2013 European Control Conference (ECC)
July 17-19, 2013, Zürich, Switzerland.

Analysis of Controlled Biological Switches via Stochastic Motion Planning

Peyman Mohajerin Esfahani, Andreas Milias-Argeitis, and Debasish Chatterjee

Abstract— **We consider the problem of regulating a genetic toggle switch by means of an in silico feedback control loop. To achieve this, we first introduce two basic notions of motion planning, and then establish a connection to a class of stochastic optimal control problems concerned with sequential stopping-times. To characterize the desired set of initial conditions, in the context of controlled diffusion processes, we propose a sequence of partial differential equations for which the first one has a known boundary condition, while the boundary conditions of the subsequent ones are determined by the solutions to the preceding steps. We then formulate the control of a bistable system as stochastic motion planning problem, and derive the closed-loop control law that maintains the system inside a prespecified region of its state space. Finally, to provide an autonomous feedback policy, we establish a connection to an eigenvalue problem that describes the asymptotic exit-time of the diffusion process.**

I. Introduction and Problem Statement

The advances in single-cell experimental techniques during the last decade have revealed that bistability underlies many biological processes related to different aspects of cellular decision-making. A bistable system can be found in two distinct and mutually exclusive states, while being able to switch from one state to the other under the influence of a transient external signal. Due to the presence of stochastic fluctuations in cellular components, biological bistable systems can also switch states randomly under the influence of molecular noise. Since the state of each individual cell is randomly determined, such systems can give rise to phenotypic heterogeneity within isogenic cell populations [1].

Control of cellular behavior has recently gained popularity as an approach to generating cells with prescribed target functions but, most importantly, as a means to better understanding cellular processes that are partially known [2]. The technique of *in silico* control of biological systems [3], [4], [5] has been proposed as a complement to synthetic feedback control schemes implemented inside cells, and makes possible the precise and fast manipulation of intracellular states to achieve various control objectives. In this work we examine the applicability of *in silico* feedback on single cells for the control of a small bistable biological system with stochastic dynamics. Control of naturally occurring bistable systems can provide useful insights into their functional organization, as well as enable the generation and study of completely different phenotypes, by driving the cells away from the commonly observed states.

We consider a toy system model inspired by the first synthetically engineered two-gene toggle switch [6], which can be thought of as an abstraction of several cellular (and more complex) bistable systems. Assuming that the molecular populations involved are relatively small, our system is governed by stochastic dynamics that enable it to flip randomly between the basins of attraction of the stable equilibria of the deterministic equations. Our feedback control objective is to keep the system for a given amount of time away from both of these equilibria, in a region that contains the unstable deterministic equilibrium. To achieve this, we use tools from the theory of stochastic motion planning for systems described by stochastic differential equations, recently introduced in [7], and described in more detail in Section II. After the presentation of the main theoretical tools, we evaluate the performance of our controllers on the toggle switch model in Section III, and discuss the strengths and weaknesses of our approach.

Research supported by the European Commission under the project MoVeS (Grant Number 257005) and the HYCON2 Network of Excellence (FP7-ICT-2009-5).

PME and AM-A are with the Automatic Control Laboratory, ETH Zürich, 8092 Zürich, Switzerland; DC is with the Systems & Control Engineering, IIT-Bombay, Powai, Mumbai 400076, India. Emails: {mohajerin,milias}@control.ee.ethz.ch, chatterjee@sc.iitb.ac.in

II. Stochastic Motion Planning

The aim of this section is to introduce different stochastic motion planning problems, which involve a controlled process visiting certain subsets of the state-space while avoiding others in a sequential fashion. The main objective is to determine the set of initial conditions for which there exists an admissible policy to execute the desired maneuver with probability at least as much as some pre-specified value. To this end, we first establish a connection from this initial condition set to a class of stochastic optimal control problems. In the following, we propose a PDE characterization of the corresponding value functions, which allows us to invoke existing PDE solvers to numerically compute the desired initial sets. Here we skip all the proofs and refer interested readers to our recent work [7] for detailed analysis and more generalized settings.

A. General Setting and Definitions

Consider a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ whose filtration $\mathbb{F} := (\mathcal{F}_s)_{s\geq 0}$ is generated by a d-dimensional Brownian motion $(W_s)_{s\geq 0}$ and enlarged by its right-continuous $\mathbb{P}$-completion; – the usual conditions of completeness and right continuity [8, p. 48]. Let $\mathbb{U} \subset \mathbb{R}^m$ be a control set, and let $\mathcal{U}_t$ denote the subset of $\mathbb{F}$-progressively measurable maps into $\mathbb{U}$ that is independent of the Brownian motion up to time t ($W_{[0,t]}$). The basic object of our study concerns the $\mathbb{R}^n$-valued stochastic differential equation (SDE)

$$\mathrm{d}X_s = f(X_s, u_s)\,\mathrm{d}s + \sigma(X_s, u_s)\,\mathrm{d}W_s, \quad s \geq t, \quad (1)$$

where $X_t = x$ given, $f : \mathbb{R}^n \times \mathbb{U} \longrightarrow \mathbb{R}^n$ and $\sigma : \mathbb{R}^n \times \mathbb{U} \longrightarrow \mathbb{R}^{n\times d}$ are measurable maps, and $\boldsymbol{u} := (u_s)_{s\geq 0} \in \mathcal{U}_t$.

It is known that under some mild assumptions (Lipschitz continuity of f and σ) there exists a unique strong solution to SDE (1). Let us denote it by $\left(X_s^{t,x;\boldsymbol{u}}\right)_{s\geq t}$ [9].

Given sets $(W_i, G_i) \in \mathfrak{B}(\mathbb{R}^d) \times \mathfrak{B}(\mathbb{R}^d)$ for $i \in \{1, \cdots, n\}$, we are interested in a set of initial conditions (t, x) in $\mathbb{S} := [0, T] \times \mathbb{R}^d$ such that there exists an admissible strategy $\boldsymbol{u} \in \mathcal{U}_t$ steering the process $X^{t,x;\boldsymbol{u}}$ through the sets $(W_i)_{i=1}^n$ while visiting $(G_i)_{i=1}^n$ in a pre-assigned order. In fact, W_i and G_i stand for *"Way"* and *"Goal"* respectively. One may pose this objective from different perspectives

978-3-952-41734-8/©2013 EUCA

based on different time scheduling for the excursions between the sets. We formally introduce some of these notions which will be addressed throughout this article.

