

Intervention in Gene Regulatory Networks via a Stationary Mean-First-Passage-Time Control Policy

Golnaz Vahedi*, *Student Member, IEEE*, Babak Faryabi, *Student Member, IEEE*,
Jean-Francois Chamberland, *Member, IEEE*, Aniruddha Datta, *Senior Member, IEEE*,
and Edward R. Dougherty, *Member, IEEE*

Abstract—A prime objective of modeling genetic regulatory networks is the identification of potential targets for therapeutic intervention. To date, optimal stochastic intervention has been studied in the context of probabilistic Boolean networks, with the control policy based on the transition probability matrix of the associated Markov chain and dynamic programming used to find optimal control policies. Dynamical programming algorithms are problematic owing to their high computational complexity. Two additional computationally burdensome issues that arise are the potential for controlling the network and identifying the best gene for intervention. This paper proposes an algorithm based on mean first-passage time that assigns a stationary control policy for each gene candidate. It serves as an approximation to an optimal control policy and, owing to its reduced computational complexity, can be used to predict the best control gene. Once the best control gene is identified, one can derive an optimal policy or simply utilize the approximate policy for this gene when the network size precludes a direct application of dynamic programming algorithms. A salient point is that the proposed algorithm can be model-free. It can be directly designed from time-course data without having to infer the transition probability matrix of the network.

Index Terms—Dynamic programming, genetic regulatory networks, mean first-passage time, probabilistic Boolean networks, stochastic optimal control.

I. INTRODUCTION

AN ULTIMATE objective of modeling genetic regulatory networks is the identification of potential targets for therapeutic intervention [1]. For instance, in cancer, one can consider correlation between metastasis and the abundances of messenger ribonucleic acid (mRNA) for certain genes. In this vein, the abundance of mRNA for the gene WNT5A has been found to be highly discriminating between cells with properties typically associated with high versus low metastatic competence [2].

Manuscript received August 14, 2007; revised March 8, 2008. First published June 10, 2008; current version published September 26, 2008. This work was supported in part by the National Science Foundation under Grant CCF-0514644 and Grant ECS-0355227. *Asterisk indicates corresponding author.*

*G. Vahedi is with the Department of Electrical and Computer Engineering, Texas A&M University, College Station, TX 77843 USA (e-mail: golnaz@tamu.edu).

B. Faryabi, J.-F. Chamberland, and A. Datta are with the Department of Electrical and Computer Engineering, Texas A&M University, College Station, TX 77843 USA (e-mail: bfaryabi@tamu.edu; chmbrlnd@tamu.edu; datta@ece.tamu.edu).

E. R. Dougherty is with the Department of Electrical and Computer Engineering, Texas A&M University, College Station, TX 77843 USA and also with the Translational Genomics Research Institute, Phoenix, AZ 85004 USA (e-mail: edward@tamu.edu).

Digital Object Identifier 10.1109/TBME.2008.925677

Appropriate alteration in the expression of WNT5A can be perceived therapeutically, and it can therefore be used to search for an optimal intervention strategy [3].

To date, optimal regulatory intervention has been studied in the context of probabilistic Boolean networks (PBNs), in particular, with respect to the dynamics determined by the probability transition matrix of the associated Markov chain [4]. Major efforts have focused on manipulating external (control) variables to desirably affect dynamical evolution over a finite time horizon [5], [6]. These short-term policies have been shown to change the dynamical performance of regulatory networks over a small number of stages; however, they are not necessarily effective in changing long-run network behavior. To address this issue, stochastic control has been employed via dynamic programming algorithms to find stationary control policies that affect the steady-state distributions of PBNs [7].

Study of infinite-horizon intervention strategies poses two fundamental questions. First, is it possible to beneficially affect a network by applying the optimal stationary control policy? This translates into assessing the *controllability* of the network. In practice, a physician would like to predict the effectiveness of a certain treatment at different stages of a disease and on different patients. Investigating the effect of a certain type of control on various networks is equivalent to questioning the controllability of the network. To date, there has been no investigation on this important topic in the context of gene regulatory networks. Second, can we identify the best intervening gene? In other words, which gene is the best potential “lever point,” to borrow the terminology from [8], in the sense of having the greatest possible impact on the desired network behavior? In principle, solving an optimal control problem for each candidate gene and comparing the performance of the system for these various controls would answer these questions; however, this process is a computationally demanding procedure. The complexity of dynamic programming algorithms is vast and increases exponentially with the number of genes [9].

In their early papers, Shmulevich *et al.* employ two methods for selecting a candidate gene for intervention: *mean first-passage time (MFPT)* and *influence* [1], [4]. The following biological example, borrowed from [10], explains the intuition behind using MFPTs for selecting the best control gene. In biology, there are numerous cases where the (in)activation of a certain gene or protein can lead more quickly (or with higher

probability) to a particular cellular functional state or phenotype than the (in)activation of another gene or protein. For instance, in a stable cancer cell line, in the absence of intervention, the cells will keep proliferating. This behavior can be reversed by controlling the expression of certain genes. Assume that the goal of the intervention is to push the cell into programmed cell death (apoptosis). Further assume that we can achieve this intervention with two candidate genes: p53 and telomerase. The p53 gene is the most well-known tumor suppressor gene [11]–[13]. The telomerase gene encodes telomerase, which maintains the integrity of the end of chromosomes (telomeres) in germ cells. Germ cells are responsible for propagating the complete genetic material to the following generation. Telomerase also maintains the integrity of the end of chromosomes in progenitor cells. Progenitor cells are responsible for replenishing cells during the normal cell turnover (homeostasis). In somatic cells, the telomerase gene is turned off, resulting in telomere shortening each time the cell divides—a key reason for the limited life span of normal cells [14]. In the majority of tumor cells, telomerase is activated, which is believed to contribute to the prolonged life span of the tumor cells [15]. This worsens prognosis for cancer patients [16], [17]. Extensive experimental results indicate that when p53 is activated in the cells, for example, in response to radiation, the cells undergo rapid growth inhibition and apoptosis in as short as a few hours [18]. In contrast, inhibition of the telomerase gene also leads to cell growth inhibition, differentiation, and cell death, but only after cells go through a number of cell divisions (allowing telomere shortening). Cell death takes a longer time through this latter process than via p53 activation. The use of MFPTs for finding the best control gene is intuitive; however, it focuses on a one-step control scenario.

The influence method distinguishes genes that have a major impact on a predictor function from those that have only a minor impact. This method was introduced to reflect the extent to which a set of genes is capable of determining the value of a target gene [4]. It has been used as a criterion to select a control gene with the suggestion that the gene with the largest influence on the target gene is likely to be a good control gene in the finite- and infinite-horizon control of PBNs [6], [7]; however, no research has been done on the overall performance of this heuristic measure.

Capitalizing on the biological intuition behind MFPT, we propose an algorithm based on MFPTs that assigns a stationary control policy for each candidate gene. We call this algorithm the *MFPT* algorithm and refer to the corresponding stationary control policy as the *MFPT* control policy. The proposed algorithm selects the MFPT control policy based on two heuristics: 1) it is preferable to *reach* desirable states as early as possible and 2) it is preferable to *leave* undesirable states as early as possible. The MFPT algorithm can be employed in four main applications.

First, the MFPT algorithm can be used for predicting the best control gene. The MFPT algorithm enables the computation of the MFPT control policies for all the genes in the network with a manageable complexity. The control gene with the highest desirable effect on the long-run behavior of the network upon

the application of the corresponding MFPT control policy is likely the most effective gene for controlling the biological system. Second, to reduce the complexity of the optimal stochastic control problem, the MFPT control policy can be used as an approximate solution. Contrary to optimal algorithms, the MFPT algorithm finds policies with constant complexity. Third, the MFPT algorithm can be used to measure the controllability of a network. Since the MFPT control policy is an approximation for the optimal control policy, one can define a network to be controllable if the effect of the MFPT control policy is greater than a desired threshold.

Finally, the MFPT algorithm can be used to design a control policy without requiring network inference. The optimal stochastic control policies proposed thus far require perfect knowledge of the probability transition matrix governing the network, which must be derived from the PBN or inferred directly. This procedure is prone to modeling errors and suffers from problems of computational complexity for both network inference and finding the optimal control solutions. To achieve model-free intervention, the MFPT control policy can be designed based on estimates of the MFPTs. The model-free intervention method has low complexity, is robust to modeling errors, and adapts to changes in the underlying biological system.

Our focus in this paper is on binary PBNs; however, the MFPT algorithm applies without change to a PBN having any discrete range of values. More generally, the MFPT algorithm can be applied to any Markovian regulatory network and should be mathematically viewed in this manner. For instance, the MFPT algorithm applies to dynamic Bayesian networks (DBNs) [19]. The proposed method can actually be viewed for DBNs in the same manner as with PBNs because any DBN can be represented by a probabilistically equivalent PBN [20]. Concentrating on the latter framework, the difficulty with PBNs possessing more than two values is that the size of the state space expands dramatically. Conceptually, networks with finer quantization can be analyzed using the same tools; indeed, the original application of automatic control considered a ternary network arising from complementary deoxyribonucleic acid (cDNA)-microarray data quantized into three values: -1 (down-regulated), $+1$ (up-regulated), and 0 (invariant) [5]. Here, let us make two points. First, from a theoretical perspective, with finer quantization, the MFPT algorithm still provides significant computational benefits over ordinary dynamic programming algorithms. Second, from a practical perspective, as with the melanoma example presented in this paper, our focus is on the up–down regulation model. In this scenario, the PBN is binary; in the case of quantization based on down-regulation, up-regulation, or invariance, the network is ternary.

This paper is structured in the following manner. Necessary definitions are provided in Section II. The MFPT algorithm and its applications are explained in Sections III and IV. We analyze the complexity of the MFPT algorithm in Section V. In Section VI, we corroborate our claims using extensive simulations for four applications: comparison of the optimal and MFPT control policies, finding the best control gene, quantifying controllability, and model-free control.

We then compare optimal and MFPT control policies for the network obtained using melanoma gene-expression data. Although our focus is on applications in the framework of PBNs, the MFPT algorithm applies to any Markovian regulatory network.

