

RESEARCH ARTICLE

 OPEN ACCESS

Investigating the Effect of Changes in Model Parameters on Optimal Control Policies, Time to Absorption, and Mixing Times

Kathleen Johnson,^a Daniel Plaughner,^b David Murrugarra^a^aDepartment of Mathematics, University of Kentucky, Lexington, KY; ^bDepartment of Toxicology and Cancer Biology, University of Kentucky, Lexington, KY**ABSTRACT**

Many processes in biology and medicine have been modeled using Markov decision processes (MDP) which provides a rich algorithmic theory for model analysis and optimal control. We focus on the class of MDP that is obtained by considering stochastic Boolean networks equipped with control actions. We study the effect of changes in model parameters on optimal control policies. We conducted a sensitivity analysis on optimal control policies for a model of the T-cell large granular lymphocyte. We quantified how the choice of propensity parameters impacts the effectiveness of the optimal policy and then we provide thresholds at which the effectiveness is guaranteed. We also examined the effect on the optimal control policies of the level of noise that is usually added for simulations. Finally, we studied the effect on changes in the propensity parameters on the time to absorption and the mixing time for a Boolean model of the Repressilator.

ARTICLE HISTORYReceived January 23, 2023
Accepted December 10, 2023**KEYWORDS**Boolean networks,
Markov decision processes,
Optimal control,
Time to absorption,
Mixing time

1 Introduction

Boolean network (BN) models are becoming increasingly popular and have been successfully used to model important biological systems such as the bistable behavior of the *lac operon* (Veliz-Cuba and Stigler, 2011), the yeast-to-hyphal transition of the yeast *Candida albicans* (Wooten et al., 2021), the interaction of pancreatic cancer cells with their microenvironment (Plaughner and Murrugarra, 2021; Plaughner et al., 2022), and they have been proposed as appropriate models for understanding the evolution of cellular differentiation (Márquez-Zacarías et al., 2021). Furthermore, there is a growing collection of BN models in the *Cell Collective* database (Helikar et al., 2012) covering processes such as signaling, regulation, and cancer.

To account for the stochasticity in various molecular processes, several versions of stochastic Boolean networks exist including Probabilistic Boolean networks (PBN) (Shmulevich et al., 2002; Shmulevich and Dougherty, 2010), perturbed Boolean networks (BNp) (Golinelli and Derrida, 1989; Trairatphisan et al., 2013), and Probabilistic Edge Weights operators (PEW) (Derritei et al., 2022). This paper will focus on the stochastic framework that was introduced by Murrugarra et al. (2012) referred here as Stochastic Discrete Dynamical Systems (SDDS). The stochasticity usually originates from the updating schedule or from an added perturbation to the functions. Standard updating schedules include synchronous update where all nodes are updated at the same time, and asynchronous update, where a random set of nodes are updated at each time step. Consequently, a synchronous update produces deterministic dynamics, while an asynchronous update generates stochastic dynamics. The SDDS framework introduces stochasticity by assigning two propensity parameters to each function in the Boolean network. The PBN framework considers multiple functions for each node where one function is selected at each step using a probability distribution.

A Markov decision process (MDP) is a mathematical modeling framework for decision making (Sutton and Barto, 1998; Bertsekas, 2019). An optimal control policy for an MDP provides an action for each state of the system such that a certain optimality criterion is achieved. Many methods exist for obtaining optimal control policies such as the infinite-horizon method with a discount factor (Yousefi et al., 2012; Aguilar et al., 2020). The optimal control policies are usually obtained using the value iteration method for approximating the solution of Bellman's equation (Bertsekas, 2005; Murrugarra and Aguilar, 2019).

The main contributions of this paper will be presented in the [Results section](#), which is divided in two parts. In the [first part](#), we study the effects of changes in the propensity parameters of the SDDS framework on the optimal control policies for a Boolean model of the T-cell large granular lymphocyte, or *T-LGL* ([Saadatpour et al., 2011](#)). For this model, we show how the choice of propensity parameters impacts the effectiveness of the optimal policy and then we provide thresholds for which the effectiveness is guaranteed. We also examined the effect of the level of noise that is usually added for simulations on the optimal control policies. In the [second part of the Results section](#), we discuss the effects of the propensities on the long-term dynamics of a Boolean model of the Repressilator. For this model, we studied the effect of changes in the propensity parameters on the time to absorption and the mixing time.

This paper is structured in the following way. In [Section 2](#), we introduce the modeling framework SDDS and their associated Markov decision processes. In [Section 3.1](#) we present the results of the effects of changes in the propensities on control policies for the T-LGL model. In [Section 3.2](#), we discuss the long term dynamics effects of propensities on a Boolean model of the Repressilator. In [Section 4](#) we discuss our main findings and provide some conclusions.

2 Methods

In this section we describe the stochastic modeling framework SDDS and the optimal control method that we will use in this paper.

2.1 Stochastic discrete dynamical systems

Consider n variables x_1, \dots, x_n that can take values in finite sets X_1, \dots, X_n . Let $X = X_1 \times \dots \times X_n$ be the Cartesian product. An SDDS in the variables x_1, \dots, x_n is defined as a collection of n triples

$$\hat{F} = \left\{ f_k, p_k^\uparrow, p_k^\downarrow \right\}_{k=1}^n, \quad (1)$$

where for $k = 1, \dots, n$,

- $f_k: \{0, 1\}^n \rightarrow \{0, 1\}$ is the update function for x_k ,
- $p_k^\uparrow \in [0, 1]$ is the activation propensity,
- $p_k^\downarrow \in [0, 1]$ is the degradation propensity.

The activation parameter p_k^\uparrow represents the probability that an activation of x_k occurs, given that f_k would activate x_k . That is, if $x_k(t) = 0$ and $f_k(x_1(t), \dots, x_n(t)) = 1$, then $x_k(t+1) = 1$ with probability p_k^\uparrow . The degradation propensity p_k^\downarrow is defined similarly.

