Robust Intervention in Probabilistic Boolean Networks

Ranadip Pal, *Member, IEEE*, Aniruddha Datta, *Senior Member, IEEE*, and Edward R. Dougherty, *Member, IEEE*

***Abstract*—Probabilistic Boolean networks (PBNs) have been recently introduced as a paradigm for modeling genetic regulatory networks. One of the objectives of PBN modeling is to use the network for the design and analysis of intervention strategies aimed at moving the network out of undesirable states, such as those associated with disease, and into desirable ones. To date, a number of intervention strategies have been proposed in the context of PBNs. However, all these techniques assume perfect knowledge of the transition probability matrix of the PBN. Such an assumption cannot be satisfied in practice since the presence of noise and the availability of limited number of samples will prevent the transition probabilities from being accurately determined. Moreover, even if the exact transition probabilities could be estimated from the data, mismatch between the PBN model and the actual genetic regulatory network will invariably be present. Thus, it is important to study the effect of modeling errors on the final outcome of an intervention strategy and one of the goals of this paper is to do precisely that when the uncertainties are in the entries of the transition probability matrix. In addition, the paper develops a robust intervention strategy that is obtained by minimizing the worst-case cost over the uncertainty set.**

***Index Terms*—Control of biological networks, estimation errors, robust dynamic programming, robust minimax control, perturbation bounds.**

I. Introduction

THE sequencing of various genomes over the last decade has given a remarkable boost to genomic studies. The improved understanding of the genomes of various organisms, along with advances in microarray technology, have provided us with enormous opportunities for the mathematical modeling of biological networks. There are two major objectives for modeling of genetic regulatory networks: first, to better understand the intergene interactions and relationships on a holistic level, thereby facilitating the diagnosis of disease; and second, to design and analyze therapeutic intervention strategies for shifting the state of a diseased network from an undesirable location to a desirable one. The first objective falls within the scope of the field known as *Systems Biology*, while the second objective falls within the scope of the field known as *Systems Medicine*. Systems medicine approaches that make use of genome based systems engineering fall within the scope of the field known as *Translational Genomics*. This paper focuses on a problem that arises in Translational Genomics.

In order to set the stage for introducing this problem, we next present a broad overview of the steps involved in the modeling and control of genetic networks. These steps are shown in Fig. 1. The first step consists of data extraction, which basically involves signal acquisition, the signals in this case being the expression levels of various genes of interest. The next step denoted by A_1 involves the discretization of these gene expression levels. Obviously, this step is not required if we are interested in arriving at a continuous model such as a differential equation model. On the other hand, this step would be crucial for discrete models such as Boolean networks (BNs) [1]–[3], probabilistic Boolean networks (PBNs) [4], [5] and Bayesian networks [6], all of which have been proposed as models for genetic regulatory networks. The next step denoted by B involves the selection of a small set of genes to be used in constructing the genetic regulatory network. This step is necessary because of at least two reasons: first, building a network of thousands of genes would require an inordinately large amount of data for inference purposes, not to mention the computational intractability of the resulting network; and second, while modeling a particular biological pathway only a few genes may be playing an important role. Having selected the genes, the next step is the actual construction of the network, and a number of approaches can be used to carry out such construction [7]–[9].

Since systems biology is focussed on understanding the detailed molecular interactions that contribute to the functioning of a cell, a genetic regulatory network designed for facilitating such an understanding must necessarily mimic the actual biological interactions in as much detail as possible. On the other hand, in translational genomics the focus is on developing therapeutic interventions, and the network used for this purpose can be a coarse representation of the biological phenomena occurring at the molecular level as long as it has the capability to faithfully capture the overall effects of intervention that are manifested at the phenotypic (observational) level. Such a coarse model can then be used to develop and evaluate suitable (control) strategies for therapeutic intervention. PBNs (PBNs), which constitute one class of coarse models, have been used in recent years to carry out such intervention studies [10]–[14].

From Fig. 1, it is clear that errors made during data extraction, discretization, gene selection and network generation will

Manuscript received September 22, 2006; revised July 25, 2007. This work was supported in part by the National Science Foundation under Grants ECS0355227 and CCF0514644 and by the National Cancer Institute under Grant CA90301. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Elias S. Manolakos.

R. Pal is with the Department of Electrical and Computer Engineering, Texas Tech University, Lubbock, TX 79409 USA (e-mail: ranadip.pal@ttu.edu).

A. Datta is with the Department of Electrical and Computer Engineering, Texas A&M University, College Station, TX 77843 USA.

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.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TSP.2007.908964

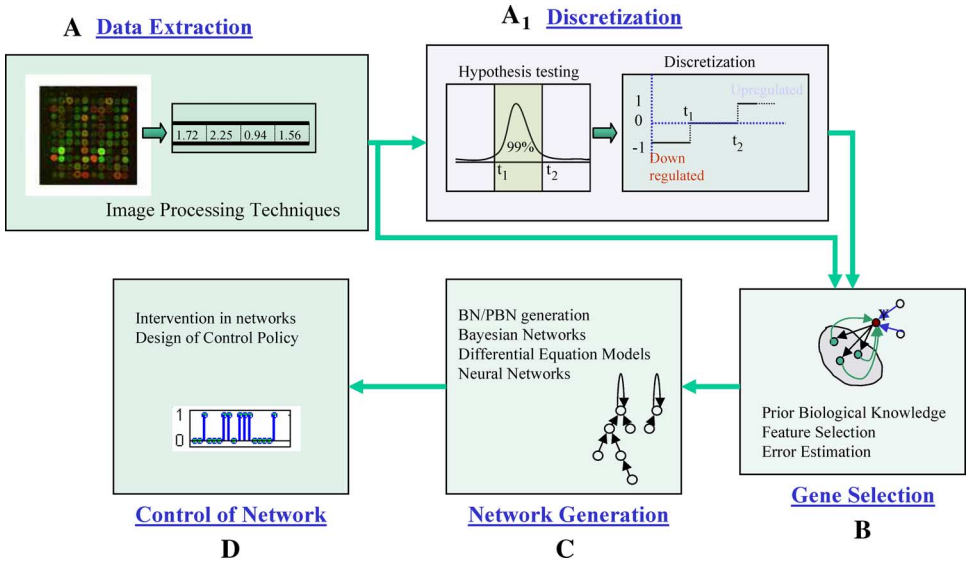

Fig. 1. Basic steps involved in modeling and control of genetic networks. (A) Extraction of gene expression data (A_1) discretization of the Data. (B) Selection of genes to build the network. (C) Generate network from the available data and prior biological knowledge. (D) Intervention in the network with the objective of moving the network from undesirable to desirable states.

all propagate downstream and impact the actual success of the designed intervention strategy. Indeed, if the designed intervention approach is to have any hope of succeeding in practice, its outcome must possess some degree of "robustness" or insensitivity to the errors that will invariably propagate down to the intervention design stage from steps further upstream. The study of the effect on intervention[1] outcome of the errors propagating from the different upstream steps is an important open problem in translational genomics. In this paper, we focus on a special subproblem where it is assumed that the combined effect of the errors propagating from the different stages manifests itself as uncertainty in the transition probabilities of the network, and the robustness of the intervention strategies is to be studied with respect to this uncertainty. With respect to Fig. 1, this corresponds to determining how the uncertainties in Step C impact the outcome of the intervention strategy designed in Step D. Besides error propagation, uncertainties arise due to the inverse problem of system identification being an ill-posed problem. We will also examine how one can go about designing an intervention strategy that optimizes the worst-case performance index over a given uncertainty class.

To date, a number of approaches have been proposed in the literature for carrying out interventions in PBNs [10]–[15]. Of these, the one proposed in [14] is of particular relevance for translational genomics since it seeks to shift the steady-state mass of the PBN from undesirable states to desirable ones. Since it is believed that the steady-state behavior of a PBN is indicative of the phenotype [3], it is likely that alterations in the steady-state behavior of the PBN would translate into changes at the phenotypic level. Moreover, the scheme of [14] makes use of stationary policies which are more easily impementable. In that paper, we have used gene expression data from melanoma studies to demonstrate the feasibility of altering the steady-state distribution of a PBN from states assumed to be representative of metastatic cancer to nonmetastatic states. However, such alteration in steady-state behavior was achieved under the assumption that the transition probability matrix of the PBN is known. Such an assumption will not be satisfied for reasons that we have already articulated. Instead, while the intervention strategy would have to be designed based on an estimated network with a transition probability matrix P, in practice it would be applied to the actual network whose transition probability $\tilde{P}$ differs from P. One of the goals of this paper is to examine how, for a given intervention policy, the mismatch between P and $\tilde{P}$ affects the steady-state distribution of the *controlled* PBN. A second objective is to examine how, given a characterization of the mismatch between P and $\tilde{P}$, one can design an optimal intervention strategy that minimizes the cost (or penalty function) corresponding to the worst-case uncertainty from the particular uncertainty set. Although worst-case (*minimax*) designs tend to be inherently conservative, a worst-case approach is reasonable here since in translational genomics, one is dealing with life and death situations.

The paper is organized as follows. In Section II, we formally define BNs and PBNs. In Section III, we examine how the perturbations in the probability transition matrix of an uncontrolled PBN affect the steady-state distribution resulting from the application of a given stationary intervention policy to that PBN. In one case, an analytical perturbation bound is derived, and in other cases, studies are carried out via simulations. In Section IV, we consider context-sensitive PBNs with perturbation and examine a number of different possibilities by which uncertainties can arise in the probability transition matrix of such a PBN. Section V presents the design of optimal intervention strategies for the worst-case scenario. Three different characterizations of the uncertainty are considered. Section VI presents two examples of minimax (worst-case) intervention design. Finally, Section VII contains some concluding remarks.

[1]The goal of the intervention strategy is to move the network from undesirable states (such as states representative of cancer) to normal states.

II. BNs and PBNs

In this section, we briefly define BNs and PBNs. For a more detailed and motivated development, the reader is referred to [3], [5].

