Stochastic Reachability Analysis in Complex Biological Networks

H. El-Samad, M. Fazel, X. Liu, A. Papachristodoulou, S. Prajna.

Abstract— **The impact of noise on cellular networks and its interplay with their rich dynamics are increasingly being characterized as important phenomena that must be thoroughly investigated for a useful understanding of biological dynamics. At the same time, the mathematical modeling and analysis of these networks in a stochastic setting presents a number of challenges, such as the need for a large number of computationally expensive stochastic simulations to collect statistics about the occurrence of important events or correlate their occurrence with the noise intensity. In this paper, we demonstrate the use of new techniques of stochastic reachability analysis to address these problems. Specifically, we study the problem of computing bounds on the probability of a biological stochastic process to reach certain parts of the state space in a finite time. The techniques presented are based on the algorithmic construction of barrier certificates using convex optimization, and are illustrated through the use of a biologically important system: the bacteriophage λ genetic switch.**

I. INTRODUCTION

The cellular networks that implement life's programs operate in the uncertain environment of the cell, and are continuously exposed to environmental noise, biochemical noise resulting from the discrete nature of the molecular populations, and thermal fluctuations that lead to stochastic fluctuations in their molecular species. The effects of such discreteness and stochasticity become especially noticeable when the molecular populations of some reactant species are very small, or if the dynamic structure of the system makes it susceptible to noise amplification. Physical evidence pointing to purely stochastic phenomena in a number of genetic networks is abundant [1], [2], [3]. Noise dependent phenomena become particularly interesting in oscillatory systems and multi-stable systems, where the interaction of large noise intensities could lead to random switching from one cellular state to the other [4]. Even in the presence of small noise intensities, switching events can still happen on long time scales. It has been hypothesized that such rare large deviation events can be responsible for important failures of cellular machinery, leading to cancer [5] and the release of latent viruses from their quiescent state [7], [6]. The characterization of these events and the statistics of their occurrence is therefore of great importance, but also of great

H. El-Samad is with the Biochemistry/Biophysics department at UC, San Francisco and the California Institute for Quantitative Biomedical Research. Email: helsamad@biochem.ucsf.edu

M. Fazel, X. Liu, and S. Prajna are with Control and Dynamical Systems, California Institute of Technology, Pasadena, CA 91125. E-mails: {maryam,xinliu,prajna}@cds.caltech.edu

A. Papachristodoulou is with the department of Engineering Science at Oxford University. Email: antonis@eng.ox.ac.uk

difficulty because of the inefficient simulation-dependent stochastic methods that are used. The computational burden of investigating rare stochastic events is exacerbated by the necessity of running extensive simulations in order to capture even one occurrence of the event of interest, and prohibitively large numbers of simulations in order to compile statistics. In this paper, we illustrate the use of new methods based on sum of squares optimization, linear programming, and the SOSTOOLS to tackle such stochastic reachability analysis problems algorithmically. Specifically, we demonstrate the computation of bounds on the probability of occurrence of noise-driven biological events, in addition to the dependence of this probability on noise intensity. As an illustration, we present a thorough analysis of the noise-induced switching between the two stable equilibria of an important genetic network implemented in the bacteriophage λ.

II. STOCHASTIC MODELING OF BIOLOGICAL NETWORKS

The stochastic investigation of biological systems mainly relies on the formulation of the Chemical Master Equation (CME), a differential equation for the time evolution of probabilities. Here, we only give the expression for the CME and refer the reader to [9] or [10] for a more detailed account.

In a chemically reacting system involving N molecular species $S_1, \ldots S_N$ reacting through M reaction channels $R_1 \ldots R_M$, we define the state vector as $Z(t) = [Z_1(t) \ldots Z_N(t)]^T$, where $Z_i(t)$, $i = 1, 2 \ldots, N$ is a random number that defines the number of molecules of species S_i at time t. We assume that the system is well stirred and in thermal equilibrium. Under these circumstances, each reaction channel R_k is characterized by a *propensity function a_k* and an N-dimensional *state change vector $s_k = [s_{1k} \ldots s_{Nk}]^T$*. The vector s_k represents the stoichiometric change of the molecular species by an R_k reaction, while $a_k(z)dt$ is the probability that one R_k reaction will occur in state $Z(t) = z$ during the next infinitesimal time interval dt. We let $a(\cdot) = [a_1(\cdot), \ldots, a_M(\cdot)]^T$ be the vector of propensity functions and $S = \{s_k\}_{k=1,\ldots,M}$ be the stoichiometry matrix. The CME for the system is then given by $\frac{\partial P(z,t|z_0,t_0)}{\partial t} = \sum_{k=1}^{M} [a_k(z - s_k)P(z - s_k, t|z_0, t_0) - a_k(z)P(z, t|z_0, t_0)]$, where $P(z,t|z_0,t_0)$ is the probability that at time t, $Z(t) = z$ given that $Z(t_0) = z_0$ (z and z_0 are integers). The CME also assumes a Markov Chain interpretation. Essentially, given N possible molecular species of interest, the set of all possible states is $\mathbb{N}^N$. One can *a priori* fix a sequence $x_1, x_2, \ldots$ of elements in $\mathbb{N}^N$ and define $\mathcal{X} := [x_1, x_2, \ldots]^T$. The particular sequence may be chosen to visit every element of $\mathbb{N}^N$. In this case, the choice of $\mathcal{X}$ corresponds to a particular enumeration of the space $\mathbb{N}^N$. Once $\mathcal{X}$ is selected, the CME can be written as