We can represent an SDDS as a Markov chain with the transition probability from $x = (x_1, \dots, x_n)$ to $y = (y_1, \dots, y_n)$ is given by

$$a_{x,y} = \prod_{i=1}^n \text{Prob}(x_i \rightarrow y_i), \quad (2)$$

where $\text{Prob}(x_i \rightarrow f_i(x))$ is the probability that x_i will change its value and is given by

$$\text{Prob}(x_i \rightarrow f_i(x)) = \begin{cases} p_k^\uparrow, & \text{if } x_i < f_i(x), \\ p_k^\downarrow, & \text{if } x_i > f_i(x), \\ 1, & \text{if } x_i = f_i(x). \end{cases} \quad (3)$$

The probability that x_i will maintain its current value is given by

$$\text{Prob}(x_i \rightarrow x_i) = \begin{cases} 1 - p_k^\uparrow, & \text{if } x_i < f_i(x), \\ 1 - p_k^\downarrow, & \text{if } x_i > f_i(x), \\ 1, & \text{if } x_i = f_i(x), \end{cases} \quad (4)$$

for all $i = 1, \dots, n$. Note that $\text{Prob}(x_i \rightarrow y_i) = 0$ for all $y_i \notin \{x_i, f_i(x)\}$. The transition matrix is

$$A = (a_{x,y})_{x,y \in S}. \quad (5)$$

2.2 Boolean networks

Boolean networks are discrete dynamical systems. Specifically, consider n variables x_1, \dots, x_n each of which can take values in $\{0, 1\}$. Then, a synchronously updated Boolean network is a function $F = (f_1, \dots, f_n): \{0, 1\}^n \rightarrow \{0, 1\}^n$, where each coordinate function is a discrete function on a subset of $\{x_1, \dots, x_n\}$, which represents how the future value of the i th variable depends on the present values of the variables. We note that setting all propensity parameters p_k^\uparrow and p_k^\downarrow to 1.0 in the SDDS framework will result in a synchronous Boolean network.

2.3 Markov chains

A discrete-time Markov chain is a sequence of random variables $\mathcal{X}_1, \dots, \mathcal{X}_y, \dots$ with the Markov property: $\text{Prob}(\mathcal{X}_{y+1} = x \mid \mathcal{X}_1 = x_1, \dots, \mathcal{X}_y = x_n) = \text{Prob}(\mathcal{X}_{y+1} = x \mid \mathcal{X}_y = x_n)$. That is, the probability of moving to the next state depends only on the current state and not on the previous states. We note that Boolean networks are Markov chains.

2.4 Markov decision processes

Here we will describe a Markov Decision Process associated to the SDDS framework (Murrugarra and Aguilar, 2019). To use tools from Markov Decision Processes with SDDS, we consider the underlying Boolean network F and corresponding wiring diagram G associated with the SDDS \hat{F} . The wiring diagram G has nodes x_1, \dots, x_n , and there is a directed edge from x_i to x_j if f_j depends on x_i . The edge and node controls can be encoded in the Boolean network $F = (f_1, \dots, f_n): \mathbb{F}_2^n \rightarrow \mathbb{F}_2^n$.

Let U be the set of all possible control inputs, where 0 represents no control. We encode the edge and node controls using the function $\mathcal{F}: \mathbb{F}_2^n \times U \rightarrow \mathbb{F}_2^n$ such that $\mathcal{F}(x, 0) = F(x)$ returns the original Boolean network. To include control in an SDDS, we replace the functions of f_k with $\mathcal{F}_k: \{0, 1\}^n \times U \rightarrow \{0, 1\}$. With edge control, $u_{i,j}$ refers to control of the edge $x_i \rightarrow x_j$. For node control, u_i refers to node x_i . In both edge and node control, setting u to zero represents an inactive control and setting u to one represents removal of the corresponding edge or node.

A control action a is a combination of control edges and nodes being applied to the Boolean network simultaneously at a given time step. A set E of control edges and set V of control nodes can be used to define a as an array of binary elements of size $|U| = |E| + |V|$. Let $A = \{0, 1\}^{|U|}$ be the set of all possible actions.

Now, we define an MDP for the SDDS \hat{F} using these control actions. A Markov decision process for SDDS has four components: the set of states S , the set of control actions A , the transition probability $P_{x,y}^a$ associated with action a at state x that will lead to state y , and the associated costs $C(x, a, y)$.

The new SDDS resulting from the application of action a is $\hat{F}_a' = \{\mathcal{F}_k^a, p_k^\uparrow, p_k^\downarrow\}_{k=1}^n$. The transition probability $P_{x,y}^a$ from state x to y associated with action a is computed with equation 2 where f_k is replaced by \mathcal{F}_k^a . The cost function $C(x, a, y)$ describes the penalties for going from state x to y using action a . Thus the cost is given as the sum of the cost for the action and one for the state:

$$C(x, a, y) = C_a + C_y. \quad (6)$$

The action cost $C_a = c_v N_v + c_e N_e$ where c_v and c_e are the penalties for the control nodes and edges respectively, and N_v and N_e are the number of applied control nodes and edges in the given action a . The state cost

$$C_y = \sum_{k=1}^N w_k |y_k - s_k^*|, \quad (7)$$

where s^* is the desired state and w_k are the weights of each edge.

2.5 Optimal control for SDDS

This section reviews a method for optimal control for Markov decision processes called the infinite-horizon method with a discounting factor as presented by Murrugarra and Aguilar (2019).

