33rd Annual International Conference of the IEEE EMBS
Boston, Massachusetts USA, August 30 - September 3, 2011

Robustness of Inverse Perturbation for Discrete Event Control

Nidhal Bouaynaya[1], Roman Shterenberg[2], and Dan Schonfeld[3]

[1]Dept. of Systems Engineering, University of Arkansas at Little Rock
[2]Dept. of Mathematics, University of Alabama at Birmingham
[3]Dept. of Electrical and Computer Engineering, University of Illinois at Chicago

***Abstract*—We study the robustness of the inverse perturbation solution in discrete-time systems modeled by homogeneous Markov chains. We cast the optimal inverse perturbation control as a strictly convex optimization problem, which admits a unique global solution. We show that the optimal inverse perturbation control is robust to estimation errors in the original network. The derived results are applied to the Human melanoma gene regulatory network, where the aim is to force the network to converge to a desired steady-state distribution of gene regulation.**

***Index Terms*—Finite Markov chains; Perturbation theory; Inverse perturbation; stability; robustness.**

I. Introduction

Perturbation analysis examines the effect of small perturbations on system parameters. Specifically, it has been proved that the the infinitesimal perturbation analysis provides unbiased or strongly consistent estimates of the performance for many systems [1], [2]. In this paper, we consider discrete-time systems, which can be modeled by finite-state Markov chains. Examples include queueing networks, resource allocation, social and biological networks, and machine replacement. Finite state Markov chains are among the most widely used probabilistic models of discrete event stochastic phenomena. It is therefore of interest to obtain their performance sensitivities under various kinds of perturbations. Of particular importance is the sensitivity of the steady-state distribution to perturbations in the probability transition matrix of the chain. The steady-state perturbation problem in Markov chains can be formulated as follows: Given an initial Markov chain characterized by its probability transition matrix P_0, consider linear perturbations of P_0 as $P(\epsilon) = P_0 + \epsilon C$, where C is a fixed zero-row sum matrix and $0 < \epsilon < \epsilon_{\max}$. What is the relationship between the stationary distribution $\pi(\epsilon)$ of the perturbed chain and the original (unperturbed) stationary distribution $\pi(0)$. Schweitzer [3] solved this problem in the case of irreducible Markov chains.

In this paper, we harness the perturbation problem into a control strategy that forces the chain to converge towards a desired steady-state distribution. Specifically, we address the following question: "Given a Markov chain characterized by its probability transition matrix P_0 and given a desired steady-state distribution π_d, can we find a perturbation matrix C that forces the chain $(P_0 + C)$ to converge towards the desired distribution π_d?" We call this formulation *the inverse perturbation problem* because it aims at finding a perturbation, which changes the dynamics of the network in a desirable way, rather than finding the new steady-state given a known perturbation of the system [4]. Observe that, as formulated, the inverse perturbation problem is a feasibility problem. We further propose to find the optimal perturbation matrix, which minimizes the "energy" between the initial and perturbed networks as measured by the Frobenius norm of the perturbation matrix. We cast the optimal inverse perturbation problem as a strictly convex optimization problem, thus admitting a unique global solution, which can be efficiently computed for large networks using convex optimization solvers [5].

We further investigate the robustness of the optimal perturbation matrix to estimation errors in the original probability transition matrix. In practical applications, errors made during data extraction, feature selection, and network inference will propagate and impact the actual success of the designed control. An efficient intervention approach must possess some degree of "robustness" or insensitivity to data and estimation errors. In this paper, we show that the optimal inverse perturbation control is robust to errors in the original probability transition matrix, in the sense that the estimation error of the optimal control solution is bounded by the estimation error of the probability transition matrix.

The proposed inverse perturbation control is applied to the Human melanoma gene regulatory network. Biological evidence suggests that steady-state distributions of molecular networks reflect the phenotype of the cell. In other words, both malignant (e.g. cancer) and benign phenotypes correspond to steady-state distributions in dynamic system models of gene regulatory networks. The inverse perturbation control introduces an isolated, one-time intervention by optimal perturbation in order to ensure that the network converges to a desired steady-state distribution of gene regulation. The ultimate goal is to develop engineering methods that will be employed to intervene in living organisms to drive cells away from a malignant state and into a benign one.