a single linear (but generally infinite) expression $\dot{\mathbf{P}}(\mathscr{X};t) = \mathbf{B} \cdot \mathbf{P}(\mathscr{X};t)$, where $\mathbf{P}(\mathscr{X};t) := [P(x_1,t), P(x_2,t), \ldots]^T$, is the complete probability density state vector at time t, and $\mathbf{B}$ is the *state reaction matrix*. The columns and rows of $\mathbf{B}$ are uniquely defined by the system's stoichiometry and the choice of $\mathscr{X}$. Beginning at any state, $\mathbf{x}_i$, there can be a maximum of $\mathbf{M}$ possible reactions; each reaction leads to a different state: $\mathbf{x}_j = \mathbf{x}_i + s_j$. We will refer to this formulation as the "Discretized CME".

Despite the fairly straight-forward form and interpretation of the CME, it is not solvable exactly for any but systems that are "closed", in the sense that reactants are modeled to exist in a fixed volume and are not permitted to interact beyond the confines of that volume. However, models of cellular networks tend not to be closed, due to the prevalence of reactions that re-introduce into the reaction volume new copies of reactants through events such as new protein synthesis. An alternative to solving the CME is to generate stochastic sample realizations of the process whose probability density function is described by the CME. This can be done using Monte Carlo techniques such as the Gillespie stochastic simulation algorithm (SSA) [11]. However, the SSA becomes prohibitively inefficient in many situations, such as when the problem is stiff. To circumvent this stiffness, it is common to aggregate many elementary reactions into elementary-complex reactions. One common method to do that is to reduce the deterministic description of the system using Michaelis-Menten approximations (singular perturbation), then use the resulting expressions to define the complex reactions and their propensities [12]. With this elementary-complex propensity-based reduced stochastic description, one can carry out a diffusion approximation that transforms the continuous-time discrete-state Markov process into a stochastic differential equation (SDE) called the Chemical Langevin Equation (CLE) [13]. Here, we only state the final result. If $a_i(Z)$, a and S are as defined before, then the deterministic description of the system (also know as the reaction-rate equations) is given by $\frac{dZ}{dt} = Sa(Z)$ and the SDE generated by the diffusion approximation is

$$dZ = Sa(Z)dt + S\sqrt{\mathrm{diag}(a(Z))}dw(t) \qquad (1)$$

III. Algorithms for Stochastic Analysis

A. Stochastic Reachability Analysis for the Discretized CME Using Linear Programming

To formulate the stochastic reachability analysis problem in a general setting, we consider a discrete Markov chain with L states, represented as the set $\mathscr{X} = \{1,\ldots,L\}$, and transition probability P, where the (i,j)-th entry $p_{i,j}$ represents the transition probability from state i to state j. We denote a realization of the stochastic process by $x[n]$, where $n = 0,1,2,\ldots$. In addition, we let $\mathscr{X}_0 \subseteq \mathscr{X}$ be the initial set, and $\mathscr{X}_u \subseteq \mathscr{X}$ be some "decisive/unsafe" set. For this system, we would like to estimate, or upper-bound the probability that a sample path $x[n]$ starting from $\mathscr{X}_0$ reaches $\mathscr{X}_u$ at some future time. In other words, we would like to find $\gamma \in [0,1]$ such that $P\{x[n] \in \mathscr{X}_u$ for some $n \geq 0 \mid x[0] \in \mathscr{X}_0\} \leq \gamma$. It is

clear that when the graph corresponding to the Markov chain is connected and there is no absorbing state (i.e., a state i such that $p_{i,i} = 1$) then such a probability will be equal to one. Thus, to get a meaningful probabilistic interpretation, some states should be made absorbing. In this case, this probability can be interpreted as the probability of reaching $\mathscr{X}_u$ before reaching the absorbing states.

To solve the above problem, we propose to find an appropriate barrier certificate $B(x)$ where $x \in \mathscr{X}$, from which we can deduce an upper bound γ. The idea is to find $B(x)$ such that for all $x[0] \in \mathscr{X}_0$ the process $B(x[n])$ is a nonnegative supermartingale, i.e., its expected value does not increase along time. The justification for this choice comes from a general result for discrete time nonnegative supermartingales [14] that we state without proof, and a theorem that provides the desired upper bound.

