Journal of Theoretical Biology 261 (2009) 540–547

Contents lists available at ScienceDirect

Journal of Theoretical Biology

journal homepage: www.elsevier.com/locate/yjtbi

Sampling-rate-dependent probabilistic Boolean networks

Golnaz Vahedi [a,*], Babak Faryabi [a], Jean-Francois Chamberland [a], Aniruddha Datta [a], Edward R. Dougherty [a,b]

[a] Department of Electrical and Computer Engineering, Texas A&M University, College Station, TX 77843, USA
[b] Translational Genomics Research Institute, 400 North Fifth Street, Suite 1600, Phoenix, AZ 85004, USA

ARTICLE INFO

Article history:
Received 23 December 2008
Received in revised form
18 June 2009
Accepted 17 August 2009
Available online 28 August 2009

Keywords:
Probabilistic Boolean network
Gene regulatory network
Optimal control
Intervention

ABSTRACT

External control of a genetic regulatory network is used for the purpose of avoiding undesirable states, such as those associated with a disease. To date, intervention has mainly focused on the external control of probabilistic Boolean networks via the associated discrete-time discrete-space Markov processes. Implementation of an intervention policy derived for probabilistic Boolean networks requires nearly continuous observation of the underlying biological system since precise application requires the observation of all transitions. In medical applications, as in many engineering problems, the process is sampled at discrete time intervals and a decision to intervene or not must be made at each sample point. In this work, sampling-rate-dependent probabilistic Boolean network is proposed as an extension of probabilistic Boolean network. The proposed framework is capable of capturing the sampling rate of the underlying system.

1. Introduction

A major concern of translational genomics is to use the knowledge of gene regulation to design therapeutic strategies. In biology, there are numerous examples where the (in)activation of one gene or protein can lead to a certain cellular functional state or phenotype. For instance, in a stable cancer cell line, the reproductive cell cycle is repeated and cancerous cells proliferate with time in the absence of intervention. One can use the p53 gene if the intervention goal is to push the cells into apoptosis, or programmed cell death, to arrest the cell cycle. The p53 gene is the most well-known tumor suppressor gene, encoding a protein that regulates the expression of several genes such as Bax and Fas/APO1, which function is to promote apoptosis (Miyashita and Reed, 1995; Owen-Schaub et al., 1995). In cultured cells, extensive experimental results indicate that when p53 is activated, e.g. in response to radiation, it leads to cell growth inhibition or cell death (El-Deiry et al., 1993). The p53 gene is also used in gene therapy, where the target gene (p53 in this case) is cloned into a viral vector. The modified virus serves as a vehicle to transport the p53 gene into tumor cells to generate intervention (Swisher et al., 1999; Bouvet et al., 1998). As this and many other examples suggest, it is prudent to use gene regulatory models to design therapeutic interventions that expediently modify the cell's dynamics via external signals. These system-based intervention methods can be useful in identifying potential drug targets and discovering treatments to disrupt or mitigate the aberrant gene functions contributing to the pathology of a disease.

Gene network modeling facilitates this effort by producing dynamical systems to serve as the mathematical basis for the derivation of optimal intervention strategies over time. Various approaches have been proposed for modeling gene regulatory networks. The discrete formalism of rule-based regulatory networks is plausible for many classes of biological systems. Strong evidences suggest that the input–output relations of regulatory interactions are signomial and can be well approximated by step functions (Huang, 1999; Thomas, 1979). One popular approach is based on Boolean networks (Kauffman, 1993). The Boolean network framework possesses certain inherent limitations. A class of continuous-time binary networks has recently been introduced that include traditional Boolean networks as a special case but are not constrained to exhibit periodic responses to constant inputs (Oktem et al., 2003). Synchronous abstraction in Boolean networks is used under implicit assumption that asynchronous updating will not unduly alter the properties of a system central to the application of interest (Harvey and Bossomaier, 1997). Various potential issues with synchronous networks have been noted. For instance, synchronous abstraction may produce spurious attractors in rule-based networks (Deng et al., 2007). In the same vein, deviation from synchronous updating modifies the attractor structure of Boolean networks (Greil and Drossel, 2005) and can change their long-run behavior (Gershenson, 2002).

To date, intervention has mainly focused on the external control of probabilistic Boolean networks (PBNs) via the associated

* Corresponding author. Tel.: +1 979 862 8896.
E-mail addresses: golnaz_vahedi@ieee.org (G. Vahedi), bfariabi@ieee.org (B. Faryabi), chmbrlnd@tamu.edu (J.-F. Chamberland), datta@ece.tamu.edu (A. Datta), edward@ece.tamu.edu (E.R. Dougherty).

0022-5193/$ - see front matter © 2009 Elsevier Ltd. All rights reserved.
doi:10.1016/j.jtbi.2009.08.026

discrete-time discrete-space Markov processes (Pal et al., 2006). Given the accuracy of the model, there are two practical impediments to PBN-based intervention, both related to temporal issues. One of these concerns the lack of information regarding the sojourn time in any given state and the other concerns the practical problem of sampling. The first issue, the effect of sojourn time on the control, has been studied in Faryabi et al. (2008). In this work, we focus on the effect of discrete sampling.

While the physical evolution of the biological gene network occurs over continuous time, the PBN records only state transitions and contains no information on the time between transitions. The PBN model inherits this property from the original Boolean model, from which it was generalized (Kauffman, 1993). Hence, the problem can be explained in the framework of the Boolean model. Fig. 1 shows the directed graph of a 3-gene Boolean network, where each 3-gene state corresponds to a gene-activity profile (GAP). Fig. 2 shows two continuous-time realizations that are equivalent from the perspective of the model of Fig. 1. In both Fig. 2(a) and (b), the initial state is "100". We observe the evolution "100"→"010"→"001", at which point there are no other changes because "001" is an attractor of the network. An attractor of a Boolean network is a state in which the network will eventually enter and repeatedly cycle forever. While equivalent from the perspective of the Boolean model, from the perspective of continuous time, the realizations of Fig. 2(a) and (b) are not the same. For instance, in the second realization, the sojourn time in state "010" is much longer than in the first realization. If we are only interested in tracking the transitions, this may be of no concern. On the other hand, suppose we are considering intervention and penalizing undesirable states. Then, if "010" is an undesirable state, the penalty should be greater in the second realization; that is, the penalty needs to consider the sojourn time in a state. This problem has been addressed in the framework of asynchronous PBNs by considering the process to be defined over continuous time and treating it as a semi-Markov process (Faryabi et al., 2008).