II. Mathematical Notation

We denote by $\mathcal{M}$ the Euclidean space of $n \times n$ square matrices equipped with the inner product $< A, B >= \mathrm{Tr}(A^t B)$, where $\mathrm{Tr}(A)$ denotes the trace of $A \in \mathcal{M}$. $\mathbf{1}$ denotes a vector all of whose components are equal to one. The equality and inequality symbols, $=, \leq$ and $\geq$ denote component-wise equality and inequality, respectively, for arrays of the same size.

III. Inverse Perturbation Control

Definition 1: A row probability vector μ^t is called a stationary distribution or a steady-state distribution for P_0 if $\mu^t P_0 = \mu^t$.

Because P_0 is stochastic (i.e., its rows sum up to 1), the existence of stationary distributions is guaranteed. In this paper, we assume that P_0 is irreducible, thus having a unique steady-state distribution π_0. In practice, there are several fast algorithms for checking irreducibility in graphs [6].

We wish to alter the stationary distribution π_0 by linearly perturbing the probability transition matrix P_0. Specifically, we consider the perturbed stochastic matrix

$$P = P_0 + C, \tag{1}$$

where C is a zero row-sum perturbation matrix. The zero row-sum condition is necessary to ensure that the perturbed matrix P is stochastic. Let us denote by π_d the desired stationary distribution. We seek to design an optimal zero row-sum perturbation matrix C such that the perturbed matrix P is irreducible and converges to the desired steady-state distribution π_d.

The Inverse Perturbation Problem: The set of perturbation matrices C, which force the chain to transition from π_0 to π_d satisfy the following constraints:

$$(i)\ \pi_d^t = \pi_d^t(P_0 + C), \quad (ii)\ C\mathbf{1} = \mathbf{0}, \quad (iii)\ P_0 + C \geq 0. \tag{2}$$

Constraint (i) guarantees that the perturbed chain or network converges towards π_d. Constraints (ii) and (iii) ensure that the perturbed matrix P is a proper probability transition matrix: constraint (ii) imposes that the perturbation matrix C is zero-row sum, and constraint (iii) requires the matrix P to be element-wise non-negative. Let $\mathcal{D}$ denote the feasible set of perturbation matrices , i.e.,

$$\mathcal{D} = \{C \in \mathbb{R}^{n\times n} : \pi_d^t = \pi_d^t(P_0 + C), C\mathbf{1} = \mathbf{0}, P_0 + C \geq 0\}. \tag{3}$$

$\mathcal{D}$ is a polyhedra as the solution of a finite number of linear equalities and inequalities [5]. In particular, $\mathcal{D}$ is convex [5]. Moreover, $\mathcal{D}$ is non-empty because it contains the perturbation matrix $C_0 = \mathbf{1}\pi_d^t - P_0$. Therefore, at least one feasible solution exists.

Proposition 1: Let $C_0 = \mathbf{1}\pi_d^t - P_0$. Any feasible solution $C \in \mathcal{D}$ can be written as

$$C = C_0 + V, \tag{4}$$

where $V \in \Gamma$ and

$$\Gamma = \{V \in \mathcal{M} : V\mathbf{1} = \mathbf{0}, V^t\pi_d = \mathbf{0}, V \geq -\mathbf{1}\pi_d^t\}. \tag{5}$$

Observe that $V = \mathbf{0} \in \Gamma$. In particular, there may be numerous (possibly infinite) perturbation matrices C which can force the chain to transition from an undesirable steady-state to a desirable one. All such perturbations, in principle, constitute plausible control strategies and can therefore be used to drive the network from one steady-state to another. In many applications, however, we would like to find the optimal perturbation matrix, in a specified sense. One optimality criterion, of interest, is to minimize the "energy" between the original and perturbed networks, as measured by the Frobenius norm of the perturbation matrix. The minimum energy constraint can be imposed to limit the structural changes in the network before and after control.