Definition 2.1 (Motion-Planning Events): Consider a fixed initial condition $(t,x) \in \mathbb{S}$ and admissible policy $\boldsymbol{u} \in \mathcal{U}_t$. Given a sequence of pairs $(W_i, G_i)_{i=1}^n \subset \mathfrak{B}(\mathbb{R}^d) \times \mathfrak{B}(\mathbb{R}^d)$ and horizon times $(T_i)_{i=1}^n \subset [t,T]$, we introduce the following **motion-planning events**:

$$\Big\{X_\cdot^{t,x;\boldsymbol{u}} \models \big[(W_1 \rightsquigarrow G_1) \circ \cdots \circ (W_n \rightsquigarrow G_n)\big]_{\le T}\Big\} := \tag{2a}$$
$$\Big\{\exists (s_i)_{i=1}^n \subset [t,T] \;\big|\; X_{s_i}^{t,x;\boldsymbol{u}} \in G_i \text{ and }$$
$$X_r^{t,x;\boldsymbol{u}} \in W_i \setminus G_i,\ \forall r \in [s_{i-1}, s_i[, \quad \forall i \le n\Big\},$$

$$\Big\{X_\cdot^{t,x;\boldsymbol{u}} \models (W_1 \xrightarrow{T_1} G_1) \circ \cdots \circ (W_n \xrightarrow{T_n} G_n)\Big\} := \tag{2b}$$
$$\Big\{X_{T_i}^{t,x;\boldsymbol{u}} \in G_i \text{ and } X_r^{t,x;\boldsymbol{u}} \in W_i, \quad \forall r \in [T_{i-1}, T_i], \quad \forall i \le n\Big\},$$

where in the above definitions $s_0 = T_0 := t$.

The set in (2a), roughly speaking, contains the events in which the trajectory $X_\cdot^{t,x;\boldsymbol{u}}$, initialized at $(t,x) \in \mathbb{S}$ and controlled via $\boldsymbol{u} \in \mathcal{U}_t$, succeeds in visiting the sets $(G_i)_{i=1}^n$ in a certain order, while the entire duration between the two visits to G_{i-1} and G_i is spent in W_i, all within the time horizon T. In other words, the journey from G_{i-1} to the next destination G_i must belong to the way W_i for all i. In the case of (2b), the set of paths is more restricted in comparison to (2a). Indeed, not only is the trajectory confined to the ways W_i, but also there is a time schedule $(T_i)_{i=1}^n$ that a priori forces the process to be at the goal sets G_i at the specific times $(T_i)_{i=1}^n$.

From the technical standpoint, if the target set G_i is not closed, then it is not difficult to see that there could be some continuous transitions through the boundary of the goal G_i that are not admissible in view of the definition (2a) since the trajectory must reside in $W_i \setminus G_i$ for the whole interval $[s_{i-1}, s_i[$ and just hit the set G_i at the time s_i. Notice that we do not need to consider this issue for the set in definition (2b) since in this case the trajectory only visits the sets G_i at the specific times T_i while any continuous transition and maneuver inside the target sets G_i are allowed. In order to address the aforementioned issue, we may impose the sets $(G_i)_{i=1}^n \subset \mathfrak{B}(\mathbb{R}^d)$ are all closed.

The events introduced in Definition 2.1 depend, of course, on the control policy $\boldsymbol{u} \in \mathcal{U}_t$ and initial condition $(t,x) \in \mathbb{S}$. The central objective of this work is to determine the set of initial conditions $x \in \mathbb{R}^d$ for which there exists an admissible policy $\boldsymbol{u}$ such that the probability of the above path-planning events is higher than a certain threshold. To this end, we formally introduce these sets as follows:

Definition 2.2 (Motion-Planning Initial Set): Consider a fixed initial time $t \in [0,T]$. Given a sequence of set pairs $(W_i, G_i)_{i=1}^n \subset \mathfrak{B}(\mathbb{R}^d) \times \mathfrak{B}(\mathbb{R}^d)$ and horizon times $(T_i)_{i=1}^n \subset [t,T]$, we define the following **motion-planning initial sets**:

$$\mathrm{PP}\big(t,p;\ (W_i,G_i)_{i=1}^n, T\big) := \Big\{x \in \mathbb{R}^d \;\big|\; \exists \boldsymbol{u} \in \mathcal{U}_t : \tag{3a}$$
$$\mathbb{P}\big\{X_\cdot^{t,x;\boldsymbol{u}} \models \big[(W_1 \rightsquigarrow G_1) \circ \cdots \circ (W_n \rightsquigarrow G_n)\big]_{\le T}\big\} > p\Big\},$$
$$\widetilde{\mathrm{PP}}\big(t,p;\ (W_i,G_i)_{i=1}^n, (T_i)_{i=1}^n\big) := \Big\{x \in \mathbb{R}^d \;\big|\; \exists \boldsymbol{u} \in \mathcal{U}_t : \tag{3b}$$
$$\mathbb{P}\big\{X_\cdot^{t,x;\boldsymbol{u}} \models (W_1 \xrightarrow{T_1} G_1) \circ \cdots \circ (W_n \xrightarrow{T_n} G_n)\big\} > p\Big\}.$$

B. Connection to Stochastic Optimal Control Problems

In this subsection we establish a connection from stochastic motion-planning initial sets PP and $\widetilde{\mathrm{PP}}$, defined in Definition 2.2, and a class of stochastic optimal control problems involving stopping times. For this purpose let us introduce a sequence of random times that corresponds to the times that the process $X_\cdot^{t,x;\boldsymbol{u}}$ for the first time exits from the sequence of sets one after another in a certain order:

Definition 2.3: Given an initial condition $(t,x) \in \mathbb{S}$ and a sequence of measurable sets $(A_i)_{i=k}^n \subset \mathfrak{B}(\mathbb{R}^d)$, the sequence of random times $\big(\Theta_i^{A_{k:n}}\big)_{i=k}^n$ is called the **sequential exit-time** through the set A_k to A_n, and defined[1] by

$$\Theta_i^{A_{k:n}}(t,x) := \inf\big\{r \ge \Theta_{i-1}^{A_{k:n}}(t,x) \;:\; X_r^{t,x;\boldsymbol{u}} \notin A_i\big\},$$

where the initial random time is $\Theta_{k-1}^{A_{k:n}}(t,x) := t$.

