

Toggling a Genetic Switch Using Reinforcement Learning

Aivar Sootla¹, Natalja Strelkova², Mauricio Barahona³, Damien Ernst⁴, Guy-Bart Stan¹

¹ Centre for Synthetic Biology and Innovation and the Department of Bioengineering, Imperial College London, UK {a.sootla, g.stan}@imperial.ac.uk

² Boehringer-Ingelheim Pharma GmbH & Co KG., Germany
natalja.strelkova@boehringer-ingelheim.com

³ Department of Mathematics, Imperial College London, UK m.barahona@imperial.ac.uk

⁴ Montefiore Institute, University of Liège, Belgium dernst@ulg.ac.be.

Abstract : In this paper, we consider the problem of optimal exogenous control of gene regulatory networks. Our approach consists in adapting an established reinforcement learning algorithm called the fitted Q iteration. This algorithm infers the control law directly from the measurements of the system's response to external control inputs without the use of a mathematical model of the system. The measurement data set can either be collected from wet-lab experiments or artificially created by computer simulations of dynamical models of the system. The algorithm is applicable to a wide range of biological systems due to its ability to deal with nonlinear and stochastic system dynamics. To illustrate the application of the algorithm to a gene regulatory network, the regulation of the toggle switch system is considered. The control objective of this problem is to drive the concentrations of two specific proteins to a target region in the state space.

1 Introduction

Synthetic biology aims at the (re-)design of biological functions in living organisms for their use in various applications such as bioengineering, bioremediation and energy (Purnick & Weiss (2009)). This is typically realised via the insertion of foreign genes inside a host cell (e.g., a bacterium *E. coli*). The expression of the foreign genes inside the host cells imposes *de facto* a burden on the native processes of the host cells. A high burden induces severe intracellular perturbations and can decrease cellular growth rate. This in turn disrupts the intended behaviour of synthetic biology gene networks (Tan *et al.* (2009)). Hence, it is highly desirable to develop means for controlling gene networks so as to efficiently enable the designed behaviour while simultaneously minimising the burden induced by this behaviour on the host cells.

The current biotechnology state-of-the-art allows us to quantitatively measure and interact with gene regulatory networks. Quantitative *in vivo* estimates of gene networks' states (outputs) can be obtained via fluorescent markers (Cai *et al.* (2006); Bennett & Hasty (2009)) (e.g., green fluorescent protein, GFP or red fluorescent protein, mCherry). A typical input is a targeted induction of the gene expression, which can be achieved by, e.g., conditional gene knock outs (Ivanova *et al.* (2006); Liu *et al.* (2007)), heat shocks (Mettetal *et al.* (2008)) or monochromatic light pulses (Shimizu-Sato *et al.* (2002); Levskaya *et al.* (2009)). This means that feedback control is technologically feasible *in vivo*. The objective of the control method can be minimal time control (i.e., driving the system as fast as possible to a target region in the state-space), minimal burden control (minimal expression of heterologous proteins), or a trade-off between the two, as considered in this paper. The control method must reach the objective, while maintaining the designed functions of a synthetic gene regulatory network.

Some control problems in gene regulatory networks were successfully addressed (Menolascina *et al.* (2011); Uhlenendorf *et al.* (2012); Miliadis-Argeitis *et al.* (2011)). In all those papers, the authors used classical control methods, which infer the control law (or the control policy) based on a mathematical model of the system. One of the bottlenecks of these approaches is the modelling part, which for large gene regulatory networks is an extremely hard and lengthy process. Moreover, there are other challenges such as stochasticity. Stochasticity is expressed in the form of the intrinsic and extrinsic noise during gene expression (Swain *et al.* (2002)). Transcription and translation processes typically involve a few randomly interacting molecules, thus adding thermodynamic stochasticity to biochemical interactions.

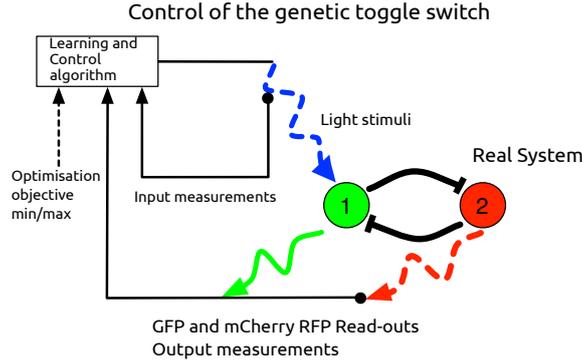


Figure 1: A schematic depiction of the exogenously controlled genetic toggle switch. The green circle represents the *lacI* gene and the red circle represents the *tetR* gene. The arrows with flat ends represent repression of one gene by another. In the steady-state only one of the genes can be upregulated (or switched on). The goal is to toggle one of the genes, i.e., drive this gene from its downregulated mode to its upregulated one.

The problems with modelling and stochasticity point towards the use of reinforcement learning methods (Sutton & Barto (1998); Buşoniu *et al.* (2010)), which infer the control policy based solely on interactions with the real system. These methods do not require a physical model. Moreover, very few assumptions on the structure of the controlled system are made. However, the major advantage of the reinforcement learning methods is to some extent their drawback. Indeed, these methods require interactions with the real system, which implies numerous costly and lengthy wet-lab experiments. A solution would be a reinforcement learning method, which learns the policy using a single experiment. For systems relevant to this paper, however, such a method will not be efficient. Indeed, a control policy, which tries to learn and control such systems in a single experiment, is generally not better than a random control policy (Castronovo *et al.* (2012)). In order to address these concerns, a hybrid approach is proposed. First, an initial control policy is computed using past experimental data and/or a mathematical model. After that the control policy is updated during the experiment using reinforcement learning methods. This approach will be applied to the regulation of the toggle switch system schematically depicted in Figure 1. The control objective of this problem is to drive the concentrations of two specific proteins to a target set in the state space and remain in this set.

