Bayesian Robustness in the Control of Gene Regulatory Networks

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

***Abstract*—The errors originating in the data extraction process, gene selection and network inference prevent the transition probabilities of a gene regulatory network from being accurately estimated. Thus, it is important to study the effect of modeling errors on the final outcome of an intervention strategy and to design robust intervention strategies. Two major approaches applied to the design of robust policies in general are the minimax (worst case) approach and the Bayesian approach. The minimax control approach is typically conservative because it gives too much importance to the scenarios which hardly occur in practice. Consequently, in this paper, we formulate the Bayesian approach for the control of gene regulatory networks. We characterize the errors emanating from the data extraction and inference processes and compare the performance of the minimax and Bayesian designs based on these uncertainties.**

***Index Terms*—Bayesian robustness, gene regulatory networks, intervention, parameter estimation, robust control.**

I. Introduction

SINCE systems biology is focused on understanding the detailed molecular interactions that contribute to cell functioning, a genetic regulatory network designed for facilitating such an understanding should mimic the actual biological interactions in as much detail as possible, at least to the extent that it is justified by the available data. On the other hand, in translational genomics the focus is on developing therapeutic interventions and a 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 sufficiently capture the effects of intervention manifested at the phenotypic (observational) level. Such a coarse model can be used to develop and evaluate suitable (control) strategies for therapeutic intervention. Coarsely quantized Probabilistic Boolean Networks (PBNs) have been used in recent years to carry out such intervention studies [1]–[3].

Manuscript received December 07, 2008; accepted April 17, 2009. First published May 12, 2009; current version published August 12, 2009. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Xiaodong Cai. This research was supported in part by the National Science Foundation under Grants ECS-0355227, CCF-0514644, and ECCS-0701531 and in part by the National Cancer Institute under Grant CA90301.

R. Pal is with the Department of Electrical and Computer Engineering Texas Tech University, Lubbock, TX 79409-3102 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 77843-3128 USA (e-mail: datta@ece.tamu.edu).

E. R. Dougherty is with the Department of Electrical and Computer Engineering, Texas A&M University, College Station, TX 77843-3128 USA, and also with the Translational Genomics Research Institute, Phoenix, AZ, 85004 USA (e-mail: edward@ece.tamu.edu).

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.2009.2022872

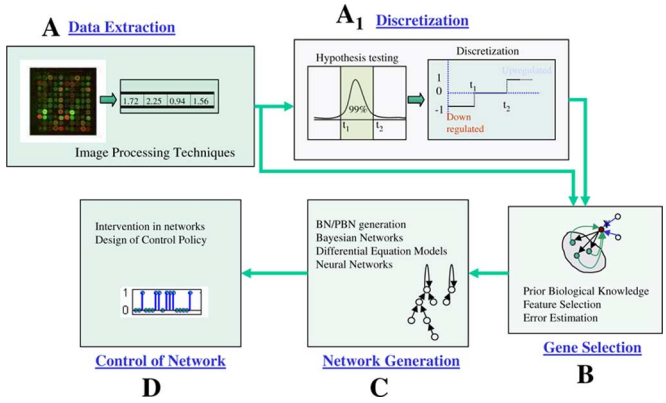

Fig. 1. Basic Steps involved in Modeling and Control of Genetic Networks: (A)Extraction of gene expression data (A_1) Discretization of the Data; (B) Selection of genes to build the network; (C) Network Generation from the available data and prior biological knowledge; and (D) Intervention in the network with the objective of moving the network from undesirable to desirable states.

A broad overview of the steps involved in the modeling and control of genetic networks is 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. This step is not required when using a continuous model such as a differential equation model. On the other hand, this step is crucial for discrete models such as Boolean Networks (BNs) [4], Probabilistic Boolean Networks (PBNs) [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 for at least two reasons: i) 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 ii) 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 for construction [1].

From Fig. 1, it is clear that errors made during data extraction, discretization, gene selection and network generation will 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. Studying the effects on intervention outcome of the errors propagating from the different upstream steps is an important open problem in translational genomics. In this paper, we will assume that the noise during data extraction and discretization manifests itself as uncertainties in the transition probability matrix.

To date, a number of approaches have been proposed in the literature for carrying out interventions in Probabilistic Boolean networks [7], [2], [3], [8], [9]. Of these, the one proposed in [3] is particularly relevant for translational genomics since it seeks to shift the steady-state mass of the PBN from undesirable states to desirable ones. Since it has been conjectured that the steady-state behaviour of a PBN is indicative of the phenotype [4], it is likely that alterations in the steady-state behaviour of the PBN would translate into changes at the phenotypic level. In the design methods just cited, including [3], the intervention policy is derived under the assumption that the transition probability matrix of the PBN is exactly known. Since this assumption is not likely to be satisfied in practice, one may wish to derive a policy taking into account uncertainties in the model.

The minimax approach to robust policy design in the context of PBNs has been explored in [10]. Minimax policies are typically conservative because they are designed to achieve the best worst-case performance. Some previous studies have considered minimax control policies in the context of Markov decision processes [11], [12]. More generally, the minimax approach to filter design has a long history in signal processing [13]–[16]. A more recent approach to robust filtering achieves robustness by taking into account the distribution of the uncertainties [17], [18]. This results in a Bayesian design where the optimal filter depends on the uncertainty distribution. This approach avoids the detrimental effects of extreme, but rare, states on minimax design, the goal being to produce the best average performance across the uncertainty distribution. The Bayesian methodology has been applied to the design of classifiers in the face of uncertainty, a common problem in gene-expression-based classification [19]. In this paper, we take the Bayesian approach to design robust control policies for PBNs and compare the results to minimax design and to a global design procedure [17], [19] based on the mean of the uncertainty distribution.

The paper is organized as follows. Section II provides an overview of Probabilistic Boolean Networks. Section III characterizes possible errors from the data extraction process and inference of PBN. In Section IV, we discuss robust intervention designs to satisfy various robustness criteria. Section V uses simulations to carry out a comparative study of the different robust design strategies proposed. Finally, Section VI contains some concluding remarks.

II. Probabilistic Boolean Networks

A *Boolean Network* (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 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 *Probabilistic Boolean Network* (PBN) consists of a set of nodes/genes $V = \{x_1, \ldots, x_n\}$, $x_i \in \{0,1,\ldots,d\}$, $i = 1, \ldots, n$, 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 $f_{ji} : \{0,1,\ldots,d\}^n \to \{0,1,\ldots,d\}$, $i = 1, \ldots, n$. 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 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 function $\mathbf{f}_j$ determines a constituent network, or *context*, and the PBN behaves as a fixed constituent network 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 $\mathbf{f}_1, \mathbf{f}_2, \ldots, \mathbf{f}_k$. Note that a decision to switch does not necessarily mean that a different network function is selected, since, if the current network function is $\mathbf{f}_j$, then with probability c_j the same function may be selected. If $q < 1$, the PBN is said to be *context-sensitive*; if $q = 1$, a network function is randomly selected at every time point and the PBN is said to be *instantaneously random*. We consider PBNs with perturbation, meaning that at each time point there is a perturbation probability p of any gene flipping its value uniformly randomly. To date, PBNs have been applied with $d = 1$ or $d = 2$, the former producing constituent networks that are BNs and the latter modeling gene expression as down-regulated (0), up-regulated (2), or invariant (1). Here we confine ourselves to binary $(d = 1)$ PBNs. This eases notation without loss of generality regarding robustness theory.

Associated with a PBN is a Markov chain describing the transition probabilities. For a context-sensitive PBN, at any time point the network state is described by a pair, $(\mathbf{f}, x)$, where $\mathbf{f}$ is the operative network function determining a context and x is the vector of gene expressions, called the *gene activity profile* (GAP). In short, the states of the Markov chain corresponding to a context-sensitive PBN are (context, GAP) pairs. For an instantaneously random PBN, the chain states are just GAPs. We focus on instantaneously random PBNs. This is consistent with [3], eases the mathematical description, and fulfills the goal of elucidating Bayesian robust design.

III. Quantification of Possible Network Errors

In this section, we quantify the errors originating from the data extraction and inference processes. Our objective is to examine how these errors propagate into the transition probabilities for the resulting Markov chain. In other words, we wish to represent the final Markov chain as a function of the noise parameters.

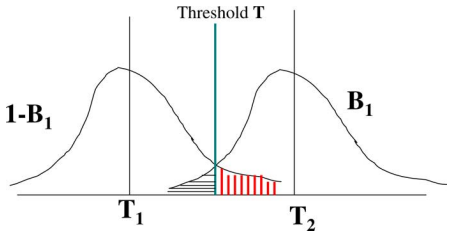

Fig. 2. Noise distributions.