A BN $B = (V, F)$ on n genes is defined by a set of nodes/genes $V = \{x_1, \ldots, x_n\}$, $x_i \in \{0,1\}$, $i = 1, \ldots, n$, and a list $F = (f_1, \ldots, f_n)$, of Boolean functions, $f_i : \{0,1\}^n \to \{0,1\}$, $i = 1, \ldots, n$. Each node x_i represents the state/expression of the gene x_i, where $x_i = 0$ means that gene i is OFF and $x_i = 1$ means that gene i is ON. The function f_i is called the *predictor function* for gene i. Updating the states of all genes in B is done synchronously at every time step according to their predictor functions. A PBN consists of a set of nodes/genes $V = \{x_1, \ldots, x_n\}$, $x_i \in \{0, 1, \ldots, Y\}$, $i = 1, \ldots, n$, $Y \in \mathbb{N}$ and a set of vector-valued network functions, $\mathbf{f}_1, \mathbf{f}_2, \ldots, \mathbf{f}_k$, governing the state transitions of the genes, each network function being of the form $\mathbf{f}_j = (f_{j1}, f_{j2}, \ldots, f_{jn})$, where $f_{ji} : \{0, 1, \ldots, Y\}^n \to \{0, 1, \ldots, Y\}$, $i = 1, \ldots, n$. In most applications, the discretization is either binary or ternary and in this paper we will use binary, i.e., $Y = 1$. The choice of which network function $\mathbf{f}_j$ to apply is governed by a selection procedure. Specifically, at each time point a random decision is made as to whether to switch the network function for the next transition, with the probability q of a switch being a system parameter. If a decision is made to switch the network function, then a new function is chosen from among $\mathbf{f}_1, \mathbf{f}_2, \ldots, \mathbf{f}_k$, with the probability of choosing $\mathbf{f}_j$ being the selection probability c_j. In other words, each network function $\mathbf{f}_j$ determines a BN and the PBN behaves as a fixed BN until a random decision (with probability q) is made to change the network function according to the probabilities $c_1, c_2, \ldots, c_k$ from among $\mathbf{f}_1, \mathbf{f}_2, \ldots, \mathbf{f}_k$. The PBN just described is called a *context-sensitive PBN*. In the special case when $q = 1$, the network function is switched at every time point and the PBN is called an *instantaneously random PBN*. We consider context-sensitive PBNs with perturbation, meaning that at each time point there is a probability p of any gene flipping its value uniformly randomly. Since there are n genes, the probability of there being a random perturbation at any time point is $1 - (1 - p)^n$. The state space S of the network together with the set of network functions, in conjunction with transitions between the states and network functions, determines a Markov chain. The random perturbation makes the Markov chain irreducible, meaning that it has the possibility of reaching any state from any other state and that it possesses a steady-state distribution.

The state vector $x(t)$ at any time step t is essentially an n-digit binary number $[x_1 x_2 \cdots x_n]$ whose decimal equivalent is given by

$$z(t) = \sum_{j=1}^{n} 2^{n-j} x_j(t). \tag{1}$$

As $x(t)$ ranges from $000\cdots0$ to $111\cdots1$, $z(t)$ takes on all values from 0 to $2^n - 1$. Instead of the vector $x(t)$, one could equivalently consider $z(t)$ to be the state of the network at time t so that the new state space becomes $S = \{0, 1, 2, \cdots, 2^n - 1\}$.

III. Perturbations for the Steady-State Distribution of a Controlled PBN

We are interested in studying the effect of mismatch between P and $\tilde{P}$ on the steady state of the controlled PBN. Consequently, we would like to derive perturbation bounds, which in this context, refer to bounds on the shift of the steady state distribution resulting from the mismatch. Before trying to derive any perturbation bounds, let us make a simple observation concerning the probability transition matrix of a controlled PBN. Note that the only kind of interventions that have been proposed to date in the literature are restricted to flipping the expression status of one or more control genes. For such intervention strategies, it is always possible to relate the transition probability matrices of the controlled and uncontrolled PBNs, via a linear transformation, as we explain next.

Let P denote the estimated probability transition matrix corresponding to the PBN of interest and suppose this PBN has m binary control inputs $a_1, a_2 \cdots a_m$ where a_i refers to the status of the ith control gene with $a_i = 1$ signifying that the ith control gene is to be flipped.[2] If we apply a stationary policy, i.e., a policy dependent only on the current state and not on the time, to the Markov Chain P, the rows of the controlled transition probability matrix P_c will be a collection of selected rows from P. This is due to the fact that the flipping of genes actually forces the Markov Chain to start from another initial state. To clearly understand this, let us look at a concrete example.

Example III.1: Suppose we have a network with seven genes, three of which, namely genes 1, 2, and 3 are control genes. This means that $m = 3$ here. Suppose that the stationary policy for state 0000001 (corresponding to the decimal number 1) is 101, i.e., flip gene 1, leave gene 2 as is, and flip gene 3. This implies that if we are currently at state 0000001, application of the stationary policy will reinitialize the state to 1010001 (corresponding to the decimal number 81). Therefore, in the controlled transition probability matrix P_c, the transition probabilities of going from state 1 to each of the other states will be the same as the transition probabilities of going from state 81 to each of those states in the original uncontrolled network with transition probability matrix P.

From the above example, it is clear that when the class of allowed interventions is restricted to the flipping of genes, the application of a stationary policy converts the uncontrolled transition probability matrix P to a controlled transition probability matrix P_c where P_c and P are related by $P_c = TP$ and T represents a matrix which has only one nonzero entry of 1 in each row. If the stationary policy is of no control, then clearly $T = I$, the identity matrix. For Example III.1, the second row of T consists of all zeros except for the 82nd entry which is 1.

Let π and π_c denote the stationary distribution vectors corresponding to the transition matrices P and P_c, respectively. Since the probability transition matrix P has been estimated from data, there can be some errors in estimation. Let $\tilde{P}$ denote the actual transition matrix of the genetic network and let $\tilde{P}_c$ denote the controlled transition probability matrix that results from

[2]Here flipping refers to toggling the value of the gene, i.e., changing 1 to 0 and 0 to 1.

the application of the stationary policy T on $\tilde{P}$. Let $\tilde{\pi}$ and $\tilde{\pi}_c$ denote the stationary distributions of $\tilde{P}$ and $\tilde{P}_c$, respectively. Our goal is to study the change $\tilde{\pi}_c - \pi_c$ based on the knowledge of P and some characterization of the estimation error $E \triangleq P - \tilde{P}$. Let us summarize the notation and relationships introduced so far:

(i) $\pi P = \pi$
(ii) $\tilde{\pi}\tilde{P} = \tilde{\pi}$
(iii) $\pi_c P_c = \pi_c$
(iv) $\tilde{\pi}_c \tilde{P}_c = \tilde{\pi}_c$
(v) $E = P - \tilde{P}$
(vi) $E_c \triangleq P_c - \tilde{P}_c$.

For two Markov Chains with transition probabilities P and $\tilde{P}$ and sharing a common state space, the difference between the two stationary distributions can be bounded by $|\pi - \tilde{\pi}|_q \leq k\|E\|_\infty$ where $q = 1$ or ∞ and $k > 0$ are some constants and $|\pi - \tilde{\pi}|_q$ refers to the qth norm of the vector $\pi - \tilde{\pi}$ and $\|E\|_\infty$ denotes the ∞ norm of the error matrix E which is equivalent to the maximum absolute row sum of E. The constants K are usually referred to as *condition numbers* and several of them have been studied in the literature. Obviously, some of the condition numbers will yield tighter bounds than the others and [16] gives a nice comparison of the available bounds. Initial studies of steady-state distributions of PBNs using condition numbers were carried out in [10] but steady-state distributions under control were not considered in that reference. Here, we will prove a theorem for a particular condition number studied by Seneta [17]. For a given transition probability matrix P, this condition number called the *ergodicity coefficient* $\tau_1(P)$ is defined by

$$\tau_1(P) = \sup_{\substack{|x^T|_1 = 1 \\ x^T \mathbf{1}_n = 0}} |x^T P|_1 \tag{2}$$

where $\mathbf{1}_n$ denotes the n-dimensional column vector having all entries equal to one. Equivalent definitions are

$$\tau_1(P) = \frac{1}{2} \max_{i,j} \sum_{s=1}^{n} |p_{is} - p_{js}| \tag{3}$$

and

$$\tau_1(P) = 1 - \min_{i,j} \sum_{s=1}^{n} \min(p_{is}, p_{js}) \tag{4}$$

where p_{ij} refers to the ith row and jth column entry of matrix P. These two definitions are more useful for the purpose of computational evaluation. In [17], the ergodicity coefficient was used to obtain a bound on the perturbation in the steady-state distribution due to perturbations in the transition probability matrix. More specifically, it was shown that if $\tau_1(P) \neq 1$, then

$$|\pi - \tilde{\pi}|_1 \leq \frac{1}{1 - \tau_1(P)} \|E\|_\infty. \tag{5}$$

Here, we will use the above result to obtain an analytical bound on the perturbations in the *controlled* steady-state distributions that could result from perturbations in the uncontrolled probability transition matrix.

A. Analytical Result Involving the Ergodicity Coefficient

Theorem III.2: Let P and $\tilde{P}$ be two compatible probability transition matrices with $\tau_1(P) \neq 1$. Then

$$|\pi_c - \tilde{\pi}_c|_1 \leq \frac{1}{1 - \tau_1(P)} \|E\|_\infty. \tag{6}$$

Proof: The proof is accomplished by showing

(i) if $\tau_1(P) \neq 1$ then $1/(1 - \tau_1(P_c)) \leq 1/(1 - \tau_1(P))$, i.e., $\tau_1(P_c) \leq \tau_1(P)$; and
(ii) $\|E_c\|_\infty \leq \|E\|_\infty$.

From our earlier discussion, for the class of interventions that have been used for PBNs, we can write $P_c = TP$ where T is a stochastic matrix with each row containing only a single nonzero entry of 1. According to [18]

$$\tau_1(P_1 P_2) \leq \tau_1(P_1)\tau_1(P_2). \tag{7}$$

Thus, in our case $\tau_1(P_c) \leq \tau_1(T)\tau_1(P)$. From (4), it is clear that ergodicity coefficient of a stochastic matrix is less than or equal to 1 and, hence

$$\tau_1(P_c) \leq \tau_1(P). \tag{8}$$

Thus, from (5), it follows that:

$$|\pi_c - \tilde{\pi}_c|_1 \leq \frac{1}{1 - \tau_1(P)} \|E_c\|_\infty. \tag{9}$$

To prove the second part, we consider

$$E_c = P_c - \tilde{P}_c \tag{10}$$
$$= TP - T\tilde{P} \tag{11}$$
$$= T(P - \tilde{P}) \tag{12}$$
$$= TE. \tag{13}$$

In view of (13), it follows that the rows of E_c are selected from the rows of E and, hence, $\|E_c\|_\infty$ (maximum absolute row sum of E_c) $\leq \|E\|_\infty$. Thus, from (9), it follows that (6) holds, and this completes the proof.