A deterministic control policy π is a set $\pi = \{\pi_0, \pi_1, \dots\}$ where each π_t associates a state $x(t)$ to an action a at time step t . Given a state $x \in S$, a control policy π and a discounting factor $\gamma \in (0, 1)$, the cost function V^π associated with π is defined as

$$V^\pi(x) = \mathbf{E} \left[\sum_{t=0}^{\infty} \gamma^t C(x(t), a) \right], \quad (8)$$

where $C(x(t), a)$ represents the expected cost at time step t for executing the policy from state x . The cost at step t is defined by

$$C(x(t), a) = \mathbf{E} [C(x, a, y)] = \sum_{y \in S} P_{x,y}^a C(x, a, y). \quad (9)$$

The optimal cost function $V^*(x)$ associated with the optimal policy π^* is

$$V^*(x) = \min_{a \in \mathcal{A}} \left\{ C(x, a) + \gamma \mathbf{E} [V^*(y)] \right\}. \quad (10)$$

Because this function satisfies Bellman's equation (Sutton and Barto, 1998; Bertsekas, 2019), we can use the value iteration algorithm to calculate the optimal control policy. The equation for the value iteration algorithm is

$$V_t(x) = \min_{a \in \mathcal{A}} \left[C(x, a) + \gamma \sum_{y \in \mathcal{S}} P_{x,y}^a V_{t-1}(x) \right]. \quad (11)$$

For any initial bounded cost function V_0 , the recursion will converge to V^* .

2.6 Time to absorption

Consider the transition matrix of a Markov chain in canonical form:

$$A = \left(\begin{array}{c|c} Q & \mathbf{0} \\ \hline R & I \end{array} \right)$$

where $\mathbf{0}$ is the zero block matrix and I is the identity submatrix. The fundamental matrix N is defined as the inverse of $(I - Q)$. That is, $N = (I - Q)^{-1}$. The time to absorption for a transient state j is defined as the expected number of steps before absorption and can be calculated as the sum of the j th row of N (see Grinstead and Snell, 1997, Theorem 11.5).

2.7 Mixing time

The mixing time t_{mix} of a Markov chain is the number of steps required for the chain to be within a fixed threshold of its stationary distribution (Levin and Peres, 2017). More precisely, we define

$$t_{\text{mix}}(\varepsilon) = \min \{ t : d(t) \leq \varepsilon \}$$

where $d(t)$ is the total variation distance which is defined by $d(t) = \max_{x \in \mathcal{S}} \| P^t(x, \cdot) - \pi \|_{TV}$. The maximum in the previous formula is taken over all states in the state space \mathcal{S} , P represents the transition matrix, and π denotes the stationary distribution.

Several estimates and bounds exist for mixing times of Markov chains. For irreducible Markov chains, one can show that (see Levin and Peres, 2017, Chapter 12)

$$t_{\text{mix}}(\varepsilon) \log \left(\frac{1}{\lambda_*} \right) \geq \log \left(\frac{1}{2\varepsilon} \right) \quad (12)$$

where λ_* is the eigenvalue with the second highest absolute value. In this paper, we will use the following estimate which can be obtained from Equation 12:

$$t_{\text{mix}} \approx -1/\log(\lambda_*).$$

3 Results

In this section we provide results from our analysis using a Boolean model where optimal control policies have been calculated. We will apply the optimal control method described in Section 2.5 and study its effectiveness at different propensities and discount factors. We will also study the effect of changes in the level of noise for simulations.

3.1 T-LGL model

Here we consider the reduced model of the T-cell large granular lymphocyte, or *T-LGL* (Saadatpour et al., 2011) with nodes and functions described in Figure 1. This model has two fixed points: 000001 and 110000, where we represented a state of the system by $x_1 x_2 x_3 x_4 x_5 x_6$ for each $x_i \in \{0, 1\}$ for $i = 1, \dots, 6$. Here, a state of 0 means that the corresponding gene is inactive while the state 1 indicates that the corresponding gene is active. Thus, the state 000001 represents the desired state (with the apoptosis node ON, $x_6 = 1$) and 110000 represents the disease state (with the apoptosis node OFF, $x_6 = 0$). The optimal control problem here is that we want to obtain a policy to ensure the system transitions to the desired state. Note that this system has 64 possible states (since $2^6 = 64$) as indicated in the x -axis of Figure 2.

Given a Boolean network model, different methods for phenotype control can be used to identify edge and node controls, including those of Murrugarra et al. (2016); Zañudo and Albert (2015); Zañudo et al. (2017). For the *T-LGL* model discussed in

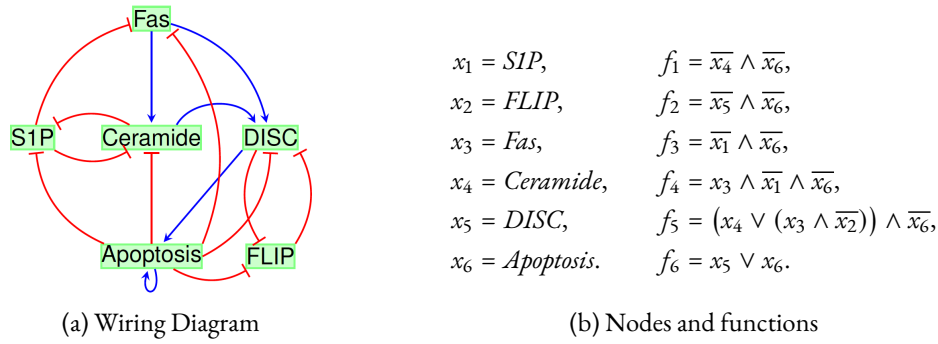


Figure 1: Reduced *T-LGL* network adapted from Saadatpour et al. (2011).




Figure 2: Heatmap of stationary distributions given all propensities are 0.9 using single-step optimal policies from the value iteration method. The states are on the *x*-axis and the *y*-axis gives the discount used to find the policy. At a discount of 0.8 or higher, the stationary distribution consistently goes to state 2 (000001), so this is when the policy becomes effective. The numbers (1 or 2) shown in each row correspond to the action specified by the policy for the corresponding state below the number. The two actions are: (1) no intervention and (2) knock-out of *S1P*. The shading gradient on the right of the heatmap correspond to the proportion of time the process ends in one of the possible states.