A. Noise Quantification for Data Extraction

Gene-expression data are usually extracted from bulk scale methods such as microrrays [20] or individual cell specific methods employing fluorescent reporter levels [21]. These approaches are dependent on the analysis of images such as of microarray slide images, images of fluorescent cells, etc. Removal of background noise or identifying the signal location is usually the starting procedure. It can be approached by segmentation algorithms using global or local thresholding. Our interests lie in discrete networks and thus we are concerned with the probability of discrete values being changed due to noise. Hypothesis tests on the expression levels of housekeeping genes or duplicate genes are often employed to calculate the thresholds for discretization [20]. For the purpose of binarization, the assumption is that some genes will be expressed more than other genes and the more expressed genes will be given value 1 while the less expressed or unexpressed genes will be given value 0. The demarcation in the ideal case will be by a threshold above which there will be a cluster of expressed genes and below which there will be a cluster of unexpressed genes. Due to variability and errors in the extraction process, there will, in fact, be overlapping clusters. Let us assume that the data distribution after normalization is bimodal, consisting of a zero mean Gaussian distribution with variance σ_1^2 centered around T_1, signifying the unexpressed genes, and a zero mean Gaussian distribution with variance σ_2^2 centered around T_2, signifying the expressed genes (Fig. 2). A heuristic for the prior probabilities of the expressed and the unexpressed genes can be obtained from the Boolean network parameter commonly known as bias [22]. This bias B is different from the biases in the microarray. B refers to the probability that a Boolean function takes on the value 1. We will assume that $B_1 = B$ is the prior probability of the expressed genes and $1 - B_1$ refers to the prior for the unexpressed genes. The optimal threshold for discretization based on Bayes' rule is T, where T satisfies the equation

$$\frac{1}{\sqrt{2\pi}\sigma_2} B_1 e^{-\frac{(T-T_2)^2}{2\sigma_2^2}} = \frac{1}{\sqrt{2\pi}\sigma_1}(1-B_1)e^{-\frac{(T-T_1)^2}{2\sigma_1^2}}. \quad (1)$$

The threshold T can be calculated by solving the following quadratic equation:

$$T^2\left(\sigma_2^2-\sigma_1^2\right)+2T\left(\sigma_1^2T_2-\sigma_2^2T_1\right)-2\sigma_1^2\sigma_2^2\ln\left(\frac{\sigma_2(1-B_1)}{\sigma_1 B_1}\right) +\sigma_2^2T_1^2-\sigma_1^2T_2^2=0. \quad (2)$$

For $\sigma_1 = \sigma_2$, we get a simplified expression for T:

$$T = \frac{T_1+T_2}{2} + \frac{\sigma^2 \ln\left(\frac{1-B_1}{B_1}\right)}{T_2 - T_1}. \quad (3)$$

Let w be the probability of changing the value of a particular gene due to noise (w referring to the shaded area in Fig. 2):

$$\begin{aligned} w &= \text{Probability of denoting a} \\ &\quad \text{not expressed gene as expressed} \\ &\quad + \text{Probability of denoting an} \\ &\quad \text{expressed gene as not expressed} \\ &= \text{vertical shaded area} \\ &\quad + \text{horizontal shaded area} \\ &= (1-B_1)\frac{1}{\sqrt{2\pi}\sigma_1}\int_T^{\infty} e^{-\frac{(x-T_1)^2}{2\sigma_1^2}}dx \\ &\quad + B_1\frac{1}{\sqrt{2\pi}\sigma_2}\int_{-\infty}^{T} e^{-\frac{(x-T_2)^2}{2\sigma_2^2}}dx \\ &= (1-B_1)\frac{1}{\sqrt{\pi}}\int_{\frac{T-T_1}{\sqrt{2}\sigma_1}}^{\infty} e^{-u^2}du \\ &\quad + B_1\frac{1}{\sqrt{\pi}}\int_{-\infty}^{\frac{T-T_2}{\sqrt{2}\sigma_2}} e^{-u^2}du \\ &= (1-B_1)\left(\frac{1}{2}-\frac{1}{\sqrt{\pi}}\int_0^{\frac{T-T_1}{\sqrt{2}\sigma_1}} e^{-u^2}du\right) \\ &\quad + B_1\left(\frac{1}{2}-\frac{1}{\sqrt{\pi}}\int_0^{\frac{T_2-T}{\sqrt{2}\sigma_2}} e^{-u^2}du\right) \\ &= \frac{1}{2}-\frac{1-B_1}{2}\text{erf}\left(\frac{T-T_1}{\sqrt{2}\sigma_1}\right) \\ &\quad - \frac{B_1}{2}\text{erf}\left(\frac{T_2-T}{\sqrt{2}\sigma_2}\right) \end{aligned} \quad (4)$$

where erf refers to the error function, defined by $\text{erf}(x) = (2)/(\sqrt{\pi})\int_0^x e^{-w^2}dw$.

The calculation of the probability transition matrix in the presence of noise is presented next. Suppose we have a probability transition matrix P_1 without noise, where $P_1(i,j)$ denotes the probability of going from state (GAP) i to state (GAP) j. The states are represented in decimal form instead of the binary forms of genes, i.e., $00\ldots0$ to $11\cdots1$ are represented as 0 to 2^n-1 assuming n genes. The noise introduction characterized by w will change the probability transition matrix from P_1 to P_2. Let us denote the observed states corresponding to transition matrix P_2 as Z^{ob} and the actual states corresponding to P_1 as Z^{ac}. Then (5) shown at the bottom of the next page relates P_2 to P_1, where $h_{k,j}$ represents the hamming distance between k and j (number of bits different in the binary representation of k and j) and $(1-w)^{n-h_{i,i'}}w^{h_{i,i'}}$ is the probability of observing state i when the actual state is i' due to noise characterized by w.

To represent (5) in a matrix form, we will introduce the following notations. Let η_t represent the actual state probability

distribution vector at time t, i.e., $\eta_t = [P(Z_t^{\text{ac}} = 0) \quad P(Z_t^{\text{ac}} = 1) \ldots\ldots\ldots P(Z_t^{\text{ac}} = 2^n - 1)]$. Let us define a matrix $Q_{n,w}$ of dimension $2^n \times 2^n$ as

$$Q_{n,w} = (1-w)^n \times \begin{bmatrix} 1 & \frac{w}{1-w} & \cdots & \left(\frac{w}{1-w}\right)^n \\ \frac{w}{1-w} & 1 & \cdots & \left(\frac{w}{1-w}\right)^{n-1} \\ \frac{w}{1-w} & \left(\frac{w}{1-w}\right)^2 & \cdots & \left(\frac{w}{1-w}\right)^{n-1} \\ \cdots & \cdots & \cdots & \cdots \\ \cdots & \cdots & \cdots & \cdots \\ \left(\frac{w}{1-w}\right)^n & \left(\frac{w}{1-w}\right)^{n-1} & \cdots & 1 \end{bmatrix} \quad (6)$$

where the individual $2^n \times 2^n$ terms of $Q_{n,w}$ for $i = 0, 1, \ldots, 2^n - 1$ and $j = 0, 1, \ldots, 2^n - 1$ are

$$Q_{n,w}(i,j) = Q_{n,w}(j,i) = (1-w)^n \left(\frac{w}{1-w}\right)^{h_{i,j}}. \quad (7)$$

We define a matrix Π_{n,t,η_t} of dimension $2^n \times 2^n$ as

$$\Pi_{n,t,\eta_t} = \begin{bmatrix} \eta_t(0) & 0 & \cdots & 0 \\ 0 & \eta_t(1) & \cdots & 0 \\ 0 & 0 & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots \\ \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & \cdots & \eta_t(2^n - 1) \end{bmatrix}. \quad (8)$$

The matrix Ψ_{n,w,t,η_t} of dimension $2^n \times 2^n$ is defined as (9) shown at the bottom of the page, where the individual $2^n \times 2^n$ terms of Ψ_{n,w,t,η_t} for $i = 0, 1, \ldots, 2^n - 1$ and $j = 0, 1, \ldots, 2^n - 1$ are

$$\Psi_{n,w,t,\eta_t}(i,j) = \begin{cases} \sum_{i'=0}^{2^n-1} (1-w)^{n-h_{i,i'}} w^{h_{i,i'}} \eta_t(i') & \text{if } i = j \\ 0 & \text{otherwise.} \end{cases}$$

With the help of the above definitions, (5) can be represented in the matrix form as

$$\Psi_{n,w,t,\eta_t} P_2 = Q_{n,w} \Pi_{n,t,\eta_t} P_1 Q_{n,w}. \quad (10)$$

Our objective is to infer the actual transition probability matrix P_1 from the observed transition probability matrix P_2. When $Q_{n,w}^{-1}$ and Π_{n,t,η_t}^{-1} exist, P_1 can be found from P_2 using the following equation:

$$P_1 = \Pi_{n,t,\eta_t}^{-1} Q_{n,w}^{-1} \Psi_{n,w,t,\eta_t} P_2 Q_{n,w}^{-1}. \quad (11)$$