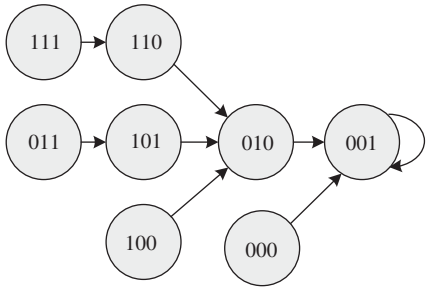

Fig. 1. Presentation of a directed graph for an arbitrary 3-gene Boolean network.

Whether one considers the original synchronous PBNs or asynchronous PBNs, implementation of the intervention policy requires nearly continuous observation because precise application requires the observation of all transitions. However, this is not generally the case in medical applications; rather, as with many engineering problems, the process is sampled at discrete time intervals and a decision to intervene or not must be made at each sample point. Since the process is not observed outside the sample points, it is impossible to know if, or how many, transitions have taken place between consecutive sample points.

In Fig. 2, the discrete-time process $\{Y_n, n \geq 0\}$ given by $Y_n = Z_{t_n}$ is called the jump chain of the continuous-time process $\{Z_t, t \geq 0\}$. Both synchronous and asynchronous PBNs deal with the jump chain under the assumption that the jumps (i.e. $t_0, t_1, \ldots$) are observed. The jump chains corresponding to realizations of Fig. 2(a) and (b) are equivalent. Fig. 2 also shows the sampled processes corresponding to each realization. The sampled process corresponding to Fig. 2(a) is "100"→"100"→"001"→"001"→"001"; for Fig. 2(b), it is "010"→"010"→"001"→"001"→"001". On account of sampling, "010" is missed in Fig. 2(a) and "100" is missed in Fig. 2(b). Whereas in a standard Boolean network self-transitions only occur for singleton attractors, the sampled process has self-transitions. Moreover, the estimated sojourn time is implicitly contained in the sampled process on account of these self-transitions. As with any sampling procedure, the sampling rate is crucial. The faster the rate, the less transitions will be missed and the more accurate will be the sojourn time estimates; the slower the rate, the more samples will be missed and the less accurate will be the sojourn time estimates. In any event, in the presence of sampling, neither the synchronous or asynchronous PBN models will adequately reflect the dynamics of the network from the perspective of the decision process required for intervention. In this paper, we propose a framework for gene regulatory networks, a sampling-rate-dependent PBN (SRD-PBN), that is capable of incorporating the sampling rate of the temporal profile. Below, we mathematically define SRD-PBNs and expose a methodology to obtain optimal intervention strategies for such systems. To set the stage, we first provide necessary definitions in Section 2. We introduce SRD-PBNs in Section 3. In Section 4, we derive an optimal policy for SRD-PBNs with various properties for synthetic networks.

2. Background

2.1. Probabilistic Boolean networks

A probabilistic Boolean network (PBN) (Shmulevich et al., 2002) consists of a sequence $V = \{x_i\}_{i=1}^{n}$ of n nodes with $x_i \in \{0, \ldots, d-1\}$, together with a sequence $\{\mathbf{f}_c\}_{c=1}^{k}$ of vector-

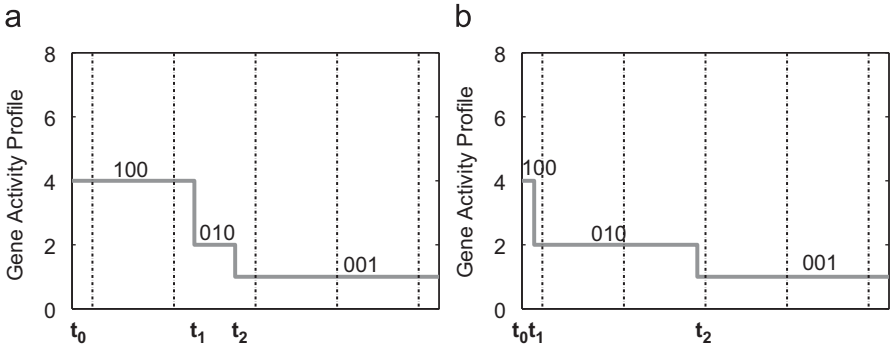

Fig. 2. Two examples of temporal gene activity profiles (GAP) for the oriented graph of Fig. 1. The dash-dot vertical lines represent the sampling times.

valued *network functions*. In the framework of gene regulation, each element x_i represents the expression level of a gene. It is common to mix the terminology by referring to x_i as the i th gene. The state vector $\mathbf{x}(t)=(x_1(t),\ldots,x_n(t))$ is called the *gene-activity profile* (GAP) at time t. Each network function $\mathbf{f}_c=(f_{c1},\ldots,f_{cn})$ determines a constituent network (*context*) of the regulatory network. The function $f_{ci}:\{0,\ldots,d-1\}^n\to\{0,\ldots,d-1\}$ is the predictor of gene i, whenever network c is selected. At each updating epoch a decision is made whether to switch the constituent network. This decision is based on a binary random variable ξ with $P(\xi=1)=q$. If $\xi=0$, then the network is not switched, the model behaves like a fixed network and the values of all genes are synchronously updated according to the current constituent network. If $\xi=1$, then a constituent network is randomly selected from among all constituent networks, including the current one, according to the selection probability distribution $\{p_c\}_{c=1}^{k}$ and, after selecting $\mathbf{f}_c$, the values of all genes are updated accordingly. If $q=1$, so that a switch is permitted at every time point, the network is said to be *instantaneously random*; if $q<1$, then the PBN will remain in a constituent network so long as ξ remains equal to 0, and the PBN is said to be *context-sensitive*.