Lemma 1: [14] Let $B(x[n])$ be a nonnegative supermartingale. Then, for any initial condition $x[0] \in \mathscr{X}$,

$$P\left\{ \sup_{0 \leq n < \infty} B(x[n]) \geq \lambda \;\middle|\; x[0] = x_0 \right\} \leq \frac{B(x_0)}{\lambda}. \qquad (2)$$

Theorem 2: Let the discrete Markov chain P and the sets $\mathscr{X}_0$, $\mathscr{X}_u$ be given. Suppose there exist real scalars $B(1)$, $B(2)$, ..., $B(N)$, such that

$$B(x) \geq 0 \quad \forall x \in \mathscr{X}, \qquad (3)$$
$$B(x) \geq 1 \quad \forall x \in \mathscr{X}_u, \qquad (4)$$
$$B(x) \leq \gamma \quad \forall x \in \mathscr{X}_0, \qquad (5)$$
$$B(x) \geq \sum_{j=1}^{N} p_{x,j} B(j), \quad \forall x \in \mathscr{X} \qquad (6)$$

then $P\{x[n] \in \mathscr{X}_u$ for some $n \geq 0 \mid x[0] \in \mathscr{X}_0\} \leq \gamma$.

The proof of the theorem follows in the same manner as the proof of Theorem 7 in [17]. Note that γ can be chosen at most equal to one, since when $\gamma = 1$, the function $B(x) = 1$ will satisfy (3)–(6). The intuitive idea behind the theorem is as follows. The function $B(x)$ is a supermartingale of the process, and therefore its value is likely to stay constant or decrease as time increases. When we start from a lower initial value of $B(x)$ (i.e., as γ gets smaller) it becomes less likely for the trajectory to reach the unsafe set, on which the value of $B(x)$ is greater than or equal to one. This is quantified by Lemma 1, which provides a Chebyshev-like inequality for bounding the probability of the distribution tail.

With regard to computation, notice that the conditions (3)–(6) are nothing but a linear programming (LP) problem, where the decision variables are the scalars $B(.)$. Furthermore, γ can be chosen as the objective function whose value is to be minimized, to get an upper bound that is as tight as possible. Indeed, let $b = [\, B(1) \; \ldots \; B(L^2) \,]^T \in \mathbb{R}^{L^2}$ denote the barrier certificate, where L is the number of molecules of each species, and let P be the probability transition matrix. The reachability problem can then be written as

$$
\begin{aligned}
\text{minimize} \quad & \gamma \\
\text{subject to} \quad & b_{\mathrm{init}} \leq \gamma, \quad b_{\mathrm{unsafe}} \geq 1 \\
& b \geq 0, \quad Pb \leq b,
\end{aligned} \qquad (7)
$$

where b_{init} is the vector of values of the barrier in the initial set, b_{unsafe} are the values in the unsafe set, and γ is the probability bound.

B. Stochastic Reachability Analysis of SDEs

For our exposition in the SDE case, we consider a complete probability space $(\Omega, \mathscr{F}, P)$ and a standard $\mathbb{R}^m$-valued Wiener process w defined on this space [15]. We mostly focus on stochastic differential equations of the form

$$dx(t) = f(x(t))dt + g(x(t))dw(t) \tag{8}$$

where $x(t) \in \mathbb{R}^n$ for all $t \geq 0$. We again denote the state space, the initial set, and the decisive/unsafe set respectively by $\mathscr{X}, \mathscr{X}_0$, and $\mathscr{X}_u$, all of which are subsets of $\mathbb{R}^n$ and assumed compact. To guarantee the existence and uniqueness of solution, we will also assume that both f and g satisfy the local Lipschitz condition and the linear growth condition on $\mathscr{X}$. For bounded $\mathscr{X}$, the last condition can be replaced by the boundedness of f and g on $\mathscr{X}$.

Since in general the process $x(t)$ is not guaranteed to always lie inside the set $\mathscr{X}$, we define the stopped process corresponding to $x(t)$ and $\mathscr{X}$ as follows.

Definition 3: Suppose that τ is the first time of exit of $x(t)$ from the open set $\text{Int}(\mathscr{X})$. The stopped process $\tilde{x}(t)$ is defined by $\tilde{x}(t) \triangleq x(t)$ for $t < \tau$, and $\tilde{x}(t) \triangleq x(\tau)$ for $t \geq \tau$.

The stopped process $\tilde{x}(t)$ satisfies various properties. For example, it inherits the right continuity and strong Markovian property of $x(t)$. Furthermore, in most cases the infinitesimal generator corresponding to $\tilde{x}(t)$ is identical to the one corresponding to $x(t)$ on the set $\text{Int}(\mathscr{X})$, and is equal to zero outside of the set [16]. This will be implicitly assumed throughout the analysis.