$$\begin{aligned} P_2(i,j) &= P\left(Z_{t+1}^{\text{ob}} = j \,|\, Z_t^{\text{ob}} = i\right) \\ &= \frac{P\left(Z_{t+1}^{\text{ob}} = j, Z_t^{\text{ob}} = i\right)}{P\left(Z_t^{\text{ob}} = i\right)} \\ &= \frac{\sum_{j'=0}^{2^n-1} \sum_{i'=0}^{2^n-1} P\left(Z_{t+1}^{\text{ob}} = j, Z_t^{\text{ob}} = i, Z_{t+1}^{\text{ac}} = j', Z_t^{\text{ac}} = i'\right)}{\sum_{i'=0}^{2^n-1} P\left(Z_t^{\text{ob}} = i, Z_t^{\text{ac}} = i'\right)} \\ &= \frac{\sum_{j'=0}^{2^n-1} \sum_{i'=0}^{2^n-1} P\left(Z_{t+1}^{\text{ob}} = j \,|\, Z_{t+1}^{\text{ac}} = j', Z_t^{\text{ob}} = i, Z_t^{\text{ac}} = i'\right) P\left(Z_{t+1}^{\text{ac}} = j', Z_t^{\text{ob}} = i, Z_t^{\text{ac}} = i'\right)}{\sum_{i'=0}^{2^n-1} P\left(Z_t^{\text{ob}} = i \,|\, Z_t^{\text{ac}} = i'\right) P\left(Z_t^{\text{ac}} = i'\right)} \\ &= \frac{\sum_{j'=0}^{2^n-1} \sum_{i'=0}^{2^n-1} P\left(Z_{t+1}^{\text{ob}} = j \,|\, Z_{t+1}^{\text{ac}} = j'\right) P\left(Z_{t+1}^{\text{ac}} = j' \,|\, Z_t^{\text{ob}} = i, Z_t^{\text{ac}} = i'\right) P\left(Z_t^{\text{ob}} = i, Z_t^{\text{ac}} = i'\right)}{\sum_{i'=0}^{2^n-1} P\left(Z_t^{\text{ob}} = i \,|\, Z_t^{\text{ac}} = i'\right) P\left(Z_t^{\text{ac}} = i'\right)} \\ &\quad \left[\text{As } P\left(Z_{t+1}^{\text{ob}} = j \,|\, Z_{t+1}^{\text{ac}} = j', Z_t^{\text{ob}} = i, Z_t^{\text{ac}} = i'\right) = P\left(Z_{t+1}^{\text{ob}} = j \,|\, Z_{t+1}^{\text{ac}} = j'\right)\right] \\ &= \frac{\sum_{j'=0}^{2^n-1} \sum_{i'=0}^{2^n-1} P\left(Z_{t+1}^{\text{ob}} = j \,|\, Z_{t+1}^{\text{ac}} = j'\right) P\left(Z_{t+1}^{\text{ac}} = j' \,|\, Z_t^{\text{ac}} = i'\right) P\left(Z_t^{\text{ob}} = i \,|\, Z_t^{\text{ac}} = i'\right) P\left(Z_t^{\text{ac}} = i'\right)}{\sum_{i'=0}^{2^n-1} (1-w)^{n-h_{i,i'}} w^{h_{i,i'}} P\left(Z_t^{\text{ac}} = i'\right)} \\ &\quad \left[\text{As } P\left(Z_{t+1}^{\text{ac}} = j' \,|\, Z_t^{\text{ob}} = i, Z_t^{\text{ac}} = i'\right) = P\left(Z_{t+1}^{\text{ac}} = j' \,|\, Z_t^{\text{ac}} = i'\right)\right] \\ &= \frac{\sum_{j'=0}^{2^n-1} \sum_{i'=0}^{2^n-1} (1-w)^{n-h_{j,j'}} w^{h_{j,j'}} (1-w)^{n-h_{i,i'}} w^{h_{i,i'}} P_1(i',j') P\left(Z_t^{\text{ac}} = i'\right)}{\sum_{i'=0}^{2^n-1} (1-w)^{n-h_{i,i'}} w^{h_{i,i'}} P\left(Z_t^{\text{ac}} = i'\right)} \end{aligned} \quad (5)$$

$$\Psi_{n,w,t,\eta_t} = \begin{bmatrix} \sum_{i'=0}^{2^n-1} (1-w)^{n-h_{0,i'}} w^{h_{0,i'}} \eta_t(i') & 0 & \cdots \\ 0 & \sum_{i'=0}^{2^n-1} (1-w)^{n-h_{1,i'}} w^{h_{1,i'}} \eta_t(i') & \cdots \\ 0 & 0 & \cdots \\ \cdots & \cdots & \cdots \\ \cdots & \cdots & \cdots \\ 0 & 0 & \cdots \end{bmatrix} \quad (9)$$

The existence of Π_{n,t,η_t}^{-1} is guaranteed for non-zero state probabilities and is given by

$$\Pi_{n,t,\eta_t}^{-1} = \begin{bmatrix} 1/\eta_t(0) & 0 & \cdots & 0 \\ 0 & 1/\eta_t(1) & \cdots & 0 \\ 0 & 0 & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots \\ \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & \cdots & 1/\eta_t(2^n-1) \end{bmatrix}. \quad (12)$$

In case of instantaneously random PBNs with perturbations, we will have non-zero steady state probabilities for each state.[1]

Let $\eta_t^{\text{ob}} = [P(Z_t^{\text{ob}}=0) \quad P(Z_t^{\text{ob}}=1) \ldots\ldots\ldots P(Z_t^{\text{ob}} = 2^n-1)]$ denote the observed state probability distribution vector at time t. As $P(Z_t^{\text{ob}} = i) = \sum_{i'=0}^{2^n-1} P(Z_t^{\text{ob}} = i \,|\, Z_t^{\text{ac}} = i')P(Z_t^{\text{ac}} = i')$, the relationship between η_t^{ob} and η_t is as follows:

$$\eta_t^{\text{ob}}(i) = \sum_{i'=0}^{2^n-1} (1-w)^{n-h_{i,i'}} w^{h_{i,i'}} \eta_t(i'). \quad (13)$$

The matrix equivalent of (13) is

$$\eta_t^{\text{ob}} = \eta_t Q_{n,w}. \quad (14)$$

The following theorem states the condition on w for the existence of $Q_{n,w}^{-1}$.

Theorem III.1: $Q_{n,w}$ is non-singular for $w \neq 0.5$

Proof: We will prove by induction that $Q_{n,w}$ is non-singular if $Q_{n-1,w}$ is non-singular. For $n=1$,

$$Q_{1,w} = \begin{bmatrix} (1-w) & w \\ w & (1-w) \end{bmatrix}$$

has determinant $1-2w$ which is *nonzero* for $w \neq 0.5$. Thus, $Q_{1,w}^{-1}$ exists.

For $n > 1, Q_{n,w}$ can be written in terms of $Q_{n-1,w}$ in the following manner

$$Q_{n,w} = \begin{bmatrix} (1-w)Q_{n-1,w} & wQ_{n-1,w} \\ wQ_{n-1,w} & (1-w)Q_{n-1,w} \end{bmatrix}. \quad (15)$$

Let us define $U_{n,w}$ as

$$\frac{1}{1-2w} \begin{bmatrix} (1-w)Q_{n-1,w}^{-1} & -wQ_{n-1,w}^{-1} \\ -wQ_{n-1,w}^{-1} & (1-w)Q_{n-1,w}^{-1} \end{bmatrix}. \quad (16)$$

It is easy to see that $Q_{n,w}U_{n,w} = I$. As $Q_{n-1,w}$ is non-singular due to the induction hypothesis, $U_{n,w} = Q_{n,w}^{-1}$ exists and has finite values for $1-2w \neq 0$.

For the simulations at the end of the paper, we will consider the case that the state transition observations are conducted once the network has reached a steady state. Thus, η_t will approach the steady state probability distribution vector and from now onwards, η without any subscript will denote the actual steady state probability distribution. At steady state, the time dependence of Π_{n,t,η_t} and Ψ_{n,w,t,η_t} will be removed. Consequently, $\Pi_{n,\eta}$ and $\Psi_{n,w,\eta}$ will refer to $\Pi_{n,\infty,\eta_\infty}$ and $\Psi_{n,w,\infty,\eta_\infty}$, i.e., their steady state values.

For a gene regulatory network, we observe the transition probability matrix P_2 and the actual transition probability matrix P_1 is calculated from P_2 using (11). Any random transition probability matrix P_2 will not necessarily produce a valid transition probability matrix P_1 using (11). We are interested in knowing the constraints on P_2 and w such that $\Pi_{n,\eta}^{-1}Q_{n,w}^{-1}\Psi_{n,w,\eta}P_2Q_{n,w}^{-1}$ is a valid transition probability transition matrix, i.e., all entries of $\Pi_{n,\eta}^{-1}Q_{n,w}^{-1}\Psi_{n,w,\eta}P_2Q_{n,w}^{-1}$ are non negative and the sum of each row is equal to 1.

For the purpose of simulations, let us assume that a valid transition probability matrix P_2 has been generated from a transition probability matrix P using the following equation

$$P_2 = \Psi_{n,a,\eta}^{-1} Q_{n,a} \Pi_{n,\eta} P Q_{n,a}. \quad (17)$$

There are no constraints on P other than being a valid transition probability matrix. Let η denote the steady state probability distribution vector of the starting PBN P and η^{est} denote the estimation of η from observed transition probability matrix P_2 assuming the noise probability to be w. When $w=a$, then $\eta^{\text{est}} = \eta$. Let η^{ob} denote the observed steady state probability vector. The different η's are related by the following equations:

$$\eta^{\text{ob}} = \eta Q_{n,a} \quad (18)$$

$$\eta^{\text{est}} = \eta^{\text{ob}} Q_{n,w}^{-1} \quad (19)$$

For P_2 generated as in (17), the following theorem provides a sufficient condition for $\Pi_{n,\eta^{\text{est}}}^{-1}Q_{n,w}^{-1}\Psi_{n,w,\eta^{\text{est}}}P_2Q_{n,w}^{-1}$ to be a valid transition probability matrix.