II. BACKGROUND

A. Probabilistic Boolean Networks

A PBN consists of a sequence $V = \{x_i\}_{i=1}^n$ of n nodes where $x_i \in \{0, \dots, d-1\}$, together with a sequence $\{f_l\}_{l=1}^k$ of vector-valued functions called *predictor functions*. In the framework of gene regulation, each element x_i represents the expression value of a gene. It is common to mix the terminology by referring to x_i as the i th gene. Each vector-valued function $f_l = (f_{l1}, \dots, f_{ln})$ determines a constituent network of the PBN. The function $f_{li} : \{0, \dots, d-1\}^n \rightarrow \{0, \dots, d-1\}$ is the predictor of gene i whenever network l is selected. The number of quantization levels is denoted by d . At each step, a predictor function is randomly selected according to probability distribution $\{p_l\}_{l=1}^k$. After selecting the predictor function f_l , the values of genes are updated accordingly, that is, in conformity with the network determined by f_l . We consider PBNs with perturbation, in which each gene may change its value with a small perturbation probability p at each time unit. The dynamics of a PBN can be represented via a Markov chain, and as a consequence of the perturbation, the Markov chain is ergodic and possesses a steady-state distribution.

Two quantization levels have thus far been used in practice. If $d = 2$ (binary), then the constituent networks are Boolean networks with 0 and 1 meaning OFF and ON, respectively. The case $d = 3$ (ternary) arises when we consider a gene to be down-regulated (-1), up-regulated (1), or invariant (0). This situation commonly occurs with cDNA microarrays, where a ratio is taken between the expression values on the test channel (red) and the base channel (green). In this paper, we will develop the methodology for $d = 2$, so that gene values are either zero or one; however, the methodology is applicable to any finite number of levels.

The gene-activity profile (GAP) is an n -digit binary vector $\mathbf{x}(t) = (x_1(t), \dots, x_n(t))$ giving the expression values of the genes at time t , where $x_i(t) \in \{0, 1\}$. We note that there is a natural bijection between the GAP $\mathbf{x}(t)$ and its decimal representation $z(t)$, which takes values in $\mathcal{S} = \{0, 1, \dots, 2^n - 1\}$. This bijection will be used later to present data.

In the presence of external controls, we suppose that the PBN has m binary inputs, $c_1(t), \dots, c_m(t)$, which specify the interventions on control genes g_1, \dots, g_m . A control $c_i(t)$, which can take values zero or one at each step t , specifies the action on the control gene g_i . The decimal bijection of the control vector $u_{g_1, \dots, g_m}(t) \in \mathcal{C} = \{0, 1, \dots, 2^m - 1\}$ describes the complete status of all the control inputs. As in previous applications, we focus on a single control gene $u_g(t) \in \mathcal{C} = \{0, 1\}$. The treatment alters the status of the control gene g , which can be selected among all the genes in the network. If the control at time step t is on, $u_g(t) = 1$, then the state of the control gene

g is toggled; if $u_g(t) = 0$, then the state of the control gene g remains unchanged.

System evolution is represented by a stationary discrete-time equation

$$z(t+1) = f(z(t), u_g(t), w(t)) \quad \forall t = 0, 1, \dots$$

where state $z(t)$ is an element of the state-space \mathcal{S} . The disturbance $w(t)$ is the manifestation of uncertainties in the PBN. It is assumed that both the gene perturbation distribution and the network switching distribution are independent and identical for all time steps t . The state $z(t)$ at any time step t is a GAP. Originating from a state i , the successor state j is selected randomly within the set \mathcal{S} according to the transition probability

$$p_{ij}(u) \triangleq P(z(t+1) = \frac{j}{z(t)} = i, u_g(t) = u)$$

for all i and j in \mathcal{S} , and for all u in \mathcal{C} . Gene perturbation ensures that all the states in the Markov chain communicate with each other. Hence, the finite-state Markov chain has a unique steady-state distribution [1].

B. Optimal Intervention

The problem of optimal intervention for PBNs is formulated as an optimal stochastic control problem. A cost-per-stage $r(i, u, j)$ is associated to each intervention in the system. In general, a cost-per-stage may depend on the origin state i , the successor state j , and the control input u . We assume that the cost-per-stage is stationary and bounded for all i, j in \mathcal{S} , and u in \mathcal{C} . We define the expected immediate cost in state i when control u is selected by

$$\bar{r}(i, u) = \sum_{j \in \mathcal{S}} p_{ij}(u) r(i, u, j).$$

We consider the discounted formulation of the expected total cost. The discounting factor $\alpha \in (0, 1)$ ensures the convergence of the expected total cost over the long run [21]. In the case of cancer therapy, the discounting factor emphasizes that obtaining treatment at an earlier stage is favored over later stages. The expected total discounted cost, given a policy π_g , an initial state i , and control gene g , is denoted by

$$J_{\pi_g}(i) = \lim_{N \rightarrow \infty} E \left\{ \sum_{t=0}^{N-1} \alpha^t r(z(t), \mu_g(z(t)), z(t+1)) \mid i \right\}. \quad (1)$$

A policy $\pi_g = \{\mu_g(0), \mu_g(1), \dots\}$ is a sequence of decision rules $\mu_g(t) : \mathcal{S} \rightarrow \mathcal{C}$, for each time step t , when the control gene is g . The vector \mathbf{J}_{π_g} of the expected total costs is called the value function. In a stochastic control problem, we seek an intervention strategy π_g^* among all the admissible intervention strategies Π_g that minimizes the value function for each state i , i.e.,

$$\pi_g^*(i) = \arg \min_{\pi_g \in \Pi_g} J_{\pi_g}(i) \quad \forall i \in \mathcal{S}. \quad (2)$$

A stationary intervention strategy for the control gene g is an admissible intervention strategy of the form $\pi_g = \{\mu_g, \mu_g, \dots\}$. It is known that an optimal intervention strategy exists for the

discounted optimal stochastic control problems. The optimal cost function J^* satisfies

$$J^*(i) = \min_{u \in \mathcal{C}} \left[\bar{r}(i, u) + \alpha \sum_{j=0}^{2^n-1} p_{ij}(u) J^*(j) \right] \quad \forall i \in \mathcal{S}. \quad (3)$$

Furthermore, J^* is the unique solution of this equation within the class of bounded functions. Equation (3) is known as the *Bellman optimality equation*. The optimal control policy attains the minimum on the right-hand side of the Bellman optimality equation for all i . Moreover, an optimal policy determined by the Bellman optimality equation is also a stationary policy. For the proofs of the earlier statements and more details, one can refer to [21]. Standard dynamic programming algorithms can be used to find a fixed point of the Bellman optimality equation.

C. Influence

Influence is a method for quantifying the relative impact of genes on other genes within the context of PBNs [4]. The influence $I_j(f)$ of gene x_j on the function f , with respect to the probability distribution $D(x)$, $x \in \{0, 1\}^n$, is defined as

$$I_j(f) = E_D \left[\frac{\partial f(x)}{\partial x_j} \right] \quad (4)$$

where $E_D[\cdot]$ is the expectation operator with respect to the distribution D , $\partial f(x)/\partial x_j = f(x^{(j,0)}) \oplus f(x^{(j,1)})$ is the partial derivative of the Boolean function f , the symbol \oplus is addition modulo 2 (EXCLUSIVE OR), and $x^{(j,k)} = (x_1, \dots, x_{j-1}, k, x_{j+1}, \dots, x_n)$ for $k \in \{0, 1\}$. In other words, (4) gives the influence as the probability [under the distribution $D(x)$] that a toggle of the j th variable changes the value of the function. In the context of PBNs, the influence of gene x_k on gene x_i is given by $I_k(x_i) = \sum_{j=1}^{l(i)} I_k(f_j^{(i)}) \cdot p_j^{(i)}$ where $\{p_j^{(i)}\}_{j=1}^{l(i)}$ are the selection probabilities of the predictor functions of gene i and $l(i)$ represents the number of predictor functions of gene i [4]. To quantify the long-run influence, $D(x)$ is the stationary distribution of the PBN.

III. MFPT ALGORITHM

In this section, we first elaborate on how the MFPT algorithm is designed based on the MFPT. We then summarize the MFPT algorithm. Application of the MFPT algorithm requires the designation of desirable and undesirable states, and this depends upon the existence of relevant biological knowledge. Intervention is performed by flipping (toggling) the expression status of a particular gene from ON to OFF or vice versa, the intent being to externally guide the time evolution of the network toward more desirable states. If g is the control gene, then applying the control (intervention) in state \mathbf{x} translates into flipping the value of g in that state (the control gene g changes to zero if its value is one and vice versa). Consequently, the network resumes its transition from the new state $\tilde{\mathbf{x}}$, which we call the *flipped state*. In the context of therapy, the state-space of a PBN can be partitioned into desirable and undesirable states. Given a control gene, when

a desirable state reaches the set of undesirable states, on average, faster than its flipped state, it is reasonable to intervene and transition into the flipped state. Similarly, if an undesirable state reaches the set of desirable states, on average, faster than its flipped state, it is reasonable not to intervene. These insights motivate the use of MFPTs for designing intervention strategies.

Without loss of generality, we can assume that the transition probability matrix \mathbf{P} of the Markov chain (representing a PBN) is partitioned according to

$$\mathbf{P} = \begin{pmatrix} P_{\mathcal{D},\mathcal{D}} & P_{\mathcal{D},\mathcal{U}} \\ P_{\mathcal{U},\mathcal{D}} & P_{\mathcal{U},\mathcal{U}} \end{pmatrix}$$

where \mathcal{D} and \mathcal{U} are the subsets of desirable and undesirable states, respectively. The MFPTs are computed by solving the following systems of linear equations [22]:

$$K_{\mathcal{U}} = e + P_{\mathcal{D},\mathcal{D}} \cdot K_{\mathcal{U}} \quad (5)$$

$$K_{\mathcal{D}} = e + P_{\mathcal{U},\mathcal{U}} \cdot K_{\mathcal{D}} \quad (6)$$

where e is a column vector of 1's with appropriate length, $K_{\mathcal{U}}$ is a vector containing the MFPTs from each state in the subset of desirable states \mathcal{D} to undesirable states in set \mathcal{U} , and $K_{\mathcal{D}}$ is a vector containing the MFPTs from each state in the subset of undesirable states \mathcal{U} to the desirable states in set \mathcal{D} .

