BIOINFORMATICS

Vol. 20 no. 6 2004, pages 924–930
DOI: 10.1093/bioinformatics/bth008

External control in Markovian genetic regulatory networks: the imperfect information case

Aniruddha Datta[1], Ashish Choudhary[1], Michael L. Bittner[2] and Edward R. Dougherty[1,3,]*

[1]Department of Electrical Engineering, Texas A & M University, College Station, TX 77843-3128, USA, [2]TGEN, 400 North Fifth Street, Suite 1600, Phoenix, AZ 85004, USA and [3]Department of Pathology, University of Texas M. D. Anderson Cancer Center, 1515 Holcombe Blvd, Houston, TX 77030, USA

Received on 11 September 2003; accepted on 31 October 2003
Advance Access publication January 29, 2004

ABSTRACT

Probabilistic Boolean Networks, which form a subclass of Markovian Genetic Regulatory Networks, have been recently introduced as a rule-based paradigm for modeling gene regulatory networks. In an earlier paper, we introduced external control into Markovian Genetic Regulatory networks. More precisely, given a Markovian genetic regulatory network whose state transition probabilities depend on an external (control) variable, a Dynamic Programming-based procedure was developed by which one could choose the sequence of control actions that minimized a given performance index over a finite number of steps. The control algorithm of that paper, however, could be implemented only when one had perfect knowledge of the states of the Markov Chain. This paper presents a control strategy that can be implemented in the imperfect information case, and makes use of the available measurements which are assumed to be probabilistically related to the states of the underlying Markov Chain.

Contact: edward@ee.tamu.edu

1 INTRODUCTION

Probabilistic Boolean Networks (PBNs) have been recently proposed as a paradigm for studying gene regulatory networks (Shmulevich *et al*., 2002c). These networks which allow the incorporation of uncertainty into the inter-gene relationships, are essentially probabilistic generalizations of the standard Boolean networks introduced by Kaufmann (1993). Given a PBN, the transition from one state to the next takes place in accordance with certain transition probabilities. Indeed, as shown in Shmulevich *et al*. (2002c), the states of a PBN form a homogeneous Markov chain with finite state space. Thus the PBNs form a subclass of the general class of Markovian Genetic Regulatory Networks.

One of the objectives of PBN modeling is to use the PBN to design and evaluate different approaches for affecting the evolution of the gene activity profile (GAP) of the network. To date such intervention studies have been attempted using three different approaches: (i) resetting the state of the PBN, as necessary, to a more desirable initial state and letting the network evolve from there (Shmulevich *et al*., 2002b); (ii) changing the steady-state (long-run) behavior of the network by minimally altering its rule-based structure (Shmulevich *et al*., 2002a); and (iii) manipulating external (control) variables that alter the transition probabilities of the network and can, therefore, be used to desirably affect its dynamic evolution (Datta *et al*., 2003a).

In this paper, we continue with the line of research initiated in Datta *et al*. (2003a). The control problem in Datta *et al*. (2003a) was posed in terms of minimizing a certain performance index over a finite number of time steps. The problem formulation was based on the assumption that the state of the PBN was available for observation. Such an assumption may not hold in many real world situations. For instance, in a cancer treatment application, we may be able to track the expression status of only a limited number of genes and not necessarily all the ones appearing in the PBN of interest. This may be necessitated by cost, accessibility or other considerations. Clearly, in such a case, the control algorithm of Datta *et al*. (2003a) cannot be implemented.

The principal objective of this paper is to present a control strategy that can be used when perfect information about the state of the Markov Chain is not available. The control strategy in this case will be based on the feedback of available measurements which, though different from the exact internal states, are probabilistically related to the latter. The problem of control of Markov Chains in the presence of imperfect state information arises in many real world problems. Therefore, it is not surprising that solutions to this problem already exist in the current control literature (e.g. Bertsekas, 1976); indeed, the development here is motivated by some of this earlier work. However, we do believe that this paper represents the first attempt at adapting these results to the context of Genomic Control.

*To whom correspondence should be addressed.