Theorem III.2: Let P be a transition probability matrix of dimension $2^n \times 2^n$ and $P_2 = \Psi_{n,a,\eta}^{-1} Q_{n,a} \Pi_{n,\eta} P Q_{n,a}$. A sufficient condition for $\Pi_{n,\eta^{\text{est}}}^{-1}Q_{n,w}^{-1}\Psi_{n,w,\eta^{\text{est}}}P_2Q_{n,w}^{-1}$ to be a valid transition probability matrix is $w \leq a < 0.5$.

Proof: $\Pi_{n,\eta^{\text{est}}}^{-1}Q_{n,w}^{-1}\Psi_{n,w,\eta^{\text{est}}}P_2Q_{n,w}^{-1} = \Pi_{n,\eta^{\text{est}}}^{-1}Q_{n,w}^{-1}\Psi_{n,w,\eta^{\text{est}}}\Psi_{n,a,\eta}^{-1}Q_{n,a}\Pi_{n,\eta}PQ_{n,a}Q_{n,w}^{-1}$ Since,

$$\Psi_{n,w,\eta^{\text{est}}}(i,j) = \begin{cases} \eta^{\text{ob}}(i) & \text{if } i=j \\ 0 & \text{otherwise} \end{cases}$$

and

$$\Psi_{n,a,\eta}(i,j) = \begin{cases} \eta^{\text{ob}}(i) & \text{if } i=j \\ 0 & \text{otherwise,} \end{cases}$$

$\Psi_{n,w,\eta^{\text{est}}} = \Psi_{n,a,\eta}$.

Thus $\Psi_{n,w,\eta^{\text{est}}}\Psi_{n,a,\eta}^{-1} = I$ where I is the Identity matrix.

We next show that $Q_{n,w}^{-1}Q_{n,a}$ is a valid transition probability matrix for $w \leq a < 0.5$.

From Theorem III.1, we have

$$Q_{n,w}^{-1} = \frac{1}{1-2w} \begin{bmatrix} (1-w)Q_{n-1,w}^{-1} & -wQ_{n-1,w}^{-1} \\ -wQ_{n-1,w}^{-1} & (1-w)Q_{n-1,w}^{-1} \end{bmatrix}$$

and

$$Q_{n,a} = \begin{bmatrix} (1-a)Q_{n-1,a} & aQ_{n-1,a} \\ aQ_{n-1,a} & (1-a)Q_{n-1,a} \end{bmatrix}.$$

[1]In the current formulation of the PBN with perturbation, all the transition probabilities will be non-zero except possibly from state i to state i. Thus, starting from any state probability distribution vector, there is a possibility that in the next stage, one of the state probabilities is zero. But for $t > 2$, η_t has all non-zero entries in the current formulation.

Thus [see the equation shown at the bottom of the page], for $n = 1$, we have

$$Q_{1,w}^{-1}Q_{1,a} = \frac{1}{1-2w}\begin{bmatrix}(1-w-a) & (a-w)\\ (a-w) & (1-w-a)\end{bmatrix}$$

which has all non negative entries for $w \leq a < 0.5$ and each row summing up to 1. By induction, $Q_{n,w}^{-1}Q_{n,a}$ also has all non negative entries and the sum of each row equal to 1. We can similarly show that $Q_{n,a}Q_{n,w}^{-1}$ is a valid transition probability matrix.

Let R_2 denote the matrix $Q_{n,w}^{-1}Q_{n,a}\Pi_{n,\eta}$. From the relationship between the different steady state vectors η's shown earlier, we have

$$\eta^{\text{est}} = \eta^{\text{ob}}Q_{n,w}^{-1} = \eta Q_{n,a}Q_{n,w}^{-1}$$

As Q's are symmetric, $(\eta^{\text{est}})^T = Q_{n,w}^{-1}Q_{n,a}\eta^T$ where T denotes the transpose. From this relationship, we can deduce that the sum of the ith row for the matrix R_2 will be equal to $\eta^{\text{est}}(i)$. Thus, $R_1 = \Pi_{n,\eta^{\text{est}}}^{-1}R_2$ will have each row summing up to 1 as $\Pi_{n,\eta^{\text{est}}}^{-1}(i,i) = 1/\eta^{\text{est}}(i)$ and $\Pi_{n,\eta^{\text{est}}}^{-1}(i,j) = 0$ for $i \neq j$. Furthermore, each of the entries in the matrices are positive and hence their multiplication will produce positive entries. Finally, as multiplication of probability transition matrices results in a probability transition matrix, $\Pi_{n,\eta^{\text{est}}}^{-1}Q_{n,w}^{-1}\Psi_{n,w,\eta^{\text{est}}}\Psi_{n,a,\eta}^{-1}Q_{n,a}\Pi_{n,\eta}PQ_{n,a}Q_{n,w}^{-1} = R_1PQ_{n,a}Q_{n,w}^{-1}$ is a valid transition probability matrix.

B. Uncertainty in the Network Selection Probabilities of a PBN

We consider an instantaneously random PBN consisting of n genes and composed of k Boolean Networks with perturbation probability p and selection probabilities $c_1, c_2, \ldots, c_k$. We will assume that we are able to accurately infer the perturbation probability p and the transitions of the constituent Boolean Networks, but there are uncertainties in the network selection probabilities c_i's. The robustness analysis will be based on different distributions of c_i's.

IV. Robust Intervention

In this section, our goal is to discuss various forms of optimal robust policies, minimax, Bayesian, and global, the main focus being on Bayesian. We first briefly describe optimal infinite-horizon intervention for the nominal case, i.e., in the absence of uncertainty [3].

A. Nominal Case

As shown in [23] and [3], a PBN with control can be modeled as a stationary discrete-time dynamic system

$$z_{t+1} = f(z_t, u_t, d_t),\ t = 0, 1, \ldots, \tag{20}$$

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 d_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 on n genes, $S = \{0, 1, 2, \ldots\ldots, 2^n - 1\}$ 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 binary control inputs. The disturbance d_t is manifested in terms of change of network due to latent variables or other external disturbances. The random disturbances $d_t, t = 0, 1, \ldots$ have identical statistics and are characterized by probabilities $P(\cdot|x_t, u_t)$ defined on D, where $P(\cdot|x_t, u_t)$ is the probability of occurrence of d_t, when the current state and control are x_t and u_t, respectively. d_t is independent of prior disturbances $d_0, d_1 \ldots d_{t-1}$.

An equivalent way to represent the dynamical system (20) is as a finite state Markov chain described by the control-dependent one-step transition probability $p_{ij}(u)$, where for any $t = 0, 1, 2, \ldots; i, j \in S$ and $u \in C$,

$$p_{ij}(u) := P(z_{t+1} = j | z_t = i, u_t = u). \tag{21}$$

In this paper, we will interchangeably use either representation (20) or (21) depending on their suitability for a particular situation or a particular derivation.

We will consider a discounted cost infinite horizon approach as in [3]. The transition probability matrix P_w is dependent on the parameter w which is a function of the parameters $T_1, T_2, \sigma_1, \sigma_2$ and B_1 as shown in (4) and (5). For the simulation examples to be provided later in this paper, some of the parameters will follow a beta distribution which in turn will produce a distribution of possible transition probability matrices.

When the transition probabilities are exactly known, the states make transitions according to $\theta := (P^u)_{u \in C}$. 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 Υ represent the set of all possible μ's, i.e., the set of all possible control policies.

Let $J_{\mu,\theta}$ denote the expected total cost for the discounted cost infinite-horizon problem under control policy μ and transitions θ:

$$J_{\mu,\theta}(z_0) = \lim_{M \to \infty} E_{d_t, t=0,1,\ldots}\left\{\sum_{t=0}^{M-1} \alpha^t \tilde{g}(z_t, u_t, d_t)\right\}, \tag{22}$$

$$Q_{n,w}^{-1}Q_{n,a} = \frac{1}{1-2w}\begin{bmatrix}(1-w-a)Q_{n-1,w}^{-1}Q_{n-1,a} & (a-w)Q_{n-1,w}^{-1}Q_{n-1,a}\\ (a-w)Q_{n-1,w}^{-1}Q_{n-1,a} & (1-w-a)Q_{n-1,w}^{-1}Q_{n-1,a}\end{bmatrix}.$$

where $0 < \alpha < 1$ denotes the discount factor and $\tilde{g}(z_t, u_t, d_t)$ represents the cost of going from state z_t to z_{t+1}[2] under the control action u_t. $\tilde{g}$ 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 expectation E is taken over the disturbance d which is characterized by the state transitions θ.

The *nominal* problem here corresponds to minimizing the cost in (22). Consequently, the optimal infinite-horizon discounted cost for the *nominal* problem is given by:

$$\Phi(\Upsilon, \theta, z_0) := \min_{\mu \in \Upsilon} J_{\mu,\theta}(z_0). \tag{23}$$

B. Robust Policies

Owing to uncertainties in the transition probability matrix, there can be various possible transitions corresponding to any given choice of control u and, therefore, various θ's. In case of uncertainties, we will parameterize the class of transitions as $\Omega := (P_a^u)_{u \in C, a \in F_a}$, where F_a is the noise parameter distribution. In Section III.A, F_a refers to the joint distribution of $T_1, T_2, \sigma_1, \sigma_2$ and B_1. Similar to the robustness analysis of classifiers in [19], we will consider three different forms of robust intervention in genetic networks.