Two quantization levels have thus far been used in practice. If $d=2$ (binary), then the constituent networks are Boolean networks with 0 or 1 meaning OFF or ON, respectively. The case $d=3$ (ternary) arises when we consider a gene to be 0 (down-regulated), 2 (up-regulated), and 1 (invariant). This situation commonly occurs with cDNA microarrays, where a ratio is taken between the expression values on the test channel (red) and the base channel (green). In this paper, we will develop the methodology for $d=2$, so that gene values are either 0 or 1; however, the methodology is applicable to any finite number of levels. For binary PBNs there is a natural bijection between the GAP $\mathbf{x}(t)$ and its integer representation, $x(t)$, which takes values in $\mathcal{W}=\{0,1,\ldots,2^n-1\}$. We consider a PBN with perturbation, meaning that there is a binary random vector $\gamma=(\gamma_1,\gamma_2,\ldots,\gamma_n)$, independent of ξ, such that $P(\gamma_i=1)=p$, $\gamma_1,\gamma_2,\ldots,\gamma_n$ are independent. If $\gamma=\mathbf{0}$ the network transitions according to the network function, and if $\gamma\neq\mathbf{0}$ the value x_i flips if and only if $\gamma_i=1$.

The dynamic behavior of a context-sensitive PBN can be modeled by a Markov chain whose states consist of (context, GAP) ordered pairs. Let $\mathbf{P}=(p_{ij};i,j\in\mathcal{S})$ denote the transition probability matrix of the Markov chain, where the state-space is

$$\mathcal{S}=\{(c,x):c\in\{1,\ldots,k\}\wedge x\in\mathcal{W}\}.$$

The time-evolution of the network can be modeled by a stationary discrete-time equation

$$z(t+1)=f(z(t),w(t))\quad\text{for } t=0,1,\ldots,$$

where state $z(t)\in\mathcal{S}$. The disturbance $w(t)$ is the manifestation of uncertainties, due to either network switching or a change in gene-activity profile resulting from a random gene perturbation. Gene perturbation insures that all states in the Markov chain communicate with one another. Hence, the finite-state Markov chain has a unique steady-state distribution.

2.2. Optimal intervention in Markovian gene regulatory networks

In the following, we describe how to devise an optimal control policy for a Markovian gene regulatory network such as context-sensitive PBN or sampling-rate-dependent PBN. Let $\mathbf{P}=(p_{ij};i,j\in\mathcal{S})$ denote the transition probability matrix of the Markov chain corresponding to the Markovian gene regulatory network. In the presence of external control, we suppose that there exists a binary control input, $u(t)\in\mathcal{C}=\{0,1\}$. A control $u(t)$, which can take values 0 or 1 at each updating epoch t, specifies the action on the control gene. Treatment alters the status of the control gene, which can be selected among all genes in the network. If the control at updating epoch t is on, $u(t)=1$, then the state of the control gene is toggled; if $u(t)=0$, then the state of the control gene remains unchanged. In the presence of external control, the system evolution is represented by a stationary discrete-time equation

$$z(t+1)=f(z(t),u(t),w(t))\quad\text{for } t=0,1,\ldots,$$

where state $z(t)$ is an element of the state-space $\mathcal{S}$; and $w(t)$ is the manifestation of uncertainties in the model. The probability of transitioning from state i to state j under control u is denoted by $p_{ij}(u)$, where $i,j\in\mathcal{S}$.

The problem of optimal intervention for a Markovian gene regulatory network is formulated as an optimal stochastic control problem. A cost-per-stage, $r(i,u,j)$, is associated to each intervention in the system. In general, a cost-per-stage may depend on the origin state i, the successor state j, and the control input u. We assume that the cost-per-stage is stationary and bounded for all i, j in $\mathcal{S}$, and u in $\mathcal{C}=\{0,1\}$. We define the expected immediate cost in state i, when control u is selected, by

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

We consider the discounted formulation of the expected total cost. The discounting factor, $\alpha\in(0,1)$, ensures convergence of the expected total cost over the long-run (Bertsekas, 2005). In the case of cancer therapy, the discounting factor emphasizes that obtaining treatment at an earlier stage is favored over later stages. The expected total discounted cost, given a policy π and an initial state i, is denoted by

$$J_\pi(i)=\lim_{N\to\infty}E\left\{\sum_{t=0}^{N-1}\alpha^t r(z(t),\mu_t(z(t)),z(t+1))|z(0)=i\right\},\tag{1}$$

where $z(t),\ i\in\mathcal{S}$. A policy $\pi=\{\mu_0,\mu_1,\ldots\}$ is a sequence of decision rules $\mu_t:\mathcal{S}\to\mathcal{C}$, for each time step t. The vector $\mathbf{J}_\pi$ of the expected total costs is called the value function. In a stochastic control problem, we seek an intervention strategy π^* among all the admissible intervention strategies Π that minimizes the value function for each state i, i.e.

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

A stationary intervention strategy is an admissible intervention strategy of the form $\pi=\{\mu,\mu,\ldots\}$. It is known that an optimal intervention strategy exists for the stochastic control problem presented in (1). Furthermore, this optimal policy is stationary. The optimal cost function J^* satisfies

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

where M represents the cardinality of $\mathcal{S}$. This cost function J^* is the unique solution to (3) within the class of bounded functions. This equation is known as the *Bellman optimality equation*. The optimal control policy attains the minimum in the right-hand side of the Bellman optimality equation for all i. Proofs of these statements along with a more complete expansion of stochastic control can be found in Bertsekas (2005). Standard dynamic programming algorithms can be employed to find the fixed-point of the Bellman optimality equation.

2.3. Continuous-time Markov chain

Consider a continuous-time discrete-space stochastic process $\{Z(t),\ t\geq 0\}$ taking on values in the set of nonnegative integers $\mathcal{Z}^+$. In analogy with a discrete-time Markov chain, we say that the process $\{Z(t),\ t\geq 0\}$ is a continuous-time Markov chain if $\forall s,t\geq 0$,

and nonnegative integers $i,j,k \in \mathcal{Z}^+$, $0 \le v < s$,

$$\Pr\{Z(t+s)=j|Z(s)=i, Z(v)=k, 0 \le v < s\} = \Pr\{Z(t+s)=j|Z(s)=i\}.$$

In other words, a continuous-time Markov chain is a stochastic process with the Markovian property. This means that the conditional distribution of the future state at time $t+s$, given the present state at time s and all the preceding states, depends only on the present state and is conditionally independent of the states prior to the current state. The past given the present does not provide more information about the future.