Figure 3: Heatmap of the stationary distribution given all propensities are 0.5 using single-step optimal policies from the value iteration method. The states are on the x -axis and the y -axis gives the discount used to find the policy. The policies for these choice of parameters are always ineffective.

this section, we use the approach of Murrugarra et al. (2016) to identify control nodes that can stabilize the system into one of its attractors. Specifically, we will use the control node $S1P = OFF$. This control was predicted to be 100% effective by Saadatpour et al. (2011). That is, the application of this control (i.e., setting the node $S1P$ to 0) will result in a controlled network with a single desired fixed point (i.e., 000001). Thus, at each state of the system, we will get access to two actions: (1) no intervention and (2) knock-out of $S1P$. Here we say that a policy is effective if this always leads the model to the fixed point 000001. In what follows, we study the effect of changes in the propensity parameters p_k^\uparrow and p_k^\downarrow of the SDDS framework on optimal control policies. Two commonly used set of values for the propensities are 0.9 and 0.5. Setting all propensity parameters p_k^\uparrow and p_k^\downarrow to 0.9 gives equal (90%) chance of update to each variable and therefore this model will highlight the synchronous transitions. Likewise, setting all propensities p_k^\uparrow and p_k^\downarrow to 0.5 gives equal (50%) chance of update to each variable and therefore this model will be closer to the fully asynchronous case.

In all heatmap figures we consider that state 2 (i.e., 000001) as the desired state while state 49 (i.e., 110000) as the undesired state. The numbers (1 or 2) shown in each row of the heatmap correspond to the action specified by the policy for the corresponding state below the number. The two actions are: (1) no intervention and (2) knock-out of $S1P$. The shading gradient on the right of the heatmaps correspond to the proportion of time the process ends in one of the possible states.

In our first result, shown in Figure 2, we have set all propensities p_k^\uparrow and p_k^\downarrow to 0.9. Allowing the discounting factor to vary, we see the policy becomes effective when the discount is at least 0.8. This coincides with a change in the action at the disease node 110000 (state 49 in the horizontal axis) that indicates that the intervention is active.

Setting the propensities p_k^\uparrow and p_k^\downarrow to 0.5 as in Figure 3, we see that the policy is never effective. Note that the action at state 110000 (state 49 in the horizontal axis) does not change to include the control node (i.e., action number 2) for any discounting factor.

For our third case, we considered a set of propensity parameters given by Murrugarra and Aguilar (2019, Table 5.3). (We reproduced these propensities here in Table 1.) These propensities were obtained from a genetic algorithm to maximize the probability of reaching state 110000. When the value iteration is applied to the model with these propensities as in Figure 4, we see that the policy is never effective. We also note that the action at state 110000 (state 49 in the horizontal axis) does not change to include both interventions for any discounting factor.

Next, we examined the effect of the noise in the results shown in Figures 2–4. In Figures 5–7 we considered noise levels of 1%, 5%, 10%, and 15% when calculating the optimal policies. In the first case, when we set all propensities p_k^\uparrow and p_k^\downarrow to 0.9,

	x_1	x_2	x_3	x_4	x_5	x_6
p_i^\uparrow	0.4792	0.4826	0.7098	0.3911	0.0549	0.0314
p_i^\downarrow	0.0274	0.9602	0.5195	0.7758	0.8373	1.0000

Table 1: Propensity values computed with a genetic algorithm by Murrugarra and Aguilar (2019).




Figure 4: Heatmap of the stationary distribution given propensities from Table 1 using single-step optimal policies from the value iteration method. The states are on the x -axis and the y -axis gives the discount used to find the policy. The policies for these choice of parameters are always ineffective.

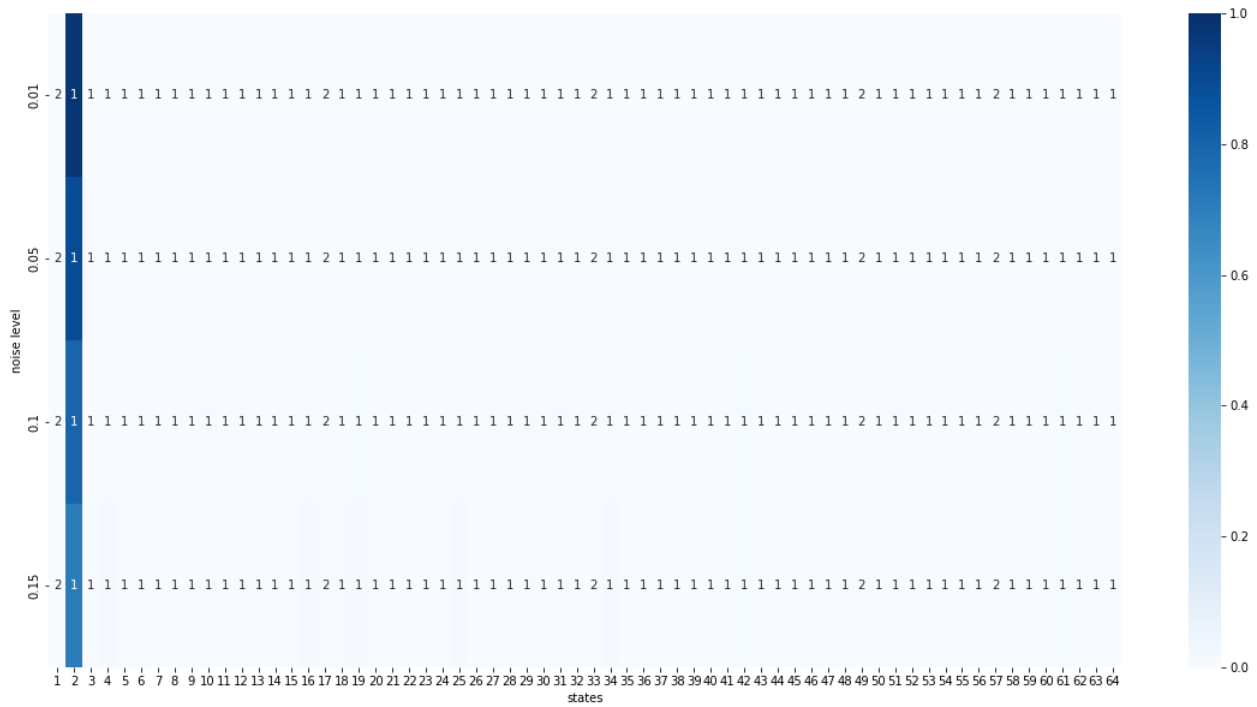


Figure 5: Heatmap of stationary distributions given all propensities are 0.9 using single-step optimal policies from the value iteration method. The states are on the x -axis and the y -axis gives the level of noise used to find the optimal policy. The discount factor was fixed to 0.9 for all policies, the stationary distribution consistently goes to state 2 (000001), so changing the noise in this case does not affect the effectiveness of the policies. The shading gradient on the right of the heatmap correspond to the proportion of time the process ends in one of the possible states.