There are other available perturbation bounds in the literature and some of them are tighter than the ergodicity coefficient bound. The reason for emphasizing the ergodicity coefficient bound here is that the kind of analytical result proved in the above theorem can be derived only for this bound. We will show with the help of simulations that the most effective perturbation bound (to be defined shortly) for the steady-state distribution of the controlled probability transition matrix can sometimes be greater than the corresponding perturbation bound for the steady-state distribution of the original uncontrolled probability transition matrix. The inequality in (8) implies that if the Markov Chain corresponding to an uncontrolled genetic network has a small ergodicity coefficient bound, then the corresponding controlled Markov Chain will also have an ergodicity coefficient that is bounded by the same bound. Consequently, if a stationary policy is designed from an estimated Markov Chain that is "close" to the actual one for the network, then this policy when applied to the actual network will produce results that are

close to the desired outcome, as far as the steady-state behaviour is concerned.

As already mentioned, there are several condition numbers other than the ergodicity coefficient that can be found in the literature. These perturbation bounds are mostly stated in terms of the *fundamental matrix* or the *group inverse* of $A := I - P$. The fundamental matrix of the Markov Chain with transition probability matrix P is defined by

$$Z = (A + e\pi^T)^{-1} \tag{14}$$

where $e = [1\ 1\ 1\ \cdots\ 1]^T$. The group inverse of A is the unique square matrix $A^{\#}$ satisfying the relationships

$$AA^{\#}A = A, A^{\#}AA^{\#} = A^{\#}, \text{ and } AA^{\#} = A^{\#}A. \tag{15}$$

The currently available condition numbers for bounding the 1 and ∞ norms of the perturbations in the steady-state distributions are [16]

$$k_1 = \|Z\|_\infty \quad q = 1 \tag{16}$$

$$k_2 = \|A^{\#}\|_\infty \quad q = 1 \tag{17}$$

$$k_3 = \frac{max_j \left(a^{\#}_{jj} - \min_i a^{\#}_{ij}\right)}{2} \quad q = \infty \tag{18}$$

$$k_4 = \max_{i,j} \left|a^{\#}_{ij}\right| \quad q = \infty \tag{19}$$

$$k_5 = \frac{1}{1 - \tau_1(P)} \quad q = 1 \tag{20}$$

$$k_6 = \tau_1(A^{\#}) = \tau_1(Z) \quad q = 1 \tag{21}$$

$$k_7 = \frac{\min_j \left\|A^{-1}_{(j)}\right\|_\infty}{2} \quad q = \infty. \tag{22}$$

Here, the bound k_5 involves the ergodicity coefficient.

B. Simulation Studies for Different Perturbation Bounds

Some of the large transition matrices encountered in genomics tend to be sparse and the perturbation bound k_5 based on the ergodicity coefficient is not very sharp for them. Accordingly, we will first report some simulation results for perturbation bounds using smaller networks of four genes (i.e., networks having $2^4 = 16$ states). For generating the networks for these simulations, we have used the data from melanoma cell lines which were previously used by us in several papers, e.g., [11]–[14]. The seven gene networks considered in those references were reduced to four gene networks using the reduction mapping algorithm given in [19].

For the simulations, we generated a number of PBNs consisting of four genes and calculated their perturbation bounds. The PBNs were then operated upon by a random stationary policy matrix T and the new perturbation bounds were calculated. In all the cases, the perturbation bound k_5 was found to be smaller for the controlled transition matrices as compared to that for the original uncontrolled transition matrix.

In Fig. 2, we show the ergodicity coefficient perturbation bounds (k_5) of the original uncontrolled PBNs as gray bars for 10 different generated PBNs. The blue stars represent the ergodicity coefficient perturbation bounds for TP where T is a randomly generated stationary policy matrix.

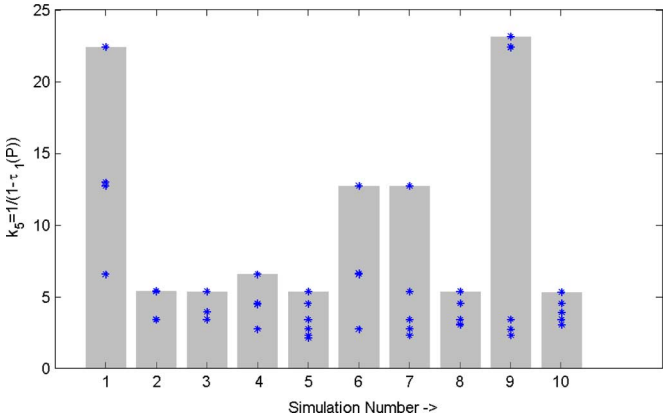

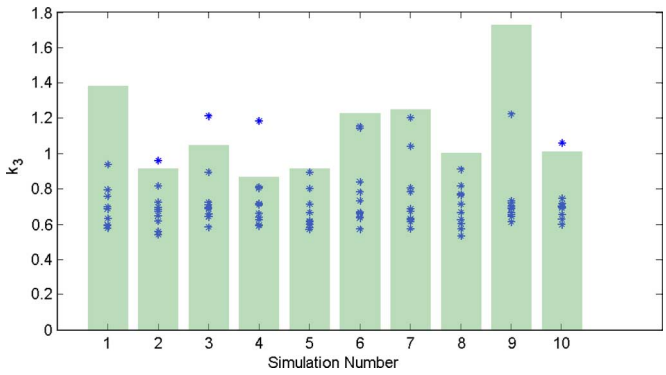

As shown in [16], the perturbation bound k_3 is one of the tighest bounds.

Fig. 3 shows the perturbation bound k_3 for different simulations. From Fig. 3, it is clear that in this case, the perturbation bound for some randomly generated stationary policy matrices (TP) can exceed the corresponding bound for the original uncontrolled probability transition matrix (P), although this situation is not very common.

The simulation studies for validating a number of control approaches that we proposed earlier were performed on a network of 7 genes containing WNT5A. To maintain uniformity, we built PBNs for the same seven genes and Fig. 4 shows the perturbation bound k_3 for a particular set of 80 simulations. The bars represent the perturbation bounds for the uncontrolled transition matrix P and the stars represent the perturbation bounds for the controlled transition matrix TP. The stationary policy corresponding to T represents the same objective as in [14], i.e., downregulating the gene WNT5A and using a discounted cost infinite horizon approach. We should note that the perturbation bounds k_3 for the uncontrolled PBN with transition probability matrix P and the controlled PBN with transition probability matrix TP are quite similar. We performed a number of other simulations and all of them led to the same conclusions.

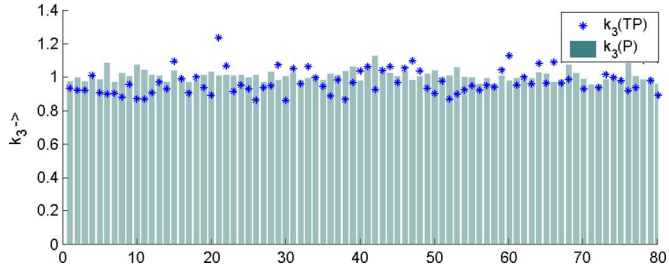

Fig. 4. Perturbation bound k_3 for 80 different PBNs (represented in green) and the stars represent perturbation bounds for stationary policies applied to the PBNs. Here the number of genes is 7.

The reason behind highlighting these perturbation bounds is that they give us a measure of the maximum change in the steady-state distributions. If for instance, the estimated transition matrix (P) of a gene regulatory network has a small perturbation bound k_5, then we can rest assured that the steady-state $(\tilde{\pi}_c)$ of the actual gene regulatory network $(\tilde{P})$ controlled by a stationary policy T will be close to the steady-state (π_c) of the gene regulatory network (P) controlled by the same stationary policy T . Thus for intervention strategies where the steady-state distribution is a measure of the effectiveness of the intervention, the perturbation bounds can provide a good estimate of the outcome of the control strategy. In the case of a PBN whose perturbation bound k_5 is high, the perturbation bound k_3 can be used to give us some idea of the uncertainty involved. As shown by simulations, the difference between the values of k_3 for the original uncontrolled probability transition matrix P and the controlled probability transition matrix TP is quite small, and for this reason the uncertainty in the steady-state distribution after application of the stationary policy will be approximately the same as the uncertainty in the steady-state distribution of the original uncontrolled PBN.

Furthermore, the perturbation bounds can be used as a kind of measure for network selection. In general, genetic networks are quite stable or, in other words, robust to small perturbations. Hence a transition probability matrix representing a genetic network should necessarily be robust to perturbations and the alteration in its steady state for small changes in the transition probabilities should be minimal. In this context, it is appropriate to mention that in the field of genomics, it is still not clear as to what metric[3] should be used to carry out network selection. The available data in genomics studies are quite limited and this can give rise to a number of possible networks that fit the data. The selection among these different networks is a very important issue and some initial approaches for doing this have been proposed in the literature, e.g. [8], [19]. The perturbation bound combined with other metrics could provide yet an alternative approach for doing so.

IV. Uncertainty Characterization for PBNs

Although the results in the last section were developed with PBNs in mind, really they were focussed on how the uncertainties in the probability transition matrix of a Markov Chain propagate to its stationary distribution. Since PBNs can be modeled using Markov Chains, such results are applicable to PBNs. In this section, we focus on relating the results of the last section more closely to PBNs by deriving explicit expressions to show how different classes of possible uncertainties for a PBN would translate into uncertainties in the transition probability matrix for the corresponding Markov Chain. In Section V, we will design worst-case optimal policies for the specific uncertainties derived in this section.

We consider a context-sensitive PBN consisting of n genes and composed of k BNs with selection probabilities $c_1, c_2, \cdots, c_k$. Furthermore, let q be the network switching probability and let p be the perturbation probability. According to the definition of Context-Sensitive PBN [13], the following mutually exclusive sequence of events can occur at any time point t: 1) the current network function is applied, the PBN transitions accordingly, and the network function remains the same for the next transition; 2) the current network function is applied, the PBN transitions accordingly, and a new network function is selected for the next transition; 3) there is a random perturbation and the network function remains the same for the next transition; 4) there is a random perturbation and a new network function is selected for the next transition.