Having defined the system and the stopped process $\tilde{x}(t)$, we state the continuous stochastic verification problem as follows: given the system (8) and the sets $\mathscr{X}, \mathscr{X}_0$ and $\mathscr{X}_u$, find $\gamma \in [0, 1]$ such that $P\{\tilde{x}(t) \in \mathscr{X}_u \text{ for some } t \geq 0 | \tilde{x}(0)\} \leq \gamma$ for all $\tilde{x}(0) \in \mathscr{X}_0$. Again, our approach will be based on finding an appropriate barrier certificate $B(x)$ from which an upper bound γ is deduced.

Theorem 4 ([17]): Let the SDE (8) and the sets $\mathscr{X}, \mathscr{X}_0, \mathscr{X}_u$ be given, and consider the stopped process $\tilde{x}(t)$ starting at some $\tilde{x}(0) \in \mathscr{X}_0$. Suppose there exists a twice continuously differentiable function $B : \mathscr{X} \rightarrow \mathbb{R}$, such that

$$B(x) \geq 0 \ \forall \ x \in \mathscr{X} \tag{9}$$

$$B(x) \geq 1 \ \forall \ x \in \mathscr{X}_u \tag{10}$$

$$B(x) \leq \gamma \ \forall \ x \in \mathscr{X}_0 \tag{11}$$

$$\frac{\partial B}{\partial x} f(x) + \frac{1}{2} \text{Trace} \left(g^T(x) \frac{\partial^2 B}{\partial x^2} g(x) \right) \leq 0 \ \forall \ x \in \mathscr{X} \tag{12}$$

then $P\{\tilde{x}(t) \in \mathscr{X}_u \text{ for some } t \geq 0 | \tilde{x}(0)\} \leq \gamma$.

The proof of Theorem 4 can be found in [17]. An upper bound γ and a barrier certificate $B(x)$ which certifies the upper bound can be computed by formulating conditions (9)–(12) as a sum of squares optimization problem (see [17]). This problem can be solved using the software SOSTOOLS [18]. T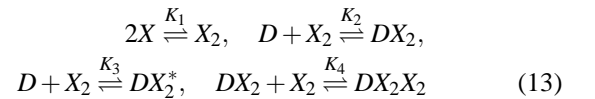he minimum value of γ obtained from the optimization will be the tightest upper bound for a given set of barrier certificates. Obviously, as we include more candidate barrier certificates in this set, we may get a better bound, although there is a trade-off between using a larger set and the computational complexity of finding a barrier certificate within the set.

IV. THE BACTERIOPHAGE λ GENETIC SWITCH

One of the best studied examples of multistability in genetic systems is the bacteriophage λ system [3]. Phages are viral organisms that can either be in the lysogenic (latent) or lytic (active) state. If following its infection of *E. coli* the λ-phage virus enters the lysogenic pathway, it represses its own developmental functions and integrates its DNA into the host chromosome. Otherwise, it enters the lytic pathway and is active. The dynamics of the transcriptional network underlying the formation of these states are very complex, and have been thoroughly studied. Instead of describing it in full complexity, we limit ourselves here to a simplified model that still captures the essence of its bistability, in addition to the interplay between noise and dynamics in its operation.

A. The Bacteriophage λ Model

The simple model we adopt was first discussed in [8], and is intended to be a minimal model that captures the bistable nature of the λ system. The model describes the dynamics of the *CI* protein, the product of the *cI* gene. Acting as a dimer, *CI* can regulate its own synthesis by binding to the *cI* gene promoter region OR_2 and increasing transcription, or to gene promoter OR_3 and repressing transcription. It can also bind to a third gene promoter OR_1, which is ignored in this model for simplicity. The fast binding reactions (e.g., binding and dissociation) are assumed to be in equilibrium with respect to the slow reactions (e.g., protein synthesis and degradation). Under this assumption, and following [8], we let X, X_2, and D denote the repressor, repressor dimer, and DNA promoter site respectively. The equilibrium reactions can then be written as

$$2X \overset{K_1}{\rightleftharpoons} X_2, \quad D + X_2 \overset{K_2}{\rightleftharpoons} DX_2,$$
$$D + X_2 \overset{K_3}{\rightleftharpoons} DX_2^*, \quad DX_2 + X_2 \overset{K_4}{\rightleftharpoons} DX_2X_2 \tag{13}$$

where the DX_2 and DX_2^* complexes denote binding to the OR_2 or OR_3 sites, respectively, DX_2X_2 denotes binding to both sites, and the K_i are forward equilibrium constants. We further set $K_3 \simeq K_2$ and $K_4 \simeq 5K_2$.