Note that the sequential exit-time $\Theta_i^{A_{k:n}}$ depends on the control policy $\boldsymbol{u}$ in addition to the initial condition (t,x), but here and later in the sequel we shall suppress this dependence. For notational simplicity, we also drop (t,x) in the subsequent sections. Consider the value functions $V, \widetilde{V} : \mathbb{S} \to [0,1]$ as follows:

$$V(t,x) := \sup_{\boldsymbol{u} \in \mathcal{U}_t} \mathbb{E}\left[\prod_{i=1}^n \mathbb{1}_{G_i}\big(X_{\eta_i}^{t,x;\boldsymbol{u}}\big)\right], \tag{4a}$$
$$\widetilde{V}(t,x) := \sup_{\boldsymbol{u} \in \mathcal{U}_t} \mathbb{E}\left[\prod_{i=1}^n \mathbb{1}_{G_i \cap W_i}\big(X_{\widetilde{\eta}_i}^{t,x;\boldsymbol{u}}\big)\right], \tag{4b}$$

with the stopping times $\eta_i := \Theta_i^{B_{1:n}} \wedge T$, $B_i := W_i \setminus G_i$, $\widetilde{\eta}_i := \Theta_i^{W_{1:n}} \wedge T_i$, where $\Theta_i^{W_{1:n}}, \Theta_i^{B_{1:n}}$ are the sequential exit-times in the sense of Definition 2.3, and $\wedge$ is the minimum operator. The main result of this subsection, Proposition 2.4 below, establishes a connection from the sets $\mathrm{PP}, \widetilde{\mathrm{PP}}$ and superlevel sets of the value functions V and $\widetilde{V}$.

Proposition 2.4: Fix a probability level $p \in [0,1]$, a sequence of set pairs $(W_i, G_i)_{i=1}^n \subset \mathfrak{B}(\mathbb{R}^d) \times \mathfrak{B}(\mathbb{R}^d)$, an initial time $t \in [0,T]$, and intermediate times $(T_i)_{i=1}^n \subset [t,T]$. Then

$$\mathrm{PP}\big(t,p;(W_i,G_i)_{i=0}^n, T\big) = \big\{x \in \mathbb{R}^d \;\big|\; V(t,x) > p\big\}. \tag{5}$$

Moreover, suppose $(W_i)_{i=1}^n \subset \mathfrak{B}(\mathbb{R}^d)$ are open. Then,

$$\widetilde{\mathrm{PP}}\big(t,p;(W_i,G_i)_{i=0}^n, (T_i)_{i=1}^n\big) = \big\{x \in \mathbb{R}^d \;\big|\; \widetilde{V}(t,x) > p\big\}, \tag{6}$$

where the value functions V and $\widetilde{V}$ are as defined in (4).

Remark 2.5 (Mixed Motion-Planning Events): In this article we focus on two sets of events as introduced in Definition 2.1, however, it might be of interest to consider an event that consists of a mixture of events in (2), e.g., $(W_1 \rightsquigarrow G_1)_{\le T_1} \circ (W_2 \xrightarrow{T_2} G_2)$. One can observe that essentially the same analytical techniques as the ones proposed here can be employed to address these mixed motion planning objectives, and establish a connection to a class of optimal control problems with some appropriate sequential stopping times. We shall provide an example of this nature in Section III.

C. PDE Characterization

In the preceding subsections a link from the stochastic motion-planning initial sets, Definition 2.2, and the superlevel sets of the value functions (4) was established. In this part, under some technical assumptions, we propose a PDE characterization of the value function (4) so as to numerically compute the desired initial sets PP and $\widetilde{\mathrm{PP}}$. The approach

[1] By convention, $\inf \emptyset = \infty$.

results in a series of Hamilton-Jacobi-Bellman PDE's, where each PDE is understood in the discontinuous viscosity sense with some boundary conditions; for the general theory of viscosity solutions we refer to [10] and [11]. To this end we proceed with a more abstract setting for which both value functions (4) can be addressed.

Let $(T_i)_{i=1}^n \subset [0,T]$ be a sequence of times, $(A_i)_{i=1}^n \subset \mathfrak{B}(\mathbb{R}^n)$ be a family of open sets, and payoff functions $\ell_i : \mathbb{R}^r \to \mathbb{R}$ that are measurable and bounded, $i = 1, \ldots, n$. We define the sequence of value functions $V_k : [0,T] \times \mathbb{R}^d \to \mathbb{R}^d$, $k = 1, \ldots, n$,

$$V_k(t,x) := \sup_{\boldsymbol{u} \in \mathcal{U}_t} \mathbb{E}\Big[\prod_{i=k}^{n} \ell_i\big(X_{\tau_i^k}^{t,x;\boldsymbol{u}}\big)\Big], \tag{7}$$

where $\tau_i^k(t,x) := \Theta_i^{A_{k:n}}(t,x) \wedge T_i$ in the sense of Definition 2.3. Notice that the sequential exit-times of the value function V_k correspond to an excursion through the sets $(A_i)_{i=k}^n$ irrespective of the first $(k-1)$ sets. It is straightforward to observe that the value functions V and $\widetilde{V}$ in (4) are particular cases of the value function V_1 defined as in (7) for an appropriate selection of the sets $(A_i)_{i=1}^n$, functions $(\ell_i)_{i=1}^n$, and intermediate times $(T_i)_{i=1}^n$.