A control policy μ_g corresponding to control gene g is a vector of size 2^n , the number of states in the network. The decision rule $\mu_g : \mathcal{S} \rightarrow \mathcal{C}$ specifies the control action for each state \mathbf{x} in \mathcal{S} . Having $\mu_g(\mathbf{x}) = 0$ for state \mathbf{x} means that whenever the network reaches state \mathbf{x} , no control is applied and the system continues its transition based on the transition probabilities of state \mathbf{x} . On the other hand, having $\mu_g(\mathbf{x}) = 1$ implies that, whenever the network reaches state \mathbf{x} , the control is applied and the system continues its evolution based on the transition probabilities of state $\tilde{\mathbf{x}}$, the flipped state of \mathbf{x} .

The goal of the MFPT algorithm is to design the MFPT control policies $\{\hat{\mu}_g\}_{g=1}^n$. The objective is to choose a control value u for every state in \mathcal{S} such that the network evolves toward more desirable states. The MFPT algorithm selects the control policy for control gene g in the following manner. Assume that state \mathbf{x} is an undesirable state. We compare the MFPTs from state \mathbf{x} to \mathcal{D} and from the flipped state $\tilde{\mathbf{x}}$ to \mathcal{D} . In other words, we would like to know, on average, which one of these two states, \mathbf{x} and $\tilde{\mathbf{x}}$, *hits* the set of desirable states for the first time faster than the other one. The algorithm chooses $\hat{\mu}_g(\mathbf{x}) = 1$ if the difference between the MFPTs of state \mathbf{x} and the flipped state $\tilde{\mathbf{x}}$ to the set of desirable states, i.e., $K_{\mathcal{D}}(\mathbf{x}) - K_{\mathcal{D}}(\tilde{\mathbf{x}})$, is greater than a tuning parameter γ (to be discussed). Otherwise, $\hat{\mu}_g(\mathbf{x}) = 0$. Analogously, if state \mathbf{x} is desirable, then $\hat{\mu}_g(\mathbf{x}) = 1$ if the difference between the MFPTs of state \mathbf{x} and the flipped state $\tilde{\mathbf{x}}$ to undesirable states, i.e., $K_{\mathcal{U}}(\tilde{\mathbf{x}}) - K_{\mathcal{U}}(\mathbf{x})$, is greater than γ . Otherwise, $\hat{\mu}_g(\mathbf{x}) = 0$. These comparisons are repeated for all states. Algorithm 1 summarizes the proposed procedure.

The threshold γ in the MFPT algorithm is a tuning parameter chosen based on the ratio of the cost of control to the cost of undesirable states. When the cost of applying treatment in a state is high compared to the cost of undesirable states,

Algorithm 1 MFPT algorithm

```

Partition the state-space into undesirable  $\mathcal{U}$  and desirable  $\mathcal{D}$ 
subsets.
Compute  $K_{\mathcal{U}}$  and  $K_{\mathcal{D}}$ .
 $g \leftarrow 1$ .
repeat
  for All states  $\mathbf{x}$  in  $\mathcal{U}$  do
     $\tilde{\mathbf{x}} \leftarrow$  flip control gene  $g$  in  $\mathbf{x}$ .
    if  $K_{\mathcal{D}}(\tilde{\mathbf{x}}) - K_{\mathcal{D}}(\mathbf{x}) > \gamma$  then
       $\mu_g(\mathbf{x}) = 1$ ;
    else
       $\mu_g(\mathbf{x}) = 0$ ;
    end if
  end for
  for All states  $\mathbf{x}$  in  $\mathcal{D}$  do
     $\tilde{\mathbf{x}} \leftarrow$  flip control gene  $g$  in  $\mathbf{x}$ .
    if  $K_{\mathcal{U}}(\tilde{\mathbf{x}}) - K_{\mathcal{U}}(\mathbf{x}) > \gamma$  then
       $\mu_g(\mathbf{x}) = 1$ ;
    else
       $\mu_g(\mathbf{x}) = 0$ ;
    end if
  end for
   $g \leftarrow g + 1$ .
until  $g >$  number of genes

```

an optimal control policy is less likely to apply the control frequently. Thus, γ is set to a larger value when this ratio is higher, the intent being to apply control less frequently. We explain after the following definitions how one can set this parameter.

An effective control policy reduces the likelihood of visiting undesirable states compared to a network without intervention by modifying the long-run behavior of the network. The effectiveness of a control policy can be measured by the amount of change (*shift*) in the aggregated probability of undesirable states before and after the intervention. As a performance measure we define

$$\Delta P_g = \frac{\sum_{i \in \mathcal{U}} \pi_i - \sum_{i \in \mathcal{U}} \pi_i^g}{\sum_{i \in \mathcal{U}} \pi_i}$$

where π_i^g is the probability of being in undesirable state i in the long run after intervening with control gene g and π_i is the probability of being in undesirable state i in the long run when there is no intervention. The ratio ΔP_g measures the proportion of reduction in the total probability of undesirable states in the steady state when the control gene g is selected. We denote this proportion by ΔP_g^{opt} when an optimal control policy μ_g^* is applied. In other words, in the optimal case, one can shift the aggregated probability of undesirable states to desirable states by ΔP_g^{opt} through appropriately altering the expression of the control gene g . Similarly, the shift obtained by the MFPT control policy $\hat{\mu}_g^\gamma$ is denoted by $\Delta P_g^{\text{MFPT}(\gamma)}$, where γ is the tuning parameter.

We define the probability of the execution of control as

$$\Gamma_g = \sum_{j=0}^{2^n - 1} \pi_j \cdot I(\mu_g(j) = 1) \quad (7)$$

where n is the number of genes, π_j is the steady-state probability of state $j \in \mathcal{S}$, $\mu_g(j)$ is the value of the control policy in state

j , and I is the indicator function. The purpose of introducing this probability is to have a fair evaluation of the performance of the MFPT control policy in terms of the number of control executions, which for the optimal policy is related to the cost of control. For each control gene g , one can define Γ_g^{opt} as the probability of the execution of control when the optimal control policy is applied. Similarly, $\Gamma_g^{\text{MFPT}(\gamma)}$ is the probability of the execution of control when the MFPT control policy with the parameter γ is applied.

We numerically find the value of the parameter γ for each control cost. We generate random intervention problems and calculate the averages of Γ_g^{opt} and $\Gamma_g^{\text{MFPT}(\gamma)}$. These averages are taken over random intervention problems with fixed control cost. Starting from $\gamma = 0$, we increase the value of γ . For each control cost, the desired value of γ is the minimal one for which, on average, $\Gamma_g^{\text{opt}} > \Gamma_g^{\text{MFPT}(\gamma)}$. This condition guarantees that, on average, the MFPT control policy applies no more control than the optimal control policy. Since the values of the parameter γ are found from random intervention problems, in practice, one can have a conservative approach and choose the parameter γ to be greater than the proposed value. The conservative approach can assure a high probability that $\Gamma_g^{\text{opt}} > \Gamma_g^{\text{MFPT}(\gamma)}$. On the other hand, the deviation of $\Delta P_g^{\text{MFPT}(\gamma)}$ from ΔP_g^{opt} becomes larger.

IV. APPLICATIONS OF MFPT ALGORITHM

We devise solutions according to the MFPT algorithm for four intervention applications.

1) *Identification of the Best Control Gene*: Recalling the example of p53 and telomerase in Section I, it is important to select the most effective control gene in a therapeutic intervention. The best control gene g^* can be found by directly solving a dynamic programming algorithm and computing $\{\Delta P_g^{\text{opt}}\}_{g=1}^n$ for all the genes g in the network. In short, g^* is given by

$$g^* = \arg \max_{g=1, \dots, n} \Delta P_g^{\text{opt}}. \quad (8)$$

However, this optimal method to find the best control gene is computationally prohibitive. On the other hand, the MFPT algorithm enables the computation of the MFPT control policies $\{\hat{\mu}_g^\gamma\}_{g=1}^n$ for all the genes in the network with an acceptable complexity. Taking this approach, the MFPT algorithm predicts the best control gene to be

$$\hat{g} = \arg \max_{g=1, \dots, n} \Delta P_g^{\text{MFPT}(\gamma)}. \quad (9)$$

We will show that $\hat{g} = g^*$ with high probability and that $\Delta P_g^{\text{opt}} - \Delta P_g^{\text{MFPT}(\gamma)}$ is small whenever $\hat{g} \neq g^*$. In this context, we are using the MFPT algorithm to find the control gene. Once the best gene candidate is identified, an optimal control policy can be obtained using dynamic programming algorithms. We will also show that the MFPT-based prediction of the best control gene significantly outperforms the influence method.

2) *Approximation of the Optimal Control Policy*: The MFPT algorithm can devise an intervention strategy as an approximation of the optimal intervention strategy. To this end, we numerically find the value of the parameter γ for each control cost so that, on average, $\Gamma_g^{\text{opt}} > \Gamma_g^{\text{MFPT}(\gamma)}$. To assess the accuracy of the

approximation, we show that the average of $\Delta P_{g^*}^{\text{opt}} - \Delta P_{g^*}^{\text{MFPT}(\gamma)}$ over random intervention problems with fixed control cost is small. Note that, so as not to confound approximation accuracy with the MFPT algorithm's ability to find a control gene, we apply both the optimal and MFPT methods using the optimal control gene g^* .

3) *Controllability*: An important aspect of prognosis is quantifying the possibility of recovery. In our framework, this amounts to quantifying the *controllability* of a gene regulatory network, a concept borrowed from traditional control theory. Can the network be sufficiently controlled to provide meaningful recovery? We desire a controllability measure where the objective of the control is to reduce the likelihood of observing the undesirable states in the long run. An optimal control strategy takes into account the cost of control, but here, we want to focus only on the possibility of sufficient control, absent concerns with costs, either medical or financial. To this end, we choose the cost of control to be zero. The zero control–cost strategy admits any number of states with active control. Our point here (one that is certainly debatable) is that we desire a measure of controllability with no restrictions on the number of times the control might be applied. Thus, a possible approach is to set the cost of control to zero and compute $\Delta P_{g^*}^{\text{opt}}$. To reduce the computational burden, we propose $\Delta P_{g^*}^{\text{MFPT}(0)}$ ($\gamma = 0$) as a controllability measure. Our simulations show that the $\Delta P_{g^*}^{\text{MFPT}(0)}$ is a highly accurate approximation of $\Delta P_{g^*}^{\text{opt}}$ when the cost of control is zero. Therefore, the MFPT algorithm can be employed to determine the controllability of a network. For example, if $\Delta P_{g^*}^{\text{MFPT}(0)}$ is very small, we conclude that the network is not controllable. If $\Delta P_{g^*}^{\text{MFPT}(0)} = 0.5$, then we conclude that it is possible to shift 50% of the probability mass of the undesirable states to desirable ones in the long run, given the application of the control has zero cost.