Figure 6: Heatmap of the stationary distribution given all propensities are 0.5 using single-step optimal policies from the value iteration method. The states are on the x -axis and the y -axis gives the level of noise used to find the optimal policy. The discount factor was fixed to 0.9 for all policies. For noise level of 0.05 or smaller, the stationary distribution consistently goes to state 2 (000001), so small levels of noise the policies becomes effective.




Figure 7: Heatmap of the stationary distribution given propensities from Table 1 using single-step optimal policies from the value iteration method. The states are on the x -axis and the y -axis gives the level of noise used to find the optimal policy. The policies for these choice of parameters are always ineffective, so changing the noise in this case does not affect the ineffectiveness of the policies.

allowing the noise level to vary while maintaining the discount factor fixed at 0.9, we see that the policies are still effective as shown in Figure 5. Thus, changing the noise level in this case does not affect the effectiveness of the policies.

In the second case, when all propensities are set to 0.5, allowing the noise level to vary while maintaining the discount factor fixed at 0.9, we see that some policies become effective while other are still ineffective as shown in Figure 6. Thus, changing the noise level in this case does affect the effectiveness of the policies. For noise level of 0.05 or smaller, the stationary distribution consistently goes to state 2 (000001), so for small levels of noise the policies becomes effective.

Finally, when we used propensities from Table 1, allowing the noise level to vary while maintaining the discount factor fixed at 0.9, we see that the policies for these choice of parameters are always ineffective. Thus, changing the noise level in this case does not affect the ineffectiveness of the policies.

We also computed optimal policies that have a duration period of L steps that are followed by a recovery period (without intervention) of $L - W$ steps as described by Yousefi et al. (2012); Aguilar et al. (2020). In Supplemental Figures 1–12 (see Supplemental Materials) we show results for optimal policies with $L = W$ for $L = 2, 4, 6, 8$. We see that for the case of all propensities equal to 0.5, the policies become effective for a discount factor of 0.9 or higher. However, for the propensities in Table 1, the policies are still ineffective. For policies with four steps or higher, even the propensities in Table 1 will produce effective policies.

3.2 The Repressilator

The Repressilator is a gene regulatory network consisting of 3 genes, each expressing a protein that represses the next gene in the loop. Its wiring diagram is a cycle of three repressors as shown in Figure 8(a). This model has been studied using differential equations (Alon, 2019). The Repressilator has been used for the design and construction of synthetic networks to understand different cell functions (Elowitz and Leibler, 2000). Here we create a Boolean model for the Repressilator using three update functions given in Figure 8(b). Then we examine the effect the propensities have on its long term dynamics using the SDDS framework. The synchronous dynamics (i.e., without stochasticity) of this model is given by two cycles: a 2-cycle and a 6-cycle as shown in Figure 9(a). We note that the 6-cycle represents a biological attractor that describes the stable oscillations of the three repressors that was shown to exist with other model types (Elowitz and Leibler, 2000; Alon, 2019). The 2-cycle is an artificial attractor which a product of the synchronous update. Once we add stochasticity to the model using the SDDS framework, the



Figure 8: Boolean model of the Repressillator.

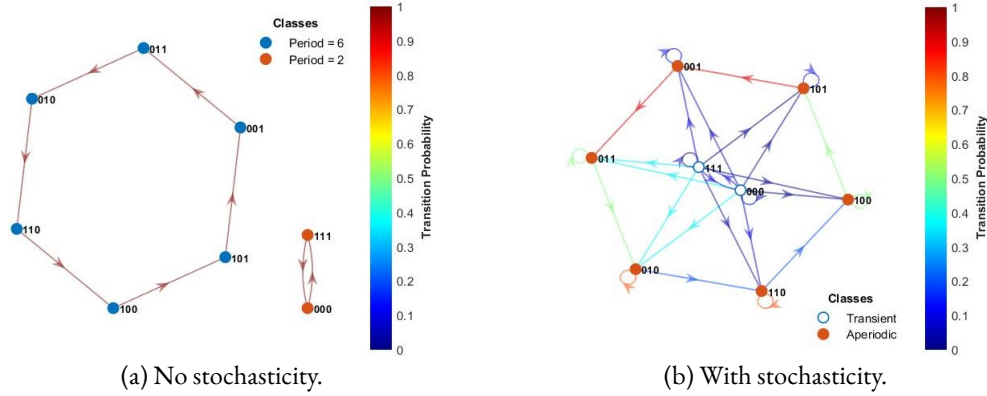


Figure 9: The synchronous dynamics are given by a 6-cycle and a 2-cycle in (a) without noise. With stochasticity using the SDDS framework, one may leave the 2-cycle for the 6-cycle, but not the other way around.

states in the 2-cycle became transient states (that is, it was possible to escape from the 2-cycle). However, the 6-cycle was still an attractor as shown in Figure 9(b). We highlight the fact that the biological attractor is robust to the update mode. Thus, we consider the 6-cycle to be the absorbing state for calculating time to absorption. We note that for this example we do not consider control actions and therefore this example is not a Markov decision process. However, we will study the effect of changes in the propensity parameters on the time to absorption and the mixing time of this model.