The paper is organized as follows. In Section 2, we review the control problem for PBNs and recall the solution for the perfect state information case presented in Datta *et al*. (2003a). In Section 3, we consider the imperfect information case and present the solution for this problem. Section 4 presents the results of applying the control strategy of this paper to a biological example while Section 5 contains some concluding remarks. Due to space limitations, all mathematical derivations are omitted. The interested reader is referred to Datta *et al*. (2003b) for all the relevant details.

2 EXTERNAL CONTROL IN PROBABILISTIC BOOLEAN NETWORKS: THE PERFECT INFORMATION CASE

In this section, we provide a brief summary of the currently available results on external control in PBNs. We will only focus on those aspects that are critical to the development in this paper. For a detailed and complete exposition, the reader is referred to Datta *et al*. (2003a) and to Shmulevich *et al*. (2002b,c) for further details on PBNs.

A PBN is a formalism that has been developed for modeling the behavior of gene regulatory networks. In such a network, each gene can take on one of two binary values, zero or one. A 0 value for a gene corresponds to the case when that particular gene is not expressed and a 1 value indicates that the corresponding gene has been turned ON. Let us assume that we are attempting to model the relationship between 'n' genes. Suppose that the activity level of gene 'i' at time step 'k' is denoted by $x_i(k)$. Thus $x_i(k) = 0$ would indicate that at the k-th time step, the i-th gene is not expressed while $x_i(k) = 1$ would indicate that the corresponding gene is expressed. The overall expression levels of all the genes in the network at time step k is given by the row vector $x(k) = [x_1(k), x_2(k), \ldots, x_n(k)]$. This vector is sometimes referred to as the GAP or state vector of the network at time k.

The state vector $x(k)$ at any time step k is essentially an n-digit binary number whose decimal equivalent is given by

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

As $x(k)$ ranges from $000\cdots0$ to $111\cdots1$, $y(k)$ takes on all values from 0 to $2^n - 1$. Now to be completely consistent with the development in Shmulevich *et al*. (2002c), define

$$z(k) = 1 + y(k). \tag{2}$$

Then as $x(k)$ ranges from $00\cdots0$ to $11\cdots1$, $z(k)$ will take on all values from 1 to 2^n. Clearly, the map from $x(k)$ to $z(k)$ is one-to-one, onto and hence invertible. Thus instead of the binary representation $x(k)$ for the state vector, one could equivalently work with the decimal representation $z(k)$.

Suppose that the PBN with n genes has m control inputs u_1, $u_2, \ldots, u_m$, each of which can take on only the binary values 0 or 1. Then at any given time step k, the row vector $u(k) \triangleq [u_1(k), u_2(k), \ldots, u_m(k)]$ describes the complete status of all the control inputs. Clearly, $u(k)$ can take on all binary values from $[0, 0, \ldots, 0]$ to $[1, 1, \ldots, 1]$. One can equivalently represent the control input status using the decimal number

$$v(k) = 1 + \sum_{i=1}^{m} 2^{m-i} u_i(k). \tag{3}$$

As $u(k)$ takes on binary values from $[0, 0 \ldots, 0]$ to $[1, 1, \ldots, 1]$, the variable $v(k)$ ranges from 1 to 2^m. We can equivalently use $v(k)$ as an indicator of the complete control input status of the PBN at time step k.

As shown in Datta *et al*. (2003a), the one-step evolution of the probability distribution vector in the case of such a PBN with control inputs takes place according to the equation:

$$w(k+1) = w(k)A(v(k)), \tag{4}$$

where $w(k)$ is the 2^n-dimensional state probability distribution vector and $A(v(k))$ is the $2^n \times 2^n$ matrix of control-dependent transition probabilities.

Since the transition probability matrix here is a function of all the control inputs $u_1(k), u_2(k), \ldots, u_m(k)$, the evolution of the probability distribution vector of the PBN with control now depends not only on the initial distribution vector but also on the values of the control inputs at different time steps. Furthermore, intuitively it appears that it may be possible to make the states of the network evolve in a desirable fashion by appropriately choosing the control input at each time step.