4) *Model-Free Intervention*: To date, the proposed intervention methods for PBNs are model-dependent, requiring at least the knowledge of the transition probability matrix. This can be derived from the PBN if the latter is known. Since in practice, PBNs are not known except via system identification from observed data, one is faced with a difficult inference problem [23]. This problem can be avoided by directly inferring the transition probability matrix; however, this is still a formidable task because the complexity of estimating the transition probabilities of a Markov chain increases exponentially with the number of genes in the model. When time-course data are available, the MFPT algorithm can be implemented by directly estimating the MFPTs. The estimated MFPTs are used to construct the matrices of the MFPTs, K_U and K_D . The MFPT algorithm can then be applied to the estimated matrices K_U and K_D to devise a *model-free* MFPT control policy.

In the following, we propose a procedure for estimating the MFPTs from time-course measurements. Assume that x is a desirable state and it is observed at time t_0 . Further assume that, starting from time t_0 , the first undesirable state occurs at time $t_0 + k_0$. In other words, it takes k_0 time points for the desirable state x to transition (reach) to an undesirable state. Similarly, assume that the next observation of state x is at time t_1 , and

since time t_1 , the first undesirable state occurs at time $t_1 + k_1$. In this example, the average first passage time from state x to the subset of undesirable states is $(k_0 + k_1)/2$. Likewise, one can define an example for an undesirable state y reaching the subset of desirable states. It is evident that for a larger number of observations, this estimation becomes more accurate. The earlier procedure needs to be implemented with a low complexity. At each time point, we update the number of steps for each state to reach the opposite subset of states and store the frequency of the occurrence of each state. One needs to update the average first passage times for a new observation. The complexity of estimating the MFPTs following our procedure is constant with respect to the number of genes for each iteration. More details regarding the implementation can be found in the supplementary materials.

An advantage of the model-free approach is that the estimated matrices K_U and K_D can be updated whenever new time-course data become available. The possibility of updating the estimated MFPTs enables the MFPT algorithm to adapt its control policy to the status of gene interactions. In other words, the model-free MFPT control method is adaptive to changes in the network model. In contrast, the control policy devised by the existing intervention methods cannot adapt to the changes in the status of gene interactions. Once the PBN is inferred from data, the model-dependent control policy is fixed.

Through numerical studies, we will exhibit the effectiveness of the model-free MFPT control policy obtained by estimating the MFPTs. On one hand, we will estimate the matrices K_U and K_D based on synthetic time-course data and use the MFPT algorithm to find the control policy; on the other hand, we will use the same time-course data to build a Markov chain representing the dynamics of the model and then find the control policy based on the estimated transition probability matrix using dynamic programming. We will observe that the MFPT control policy based on the estimated MFPTs outperforms the control policy devised from the estimated transition probabilities of the Markov chain, given the same set of time-course data, for feasible size datasets.

V. COMPLEXITY ANALYSIS OF MFPT ALGORITHM

The main objective of an effective intervention strategy is to reduce the likelihood of visiting undesirable states compared to a network without intervention by modifying the long-run behavior of the network. Given a time-course dataset, there are two possible approaches for designing a strategy for any model such that its dynamic behavior can be represented as a Markov chain (such as PBN or dynamic Bayesian network).

In the first approach, one can estimate the transition probabilities of the states from time-course measurements. Let us call this approach “model-dependent.” We require all the details about the model, i.e., the transition probabilities of the Markov chain. Various methods can be employed to design an effective intervention strategy based on the estimated model. The optimal control policy can be designed via dynamic programming techniques [7]. In favor of lower computational complexity, an

approximation of the optimal control policy can be achieved using the MFPT algorithm.

In the second approach, an effective intervention strategy can be designed directly from time-course measurements. We call this approach “model-free.” In contrast to the model-dependent approach where the transition probabilities of the Markov chain are needed, we do not require the details of the model. To this end, a model-free algorithm based on reinforcement learning has recently been introduced [24]. This method bypasses the impediment of model estimation and an effective control policy can be designed with a low complexity. We propose that the MFPT algorithm can also be considered as a model-free method. In this section, we analyze the complexity of the model-based and the model-free MFPT algorithms.

A. Model-Dependent Approach

In the previous section, we introduced the four major applications of the MFPT algorithm: identification of the best control gene, approximation of an optimal control policy, controllability, and model-free intervention. Employment of the MFPT algorithm in the first three applications is considered as a model-dependent approach since it is assumed that the transition probability matrix of the Markov chain is known. Given the model is known, let us compare the computational complexities of the dynamic programming and the MFPT algorithms.

To find an optimal control policy using value or policy iteration, one should iteratively find the value (cost) function until the algorithm reaches the fixed point of the Bellman optimality equation in (3). Once the optimal cost functions are computed, one must check which control value attains the minimum on the right-hand side of the Bellman optimality equation and this procedure should be iterated for all the states. To the best of our knowledge, there does not exist a tight upper bound on the number of iterations required to find an optimal policy using either value or policy iteration, despite recent research initiatives [25]. Given the control gene, the policy iteration algorithm has complexity $\mathcal{O}(2^{3n})$ per iteration, whereas the complete complexity of the MFPT algorithm, which consists of two matrix inversions, is $\mathcal{O}(2^{3n})$. In general, it is known that the policy iteration algorithm converges, but it is not known whether “the number of iterations in policy iteration can be bounded by a polynomial in the instance size” [25]. Even assuming that the number of iterations can be bounded by a polynomial in the number of states, the complexity of the MFPT algorithm is lower than the policy iteration algorithm because it is computed once and does not require iteration. Regarding the value iteration algorithm, the asymptotic upper bound on the number of iterations required to find an optimal policy using the value iteration algorithm is polynomial in the number of states [25]. The degree of the polynomial is determined to be greater than 2 in special cases [26], [27]. Given the complexity of each iteration in the value iteration algorithm is $\mathcal{O}(2^{2n})$, the complexity of the value iteration algorithm to find an optimal control policy is $\mathcal{O}(2^{(2+\alpha)n})$, where $\alpha > 1$. Hence, the complexity of the MFPT algorithm is also lower than the complexity of the value iteration algorithm. To find the optimal cost functions for n control

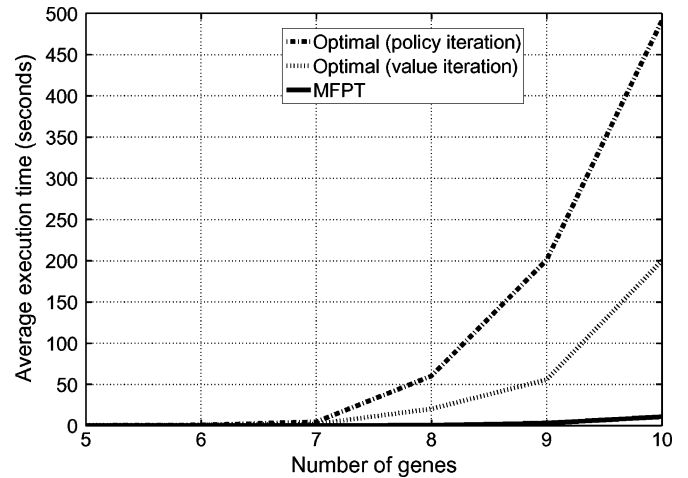


Fig. 1. Average execution time of the value and policy iteration algorithms over 1000 randomly generated intervention problems as functions of the number of genes along with the execution times of the MFPT algorithm.

genes, the complexity of a dynamic programming algorithm is n times the complexity of this algorithm for one control gene. In contrast, once the MFPT vectors are computed, they can be used to devise MFPT control policies for all control genes.

It is important to point out that for any control gene, in addition to the aforementioned complexities, the dynamic programming and the MFPT algorithms must loop over all the states to find their corresponding control policies. In dynamic programming algorithms, to obtain the optimal control policy, one must check which control value attains the minimum on the right-hand side of the Bellman optimality equation and this procedure must be iterated for all the states. In the MFPT algorithm, one must investigate which control value leads to a more favorable MFPT and this procedure must be repeated for all the states.

It is evident from the earlier analysis that the application of our proposed method is restricted to small number of genes since the complexity of the MFPT algorithm increases exponentially with the number of genes. We should point out that in our application of interest, intervention in gene regulatory networks, the goal is not to model fine-grained molecular interactions among a host of genes, but rather to model a limited number of genes, typically with very coarse quantization, whose regulatory activities are significantly related to a particular aspect of a specific disease, such as metastasis in melanoma [3]. Hence, while the asymptotic results on the complexities of optimal algorithms are encouraging, they are not our main interest; rather, our problem deals with networks with small numbers of states. Fig. 1 shows the average execution time of the value and policy iteration algorithms over 1000 randomly generated intervention problems as a function of the number of genes n along with the execution times of the MFPT algorithm. Per this figure, the execution time of the MFPT algorithm is much smaller than the execution time of the two optimal algorithms. The direct comparison has been limited to ten-gene networks on account of the high complexity of the modeling and optimal intervention algorithms. The

maximum size of the intervention problem that can be solved by our MFPT method is hardware-dependent. For instance, our current hardware configuration (single Xeon processor and 1-GB memory) can obtain MFPT intervention policy for a synthetic 15-gene regulatory network, which, given the data limits of current expression measuring technology, is sufficient for the applications in which we are now engaged. Given more memory and processing power, intervention strategies can be designed for larger networks. Should the need arise for larger networks, we will consider implementation on our Beowulf cluster at the Translational Genomics Research Institute.

B. Model-Free Approach

The model-dependent approaches yield effective solutions for large numbers of observations. However, these approaches have major drawbacks in practice. For lower numbers of observations, which correspond better to feasible experimental conditions, estimating the Markov chain yields poor results. Estimation errors may have a huge impact on finding an effective intervention strategy, which is often quite sensitive to changes in the transition probabilities [28]. Furthermore, the complexity of estimating the transition probabilities of a Markov chain increases exponentially with the number of genes in the model, $\mathcal{O}(2^{2n})$. This is in addition to the complexity of designing an effective intervention strategy. Hence, a procedure that can find an effective intervention strategy without having to know the transition probabilities is very attractive.