If we let τ_i denote the amount of time that the process stays in state i before making a transition into a different state, then the Markov property implies

$$\Pr\{\tau_i > s+t|\tau_i > s\} = \Pr\{\tau_i > t\}, \quad \forall s,t \ge 0.$$

The random variable τ_i is memoryless and must therefore be exponentially distributed. In general, a continuous-time Markov chain is defined by a $\mathbf{Q}$-matrix. A $\mathbf{Q}$-matrix on $\mathcal{Z}^+$ is a matrix $Q=(q_{ij}; i,j \in \mathcal{Z}^+)$ satisfying the following conditions (Norris, 1998):

$$\begin{aligned} &\text{(i)} \quad 0 \le -q_{ii} < \infty, \quad \forall i;\\ &\text{(ii)} \quad q_{ij} \ge 0, \quad \forall i \ne j;\\ &\text{(iii)} \quad \sum_{j \in I} q_{ij} = 0, \quad \forall i. \end{aligned} \tag{4}$$

In the above, q_{ij} is the rate of transitioning from i to j and $q_i = \sum_{i \ne j} q_{ij}$ is the rate of leaving state i. It is known that a matrix Q is a $\mathbf{Q}$-matrix on $\mathcal{Z}^+$ if and only if $\mathbf{P}(t)=e^{Qt}$ is a stochastic matrix, $\forall t \ge 0$ (Norris, 1998). In particular, the transition probability from i to j after t unit of time, the (i,j) element of $\mathbf{P}(t)$, is given by

$$\Pr(X_t = j|X_0 = i) = p_{ij}^t = [e^{Qt}]_{ij}.$$

3. Sampling-rate-dependent probabilistic Boolean networks

A context-sensitive PBN disregards the information about the sojourn time in states present in temporal data. From another point of view, a context-sensitive PBN models the jump chain corresponding to the continuous-time process of interest. This means that in an arbitrary temporal profile such as Fig. 3, the observer can only apply intervention at instants $t_0, t_1, \ldots$. However, in patient treatment, it is not known in advance when a transition (i.e. a jump) will occur. As such, a model based on applying treatment when a transition occurs may not conform with the reality and limitations of patient treatment. Time samples and state changes are unlikely to coincide perfectly and an intervention strategy must focus on the former not the latter.

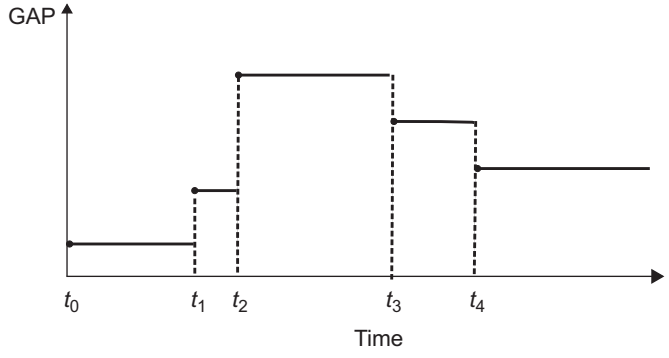

Fig. 3. An example of temporal gene activity profiles.

Our objective in this work is to propose a discrete-time discrete-space model based on context-sensitive PBNs such that (i) it can embody sojourn time of states into the network dynamics, (ii) it allows us to incorporate the sampling rate into the network's dynamics. A transition probability matrix must be derived for the state-space of an SRD-PBN under specific assumptions. Similar to other Markovian models, the transition probability matrix derived for an SRD-PBN is sufficient to describe its dynamics. The task of finding the most effective intervention strategy can then be formulated as a sequential decision making problem via the associated transition probability matrix.

Let us first briefly introduce the underlying structure of an SRD-PBN. The states of the SRD-PBN take values in $\mathcal{S}$, as we defined for a context-sensitive PBN. Logical rules of different contexts determine the probability of jumps among GAPs. To coarsely capture the rate of change in the underlying biological system, the proposed framework requires two parameters, which are either known a priori or can be estimated from temporal data. These two parameters are the maximum rate of change among GAPs and the maximum rate of change among contexts. The rate of change between any two states, i.e. the average number of transitions between these two states in every unit of time, depends on the probability of jumps between these two states, the sampling period, the maximum rate of change among GAPs, and the maximum rate of change among contexts. Employing these parameters, we construct a $\mathbf{Q}$-matrix on the state-space $\mathcal{S}$ of the SRD-PBN. This matrix is the generator of a continuous-time Markov chain. We are interested in the state of the continuous process only at discrete observation instants. The memoryless property of the continuous-time Markov chain allows us to model the dynamics of the sampled process as a discrete-time Markov chain. The transition probability matrix of this Markov chain is the transition probability matrix of the SRD-PBN. Below, we define the SRD-PBN in more details.

Given Boolean functions of context c, the probability of jumping from state (c,x) to state (c,x') is

$$P_{(c,x),(c,x')} = p^{D(x,x')}(1-p)^{n-D(x,x')} + (1-p)^n \mathbf{1}(\boldsymbol{f}_c(x) = x'), \tag{5}$$

where p is the perturbation probability in the Boolean network. The Hamming distance between GAPs x and x' is denoted by $D(x,x')$. We use $\mathbf{1}()$ to denote the indicator function. The successor state of GAP x according to the Boolean functions of context c is denoted by $\boldsymbol{f}_c(x)$. The first part of (5) corresponds to the transition probability due to gene perturbation. The probability of transitioning between GAPs x and x' based on the selected context $\boldsymbol{f}_c$ is presented as the second part of (5).

To include timing in our proposed model, given (5), we introduce matrix Q which shows the rate of transitions among states in $\mathcal{S}$. We denote the maximum rate of change among GAPs by λ and the maximum rate of change among contexts by γ. In practice, λ can be estimated from temporal data. Knowledge of the ratio λ/γ, provided by experiments, would determine the value of γ. Matrix Q is the generator of a continuous-time Markov chain. Let $Q=(q_{(c,x),(c',x')}; c,c' \in \{1,\ldots,k\}, x,x' \in \mathcal{W})$ denote the $\mathbf{Q}$-matrix of the continuous-time Markov chain $\{Z(t), t \ge 0\}$ whose state-space is $\mathcal{S}$. Elements of the $\mathbf{Q}$-matrix show the rate of change among states and can be computed in the following manner.