The Optimal Inverse Perturbation Control: The optimal inverse perturbation problem can be formulated as the following optimization problem:

$$\text{Minimize} \quad \|C\|_F^2 \quad \text{subject to} \quad C \in \mathcal{D}, \tag{6}$$

where $\|.\|_F$ denotes the Frobenius norm given by $\|C\|_F^2 = \sum_{i=1}^n \sum_{j=1}^n c_{ij}^2 = \text{Tr}(CC^t)$, and $\mathcal{D}$ is the feasible set in Eq. (3).

Observe that the optimization problem in (6) is convex, i.e., the objective function is convex, the equality constraints are affine and the inequality constraints are convex [5]. Moreover, since the square of the Frobenius norm is strictly convex, the optimization problem in (6) has a unique global solution [5]. It is important to notice, however, that the optimal perturbed matrix may not be irreducible. Nonetheless, it converges towards the desired steady-state starting from its basin of attraction.

The following proposition provides a geometrical characterization of the optimization problem in (6).

Proposition 2: Let $C_0 = \mathbf{1}\pi_d^t - P_0$. Then, the optimal perturbation matrix is given by

$$C^* = C_0 - C_0^\Gamma, \tag{7}$$

where C_0^Γ denotes the unique projection of C_0 onto the convex set Γ [1] given in Eq. (5).

In practice, one can use numerical algorithms to find the projection onto the convex set Γ [7].

IV. Robustness of the Optimal Inverse Perturbation Control

We now assume that the estimated probability transition matrix $\hat{P}_0$ is given by

$$\hat{P}_0 = P_0 + \delta P_0, \tag{8}$$

where δP_0 is a zero-row sum matrix representing noisy and missed data and estimation errors in P_0. We will show that the norm of the error in the optimal inverse perturbation matrix is bounded by the norm of the error in P_0. We first present the following Lemma.

Lemma 1: Consider a vector subspace $\mathcal{E}$, equipped with an inner product norm $\|.\|$, and a convex subset $\mathcal{C} \subseteq \mathcal{E}$. Let x_1, x_2 be two points in $\mathcal{E}$, and p_1, p_2 be their respective closest points in $\mathcal{C}$. Then, we have

$$\|p_1 - p_2\| \leq \|x_1 - x_2\|. \tag{9}$$

Proposition 3: The estimated optimal perturbation matrix, $\hat{C}^*$, satisfies $\hat{C}^* = C^* + \delta C^*$, where C^* is the optimal perturbation matrix and δC^* is an error zero-row sum matrix satisfying

$$\|\delta C^*\|_F \leq \|\delta P_0\|_F. \tag{10}$$

[1]Note that projection, here, refers to minimal distance to the set. Since Γ is not a vector subspace, we cannot define (orthogonal) projection onto Γ. However, we can determine the closest point in Γ to a given point in $\mathcal{M}$.

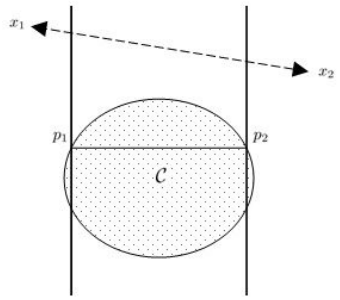

Fig. 1. Graphical illustration of the proof of Lemma 1.

V. Simulation Results

We apply the inverse perturbation control to the melanoma gene regulatory network, which is one of the most studied gene regulatory networks in the literature [8]. The control strategy is to reduces WNT5A's action in inducing a metastatic phenotype. A seven-gene probabilistic Boolean network model of the melanoma network containing the genes WNT5A, pirin, S100P, RET1, MART1, HADHB, and STC2 was derived in [8]. Note that the Human melanoma Boolean network consists of $2^7 = 128$ states ranging from $00\cdots0$ to $11\cdots1$, where the states are ordered as WNT5A, pirin, S100P, RET1, MART1, HADHB, and STC2, with WNT5A and STC2 denoted by the most significant bit (MSB) and least significant bit (LSB), respectively. The probability transition matrix of the Human melanoma network, used in this paper, is courtesy of Dr. Ranadip Pal.