Assumption 2.6: We stipulate that

a. The diffusion term σ of the SDE (1) is uniformly *non-degenerate*, i.e., there exists $\delta > 0$ such that for all $x \in \mathbb{R}^d$ and $u \in \mathbb{U}$, $\|\sigma\sigma^\top\| > \delta$.
b. The open sets A_i satisfy the *exterior cone condition*; for mathematical explanation we refer to [7, Assumption 5.2].
c. $(\ell_i)_{i=1}^n$ are all lower semicontinuous.

Definition 2.7 (Dynkin Operator): Given $u \in \mathbb{U}$, we denote by $\mathcal{L}^u$ the Dynkin operator (also known as the infinitesimal generator) associated to the controlled diffusion (1) as

$$\mathcal{L}^u\Phi(t,x) := \partial_t\Phi(t,x) + f(x,u).\partial_x\Phi(t,x) + \frac{1}{2}\mathrm{Tr}[\sigma\sigma^\top(x,u)\partial_x^2\Phi(t,x)],$$

where Φ is a real-valued function smooth on the interior of $\mathbb{S}$, with $\partial_t\Phi$ and $\partial_x\Phi$ denoting the partial derivatives with respect to t and x respectively, and $\partial_x^2\Phi$ denoting the Hessian matrix with respect to x. We refer to [12, Theorem 17.23] for more details on the above differential operator.

The following Theorem is the main result of this section and we refer to [7, Section 5] for the detailed proofs and further discussion on numerical issues.

Theorem 2.8: Consider the system (1), and suppose that Assumptions 2.6 hold. Let the value functions $V_k : \mathbb{S} \to \mathbb{R}^d$ be as defined in (7). Then, for all k in $\{1, \cdots, n\}$ with a convention $V_{n+1} \equiv 1$, we have

$$\begin{cases} -\sup_{u \in \mathbb{U}} \mathcal{L}^u V_k(t,x) = 0 & \text{on} \quad [0,T_k[\times A_k, \\ V_k(t,x) = V_{k+1}(t,x)\ell_k(x) & \text{on} \quad [0,T_k] \times A_k^c \bigcup \{T_k\} \times \mathbb{R}^d \end{cases}$$

Notice that Theorem 2.8 allows us to obtain the value function V_k, given value function V_{k+1}, by solving a certain PDE with some boundary (possibly both terminal and lateral) conditions. Hence, in light of $V_{n+1} \equiv 1$, one can infer that Theorem 2.8 suggests a series of PDE equations for which the first one has known boundary condition ℓ_n, while the boundary conditions of the subsequent steps are determined by the solution of the preceding PDE step, i.e., V_{k+1} provides boundary conditions for the PDE corresponding to the value function V_k. Let us highlight that the basic motion planning maneuver involving only two sets is effectively the same as the first step of this series of PDEs and was studied in our earlier work [13], [14].

III. Application to the Toggle Switch

As described in Section I, we consider the toggle switch model structure proposed in [7, Section 6] that describes the dynamics of a system consisting of two mutually repressing genes. Starting from a continuous time Markov chain description of the process [15], we consider its Langevin approximation given by the following set of stochastic differential equations:

$$\mathrm{d}X_t = \big(f(Y_t,\boldsymbol{u}_x) - \mu_x X_t\big)\mathrm{d}t + \sqrt{f(Y_t,\boldsymbol{u}_x)}\mathrm{d}W_t^1 + \sqrt{\mu_x X_t}\mathrm{d}W_t^2, \tag{8a}$$

$$\mathrm{d}Y_t = \big(g(X_t,\boldsymbol{u}_y) - \mu_y Y_t\big)\mathrm{d}t + \sqrt{g(X_t,\boldsymbol{u}_y)}\mathrm{d}W_t^3 + \sqrt{\mu_y Y_t}\mathrm{d}W_t^4, \tag{8b}$$

where X_t and Y_t are the concentrations of the two transcription factors with the respective degradation rates μ_x and μ_y; $(W_t^i)_{t\geq}$ are independent standard Brownian motion processes. The production rate functions f and g are defined by *Hill* functions as follows [16]:

$$f(y,u) := \frac{k_1\theta_1^{n_1}}{y^{n_1} + \theta_1^{n_1}}u, \qquad g(x,u) := \frac{k_2\theta_2^{n_2}}{x^{n_2} + \theta_2^{n_2}}u, \tag{9}$$

where θ_i are the threshold of the production rate with respective Hill exponents n_i, and k_i are the production rate scaling factors. The parameter u plays the role of an external control signal that can affect the production rate of a gene (e.g., through a light-sensitive response element [3]). The control signals u_x and u_y are assumed to take values in the sets $\mathbb{U}_x := [\underline{u}_x, \overline{u}_x]$ and $\mathbb{U}_y := [\underline{u}_y, \overline{u}_y]$ respectively, and to be fixed to a nominal value $\hat{u}$ when no control is applied. The model parameters are chosen in a way such that when $u \equiv \hat{u}$ the system without noise exhibits two stable equilibria at z_a and z_c, separated by an unstable one at z_b. Note that the model equations are well-posed only when $X_t, Y_t \geq 0$ (a common characteristic of Langevin approximations [17]). This issue will be taken into account in the design of the control objective.

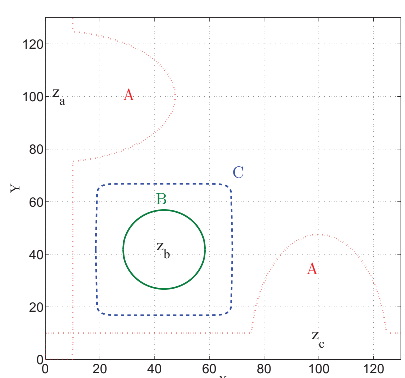

Fig. 1. The set A is an avoidance region contained in the region of attraction of the stable equilibria z_a and z_c, B is the target set around the unstable equilibrium z_b, and C is the maintenance margin.