Clearly, each BN will have 2^n states $00\cdots0$ to $11\cdots1$ and the collection of k BNs can be considered to have a set of $2^n k$ states. Let $w(t) \in \{0, 1, 2, \cdots, 2^n k - 1\}$ be the state that is occupied by the network at time t. Let $a, b \in \{0, 1, 2, \cdots, 2^n k - 1\}$ be any two states and define $r_1 = [a/2^n]$ and $r_2 = [b/2^n]$. Then the transition probability from a to b is given by [13]

$$\begin{aligned} &Pr\left(w(t+1)=b|w(t)=a\right) \\ &\quad= \left[(1-q)(1-p)^n f_{r,a,b}+(1-q)(1-p)^{n-h}p^h s(h)\right] g(a,b) \\ &\qquad+\left[q\frac{c_{r_2}}{\sum_{l=1,l\neq r_1}^{k} c_l}(1-p)^n f_{r_1,a,b}\right. \\ &\qquad\quad\left.+q\frac{c_{r_2}}{\sum_{l=1,l\neq r_1}^{k} c_l}(1-p)^{n-h}p^h s(h)\right] \\ &\qquad\times(1-g(a,b)) \end{aligned} \tag{23}$$

where h is the hamming distance between $mod(a, 2^n)$ and $mod(b, 2^n)$ represented in binary digits, i.e., the number of genes which differ between the two states

$$f_{r,a,b} = \begin{cases} 1 & \text{if } mod(a,2^n) \text{ transitions to } mod(b,2^n) \\ & \quad \text{in a single step in network } r \\ 0 & \text{otherwise} \end{cases}$$

$$g(a,b) = \begin{cases} 1 & \text{if } r_1 = r_2 = r \text{ (say)} \\ 0 & \text{otherwise} \end{cases}$$

and

$$s(h) = \begin{cases} 0 & \text{if } h = 0 \\ 1 & \text{otherwise.} \end{cases}$$

In (23), the first two terms correspond to events 1 and 3, respectively, while the next two terms correspond to events 2 and 4, respectively.

In practice, it will be almost impossible to detect the BN from which the current gene activity profile is being emitted. In most cases, we will have knowledge only of the expression status of

[3]Here metric refers to a type of distance function to differentiate networks fitting the data. For instance, if two networks fit the data, the one with the lower perturbation bound can be selected.

the genes. To handle such situations, we take the expectation of the transition probabilities over the networks. To this end, let $z(t)$ be the decimal equivalent of the binary gene activity profile at time t. Then for any $s_1, s_2 \in \{0, 1, 2, \cdots, 2^n - 1\}$, the averaged out transition probability is given by [13]

$$Pr\left[z(t+1)=s_1|(t)=s_2\right]=\sum_{r_1=1}^{k}\sum_{r_2=1}^{k} c_{r_1}.Pr\left[w(t+1)=s_1 +2^n(r_2-1)|w(t)=s_2+2^n(r_1-1)\right]. \quad (24)$$

Substituting (23) into (24) yields

$$\begin{aligned}
&\Pr\left[z(t+1)=s_1|z(t)=s_2\right]\\
&=\sum_{r_1=1}^{k} c_{r_1}[(1-q)(1-p)^n f_{r_1,s_2,s_1}\\
&\qquad+(1-q)(1-p)^{n-h}p^h s(h)\\
&\qquad+q\sum_{r_2=1,r_2\neq r_1}^{k}\frac{c_{r_2}}{\sum_{l=1,l\neq r_1}^{k} c_l}(1-p)^n f_{r_1,s_2,s_1}\\
&\qquad+q\sum_{r_2=1,r_2\neq r_1}^{k}\frac{c_{r_2}}{\sum_{l=1,l\neq r_1}^{k} c_l}(1-p)^{n-h}p^h s(h)\Bigg]\\
&=(1-p)^n\left\{\sum_{r_1=1}^{k} c_{r_1}((1-q)f_{r_1,s_2,s_1}\right.\\
&\qquad+q\frac{\sum_{r_2=1,r_2\neq r_1}^{k} c_{r_2}}{\sum_{l=1,l\neq r_1}^{k} c_l}f_{r_1,s_2,s_1}\\
&\qquad+(1-q)\left(\frac{p}{1-p}\right)^h s(h)\\
&\qquad\left.+q\frac{\sum_{r_2=1,r_2\neq r_1}^{k} c_{r_2}}{\sum_{l=1,l\neq r_1}^{k} c_l}\left(\frac{p}{1-p}\right)^h s(h)\right)\Bigg\}\\
&=(1-p)^n\left\{\sum_{r_1=1}^{k} c_{r_1}\left(f_{r_1,s_2,s_1}+\left(\frac{p}{1-p}\right)^h s(h)\right)\right\}\\
&=(1-p)^n\left\{\sum_{r_1=1}^{k} c_{r_1}f_{r_1,s_2,s_1}\right\}\\
&\quad+(1-p)^n\left(\frac{p}{1-p}\right)^h s(h) \qquad (25)
\end{aligned}$$

Let us denote by $B_1, B_2, \cdots, B_k$, the transition matrices of the individual BNs. Then B_v, $v \in [1, \ldots, k]$ represent deterministic transition matrices and hence each B_v has a single nonzero entry of 1 in each row. The second term of (25) is independent of the constituent BNs or their selection probabilities and depends only on the perturbation probability p, number of genes n and the Hamming distance h between states (which can be determined when n is known). Consequently, the probability transition matrix P of the *context-sensitive* PBN composed of k BNs is of the form

$$P=(1-p)^n\sum_{v=1}^{k} c_v B_v + D^{n,p} \qquad (26)$$

where the entries of B_v are generated from the first term in (25) and $D^{n,p}$ represents the entries corresponding to the second term of (25).

Furthermore, the matrix $D^{n,p}$ has the form shown in the equation at the bottom of the page.

The individual $2^n \times 2^n$ terms of $D^{n,p}$ for $i = 0, 1, \ldots, 2^n - 1$ and $j = 0, 1, \ldots, 2^n - 1$ are

$$D^{n,p}(i,j)=D^{n,p}(j,i)=(1-p)^n\begin{cases}0 & \text{if } i=j\\ \left(\frac{p}{1-p}\right)^{h(i,j)} & \text{otherwise}\end{cases}$$

where $h(i, j) =$ no. of bits different in the binary representation of i and j.

From (26), it is clear that the transition probability matrix of an *instantaneously random* PBN composed of BNs is the same as the transition probability matrix of an averaged *context-sensitive* PBN. Thus, in the rest of this paper, the term PBN will denote both *context-sensitive* and *instantaneously random* PBNs. For a PBN without perturbations, the transition probability matrix will have the form

$$P=\sum_{v=1}^{k} c_v B_v. \qquad (27)$$

Two possible forms of uncertainty that might be encountered during the inference of genetic networks are as follows.

Case 1) The number of BNs k along with their network transition matrices $B_1, B_2, \cdots, B_k$ are correctly inferred but the network selection probabilities

$$D^{n,p}=(1-p)^n\begin{bmatrix}
0 & \frac{p}{1-p} & \frac{p}{1-p} & \left(\frac{p}{1-p}\right)^2 & \cdots & \left(\frac{p}{1-p}\right)^n\\
\frac{p}{1-p} & 0 & \left(\frac{p}{1-p}\right)^2 & \frac{p}{1-p} & \cdots & \left(\frac{p}{1-p}\right)^{n-1}\\
\frac{p}{1-p} & \left(\frac{p}{1-p}\right)^2 & 0 & \cdot\cdot & \cdots & \left(\frac{p}{1-p}\right)^{n-1}\\
\cdot\cdot & \cdot\cdot & \cdot\cdot & \cdot\cdot & \cdots & \cdot\cdot\\
\cdot\cdot & \cdot\cdot & \cdot\cdot & \cdot\cdot & \cdots & \cdot\cdot\\
\left(\frac{p}{1-p}\right)^n & \left(\frac{p}{1-p}\right)^{n-1} & \left(\frac{p}{1-p}\right)^{n-1} & \left(\frac{p}{1-p}\right)^{n-2} & \cdots & 0
\end{bmatrix}.$$

$c_1, c_2, \cdots, c_k$ are uncertain. We might have an interval estimate of the selection probabilities, i.e., $\underline{c_v} \leq c_v \leq \overline{c_v}$.

Case 2) The network selection probabilities $c_1, c_2, \cdots, c_k$ are correctly determined but there is uncertainty in some of the functions of the BNs $B_1, B_2, \cdots, B_k$.

We separately consider each of these two cases:

Case 1 Uncertainty: For a transition from state i to state j, let $l \in \{0, 1, 2, \cdots, k\}$ denote the number of BNs allowing such a transition. Then, from (25)

$$P(i,j) = (1-p)^n \sum_{x=1}^{l} c_{v_x} + D^{n,p}(i,j). \tag{28}$$

The intervals for $P(i,j)$ corresponding to this case satisfy

$$\begin{aligned}(1-p)^n \min \left(\sum_{x=1}^{l} c_{v_x} \right) + D^{n,p}(i,j) \\ \leq P(i,j) \\ \leq (1-p)^n \max \left(\sum_{x=1}^{l} c_{v_x} \right) + D^{n,p}(i,j).\end{aligned} \tag{29}$$

Here the min (max) refers to the minimum (maximum) that the summation $\sum_{x=1}^{l} c_{v_x}$ can achieve. $\sum_{x=1}^{l} \underline{c_{v_x}}$ is a lower bound for the min while $\sum_{x=1}^{l} \overline{c_{v_x}}$ is an upper bound for the max but for specific cases, tighter lower (upper) bounds can be achieved as shown in the next paragraph.

We consider a few representative cases. For $l = 0$, there is no uncertainity and $P(i,j) = D^{n,p}(i,j)$. For $l = k$, the transition from state i to j is present in all the networks and hence, $\sum_{x=1}^{l} c_{v_x} = \sum_{x=1}^{k} c_{v_x} = 1$ which results in $P(i,j) = (1-p)^n + D^{n,p}(i,j)$ for $l = k$. When a transition is present in only one BN, i.e., $l = 1$, then the interval is $P(i,j) \in [(1-p)^n \underline{c_{v_1}} + D^{n,p}(i,j), (1-p)^n \overline{c_{v_1}} + D^{n,p}(i,j)]$. For $l = k-1$, let B_{v_k} be the BN which does not allow the transition. Then $(1-p)^n(1-\overline{c_{v_k}}) + D^{n,p}(i,j) \leq P(i,j) \leq (1-p)^n(1-\underline{c_{v_k}}) + D^{n,p}(i,j)$.