We consider the (fictitious) desired steady-state distribution where the probability of the states having WNT5A upregulated is 10^{-4} and the probability of the other states, which correspond to WNT5A downregulated is chosen randomly such that the total probability mass is equal to 1 (see Fig. 2(c)). Observe that the states from 0 to 63 have WNT5A downregulated (0) and hence are desirable states, as compared to states 64 to 127 that have WNT5A upregulated (1) and hence are undesirable. The probability transition matrices of the Human melanoma networks corresponding to the original and perturbed networks are portrayed in Fig. 2(a) and Fig. 2(b), respectively. Note that the controlled and desired steady-state distributions are identical. Figures 2(d) and 2(e) show the matrices δP_0 and δC^*, respectively. In this simulation, δP_0 is the zero-row sum matrix having only two non-zero entries in each row, equal to $\epsilon_i, -\epsilon_i,\ 1 \leq i \leq n$, where ϵ_i is chosen such that $\hat{P}_0 = P_0 + \delta P_0$ is a proper probability transition matrix, i.e., $\hat{P}_0$ is stochastic with positive entries. We verify the stability of the inverse perturbation solution, i.e., $\|\delta C^*\|_F = 2.6690 \leq \|\delta P_0\|_F = 4.8149$, as displayed in Fig. 2(f). This simulation has been implemented in MATLAB, and the (strictly) convex optimization problem in (6) has been solved using the CVX software for convex optimization [9].

VI. Conclusion

We developed a mathematical framework for the solution of the optimal inverse perturbation problem for irreducible Markov chains. Our aim was to derive a minimal-perturbation intervention control in order to change the network's dynamics and force it to converge to a desirable steady-state distribution. We further investigated the robustness of the optimal inverse perturbation solution with respect to estimation errors in the probability transition matrix of the chain. We showed that the estimation error in the optimal inverse perturbation control is bounded by the error in the probability transition matrix of the chain. We applied the proposed optimal inverse perturbation control to the Human melanoma gene regulatory network to ensure that it converges to a desired steady-state distribution of gene regulation. The ultimate goal is to develop engineering methods that will be employed to intervene in living organisms to drive cells away from a malignant state and into a benign one.

Appendix

Proof of Proposition 1: The proof is straightforward by checking that matrices of the form $C = C_0 + V$, where $V \in \Gamma$ satisfy the three conditions in Eq. (2). Conversely, given a feasible perturbation matrix $C \in \mathcal{D}$ we get that $C - C_0 \in \Gamma$. ■

Proof of Proposition 2: From Proposition 1, the optimization problem in (6) is equivalent to minimizing the Frobenius norm of $(C_0 + V)$ subject to $V \in \Gamma$. It is straightforward that the solution to this problem is

$$V^* = -C_0^{\Gamma}. \tag{11}$$

Therefore, $C^* = C_0 - C_0^{\Gamma}$. ■

Proof of Lemma 1: Consider the line (p_1, p_2). Let $\mathcal{H}_1$ (resp., $\mathcal{H}_2$) be the hyperplane orthogonal to (p_1, p_2) at p_1 (resp., p_2). Then, x_1 must be to the left of $\mathcal{H}_1$ (see Fig. 1), otherwise some point strictly inside the segment $]p_1, p_2[\subseteq \Gamma$ will be closer to x_1 than p_1. Similarly, x_2 must be to the right of $\mathcal{H}_2$. Therefore, $\|x_1 - x_2\| \geq \|p_1 - p_2\|$. ■

Proof of Proposition 3: From proposition 2, we have

$$\begin{aligned} C^* - \hat{C}^* &= (C_0 - C_0^{\Gamma}) - (\hat{C}_0 - \hat{C}_0^{\Gamma}) && (12)\\ &= (C_0 - \hat{C}_0) + (\hat{C}_0^{\Gamma} - C_0^{\Gamma}) && (13)\\ &= \delta P_0 + (\hat{C}_0^{\Gamma} - C_0^{\Gamma}). && (14) \end{aligned}$$