A. Motion-Planning Results

In this example we consider system (8) with the rate functions (9) along with the following parameters: $\theta_i = 40$, $\mu_i = 0.04$, $k_i = 4$ for both $i \in \{1,2\}$, and the exponents

$n_1 = 4$, $n_2 = 6$, where the nominal control value is $\hat{u} := 1$. Figure 1 shows the locations of the system equilibria under this parametrization. Our objective is to first steer the number of proteins towards a target set around the unstable equilibrium by synthesizing appropriate input signals $\boldsymbol{u}_x$ and $\boldsymbol{u}_y$ within a certain time horizon, say T_1. During this task we opt to avoid the region of attraction of the stable equilibria as well as low numbers for each protein; the latter justifies our model being well-posed in the region of interest. The aforementioned target and avoid sets are denoted, respectively, by the closed sets B and A in Figure 1. In the second phase of the task, once the trajectory visits the target set B, it is required to keep the molecular populations within a slightly larger margin around equilibrium z_b for some time, say T_2; Figure 1 depicts this maintenance margin by the open set C.

Therefore, the motion planning consists of two parts: reaching the target set B while avoiding the set A within the certain horizon T_1, and staying in the set C for a certain time T_2 after visiting the set B for the first time. In view of motion-planning events introduced in Definition 2.1, the first phase of the path can be expressed as $(A^c \rightsquigarrow B)_{\le T_1}$, and the second phase as $(C \xrightarrow{T_2} C)$; see (2) for detailed definitions of these symbols. By defining the joint process $Z^{t,z;\boldsymbol{u}}_{\cdot} := \big[X^{t,x;\boldsymbol{u}}_{\cdot}, Y^{t,y;\boldsymbol{u}}_{\cdot}\big]$, with the initial condition $z := (x, y)$, the desired excursion is a combination of the events studied in the preceding sections and, with a slight abuse of notation, can be expressed by

$$\Big\{ Z^{t,z;\boldsymbol{u}}_{\cdot} \models (A^c \rightsquigarrow B)_{\le T_1} \circ (C \xrightarrow{T_2} C) \Big\}.$$

The above event depends, of course, on the initial condition (t, x) and control policy $\boldsymbol{u}$, and the objective is to maximize its probability over all admissible policies $\boldsymbol{u} := [\boldsymbol{u}_x, \boldsymbol{u}_y]$. The desired path is not exactly in the framework of Definition 2.1 but, nonetheless, one can invoke the same ideas as in Section II and introduce the following value functions:

$$V_1(t,z) := \sup_{\boldsymbol{u}\in\mathcal{U}_t} \mathbb{E}\Big[\mathbb{1}_B\big(Z^{t,z;\boldsymbol{u}}_{\tau_1^1}\big)\mathbb{1}_C\big(Z^{t,z;\boldsymbol{u}}_{\tau_2^1}\big)\Big], \tag{10a}$$

$$V_2(t,z) := \sup_{\boldsymbol{u}\in\mathcal{U}_t} \mathbb{E}\Big[\mathbb{1}_C\big(Z^{t,z;\boldsymbol{u}}_{\tau_2^2}\big)\Big], \tag{10b}$$

where τ_1^1 and τ_2^2 are defined in a same spirit as (7) with given sets $A_1 := (A \cup B)^c$ and $A_2 := C$. However, the stopping time τ_2^1 requires a slight modification so as to address the combination of both motion-planning events introduced in Definition 2.1: $\tau_2^1 := \Theta_2^{A_{1:2}} \wedge (\tau_1^1 + T_2)$.

The solution of our motion planning objective is the value function V_1 in (10a), which in view of Theorem 2.8 is characterized by the Dynkin differential operator in the interior of $[0, T_1[\times(A \cup B)^c$. However, we need first to solve numerically for the auxiliary value function V_2 in (10b) in order to provide boundary conditions for the PDE corresponding to V_1 by

$$V_1(t,z) = \mathbb{1}_B(z)V_2(t,z), \tag{11}$$

for all (t, z) in $[0, T_1] \times (A \cup B) \bigcup \{T_1\} \times \mathbb{R}^n$. It is straightforward to observe that the boundary condition for the value function V_2 is

$$V_2(t,z) = \mathbb{1}_C(z),$$

for (t, z) in $[0, T_1 + T_2] \times C^c \bigcup \{T_1 + T_2\} \times \mathbb{R}^n$. Therefore, we need to solve the PDE of V_2 with the above boundary condition backward from the time $T_1 + T_2$ to the time T_1, and then at time T_1 restrict the value function V_2 onto the set B to provide boundary conditions for the value function V_1. Thus, the value function V_1 can be computed via solving the same PDE from T_1 to 0 but along with different boundary conditions provided by the preceding step. According to Definition 2.7 for any smooth function $\phi := \phi(t, x, y)$ the Dynkin operator $\mathcal{L}^u$ can be simplified to

$$\begin{aligned}
&\sup_{u\in\mathbb{U}} \mathcal{L}^u\phi(t,x,y)\\
&\quad = \max_{u\in\mathbb{U}} \Big[\partial_t\phi + \partial_x\phi\big(f(y,u_x) - \mu_x x\big) + \partial_y\phi\big(g(x,u_y) - \mu_y y\big)\\
&\qquad + \frac{1}{2}\partial_x^2\phi\big(f(y,u_x) + \mu_x x\big) + \frac{1}{2}\partial_y^2\phi\big(g(x,u_y) + \mu_y y\big)\Big]\\
&\quad = \partial_t\phi - \big(\partial_x\phi - \frac{1}{2}\partial_x^2\phi\big)\mu_x x - \big(\partial_y\phi - \frac{1}{2}\partial_y^2\phi\big)\mu_y y\\
&\qquad + \max_{u_x\in[\underline{u}_x,\overline{u}_x]} \big[f(y,u_x)\big(\partial_x\phi + \frac{1}{2}\partial_x^2\phi\big)\big]\\
&\qquad + \max_{u_y\in[\underline{u}_y,\overline{u}_y]} \big[g(x,u_y)\big(\partial_y\phi + \frac{1}{2}\partial_y^2\phi\big)\big].
\end{aligned}$$