The initial policy is obtained by the Fitted Q Iteration algorithm (Ernst *et al.* (2005)). The algorithm requires only one-step system transitions to infer the control policy. A one-step system transition is a triplet $\{\mathbf{n}, \mathbf{a}, \mathbf{n}^+\}$, where \mathbf{n}^+ denotes a successor state of the system in state \mathbf{n} subjected to input \mathbf{a} . Fitted Q Iteration can also handle nonlinear and stochastic systems and it is sample efficient. One-step transitions can be obtained by simulating the mathematical model of the system or using past experimental data. Afterwards the policy is updated by mixing the online measurements with past observations. The Exploration/Exploitation trade-off is addressed using an ε -greedy policy.

This paper is organised as follows. Mathematical preliminaries are described in Section 2. In order to make the paper self contained, the fitted Q algorithm is sketched and different aspects of modelling in gene regulatory networks are discussed. The problem of controlling the toggle switch is formulated and discussed in detail in Section 3. Finally, the simulation results are presented in Section 4. Reference trajectory tracking for the generalised repressilator system is the subject of our previous publication (Sootla *et al.* (2013)).

2 Preliminaries

2.1 Modelling in Biology

The following approach to chemical reaction modelling is described in detail in (Gillespie (1977)). Consider a well-stirred system of k species in a constant volume Ω and a thermal equilibrium. Assume the species are interacting through m reactions. Let $N^i(t)$ be the number of molecules of species i and $\nu_{ij}(t)$ be the change in the molecular concentration of species i at time t if the reaction j occurs. The bold symbols

will be used to denote vectors, e.g., \mathbf{n} stands for the vector with elements N^i . Finally, let $a_j(\mathbf{n})dt$ be the probability of reaction j occurring in the next infinitesimal interval $[t, t + dt]$, if the number of molecules at time t , $\mathbf{N}(t)$, is equal to \mathbf{n} . The functions $a_j(\cdot)$ are called propensity functions. At the cellular level chemical reactions depend on thermodynamical principles, since molecules must collide before a reaction can start. Therefore chemical reactions inside living organisms are modelled using stochastic calculus. The time evolution of the concentration of species can be modelled by a Markov stochastic process, for which:

$$\frac{\partial \Pr(\mathbf{n}, t | \mathbf{n}_0, t_0)}{\partial t} = \sum_{j=1}^m a_j(\mathbf{n} - \boldsymbol{\nu}_j) \Pr(\mathbf{n} - \boldsymbol{\nu}_j, t | \mathbf{n}_0, t_0) - a_j(\mathbf{n}) \Pr(\mathbf{n}, t | \mathbf{n}_0, t_0) \quad (1)$$

where the probability $\Pr(\mathbf{n}, t | \mathbf{n}_0, t_0)$ stands for $\Pr(\mathbf{N}(t) = \mathbf{n} | \mathbf{N}(t_0) = \mathbf{n}_0)$. This equation is called the *Chemical Master Equation*. The propensity functions a_j depend also on the volume Ω . It can be shown that for large volumes Ω the CME (1) becomes a deterministic equation

$$\frac{d\mathbf{n}(t)}{dt} = \sum_{j=1}^m \boldsymbol{\nu}_j \tilde{a}_j(\mathbf{n}(t)), \quad (2)$$

where the propensities \tilde{a}_j are independent of the volume Ω . For small volumes Ω , the stochastic model (1) describes better the behaviour of the cells than the deterministic model (2). Hence, in synthetic biology setting using a stochastic model is preferable. Nevertheless the deterministic model can be still useful for small volumes in order to provide some idea of the system behaviour, since stochastic models are harder to simulate and analyse.

2.2 Formulation of the Optimal Control Problem

Consider a deterministic discrete-time dynamical system

$$\mathbf{n}_{t+1} = f(\mathbf{n}_t, \mathbf{a}_t) \quad (3)$$

where \mathbf{a}_t is the control input at time t , which belongs to a compact set \mathbf{A} for every t . In the stochastic case, Markov decision processes (MDPs) are typically employed, for which

$$\Pr(\mathbf{n}_{t+1} \in \mathbf{N}_{t+1} | \{\mathbf{n}_k\}_{k=0}^t, \{\mathbf{a}_k\}_{k=0}^t) = \Pr(\mathbf{n}_{t+1} \in \mathbf{N}_{t+1} | \mathbf{n}_t, \mathbf{a}_t).$$

The above relationship means that the probability of the state \mathbf{n}_{t+1} belonging to the set \mathbf{N}_{t+1} does not depend on the entire history of the realisation of the states $\{\mathbf{n}_k\}_{k=0}^t$ and control signals $\{\mathbf{a}_k\}_{k=0}^t$, but depends only on the current values \mathbf{n}_t and \mathbf{a}_t . Under the above Markovian assumption, dynamical stochastic systems can be modelled as

$$\Pr(\mathbf{n}_{t+1} \in \mathbf{N}_{t+1} | \mathbf{n}_t, \mathbf{a}_t) = \int_{\mathbf{N}_{t+1}} f(\mathbf{n}_t, \mathbf{a}_t, x) dx,$$

or in a compact form

$$\mathbf{n}_{t+1} \sim f(\mathbf{n}_t, \mathbf{a}_t, \cdot).$$

Here, we slightly abuse the notation by using again the symbol f as in the deterministic system (3). This is done, in order to signify that these functions describe the dynamics of the system whether it is stochastic or deterministic.