For other values of l, a conservative estimate of the interval is

$$\begin{aligned}(1-p)^n \sum_{x=1}^{l} \underline{c_{v_x}} + D^{n,p}(i,j) \\ \leq P(i,j) \\ \leq min \left(1, \left((1-p)^n \sum_{x=1}^{l} \overline{c_{v_x}} + D^{n,p}(i,j) \right) \right).\end{aligned} \tag{30}$$

This form of interval uncertainty in network selection probabilities can be mapped to interval uncertainties in the transition probabilities. We refer to this uncertainty as being characterized by an *Interval Matrix Constraint*.

The perturbation bounds in the last section were all derived in terms of the induced ∞-norm of the difference between the estimated and actual probability transition matrices. In order to make these results applicable to the interval matrix uncertainty case, we can map the interval matrix uncertainty into the ∞-norm of the uncertainty in the probability transition matrix.

Let β denote the upper bound on the induced infinity norm of the uncertainty in the probability transition matrix. If $k < 2^n$, i.e., the number of networks is less than the number of states, then the maximum possible value for β is $(1-p)^n \sum_{v=1}^{k} (\overline{c_v} - \underline{c_v})$. This is because from any state i, a transition will occur to one of the states $[0, 2^n - 1]$ and the upper bound on the uncertainty is largest when the states that i transitions to in each of the k different BNs are distinct. This makes intuitive sense as dissimilar transitions from a particular state i in the different networks imply that considerable uncertainty would be involved in the inference of the transitions from that state.

Case 2 Uncertainty: The ambiguity in the functions forming a BN will give rise to a number of possible BNs and consequently different PBNs. For a PBN composed of k BNs, let us suppose that the function for the first gene in BN1 is uncertain and n_d different functions are possible. This will give rise to n_d different BNs and consequently n_d different PBNs. Note that the uncertainty here is characterized by the existence of a finite number of transition matrices and we will refer to this uncertainty as being characterized by a *Finite Number of Transition Matrices Constraint*.

V. Optimal Intervention for the Worst-Case Scenario

In this section, our goal is to develop infinite horizon PBN intervention policies that are optimal for the worst-case uncertainty within a given uncertainty class. To do so, we will first recall some facts and notation from [14] where we considered optimal infinite horizon intervention for the nominal case, i.e., in the absence of uncertainty. As shown in [13] and [14], a PBN with control can be modeled as a stationary discrete-time dynamic system

$$z_{t+1} = f(z_t, u_t, w_t), \quad t = 0, 1, \ldots. \tag{31}$$

where for all t, the state z_t is an element of a space S , the control input u_t is an element of a space C, the disturbance w_t, which captures the randomness due to different sources, is an element of a space D and $f : S \times C \times D \mapsto S$. In the particular case of PBNs of n genes composed of k BNs with perturbation probability p and network transition probability q, $S = [0, 1, 2, \ldots\ldots, 2^n - 1]$ and the control input u_t is constrained to take values in the space $C = [0, 1, \ldots\ldots, 2^m - 1]$, where m is the number of control inputs. We next explain the basis for defining C as above. Suppose the PBN has m control inputs $v_1, v_2, \cdots, v_m$, each of which can take on only the binary values 0 or 1. Then at any given time step t, the row vector $v(t) \triangleq [v_1(t), v_2(t), \cdots, v_m(t)]$ describes the complete status of all the control inputs. $v(t)$ can take on all binary-vector values

from $[0,0,\cdots,0]$ to $[1,1,\cdots,1]$. One can equivalently represent the control input status using the decimal number

$$u_t = \sum_{i=1}^{m} 2^{m-i} v_i(t). \tag{32}$$

As $v(t)$ takes on binary values from $[0,0\cdots,0]$ to $[1,1,\cdots,1]$, the variable u_t ranges from 0 to $2^m - 1$. We can equivalently use u_t as an indicator of the complete control input status of the probabilistic BN at time step t. Therefore, $C = [0,1,2,3\ldots\ldots 2^m - 1]$.

The disturbance w_t is manifested in terms of change of network based on the network switching probability q or change of state due to perturbation probability p. The random disturbances w_t, $t = 0,1,\ldots$ have identical statistics and are characterized by probabilities $P(.|x_t,u_t)$ defined on D, where $P(.|x_t,u_t)$ is the probability of occurrence of w_t, when the current state and control are x_t and u_t, respectively. w_t is independent of prior disturbances $w_0, w_1 \ldots . w_{t-1}$.

Another equivalent way to represent the dynamical system (31) is as a finite state *Markov Chain* described by the control-dependent one-step transition probability $p_{ij}(u)$ where for any $t = 0,1,2,\cdots; i,j \in S$ and $u \in C$,

$$p_{ij}(u) := P(z_{t+1} = j | z_t = i, u_t = u). \tag{33}$$

In this paper, we will interchangeably use either representation (31) or (33) depending on their suitability for a particular context or a particular derivation.

We will consider a discounted cost infinite horizon approach as in [14]. Similar to Section III, let us define P to be the estimated transition probability matrix and $\tilde{P}$ to be the actual transition matrix of the PBN. So far, we have considered three different classes of uncertainty: 1) an upper bound on the induced ∞-norm of the difference of the probability transition matrices; 2) transition probabilities in an interval; and 3) transition probabilities selected from a finite set of transition matrices. We will first carry out the development for a general description of the uncertainty. Specific solution details for the three uncertainty classes will be provided later.

When the transition probabilities are exactly known, the states make transitions according to $\varpi := (P^u)_{u\in C}$. In this case, once a control input is chosen, the resulting controlled transition probability matrix is uniquely determined. Let $\mu = (u_1, u_2, \ldots.)$ represent a generic control policy and Π represent the set of all possible μs, i.e., the set of all possible control policies.

Let $J_{\mu,\varpi}$ denote the expected total cost for the discounted cost infinite horizon problem under control policy μ and transitions ω, i.e.

$$J_{\mu,\varpi}(z_0) = \lim_{M\to\infty} E_{w_t, t=0,1,\ldots} \left\{ \sum_{t=0}^{M-1} \alpha^t \tilde{g}(z_t, u_t, w_t) \right\} \tag{34}$$

where $0 < \alpha < 1$ denotes the discount factor and $\tilde{g}(z_t,u_t,w_t)$ represents the cost of going from state z_t to z_{t+1}[4] under the control action u_t. As will be shown in the examples presented in Section VI, $\tilde{g}$ is higher for undesirable destination states. For the same destination states, $\tilde{g}$ is higher when the control is active versus when it is not.

The *nominal* problem here corresponds to minimizing the cost in (34), and consequently, the optimal infinte-horizon discounted cost for the *nominal* problem is given by:

$$\Phi(\Pi, \varpi, z_0) := \min_{\mu\in\Pi} J_{\mu,\varpi}(z_0). \tag{35}$$

Due to uncertainties in the transition probability matrix, there can be various possible transitions corresponding to any given choice of control u and, therefore, various ϖ's. Let us assume that for each control action u and time t, the corresponding transition matrix P_t^u lies in some uncertainty set $\mathcal{P}^u$. This set can be characterized in different ways. For instance, there could be *interval* uncertainty associated with each entry of the controlled transition matrix, or a bound could be specified for the maximum deviation of each entry of the controlled transition matrix from the corresponding entry of some reference transition matrix. Yet another uncertainty characterization could be one where the controlled probability transition matrix has to be chosen from a finite set of probability transition matrices.

With respect to time, two models of uncertainty are possible: (i) *Stationary (time-invariant) uncertainty* model where the control-dependent transition matrices are chosen by *nature*[5] once and for all, and remain fixed thereafter; (ii) *Time-varying uncertainty* where the control-dependent transition matrices can vary arbitrarily with time, within some prescribed bounds.

Mathematically, a *policy of nature* signifies a specific collection of time-dependent controlled transition matrices $\omega := (P_t^u)_{u\in C, t\in\mathcal{N}}$ chosen by nature, and the set of admissible policies of nature is $\Omega := (\mathbf{X}_{u\in C}\mathcal{P}^u)^{|\mathcal{N}|}$ where $\mathbf{X}$ denotes the Cartesian product. We denote by Ω_s the set of stationary controlled transition matrices:

$$\Omega_s = \Big\{ \omega := (P_t^u)_{u\in C, t\in\mathcal{N}} \in \Omega : (P_t^u) = (P_s^u) \text{ for every } t,s \in \mathcal{N}, u \in C \Big\}. \tag{36}$$

For designing an optimal control policy based on the worst case scenario, we would like to maximize the cost with respect to the uncertainties and minimize the cost using control actions. The resulting optimal cost for the stationary uncertainty model can be represented as

$$\Phi(\Pi, \Omega_s, z_0) := \min_{\mu\in\Pi} \max_{\omega\in\Omega_s} J_{\mu,\omega}(z_0). \tag{37}$$

In contrast, the time-varying uncertainty model leads to the following optimal cost:

$$\Phi(\Pi, \Omega, z_0) := \min_{\mu\in\Pi} \max_{\omega\in\Omega} J_{\mu,\omega}(z_0). \tag{38}$$

Since $\Omega_s \subset \Omega$, it follows that:

$$\Phi(\Pi, \Omega_s, z_0) \le \Phi(\Pi, \Omega, z_0). \tag{39}$$

[4]Note that a Markov Chain can be modeled by $z_{t+1} = w_t$ [20]. Hence the destination state is the same as the disturbance.

[5]Here nature refers to the concept of uncertainty or in other words, modeling over which we have no control.

The stationary uncertainty model is appropriate when we know that the original genetic network is a time-invariant one but our inference procedure can only provide estimates of its transition probabilities within some intervals. The time-varying uncertainty model is appropriate for the instances when we know that the transition probabilities can change with time, although within certain prescribed bounds. For infinite horizon problems, it can be shown [21] that both the stationary and time-varying uncertainty models produce the same optimal cost. Accordingly, in this paper, we will stick to the stationary uncertainty model.