Let us introduce $\delta_x^i := \partial_x V_i + \frac{1}{2}\partial_x^2 V_i$ and $\delta_y^i := \partial_y V_i + \frac{1}{2}\partial_y^2 V_i$. On account of Theorem 2.8 and linearity of the drift terms in u, one can propose an optimal policy in terms of derivatives of the value functions V_1 and V_2 in (10), respectively, for the first and second phase of the motion:

$$\boldsymbol{u}_x^*(t,x,y) = \begin{cases} \overline{u}_x(t,x,y) & \text{if} \quad \delta_x^i(t,x,y) > 0,\\ \underline{u}_x(t,x,y) & \text{if} \quad \delta_x^i(t,x,y) \le 0,\end{cases} \tag{12a}$$

$$\boldsymbol{u}_y^*(t,x,y) = \begin{cases} \overline{u}_y(t,x,y) & \text{if} \quad \delta_y^i(t,x,y) > 0,\\ \underline{u}_y(t,x,y) & \text{if} \quad \delta_y^i(t,x,y) \le 0,\end{cases} \tag{12b}$$

where $i \in \{1, 2\}$ corresponds to the phase of the path.

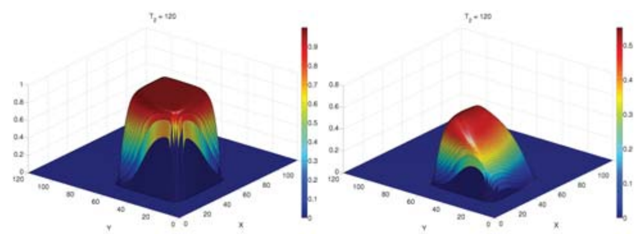

(a) V_2 in case of full controllability over both production rates.

(b) V_2 in case of half controllability over only the production rate of protein x.

Fig. 2. The value function V_2 as defined in (10b) corresponding to probability of staying in C for 120 time units.

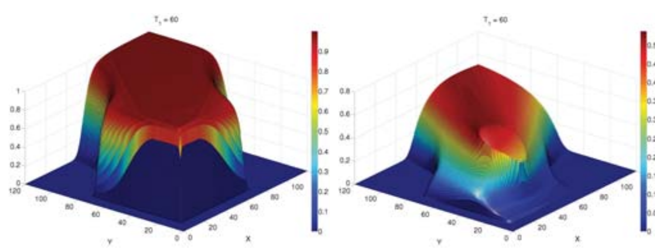

(a) V_1 in case of full controllability over the production rates.

(b) V_1 in case of inability to increase the production rate of protein y.

Fig. 3. The value function V_1 as defined in (10a) corresponding to probability of staying in C for 120 time units, once it reaches B while avoiding A within 60 time units.

In the sequel we investigate two scenarios: first, when full control over both production rates is possible, i.e., $\underline{u}_x =$

$\underline{u}_y = 0$ and $\overline{u}_x = \overline{u}_y = 2$; second, when we only have access to the production rate of protein x, i.e., $\overline{u}_y = \underline{u}_y = \hat{u}$. Figure 2 depicts the probability distribution of staying in set C within the time horizon $T_2 = 120$ time units [2] in terms of the initial conditions $(x, y) \in \mathbb{R}^2$. V_2 is zero outside set C, as the process has obviously left C if it starts outside it. Figures 2(a) and 2(b) demonstrate the first and second control scenarios, respectively. Note that in the second case the probability of success dramatically decreases in comparison with the first. This result indicates the importance of full controllability of the production rates for the achievement of the given control objective.

Figure 3 depicts the probability of successively reaching set B within the time horizon $T_1 = 60$ time units and staying in set C for $T_2 = 120$ time units thereafter. Since the objective is to avoid set A, the value function V_1 takes zero value on A. Figures 3(a) and 3(b) demonstrate the first and second control scenarios, respectively. It is easy to observe the non-smooth behavior of the value function V_1 on the boundary of set B in Figure 3(b). This is indeed a consequence of the boundary condition explained in (11). All simulations in this subsection were obtained using the Level Set Method Toolbox [18] (version 1.1), with a grid 121×121 in the region of simulation.

B. Autonomous Feedback Policy

As pointed out earlier, in the process of solving the corresponding PDEs of value functions V_i, one can compute the optimal time-varying feedback controls (12). However, it is of practical interest to introduce a suboptimal but autonomous (time-invariant) feedback control for the respective objectives. For this purpose, during the first phase of the path, $(A^c \rightsquigarrow B)_{\le T_1}$, one can employ a heuristic approach which basically assigns a control action to the state (x, y) for which the value function V_1 crosses a certain probability level $p_0 \in (0, 1)$ for the first time while the corresponding PDE is computed backward in time. Notice that the threshold level p_0 essentially can be viewed as a design parameter that indicates the trade-off between a direct way towards the goal versus the risk of hitting the obstacles; for more details on this idea see [19].

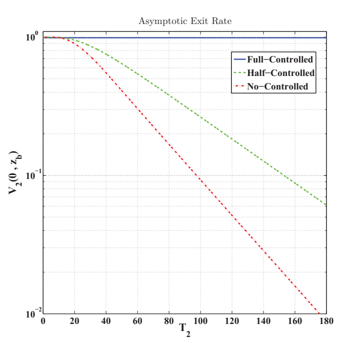

Fig. 4. Asymptotic exit-rate λ^* introduced in (13) for different scenarios

The objective of the second phase, $(C \xrightarrow{T_2} C)$, is to stay in set C for a given amount of time. The probability of such an event obviously tends to zero in long-run (as T_2 goes to infinity), due to the non-degeneracy of the diffusion terms (Assumption 2.6.a.) and the unboundedness of the noise. In order to synthesize an autonomous feedback policy to