In Figure 9(b) we can see the effect the propensities have on the transition from one state to another. In this particular example the propensities are

$$\begin{pmatrix} 0.2 & 0.9 & 0.5 \\ 0.9 & 0.2 & 0.5 \end{pmatrix},$$

where the top row is the activation propensity for each function and the bottom row is the degradation propensity for each function. We note that in this selection of propensities we are specifying a slow activation rate for the variable x_1 and a fast degradation rate for variable x_1 . Inversely, we are specifying a fast activation rate for the variable x_2 and a slow degradation rate for variable x_2 . The mixing time is 4.5547 steps with the repressillator spending 6.85% of the time in each state 001 and 101, 12.33% of the time in each state 011 and 100, and 30.82% of the time in each state 010 and 110. Negligible time is spent in states 000 and 111. We want to investigate the range of mixing times for the Repressillator.

First, we examine the mixing time when both propensities for all variables are the same. In Figure 10, the top line in blue represents the mixing times when all propensities are the same. We find the lowest mixing time when all of the propensities are 0.5. The red line represents when we only change the propensities inversely for p_2^\downarrow and p_3^\uparrow while keeping all other propensities 0.5. In this case, we find that the mixing time is lowest at the ends of the graph. That is when $p_2^\downarrow = 0.9$ and $p_3^\uparrow = 0.1$ and vice versa, we achieve the lowest mixing time.

It is intuitive that when all of the propensities are low, the states in the Repressillator are less likely to change their current states. In particular, since the 2-cycle is less likely to change states, the time to absorption becomes larger which causes the mixing time to be higher as well—see Figure 11(a). A similar effect is seen with high propensities where the 000 and 111 alternate before being absorbed as in Figure 11(b). The smallest mixing time appears to occur when the propensities are at 0.5 when the system has the highest change to leave the 2-cycle and be absorbed—see Figure 12(a). However this mixing time is 6.9521, which is higher than the mixing time for the starting example. The minimum mixing time seems to occur when the propensities are not equal.

To further analyze the effect of changes in propensities, we start by setting all propensities to 0.5 and allowing one propensity to differ. The lowest mixing time for propensities tested from 0.1 to 0.9 is 4.8887 at 0.1 and remains the same for all variables.




Figure 10: A comparison of the mixing times. The x -axis represents the value of the propensities. The y -axis represents the mixing time. The blue line represents the mixing times when all propensities are the same. We find the lowest mixing time when all of the propensities are 0.5. The red line represents when we only change the propensities inversely for p_2^\downarrow and p_3^\uparrow while keeping all other propensities 0.5. In this case, we find that the mixing time is lowest at the ends of the graph. That is when $p_2^\downarrow = 0.9$ and $p_3^\uparrow = 0.1$ and vice versa, we achieve the lowest mixing time.

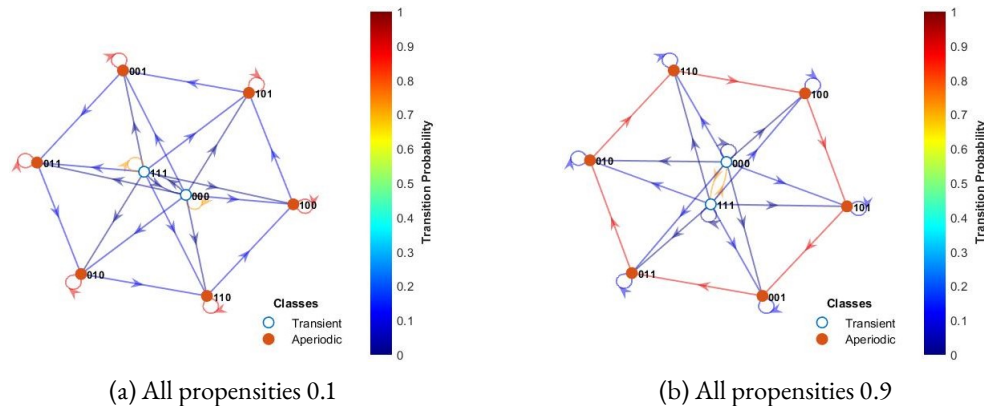


Figure 11: Same propensities comparison: When the propensities are all 0.1, the network has a tendency to stay at the state it started on and when the propensities are all 0.9, it has a tendency to go to the next state in the cycle more quickly. In both cases, the probability to leave the 2-cycle is low.

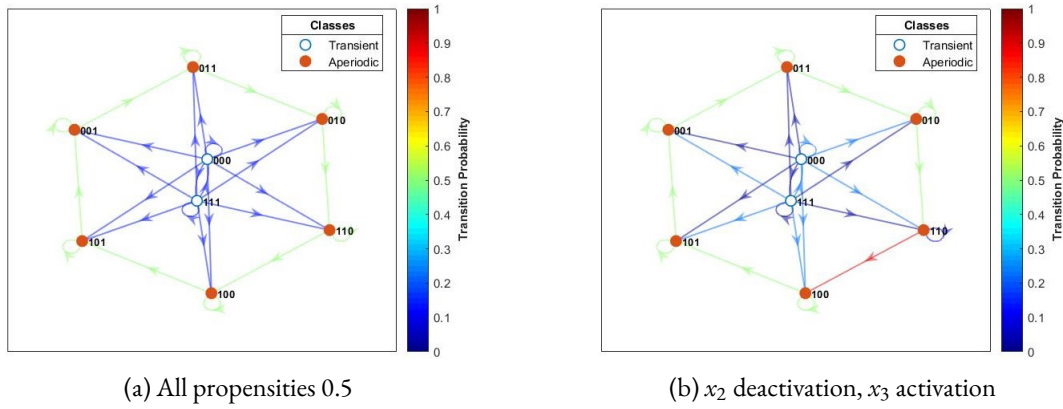


Figure 12: When all of the propensities are 0.5, the system has the fastest mixing time of all same propensities. When $p_2^\downarrow = 0.9$ and $p_3^\uparrow = 0.1$, we see that state 001 has a much higher probability to go to state 100 in comparison to all other probabilities in the network. This is when the mixing time is fastest overall.