The solution to the problem of minimizing the infinite horizon nominal discounted cost in (34) has been presented in [14]. The optimal cost (or *value function*) is the unique solution of a *Bellman Equation* and can be iteratively determined using a recursive procedure called *Value Iteration*. The optimal control is obtained as a byproduct of the steps involved in the value iteration process. These results form the basis of the optimal control solution even in the presence of uncertainty and, therefore, we will briefly recall them here. To this end, let us introduce a couple of definitions. For a given stationary controlled transition probability matrix P^u and any state $z \in S$, define $(P_z^u)^T$ to be the row of P^u that defines the transition probabilities from state z. Furthermore, for any initial state z_0, define the value function $v : S \to R$ by $v(z_0) = \min_{\mu \in \Pi} J_{\mu,\omega}(z_0)$, and the vector $\mathbf{v}$ by $\mathbf{v} = [v(0), v(1), \cdots, v(2^n - 1)]^T$. Then, as shown in [14], the value function v satisfies the Bellman Equation[6]

$$v(z) = \min_{u \in C} \left[g(z,u) + \alpha \left(P_z^u\right)^T \mathbf{v} \right] \quad \forall z \in S. \tag{40}$$

Furthermore, the value function v is the unique limit value of the convergent vector sequence

$$v_{k+1}(z) = \min_{u \in C} \left[g(z,u) + \alpha \left(P_z^u\right)^T \mathbf{v_k} \right], \quad z \in S \text{ and } k = 1, 2, \cdots, . \tag{41}$$

We are now ready to state the solution to the infinite horizon discounted cost problem in the presence of stationary uncertainty. Intuitively, the robust versions of the Bellman Equation (40) and the value iteration (41) can be derived by noting that, in this case, we are focussing on the worst-case scenario and, therefore, P_z^u on the right hand sides of (40) and (41) must be chosen to maximize $(P_z^u)^T \mathbf{v}$ and $(P_z^u)^T \mathbf{v_k}$, respectively. These ideas have been formalized in [21] and [22] and are summarized in the theorem stated here.

Theorem V.1: For the infinite horizon robust control problem (37) with stationary uncertainties on the transition matrices, stationary control policies and a discounted cost function with discount factor $\alpha \in [0, 1)$, perfect duality holds, i.e.

$$\begin{aligned} \Phi(\Pi, \Omega_s, z_0) &= \min_{\mu \in \Pi} \max_{\omega \in \Omega_s} J_{\mu,\omega}(z_0) \\ &= \max_{\omega \in \Omega_s} \min_{\mu \in \Pi} J_{\mu,\omega}(z_0) := \Psi(\Pi, \Omega_s, z_0). \end{aligned} \tag{42}$$

[6]As shown in [14], the function $\bar{g}(z_t, u_t, w_t)$ is replaced by a new function $g(z_t, u_t)$ by taking the expectation over the destination states.

The optimal value is given by $\Phi(\Pi, \Omega_s, z_0) = v(z_0)$, where z_0 is the initial state, and the value function v satisfies the optimality conditions

$$v(z) = \min_{u \in C} \left(g(z,u) + \alpha \sigma_{\mathcal{P}_z^u}(\mathbf{v}) \right), \quad z \in S \tag{43}$$

where

$$\sigma_{\mathcal{P}_z^u}(\mathbf{v}) := \sup_{P_z^u \in \mathcal{P}_z^u} \left[\left(P_z^u\right)^T \mathbf{v} \right]. \tag{44}$$

Furthermore, the value function is the unique limit value of the convergent vector sequence defined by

$$v_{k+1}(z) = \min_{u \in C} \left(g(z,u) + \alpha \sigma_{\mathcal{P}_z^u}(\mathbf{v_k}) \right), \quad z \in S, k = 1, 2, \ldots. \tag{45}$$

A stationary optimal control policy $\mu = (u^*, u^*, \ldots.)$ is obtained as the u that minimizes the right hand side of (45) once the iteration has converged.

The theorem stated above supplies us with a *value iteration* procedure [20] to construct control policies which are optimal for the worst case scenario. The theorem is quite general and does not focus on a specific characterization of the uncertainty. However, to apply the theorem to obtain worst-case optimal strategies, one will have to repeatedly carry out the maximization in (44). This maximization problem can present different levels of difficulty depending on the characterization of the uncertainty. Following the discussion in Section IV, we will address the solution to the maximization problem in (44) for 3 specific cases: (A) $||\ ||_\infty$ *norm constraint*: The uncertainty will be characterized by an upper bound on the $||\ ||_\infty$ norm of the difference between the actual and estimated transition matrices; (B) *Interval Matrix* constraint (**Case 1 Uncertainty**): The transition probabilities are assumed to lie between two reference matrices or in other words, the transition probabilities are in an interval; and (C) *Finite number of transition matrices* constraint (**Case 2 Uncertainty**): The transition probabilities belong to a finite set of matrices.

The solution to the maximization problem in (44) for the 3 cases are provided next. The proofs for Cases A and B are provided in the website http://www.gsp.tamu.edu/publications/robust/supplement.htm.

Theorem V.2. Case A: The maximization problem for the infinity norm takes the form: evaluate B^* where

$$B^* = \max_p v^T p \tag{46}$$

subject to the constraints $p^T \mathbf{1} = 1$, $||p - q||_1 \leq \beta$. Here, the vector p^T represents the appropriate row of a probability transition matrix, and the vector q^T represents the corresponding reference row about which the perturbation must satisfy a 1-norm constraint. Let us arrange the values of the components of v in ascending order and denote them by $\tilde{v}(i) : 1 \leq i \leq 2^n$. The

same order is used to sort the values of q represented as $\tilde{q}$. Then the solution to (46) is

$$B^* = \min_{1\le j\le 2^n} \left(\sum_{i>j} \tilde{q}(i)\tilde{v}(i) + \tilde{v}(j)\left(\sum_{i\le j} \tilde{q}(i) - \frac{\beta}{2} \right) + \tilde{v}(2^n)\frac{\beta}{2} \right). \quad (47)$$

Theorem V.3. Case B: The maximization problem for *interval matrix constraint* takes the form: evaluate B^* where

$$B^* = \max_{p\ge 0} v^T p \quad (48)$$

subject to the constraints $p^T\mathbf{1} = 1$, $p1 \le p \le p2$. Here the vector p^T represents the row of a probability transition matrix, and the nonnegative vectors p_1^T and p_2^T characterize the interval uncertainty associated with that row. The corresponding solution is:

$$B^* = \min_{\zeta\in 0, v_1, \ldots v_{2^n}} v^T p2 + \zeta(1 - p2^T\mathbf{1}) + (p2 - p1)^T(\zeta\mathbf{1} - v)^+ \quad (49)$$

where x^+ stands for the positive part of vector x.

Case C: (*Finite number of transition matrices constraint*) If the uncertainty is modeled as a finite collection of transition matrices, then for every control action u, $\mathcal{P}^u = \{P^{u,1}, \cdots, P^{u,L}\}$ where $P^{u,x}$ represents a possible transition matrix. Let $(P_i^{u,x})^T$ denote the ith row of the transition matrix $P^{u,x}$. Then, the maximization problem corresponding to $z = i$ in (44) is:

$$B^* = \max_{p\in\{P_i^{u,1},\cdots,P_i^{u,L}\}} \mathbf{v}^T p = \max_{1\le x\le L} \mathbf{v}^T P_i^{u,x}. \quad (50)$$

The complexity of each step of the Bellman recursion for the nominal problem (without considering any uncertainties) is $O(2^m 2^n)$. For the robust versions, the one step complexities are $O(2^m 2^n 2^n)$ for Cases A and B and $O(2^m 2^n L)$ for Case C.

VI. Worst-Case Design Examples

A. Design for a Simple "Toy" Example

Let us consider 2 BNs each consisting of three genes. Their state transition diagrams are shown in Figs. 5 and 6. Both of them have the same attractors ('000' and '101') but their basins are different. The states represented by decimal 0 to 3 ('000,' '001,' '010,' and '011') are assumed to be desirable with a cost of 0 while states 4 to 7 ('100,' '101,' '110,' and '111') are marked as undesirable and each is associated with a cost of five units. We will consider an infinite-horizon cost with a discount factor of $\alpha = 0.9$. For clarity of presentation, in this example the cost of control is assumed to be 0. The perturbation probability p is set to be 0.01. The control action allowed is the flipping of gene 3. For example, when we are at state "010," we can flip gene 3 to move to "011" or suppose in case of being at state "111," application of control can reset the state to "110."

Case 1: We will assume for this case that the PBN is formed from BN1 and BN2 but their selection probabilities are uncertain. The range of c_1 is assumed to be [0.2 0.6] and the range of c_2 is [0.4 0.8]. Fig. 7 shows the robust control policy obtained using these intervals along with the optimal policies for two specific PBNs, PBN1, and PBN2. PBN1 has selection probabilities $c_1 = 0.6$ and $c_2 = 0.4$, i.e. $P_{PBN1} = (1-p)^3(0.6BN_1 + 0.4BN_2) + D^{3,0.01}$ while PBN2 has selection probabilities $c_1 = 0.2$ and $c_2 = 0.8$, i.e. $P_{PBN2} = (1-p)^3(0.2BN_1 + 0.8BN_2) + D^{3,0.01}$.

Table I shows the average cost per state for a number of conditions. Column 2 represents the uncontrolled cost, while columns 3 and 4 represent the costs of optimal policies designed based on PBN1 and PBN2, respectively. Column 5 denotes the cost of the robust control policy designed based on the aforementioned intervals. Rows 2 and 3 contain the entries for PBN1

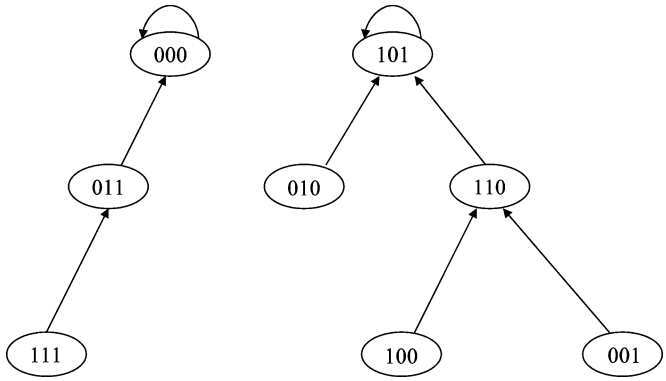

Fig. 5. BN 1.