These ideas were formalized in Datta *et al*. (2003a) to arrive at the following finite horizon optimization problem:[1]
Given an initial state z_0,

$$\min_{\mu_0, \mu_1, \ldots, \mu_{M-1}} E\left[\sum_{k=0}^{M-1} C_k(z_k, \mu_k(z_k)) + C_M(z_M)\right] \tag{5}$$

subject to

$$\Pr\{z_{k+1} = j | z_k = i, v_k\} = a_{ij}(v_k), \tag{6}$$

where

- $a_{ij}(v_k)$ is the i-th row, j-th column entry of the matrix $A(v_k)$;
- M represents the treatment/intervention window;
- $\mu_k : [1, 2, 3, \ldots, 2^n] \to [1, 2, 3, \ldots, 2^m]$, $k = 0, 1, 2, \ldots, M-1$ are functions mapping the state space into the control space;
- $C_k(z_k, v_k)$ is the one-step cost of applying the control v_k at state z_k;
- and $C_M(z_M)$ is the terminal cost associated with the state z_M.

[1] In the rest of this paper, we will be denoting $w(k)$, $z(k)$, $v(k)$ by w_k, z_k, v_k respectively, i.e. time dependence will be denoted by using a subscript. Since we will no longer focus attention on the individual elements of the vectors w_k, z_k, this will permit us to simplify the notation while avoiding any potential confusion.

As discussed in Datta *et al*. (2003a), the consideration of such an optimization problem could be naturally motivated in the context of cancer treatment applications where one must choose between a number of alternative treatments to be applied over a finite horizon of time. Once input from biologists/clinicians has been used to select an appropriate cost function and an appropriate treatment window, the control problem is essentially reduced to that of controlling a Markov Chain over a finite horizon.

The dynamic programming solution to (5), (6) is given by Bertsekas (1976), Datta *et al*. (2003a)

$$J_M(z_M) = C_M(z_M), \tag{7}$$

$$J_k(z_k) = \min_{v_k \in \{1,2,\ldots,2^m\}} \left[C_k(z_k, v_k) + \sum_{j=1}^{2^n} a_{z_k,j}(v_k) \cdot J_{k+1}(j) \right],$$
$$k = 0, 1, 2, \ldots, M-1. \tag{8}$$

Furthermore, if $v_k^* = \mu_k^*(z_k)$ minimizes the right-hand side of (8) for each z_k and k, the control law $\pi^* = \{\mu_0^*, \mu_1^*, \ldots, \mu_{N-1}^*\}$ is optimal.

3 EXTERNAL CONTROL IN PROBABILISTIC BOOLEAN NETWORKS: THE IMPERFECT INFORMATION CASE

The control law that emerges from the solution of the dynamic programming problem (7), (8) in the last section takes the form of a state feedback

$$v_k = \mu_k(z_k), \quad k = 0, 1, 2, \ldots, M-1. \tag{9}$$

When the state vector z_k of the PBN is not available for measurement, such a control law cannot be implemented. In that case, we will assume that when the PBN is in the state z_k, it emits q *measurable* outputs, each of which could take on the value 0 or 1. Thus the output status of the PBN at any time k could be captured by a q digit binary number or alternatively, by its decimal equivalent plus one, which we shall call θ_k. Clearly, as the outputs range over all possible binary values, θ_k takes on all values from 1 to 2^q.

The design of the optimal control in this case can make use of only the signals available to the controller. In other words, at time k, the controller tries to design the control input v_k using all the available signals $\theta_0, \theta_1, \ldots, \theta_k, v_0, v_1, \ldots, v_{k-1}$. Although the state z_k evolves according to (6) and is not available for measurement, we assume that the output θ_k at time k is probabilistically related to the state z_k at time k and the input v_{k-1} through the known conditional probability measure $\Pr_{\theta_k}(\cdot|z_k, v_{k-1})$ defined by

$$\Pr\{\theta_k = j | z_k = i, v_{k-1} = v\} = r_{ij}^v. \tag{10}$$

Let I_k denote the total information that is available for control at time k. Then clearly $I_k = [\theta_0, v_0, \theta_1, v_1, \ldots, v_{k-1}, \theta_k]^{\mathrm{T}}$. Furthermore, I_k can be generated recursively using the equation

$$I_{k+1} = [I_k^{\mathrm{T}}, v_k, \theta_{k+1}]^{\mathrm{T}}, \quad I_0 = \theta_0. \tag{11}$$