A minimax intervention policy is defined as a policy whose worst performance over the uncertainty class Ω is best among all admissible policies. The *minimax robust policy*, denoted μ_{mm}, is the one that satisfies

$$\Phi(\Upsilon, \Omega, z_0) := \min_{\mu \in \Upsilon} \max_{\theta \in \Omega} J_{\mu,\theta}(z_0). \tag{24}$$

There is no guarantee that a minimax robust solution will exist for all distributions. For some uncertainty classes, analytical solutions are feasible. In [10], we have worked out the minimax robust solution for three different uncertainty classes.

A Bayesian robust policy minimizes the average cost over the uncertainty class Ω. Let μ_b denote the policy designed to be optimal at point b of the parameter distribution i.e., μ_b minimizes $J_{\mu,\theta_b}(z_0)$. Here, θ_b refers to the probability transition matrix at point b.

Let $\gamma_a(\mu_b)$ denote the expected cost per state at point a of the parameter distribution for the intervention policy μ_b:

$$\gamma_a(\mu_b) = E_{z_0}[J_{\mu_b,\theta_a}(z_0)] \tag{25}$$

where the expectation over the states z_0 of the network is taken to arrive at a single value for representing the cost of a policy. Consequently, the expected cost of $\gamma_a(\mu_b)$ over the distribution F_a is given by

$$E_a[\gamma_a(\mu_b)] = E_a[E_{z_0}[J_{\mu_b,\theta_a}(z_0)]] \tag{26}$$

where E_a denotes expectation relative to the parameter distribution. A Bayesian control policy, denoted by μ_{b^*}, is one that minimizes $E_a[\gamma_a(\mu_b)]$. The parameter value b^* that achieves the minimum is called a *maximally robust parameter value* and the corresponding policy is called the *Bayesian robust policy*. If the actual parameter for the distribution is b and one optimizes at b, then there is no gain, and perhaps a loss, in using the Bayesian robust policy because $\gamma_b(\mu_b) \leq \gamma_b(\mu_{b^*})$; however, in the face of uncertainty regarding the actual parameter, the gain in applying the Bayesian robust policy is $E_a[\gamma_a(\mu_b)] - E_a[\gamma_a(\mu_{b^*})]$, which is always nonnegative and can be substantial for parameter values far from the maximally robust one. We should note that here, we are minimizing the expected cost over the set of control policies that are optimal for various points in the distribution. A policy χ can exist that is not optimal for any point in the distribution but with $E_a[\gamma_a(\chi)] \leq E_a[\gamma_a(\mu_{b*})]$. We are not aiming to find the policy χ as no analytical methods exist for generation of such a policy and simulation in that case is prohibitive due to the large number of possible policies. The problem with a dynamic programming based value iteration approach for directly calculating either χ or μ_{b*} is that the principle of optimality is not valid for the current setting.

Rather than trying to design robust optimal policies at all parameter values and then select the maximally robust policy, we can instead design a single global optimal policy by using a centrality parameter, such as the mean or median, to represent the mass of the uncertainty distribution. Here we will use the mean of the distribution as the global parameter value for designing the intervention strategy. The *global policy*, denoted by $\mu_\bullet$, is the optimal policy for the global parameter value.

For the simulations, we will generate n transition probability matrices $P_1, P_2, \ldots P_n$ corresponding to n points $b_1, b_2, \ldots\ldots b_n$, on the parameter distribution. The Bayesian policy is selected from one of the optimal policies for points $b_1, b_2, \ldots\ldots b_n$. To have a consistent comparison with the Bayesian approach, the set of admissible policies Υ for the minimax policy will be $[\mu_{b_1}, \mu_{b_2}, \ldots\ldots \mu_{b_n}]$.

V. Performance Comparison

In this section, we compare the performances of robust policy designs in the framework of errors emanating from noise in the data extraction process and errors in the model parameter estimation. We vary various parameters as described in Section III-A and plot the following nine cost curves:

$\gamma_a(\mu_b) = E_{z_0}[J_{\mu_b,\theta_a}(z_0)]$—Cost at parameter value a using the optimal policy designed at b.

$\gamma_b(\mu_b) = E_{z_0}[J_{\mu_b,\theta_b}(z_0)]$—Cost at parameter value b using the optimal policy designed at b.

$\gamma_b(\mu_{b^*}) = E_{z_0}[J_{\mu_{b^*},\theta_b}(z_0)]$—Cost of the Bayesian robust policy at parameter value b.

$\gamma_b(\mu_\bullet) = E_{z_0}[J_{\mu_\bullet,\theta_b}(z_0)]$—Cost of the global robust policy at parameter value b.

$\gamma_b(\mu_{\mathrm{mm}}) = E_{z_0}[J_{\mu_{\mathrm{mm}},\theta_b}(z_0)]$—Cost of the minimax robust policy at parameter value b.

$E_a[\gamma_a(\mu_b)]$—Expected cost of the optimal policy designed at b across the parameter space.

$E_a[\gamma_a(\mu_{b^*})]$—Expected cost of the Bayesian robust policy across the parameter space.

$E_a[\gamma_a(\mu_\bullet)]$—Expected cost of the global robust policy across the parameter space.

$E_a[\gamma_a(\mu_{\mathrm{mm}})]$—Expected cost of the minimax robust policy across the parameter space.

[2]Note that a Markov chain can be modeled by $z_{t+1} = d_t$ [24]. Hence, the destination state is the same as the disturbance.

A. Noise in the Data Extraction Process

To compare different robust policies relative to data-extraction noise, we derive an instantaneously random PBN from gene expression data collected in a study of metastatic melanoma [25]. 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 by the standard *in vitro* assays for metastasis [26]. 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 [2], [27], [23], [3] for demonstrating earlier intervention strategies.

We consider a seven-gene network containing the genes WNT5A, pirin, S100P, RET1, MART1, HADHB, and STC2. To obtain the PBN, we have used the algorithms described in [28] to construct three probable Boolean networks to use as the constituent Boolean networks in the PBN. The three BNs are provided in the companion website http://cvial.ece.ttu.edu/ranadippal/robust/supplement.htm. Each constituent network is assumed to be derived from steady-state gene-expression data (a common assumption—see [28]). The states are ordered as WNT5A, pirin, S100P, RET1, MART1, HADHB, and STC2, with WNT5A as the most significant bit (MSB) and STC2 as the least significant bit (LSB). An instantaneously random PBN P with perturbation probability $p = 0.01$ is constructed from the 3 BNs. Let η denote the steady state probability distribution vector for the PBN P. The observed PBN P_2 for the simulations is taken to be $P_2 = \Psi^{-1}_{7,.4,\eta} Q_{7,.4} \Pi_{7,\eta} P Q_{7,.4}$. The distribution of the possible networks is generated from the observed transition probability matrix P_2 using (11). The w in (11) incorporates the uncertainties. As shown earlier, w is dependent on noise variances σ_1 and σ_2, bimodal distribution center points T_1 and T_2, and bias B. Thus, n different sets of parameter values $[\sigma_{1,i}, \sigma_{2,i}, T_{1,i}, T_{2,i}, B_{1,i}]$, $i = 1, \ldots n$, produce n possible PBNs generating the observed transition probability matrix P_2.

The control problem formulation is similar to the one given in [3]. The states 0 to 63 are assumed to be the desirable states while the states 64 to 127 are assumed to be the undesirable ones. The second gene is used as the control gene. The cost of control is assumed to be 1 and the states are assigned penalties as follows:

$$\tilde{g}(i,u,j) = \begin{cases} 0 & \text{if } u = 0 \text{ and } 0 \le j \le 63 \\ 1 & \text{if } u = 1 \text{ and } 0 \le j \le 63 \\ 5 & \text{if } u = 0 \text{ and } 64 \le j \le 127 \\ 6 & \text{if } u = 1 \text{ and } 64 \le j \le 127 \end{cases} \tag{27}$$

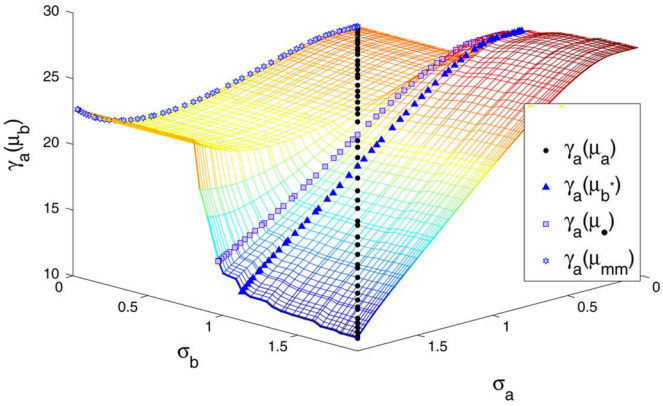

Fig. 3. Cost for policies designed at σ_b over the set of all possible σ_a's.