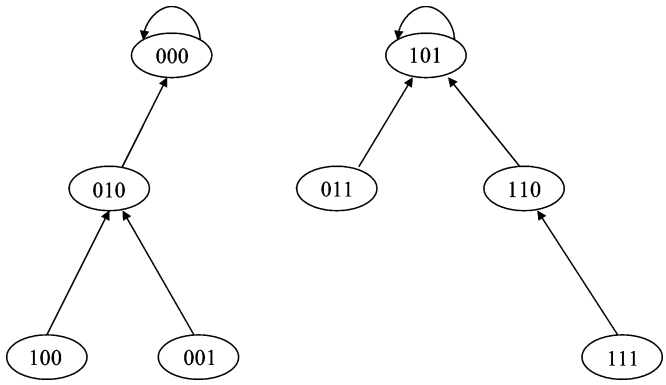

Fig. 6. BN 2.

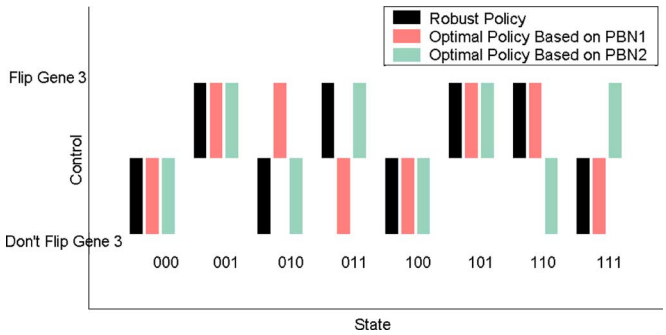

Fig. 7. Control policies.

TABLE I
AVERAGE COST TABLE

PBN no.	Uncontrolled Cost	PBN1 OP cost	PBN2 OP cost	R policy cost
PBN1	32.390343	8.774967	17.611064	11.846590
PBN2	27.321924	17.356684	6.512496	8.546661
worst case	36.372613	22.554549	17.613279	17.494679

TABLE II
AVERAGE COST TABLE

PBN no.	Uncontrolled Cost	OP1 cost	OP2 cost	R policy cost
PBN1	32.724398	7.293911	22.128735	8.893253
PBN2	29.167868	14.125852	8.132147	10.217805

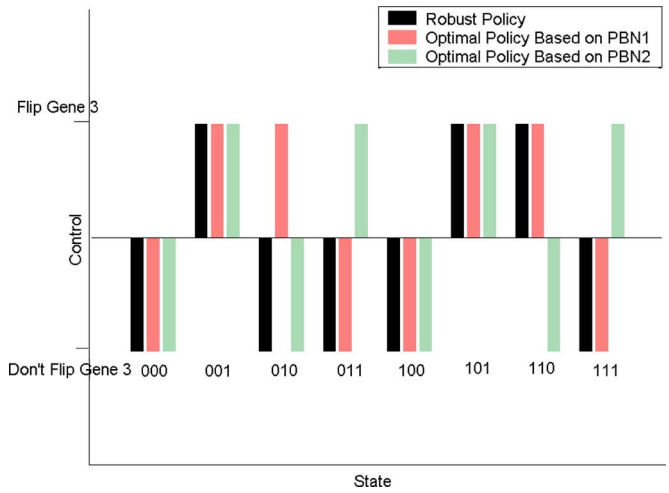

Fig. 8. Control policies.

and PBN2, respectively, while row 4 contains the worst case cost for each policy based on interval uncertainty.

We should note that the application of policy PBN2 OP[7] on PBN1 generates a reasonably high cost (17.611064) whereas the cost with the optimal robust policy is much lower (11.846590). The optimal robust policy is not the optimal policy for PBN1 as shown by the PBN1 OP cost (8.774967) but it minimizes the worst case cost for both PBN1 and PBN2. The entries in the fifth column are less than the maximum of the entries in columns 3 and 4, which agrees with the intuitive expectation that the cost with the optimal robust policy will be lower than the worst-case cost with any other policy. This example shows that the robust policy design will be of considerable benefit in cases of interval uncertainty when our objective is to minimize the associated worst-case cost.

Case 2: For clarity of presentation, let us assume that the PBN consists of a single BN which is either BN1 (Fig. 5) or BN2 (Fig. 6). Thus, PBN1 is formed of BN1 ($P_{PBN1} = (1-p)^3 BN1 + D^{3,0.01}$) and PBN2 is formed of BN2 ($P_{PBN1} = (1-p)^3 BN2 + D^{3,0.01}$). The optimal policy for each individual PBN is shown in Fig. 8 along with the robust policy based on the *finite number of transition matrices* model. Let us consider a particular state "110." Observing the attractor landscape of BN1 suggests that it is better to flip gene 3 at this state because flipping gene 3 will take us to "111" which is in the basin of a desirable attractor "000." But for BN2, intuitively it appears that flipping gene 3 is not beneficial as the state "111," that results from flipping, is in the same basin as before. The robust policy shown in Fig. 8 suggests flipping of gene 3 while at state "110," and this is based on the worst case scenario among the two PBNs. Intuitively, this robust policy makes sense since the cost of flipping has been assumed to be zero, and flipping gene 3 in BN2, although certainly not beneficial, is not detrimental either.

In terms of numerical values, Table II shows the average cost per state while applying different stationary policies. The second column under heading Uncontrolled cost denotes the average cost per state when no control is applied. The third column refers to the average cost per state when the optimal policy based on PBN1 is applied to both the PBNs. Similarly the fourth column denotes the average cost per state when the optimal policy based on PBN2 is applied to both the PBNs. R policy Cost refers to the average cost per state when the robust policy is applied to the individual PBNs. The table shows that the expected per state cost for application of the robust policy on PBN1 (8.893253) is much lower than the expected cost (22.128735) for application of policy OP2 (which is optimal for PBN2) on PBN1. Similarly for PBN2, the robust policy expected cost (10.217805) is much lower than the expected cost (14.125852) on application of OP1 on PBN2.

B. Design for an Example Based on Gene Expression Data

Simulations to study the robust policy design approach were performed on PBNs derived from gene expression data collected in a study of metastatic melanoma [23]. In this study, the abundance of mRNA for the gene WNT5A was found to be highly discriminating between cells with properties typically associated with high versus low metastatic competence. 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

[7]This policy is optimal for PBN2.

TABLE III
AVERAGE COST PER STATE TABLE FOR ∞ NORM CONSTRAINT

PBN No.	β	UC	WCUC	RWCC	NWCC
PBN1	0.05	29.986091	31.873921	25.6404	25.640452
	0.10	29.986091	33.638112	27.403454	27.403505
	0.15	29.986091	35.336042	29.14252	29.142322
PBN2	0.05	27.354292	29.549255	24.606525	24.6066
	0.10	27.354292	31.6137	26.530361	26.533224
	0.15	27.354292	33.581526	28.442474	28.444082

TABLE IV
AVERAGE COST PER STATE TABLE FOR INTERVAL CONSTRAINT

PBN No.	Interval	UC	WCUC	RWCC	NWCC
PBN1	0.05	29.986091	30.938946	24.805602	24.807347
	0.10	29.986091	31.878338	25.779872	25.784878
	0.15	29.986091	32.803885	26.756556	26.766682
PBN2	0.05	27.354292	28.489167	24.805602	24.807347
	0.10	27.354292	29.61195	25.779872	25.784878
	0.15	27.354292	30.720951	26.756556	26.766682

by the standard *in vitro* assays for metastasis [24]. Furthermore, it was found that an intervention that blocked the Wnt5a protein from activating its receptor, the use of an antibody that binds the Wnt5a protein, could substantially reduce Wnt5a's ability to induce a metastatic phenotype. This suggests a control strategy that reduces the WNT5A gene's action in affecting biological regulation, since the available data suggest that disruption of this influence could reduce the chance of a melanoma metastasizing, a desirable outcome. PBNs derived from the same expression data have been used in [11]–[14] for demonstrating earlier intervention strategies.

We consider seven gene networks containing the genes WNT5A, pirin, S100P, RET1, MART1, HADHB, and STC2. Note that the limited number of samples along with the absence of time-series data in human cancers restricts the ability to infer a network with high precision. Reverse engineering can produce a multitude of feasible PBNs instead of a unique one. We will report results of simulation studies where it is assumed that the actual transition matrix of the biological network varies within a range of the inferred transition probabilities. Currently available steady-state data may not support this kind of modeling but novel experimental techniques enabling the monitoring of time-series data for the governing regulatory proteins using imaging [25] might provide us with much better approximations to the transition probabilities and, consequently, the bounds on the intervals and the ∞-norms can be made quite small.

The robust control strategies described in Section V and the nominal (nonrobust) policies studied earlier in [14] have been applied to the generated PBNs with pirin chosen as the control gene ($u = 1$ signifying the state of pirin is reversed and $u = 0$ signifying no intervention) and $p = q = 0.01$.

The cost of control is assumed to be 1 and the states are assigned penalties as follows:

$$\tilde{g}(i,u,j) = \begin{cases} 5 & \text{if } u = 0 \text{ and WNT5A is 1 for state } j \\ 6 & \text{if } u = 1 \text{ and WNT5A is 1 for state } j \\ 1 & \text{if } u = 1 \text{ and WNT5A is 0 for state } j \\ 0 & \text{if } u = 0 \text{ and WNT5A is 0 for state } j. \end{cases}$$

The penalty assignment is based on the fact that for infinite-horizon problems, there is no terminal penalty; instead, the cost per stage $\tilde{g}$ contains the penalties of each state. Since our objective is to downregulate the WNT5A gene, a higher penalty is assigned for destination states having WNT5a upregulated. 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.

We performed a number of simulations for robust policies with both ∞-norm and interval matrix constraints and also generated original nonrobust policies. Simulation results suggest that for smaller values of β in the ∞-norm constraint, the policies obtained from robust design are very similar to the policies obtained without taking uncertainties into consideration. Furthermore, the difference in the infinite horizon costs on application of robust policy and the nominal policy was observed to be minimal. Here small β refers to $\beta < 0.15$. Similar results were obtained for the interval matrix constraint where the range of intervals was less than 0.15. We next provide tables containing the infinite-horizon costs for different conditions. Due to space constraints, we will limit the enumeration to two PBNs. Table III shows the average cost per state of two different PBNs obtained from the Melanoma Data for different values of β in the

∞-norm constraint. UC refers to the uncontrolled cost for the estimated PBN, $WCUC$ indicates the worst case uncontrolled cost, $RWCC$ refers to the worst case cost when Robust Policy is applied while $NWCC$ denotes the worst case cost when the nominal policy is applied. We notice that the worst case cost for the nominal and robust policies are very similar for small β. Table IV refers to the average cost per state for the same two PBNs but with uncertainties represented as interval matrices. The value of the interval denotes the percentage range of the transition probabilities. For example, an interval value of 0.10 in Table IV specifies that the probabilities $p_{i,j}$ are between $0.9p_{i,j}$ and $1.1p_{i,j}$. We observe similar results here, the nominal policy is quite robust and the fifth and sixth column entries are almost equal.