The model-free-based MFPT algorithm (fourth application) estimates the MFPTs from time-course measurements. The complexity of estimating these vectors following the proposed procedure in the previous section is constant with respect to n for each iteration, where n denotes the number of genes. In other words, we devise an effective intervention strategy by learning about the MFPTs directly from the data.

The highlight of this paper is the possibility of employing the MFPT algorithm in a model-free approach. To this end, we summarize the two main benefits of our proposed model-free method: 1) the complexity of the modeling and intervention is significantly less than that of the model-dependent methods and 2) in contrast to the optimal control problem approach, which is sensitive to changes in the system, the MFPT algorithm needs the average behavior of the system and is expected to be more appealing for smaller number of observations. We corroborate this claim in Section VI by comparing the model-free MFPT method with the model-dependent optimal control method.

VI. RESULTS AND DISCUSSION

In this section, we first study the performance of the MFPT algorithm for each of the aforementioned applications through extensive simulations of random PBNs. We then compare the performance of the MFPT algorithm and the influence method for the network obtained from a melanoma gene-expression dataset.

A. Synthetic Networks

We postulate the following cost-per-stage:

$$r(u, j) = \begin{cases} 0, & \text{if } u = 0 \text{ and } j \in \mathcal{D} \\ 10, & \text{if } u = 0 \text{ and } j \in \mathcal{U} \\ c, & \text{if } u = 1 \text{ and } j \in \mathcal{D} \\ 10 + c, & \text{if } u = 1 \text{ and } j \in \mathcal{U} \end{cases}$$

where c is the cost of control. The target gene is chosen to be the most significant gene in the GAP. We assume that the up-regulation of the target gene is undesirable. Consequently, the state-space is partitioned into desirable states $\mathcal{D} = \{0, \dots, 2^{n-1} - 1\}$, and undesirable states $\mathcal{U} = \{2^{n-1}, \dots, 2^n - 1\}$, where n is the number of genes. The cost values have been chosen in accord with an earlier study [7]. Since our objective is to down-regulate the target gene, a higher cost is assigned to destination states having an up-regulated target gene. Moreover, for a given status of the target gene for a destination state, a higher cost is assigned when the control is applied, versus when it is not. In practice, the cost values will have to mathematically capture the benefits and costs of intervention and the relative preference of states. These cost values will eventually be set with the help of physicians in accordance with their clinical judgment. Although this is not feasible within current medical practice, we do believe that such an approach will become feasible when engineering approaches are integrated into translational medicine. In order to investigate the effect of the cost of control in our algorithm, we vary its value from zero to ten, which is the cost of the undesirable states.

We generate random PBNs in the following manner. Each PBN consists of ten constituent BNs. Each BN is randomly generated with a specific bias b , the bias being the probability that a randomly generated Boolean function takes on the value one. Since the bias affects the dynamical properties of the randomly generated BNs [29], we take it as a parameter in our simulations. We randomly select the bias b of a BN from a beta distribution. We vary the mean of the beta distribution from 0.3 to 0.7 with step size 0.1. The variance σ^2 of the beta distribution is set to a constant value 0.0001. This provides random biases from low (0.3) to high (0.7). We generate 1000 random PBNs for each bias mean. For each PBN, the transition probabilities of the corresponding Markov chain are estimated. The earlier procedure is repeated for networks of five to ten genes. Due to the computational complexity of the optimal stochastic control problem and the estimation of the transition probabilities of the corresponding Markov chain, the study of a large number of networks beyond ten genes is outside our current computational capacity.

1) *Identification of the Best Control Gene:* We first show the performance of the MFPT algorithm and the influence method when they are employed to predict the best control gene. It is assumed that the cost of control c is equal to 1. In Tables I–IV, we compare the performances of the MFPT algorithm and the influence method for predicting the best control gene. First, the optimal control policy for each control gene is obtained by a dynamic programming algorithm. The best control gene g^* is

TABLE I
PROBABILITY OF FINDING THE BEST CONTROL GENE
WITH MFPT ALGORITHM WHEN $c = 1$ FOR NETWORKS
WITH DIFFERENT NUMBER OF GENES

Bias	0.3	0.4	0.5	0.6	0.7
(A) 5 genes	0.9850	0.9640	0.9570	0.9600	0.9720
(B) 6 genes	0.9430	0.9700	0.9760	0.9580	0.9870
(C) 7 genes	0.9440	0.9680	0.9670	0.9700	0.9570
(D) 8 genes	0.9660	0.9740	0.9860	0.9790	0.9710
(E) 9 genes	0.9132	0.9233	0.9741	0.9812	0.9812
(F) 10 genes	0.9470	0.9570	0.9860	0.9690	0.9610

TABLE II
AVERAGE DIFFERENCE BETWEEN PROPORTIONS OF REDUCTION IN TOTAL
PROBABILITY OF UNDESIRABLE STATES OBTAINED BY THE BEST CONTROL
GENE g^* AND PREDICTED CONTROL GENE OBTAINED BY MFPT ALGORITHM \hat{g}
FOR NETWORKS WITH VARIOUS NUMBER OF GENES

Bias	0.3	0.4	0.5	0.6	0.7
(A) 5 genes	0.0000	0.0000	0.0001	0.0001	0.0001
(B) 6 genes	0.00016	0.00010	0.00003	0.00006	0.00006
(C) 7 genes	0.00013	0.00013	0.00006	0.00005	0.00005
(D) 8 genes	0.0001	0.00008	0.00005	0.00002	0.00003
(E) 9 genes	0.0002	0.00001	0.00001	0.00004	0.00001
(F) 10 genes	0.0001	0.00008	0.00003	0.00002	0.00005

TABLE III
PROBABILITY OF FINDING THE BEST CONTROL GENE WITH THE INFLUENCE
METHOD WHEN $c = 1$ FOR NETWORKS WITH DIFFERENT NUMBER OF GENES

Bias	0.3	0.4	0.5	0.6	0.7
(A) 5 genes	0.6660	0.6240	0.5480	0.5670	0.5740
(B) 6 genes	0.5630	0.5320	0.4790	0.5070	0.5340
(C) 7 genes	0.5470	0.5550	0.5320	0.5460	0.5060
(D) 8 genes	0.5190	0.5290	0.5290	0.5780	0.5600
(E) 9 genes	0.5086	0.5186	0.5186	0.5676	0.5496
(F) 10 genes	0.5480	0.5230	0.5030	0.4010	0.4610

TABLE IV
AVERAGE DIFFERENCE BETWEEN PROPORTIONS OF REDUCTION IN TOTAL
PROBABILITY OF UNDESIRABLE STATES OBTAINED BY THE BEST CONTROL
GENE g^* AND PREDICTED CONTROL GENE OBTAINED BY INFLUENCE METHOD
 \check{g} FOR NETWORKS WITH VARIOUS NUMBER OF GENES

Bias	0.3	0.4	0.5	0.6	0.7
(A) 5 genes	0.0079	0.0109	0.0102	0.0133	0.0134
(B) 6 genes	0.0081	0.0107	0.0140	0.0207	0.0158
(C) 7 genes	0.0086	0.0108	0.0104	0.0115	0.0130
(D) 8 genes	0.0100	0.0137	0.0151	0.0180	0.0131
(E) 9 genes	0.0016	0.0228	0.0134	0.0415	0.0130
(F) 10 genes	0.0104	0.0097	0.0152	0.0178	0.0211

found based on (8). Similarly, the MFPT control policy for each control gene is computed and the predicted control gene \hat{g} is found based on (9). The influence method is also employed to predict the best control gene. The predicted best control gene by the influence method is called \check{g} . We define the probability of the correct prediction of each method to be the number of PBNs for which the method correctly predicts the best control gene divided by the total number of PBNs in the experiment. The probabilities of correctly predicting the best control gene by the MFPT algorithm and the influence method are shown in Tables I and III. The average differences between proportions of reduction in the total probability of undesirable states cor-

TABLE V
PROBABILITY OF FINDING THE BEST CONTROL GENE
WITH MFPT ALGORITHM

Bias	0.3	0.4	0.5	0.6	0.7
(A) $c=2$	0.9034	0.9121	0.8983	0.8848	0.9085
(B) $c=4$	0.8614	0.8897	0.8839	0.8701	0.8035

TABLE VI
AVERAGE DIFFERENCE BETWEEN PROPORTIONS OF REDUCTION IN TOTAL
PROBABILITY OF UNDESIRABLE STATES OBTAINED BY THE BEST CONTROL
GENE g^* AND PREDICTED CONTROL GENE OBTAINED BY MFPT ALGORITHM \hat{g}
WITH VARIOUS COST VALUES

Bias	0.3	0.4	0.5	0.6	0.7
(A) $c=2$	0.0004	0.0005	0.0006	0.0008	0.0005
(B) $c=4$	0.0020	0.0013	0.0014	0.0020	0.0022

TABLE VII
PROBABILITY OF FINDING THE BEST CONTROL GENE
WITH INFLUENCE METHOD

Bias	0.3	0.4	0.5	0.6	0.7
(A) $c=2$	0.6432	0.6670	0.5950	0.5755	0.6050
(B) $c=4$	0.6151	0.6247	0.6616	0.6321	0.6533

TABLE VIII
AVERAGE DIFFERENCE BETWEEN PROPORTIONS OF REDUCTION IN
TOTAL PROBABILITY OF UNDESIRABLE STATES OBTAINED
BY THE BEST CONTROL GENE g^* AND PREDICTED CONTROL GENE OBTAINED
BY INFLUENCE METHOD \check{g} WITH VARIOUS COST VALUES

Bias	0.3	0.4	0.5	0.6	0.7
(A) $c=2$	0.0098	0.0102	0.0120	0.0133	0.0144
(B) $c=4$	0.0103	0.0190	0.0151	0.0190	0.0115

responding to the gene predicted by each method and the best control gene, i.e., $(\Delta P_{g^*}^{\text{opt}} - \Delta P_{\hat{g}}^{\text{opt}})$ and $(\Delta P_{g^*}^{\text{opt}} - \Delta P_{\check{g}}^{\text{opt}})$, are shown in Tables II and IV. In our experiments, the probability of the correct prediction by the MFPT algorithm is always greater than 0.94. Table II shows that $\Delta P_{g^*}^{\text{opt}} - \Delta P_{\hat{g}}^{\text{opt}}$ on average is less than 0.0002.