Since the state z_k is not available, it seems reasonable to replace the state feedback control (9) by the information feedback control

$$v_k = \mu_k(I_k), \quad k = 0, 1, 2, \ldots, M-1 \tag{12}$$

and search for the optimal μ_k over the space of all functions μ_k mapping the space of information vectors I_k into the control space $\{1, 2, 3, \ldots, 2^m\}$. Thus the counterpart to the optimization problem (5) for this case becomes (Bertsekas, 1976; Datta *et al*., 2003b):

$$\min_{\mu_0, \mu_1, \ldots, \mu_{M-1}} E_{\substack{z_0, d_0, d_1, \ldots, d_{M-1} \\ \theta_0, \theta_1, \ldots, \theta_{M-1}}} \left\{ \sum_{k=0}^{M-1} C_k(z_k, \mu_k(I_k), d_k) + C_M(z_M) \right\} \tag{13}$$

subject to

$$z_{k+1} = d_k, \tag{14}$$

$$\Pr\{d_k = j | z_k = i, v_k\} = a_{ij}(v_k), \tag{15}$$

$$I_{k+1} = [I_k^T, v_k, \theta_{k+1}]^{\mathrm{T}}, \quad I_0 = \theta_0. \tag{16}$$

The dynamic programming algorithm for the above problem is given by Bertsekas (1976), Datta *et al*. (2003b)

$$J_{M-1}(I_{M-1}) = \min_{v_{M-1} \in \{1,2,\ldots,2^m\}} \left\{ E_{z_{M-1}, d_{M-1}} [C_M(d_{M-1}) + C_{M-1}(z_{M-1}, v_{M-1}, d_{M-1}) | I_{M-1}, v_{M-1}] \right\} \tag{17}$$

$$J_k(I_k) = \min_{v_k \in \{1,2,\ldots,2^m\}} \left\{ E_{\theta_{k+1}, z_k, d_k} [\{C_k(z_k, v_k, d_k) + J_{k+1}([I_k^{\mathrm{T}}, v_k, \theta_{k+1}]^{\mathrm{T}})\} | I_k, v_k] \right\},$$
$$k = 0, 1, 2, \ldots, M-2 \tag{18}$$

and the optimal control input is obtained from the values minimizing the right-hand side of (17) and (18). Using this algorithm, we will ultimately arrive at $J_0(I_0) = J_0(\theta_0)$. The optimal cost J^* can be obtained by taking the expectation of this quantity with respect to θ_0, i.e.

$$J^* = E_{\theta_0}[J_0(\theta_0)]. \tag{19}$$

4 A REAL WORLD EXAMPLE BASED ON GENE EXPRESSION DATA

In this section, we apply the methodology of this paper to derive an optimal intervention strategy for a particular gene

regulatory network. The network chosen as an example of how control might be applied is one developed from data collected in a study of metastatic melanoma (Bittner *et al*., 2000). In this expression profiling study, the abundance of messenger RNA for the gene WNT5A was found to be a highly discriminating difference between cells with properties typically associated with high metastatic competence versus those with low metastatic competence. These findings were validated and expanded in a second study (Weeraratna *et al*., 2002). In this study, 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. A further finding of interest in the current study was that an intervention that blocked the WNT5A protein from activating its receptor, the use of an antibody that binds WNT5A protein, could substantially reduce WNT5A's ability to induce a metastatic phenotype. This of course suggests a study of control based on interventions that alter the contribution of the WNT5A gene's action to biological regulation, since the available data suggests that disruption of this influence could reduce the chance of a melanoma metastasizing, a desirable outcome.

The methods for choosing the genes involved in a small local network that includes the activity of the WNT5A gene and the rules of interaction have been described in Kim *et al*. (2002). As discussed in that paper, the WNT5A network was obtained by studying the predictive relationship between 587 genes. The expression status of each gene was quantized to one of three possible levels: -1 (down-regulated), 0 (unchanged) and 1 (up-regulated). Thus in this case, the gene activity profile at any time step is not a binary number but a ternary one. However, the PBN formulation and the associated control strategy can be developed exactly as described in Sections 2 and 3 with the only difference that now for an n-gene network, we will have 3^n states instead of the 2^n states encountered earlier. In this context, it is appropriate to point out that to apply the imperfect information control algorithm of this paper, it is not necessary to actually construct a PBN; all that is required are the transition probabilities between the different states under the different controls and the probabilities of the different observations conditioned on the current state and the previous control.