Let $\mathcal{E}$ be the subspace defined by the two equality constraints, i.e., $\mathcal{E} = \{V \in \mathcal{M} : V\mathbf{1} = \mathbf{0}, V^t \pi_d = \mathbf{0}\}$. Let us consider the decomposition of δP_0 as

$$\delta P_0 = \delta P_0^{\mathcal{E}} + \delta P_0^{\mathcal{E}^{\perp}}, \tag{15}$$

where $\delta P_0^{\mathcal{E}}$ and $\delta P_0^{\mathcal{E}^{\perp}}$ are, respectively, the projections of δP_0 onto the subspace $\mathcal{E}$ and its orthogonal $\mathcal{E}^{\perp}$. Then Eq. (14) becomes

$$C^* - \hat{C}^* = \delta P_0^{\mathcal{E}} + \delta P_0^{\mathcal{E}^{\perp}} + (\hat{C}_0^{\Gamma} - C_0^{\Gamma}). \tag{16}$$

Let us further decompose $\delta P_0^{\mathcal{E}}$ into the subspaces parallel and orthogonal to $\hat{C}_0^{\Gamma} - C_0^{\Gamma}$, i.e.,

$$\delta P_0^{\mathcal{E}} = \delta P_{0\parallel}^{\mathcal{E}} + \delta P_{0\perp}^{\mathcal{E}}. \tag{17}$$

We have

$$\begin{aligned} \|\delta P_0^{\mathcal{E}} + \hat{C}_0^{\Gamma} - C_0^{\Gamma}\|^2 &= \|\delta P_{0\parallel}^{\mathcal{E}} + \delta P_{0\perp}^{\mathcal{E}} + \hat{C}_0^{\Gamma} - C_0^{\Gamma}\|^2 \\ &= \|\delta P_{0\parallel}^{\mathcal{E}} + \hat{C}_0^{\Gamma} - C_0^{\Gamma}\|^2 + \|\delta P_{0\perp}^{\mathcal{E}}\|^2. \end{aligned}$$

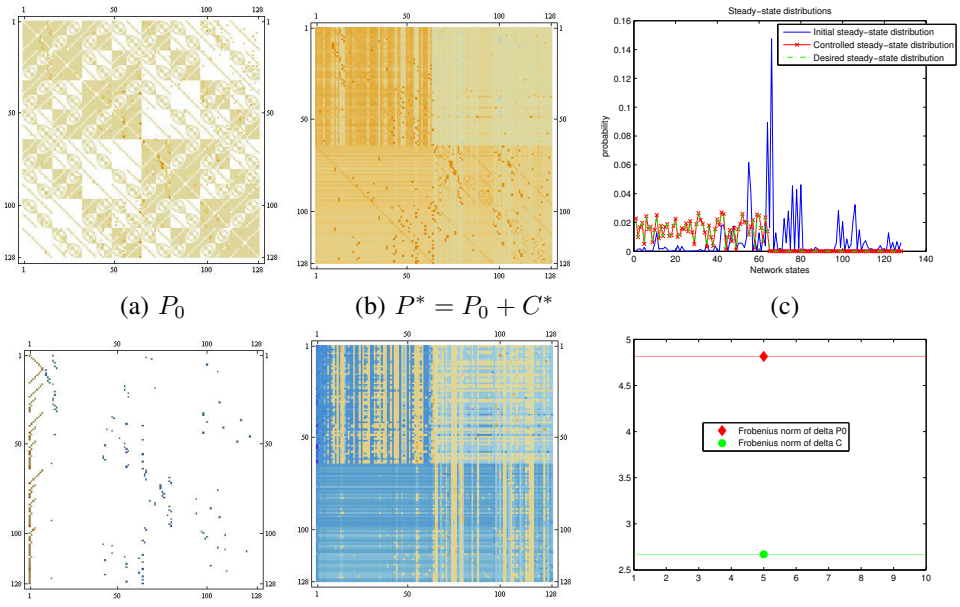

(a) P_0 (b) $P^* = P_0 + C^*$ (c)

(d) δP_0 (e) δC^* (f)

Fig. 2. Initial and controlled probability transition matrices for the Human melanoma gene regulatory network: The matrix plots are obtained using the function *MatrixPlot* in MATHEMATICA. The color of entries varies from white to red corresponding to the values of the entries in the range of 0 to 1: (a) the probability transition matrix of the original melanoma network P_0; (b) the optimal perturbed probability transition matrix P^*; (c) The original (red line), desired (blue line), and minimal-perturbation energy controlled (green line) steady-state distributions of the Human melanoma gene regulatory network. The x-axis represents the 128 states of the network, and the y-axis indicates the probability of each state; (d) The original error δP_0; (e) the inverse perturbation solution error δC^*; (f) The Frobenius norms of δP_0 and δC^*.