The performance of the influence method is also shown in Tables III and IV. These tables suggest that approximately 0.60 of the time the influence method's prediction is correct. In general, $\Delta P_{g^*}^{\text{opt}} - \Delta P_{\check{g}}^{\text{opt}}$ is greater than 0.001. Tables V–VIII show the performance of the MFPT algorithm for higher values of c . Although the correct prediction of the MFPT algorithm slightly degrades for higher values of the control cost c , it still outperforms the influence method.

2) *Approximation of the Optimal Control Policy*: Once the best control gene g^* is known, the corresponding MFPT control policy $\hat{\mu}_{g^*}^{\gamma}$ can be used as an approximate solution to the optimal stochastic control problem. As previously explained, the parameter γ depends on the ratio of the cost of control to the cost of undesirable states. We numerically find the minimal value of the parameter γ for each control cost so that, on average, $\Gamma_{g^*}^{\text{opt}} > \Gamma_{g^*}^{\text{MFPT}(\gamma)}$. It is shown that the average of $\Delta P_{g^*}^{\text{opt}} - \Delta P_{g^*}^{\text{MFPT}(\gamma)}$ over random intervention problems with fixed control cost is small. We generate random PBNs following the procedure explained earlier. The cost of undesirable states is fixed. For the PBNs with identical bias-mean, we

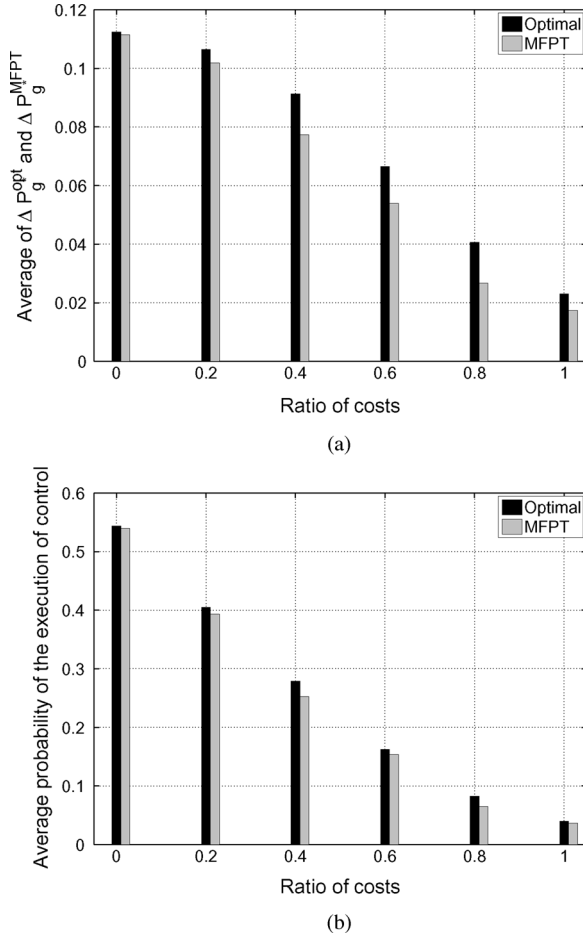


Fig. 2. (a) Average of $\Delta P_{g^*}^{\text{MFPT}(\gamma)}$ and $\Delta P_{g^*}^{\text{opt}}$. (b) Average of $\Gamma_{g^*}^{\text{MFPT}(\gamma)}$ and $\Gamma_{g^*}^{\text{opt}}$. Horizontal axis shows the ratio of the cost of control to the cost of undesirable states. Values of γ are chosen from Table IX.

TABLE IX
VALUE OF PARAMETER γ AS FUNCTION OF RATIO OF COST OF CONTROL TO COST OF UNDESIRABLE STATES

Ratio of costs	0	0.2	0.4	0.6	0.8	1
γ	0	0.29	0.61	0.91	1.5	1.94

formulate the intervention problems with various costs of control, which are varied such that the ratio of the cost of control to the cost of undesirable states changes from 0 to 1. For PBNs with each bias mean and cost of control, we compute the averages of $\Delta P_{g^*}^{\text{opt}}$ and $\Gamma_{g^*}^{\text{opt}}$. The averages are taken over 1000 intervention problems with PBNs whose bias means are fixed. Similarly, the averages of $\Delta P_{g^*}^{\text{MFPT}(\gamma)}$ and $\Gamma_{g^*}^{\text{MFPT}(\gamma)}$ are found. Furthermore, we compute the average of these averages over all bias means. The parameter γ is determined such that $\Gamma_{g^*}^{\text{MFPT}(\gamma)} < \Gamma_{g^*}^{\text{opt}}$. For each given control cost, we show the behavior of $\Delta P_{g^*}^{\text{opt}}$ and $\Gamma_{g^*}^{\text{opt}}$ ($\Delta P_{g^*}^{\text{MFPT}}$ and $\Gamma_{g^*}^{\text{MFPT}(\gamma)}$). As seen in Fig. 2(a), both $\Delta P_{g^*}^{\text{MFPT}(\gamma)}$ and $\Delta P_{g^*}^{\text{opt}}$ decrease when the ratio of the cost of control to the cost of undesirable states increases. We observe that, on average, the difference between $\Delta P_{g^*}^{\text{opt}}$ and $\Delta P_{g^*}^{\text{MFPT}(\gamma)}$ is less than 0.02. As Fig. 2(b) shows, the probability of the execution of control for both policies decreases as the cost of control increases. Table IX

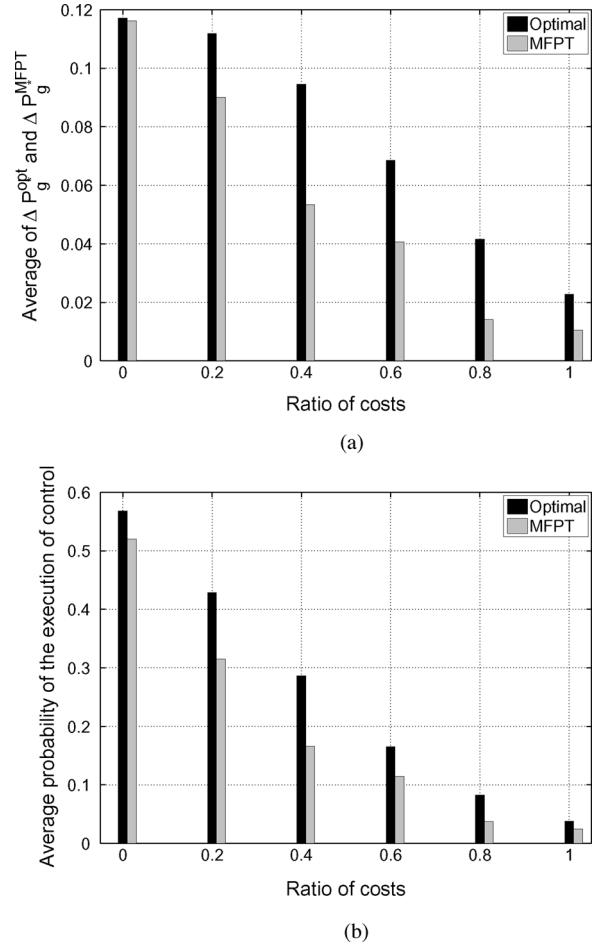


Fig. 3. (a) Average of $\Delta P_{g^*}^{\text{MFPT}(\gamma)}$ and $\Delta P_{g^*}^{\text{opt}}$. (b) Average of $\Gamma_{g^*}^{\text{MFPT}(\gamma)}$ and $\Gamma_{g^*}^{\text{opt}}$. Horizontal axis shows the ratio of the cost of control to the cost of undesirable states. Values of γ are chosen conservatively from Table X.

TABLE X
CONSERVATIVE VALUE OF THE PARAMETER γ AS A FUNCTION OF THE RATIO OF THE COST OF CONTROL TO THE COST OF UNDESIRABLE STATES

Ratio of costs	0	0.2	0.4	0.6	0.8	1
γ	0.05	0.5	0.9	1.1	1.9	2.3

shows the relation of the parameter γ with the ratio of the cost of control to the cost of undesirable states found in the earlier experiment. The extensive results of the simulations at various bias means can be found in the supplementary materials.¹

Since the values in Table IX are found from random PBNs, one can have a conservative approach and choose the parameter γ to be greater than the proposed values. To this end, $\Gamma_{g^*}^{\text{MFPT}(\gamma)}$ is smaller than $\Gamma_{g^*}^{\text{opt}}$ in each intervention problem. Fig. 3 and Table X show the outcomes of the same experiment explained earlier when the parameter γ is chosen conservatively. In all the intervention problems of this experiment, $\Gamma_{g^*}^{\text{MFPT}(\gamma)} < \Gamma_{g^*}^{\text{opt}}$ and the deviation of $\Delta P_{g^*}^{\text{MFPT}(\gamma)}$ from $\Delta P_{g^*}^{\text{opt}}$ is smaller than

¹This paper has supplementary downloadable material available at <http://ieeexplore.ieee.org>, provided by the authors. This includes a pdf file of extensive simulation results. This material is 64 kB in size.

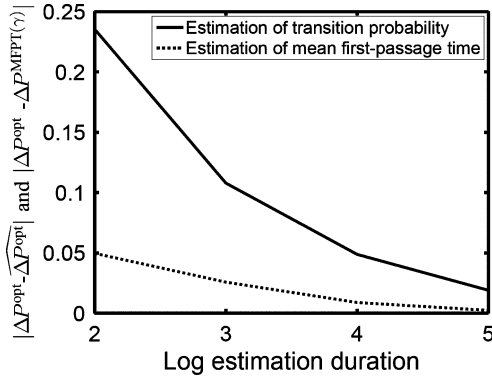


Fig. 4. Average of $|\Delta P_{g^*}^{\text{opt}} - \widehat{\Delta P_{g^*}^{\text{opt}}}|$ (solid) and $|\Delta P_{g^*}^{\text{opt}} - \Delta P_{g^*}^{\text{MFPT}(\gamma)}|$ (dash) over 1000 trials as a function of the logarithm of estimation duration.