In both cases, consider an optimal control problem, which is defined through the minimisation of an infinite sum of discounted costs $\mathbf{c}(\mathbf{n}, \mathbf{a})$. In the deterministic case the problem is defined as

$$V(\mathbf{n}_t) = \min_{\pi(\cdot): \pi(\mathbf{n}_i) = \mathbf{a}_i} \sum_{i=t}^{\infty} \gamma^{i-t} \mathbf{c}(\mathbf{n}_i, \mathbf{a}_i)$$

and in the stochastic case as

$$V(\mathbf{n}_t) = \min_{\pi(\cdot): \pi(\mathbf{n}_i) = \mathbf{a}_i} \lim_{K \rightarrow \infty} \mathbb{E}_{\mathbf{n}_{t+1} \sim f(\mathbf{n}_t, \mathbf{a}_t, \cdot)} \sum_{i=t}^K \gamma^{i-t} \mathbf{c}(\mathbf{n}_i, \mathbf{a}_i)$$

Algorithm 1 Fitted Q iteration algorithm**Inputs:** Set of triplets $\mathcal{F} = \{\mathbf{n}_l, \mathbf{a}_l, \mathbf{n}_l^+\}_{l=1}^{\#\mathcal{F}}$, stopping criterion, cost function $c(\cdot, \cdot)$ **Outputs:** Policy $\hat{\pi}^*(\mathbf{n})$ $k \leftarrow 0$ $\hat{Q}_0(\cdot, \cdot) \leftarrow c(\cdot, \cdot)$ **repeat** $k \leftarrow k + 1$ In order to obtain the values of $\hat{Q}_k(\cdot, \cdot)$ for all $\{\mathbf{n}_l, \mathbf{a}_l\}$ in \mathcal{F} compute:

$$\hat{Q}_k(\mathbf{n}_l, \mathbf{a}_l) = c(\mathbf{n}_l, \mathbf{a}_l) + \gamma \min_{\mathbf{a} \in \mathbf{A}} \hat{Q}_{k-1}(\mathbf{n}_l^+, \mathbf{a}) \quad (5)$$

Estimate the function $\hat{Q}_k(\mathbf{n}, \mathbf{a})$ using a regression algorithm with input pairs $(\mathbf{n}_l, \mathbf{a}_l)$ and function values $\hat{Q}_k(\mathbf{n}_l, \mathbf{a}_l)$.**until** the stopping criterion is satisfiedCompute $\hat{\pi}^*(\mathbf{n}) = \operatorname{argmin}_{\mathbf{a} \in \mathbf{A}} \hat{Q}_k(\mathbf{n}, \mathbf{a})$

where $V(\mathbf{n}_t)$ is called the value function and $\pi(\cdot)$ is a mapping from \mathbf{n} to \mathbf{a} , which is called the control policy. The cost function c specifies the objective of the control problem, which in our case is driving the system to a specific region in the state-space. In our setting, the control policy should be inferred based only on realisations of one-step system transitions $\{\mathbf{n}_l, \mathbf{a}_l, \mathbf{n}_l^+\}$, where \mathbf{n}_l^+ is a successor state of the system in the state \mathbf{n}_l and subjected to the input \mathbf{a}_l (in the deterministic case, if the function $f(\cdot, \cdot)$ is known \mathbf{n}_l^+ is equal to $f(\mathbf{n}_l, \mathbf{a}_l)$). For the purpose of this paper, the function $c(\cdot, \cdot)$ is assumed to be known in advance.

2.3 Fitted Q Iteration

A central object of the fitted Q algorithm is the Q function, which is introduced as follows:

$$Q(\mathbf{n}_t, \mathbf{a}_t) = c(\mathbf{n}_t, \mathbf{a}_t) + \min_{\pi(\cdot)} \sum_{i=t}^{\infty} \gamma^{i-t} c(\mathbf{n}_i, \pi(\mathbf{n}_i))$$

Once a Q function is computed, the optimal feedback control policy is given as:

$$\pi^*(\mathbf{n}) = \operatorname{argmin}_{\mathbf{a} \in \mathbf{A}} Q(\mathbf{n}, \mathbf{a})$$

Under certain conditions, the Q function can be obtained as the unique solution of the following iterative procedure:

$$Q_k(\mathbf{n}, \mathbf{a}) = c(\mathbf{n}, \mathbf{a}) + \gamma \min_{\mathbf{a}' \in \mathbf{A}} Q_{k-1}(f(\mathbf{n}, \mathbf{a}), \mathbf{a}') \quad (4)$$

where Q_0 is equal to c . However, (4) is hard to solve in general, especially if only the triplets \mathcal{F} are given. Therefore an approximation \hat{Q} of the Q function is computed using an iterative procedure. Let $\hat{Q}_0 = c$ and for every $(\mathbf{n}_l, \mathbf{a}_l, \mathbf{n}_l^+)$ in \mathcal{F} compute:

$$\hat{Q}_1(\mathbf{n}_l, \mathbf{a}_l) = c(\mathbf{n}_l, \mathbf{a}_l) + \gamma \min_{\mathbf{a} \in \mathbf{A}} \hat{Q}_0(\mathbf{n}_l^+, \mathbf{a})$$

This expression gives \hat{Q}_1 only for $\mathbf{n}_l, \mathbf{a}_l$ in \mathcal{F} , while the entire function $\hat{Q}_1(\cdot, \cdot)$ is estimated by a regression algorithm (e.g., EXTRA Trees by Geurts *et al.* (2006)). This can be generalised to an iterative procedure, which can be used to obtain a near-optimal control policy as outlined in Algorithm 1. The stopping criterion can be simply the maximum number of iterations N_{it} , which is chosen such that the number $\gamma^{N_{\text{it}}}$ is sufficiently small and the values $\hat{Q}_k(\mathbf{n}_l, \mathbf{a}_l)$ are not modified significantly for k larger than N_{it} . Other criteria are described in (Ernst *et al.* (2005)). Note that Algorithm 1 can be extended to handle the stochastic case as well (Ernst *et al.* (2005)).

3 System Description and Problem Setting

3.1 Models

First, we briefly describe our benchmark problem - regulation of the toggle switch system (Gardner *et al.* (2000)). The original genetic toggle switch system consists of the *lacI* and *tetR* genes mutually repressing each other (see Figure 1). We consider a generic toggle switch model; therefore, we will use numeric references for genes and proteins, that is, gene 1 and 2 instead of *lacI* and *tetR* genes. We will refer to the protein products of genes 1 and 2 as proteins 1 and 2, respectively. We assume that for both genes the protein concentrations are given as readouts via fluorescent markers. We also assume that the control inputs are implemented as light pulses activating a photo-sensitive promoter controlling the expression of gene 1 (Shimizu-Sato *et al.* (2002)). When this photo-sensitive promoter is activated through a light pulse the concentration of protein 1 is increased by a small amount through the expression of gene 1.