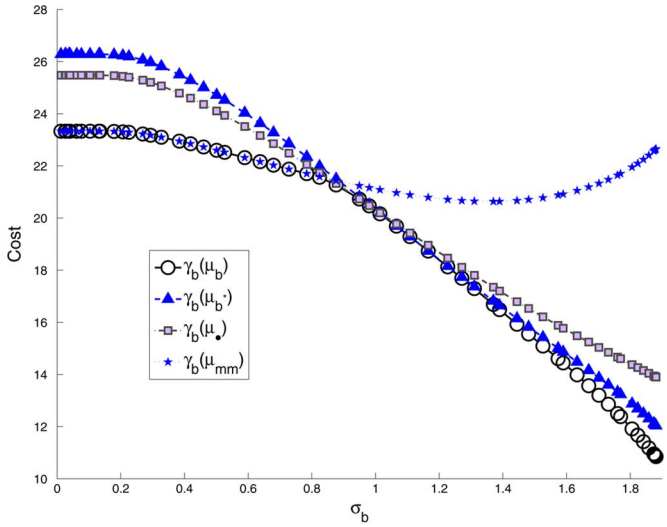

Fig. 4. Cost for various policies.

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. The discount factor α for the following simulations is taken to be 0.9.

1) Uncertainty in σ With Rest of the Parameters Constant: We consider equal noise variances for the expressed and non-expressed gene data, $\sigma_1 = \sigma_2 = \sigma$, and σ varies from 0 to 1.9 with $\sigma/1.9$ following a beta distribution having parameters $\alpha = 0.5$ and $\beta = 0.5$. The remaining parameters are fixed with the following values: $B_1 = 0.5, T_1 = 0, T_2 = 1$. Fig. 3 shows the cost for the different points in the distribution (σ_a) when the policy is designed at particular points in the distribution (σ_b) and applied at σ_a. Figs. 4 and 5 show the costs and expected costs, respectively, as functions of the parameter σ_b ranging from 0 to 1.9.

Fig. 3 shows that for a network with parameter value σ_a, the cost is minimum when the policy designed at parameter value σ_a is applied to it. The minimum for each σ_a is represented by black dots in the figure. Fig. 3 also shows the cost for the Bayesian (triangles), minimax (stars) and global (squares) robust policies.

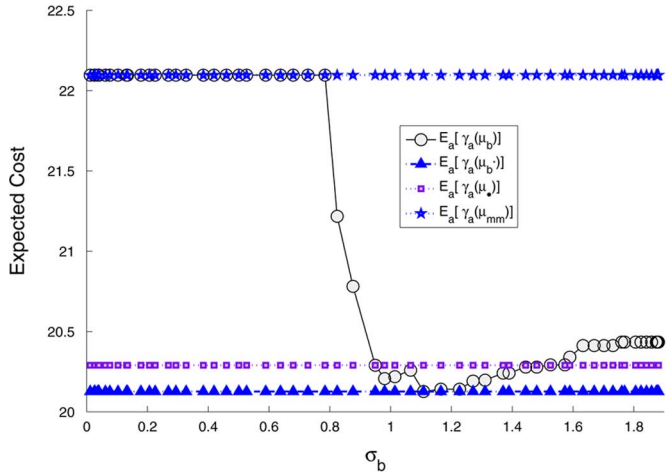

Fig. 5. Expected cost for various policies.

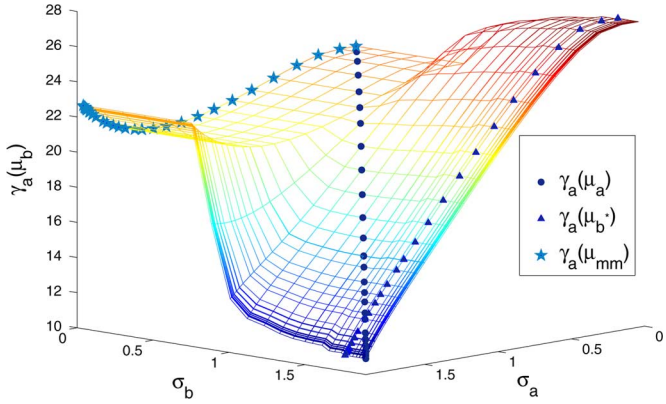

Fig. 6. Cost for policies designed at σ_b over the set of all possible σ_a's.

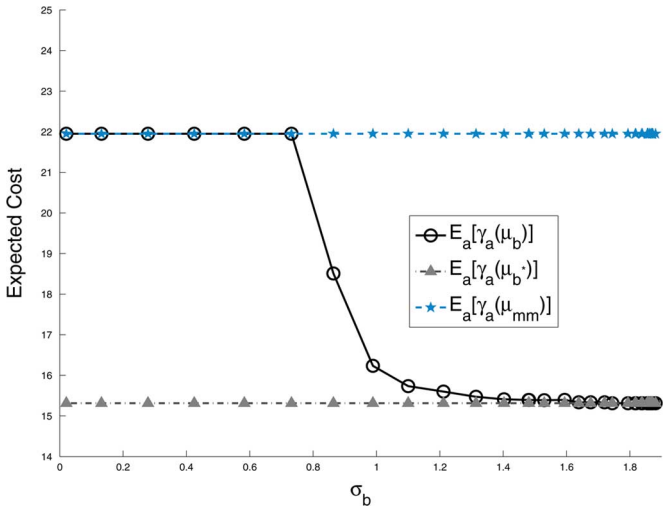

Fig. 7. Expected cost for various policies.

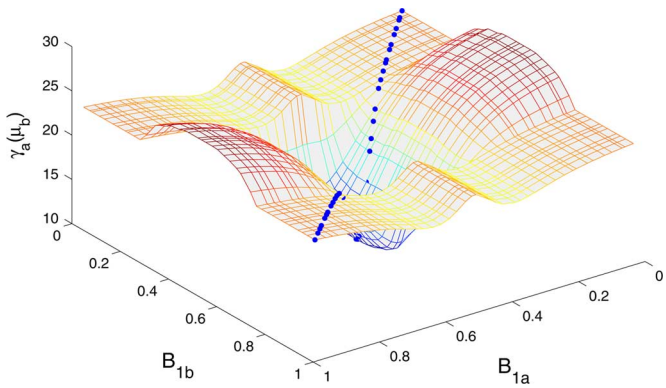

Fig. 8. Cost for policies designed at B_{1b} over the set of all possible B_{1a}'s.

We notice that as the objective of the minimax policy is to minimize the maximum cost, its selects the policy designed at $\sigma_b = 0$ which has a lower maximum but has a higher average value. The global robust policy is designed at the mean of the distribution $1.9\alpha/(\alpha+\beta) = 1.9/2 = 0.95$. The maximally robust parameter value is $\sigma_b = 1.11$ as denoted by the triangles in the figure. The costs of the different policies (optimal, Bayesian, global and minimax) when applied to networks with different σ's are further plotted in 2 dimension in Fig. 4 for easier comparison. Note that $\gamma_b(\mu_b) \leq \gamma_b(\mu_{b^*})$, with $\gamma_b(\mu_b) = \gamma_b(\mu_{b^*})$ for $\sigma = 1.11$ $(b = b^*)$, the maximally robust parameter value.

Fig. 4 shows that the cost for the Bayesian policy is much less than the cost for the minimax Policy for higher σ's but for smaller $\sigma's$, the minimax is performing better. But as the difference between the minimax and the Bayesian policy costs are more for higher σ's compared to the lower σ's, the expected cost for the Bayesian Policy is much less than the expected cost for the minimax policy as shown in Fig. 5. The expected cost curve, $E_a[\gamma_a(\mu_b)]$, minimizes at the maximally robust parameter value $\sigma = 1.11$. Except for that value, there is a gain, $E_a[\gamma_a(\mu_b)] - E_a[\gamma_a(\mu_{b^*})]$, from using the Bayesian robust policy and this gain can be significant, especially for smaller σ. The expected cost, $E_a[\gamma_a(\mu_{b^*})]$, using Bayesian robust policy is less than that of using either the global or minimax robust policies, these being given by $E_a[\gamma_a(\mu_\bullet)]$ and $E_a[\gamma_a(\mu\text{mm})]$, respectively.

To elucidate the conservativeness of the minimax policy, another simulation was conducted with the parameter $\sigma/1.9$ following a β distribution with $\alpha = 0.8$ and $\beta = 0.2$. This set of parameters for the beta distribution produces low probability for lower σ's and most of the probability mass is concentrated near $\sigma = 1.9$. The cost for the different points in the distribution (σ_a) when the policy is designed at particular points in the distribution (σ_b) is shown in Fig. 6. The expected cost curves are plotted in Fig. 7. When compared to the previous case shown in Figs. 5, 7 shows that the expected cost for the minimax policy has not changed much but the expected cost for the Bayesian Policy has decreased by around 25%. The difference in the expected cost for the minimax and the Bayesian Policy has increased from $22.1 - 20.13 = 1.97$ to $21.95 - 15.31 = 6.64$. This highlights the case that minimax policies give equal importance to cases with low probabilities of occurrence (here low σ's) and therefore produce an overly conservative policy, as shown in Fig. 7.