At any updating epoch, there are two independent processes: (i) a process that updates the GAP in the current context, (ii) a process that updates the context. There are null probabilities for both processes to occur at the same time. For the first process, we can compute the rate of change among GAPs x and x' in context c as the product of λ, the maximum rate of change between GAPs, times the probability to jump from GAP x to x', i.e. $P_{(c,x),(c,x')}$. For the second process, we can compute the rate of change between contexts c and c' as the product of γ, the maximum rate of change

between contexts, times the selection probability of context c'. Furthermore, in order to have a valid **Q**-matrix (4), all diagonal elements of Q should be defined such that the sum of elements in each row is zero. Thus, the rate of change between any two states (c,x) and (c',x') in $\mathcal{S}$ is defined as

$$q_{(c,x),(c',x')} = \begin{cases} \lambda P_{(c,x),(c,x')} & \text{if } c = c' \text{ and } x \neq x', \\ \gamma p_{c'} & \text{if } c \neq c' \text{ and } x = x', \\ 0 & \text{if } c \neq c' \text{ and } x \neq x', \\ -\sum_{X \neq x} \sum_{C \neq c} q_{(c,x),(C,X)} & \text{if } c = c' \text{ and } x = x', \end{cases} \quad (6)$$

where $p_{c'}$ is the selection probability of context c'.

We define p_{ij}^t to be the probability that the continuous-time process $\{Z_t, t \geq 0\}$ associated to the SRD-PBN makes a transition from current state i to successor state j after t units of time. Using this notation, p_{ij}^t corresponds to (i,j) entry in matrix $\mathbf{P}(t)$, where $\mathbf{P}(t) = e^{Qt}$.

From the intervention perspective, we are interested in the dynamical behavior of the SRD-PBN at discrete observation instants, i.e. every T units of time. Such a discrete-time model yields more information for the decision making process. Employing the memoryless property of the continuous-time Markov chain, we obtain a discrete-time Markov chain by taking samples from the continuous-time Markov chain at every T units of time. This discrete-time model describes the dynamics of the SRD-PBN. For a given sampling period T, the transition probability matrix that expresses the dynamics of the SRD-PBN is computed as

$$\mathbf{P}(T) = e^{QT}, \quad (7)$$

where elements of Q are defined in (6). We note that the transition probability matrix associated to the SRD-PBN is a function of the sampling period T. Optimal intervention strategies, as described in Section 2, can then be derived for this SRD-PBN using the corresponding transition probability matrix.

Example. To illustrate the details of an SRD-PBN, we produce a simple 3-gene, 2-context example. Given the logical rules of each constituent Boolean network, one can draw the directed graphs corresponding to each Boolean network. Fig. 4 shows the directed graphs of the constituent Boolean networks in our simple example. The transition probability matrix corresponding to the context-sensitive PBN constructed based on these Boolean networks is shown in Fig. 5(a). This figure is a heat-map where larger transition probabilities are brighter and the smaller ones are darker. This transition probability matrix is computed

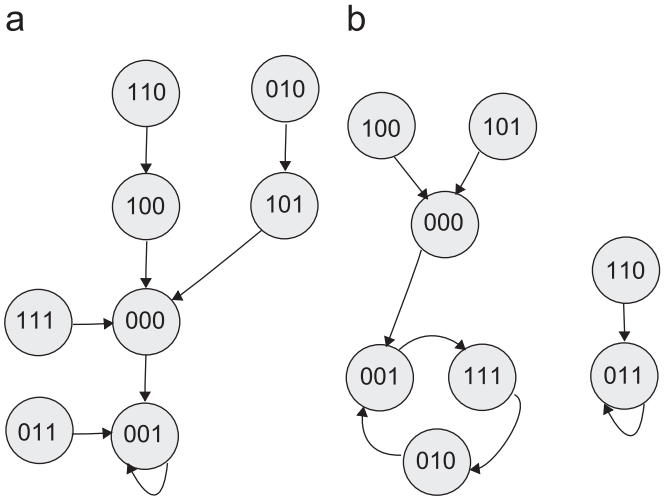

Fig. 4. Directed graphs of Boolean networks corresponding to the toy example.

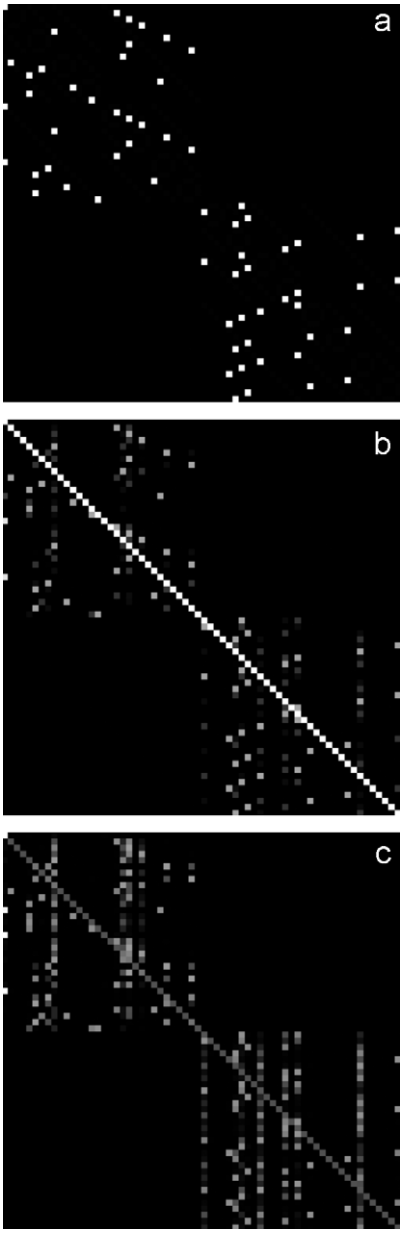

Fig. 5. Heat-map of transition probability matrix. The brightness corresponds to probability. (a) Context-sensitive PBN, (b) SRD-PBN for sampling period $T = 2$, (c) SRD-PBN for sampling period $T = 4$.

following the methodology described in Faryabi et al. (2009). The switching probability q is chosen to be 0.01 and there exists a gene perturbation probability of 0.01. It is clear that most of the states have zero self-transition probabilities. To construct the transition probability matrix of the SRD-PBN model, we first select λ and γ to be 0.1 and 0.05, respectively. The rate matrix Q is computed based on (6). The transition probability matrix of the SRD-PBN corresponding to this matrix Q for sampling period of $T = 2$ is computed based on (7) and the corresponding heat-map is shown in Fig. 5(b). A similar procedure is repeated for $T = 4$ and the heat-map corresponding to the transition probability matrix of the SRD-PBN is shown in Fig. 5(c). It is evident that the self-transition probabilities in Fig. 5(b) and (c) are not zero. Since the diagonal elements are brighter, these values are different for $T = 2$ and 4. Intuitively, we expect a higher self-transition probability for a smaller sampling period and a lower self-transition probability for a larger sampling period. It can be seen that

self-transition probabilities are larger in Fig. 5(b) for $T=2$ compared to Fig. 5(c) for $T=4$.