Basic mass-action kinetics of the toggle switch result in a high-order model, which is typically reduced to a two state model using quasi steady state approximation (Guantes & Poyatos (2006)). This can be done because most of the reactions (including the mRNA dynamics and the light-induction of the promoter) occur on a fast time scale (order of seconds) in comparison with the gene expression time scale (order of minutes or even hours). The reduced order model of the toggle switch system has two states, which are the two protein concentrations:

$$\begin{aligned} n_{t+1}^1 &= \beta_1 + \frac{c_1}{1 + (n_t^2)^{\alpha_2}} - c_2 n_t^1 + b u_t \\ n_{t+1}^2 &= \beta_2 + \frac{c_3}{1 + (n_t^1)^{\alpha_1}} - c_4 n_t^2 \end{aligned} \quad (6)$$

where n_t^i is the concentration of protein i at time t , c_1 and c_2 are the effective rate of synthesis of the repressors, α_i is the cooperativity coefficient of the repressor i , c_2 and c_4 are the degradation rates of proteins, β_i models leaky transcription during the gene expression of the gene i , and b is the increase in protein concentration produced per unit of time as a result of one light pulse. We approximate the action of light induction u_t as a discrete variable in the set $U = \{0, 1\}$. A more realistic model would also have a time-delayed control action. Such an extension requires simple modifications of our control algorithm, but makes the results less transparent and harder to analyse. In our simulations we use a training model and a validation model.

Training model	Validation model
$n_{t+1}^1 = 0.1 + \frac{30}{1 + (n_t^2)^2} - n_t^1 + 20u_t$	$n_{t+1}^1 = 0.1 + \frac{60}{1 + (n_t^2)^2} - n_t^1 + 20u_t$
$n_{t+1}^2 = 0.1 + \frac{60}{1 + (n_t^1)^2} - n_t^2$	$n_{t+1}^2 = 0.1 + \frac{30}{1 + (n_t^1)^2} - n_t^2$

(7)

Both models are bi-stable toggle switches with quantitatively different behaviours. Moreover, the steady states of the validation model are relatively far from the steady states of the training model. The stable steady states of the training model are approximately at $se_1 = (0.11 \ 29.26)$ and $se_2 = (59.4 \ 0.17)$ concentration units, while the stable steady states of the validation model are at $se_1 = (59.4 \ 0.17)$ and $se_2 = (0.11 \ 29.26)$. Such a situation is possible in biological applications, for example, due to different cell behaviours within a population of cells. Moreover, even for a single cell, different experiments may produce values of parameters with a large variation.

3.2 Control Algorithm

Our goal is to develop a control algorithm, which learns how to near-optimally control the toggle switch system in a single experiment. Toggling the switch can be done experimentally in a couple of hours and the fastest measurement sampling is in the order of one minute. This gives at most 200 samples in a single trajectory. Learning a near-optimal control policy for toggling the switch with such limited amount of data is an extremely hard problem to solve. To tackle this issue, we propose to first learn a ‘‘rough approximation’’ of the control policy obtained by applying Algorithm 1 to one-step system transitions \mathcal{F} artificially generated from simulations of a mathematical model of a genetic toggle switch. Afterwards the policy

Algorithm 2 Online learning algorithm

Inputs: Set $\mathcal{F} = \{\mathbf{n}_l, \mathbf{a}_l, \mathbf{n}_l^+\}_{l=1}^{\#\mathcal{F}}$, cost function $c(\cdot, \cdot)$, function $\hat{Q}_{\text{Alg 1}}(\cdot, \cdot)$, number of iterations N , function $h(\cdot, \cdot)$

$\tilde{Q}_{\text{cur}}(\cdot, \cdot) \leftarrow \hat{Q}_{\text{Alg 1}}(\cdot, \cdot)$
 $\mathcal{F}_{\text{cur}} \leftarrow \mathcal{F}$

while new data is received **do**

$\pi(\mathbf{n}) = \min_{\mathbf{a}' \in \mathcal{A}} \tilde{Q}_{\text{cur}}(\mathbf{n}, \mathbf{a}')$

$k \leftarrow 1, i \leftarrow 1$

while $i \leq T_{\text{update}}$ **do**

compute $\mathbf{a}_i = \pi(\mathbf{n}_i)$

observe the successor state \mathbf{n}_{i+1} for the state-action pair $(\mathbf{n}_i, \mathbf{a}_i)$.

$i \leftarrow i + 1$

end while

Collect a set of new samples $\mathcal{F}_{\text{new}} = \{\mathbf{n}_m, \mathbf{a}_m, \mathbf{n}_m^+\}_{m=1}^{T_{\text{update}}}$

$\mathcal{F}_{\text{cur}} \leftarrow h(\mathcal{F}_{\text{cur}}, \mathcal{F}_{\text{new}})$

while $k \leq N$ **do**

In order to obtain the values of $\tilde{Q}_{k+1}(\cdot, \cdot)$ for all $\{\mathbf{n}_l, \mathbf{a}_l\}$ in \mathcal{F}_{cur} compute:

$$\tilde{Q}_{k+1}(\mathbf{n}, \mathbf{a}) = c(\mathbf{n}, \mathbf{a}) + \gamma \min_{\mathbf{a}' \in \mathcal{A}} \tilde{Q}_k(\mathbf{n}^+, \mathbf{a}') \quad (8)$$

Estimate the function $\tilde{Q}_{k+1}(\mathbf{n}, \mathbf{a})$ using a regression algorithm with input pairs $(\mathbf{n}_l, \mathbf{a}_l)$ and function values $\tilde{Q}_{k+1}(\mathbf{n}_l, \mathbf{a}_l)$.