A network with 587 genes will have 3^{587} states which is an intractably large number to use either for modeling or for control. Consequently, the number of genes was narrowed down to the 10 most significant ones and the resulting multivariate relationship, using the best three-gene predictor for each gene, is shown in Figure 1. These relationships were developed using the COD (Coefficient of Determination) technique (Kim *et al*., 2000) applied to the gene expression patterns across 31 different conditions and prior biological knowledge. A detailed description of this is available in Kim *et al*. (2002).

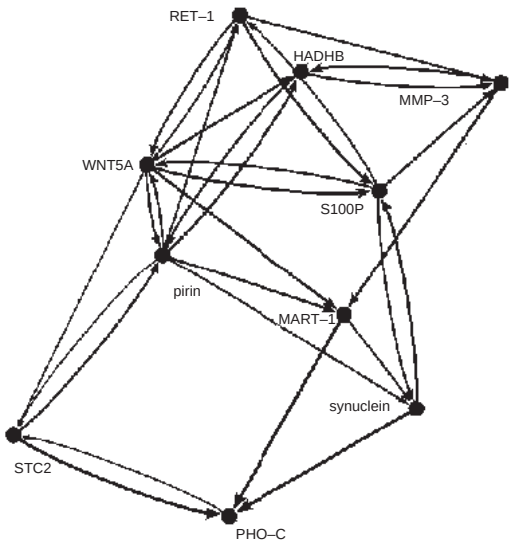

Fig. 1. Multivariate relationship between the genes of the 10-gene WNT5A Network (Kim *et al*., 2002).

The control objective for this 10-gene network is to externally down-regulate the WNT5A gene. The reason is that it is biologically known that WNT5A ceasing to be down-regulated is strongly predictive of the onset of metastasis. Controlling the 10-gene network, even for the perfect information case, would require us to design a dynamic programming based control algorithm for a system with 3^{10} $(=59\,049)$ states. Although there is nothing conceptually difficult about doing this, it was beyond the computational limits of our software, which we are currently in the process of improving. Moreover, parallel implementation is straightforward and could accomplish processing for significantly larger networks.

Accordingly, in Datta *et al*. (2003a), we further narrowed down the number of genes in the network to seven by using COD analysis on the 31 samples. Next, the $3^7 \times 3^7$ matrix of transition probabilities for the Markov Chain corresponding to the dynamic evolution of the GAP of the seven gene network was determined. Using this matrix of transition probabilities, the optimal control problem for the perfect information case was posed and solved in Datta *et al*. (2003a). The interested reader can consult that reference for all the necessary details and a discussion of the biological significance of the results of the perfect-information-based control action.

In this paper, we consider a 7-gene network which is a slight variation of the one considered in Datta *et al*. (2003a). Since implementing the imperfect information based control is computationally more intensive compared to the perfect information case, we developed a binary 7-gene network using COD analysis on the same experimental data. The resulting genes along with their multivariate relationship are shown in Figure 2. For each gene in this network, we determined their two best two-gene predictors and their corresponding

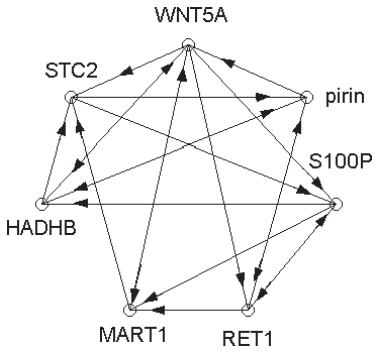

Fig. 2. Multivariate relationship between the genes of the 7-gene WNT5A network.