2) Uncertainty in B_1 With Other Parameters Constant: We vary B_1 from 0 to 1 uniformly. The other parameters are fixed at $\sigma_1 = \sigma_2 = 1.9$, $T_1 = 0$ and $T_2 = 1$. Fig. 8 shows the cost for the different points in the distribution (B_{1a}) when the policy is designed at particular points in the distribution (B_{1b}). Figs. 9

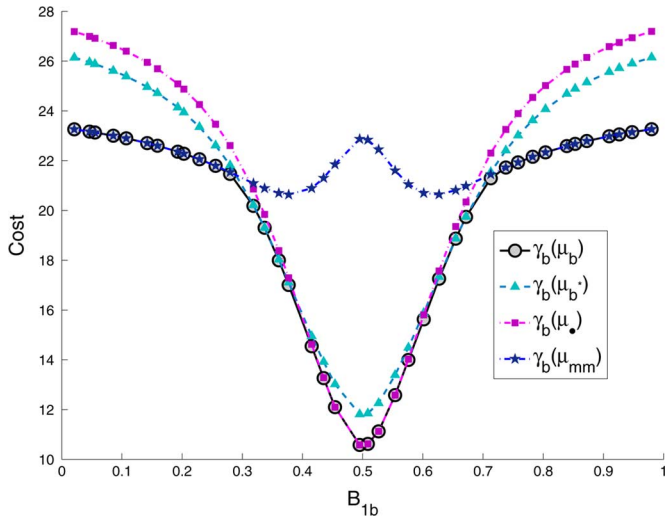

Fig. 9. Cost for various policies.

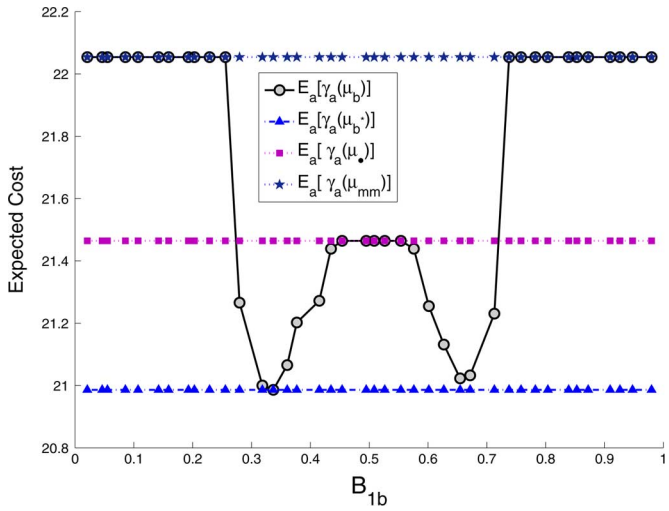

Fig. 10. Expected cost for various policies.

and 10 correspond to Figs. 4 and 5, respectively, and analogous observations apply. For instance, the Bayesian robust policy performs much better than the minimax policy, which again suffers from extreme cases in the distribution with minimal probability. The maximally robust parameter value is $B_1 = 0.34$ and the global robust policy is designed at the mean of the distribution, which is equal to 0.5.

From Fig. 8, we notice that the shape of the curve to be vastly different from Fig. 3. The reason is that here w increases in the range $B_1 = 0$ to $B_1 = 0.5$ while it decreases for $B_1 = 0.5$ to $B_1 = 1$, attaining a maximum at $B_1 = 0.5$, whereas in the previous case, it is a monotonically increasing function for σ, attaining the maximum at the end point $\sigma = 1.9$.

3) Uncertainty in T_1 With Other Parameters Constant: We vary T_1 from -0.75 to 0.75 with $(T_1 + 0.75)/1.5$ following a beta distribution with $\alpha = 5$ and $\beta = 1$. The other parameters are $T_2 = 1$, $B_1 = 0.5$ and $\sigma_1 = \sigma_2 = 0.5$. The cost curves are shown in Figs. 11, 12 and 13. As shown in Fig. 13, the expected cost is minimum for the Bayesian policy followed by the global policy, and both outperform the minimax policy. The maximally robust parameter value is $T_1 = 0.57$, which is different from the mean of the distribution, which is 0.5.

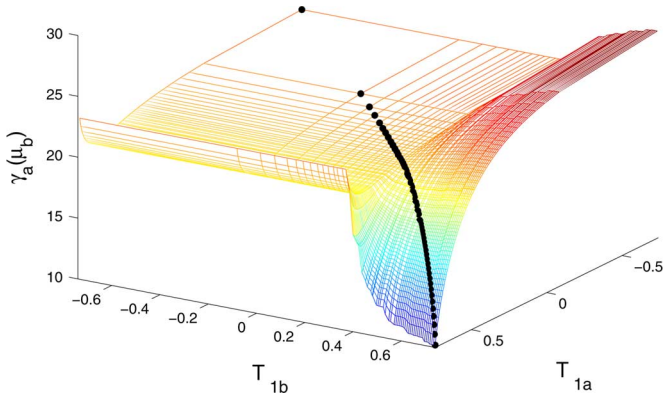

Fig. 11. Cost for policies designed at T_{1b} over the set of all possible T_{1a}'s.

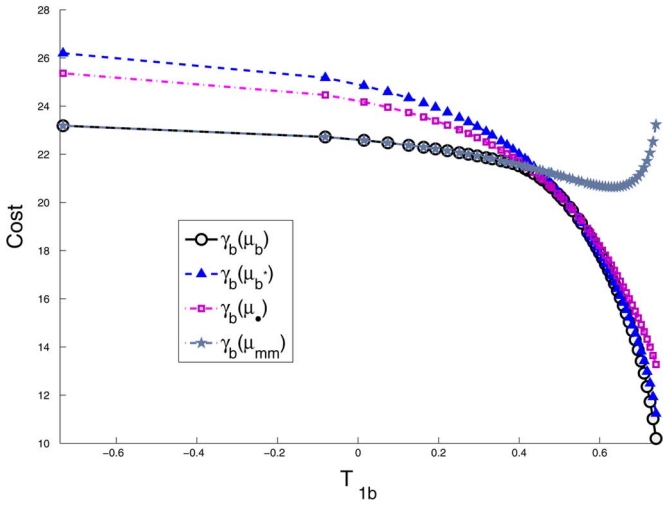

Fig. 12. Cost for various policies.

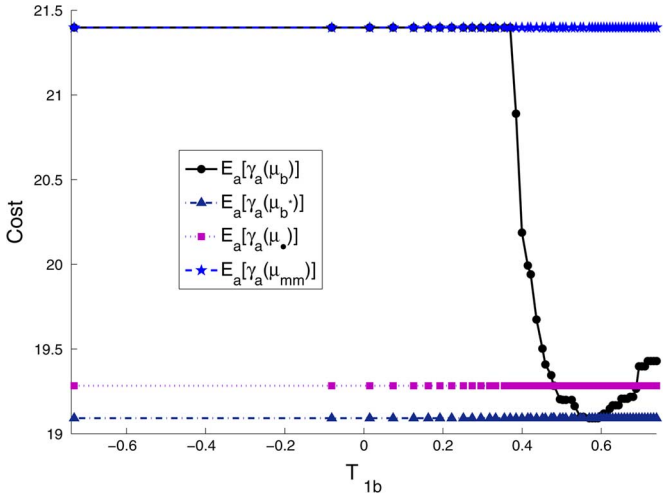

Fig. 13. Expected cost for various policies.

B. Uncertainty in the Switching Probabilities of a PBN

To avoid the complication of multiple switching probabilities and thereby better graphically illustrate robustness rela-

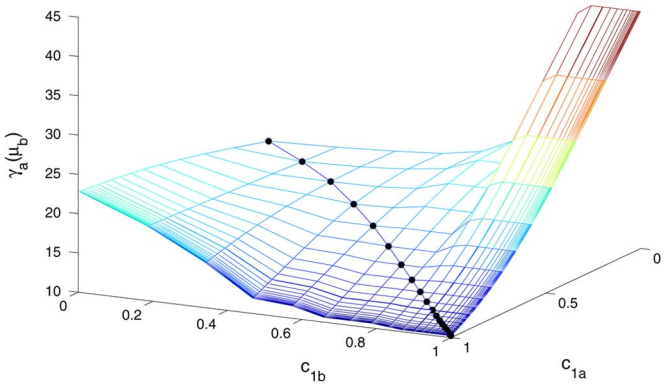

Fig. 14. Cost for policies designed at c_{1b} over the set of all possible c_{1a}'s.

tive to switching probabilities, we consider an artificial example of an instantaneously random PBN consisting of 2 constituent Boolean networks of 4 genes each. The transition diagrams of the BNs are provided in the companion website. We vary c_1 following a beta distribution having parameters $\alpha = 1$ and $\beta = 0.1$. The selection probability c_2 of the other network is calculated from the equation $c_1 + c_2 = 1$. The parameter p is fixed at $p = 0.1$. The control problem for the Markov Chain with dimension 16×16 is formulated with cost of control being 0.2 and the states assigned penalties as follows:

$$\tilde{g}(i,u,j) = \begin{cases} 0 & \text{if } u = 0 \text{ and } 0 \le j \le 7 \\ 0.2 & \text{if } u = 1 \text{ and } 0 \le j \le 7 \\ 5 & \text{if } u = 0 \text{ and } 8 \le j \le 15 \\ 5.2 & \text{if } u = 1 \text{ and } 8 \le j \le 15 \end{cases} \tag{28}$$

The penalty structure indicates that the highly desirable states have the 1st gene down-regulated (0) while the undesirable states have the first gene up-regulated (1). The discount factor α is taken to be 0.9.