$k \leftarrow k + 1$

end while

$\tilde{Q}_{\text{cur}}(\cdot, \cdot) \leftarrow \tilde{Q}_k(\cdot, \cdot)$

end while

is fine-tuned by mixing the online measurements with past observations \mathcal{F} . The Exploration/Exploitation trade-off is addressed using an ε -greedy policy.

Our approach is outlined in Algorithm 2. Let $\hat{Q}_{\text{Alg 1}}(\cdot, \cdot)$ be the approximation of the Q function obtained by Algorithm 1. This will be the initial Q function denoted as \tilde{Q}_{cur} . Then we assume that T_{update} direct interactions with the real system are performed by computing actions using \tilde{Q}_{cur} and new input-output samples are collected in \mathcal{F}_{new} . Given the set \mathcal{F}_{new} , the approximation \tilde{Q}_{cur} of the Q function is updated as prescribed in Algorithm 2. After that the new set \mathcal{F}_{new} is formed and new samples are collected.

The major challenge of Algorithm 2 is appropriately choosing the function $h(\cdot, \cdot)$, which combines the sets \mathcal{F}_{cur} and \mathcal{F}_{new} . As an example, we consider $h(\mathcal{F}_{\text{cur}}, \mathcal{F}_{\text{new}}) = \mathcal{F}_{\text{cur}} \cup \mathcal{F}_{\text{new}}$. Such a choice has some drawbacks. If the initial set \mathcal{F} contains many samples, then the updates in (8) will not result in significant changes in the policy. This happens because the algorithm appreciates equally the samples in \mathcal{F} and the new sets \mathcal{F}_{new} , even though the samples in \mathcal{F} are artificially generated using a mathematical model and the samples in \mathcal{F}_{new} are obtained from the real system.

An important task of such a learning algorithm is a trade-off between exploration and exploitation during the generation of the set \mathcal{F}_{new} . Exploration is required, since the real system is essentially unknown to the algorithm and the exploratory actions will provide new information. The trade-off policy between exploration and exploitation is defined as follows:

$$\mathbf{a}_t = \begin{cases} \operatorname{argmin}_{\mathbf{a}' \in \mathcal{A}} \tilde{Q}_k(\mathbf{n}_t, \mathbf{a}') & \text{with probability } 1 - \varepsilon_t \\ \text{random action} & \text{with probability } \varepsilon_t \end{cases}$$

where \mathbf{n}_t is the state measured at time t and $\tilde{Q}_k(\cdot, \cdot)$ is a current approximation of the Q function. In our experiment ε_t is a decreasing function of t between zero and one. During the first time samples, the need for new information is typically higher, and thus a high value of ε_t should be chosen.

3.3 Parameters of the Algorithm

The structure of the instantaneous cost $c(\mathbf{n}, \mathbf{a})$ is chosen as follows:

$$c(n^1, n^2, u) = \max(n^1/\alpha_1, n^2/\alpha_2) - \min(n^1/\alpha_1, n^2/\alpha_2) + \alpha_u u$$

where $\alpha_1, \alpha_2, \alpha_u$ are non-negative constants. The function

$$\max(n^1/\alpha_1, n^2/\alpha_2) - \min(n^1/\alpha_1, n^2/\alpha_2)$$

appears in studies on consensus theory as a Tsitsiklis Lyapunov function. The vector $(\alpha_1 \ \alpha_2)$ can be seen as the target point of the control algorithm and the function itself can be viewed as a metric. Since only the ratio between the protein concentrations and the constants α_i appears in the cost, the algorithm is robust towards changes in α_i , which are within one order of magnitude of α_i . The major requirement is that α_1 is much larger than α_2 , which forces the protein concentration n^1 to be much larger than the protein concentration n^2 . Note that instead of a Tsitsiklis Lyapunov function other functions can be used, for example, a distance in l_p , a linear Lyapunov function $n^1/\alpha_1 + n^2/\alpha_2$ etc. However, the main concern of this work is evaluating the performance of the online algorithm; therefore, the choice of the cost function will be addressed in future work. The term $\alpha_u u$ penalises the control signal and therefore attempts to minimise the burden associated with light-induced gene expression. The choice of α_u dictates the trade-off that exists between toggling the switch fast and toggling the switch with a reduced gene expression burden. We choose parameters α_u and γ by tuning. The parameter α_u is equal to one in the simulations, and the discount factor γ is set to 0.75.

Computing the control actions is a cheap procedure; however, performing the updates of the Q function is a computationally harder problem. Therefore, the online algorithm performs 10 iterations of the fitted Q algorithm every 10 time samples, in order to emulate computationally constrained controllers. The number of input-output samples used for computing the initial policy for the online algorithm is small in comparison with the purely offline algorithm. There are two reasons for such an assumption: (a) it is more realistic to assume sparse input data, if we consider input-output data from previous experiments; (b) fewer input-output samples imply computationally cheaper updates of the Q function; (c) a large amount of samples can limit the ability of the online algorithm to update the policy and the Q function efficiently.

For the results in Figures 2 and 3(a), we generated 1000 trajectories with 100 one-step transitions in each trajectory. For the simulation of the online update algorithm we generated 100 trajectories with 100 one-step transitions in each trajectories. The policy is updated every 10 time samples. The stochastic simulation is performed using the direct Gillespie stochastic simulation algorithm. At every time instance t , one hundred trajectories starting at \mathbf{n}_t are computed until the next time instance $t + 1$, and the value \mathbf{n}_{t+1} is then averaged over these trajectories. The average over these trajectories represents the average value of protein concentrations in a population of cells, which is much easier to measure.

Finally, the trade-off between exploration and exploitation is decided by choosing the ε_t function as follows:

$$\varepsilon_t = \varepsilon \cdot \frac{1}{N_{\text{update}} + 1}$$

where N_{update} is the number of times the policy was updated online. We update the policy after 10 time samples; therefore, N_{update} is equal to $O(t)$ for large t .