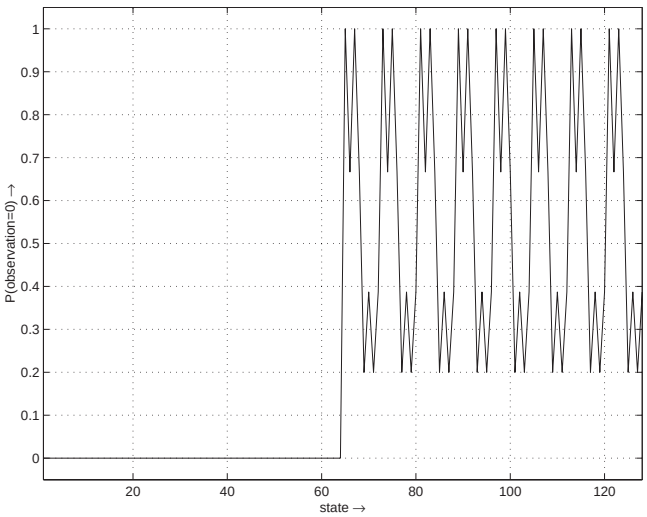

Fig. 3. Probability (observed variable = 0) versus current state.

CODs. Using the procedure discussed in (Shmulevich *et al.*, 2002c), the COD information for each of the predictors was then used to determine the $2^7 \times 2^7$ matrix of transition probabilities for the Markov Chain corresponding to the dynamic evolution of the GAP of the 7-gene network. The transition probability matrix $A(v(k))$, the probability distribution of the observations given the current state and the immediately prior control (r_{ij}^v), and the initial state probability distribution vector (P_0) together constitute the data needed for setting up the optimal control problem in the presence of imperfect state information. In our construction, the vector r_{ij}^v does not depend on the prior control input v and probabilistically relates the observation to the current state of the network. This relationship is shown in Figure 3 and it closely mimics the behavior of a gene MMP-3 which appears in the 10-gene network of Figure 1 but does not appear in the 7-gene network of Figure 2.

The optimal control problem can now be completely specified by choosing (i) the treatment/intervention window, (ii) the terminal penalty and (iii) the types of controls and the costs associated with them. For the treatment window, we arbitrarily chose a window of length 5, i.e. the control inputs would be applied only at time steps 0, 1, 2, 3 and 4. The terminal penalty at time step 5 was chosen as follows. Since our objective is to ensure that WNT5A is not up-regulated, we assigned a penalty of zero to all states for which WNT5A equals 0 and a penalty of 3 to all states for which WNT5A equals 1. Here the choice of the number 3 is somewhat arbitrary but it does reflect our attempt to numerically capture the biological notion that states where WNT5A equals 1 are less desirable than those where WNT5A equals 0.

We next discuss two possible types of control actions for various initial state probability distributions.

Case 1. WNT5A controlled directly: in this case, the control action at any given time step is to force WNT5A equal to 0, if necessary, and let the network evolve from there. Biologically such a control could be implemented by using a WNT5A inhibitory protein. In this case, the control variable is binary with 0 indicating that no WNT5A inhibitory protein is used while 1 indicates that such an intervention has been applied. The one-step cost of control is taken to be equal to the value of the control variable. Of course, whether at a given time step, such intervention takes place or not is decided by the solution to the resulting dynamic programming algorithm depending on the initial distribution P_0 and the subsequent total information vector I_k. Note that unlike the perfect information scenario considered in Datta *et al.* (2003a), we are now not in a position to determine if forcible alteration of the state takes place or not. Consequently, it is reasonable to expect that WNT5A inhibition may be used, even when not absolutely necessary, thereby contributing to a possible increase in the total optimal expected cost, compared to the perfect information case.

Next we used (17) and (18) recursively to calculate the optimal controls for certain given initial state probability distributions. The net result, in each case, was a tree of control actions corresponding to each control action and subsequent observation.

Starting with P_{data}, the distribution of states in the 31 point data set, we found the optimal expected cost based on imperfect information to be 0.4079. The corresponding optimal cost using full state observation as in Datta *et al.* (2003a) was found to be 0.3226. The expected cost incurred by not using any control was 0.9677. We computed these quantities for a few different cases of initial state distributions. The relevant quantities are tabulated in Table 1.