0.04. Results of this experiment for the networks with various numbers of genes are presented in the supplementary materials.

3) *Controllability*: To corroborate that the MFPT algorithm can be employed to determine the controllability of a network, we consider the results in Fig. 2. In this figure, when the cost of control is zero ($\gamma = 0$), $\Delta P_{g^*}^{\text{MFPT}(0)}$ is an accurate approximation of $\Delta P_{g^*}^{\text{opt}}$. The average of the difference $\Delta P_{g^*}^{\text{MFPT}(0)} - \Delta P_{g^*}^{\text{opt}}$ has a negligible value equal to 0.0007.

4) *Model-Free Intervention*: To compare the performance of the model-free MFPT control algorithm with an optimal control algorithm, where the latter includes estimation of the transition probability matrix, we generate synthetic time-course data for 100 000 time steps from an existing model. Using the synthetic time-course data, we estimate the MFPTs after each 10^k time steps, for $k = 2, \dots, 5$, and fix the cost of control to have the value 1. As the duration of estimating the MFPTs increases, $\Delta P_{g^*}^{\text{MFPT}(\gamma)}$ approaches $\Delta P_{g^*}^{\text{opt}}$. Fig. 4 shows the average of $|\Delta P_{g^*}^{\text{opt}} - \Delta P_{g^*}^{\text{MFPT}(\gamma)}|$, where $\Delta P_{g^*}^{\text{opt}}$ is obtained from the original transition probabilities, with various estimating durations over 1000 trials. For an optimal control policy based on the Markov chain estimated from the data, we denote the shift in the steady-state distribution by $\widehat{\Delta P_{g^*}^{\text{opt}}}$. Fig. 4 shows the average of $|\Delta P_{g^*}^{\text{opt}} - \widehat{\Delta P_{g^*}^{\text{opt}}}|$ with various estimating durations over 1000 trials. The graphs clearly demonstrate the superior performance of the model-free approach using the MFPT algorithm. In particular, for lower numbers of observations, which correspond better to feasible experimental conditions, estimating the Markov chain yields poor results, whereas the MFPT approximation performs quite well. We have conducted the same experiment with various costs of the control for networks with different numbers of genes, and these results can be found in the supplementary materials.

B. Melanoma Gene Expression

In this section, we compare the performances of optimal and MFPT control policies in the context of a gene network developed from steady-state data. These steady-state data were collected in a profiling study of metastatic melanoma in which the abundance of messenger RNA for the gene WNT5A was

found to be highly discriminating between cells with properties typically associated with high metastatic competence versus those with low metastatic competence [2]. These findings were validated and expanded in a second study, in which experimentally increasing the levels of the WNT5A protein secreted by a melanoma cell line via genetic engineering methods directly altered the metastatic competence of that cell as measured by the standard *in vitro* assays for metastasis [30]. A further finding of interest in this study was that an intervention that blocked the WNT5A protein from activating its receptor, the use of an antibody that binds the WNT5A protein, can substantially reduce WNT5A's ability to induce a metastatic phenotype. This suggests control based on intervention that alters the contribution of the WNT5A gene to biological regulation. Disruption of this influence can potentially reduce the chance of a melanoma metastasizing, a desirable outcome. Ten genes, including the WNT5A gene, were selected in [31] based on the predictive relationships among 587 genes: WNT5A, pirin, S100P, RET1, MMP3, PHOC, MART1, HADHB, Synuclein, and STC3. We apply the design procedure proposed in [32] to generate a PBN possessing four constituent BNs. The method of [32] generates BNs with given attractor structures and the overall PBN is designed so that the data points, which are assumed to come from the steady-state distribution of the network, are attractors in the designed PBN. The regulatory graphs of these BNs can be found in the supplementary materials. This approach is reasonable because our interest is in controlling the long-run behavior of the network. The control objective for this ten-gene network is to down-regulate the WNT5A gene, because WNT5A ceasing to be down-regulated is strongly predictive of the onset of metastasis. A number of other control studies based on the same data have aimed to down-regulate the WNT5A gene. This model has been used because the relation of WNT5A to metastasis is well established and the binary nature of the up- or down-regulation suits a binary model. A state is desirable if WNT5A = 0 and undesirable if WNT5A = 1. As we mentioned earlier, the application of the MFPT algorithm, or any of the other methods developed for the control of gene regulatory networks, requires the designation of desirable and undesirable states, and this depends upon the existence of relevant biological knowledge. In this example, the use of the state WNT5A has resulted from biological knowledge relating the state of WNT5A to metastasis in melanoma tumors. Based on our objective, the cost of control is assumed to be 1 and the states are assigned penalties according to the following scheme:

$$r(u, j) = \begin{cases} 0, & \text{if } u = 0 \text{ and } j \in \mathcal{D} \\ 5, & \text{if } u = 0 \text{ and } j \in \mathcal{U} \\ 1, & \text{if } u = 1 \text{ and } j \in \mathcal{D} \\ 6, & \text{if } u = 1 \text{ and } j \in \mathcal{U} \end{cases}$$

which is the same cost structure as assumed in [7]. Since our objective is to down-regulate the WNT5A gene, a higher penalty is assigned for destination states having WNT5A up-regulated. Also, for a given WNT5A status for the destination state, a higher penalty is assigned when the control is active versus when it is not. Note that the cost scheme reflects our objective;

TABLE XI
 ΔP_g^{OPT} AND $\Delta P_g^{\text{MFPT}(\gamma)}$ FOR ALL CONTROL GENES g IN MELANOMA
 CASE STUDY

Gene (g)	STC2	Synuclein	HADHB	MART1	PHOC	MMP3	RET1	S100P	pirin
ΔP_g^{OPT}	0.0733	0.0892	0.1453	0.1104	0.2325	0.1121	0.0529	0.1032	0.1305
$\Delta P_g^{\text{MFPT}(\gamma)}$	0.0721	0.0824	0.1437	0.1071	0.2312	0.1120	0.0507	0.1021	0.1272

TABLE XII
 COMPARISON OF CONTROL GENE RANKING BASED ON $\Delta P_{g^*}^{\text{OPT}}$, $\Delta P_{\hat{g}}^{\text{OPT}}$,
 AND $\Delta P_{\hat{g}}^{\text{OPT}}$

Rank	1	2	3	4	5	6	7	8	9
Optimal	PHOC	HADHB	pirin	MMP3	MART1	S100P	Synuclein	STC2	RET1
MFPT	PHOC	HADHB	pirin	MMP3	MART1	S100P	Synuclein	STC2	RET1
Influence	MMP3	HADHB	MART1	S100P	STC2	pirin	PHOC	RET1	Synuclein

in practice, the actual values would have to be assigned by a physician according to his or her understanding of the disease. Optimal and MFPT control policies are found for the melanoma-related PBN. Table XI summarizes the amount of the shift in the total probability mass of the undesirable states obtained by each of these two methods. We apply the influence method to predict the best control gene. We then compare the prediction of the influence method with the prediction of the MFPT algorithm and the optimal gene determined directly by the solution of a dynamic programming algorithm. Table XII shows the ranking of the genes based on direct solution of the optimal control policy, the MFPT algorithm, and the influence method. The MFPT method not only predicts the best control gene, but it also exactly predicts the ranking of the control genes. As Table XII shows, the influence method does a poor job on predicting the best control gene.

VII. CONCLUSION

To overcome the computational impediments to optimal stochastic control, we proposed an algorithm based on MFPTs to address questions regarding the controllability in the context of Markovian gene regulatory networks. We have compared its performance with the results from optimal stochastic control. The direct comparison has been limited to ten-gene networks on account of the high complexity of the optimal algorithm. The high accuracy and the low complexity of the MFPT algorithm make it appealing for various applications: identification of the best control gene, approximation of an optimal control policy, controllability, and model-free intervention.