4 Results and Discussion

As an illustration of the benefits of the proposed approach, we investigate how it handles model uncertainty. In order to do so, we specify a training model and a validation model as in (7). Both models are bi-stable toggle switches with quantitatively different behaviours. Moreover, the steady states of the validation model are relatively far from the steady states of the training model. The stable steady states of the training model are approximately at $se_1 = (0.11 \ 29.26)$ and $se_2 = (59.4 \ 0.17)$ concentration units, while the stable steady states of the validation model are at $se_1 = (59.4 \ 0.17)$ and $se_2 = (0.11 \ 29.26)$. The goal is to compute a control policy (control law), which will steer the model from the stable steady state se_1 to the stable steady state se_2 . A control policy is a binary function of a current measurement computing the current action, which is the presence or the absence of a light pulse. Due to the systems' dynamics, a larger

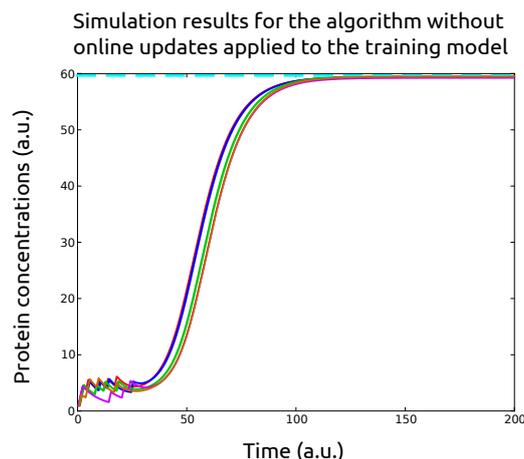


Figure 2: “Ideal” control of the toggle switch system. The policy is computed from the input-output data of a system and then applied to control the same system. This setting is unrealistic; however, it illustrates the robustness of our control objective towards errors in the choice of the expected target point. Differently coloured trajectories correspond to the different input target points for the steady state concentration of protein 1. We choose the target points with 15, 30, 45, 60, 90 units, while the real upregulated steady state is approximately 60 units (the cyan dashed line). These simulations show, that a considerable error can be made in the specification of the input target point without significant effect on the performance of the control algorithm. The schedule of light pulses is not shown due to overlapping trajectories, but all the pulses occur when the concentration of protein 1 is smaller than 10 units.

amount of light pulses is typically required to switch from se_1 to se_2 in the validation model in comparison with the training model. We are going to test our online control algorithm by computing the initial control policy from the data generated by the training model and apply this policy to the validation model. One of the challenges for an efficient control algorithm is that not only the dynamics change, but also the target steady state. This setting mimics the experimental setup, when the trajectories of the model used for the policy computation (or the training model) do not match exactly the trajectories of the real system (or the validation model).

We first consider the deterministic case and study the robustness of our algorithm towards errors in the choice of the target steady state, that is, the presumed and a priori specified value of the upregulated steady state concentration of protein 1. Therefore, we evaluate the proximity of five trajectories obtained with different target points, while the control policies are learned from and applied to the training model. In this case, we use the control algorithm without online updates. Figure 2 depicts the obtained trajectories of the concentration of protein 1 associated with the gene being upregulated. The red curve corresponds to the protein concentration obtained with the input target point equal to 15 units, the green curve corresponds to the input target point of 30 units, the blue curve to the input target point 45, the orange curve to 60, and the purple curve to 90. The upregulated steady state concentration of protein 1 is approximately equal to 60 units (the cyan dashed line). All the curves are very close to each other and hardly distinguishable, which indicates that our algorithm is robust to some perturbations in the choice of the target point. Note that all the light pulses occur when the protein concentration is smaller than 10 concentration units. Hence, the policy essentially defines a threshold in the concentration of protein 1, below which light pulses are applied and above which light pulses are not necessary since the trajectories will eventually converge to the upregulated state due to the unforced system dynamics¹. This threshold can be adjusted by modifying the parameters of the algorithm according to the control goal: faster control or smaller burden.

However, the setting when the policy is learned from a system and then used to control the same system is not entirely realistic. Typically, some model parameter variations are present. Here we model the case when the validation system has a considerable difference in parameter values in comparison with the training system. In Figure 3(a), we depict the simulated trajectories in such a situation. We run the algorithm with three different target points: 15 (the red curve), 30 (the green curve), and 45 (the blue curve). The actual

¹The optimal policy is more complicated than a simple threshold; however, the approximation of the policy by a threshold provides a general idea about the shape of the control policy