We also calculated the optimal expected costs when the initial state is deterministic. These values for all the 128 possible initial states are shown in Figure 4. Note that, as expected, the optimal cost for control with imperfect information is higher than that for control with perfect state information. The cost function, however, is a somewhat subjective quantity chosen by us to mathematically capture the underlying

Table 1. Expected costs for various initial state distributions

Initial distribution	Control using Observation	Control using Full state	No control
$P_{\text{Sample–data}}$	0.4079	0.3226	0.9677
$\left[\frac{1}{128}, \frac{1}{128}, \ldots\right]$	0.7068	0.3395	0.9990
$\left[0, \frac{1}{64}, 0, \frac{1}{64}, \ldots\right]$	0.7296	0.3395	0.9990
$\left[\frac{1}{64}, 0, \frac{1}{64}, 0, \ldots\right]$	0.5692	0.3395	0.9990

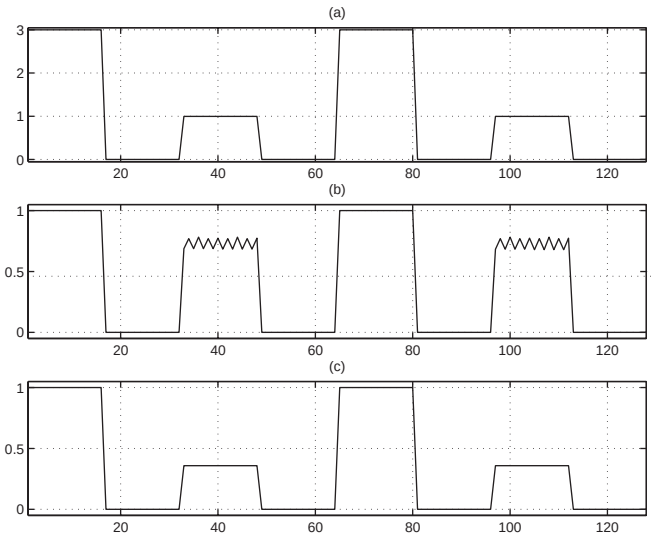

Fig. 4. Optimal expected cost versus initial states **(a)** uncontrolled, **(b)** control using imperfect information, **(c)** control using full state information.

biological objective. A more natural way to look at the performance of the control scheme would be to examine the probability of WNT5A being equal to 0 at the final time step, i.e. at $k = 5$. This quantity was computed for each (determinstic) initial state for both the uncontrolled and imperfect-information-based controlled cases. These plots are shown in Figure 5. From this figure, it is clear that the control strategy for each initial state is increasing the probability for WNT5A equal to zero at the terminal time point relative to the corresponding probability in the uncontrolled case. This is, indeed, a desirable outcome achieved by using control.

Case 2. WNT5A controlled through pirin: in this case, the control objective is the same as in Case 1, namely to keep WNT5A at 0. The only difference is that this time, we use another gene, pirin, to achieve this control. The treatment window and the terminal penalties are kept exactly the same as before. The control action consists of either using a pirin inhibitor (corresponding to a control input of 1) or not employing such an inhibitor (corresponding to a control input of 0). The one-step cost of control is taken to be equal to the value

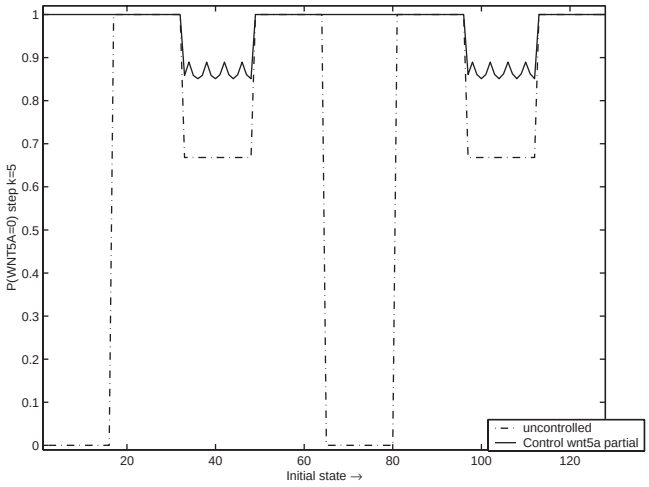

Fig. 5. Probability of WNT5A $= 0$ at the terminal time point versus the initial state for the uncontrolled and imperfect-information-based controlled cases.

of the control variable. As before, at any step, whether such intervention takes place or not is decided by the solution to the resulting dynamic programming algorithm. Having chosen these design parameters, we implemented the algorithm with pirin as the control.

