \times 10^{-5}t^2x_3x_2 - .79472 \times 10^{-3}t^2x_1x_2 \\
& - .015999x_2x_1 + .0052841x_3^2t + .0019683x_3^2,
\end{aligned}$$

in effect proving that the model without the degradation (FtsH) loop and with parameters $K_d, \alpha_0, \eta(T)$ satisfying (25) cannot possibly generate a time response that satisfies (22)–(24). This in turn suggests that an inherent mechanism is missing from this model. In practice, the data is generated by real experiments, and the invalidation of a proposed model will provide very useful information to the biologist because it suggests either missing components, missing interactions, or both. In some ways, model invalidation is more important to the experimental biologist than creating a quantitative model that is consistent with the experiments.

IV. CONCLUSIONS

In this paper we have presented two instances in which systems theory has recently been applied to analyze biological networks. In the first example, we showed how linear systems theory led to the development and improvement of the Finite State Projection (FSP) method for the solution to the Chemical Master Equation (CME). Here, we have extended that method to the case of time varying reaction propensity functions. We have also included a new systems theory based proof to show how a time scale separation reduction scheme can significantly increase the efficiency of the FSP method. These and other applications of systems theory have made the solution of the CME a computationally tractable problem for a much wider range of stochastic gene regulatory systems. In the second section, we have shown how systems theory can be used to systematically aid in the development of gene regulatory models. For a proposed model and a given set of input and output data, one can use sum of squares techniques to ascertain if there exists any possible set of parameters for which the model may be valid. Otherwise, the model is known to be insufficient. This approach has been applied to show that a particular model of the heat shock response in *E. coli* cannot possibly be valid unless it includes a key feedback mechanism. This type of reasoning can lead researchers to (1) pinpoint critical gene regulatory mechanisms; (2) theorize new mechanisms to account for discrepancies between model predictions and

experimental observations; and (3) suggest additional experiments to test these theories.

REFERENCES

- [1] van Kampen. *Stochastic Processes in Physics and Chemistry*. Elsevier, 3 edition, 2001.
- [2] D. T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.*, 81(25):2340–2360, May 1977.
- [3] D. T. Gillespie. Approximate accelerated stochastic simulation of chemically reacting systems. *J. Chem. Phys.*, 115(4):1716–1733, Jul. 2001.
- [4] T. Tian and K. Burrage. Binomial leap methods for simulating stochastic chemical kinetics. *J. Chem. Phys.*, 121(21):10356–10364, Dec. 2004.
- [5] Y. Cao, D. T. Gillespie, and L. R. Petzold. Avoiding negative populations in explicit poisson tau-leaping. *J. Chem. Phys.*, 123(054104), 2005.
- [6] M. Rathinam, L. R. Petzold, Y. Cao, and D. T. Gillespie. Stiffness in stochastic chemically reacting systems: The implicit tau-leaping method. *J. Chem. Phys.*, 119(24):12784–12794, Dec. 2003.
- [7] C. V. Rao and A. P. Arkin. Stochastic chemical kinetics and the quasi-steady-state assumption: Application to the Gillespie algorithm. *J. Chem. Phys.*, 118(11):4999–5010, Mar. 2003.
- [8] Y. Cao, D. Gillespie, and L. Petzold. The slow-scale stochastic simulation algorithm. *J. Chem. Phys.*, 122(014116), Jan. 2005.
- [9] H. Salis and Y. Kaznessis. Accurate hybrid stochastic simulation of a system of coupled chemical or biological reactions. *J. Chem. Phys.*, 112(054103), 2005.
- [10] B. Munsky and M. Khammash. The finite state projection algorithm for the solution of the chemical master equation. *J. Chem. Phys.*, 124(044104), 2006.
- [11] K. Burrage, M. Hegland, S. Macnamara, and R. Sidje. A Krylov-based finite state projection algorithm for solving the chemical master equation arising in the discrete modelling of biological systems. *Proc. of The A.A. Markov 150th Anniversary Meeting*, 2006.
- [12] B. Munsky and M. Khammash. A reduced model solution for the chemical master equation arising in stochastic analyses of biological networks. *Proc. 45th IEEE CDC*, Dec. 2006.
- [13] S. Peles, B. Munsky, and M. Khammash. Reduction and solution of the chemical master equation using time-scale separation and finite state projection. *J. Chem. Phys.*, 125(204104), Nov. 2006.
- [14] B. Munsky, S. Peles, and M. Khammash. Stochastic analysis of gene regulatory networks using finite state projections and singular perturbation. *Submitted to ACC*, Jul. 2007.
- [15] B. Munsky and M. Khammash. Modeling and analysis of a bacterial stochastic switch. *Proc. 14th Mediterranean Conference on Control and Automation MED06*, Jun. 2006.
- [16] S. Prajna. Barrier certificates for nonlinear model validation. *Automatica*, 42(1):117–126, 2006.
- [17] J. F. Sturm. Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optimization Methods and Software*, 11–12:625–653, 1999. Available at <http://fewcal.kub.nl/sturm/software/sedumi.html>.
- [18] S. Prajna, A. Papachristodoulou, and P. A. Parrilo. Introducing SOSTOOLS: A general purpose sum of squares programming solver. In *Proceedings of the 41st IEEE Conf. on Decision and Control*, 2002.
- [19] D.B. Straus, W.A. Walter, and C.A. Gross. *dnak, dnaj, and grpe* heat shock proteins negatively regulate heat shock gene expression by controlling the synthesis and stability of σ^{32} . *Gen. Dev.*, 4(12A):2202–2209, 1990.
- [20] M.T. Morita et al. Translational induction of heat shock transcription factor σ^{32} : Evidence for a built-in rna thermosensor. *Genes & Dev.*, 13(6):655–665, 1999.
- [21] H. El-Samad, M. Khammash, H. Kurata, and J.C. Doyle. Robustness Analysis of the Heat Shock Response in *E. coli*. In *Proceedings of the American Control Conference*, pages 1742–1747, 2002.
- [22] H. Kurata, H. El-Samad, T.M. Yi, M. Khammash, and J.C. Doyle. Feedback Regulation of the Heat Shock Response in *E. coli*. In *Proceedings of the 40th IEEE Conference on Decision and Control*, pages 837–842, 2001.
- [23] H. El-Samad, H. Kurata, J. Doyle, C. Gross, and M. Khammash. Surviving heat shock: Control strategies for robustness and performance. *Proc. Natl. Acad. Sci. USA*, 102(9):2736–2741, 2005.