1) The One-Dimensional Model: In addition to the fast reactions given above, the slow reactions are production and degradation of proteins given by the reactions

$$DX_2 + P \overset{k_t}{\longrightarrow} DX_2 + P + nX, \quad X \overset{k_d}{\longrightarrow} A \tag{14}$$
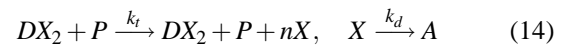
where P denotes the concentration of RNA polymerase, and n is the number of proteins per mRNA transcript. To model the system, we state variables as $x = [X], y = [X_2], d = [D], u = [DX_2], v = [DX_2^*]$, and $z = [DX_2X_2]$, where $[.]$ denotes concentration. The evolution of the concentration of the repressor x can then be described by $\dot{x} = -2k_1x^2 + 2k_{-1}y +$
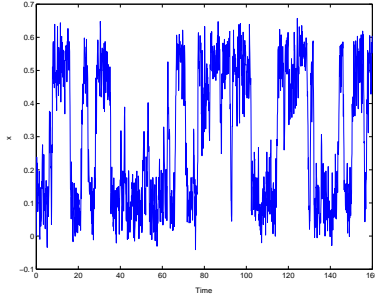
Fig. 1. Switching between equilibria in the λ system caused by noise.

Fig. 2. Transition probabilities between adjacent lattice points.

$nk_t p_0 u - k_d x + q$, where q is a basal rate of production. Because the reactions in (13) are fast, y, u, and d have algebraic expressions in terms of x. After eliminating their dynamics through singular perturbation and rescaling the repressor concentration x and time through a change of variables [8], the resulting one-dimensional model for x (with some abuse of notation) is then given by

$$\dot{x} = \frac{\alpha x^2}{1 + 2x^2 + 5x^4} - \gamma x + 1 \tag{15}$$

We will be concerned with the case where the system has three equilibria: two stable and one unstable. In this bistable case, the deterministic system trajectories will settle in one of the stable equilibria at steady-state depending on initial condition. However, the presence of biochemical noise can cause accidental switching between the two steady-states. One possible behavior obtained using the SSA is depicted in Figure 1. Later we will be concerned with the characterization of the statistics of such switching events.

2) The Two-Dimensional Model: The model for the slow reactions presented above is particularly simple. Next, we present a 2-D description of the λ system while giving a more plausible description of the biochemical noise affecting the system. We start by introducing a crucial step that was eliminated from the previous model: translation. Recall that in the previous model, X is produced in a single step process, then degraded. More realistically, it should be produced through a two step process of transcription first, then translation. The resulting 2-D model is $\dot{m} = k_t p_0 u - k_d m + q$ and $\dot{x} = nk_T m - k_m x + 2k_{-1}y - 2k_1 x^2$, where m is mRNA, and x is protein. Following [8], a change of variables is carried out and the values of the parameters are fixed to obtain:

$$\dot{\hat{m}} = \frac{50\hat{x}^2}{1 + 2\tilde{x}^2 + 5x^4} - 150\hat{m} + 1 \tag{16}$$
$$\dot{\hat{x}} = 10\hat{m} - \hat{x} \tag{17}$$

This higher order system, with the chosen parameter values, also exhibits bistability.

B. Numerical Results for the Discretized CME 2-D Model

Here we consider the λ-phage example described in the CME formulation. Rather than accounting for all of the elementary reactions occurring in the system (i.e. (13) and(14)), we use elementary complex reactions, leading to $D \xrightarrow{1+b(x)} m + D$, $m \xrightarrow{\gamma_1} \phi$, $m \xrightarrow{k} m + X$, and $x \xrightarrow{\gamma_2} \phi$, where
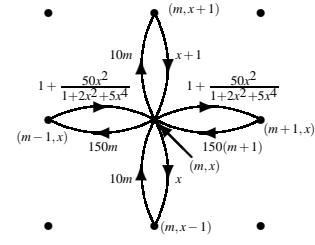
the transition probabilities $b(x) = \frac{\alpha x^2}{1 + 2x^2 + 5x^4}$, $\gamma_1 m$, km, and $\gamma_2 x$ are calculated from (16)–(17). Our state space is a 2-dimensional lattice, part of which is shown in Figure 2.

Computing probability of reaching some states on this lattice from others can be expressed as the LP (7). Thus we can address questions such as: what is the probability of reaching the unsafe state before reaching given intermediate states? This is done by making the given intermediate states "absorbing", which is the same as removing the constraint corresponding to that node from $Pb \leq b$. We solve the LP for a 3×3 lattice, and obtain the bound $\gamma = 0.0063$, which corresponds to the probability of going from $\mathscr{X}_0$ to $\mathscr{X}_u$ before first going to states in $\mathscr{X}_i$. The problem setup and resulting barrier values are shown in Figure 3.

The complexity of this method scales exponentially with the number of species, therefore it is only useful for small problems. Moreover, certain instances of this LP are fragile (ill-conditioned), which motivates us to look at stochastic descriptions of the system at lower levels of resolution.