REFERENCES

- [1] I. Shmulevich, E. R. Dougherty, and W. Zhang, "Gene perturbation and intervention in probabilistic Boolean networks," *Bioinformatics*, vol. 18, no. 10, pp. 1319–1331, 2002.
- [2] M. Bittner, P. Meltzer, Y. Chen, Y. Jiang, E. Seftor, M. Hendrix, M. Radmacher, R. Simon, Z. Yakhini, A. Ben-Dor, N. Sampas, E. Dougherty, E. Wang, F. Marincola, C. Gooden, J. Lueders, A. Glatfelter, P. Pollock, J. Carpten, E. Gillanders, D. Leja, K. Dietrich, C. Beaudry, M. Berens, D. Alberts, V. Sondak, N. Hayward, and J. Trent, "Molecular classification of cutaneous malignant melanoma by gene expression profiling," *Nature*, vol. 406, no. 6795, pp. 536–450, 2000.
- [3] A. Datta, R. Pal, A. Choudhary, and E. R. Dougherty, "Control approaches for probabilistic gene regulatory networks," *IEEE Signal Process. Mag.*, vol. 24, no. 1, pp. 54–63, Jan. 2007.
- [4] I. Shmulevich, E. R. Dougherty, S. Kim, and W. Zhang, "Probabilistic Boolean networks: A rule-based uncertainty model for gene regulatory networks," *Bioinformatics*, vol. 18, no. 2, pp. 261–274, 2002.
- [5] A. Datta, A. Choudhary, M. Bittner, and E. R. Dougherty, "External control in Markovian genetic regulatory networks," *Mach. Learn.*, vol. 52, no. 1/2, pp. 169–191, 2003.
- [6] A. Datta, R. Pal, and E. R. Dougherty, "Intervention in probabilistic gene regulatory networks," *Current Bioinformatics*, vol. 1, no. 2, pp. 167–184, 2006.
- [7] R. Pal, A. Datta, and E. R. Dougherty, "Optimal infinite-horizon control for probabilistic Boolean networks," *IEEE Trans. Signal Process.*, vol. 54, no. 6, pp. 2375–2387, Jun. 2006.
- [8] S. Huang, "Gene expression profiling, genetic networks, and cellular states: An integrating concept for tumorigenesis and drug discovery," *J. Mol. Med.*, vol. 77, no. 6, pp. 469–480, 1999.
- [9] T. Akutsu, M. Hayashida, W.-K. Ching, and M. K. Ng, "Control of Boolean networks: Hardness results and algorithms for tree structured networks," *J. Theor. Biol.*, vol. 244, no. 4, pp. 670–679, 2007.
- [10] I. Shmulevich, E. Dougherty, and W. Zhang, "Control of stationary behavior in probabilistic Boolean networks by means of structural intervention," *Biol. Syst.*, vol. 10, no. 4, pp. 431–446, 2002.
- [11] T. Miyashita and J. C. Reed, "Tumor-suppressor p53 is a direct transcriptional activator of the human bax gene," *Cell*, vol. 80, no. 2, pp. 293–299, 1995.
- [12] L. B. Owenschaub, W. Zhang, J. C. Cusack, L. S. Angelo, S. M. Santee, T. Fujiwara, J. A. Roth, A. B. Deisseroth, W. W. Zhang, E. Kruzel, and R. Radinsky, "Wild-type human p53 and a temperature-sensitive mutant induce Fas/Apo-1 expression," *Mol. Cell. Biol.*, vol. 15, no. 6, pp. 3032–3040, 1995.
- [13] W. S. Eldeiry, T. Tokino, V. E. Velculescu, D. B. Levy, R. Parsons, J. M. Trent, D. Lin, W. E. Mercer, K. W. Kinzler, and B. Vogelstein, "Waf1, a potential mediator of p53 tumor suppression," *Cell*, vol. 75, no. 4, pp. 817–825, 1993.
- [14] C. B. Harley, "Telomere loss—Mitotic clock or genetic time bomb," *Mutat. Res.*, vol. 256, no. 2–6, pp. 271–282, 1991.
- [15] N. W. Kim, M. A. Piatyszek, K. R. Prowse, C. B. Harley, M. D. West, P. L. C. Ho, G. M. Coviello, W. E. Wright, S. L. Weinrich, and J. W. Shay, "Specific association of human telomerase activity with immortal cells and cancer," *Science*, vol. 266, no. 5193, pp. 2011–2015, 1994.
- [16] E. Hiyama, K. Hiyama, T. Yokoyama, Y. Matsuura, M. A. Piatyszek, and J. W. Shay, "Correlating telomerase activity levels with human neuroblastoma outcomes," *Nature Med.*, vol. 1, no. 3, pp. 249–255, 1995.
- [17] W. Zhang, M. A. Piatyszek, T. Kobayashi, E. Estey, M. Andreeff, A. B. Deisseroth, W. E. Wright, and J. W. Shay, "Telomerase activity in human acute myelogenous leukemia: Inhibition of telomerase activity by differentiation-inducing agents," *Clin. Cancer Res.*, vol. 2, no. 5, pp. 799–803, 1996.
- [18] S. W. Lowe, E. M. Schmitt, S. W. Smith, B. A. Osborne, and T. Jacks, "p53 is required for radiation-induced apoptosis in mouse thymocytes," *Nature*, vol. 362, no. 6423, pp. 847–849, 1993.
- [19] N. Friedman, K. Murphy, and S. Russell, "Learning the structure of dynamic probabilistic networks," in *Proc. 15th Nat. Conf. Artif. Intell. AAAI 1998*, Madison, WI, pp. 139–147.
- [20] H. Lahdesmaki, S. Hautaniemi, I. Shmulevich, and O. Yli-Harja, "Relationships between probabilistic Boolean networks and dynamic Bayesian networks as models of gene regulatory networks," *Signal Process.*, vol. 86, no. 4, pp. 814–834, 2006.
- [21] D. P. Bertsekas, *Dynamic Programming and Optimal Control*. Belmont, MA: Athena Scientific, 2005.
- [22] J. R. Norris, *Markov Chains*. Cambridge, U.K.: Cambridge Univ. Press, 1998.
- [23] S. Marshall, L. Yu, Y. Xiao, and E. R. Dougherty, "Inference of a probabilistic Boolean network from a single observed temporal sequence," *EURASIP J. Bioinf. Syst. Biol.*, vol. 2007, pp. 32 454–32 569, 2007.
- [24] B. Faryabi, A. Datta, and E. Dougherty, "On approximate stochastic control in genetic regulatory networks," *IET Syst. Biol.*, vol. 1, no. 6, pp. 361–368, 2007.
- [25] V. D. Blondel and J. N. Tsitsiklis, "A survey of computational complexity results in systems and control," *Automatica*, vol. 36, no. 9, pp. 1249–1274, 2000.
- [26] O. Madani, "Complexity results for infinite-horizon markov decision processes," Ph.D. dissertation, Dept. Comp. Sci. Eng., Univ., Washington, Seattle, 2000.
- [27] P. Tseng, "Solving h-horizon, stationary markov decision-problems in time proportional to log(h)," *Oper. Res. Lett.*, vol. 9, no. 5, pp. 287–297, 1990.
- [28] E. Feinberg and A. Shwartz, *Handbook of Markov Decision Processes*. Boston, MA: Kluwer, 2002.

- [29] I. Shmulevich and E. R. Dougherty, *Genomic Signal Processing*. Princeton, NJ: Princeton Univ. Press, 2007.
- [30] A. T. Weeraratna, Y. A. Jiang, G. Hostetter, K. Rosenblatt, P. Duray, M. Bittner, and J. M. Trent, "Wnt5a signaling directly affects cell motility and invasion of metastatic melanoma," *Cancer Cell*, vol. 1, no. 3, pp. 279–288, 2002.
- [31] S. Kim, H. Li, E. R. Dougherty, N. W. Cao, Y. D. Chen, M. Bittner, and E. B. Suh, "Can Markov chain models mimic biological regulation?," *J. Biol. Syst.*, vol. 10, no. 4, pp. 337–357, 2002.
- [32] R. Pal, I. Ivanov, A. Datta, M. Bittner, and E. R. Dougherty, "Generating Boolean networks with a prescribed attractor structure," *Bioinformatics*, vol. 21, no. 21, pp. 4021–4025, 2005.

Golnaz Vahedi (S'02) received the B.Sc. degree in electrical engineering from Sharif University of Technology, Tehran, Iran, in 2001, and the M.Sc. degree in electrical and computer engineering from the University of Alberta, Edmonton, Canada, in 2004. She is currently working toward the Ph.D. degree in electrical and computer engineering at Texas A&M University, College Station.

Her current research interests include systems biology, genomic signal processing, and inference and control of genetic regulatory networks.

Babak Faryabi (S'03) received the M.Sc. degree in electrical and computer engineering from Sharif University of Technology, Tehran, Iran, in 1997. He is currently working toward the Ph.D. degree at Texas A&M University, College Station.

He was a Research Assistant at the University of Toronto, Toronto, Canada. From 2003 to 2005, he was a recipient of the Edward S. Rogers Scholarship from the Department of Electrical and Computer Engineering, University of Toronto. His current research interests include the general area of systems biology, especially in the applications of Markovian decision processes and their approximations in the modeling and control of gene regulatory networks.

Mr. Faryabi is a member of the International Society for Computational Biology.

Jean-Francois Chamberland (S'98–M'04) received the B.E. degree from McGill University, Montreal, Canada, in 1998, the M.Sc. degree from Cornell University, Ithaca, NY, in 2000, and the Ph.D. degree from the University of Illinois, Urbana-Champaign, in 2004, all in electrical engineering.

During 2004, he was with Texas A&M University, College Station, where he is currently an Assistant Professor in the Department of Electrical and Computer Engineering. His current research interests include communication systems, queueing theory, detection and estimation, and genomic signal processing.

Dr. Chamberland was the recipient of a Young Author Best Paper Award from the IEEE Signal Processing Society in 2006.

Aniruddha Datta (S'87–M'91–SM'97) received the B.Tech. degree in electrical engineering from the Indian Institute of Technology-Kharagpur, Kharagpur, India, in 1985, the M.S.E.E. degree from the Southern Illinois University, Carbondale, in 1987, and the M.S. and Ph.D. degrees in applied mathematics from the University of Southern California, Los Angeles, in 1991.

During August 1991, he was with the Department of Electrical and Computer Engineering, Texas A&M University, College Station, where he is currently a Professor and also holds the J. W. Runyon, Jr., '35 Professorship II. His current research interests include adaptive control, robust control, PID control, and genomic signal processing. From 2001 to 2003, he was engaged in cancer research as a Postdoctoral Trainee under a National Cancer Institute (NCI) Training Grant. He is the author of the book *Adaptive Internal Model Control* (Springer-Verlag, 1998) and a coauthor of the books *Structure and Synthesis of PID Controllers* (Springer-Verlag, 2000), *PID Controllers for Time Delay Systems* (Birkhauser, 2004), and *Introduction to Genomic Signal Processing With Control* (CRC, 2007).

Prof. Datta was an Associate Editor of the IEEE TRANSACTIONS ON AUTOMATIC CONTROL from 2001 to 2003 and the IEEE TRANSACTIONS ON SYSTEMS, MAN AND CYBERNETICS, PART B from 2005 to 2006.

Edward R. Dougherty (M'05) received the M.S. degree in computer science from Stevens Institute of Technology, Hoboken, NJ, the Ph.D. degree in mathematics from Rutgers University, New Brunswick, NJ, and the Doctor Honoris Causa from the Tampere University of Technology, Hervanta, Finland, in 2007.

He is currently a Professor in the Department of Electrical and Computer Engineering, Texas A&M University, College Station, where he is also the Robert M. Kennedy Chair, the Director of the Genomic Signal Processing Laboratory, and the Halliburton Professor of Dwight Look College of Engineering. He is the Director of the Computational Biology Division, Translational Genomics Research Institute, Phoenix, AZ. He is the author or coauthor of 14 books and more than 200 journal papers, and an editor of five books. He has contributed extensively to the statistical design of nonlinear operators for image processing and the consequent application of pattern recognition theory to nonlinear image processing. His current research interests include genomic signal processing that is focussed on diagnosis and prognosis based on genetic signatures and using gene regulatory networks to develop therapies based on the disruption or mitigation of aberrant gene function contributing to the pathology of a disease.

Prof. Dougherty is a Fellow of the International Society for Optical Engineers (SPIE) and a recipient of the SPIE President's Award. At Texas A&M, he received the Association of Former Students Distinguished Achievement Award in Research, and was named a Fellow of the Texas Engineering Experiment Station. He was the editor of the *SPIE/IS&T Journal of Electronic Imaging*.