4. Results and discussion

Our prime goal of modeling gene regulatory networks from temporal gene expression data is to derive effective intervention strategies and beneficially alter the long-run behavior of the inferred model. From a practical point of view, at every observation point, this strategy decides which action should be applied to the underlying biological system. Provided that the model framework captures the dynamics of gene regulatory networks accurately, the derived intervention strategy would favorably alter the behavior of aberrant cells.

In this section, through numerical studies, we provide supporting evidence for the need to extend the original PBN framework. In the following simulations, the target gene, the gene responsible for aberrant behavior of the cell, is chosen to be the most significant gene in the GAP. We assume the up-regulation of the target gene is undesirable. Consequently, the state-space is partitioned into desirable states, $\mathcal{D}$, and undesirable states, $\mathcal{U}$. Since our objective is to down-regulate the target gene, a higher cost is assigned to destination states having an up-regulated target gene. Moreover, for a given status of the target gene for a destination state, a higher cost is assigned when the control is applied, versus when it is not. In practice, the cost values will have to mathematically capture the benefits and costs of intervention and the relative preference of states. These cost values will eventually be set with the help of physicians in accordance with their clinical judgement. Although this is not feasible within current medical practice, we do believe that such an approach will become feasible when engineering approaches are integrated into translational medicine.

We postulate the following cost-per-stage in state j under control u:

$$r(u,j)=\begin{cases}0 & \text{if } u=0 \text{ and } j\in\mathcal{D},\\ 5 & \text{if } u=0 \text{ and } j\in\mathcal{U},\\ c & \text{if } u=1 \text{ and } j\in\mathcal{D},\\ 5+c & \text{if } u=1 \text{ and } j\in\mathcal{U},\end{cases} \tag{8}$$

where c denotes the cost of control. We study the effect of c in our simulations. A cost minimization framework is used to effectively trade-off the number of interventions and the likelihood of the network being in an undesirable state. An optimal control policy with regard to the cost values can be found via dynamic programming.

In our simulation studies, our objective is to show that an optimal policy derived for the current definition of context-sensitive PBN will no longer be optimal if we include the timing information of temporal data into the dynamics of gene regulatory networks. To this end, we generate synthetic SRD-PBNs and corresponding context-sensitive PBNs. We compute the cost induced by the optimal policy derived for the context-sensitive PBN and the cost induced by the optimal policy derived for the SRD-PBN, when both are applied to a sequence of data generated from the SRD-PBN. These two cost values are compared in our simulation studies. An SRD-PBN accommodates the sampling rate, which is in this simulation identical to the intervention rate. The goal of this study is to measure how costly it is to apply an optimal policy derived for a context-sensitive PBN to a sequence of data generated based on an SRD-PBN. In the following, we consider synthetically generated SRD-PBNs.

We generate SRD-PBNs in the following manner. Each SRD-PBN consists of two constituent Boolean networks. Each Boolean function of a Boolean network is randomly generated with a random bias. Given a set of Boolean networks, we generate various SRD-PBNs. We let $\gamma=0.01$. We vary the value of λ from 0.05 to 4 with step-size 0.2. We choose the gene perturbation probability of 0.01. The constituent Boolean networks are selected with equal probabilities. Furthermore, for the given set of Boolean networks, we generate the corresponding context sensitive PBNs for switching probability $q=0.01$. We let the observation period to be every 1 unit of time, i.e. $T=1$. Transition probability matrix of context-sensitive PBN is computed based on Faryabi et al. (2009).

Using dynamic programming, given the cost-per-stage defined in (8), we derive an optimal intervention policy μ^*_{srd} for an SRD-PBN. Our goal is to estimate the average total discounted cost induced by μ^*_{srd} for a sequence of data generated from the SRD-PBN. To this end, we generate synthetic time-course data for 1000 time-steps from the SRD-PBN model while μ^*_{srd} is applied. We estimate the discounted cost by accumulating the discounted cost of each state given the action at that state. This procedure is repeated 10,000 times for random initial states and the average of the induced discounted cost is computed. Likewise, an optimal policy μ^*_{cs} for a context-sensitive PBN is derived. Following a similar procedure, μ^*_{cs} is applied to the SRD-PBN, which we already described, and the average discounted cost is computed. Moreover, we compute the average discounted cost of a sequence of time-course data for an SRD-PBN in the absence of intervention.

In sum, for each set of Boolean networks, we have the following: ($\overline{C}^{\text{srd}}$) average total discounted cost induced by μ^*_{srd} on the SRD-PBN; ($\overline{C}^{\text{cs}}$) average total discounted cost induced by μ^*_{cs} on the SRD-PBN; ($\overline{C}^{\text{woc}}$) average total discounted cost induced in the absence of any intervention on SRD-PBN. The preceding procedure is repeated for 1000 random Boolean networks, thereby yielding 1000 values for each statistic: $\overline{C}^{\text{srd}}_1,\ldots,\overline{C}^{\text{srd}}_{1000}$; $\overline{C}^{\text{cs}}_1,\ldots,\overline{C}^{\text{cs}}_{1000}$; $\overline{C}^{\text{woc}}_1,\ldots,\overline{C}^{\text{woc}}_{1000}$. Using these, we compare the effect of μ^*_{srd} and μ^*_{cs} on an SRD-PBN by the empirical averages $\mathbf{M}[C^{\text{srd}}]$ of $\overline{C}^{\text{srd}}_1,\ldots,\overline{C}^{\text{srd}}_{1000}$; $\mathbf{M}[C^{\text{cs}}]$ of $\overline{C}^{\text{cs}}_1,\ldots,\overline{C}^{\text{cs}}_{1000}$; and $\mathbf{M}[C^{\text{woc}}]$ of $\overline{C}^{\text{woc}}_1,\ldots,\overline{C}^{\text{woc}}_{1000}$. We define the gain obtained by each intervention policy as the difference between the average discounted cost before and after intervention. G_{srd}, the gain of policy μ^*_{srd}, is $\mathbf{M}[C^{\text{woc}}]-\mathbf{M}[C^{\text{srd}}]$ and G_{cs}, the gain of policy μ^*_{cs} applied to an SRD-PBN, is $\mathbf{M}[C^{\text{woc}}]-\mathbf{M}[C^{\text{cs}}]$. We are interested in $\mathbf{M}[C^{\text{cs}}]-\mathbf{M}[C^{\text{srd}}]$, which we refer to as ΔG. ΔG measures how costly it is to apply an optimal control policy derived for a context-sensitive PBN to a sequence of data generated based on an SRD-PBN.