C. Numerical Results for the SDE Model

1) SDE Analysis of the 1-D Model: We start by incorporating additive noise in the deterministic rate equation of (15). Taking the variable x to represent the repressor number within a colony of cells, we can regard an additive white noise term as an external field accounting for the impact of the the environmental noise. The resulting SDE is

$$dx(t) = f(x)dt + \sigma d\xi(t) \tag{18}$$

where $f(x)$ is the right-hand side of (15), $\xi(t)$ is a Wiener process, and σ is a scalar. While in the deterministic description, trajectories that reach either steady state come to rest there, in the presence of noise they never do so. One can then ask whether there is a large probability of noise-induced switching between equilibria within a given time frame, then describe its correlation with the noise intensity. We use the
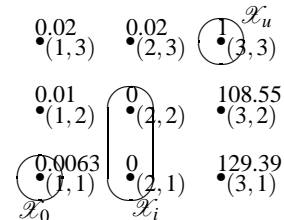

Fig. 3. Values of the barrier certificate.

$$dx = \left(\frac{50x^2}{1+2x^2+5x^4} - 15x + 1 \right) dt + d\xi(t)$$

Equilibria:
$x_1^* = 0.560$ (stable)
$x_2^* = 0.098$ (stable)
$x_3^* = 0.255$ (unstable)

$X_u = \{ x \in R \mid (x-0.098)^2 - 0.02^2 \le 0 \}$
$X_0 = \{ x \in R \mid (x-0.56)^2 - r^2 \le 0 \}$
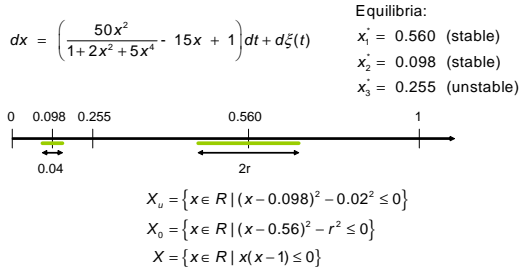$X = \{ x \in R \mid x(x-1) \le 0 \}$

Fig. 4. Setup for computing the bound of the probability of false switching because of noise, for system (18).

sum of squares machinery as an efficient algorithmic method for such investigations. We specifically illustrate this point by investigating the situation where the system given by (15) is started close to one of the two stable equilibria, and an estimate of the the probability of transition to a region around the other stable equilibrium is computed. We will refer to this as *the bound on the probability of false switching caused by noise*. The setup for this problem is shown in Figure 4, where we are interested in whether an initial trajectory from a point inside a set of variable size r centered at the high equilibrium $\mathscr{X}_0 = \{ x \in \mathbb{R} \mid (x - 0.56)^2 - r^2 \le 0 \}$ can ever reach a region around the other equilibrium, $\mathscr{X}_u = \{ x \in \mathbb{R} \mid (x - 0.098)^2 - 0.02^2 \le 0 \}$, when the total state-space of interest is $\mathscr{X} = \{ x \in \mathbb{R} \mid x(x-1) \le 0 \}$.

This question, however, is ill-posed if the time horizon over which this probability is to be computed is not bounded. A more meaningful biological question is the following. Starting from a region around one equilibrium, estimate the probability of reaching the other equilibrium *in a finite time horizon*, say from $t = 0$ to 2 non-dimensional time units. Time is now another variable in the system, and we construct a time-dependent $B(x,t)$ to estimate this probability. We also increase the state-space adequately, so that the "escape" probability is reduced significantly. As σ is varied in this modified problem, results from our methodology are compared with probability estimates obtained by direct simulations of the SDE in (18). When $\sigma = 1$, 98% of the simulations enter the region around the other equilibrium, whereas our methodology gives an upper bound of $\beta = 0.99$. When $\sigma = 0.5$, direct simulations gives 0.57, and our methodology returns $\beta = 0.6$. We see the expected result that as σ is decreased, the probability of false switching is decreased, and the upper bounds are close to the probability estimates obtained by direct simulations.

Using the same methodology, we can also address the dependence of the switching probability from one steady-state to the other on time and for a given noise intensity. The upper bound on this probability is plotted in Figure 5 as a function of time, along with probability estimates obtained from simulations. It can be seen that the two plots are satisfactorily close to each other. However, while a large number of simulations is needed to compute the probability estimates, with an increasing computational cost as the simulation time interval was increased, the barrier certificate method generated good upper bounds in a fraction of time
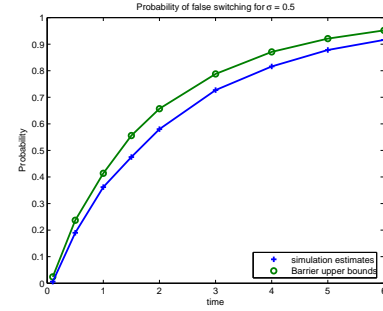

Fig. 5. Probability of false switching from the high steady state to the low steady state under the influence of noise as a function of time. The blue line depicts probability bounds computed through the use of barrier certificates while the green line depicts the probabilities computed from 2000 realizations obtained from direct simulations of the SDE.

and each run had the same computational burden. In general, probabilities computed using barrier certificates are conservative bounds on the true probabilities. Nevertheless, as illustrated here, these bounds can provide useful information.