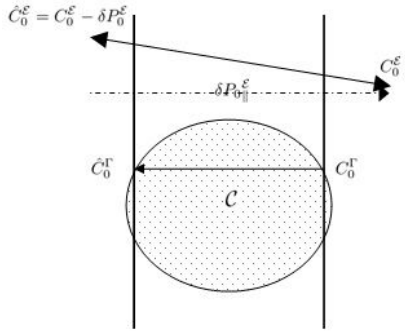

Fig. 3. Graphical illustration of the proof of Proposition 3

From Lemma 1 and given that $\delta P_{0\|}^{\mathcal{E}}$ is in the opposite direction as $\hat{C}_0^{\Gamma} - C_0^{\Gamma}$ (see Fig. 3), we have

$$\|\delta P_{0\|}^{\mathcal{E}} + \hat{C}_0^{\Gamma} - C_0^{\Gamma}\|^2 \leq \|\delta P_{0\|}^{\mathcal{E}}\|^2. \tag{18}$$

Therefore,

$$\|\delta P_0^{\mathcal{E}} + \hat{C}_0^{\Gamma} - C_0^{\Gamma}\|^2 \leq \|\delta P_{0\|}^{\mathcal{E}}\|^2 + \|\delta P_{0\perp}^{\mathcal{E}}\|^2 = \|\delta P_0{}^{\mathcal{E}}\|^2.$$

Finally, Eq. (16) becomes

$$\begin{aligned} \|C^* - \hat{C}^*\|^2 &= \|\delta P_0^{\mathcal{E}} + \delta P_0^{\mathcal{E}^\perp} + \hat{C}_0^{\Gamma} - C_0^{\Gamma}\|^2 && (19)\\ &= \|\delta P_0^{\mathcal{E}} + \hat{C}_0^{\Gamma} - C_0^{\Gamma}\|^2 + \|\delta P_0^{\mathcal{E}^\perp}\|^2 && (20)\\ &\leq \|\delta P_0^{\mathcal{E}}\|^2 + \|\delta P_0^{\mathcal{E}^\perp}\|^2 = \|\delta P_0\|^2 && (21) \end{aligned}$$

■

VII. Acknowledgment

The authors would like to extend their gratitude to Dr. Ranadip Pal from Texas Tech University for providing the Human melanoma gene regulatory network dataset.

This project is supported by Award Number R01GM096191 from the National Institute Of General Medical Sciences (NIH/NIGMS). The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institute Of General Medical Sciences or the National Institutes of Health.

References

[1] Y. C. Ho and X. R. Cao, *Perturbation Analysis of Discrete-Event Dynamic Systems*, Kluwer, Boston, MA, 1991.

[2] E. K. P. Chong and P. J. Ramadge, "Optimization of queues using infinitesimal perturbation analysis-based algorithms with general update times," *SIAM Journal on Control and Optimization*, vol. 31, pp. 698732, 1993.

[3] P J Schweitzer, "Perturbation series expansions for nearly completely-decomposable markov chains," in *Proceedings of the international seminar on Teletraffic analysis and computer performance evaluation*, 1986, pp. 319 – 328.

[4] N. Bouaynaya, R. Shterenberg, and D. Schonfeld, "Inverse perturbation for optimal intervention in gene regulatory networks," *Bioinformatics*, vol. 27, no. 1, pp. 103–110, January 2011.

[5] S. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, 2003.

[6] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, Prentice Hall, 2003.

[7] C. L. Byrne, *Signal Processing: A Mathematical Approach*, A K Peters, Ltd., 2005.

[8] R. Pal, A. Datta, and E. Dougherty, "Optimal infinite horizon control for probabilistic Boolean networks," *IEEE Transactions on Signal Processing*, vol. 54, no. 6, pp. 2375–2387, 2006.

[9] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 1.21," http://cvxr.com/cvx, May 2010.