[2]Notice that the half-life of each protein is assumed to be 17.32 time units

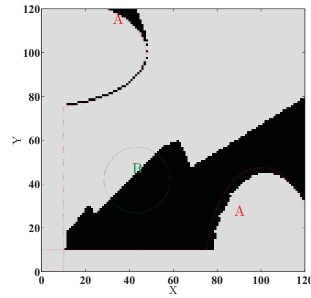

(a) Two controlled rates: controller $\boldsymbol{u}_x \in [0, 2]$

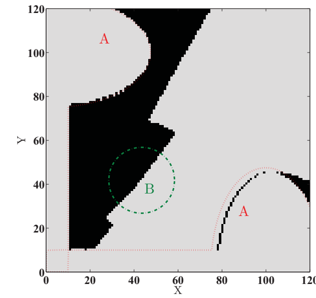

(b) Two controlled rates: controller $\boldsymbol{u}_y \in [0, 2]$

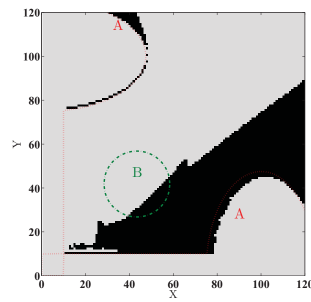

(c) One controlled rate: controller $\boldsymbol{u}_x \in [0, 2]$

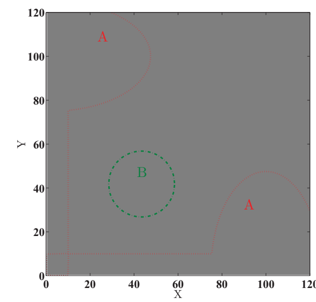

(d) One controlled rate: controller $\boldsymbol{u}_y \equiv 1$

Fig. 5. Autonomous feedback control policies for the two different control scenarios. White and black regions denote the lower and upper bounds of the control sets respectively, while grey indicates the inactive controller.

keep the process inside C, one may aim to minimize the asymptotic rate with which the process exits C. Formally speaking, this suggests minimizing

$$\lambda := -\limsup_{t \to \infty} \frac{1}{t} \log \mathbb{P}\big(\tau_C > t\big), \tag{13}$$

where τ_C is the first exit-time from the set C, and is the same as $\Theta_2^{A_{2:2}}$ used for the definition of V_2 in (10b). Notice that in general λ could be a function of initial condition. However, it can be shown that the tail of the distribution of τ_C should be independent of the starting state of the process, which has been studied in literature in the context of *quasi-stationary* distributions [20], [21]; hence, the initial state is omitted from the notation. In this sense, τ_C has the same tail as the stopping time $\Theta_2^{A_{2:2}}$ used for the definition of V_2 in (10b). As shown in [22] for controlled diffusion processes, under some mild assumptions, there exist an optimal exit-rate λ^* as defined in (13) and a corresponding twice differentiable function $\psi^* : C \to [0, 1]$ such that a Markov control $\boldsymbol{u}^*$ is optimal if and only if $\boldsymbol{u}^*$ is a measurable selector of $\operatorname{argmax}_{u \in \mathbb{U}}(\mathcal{L}^u \psi^*)$, where ψ^* is the unique solution (up to a scalar multiple) of the following eigenvalue problem:

$$-\mathcal{L}^{\boldsymbol{u}^*(x,y)} \psi^*(x, y) = \lambda^* \psi^*(x, y), \qquad \forall (x, y) \in C. \tag{14}$$

Let us highlight that the Dynkin operator introduced in Definition 2.7 is slightly different than the extended generator of [22]. However, since the eigenfunction ψ^* is time-invariant, then $\partial_t \psi^* = 0$, and consequently the implication of [22] holds for our problem in (14).

Therefore, from the asymptotic exit-rate (13) perspective, eigenfunction ψ^* can be used to determine the optimal autonomous feedback policy $\boldsymbol{u}^*$. Notice that in problem

(14) the eigenfunction ψ^* can be determined up to a scalar multiple, i.e., its shape would be enough to form the optimal feedback policy via $\boldsymbol{u}^* \in \operatorname{argmax}_{u\in\mathbb{U}} \mathcal{L}^u\psi^*$. To the best of our knowledge, there is no explicit numerical technique to tackle problem (14). Nevertheless, one can observe that if during the process of solving the PDE of Theorem 2.8 backwards in time the value function V_2 gets shape-wise saturated, then the corresponding PDE turns into the eigenvalue problem (14) as $t \to \infty$, and consequently the shape and decay rate of the solution should approximate ψ^* and λ^*, respectively. As an illustration, Figures 2 depict the eigenfunction ψ^* for the case of two and one controlled production rate scenarios described in the preceding subsection. In fact, starting from the terminal condition $\mathbb{1}_C$, the PDE solutions quickly transform into the shapes depicted in Figures 2, and thereafter only decay exponentially over time with a rate equal to the optimal exit rate λ^*. Figure 4 depicts the decay of $V_2(0, z_b)$ versus the final time T_2 for three different control scenarios: full and partial control of production rates, as defined in Subsection III-A, as well as the case in which neither of the production rates is controlled, i.e., $\boldsymbol{u}_x = \boldsymbol{u}_y \equiv 1$. As expected, the asymptotic decay rate of $V_2(0, z_b)$ is smallest when both genes can be controlled and largest when no control is used.

Returning to the autonomous feedback policy calculation, using the eigenfunctions shown on Figure 2 in conjunction with (12), one can propose a *bang-bang* feedback control for set C that only takes the extremum values of the control sets. It then suffices to restrict this policy to set B and concatenate with the feedback control introduced in the preceding part so as to derive one autonomous feedback policy for the objective of Subsection III-A. Figure 5 demonstrates the autonomous feedback policies obtained in this manner for the two controllers and the different control scenarios. The corresponding region for each extremum value is depicted by a certain color code in Figure 5. Figures 5(a) and 5(b) show the optimal feedback control for the case of two controlled production rates. Similarly, Figures 5(c) and 5(d) depict the feedback policy when only u_x can be controlled (for this reason, in Figure 5(d) the production rate of protein Y is fixed to the nominal value and uniformly colored).

IV. Conclusion and Future Directions

In this work, we have approached the control problem for a small toggle switch model within the framework of stochastic motion planning. Assuming that production rates for one or two genes can be controlled externally within certain bounds, we calculated the time-invariant feedback control laws that achieve the prescribed objective of keeping the system within a region that contains its unstable equilibrium and evaluated their performance. Our results suggest that control of both genes is necessary to steer and keep the system in the desired set with high probability for a significant amount of time, and serves as a first proof of concept for the application of feedback control on a bistable biological system. Further testing of different control configurations can be accomplished relatively easily using our approach.