The same techniques can again be used to compute probability lower bounds. For example, Figure 6 shows the upper and lower bounds for the probability of escaping a neighborhood of the high equilibrium in a finite time horizon as the noise intensity increases. It is interesting to note that the probability goes sharply from 0 to 1 (indicating a phase transition) as the noise is increased, and that this technique can be used to estimate the critical noise intensity necessary for escaping the given region around the equilibrium.

2) SDE Analysis of the 2-D Model: In the phase plane, the deterministic part of the 2-D λ-model is shown in Figure 7. Rather than introducing additive noise, we consider here more biologically plausible biochemical fluctuations. As described previously, we replace the elementary reactions by complex birth and death reactions for $\tilde{m}$ and $\tilde{x}$ whose propensities are taken from (16)–(17). With these elementary-complex propensities generating a reduced order stochastic description of the λ-phage system, we carry out a diffusion approximation that transforms the description of the system into a stochastic differential equation (i.e. a CLE). Following (1), the final CLE describing the 2-D λ system becomes $d\hat{m} = \left(\frac{50\hat{x}^2}{1+2\hat{x}^2+5\hat{x}^4} - 150\hat{m} + 1 \right) dt +$
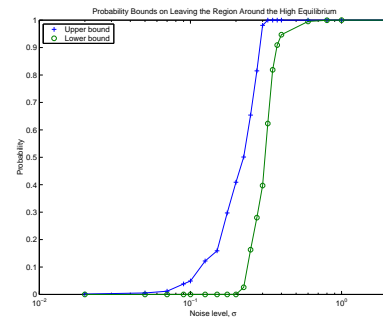

Fig. 6. Upper and lower bounds on probability of escaping a neighborhood of the high equilibrium within some finite time interval, computed using barrier certificates.
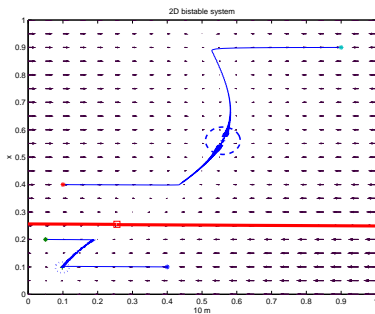
Fig. 7. Phase plane for system (16)–(17). Equilibria are shown by '+' (stable) and '□' (unstable). Solid lines are trajectories from initial conditions denoted by '*'. The solid thick line is a separatix dividing the phase plane in two, so that if the deterministic system is initialized in one region then its trajectory flows towards one equilibrium, whereas in the other region the trajectory flows to the other equilibrium.

$$\sqrt{\frac{50\hat{x}^2}{q(1+2\hat{x}^2+5\hat{x}^4)}}d\xi_1 + \sqrt{\frac{150\hat{m}}{q}}d\xi_2 + \frac{d\xi_3}{\sqrt{q}}, \quad \text{and} \quad d\hat{x} = (10\hat{m}-\hat{x})dt + \sqrt{\frac{10\hat{m}}{q}}d\xi_4 + \sqrt{\frac{\hat{x}}{q}}d\xi_5.$$

Here again, we consider whether from an initial region around the low equilibrium defined by $\mathscr{X}_0 = \{x \in \mathbb{R} | (x - 0.098)^2 + (10m - 0.098)^2 - 0.02^2 \leq 0\}$ one can reach a region around the high equilibrium defined by $\mathscr{X}_u = \{x \in \mathbb{R} | (x - 0.560)^2 + (10m - 0.560)^2 - 0.05^2 \leq 0\}$, in a certain time span, say $t \in [0,10]$ for different noise intensities. In this case, and for a fixed structure of the system, different noise intensities can be achieved by changing the number of molecules in the system. We therefore define $\hat{x} = Nx$ for equation (1). N is commonly called the system size. This change of variables scales the number of molecules without modifying the dynamics of the deterministic system. However, the associated noise intensity grows as the square root of N. With N as our varying noise intensity, we attempt to answer the false switching question. We construct certificates $B(x,t)$ algorithmically, again using the methodology described earlier. Note that the non-polynomial nature of the vector field (the way noise affects the system is through the square root of some polynomial) does not cause any problems, as these terms appear *squared* in the expressions to be tested — see condition (12) in Theorem 4. The upper bound on this probability as function of N is given in Figure 8. Notice the expected fact that as N increase, therefore yielding lower noise affecting the system, the probability of false switching decreases. As in the previous case, this result was obtained without time consuming computations and was not based on numerous simulations of the corresponding SDE.