The cost curves are shown in Figs. 14, 15 and 16. Fig. 15 shows that the cost for low values of c_1 is lower with the minimax policy compared to the Bayesian and global policies but the performance of the minimax policy deteriorates for higher values of c_1. Based on the current parameters of the Beta distribution, the probability of occurrence of low values of c_1 is much less than that of high values of c_1. This curve indicates that the minimax design gives equal importance to all cases irrespective of their probability of occurrence. With regard to expected cost, the minimax policy performs much worse than the Bayesian or global robust policies, as shown in Fig. 16. The maximally robust parameter value is $c_1 = 0.658$ and the globally robust parameter value is the mean of the distribution $\alpha/(\alpha+\beta) = 1/1.1 = 0.909$.

VI. Conclusion

For an intervention policy to work in practice, it has to be robust. As demonstrated by the simulation examples in this paper, the various forms of robustness discussed here have their advantages and disadvantages. Minimax policies are typically conservative and give too much weight to parameter values that may rarely occur in practice. In that respect, the Bayesian robust policy design is suitable as it gives better expected performance. Global policy design is advisable when we want to take account of the parameter distribution but only have limited knowledge

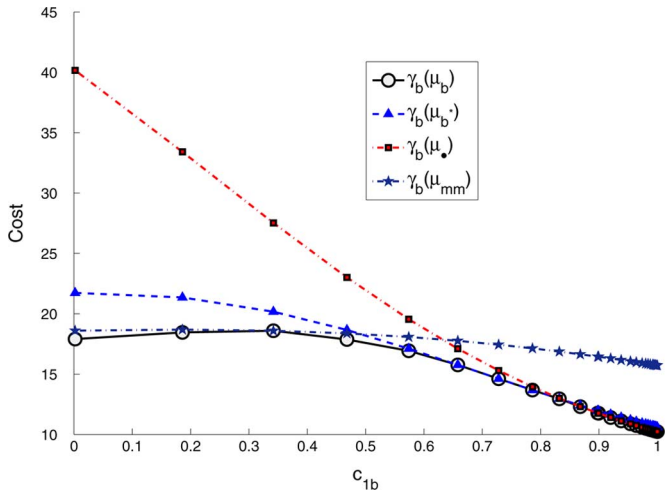

Fig. 15. Cost for various policies.

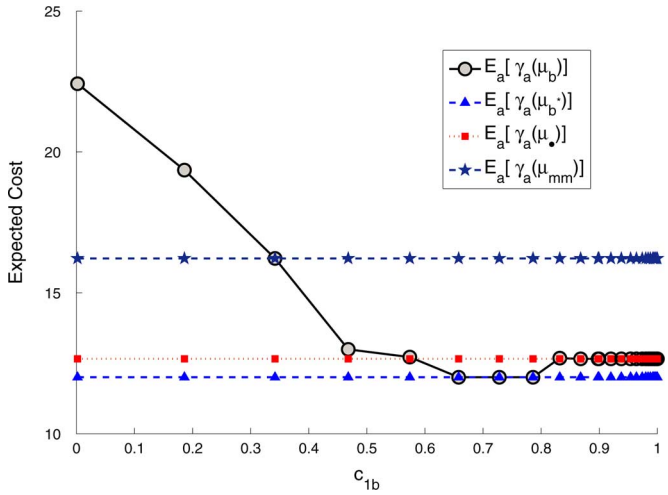

Fig. 16. Expected cost for various policies.

of it, say, an estimate of the mean or median. Based on the application, an appropriate design strategy should be selected. As a general rule, for avoiding extremely undesirable results, a minimax design is recommended. On the other hand, when our objective is to improve the expected chance of success, a Bayesian robust policy is to be preferred.

References

[1] A. Datta and E. R. Dougherty, *Introduction to Genomic Signal Processing With Control*. Boca Raton, FL: CRC Press, 2007.

[2] 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.

[3] 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.

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

[5] 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.

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

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

[8] W.-K. Ching, S.-Q. Zhang, Y. Jiao, T. Akutsu, and A. S. Wong, "Optimal finite-horizon control for probabilistic boolean networks with hard constraints," in *Int. Symp. Optimization Systems Biology*, 2007, pp. 21–28.
[9] P. C. Y. Chen and J. W. Chen, "A Markovian approach to the control of genetic regulatory networks," *BIOSYSTEMS*, vol. 90, pp. 535–545, 2007.
[10] R. Pal, A. Datta, and E. R. Dougherty, "Robust intervention in probabilistic boolean networks," *IEEE Trans. Signal Process.*, vol. 56, no. 3, pp. 1280–1294, Mar. 2008.
[11] A. Nilim and L. El Ghaoui, "Robust Markov decision processes with uncertain transition matrices," EECS Dept., Univ. of California, Berkeley, CA, Tech. Rep. UCB/ERL M04/28, 2004.
[12] G. Lyengar, "Robust dynamic programming," *Math. Oper. Res.*, vol. 30, pp. 1–21, 2005.
[13] V. P. Kuznetsov, "Stable detection when the signal and spectrum of normal noise are inaccurately known," *Telecommun. Radio Eng.*, vol. 30–31, pp. 58–64, 1976.
[14] S. A. Kassam and T. I. Lim, "Robust wiener filters," *Franklin Inst.*, vol. 304, pp. 171–185, 1977.
[15] H. V. Poor, "On robust wiener filtering," *IEEE Trans. Autom. Control*, vol. 26, pp. 531–536, 1980.
[16] K. S. Vastola and H. V. Poor, "Robust Wiener–Kolmogorov theory," *IEEE Trans. Inf. Theory*, vol. 28, pp. 316–327, 1983.
[17] A. M. Grygorian and E. R. Dougherty, "Design and analysis of robust optimal binary filters in the context of a prior distribution for the states of nature," *Math. Imag. Vis.*, vol. 11, pp. 239–254, 2001.
[18] A. M. Grygorian and E. R. Dougherty, "Bayesian robust optimal linear filters," *Signal Process.*, vol. 81, pp. 2503–2521, 2001.
[19] E. R. Dougherty, J. Hua, Z. Xiong, and Y. Chen, "Optimal robust classifiers," *Pattern Recogn.*, vol. 38, pp. 1520–1532, 2005.
[20] Y. Chen, E. R. Dougherty, and M. L. Bittner, "Ratio-based decisions and the quantitative analysis of CDNA microarray images," *J. Biomed. Opt.*, vol. 2, pp. 364–374, 1997.
[21] D. Longo and J. Hasty, "Dynamics of single-cell gene expression," *Molecul. Syst. Biol.*, vol. 2, pp. 4100–4110, 2006.
[22] I. Shmulevich and S. A. Kauffman, "Activities and sensitivities in boolean network models," *Phys. Rev. Lett.*, vol. 93, pp. 048701–4, 2004.
[23] 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.
[24] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, 2nd ed. Singapore: Athena Scientific, 2001.
[25] 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.
[26] 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.
[27] 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.
[28] 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.

Ranadip Pal (M'09) received the B.Tech. degree in electronics and electrical communication engineering from the Indian Institute of Technology, Kharagpur, India, 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 and Computer Engineering at 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 Ebsenberger/Fouraker Fellowship, a Distinguished Graduate Student Masters Research Award, and a National Instruments Fellowship while at Texas A&M. He was also an Indian National Math Olympiad Awardee and was ranked first in the Regional Math Olympiad, West Bengal, India.

Aniruddha Datta (F'09) 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. degree (in applied mathematics) and the Ph.D. degree from the University of Southern California, Los Angeles, in 1991.

In August 1991, he joined the Department of Electrical and Computer Engineering at Texas A&M University, College Station, where he is currently a Professor and holder of the J. W. Runyon, Jr., 35 Professorship II. 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* (Springer-Verlag 1998), a coauthor (with M. T. Ho and S. P. Bhattacharyya) of the book *Structure and Synthesis of PID Controllers* (Springer-Verlag, 2000), a coauthor (with G. J. Silva and S. P. Bhattacharyya) of the book *PID Controllers for Time Delay Systems* (Birkhauser, 2004), a coauthor (with E. R. Dougherty) of the book *Introduction to Genomic Signal Processing with Control* (CRC Press, 2007), and a coauthor (with S. P. Bhattacharyya and L. H. Keel) of the book *Linear Control Theory: Structure, Robustness and Optimization* (CRC Press, 2009).

Dr. Datta 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, and he is currently serving as an Associate Editor of the *EURASIP Journal on Bioinformatics and Systems Biology*.

Edward R. Dougherty (M'05–SM'09) received the M.S. degree in computer science from Stevens Institute of Technology, Hoboken, NJ, the Ph.D. degree in mathematics from Rutgers University, Piscataway, 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 at Texas A&M University, College Station, where he holds the Robert M. Kennedy 26 Chair in Electrical Engineering 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 in Phoenix, AZ, and an Adjunct Professor in the Department of Pathology of the University of Texas M. D. Anderson Cancer Center in Houston, TX. He is 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 current 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 SPIE, has received the SPIE Presidents Award, and served as the editor of the SPIE/IS&T *Journal of Electronic Imaging*. At Texas A&M, he has received the Association of Former Students Distinguished Achievement Award in Research, been named Fellow of the Texas Engineering Experiment Station, and named Halliburton Professor of the Dwight Look College of Engineering.