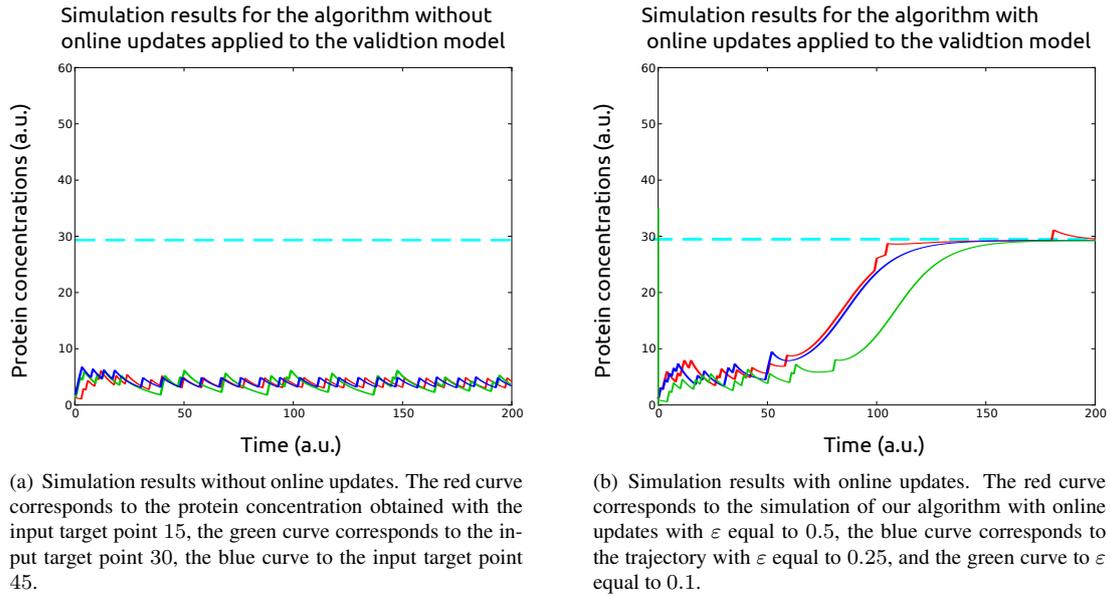


Figure 3: Simulation results of the validation model with an initial control policy computed using the input-output data generated by the training model. In the left panel no online updates are performed, in the right panel the policy is updated using the measured input-output data. In all the simulations of the algorithm without online updates the switch is not toggled (the left panel), while in the simulations of the algorithm with online updates the switch is successfully toggled (the right panel). In both panels, the upregulated steady concentration of protein 1 is approximately equal to 30 (the cyan dashed line). The algorithm with online updates has two phases: exploration and exploitation. The trade-off between these two phases is decided by the parameter ε , larger values of which imply more aggressive exploration and faster learning of the system.

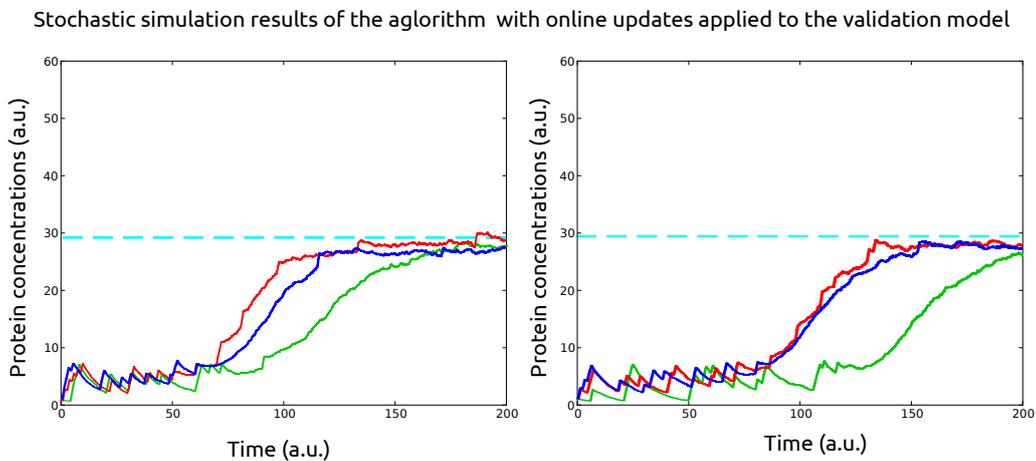


Figure 4: Stochastic simulation results of the algorithm with online updates applied to the validation model. In the simulations in the right panel, the penalty on the amount of light pulses is twice larger than in the simulations in the left panel. The red lines correspond to the value of ε equal to 0.5, the blue lines corresponds to the value of ε equal to 0.25, and the green lines to ε equal to 0.1.

upregulated protein concentration is approximately equal to 30. In all the simulations the algorithm without online updates cannot force the system into the upregulated state (the cyan dashed line). This occurs because the threshold required to ensure the switch in the validation model is higher than the one computed using the training model.

In Figure 3(b), we present simulation results of the proposed algorithm with online updates. The algo-

rithm collects new input-output samples and updates the policy at certain time intervals. The algorithm alternates between two phases: exploration and exploitation. In the exploitation phase the algorithm steers the system towards the specified goal using the control policy computed so far. In the exploration phase the algorithm generates data by randomly choosing “to apply a light pulse” or “do nothing”. Due to these random choices, data generated during the exploration phase is not correlated with the past samples. A major challenge in this algorithm is deciding the trade-off between the exploitation and exploration phases. A simple heuristic for tackling this trade-off is as follows. At time t , with a probability ε_t explore the system, and with a probability $1 - \varepsilon_t$ exploit the system. There is a bigger need in exploration in the beginning of the experiment; therefore, ε_t should be larger for small t and decrease with time. Hence, we choose ε_t as $\varepsilon \cdot \sigma(t)$, where ε is a positive constant smaller than one and $\sigma(t)$ is a monotonically decreasing function of t such that $\varepsilon \cdot \sigma(t)$ is always larger than zero and smaller than one. Larger values of ε indicate more aggressive exploration and faster learning. In these simulations, the input target point for the concentration of protein 1 is equal to 15. In Figure 3(b), the red curve corresponds to the simulation of our algorithm with online updates with ε equal to 0.5, the blue curve corresponds to the trajectory with ε equal to 0.25, and the green curve to ε equal to 0.1. The cyan dashed line represents the upregulated steady state concentration of protein 1. In all the simulations the switch is successfully toggled for the validation model, even if the initial policy is obtained by learning from the training model.

One of the biggest advantages of our approach is the ability to handle stochastic dynamics without any modifications of the algorithm. Moreover, behaviours of the controlled toggle switches in the stochastic case are qualitatively similar to the deterministic case. We present the simulation results with online updates for a similar setting as in the deterministic case in the left panel of Figure 4. Additionally, we present the simulation results in the setting with a twice as large penalty on the amount of applied light pulses in the right panel of Figure 4. In both figures, the red curves correspond to the simulation of our online algorithm with ε equal to 0.5, the blue curves correspond to ε equal to 0.25, and the green curves to ε equal to 0.1. It is noticeable that toggling the switch takes longer with a larger penalty on the amount of light pulses. However, the main outcome of these simulations is that our algorithm can be applied to systems with stochastic dynamics, and as a consequence can potentially handle wet-lab data efficiently.