V. CONCLUSIONS AND FUTURE WORK

In this paper, we have illustrated how the sum of squares technique can be used to tackle algorithmically a number of issues related to the stochastic reachability analysis of nonlinear systems. In particular, we used barrier certificates to compute certified upper bounds on the probability of reaching a region around one stable steady-state from a region around the other in a specified time interval, and as a function of the noise intensity.

While the presented example illustrates a successful use
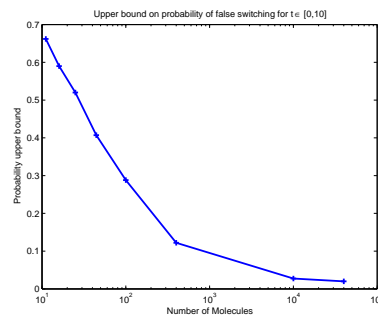

Fig. 8. Upper bound on switching probability from the higher equilibrium as a function of the number of molecules in the system

of promising new methods to study the problem, it also uncovered research challenges that must be overcome to broaden the applicability of the methods and investigate their limitations. A major challenge that must addressed is the issue of scalability. For example, in the discretized CME description, the complexity scales exponentially with the number of species in the system. Furthermore, the resulting LPs can be ill-conditioned. These issues are currently the subject of an active research effort.

REFERENCES

[1] M. A. Dingemanse *et al*, The expression of liver specific genes within rat embryonic hepatocytes is a discontinuous process, *Differentiation*, vol. 56, 1994, 153-162.

[2] D. C. Bennett, Differentiation in Mouse Melanoma Cells: Initial Reversibility and an On-Off Stochastic Model, *Cell*, Vol. 34, 1983, 445-453.

[3] A. Arkin, J. Ross, and H. H. McAdams, Stochastic Kinetic Analysis of the Developmental Pathway Bifurcation in Phase λ-Infected *Escehrichia coli* Cells, *Genetics*, vol. 149, 1998, 1633-1648.

[4] H. El-Samad, M. Khammash, L. Petzold, and D. Gillespie, Stochastic Modeling of Gene Regulatory Networks, *Int. Journal of Nonlinear and Robust Control*, 2005, To Appear.

[5] F. Benz, Cell Death and Cancer Therapy, *Current Opinion in Pharmacology*, vol. 1, 2001, pp. 337-341.

[6] O. Kobiler *et al*, Quantitative Analysis of bacteriphage λ Genetic Network, *Proc. Nat. Academy of Sci.*, vol. 102, 2005, pp. 4470-4475.

[7] M. Ptashne *A Genetic Switch: Phage λ and Higher Organisms*. Cell Press and Blackwell Scientific Publications, Cambridge, MA, 1992.

[8] J. Hasty *et al*, Noise-Based Switches and Amplifiers for Gene Expression. *Proc. Nat. Academy of Sci.*, vol. 97 , 2000, 2075-2080.

[9] N.G. van Kampen, *Stochastic Processes in Physics and Chemistry*, Elsevier Science Publishing Company, Berlin, 1992.

[10] D. Gillespie, A Rigorous Derivation of the Chemical Master Equation, *Physica A*, vol. 188, 1992, 404-425.

[11] D. T. Gillespie, Exact Stochastic Simulation of Coupled Chemical Reactions, *J. Phys. Chem.*, vol. 81, 1977, pp. 2340-2361.

[12] C. Rao and A.P. Arkin, Stochastic Chemical Kinetics and the Quasi-Steady-State Assumption: Application to the Gillespie Algorithm, *J. Chem. Phys.*, vol. 118, 2003, pp. 4999-5010.

[13] D.T. Gillespie, The chemical Langevin Equation, *J. Chem. Phys*, vol. 13, 2000, pp.297–306.

[14] G. A. Edgar and L. Sucheston, *Stopping Times and Directed Processes*, Cambridge University Press, 1992.

[15] B. Øksendal, *Stochastic Differential Equations: An Introduction with Applications*, Springer-Verlag, Berlin, 2000.

[16] H.J. Kushner, *Stochastic Stability and Control*, Academic Press, New York, 1967.

[17] S. Prajna, A. Jadbabaie, and G. J. Pappas, Stochastic Safety Verification Using Barrier Certificates, *Proc. IEEE Conf. on Decision and Control*, 2004.

[18] S. Prajna, A. Papachristodoulou, P. J. Seiler, and P. A. Parrilo, SOSTOOLS: Sum of squares optimization toolbox for MATLAB, Available at http://www.cds.caltech.edu/sostools and http://www.mit.edu/~parrilo/sostools, 2004.