Figs. 6–8 demonstrate the outcome of the above experiment for various values of cost of control c. It is evident that the intervention gains G_{srd} and G_{cs} are larger for smaller cost of intervention. The structure of a context-sensitive PBN is such that there is a transition to a new state after each unit of time, which corresponds to one change at every unit of time on average. When λ is substantially smaller or larger than 1, ΔG is larger compared to the case where λ is closer to 1, as is shown in Figs. 6–8. We should point out that the value of λ for which ΔG attains its minimum depends on many factors, such as γ, the switching probability q in context-sensitive PBN, and the cost of control. It is also observed that ΔG increases for larger cost of control.

We emphasize that this simulation study compares the gains obtained by two policies, the policy optimal for the SRD-PBN and the policy optimal for the context-sensitive PBN, when each is applied to SRD-PBN. Our objective is to show how poor the effect of an intervention policy derived for a context-sensitive PBN might be if the rate of change among observations is substantially different from 1.

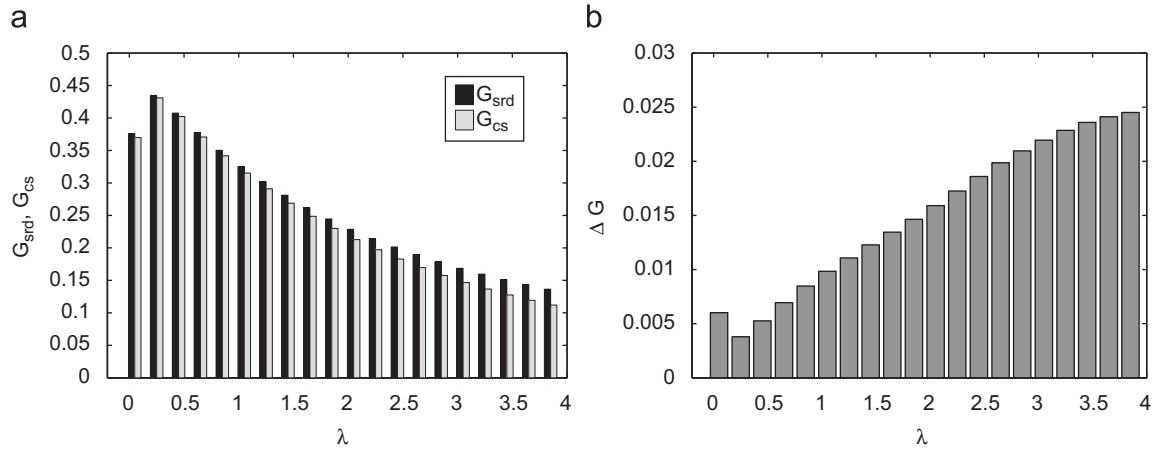

Fig. 6. (a) G_{srd}, gain obtained by the policy optimal for SRD-PBN, and G_{cs}, the gain obtained by the policy optimal for context-sensitive PBN, when both are applied to SRD-PBN for various λ. (b) Difference between the gains, ΔG, for various λ. The cost of control is 0.1.

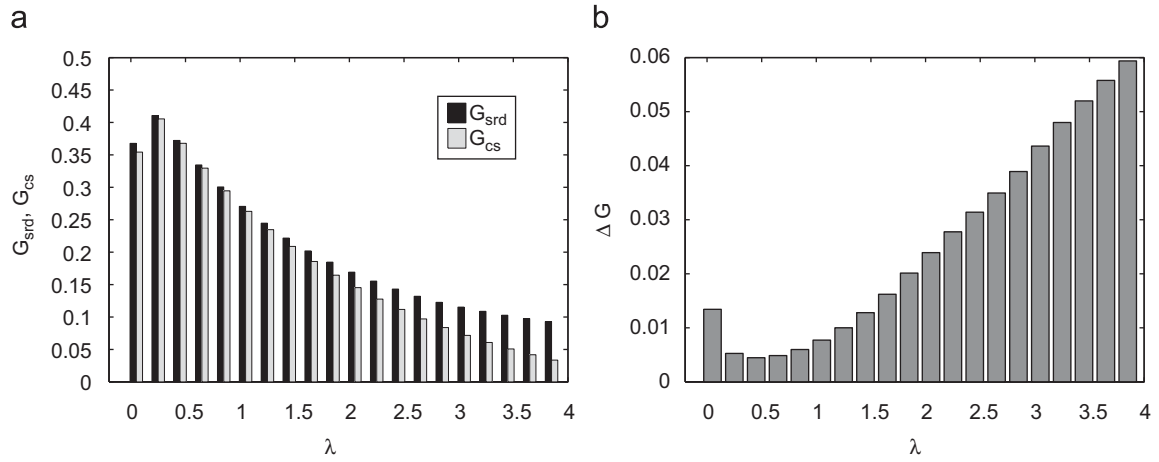

Fig. 7. (a) G_{srd}, gain obtained by the policy optimal for SRD-PBN, and G_{cs}, the gain obtained by the policy optimal for context-sensitive PBN, when both are applied to SRD-PBN for various λ. (b) Difference between the gains, ΔG, for various λ. The cost of control is 1.0.