Our algorithm, however, does not take into account the *a priori* knowledge that it is being applied to a different, but structurally similar system. A correct exploitation of structural similarity between the learned from and applied to systems may significantly improve the performance of the presented algorithm. This constitutes one of the main directions for future work that is currently under investigation.

As a final remark, we have shown that the presented framework can efficiently control a (stochastic) model of the genetic toggle switch with a parametric uncertainty. The major feature of our control algorithm is its learning nature. The algorithm computes an initial control policy using input-output data obtained from simulations of a training model, and after that updates the policy by using the input-output data obtained from the validation model (or the real system). In the presented example, despite the fact that the training and validation models had quite different quantitative behaviours the control objective was always reached using our online control method. This indicates a potential for a generalisation of this data-based control method to more complex gene regulatory networks.

Acknowledgement

Aivar Sootla and Guy-Bart Stan acknowledge the support of EPSRC through the project EP/J014214/1 and the EPSRC Science and Innovation Award EP/G036004/1. Damien Ernst acknowledges support of the Belgian Network DYSCO (Dynamical Systems, Control, and Optimization), funded by the Interuniversity Attraction Poles Programme, initiated by the Belgian State, Science Policy Office.

References

- BENNETT M. R. & HASTY J. (2009). Microfluidic devices for measuring gene network dynamics in single cells. *Nat Rev Genet*, **10**(9), 628–638.
- BUŞONIU L., BABUŠKA R., DE SCHUTTER B. & ERNST D. (2010). *Reinforcement Learning and Dynamic Programming Using Function Approximators*. CRC Pr I Llc.
- CAI L., FRIEDMAN N. & XIE X. S. (2006). Stochastic protein expression in individual cells at the single molecule level. *Nature*, **440**, 358–362.

- CASTRONOVO M., MAES F., FONTENEAU R. & ERNST D. (2012). Learning exploration/exploitation strategies for single trajectory reinforcement learning. In *Proc. Eur. Workshop Reinforcement Learn.*
- ERNST D., GEURTS P. & WEHENKEL L. (2005). Tree-based batch mode reinforcement learning. *J Mach Learn Res*, **6**, 503–556.
- GARDNER T., CANTOR C. R. & COLLINS J. J. (2000). Construction of a genetic toggle switch in *Escherichia coli*. *Nature*, **403**, 339–342.
- GEURTS P., ERNST D. & WEHENKEL L. (2006). Extremely randomized trees. *Machine Learning*, **63**(1), 3–42.
- GILLESPIE D. (1977). Exact stochastic simulation of coupled chemical reactions. *The journal of physical chemistry*, **81**(25), 2340–2361.
- GUANTES R. & POYATOS J. F. (2006). Dynamical principles of two-component genetic oscillators. *PLoS Comput Biol*, **2**(3), e30.
- IVANOVA N., DOBRIN R., LU R., KOTENKO I., LEVORSE J., DECOSTE C., SCHAFER X., LUN Y. & LEMISCHKA I. R. (2006). Dissecting self-renewal in stem cells with rna interference. *Nature*, **442**, 533–538.
- LEVSKAYA A., WEINER O. D., LIM W. A. & VOIGT C. A. (2009). Spatiotemporal control of cell signalling using a light-switchable protein interaction. *Nature*, **461**, 997–1001.
- LIU Y., ASAKURA M., INOUE H., NAKAMURA T., SANO M., NIU Z., CHEN M., SCHWARTZ R. J. & SCHNEIDER M. D. (2007). Sox17 is essential for the specification of cardiac mesoderm in embryonic stem cells. *Proc Natl Acad Sci USA*, **104**(10), 3859–3864.
- MENOLASCINA F., DI BERNARDO M. & DI BERNARDO D. (2011). Analysis, design and implementation of a novel scheme for in-vivo control of synthetic gene regulatory networks. *Automatica, Special Issue on Systems Biology*, **47**(6), 1265–1270.
- METTETAL J. T., MUZZEY D., GOMEZ-URIBE C. & VAN OUDENAARDEN A. (2008). The Frequency Dependence of Osmo-Adaptation in *Saccharomyces cerevisiae*. *Science*, **319**(5862), 482–484.
- MILIAS-ARGEITIS A., SUMMERS S., STEWART-ORNSTEIN J., ZULETA I., PINCUS D., EL-SAMAD H., KHAMMASH M. & LYGEROS J. (2011). In silico feedback for in vivo regulation of a gene expression circuit. *Nature biotechnology*.
- PURNICK P. E. M. & WEISS R. (2009). The second wave of synthetic biology: from modules to systems. *Nat. Rev. Mol. Cell Biol.*, **10**(6), 410–422.
- SHIMIZU-SATO S., HUQ E., TEPPERMAN J. M. & QUAIL P. H. (2002). A light-switchable gene promoter system. *Nat Biotech*, **20**(10), 1041–1044.
- SOOTLA A., STRELKOWA N., ERNST D., BARAHONA M. & STAN G.-B. (2013). On reference tracking using reinforcement learning with application to gene regulatory networks. In *Conf. Decision Control*, p. 4086–4091, Florence, Italy.
- SUTTON R. & BARTO A. (1998). *Reinforcement Learning, an Introduction*. MIT Press.
- SWAIN P. S., ELOWITZ M. B. & SIGGIA E. D. (2002). Intrinsic and extrinsic contributions to stochasticity in gene expression. *Proc Natl Acad Sci USA*, **99**(20), 12795–12800.
- TAN C., MARGUET P. & YOU L. (2009). Emergent bistability by a growth-modulating positive feedback circuit. *Nat. Chem. Biol.*, **5**(11), 842–848.
- UHLENDORF J., MIERMONT A., DELAVEAU T., CHARVIN G., FAGES F., BOTTANI S., BATT G. & HERSEN P. (2012). Long-term model predictive control of gene expression at the population and single-cell levels. *Proceedings of the National Academy of Sciences*, **109**(35), 14271–14276.