The generality and flexibility of the stochastic motion planning permits the formulation of quite complex control tasks, that can be translated into the solution of a sequence of PDEs, which in turn completely characterize the probabilistic properties of the resulting closed-loop systems. On the downside, the need to numerically solve these PDEs limits applicability of this approach to systems consisting of just a few states. Moreover, we should remark that our approach essentially results in *single-cell* feedback control laws, that can be practically implemented only if one is able to perform single-cell measurements of both proteins and apply the appropriate inputs to each cell separately. While the first requirement is not too restrictive, given the current measurement capabilities of fluorescently tagged proteins, the latter remains technically challenging. A first implementation of a single-cell feedback control scheme has been given in [4] using light as input. In any case, single-cell feedback control provides an "upper bound" to the performance of cellular population control using a single input signal for all cells simultaneously.

References

[1] J.-W. Veening, W. K. Smits, and O. P. Kuipers, "Bistability, epigenetics, and bet-hedging in bacteria," *Annual Review of Microbiology*, vol. 62, no. 1, pp. 193–210, 2008.

[2] C. Bashor, A. Horwitz, S. Peisajovich, and W. Lim, "Rewiring cells: synthetic biology as a tool to interrogate the organizational principles of living systems," *Annual review of biophysics*, vol. 39, p. 515, 2010.

[3] A. Milias, S. Summers, J. Stewart-Ornstein, I. Zuleta, D. Pincus, H. El-Samad, M. Khammash, and J. Lygeros, "In silico feedback for in vivo regulation of a gene expression circuit," *Nature Biotechnology*, vol. 29, no. 12, p. 11141116, Dec. 2011. [Online]. Available: http://control.ee.ethz.ch/index.cgi?page=publications;action=details;id=3903

[4] J. Toettcher, D. Gong, W. Lim, and O. Weiner, "Light-based feedback for controlling intracellular signaling dynamics," *Nature methods*, 2011.

[5] J. Uhlendorf, A. Miermont, T. Delaveau, G. Charvin, F. Fages, S. Bottani, G. Batt, and P. Hersen, "Long-term model predictive control of gene expression at the population and single-cell levels," vol. 109, no. 35, pp. 14 271–14 276, 2012.

[6] T. S. Gardner, C. R. Cantor, and J. J. Collins, "Construction of a genetic toggle switch in Escherichia coli," *Nature*, vol. 403, no. 6767, pp. 339–42, Jan. 2000.

[7] P. Mohajerin Esfahani, D. Chatterjee, and J. Lygeros, "Motion Planning via Optimal Control for Stochastic Processes," Tech. Rep., Oct. 2012. [Online]. Available: http://control.ee.ethz.ch/index.cgi?page=publications;action=details;id=4222

[8] I. Karatzas and S. Shreve, *Brownian Motion and Stochastic Calculus*, 2nd ed., ser. Graduate Texts in Mathematics. New York: Springer-Verlag, 1991, vol. 113.

[9] V. S. Borkar, "Controlled diffusion processes," *Probability Surveys*, vol. 2, pp. 213–244 (electronic), 2005. [Online]. Available: http://dx.doi.org/10.1214/154957805100000131

[10] M. G. Crandall, H. Ishii, and P. L. Lions, "User's guide to viscosity solutions of second order partial differential equations," *American Mathematical Society*, vol. 27, pp. 1–67, 1992. [Online]. Available: http://dx.doi.org/10.1214/154957805100000131

[11] W. Fleming and H. Soner, *Controlled Markov Processes and Viscosity Solution*, 3rd ed. Springer-Verlag, 2006.

[12] O. Kallenberg, *Foundations of Modern Probability*, ser. Probability and its Applications (New York). New York: Springer-Verlag, 1997.

[13] P. Mohajerin Esfahani, D. Chatterjee, and J. Lygeros, "On a problem of stochastic reach-avoid set characterization for diffusions," in *IEEE Conference on Decision and Control*, dec. 2011, pp. 7069 –7074.

[14] ——, "On A Stochastic Reach-Avoid Problem and Set Characterization," http://arxiv.org/abs/1202.4375, Feb. 2012.

[15] M. Khammash, "Modeling and analysis of stochastic biochemical networks," in *Control Theory and System Biology*.

[16] J. Cherry, "How to make a Biological Switch," *Journal of Theoretical Biology*, vol. 203, no. 2, pp. 117–133, Mar. 2000.

[17] D. J. Higham, "Stochastic ordinary differential equations in applied and computational mathematics," *IMA Journal of Applied Mathematics*, vol. 76, no. 3, pp. 449–474, 2011.

[18] I. Mitchell, "A toolbox of hamilton-jacobi solvers for analysis of nondeterministic continuous and hybrid systems," in *Hybrid systems: computation and control*, ser. Lecture Notes in Comput. Sci., M. Morari and L. Thiele, Eds. Springer-Verlag, 2005, no. 3414, pp. 480–494.

[19] T. Wood, P. Mohajerin Esfahani, and J. Lygeros, "Hybrid Modelling and Reachability on Autonomous RC-Cars," in *IFAC Conference on Analysis and Design of Hybrid Systems (ADHS)*, Eindhoven, Netherlands, jun. 2012.

[20] S. Méléard and D. Villemonais, "Quasi-stationary distributions and population processes," *Probability Surveys*, vol. 9, pp. 340–410, 2012.

[21] R. Pinsky, "On the convergence of diffusion processes conditioned to remain in a bounded region for large time to limiting positive recurrent diffusion processes," *The Annals of Probability*, vol. 13, no. 2, pp. 363–378, 1985.

[22] A. Biswas and V. S. Borkar, "On a controlled eigenvalue problem," *Systems Control Lett.*, vol. 59, no. 11, pp. 734–735, 2010. [Online]. Available: http://dx.doi.org/10.1016/j.sysconle.2010.08.009