We found that using pirin as a control is totally ineffective. The expected cost, with pirin as the control, was found to be the same as the one obtained in Table 1 with no control. Even with full state feedback we still found that pirin was as ineffective as before (data not shown). This is in stark contrast to our results in Datta *et al*. (2003a) where we demonstrated the feasibility of doing full state feedback control of WNT5A through pirin. It is possible that going from a ternary setup in Datta *et al*. (2003a) to the binary setup here may have drastically reduced our ability to control WNT5A through pirin. This suggests that the standard control theoretic notions of controllability and observability (Kalman, 1962) may have to be revisited in the context of Genetic Regulatory Networks to enable us to decide which genes can be used as effective controls and which ones can be used as meaningful observations. Equally important is the need to introduce causality into Genetic Regulatory Networks, an aspect that the current COD formalism does not capture. These topics, however, are beyond the scope of the current paper, and will have to be investigated in the future.

5 CONCLUDING REMARKS

In this paper, we have extended our earlier results on external control in Markovian genetic regulatory networks to the case where perfect information about the state of the network is not available. In such a situation, the optimal control must be designed based on the available measurements, which are

assumed to be probabilistically related to the state of the genetic regulatory network.

The optimal control results presented in this paper assume known transition probabilities and pertain to a finite horizon problem of known length. Their extension to the situation where the transition probabilities and the horizon length are unknown is a topic for further investigation.

ACKNOWLEDGEMENTS

This work was supported in part by the National Cancer Institute under Grant CA90301, by the National Human Genome Research Institute, by the National Science Foundation under Grant ECS-9903488 and by the Texas Advanced Technology Program under Grant No. 000512-0099-1999.

REFERENCES

Bertsekas,D.P. (1976) *Dynamic Programming and Stochastic Control*. Academic Press, New York.

Bittner,M., Meltzer,P., Chen,Y., Jiang,Y., Seftor,E., Hendrix,M., Radmacher,M., Simon,R., Yakhini,Z., Ben-Dor,A. *et al*. (2000) Molecular classification of cutaneous malignant melanoma by gene expression profiling. *Nature*, **406**, 536–540.

Datta,A., Choudhary,A., Bittner,M.L. and Dougherty,E.R. (2003a) External control in Markovian genetic regulatory networks. *Mach. Learning*, **52**, 169–191.

Datta,A., Choudhary,A., Bittner,M.L. and Dougherty,E.R. (2003b) External control in Markovian genetic regulatory networks: the imperfect information case. Report No. 03-09-01, Genomic Signal Processing Laboratory, Texas A and M University, College Station, Texas.

Kalman,R.E. (1962) Canonical structure of linear dynamical systems. *Proc. Natl Acad. Sci., USA*, 596–600.

Kauffman,S.A. (1993) *The Origins of Order: Self-organization and Selection in Evolution*. Oxford University Press, New York.

Kim,S., Dougherty,E.R., Bittner,M.L., Chen,Y., Sivakumar,K., Meltzer,P. and Trent,J.M. (2000) A general framework for the analysis of multivariate gene interaction via expression arrays. *Biomed. Opt.*, **4**, 411–424.

Kim,S., Li,H., Dougherty,E.R., Cao,N., Chen,Y., Bittner,M. and Suh,E. (2002) Can Markov chain models mimic biological regulation? *J. Biol. Sys.*, **10**, 437–458.

Shmulevich,I., Dougherty,E.R. and Zhang,W. (2002a) Control of stationary behavior in probabilistic Boolean networks by means of structural intervention. *Biol. Syst.*, **10**, 431–446.

Shmulevich,I., Dougherty,E.R. and Zhang,W. (2002b) Gene perturbation and intervention in probabilistic Boolean networks. *Bioinformatics*, **18**, 1319–1331.

Shmulevich,I., Dougherty,E.R., Kim,S. and Zhang,W. (2002c) Probabilistic Boolean networks: a rule-based uncertainty model for gene regulatory networks. *Bioinformatics*, **18**, 261–274.

Weeraratna,A.T., Jiang,Y., Hostetter,G., Rosenblatt,K., Duray,P., Bittner,M. and Trent,J.M. (2002) WNT5A signalling directly affects cell motility and invasion of metastatic melanoma. *Cancer Cell*, **1**, 279–288.