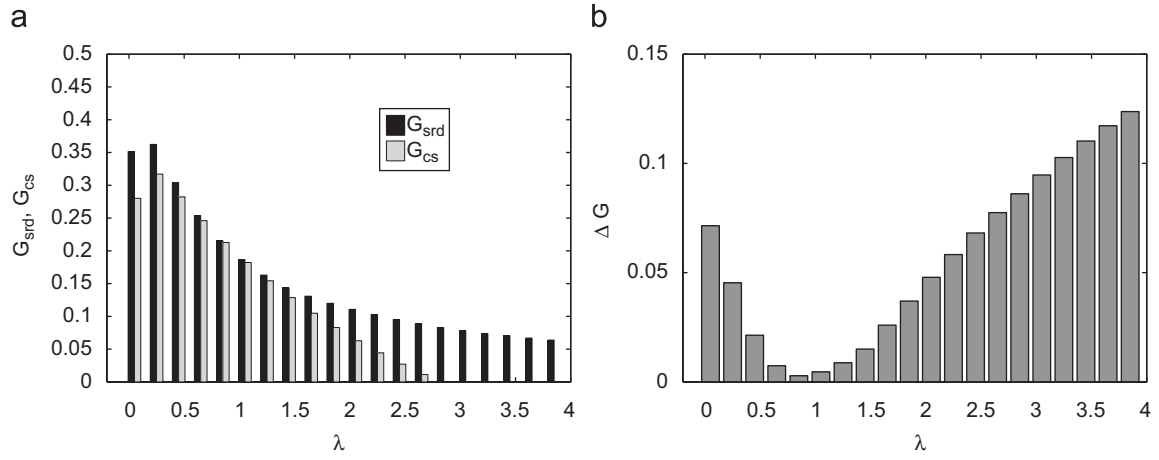

Fig. 8. (a) G_{srd}, gain obtained by the policy optimal for SRD-PBN, and G_{cs}, the gain obtained by the policy optimal for context-sensitive PBN, when both are applied to SRD-PBN for various λ. (b) Difference between the gains, ΔG, for various λ. The cost of control is 3.

5. Conclusion

Implementation of an intervention policy derived for PBNs or asynchronous PBNs requires nearly continuous observation of the underlying biological system since precise application requires the observation of all jumps. In medical applications, as in many engineering problems, this is not the case. The process is sampled at discrete time intervals and a decision to intervene or not must be made at each sample point. Our goal in this work has been to construct a framework as an

extension of PBNs that can capture the sampling rate of the underlying system.

By estimating a few temporal parameters of the underlying biological system, we construct a continuous-time process that can embody the sojourn time of each state. Our goal is to construct a discrete-time model by sampling the continuous-time process. To guarantee the Markov property for the sampled process, the continuous-time process requires the memoryless property. To this end, we select a continuous-time Markov chain which leads the sampled process to be a Markov chain. This Markov chain models the dynamics of our proposed framework. By only involving two parameters, γ and λ, in the timing model, the model is not overburdened with the estimation of a large number of timing parameters or timing distributions.

References

Bertsekas, D.P., 2005. Dynamic Programming and Optimal Control. Athena Scientific, Belmont, MA.

Bouvet, M., Bold, R.J., Lee, J., Evans, D.B., Abbruzzese, J.L., Chiao, P.J., McConkey, D.J., Chandra, J., Chada, S., Fang, B., Roth, J.A., 1998. Adenovirus-mediated wild-type p53 tumor suppressor gene therapy induces apoptosis and suppresses growth of human pancreatic cancer. Annals of Surgical Oncology 5 (8), 681–688.

Deng, X., Geng, H., Matache, M.T., 2007. Dynamics of asynchronous random Boolean networks with asynchrony generated by stochastic processes. BioSystems 88 (1–2), 16–34.

El-Deiry, W.S., Tokino, T., Velculescu, V.E., Levy, D.B., Parsons, R., Trent, J.M., Lin, D., Mercer, W.E., Kinzler, K.W., Vogelstein, B., 1993. Waf1, a potential mediator of p53 tumor suppression. Cell 75 (4), 817–825.

Faryabi, B., Chamberland, J.F., Vahedi, G., Datta, A., Dougherty, E.R., 2008. Optimal intervention in asynchronous genetic regulatory networks. IEEE Journal of Selected Topics in Signal Processing 2 (3), 412–423.

Faryabi, B., Chamberland, J.F., Vahedi, G., Datta, A., Dougherty, E.R., 2009. Intervention in context-sensitive probabilistic Boolean networks revisited. EURASIP Journal on Bioinformatics and Systems Biology 2009 (2009), Article ID 360864, 13 pages.

Gershenson, C., 2002. Classification of random Boolean networks. In: Proceedings of the Eighth International Conference on Artificial Life (Artificial Life VIII), Sydney, Australia, December 2002, pp. 1–8.

Greil, F., Drossel, B., 2005. The dynamics of critical Kauffman networks under asynchronous stochastic update. Physical Review Letters 95 (4), 048701.

Harvey, I., Bossomaier, T., 1997. Time out of joint: attractors in asynchronous random Boolean networks. In: Proceedings of the Fourth European Conference on Artificial Life (ECAL97), July 1997. MIT Press, Cambridge, MA, pp. 67–75.

Huang, S., 1999. Gene expression profiling, genetic networks, and cellular states: an integrating concept for tumorigenesis and drug discovery. Journal of Molecular Medicine 77 (6), 469–480.

Kauffman, S., 1993. The Origins of Order: Self-Organization and Selection in Evolution. Oxford University Press, New York.

Miyashita, T., Reed, J.C., 1995. Tumor suppressor p53 is a direct transcriptional activator of the human Bax gene. Cell 80 (2), 193–199.

Norris, J.R., 1998. Markov Chains. Cambridge University Press, Cambridge.

Oktem, H., Pearson, R., Egiazarian, K., 2003. An adjustable aperiodic model class of genomic interactions using continuous time Boolean networks (Boolean delay equations). Chaos 3 (4), 1167–1174.

Owen-Schaub, L.B., Zhang, W., Cusack, J.C., Angelo, L.S., Santee, S.M., Fujiwara, T., Roth, J.A., Deisseroth, A.B., Zhang, W.W., Kruzel, E., Radinsky, R., 1995. Wild-type human p5 and a temperature-sensitive mutant induce Fas/APO-1 expression. Molecular and Cellular Biology 15 (6), 3032–3040.

Pal, R., Datta, A., Dougherty, E.R., 2006. Optimal infinite-horizon control for probabilistic Boolean networks. IEEE Transactions on Signal Processing 54 (6), 2375–2387.

Shmulevich, I., Dougherty, E.R., Kim, S., Zhang, W., 2002. Probabilistic Boolean networks: a rule-based uncertainty model for gene regulatory networks. Bioinformatics 18 (2), 261–274.

Swisher, S.G., Roth, J.A., et al., 1999. Adenovirus-mediated p53 gene transfer in advanced non-small-cell lung cancer. Journal of the National Cancer Institute 91 (9), 763–771.

Thomas, R., 1979. Kinetic Logic: A Boolean Approach to the Analysis of Complex Regulatory Systems. Springer, Berlin.