$$\begin{pmatrix}
 (1 - P_1^\uparrow) (1 - P_2^\uparrow) (1 - P_3^\uparrow) & 0 & 0 & 0 & 0 & 0 & 0 & 0 & P_1^\downarrow P_2^\downarrow P_3^\downarrow \\
 (1 - P_1^\uparrow) (1 - P_2^\uparrow) P_3^\uparrow & 1 - P_2^\uparrow & 0 & 0 & 0 & P_1^\downarrow & 0 & 0 & P_1^\downarrow P_2^\downarrow (1 - P_3^\downarrow) \\
 (1 - P_1^\uparrow) P_2^\uparrow (1 - P_3^\uparrow) & 0 & 1 - P_1^\uparrow & P_3^\uparrow & 0 & 0 & 0 & 0 & P_1^\downarrow (1 - P_2^\downarrow) P_3^\downarrow \\
 (1 - P_1^\uparrow) P_2^\uparrow P_3^\uparrow & P_2^\uparrow & 0 & 1 - P_3^\uparrow & 0 & 0 & 0 & 0 & P_1^\downarrow (1 - P_2^\downarrow) (1 - P_3^\downarrow) \\
 P_1^\uparrow (1 - P_2^\uparrow) (1 - P_3^\uparrow) & 0 & 0 & 0 & 1 - P_3^\downarrow & 0 & P_2^\downarrow & 0 & (1 - P_1^\downarrow) P_2^\downarrow P_3^\downarrow \\
 P_1^\uparrow (1 - P_2^\uparrow) P_3^\uparrow & 0 & 0 & 0 & P_3^\downarrow & 1 - P_1^\downarrow & 0 & 0 & (1 - P_1^\downarrow) P_2^\downarrow (1 - P_3^\downarrow) \\
 P_1^\uparrow P_2^\uparrow (1 - P_3^\uparrow) & 0 & P_1^\uparrow & 0 & 0 & 0 & 1 - P_2^\downarrow & 0 & (1 - P_1^\downarrow) (1 - P_2^\downarrow) P_3^\downarrow \\
 P_1^\uparrow P_2^\uparrow P_3^\uparrow & 0 & 0 & 0 & 0 & 0 & 0 & 0 & (1 - P_1^\downarrow) (1 - P_2^\downarrow) (1 - P_3^\downarrow)
 \end{pmatrix}$$

Figure 13: Transition matrix of the Repressilator model.

Allowing two functions to have differing inverse propensities yields more interesting results. For example, if the propensities are

$$\begin{pmatrix} 0.5 & 0.5 & 0.1 \\ 0.5 & 0.9 & 0.5 \end{pmatrix},$$

we get a mixing time of 4.3157. The mixing times for propensities for this inverse relationship between the activation propensity for the third function and the deactivation propensity for the second function can be seen in Figure 12(b) starting at 0.1 and 0.9 respectively. Setting all other propensities to 0.9 allows an even shorter mixing time of 2.8627 at endpoints 0.1 and 0.9, but setting all other propensities to 0.1 forces the mixing time to be even longer with 11.9877 being the smallest mixing time at 0.5. Finally, it is worth noting that the time to absorption remains low even when mixing times are large. For instance, when the mixing time is 11.8977 for propensity matrix

$$\begin{pmatrix} 0.1 & 0.1 & 0.5 \\ 0.1 & 0.1 & 0.5 \end{pmatrix},$$

the time to absorption is 1.6949. In summary, we have observed that the mixing times of this model are sensitive to changes in the propensity parameters while the time to absorption remained stably low. This means that it is always possible to escape from the artificial attractor and then to transition into the biologically-relevant attractor relatively quickly, and after a number of steps given by the mixing time the stable oscillations will be achieved.

We show the transition matrix (obtained using Eq. 2) of the Boolean Repressilator model in Figure 13 and their eigenvalues in Equation 13. When all the propensities are the same (i.e., $p = p_k^\uparrow = p_k^\downarrow$ for all k), we see that the absolute value of all eigenvalues are minimized when $p = p_k^\uparrow = p_k^\downarrow = 0.5$ as shown in Figure 14.

$$\begin{aligned}
 \lambda_1 &= 1, & \lambda_3 &= 3p^2 - 3p + 1, & \lambda_5 &= 1 + \frac{3^{(1/2)} p * i}{2} - \frac{p}{2}, & \lambda_7 &= 1 + \frac{3^{(1/2)} p * i}{2} - \frac{3p}{2}, \\
 \lambda_2 &= 1 - 2p, & \lambda_4 &= 1 - \frac{3^{(1/2)} p * i}{2} - \frac{p}{2}, & \lambda_6 &= 1 - \frac{3^{(1/2)} p * i}{2} - \frac{3p}{2}, & \lambda_8 &= -2p^3 + 3p^2 - 3p + 1.
 \end{aligned} \tag{13}$$




Figure 14: Here we show the plots of eigenvalues from the transition matrix in Figure 13, where all propensities are assumed to be the same. Eigenvalues (listed in order in Equation 13) are all minimized when the propensity is $p = 0.5$, and graphical results can easily be confirmed analytically or numerically with MATLAB's `fminsearch` function.

4 Conclusion

In this paper, we studied the effect of changes in model parameters on optimal control policies, time to absorption, and mixing times by focusing on a class of Markov decision processes that is obtained by considering stochastic Boolean networks equipped with control actions. The stochastic framework that we used is the one called Stochastic Discrete Dynamical Systems (SDDS). The SDDS framework introduces stochasticity by assigning two propensity parameters to each function in a Boolean network. Using the SDDS framework and their control actions, we have performed an analysis of the effect of changes in propensity parameters on optimal control policies, time to absorption, and mixing times on specific Boolean models.

In the first part of the results, we studied how the choice of propensity parameters affects the effectiveness of optimal control policies where, depending on the range of values of certain parameters (e.g., the discount factor, level of noise, and the period of control policy duration), some policies became ineffective. We have also shown that by tuning the noise level and duration of the control policies one can still guarantee the effectiveness of some of the policies that we have considered.