These results support the practical validity of the bounds obtained in Section III. Most of the transition matrices have perturbation bound[8] $k_3 < 1$ and hence for small $\beta = ||E||_\infty$, the changes in the steady state distribution will be small for any stationary policy. Consequently, the steady-state distribution corresponding to the stationary policy designed without taking uncertainties into consideration will be close to the steady-state distribution corresponding to the robust stationary policy. Here, we have used a discounted cost approach and the cost function for this formulation is closely related to the steady-state distribution over an infinite horizon with a discount factor α close to 1.

VII. Concluding Remarks

In this paper, we have studied the robustness of some of the intervention strategies that have been recently proposed for PBNs. We specifically focussed on the infinite horizon problem and examined how uncertainties in the transition probability matrix of the uncontrolled PBN show up in the steady-state distribution of the controlled PBN. Since the steady-state distribution of a PBN is thought to characterize the phenotype, our studies essentially seek to examine the effect of network uncertainty on the phenotype that would result from the application of intervention strategies.

Through analytical derivation and simulation studies, we demonstrated that the stationary infinite horizon optimal control policies proposed to date are quite robust with respect to network uncertainty. The intervention strategies for PBNs that have been proposed thus far are all limited to flipping the expression status of one or more genes in the network, and this is dictated by what interventions are implementable with the currently available biological techniques. This limited class of interventions ensures that the controlled probability transition matrix is related to the uncontrolled probability transition matrix via a linear transformation, and this is what made it possible to establish the robustness results of this paper.

We also characterized different classes of uncertainty that can arise in a PBN and determined how those uncertainties map into uncertainties in the corresponding transition matrix. Furthermore, for each uncertainty class, we presented intervention designs that are optimal for the worst-case scenarios. Such worst-case designs are appropriate if the PBN is being used to design cancer therapy since, in this case, we wish to avoid extremely undesirable results.

In this paper, we have not examined how the uncertainties in the various experimental steps involved in genetic regulatory network modeling propagate downstream to affect the results of the intervention. Future research should be directed at examining this uncertainty propagation and also obtaining less conservative policies.

[8]For the two specific PBNs in the Tables III and IV, perturbation bound k_3 is 0.796977 for PBN1 and 0.837857 for PBN2.

References

[1] S. A. Kauffman, "Metabolic stability and epigenesis in randomly constructed genetic nets," *Theor. Biol.*, vol. 22, pp. 437–467, 1969.

[2] S. A. Kauffman, "Homeostasis and differentiation in random genetic control networks," *Nature*, vol. 224, pp. 177–178, 1969.

[3] S. A. Kauffman, *The Origins of Order: Self-Organization and Selection in Evolution*. New York: Oxford Univ. Press, 1993.

[4] I. Shmulevich, E. R. Dougherty, and W. Zhang, "From Boolean to probabilistic Boolean networks as models of genetic regulatory networks," *Proc. IEEE*, vol. 90, pp. 1778–1792, 2002.

[5] 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, pp. 261–274, 2002.

[6] N. Friedman, M. Linial, I. Nachman, and D. Pe'er, "Bayesian networks to analyze expression data," in *Proc. 4th Ann. Int. Conf. Computat. Molec. Biol.*, 2000, pp. 127–135.

[7] R. F. Hashimoto, K. Seungchan, I. Shmulevich, W. Zhang, M. L. Bittner, and E. R. Dougherty, "Growing genetic regulatory networks from seed genes," *Bioinformatics*, vol. 20, pp. 1241–1247, 2004.

[8] R. Pal, I. Ivanov, A. Datta, M. L. Bittner, and E. R. Dougherty, "Generating Boolean networks with a prescribed attractor structure," *Bioinformatics*, vol. 21, pp. 4021–4025, 2005.

[9] X. Zhou, X. Wang, R. Pal, I. Ivanov, M. L. Bittner, and E. R. Dougherty, "A bayesian connectivity-based approach to constructing probabilistic gene regulatory networks," *Bioinformatics*, vol. 20, pp. 2918–2927, 2004.

[10] I. Shmulevich, E. R. Dougherty, and W. Zhang, "Gene perturbation and intervention in probabilistic Boolean networks," *Bioinformatics*, vol. 18, pp. 1319–1331, 2002.

[11] A. Datta, A. Choudhary, M. L. Bittner, and E. R. Dougherty, "External control in Markovian genetic regulatory networks," *Mach. Learn.*, vol. 52, pp. 169–191, 2003.

[12] A. Datta, A. Choudhary, M. L. Bittner, and E. R. Dougherty, "External control in Markovian genetic regulatory networks: The imperfect information case," *Bioinformatics*, vol. 20, pp. 924–930, 2004.

[13] R. Pal, A. Datta, M. L. Bittner, and E. R. Dougherty, "Intervention in context-sensitive probabilistic Boolean networks," *Bioinformatics*, vol. 21, pp. 1211–1218, 2005.

[14] R. Pal, A. Datta, and E. R. Dougherty, "Optimal infinite horizon control for probabilistic Boolean networks," *IEEE Trans. Signal Process.*, vol. 54, no. 6, pt. 2, pp. 2375–2387, June 2006.

[15] I. Shmulevich, E. R. Dougherty, and W. Zhang, "Control of stationary behaviour in probabilistic Boolean networks by means of structural intervention," *Biolog. Syst.*, vol. 10, pp. 431–446, 2002.

[16] G. E. Cho and C. D. Meyer, "Comparison of perturbation bounds for the stationary distribution of a Markov chain," *Linear Algebra and Its Applicat.*, vol. 335, pp. 137–150, 2001.

[17] E. Seneta, "Perturbation of the stationary distribution measured by ergodicity coefficients," *Adv. Appl. Process.*, pp. 228–230, 1988.

[18] E. Seneta, "Sensitivity of finite Markov chains under perturbation," *Statist. Probab. Lett.*, vol. 17, pp. 163–168, 1993.

[19] I. Ivanov, R. Pal, and E. Dougherty, "Dynamics-preserving size-reduction mappings for probabilistic Boolean networks," *IEEE Trans. Signal Process.*, to be published.

[20] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, 2nd ed. Belmont, MA: Athena Scientific, 2001.

[21] A. Nilim and L. El Ghaoui, Robust Markov Decision Processes With Uncertain Transition Matrices EECS Dep., Univ. Calif., Berkeley, UCB/ERL M04/28, 2004.

[22] G. Iyengar, Robust Dynamic Programming 2002.

[23] 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, and V. Sondak, "Molecular classification of cutaneous malignant melanoma by gene expression profiling," *Nature*, vol. 406, pp. 536–540, 2000.

[24] A. T. Weeraratna, Y. Jiang, G. Hostetter, K. Rosenblatt, P. Duray, M. Bittner, and J. M. Trent, "Wnt5a signalling directly affects cell motility and invasion of metastatic melanoma," *Cancer Cell*, vol. 1, pp. 279–288, 2002.
[25] S. Cookson, N. Ostroff, W. L. Pang, D. Volfson, and J. Hasty, "Monitoring dynamics of single-cell gene expression over multiple cell cycles," *Molec. Syst. Biol.*, vol. 1, 2005.

Ranadip Pal (S'05–M'07) received the B.Tech. degree in electronics and electrical communication engineering from the Indian Institute of Technology, Kharagpur, in 2002 and the M.S. and Ph.D. degrees in electrical engineering from Texas A&M University, College Station, in 2004 and 2007, respectively.

In August 2007, he joined the Department of Electrical Engineering, Texas Tech University, Lubbock, as an Assistant Professor. His research areas are computational biology, genomic signal processing, and control of genetic regulatory networks.

Dr. Pal was the recipient of the Ebensberger/Fouraker Fellowship while at Texas A&M. He was also a recipient of the Distinguished Graduate Student Masters Research Award, a National Instruments Fellowship, and an Indian National Math Olympiad Awardee. He was ranked first in the Regional Math Olympiad, West Bengal, India.

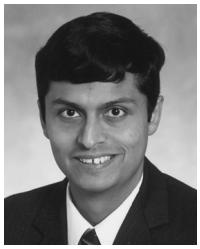

Aniruddha Datta (S'87–M'91–SM'97) received the B.Tech. degree in electrical engineering from the Indian Institute of Technology, Kharagpur, in 1985, the M.S.E.E. degree from Southern Illinois University, Carbondale, in 1987, and the M.S. (applied mathematics) and Ph.D. degrees from the University of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

In August 1991, he joined the Department of Electrical Engineering at Texas A&M University, College Station, where he is currently a Professor. His areas of interest include adaptive control, robust control, PID control, and genomic signal processing. From 2001 to 2003, he worked on 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* (New York: Springer-Verlag, 1998), a coauthor (with M. T. Ho and S. P. Bhattacharyya) of the book *Structure and Synthesis of PID Controllers* (New York: Springer-Verlag, 2000), a coauthor (with G. J. Silva and S. P. Bhattacharyya) of the book *PID Controllers for Time Delay Systems* (Boston, MA: Birkhauser, 2004), and a coauthor (with E. R. Dougherty) of the book *Introduction to Genomic Signal Processing with Control* (Boca Raton, FL: CRC, 2007).

Dr. Datta has served as 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, and the Ph.D. degree in mathematics from Rutgers University, New Brunswick, NJ, and has been awarded the *Doctor Honoris Causa* by the Tampere University of Technology, Finland.

He is a Professor in the Department of Electrical and Computer Engineering, Texas A&M University, College Station, where he holds the Robert M. Kennedy Chair and is Director of the Genomic Signal Processing Laboratory. He is also the Director of the Computational Biology Division of the Translational Genomics Research Institute, Phoenix, AZ. He is an author of 14 books, editor of five others, and author of more than 200 journal papers. 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 research in genomic signal processing is aimed at 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.

Dr. Dougherty is a Fellow of the International Society of Optical Engineering (SPIE), has received the SPIE President's Award, and served as the editor of the SPIE/IS&T *Journal of Electronic Imaging*. At Texas A&M University, he received the Association of Former Students Distinguished Achievement Award in Research. He has been named Fellow of the Texas Engineering Experiment Station, and named Halliburton Professor of the Dwight Look College of Engineering.