In the second part of the results, we have studied the effect of changes in the propensity parameters on the time to absorption and the mixing time for a Boolean model of the Repressilator. We have shown that the mixing time of Repressilator model was very sensitive to changes in the propensity parameters while the time to absorption remained stably low. This means that in this model, regardless of the propensities, it will always be possible to reach the biologically-meaningful attractor relatively quickly, and after a number of steps specified by the mixing time the stable oscillations will be achieved.

To conclude, the results of this paper will help us understand the sensitivities of model parameters in similar models and their results and might help us inform the selection of different parameters in other models.

Acknowledgments

D.P. was supported by the NIH Training Grant T32CA165990. D.M. was partially supported by a Collaboration grant (850896) from the Simons Foundation. The authors thank the anonymous reviewers for their insightful comments that have improved the manuscript.

References

- Aguilar, B., P. Fang, R. Laubenbacher, and D. Murrugarra (2020, May). A near-optimal control method for stochastic Boolean networks. *Letters in Biomathematics* 7(1), 67–80. 193, 201
- Alon, U. (2019). *An Introduction to Systems Biology*. Boca Raton: Chapman and Hall/CRC. 201

- Bertsekas, D. (2019). *Reinforcement Learning and Optimal Control*. Athena Scientific. 193, 196
- Bertsekas, D. P. (2005). *Dynamic Programming and Optimal Control*. Athena Scientific. 193
- Deritei, D., N. Kunšič, and P. Csermely (2022). Probabilistic edge weights fine-tune Boolean network dynamics. *bioRxiv*. 193
- Elowitz, M. B. and S. Leibler (2000). A synthetic oscillatory network of transcriptional regulators. *Nature* 403(6767), 335–338. 201
- Golinelli, O. and B. Derrida (1989). Barrier heights in the Kauffman model. *Journal De Physique* 50(13), 1587–1601. 193
- Grinstead, C. and J. Snell (1997). *Introduction to Probability*. American Mathematical Society. 196
- Helikar, T., B. Kowal, S. McClenathan, M. Bruckner, T. Rowley, A. Madrahimov, B. Wicks, M. Shrestha, K. Limbu, and J. A. Rogers (2012). The cell collective: toward an open and collaborative approach to systems biology. *BMC Systems Biology* 6(1), 1–14. 193
- Levin, D. A. and Y. Peres (2017). *Markov Chains and Mixing Times*, Volume 107. American Mathematical Soc. 196
- Márquez-Zacarías, P., R. M. Pineau, M. Gomez, A. Veliz-Cuba, D. Murrugarra, W. C. Ratcliff, and K. J. Niklas (2021). Evolution of cellular differentiation: from hypotheses to models. *Trends in Ecology & Evolution* 36(1), 49–60. 193
- Murrugarra, D. and B. Aguilar (2019). Modeling the stochastic nature of gene regulation with Boolean networks. In *Algebraic and Combinatorial Computational Biology*, pp. 147–173. Elsevier. 193, 195, 198, 199
- Murrugarra, D., A. Veliz-Cuba, B. Aguilar, S. Arat, and R. Laubenbacher (2012). Modeling stochasticity and variability in gene regulatory networks. *EURASIP Journal on Bioinformatics and Systems Biology* 2012(1), 5. 193
- Murrugarra, D., A. Veliz-Cuba, B. Aguilar, and R. Laubenbacher (2016, Sep). Identification of control targets in Boolean molecular network models via computational algebra. *BMC Syst Biol* 10(1), 94. 196, 198
- Plaugher, D., B. Aguilar, and D. Murrugarra (2022). Uncovering potential interventions for pancreatic cancer patients via mathematical modeling. *Journal of Theoretical Biology* 548, 111197. 193
- Plaugher, D. and D. Murrugarra (2021). Modeling the pancreatic cancer microenvironment in search of control targets. *Bulletin of Mathematical Biology* 83(11), 1–26. 193
- Saadatpour, A., R.-S. Wang, A. Liao, X. Liu, T. P. Loughran, I. Albert, and R. Albert (2011, Nov). Dynamical and structural analysis of a t cell survival network identifies novel candidate therapeutic targets for large granular lymphocyte leukemia. *PLoS Comput Biol* 7(11), e1002267. 194, 196, 197, 198
- Shmulevich, I. and E. R. Dougherty (2010). *Probabilistic Boolean Networks - The Modeling and Control of Gene Regulatory Networks*. SIAM. 193
- Shmulevich, I., E. R. Dougherty, S. Kim, and W. Zhang (2002). Probabilistic Boolean networks: a rule-based uncertainty model for gene regulatory networks. *Bioinformatics* 18(2), 261–274. 193
- Sutton, R. S. and A. G. Barto (1998). *Reinforcement Learning: An Introduction*, Volume 1. MIT press Cambridge. 193, 196
- Trairatphisan, P., A. Mizera, J. Pang, A. A. Tantar, J. Schneider, and T. Sauter (2013). Recent development and biomedical applications of probabilistic Boolean networks. *Cell Communication and Signaling* 11(1), 1–25. 193
- Veliz-Cuba, A. and B. Stigler (2011). Boolean models can explain bistability in the lac operon. *Journal of Computational Biology* 18(6), 783–794. 193
- Wooten, D. J., J. G. T. Zañudo, D. Murrugarra, A. M. Perry, A. Dongari-Bagtzoglou, R. Laubenbacher, C. J. Nobile, and R. Albert (2021). Mathematical modeling of the *Candida albicans* yeast to hyphal transition reveals novel control strategies. *PLoS Computational Biology* 17(3), e1008690. 193
- Yousefi, M. R., A. Datta, and E. R. Dougherty (2012). Optimal intervention strategies for therapeutic methods with fixed-length duration of drug effectiveness. *IEEE Transactions on Signal Processing* 60(9), 4930–4944. 193, 201
- Zañudo, J. G. T. and R. Albert (2015, Apr). Cell fate reprogramming by control of intracellular network dynamics. *PLoS Comput Biol* 11(4), e1004193. 196
- Zañudo, J. G. T., G. Yang, and R. Albert (2017). Structure-based control of complex networks with nonlinear dynamics. *Proceedings of the National Academy of Sciences* 114(28), 7234–7239. 196