Optimal Design of Personalized Prostate Cancer Therapy using Infinitesimal Perturbation Analysis[☆]

Julia L. Fleck^{a,*}, Christos G. Cassandras^a

^a*Division of Systems Engineering and Center for Information and Systems Engineering,
Boston University, Brookline, MA 02446, USA*

Abstract

The standard treatment for advanced prostate cancer is hormone therapy in the form of continuous androgen suppression (CAS), which unfortunately frequently leads to resistance and relapse. An alternative scheme is intermittent androgen suppression (IAS), in which patients are submitted to cycles of treatment (in the form of androgen deprivation) and off-treatment periods in an alternating manner. In spite of extensive recent clinical experience with IAS, the design of ideal protocols for any given patient remains a challenge. The level of prostate specific antigen (PSA) is frequently monitored to determine when patients will be taken off therapy and when therapy will resume. In this work, we propose a threshold-based policy for optimal IAS therapy design that is parameterized by lower and upper PSA threshold values and is associated with a cost metric that combines clinically relevant measures of therapy success. We use a Stochastic Hybrid Automaton (SHA) model of prostate cancer evolution under IAS and perform Infinitesimal Perturbation Analysis (IPA) to adaptively adjust PSA threshold values so as to improve therapy outcomes. We also apply this methodology to clinical data from real patients, and obtain promising results and valuable insights for personalized IAS therapy design.

Keywords: stochastic hybrid system (SHS), perturbation analysis,

[☆]The authors' work is supported in part by NSF under grants CNS-1239021, ECCS-1509084, and IIP-1430145, by AFOSR under grant FA9550-15-1-0471, and by ONR under grant N00014-09-1-1051.

*Corresponding author

Email addresses: jfleck@bu.edu (Julia L. Fleck), cgc@bu.edu (Christos G. Cassandras)

1. Introduction

Cancer is currently deemed a “disease of stages”, and tumors are believed to progress through a series of “states” before becoming malignant [1]. A case in point is prostate cancer, which is known to be a multistep process [2]. For instance, a patient diagnosed with localized prostate cancer who has had all the tumor surgically removed is considered to remain in the state of “localized disease” until he progresses to a new state. At each state, distinct therapies can be prescribed, and the time spent by the patient in any given state is a measure of the efficacy of the corresponding intervention.

The primary treatments for patients with localized prostate cancer are surgery, radiation therapy, or active surveillance [2], which can be used alone or in combination. For patients who evolve into a state of metastatic disease, standard treatment is hormone therapy in the form of continuous androgen suppression (CAS) [2]. The initial response to CAS is frequently positive, leading to a significant decrease in tumor size; unfortunately, most patients eventually develop resistance and relapse. A generally acceptable mechanism for explaining such relapse is the existence of an androgen-independent cancer cell phenotype that is resistant to secondary endocrine therapy and whose outgrowth leads to tumor recurrence [3].

Intermittent androgen suppression (IAS) therapy is an alternative treatment strategy for delaying or even preventing time to relapse. The goal of IAS is to prevent the existing tumor from progressing into an androgen-independent state. In spite of extensive recent clinical experience with IAS, the design of an ideal protocol for any given patient remains one of the main challenges associated with effectively implementing this therapy [4]. Although clinical trials [5],[6] revealed that the success of IAS ultimately depends on the ability to tailor on and off-treatment schemes to individual patients, defining optimal personalized IAS treatment schemes remains an unsolved problem.

A number of mathematical models have been proposed to explain the progression of prostate cancer in patients who are submitted to hormone therapy. In [3] a model is developed in which prostate tumors are composed of two subpopulations of cancer cells, one that is sensitive to androgen suppression and another that is not, without directly addressing the issue of IAS therapy

design. The evolution of a prostate tumor under IAS is modeled using a hybrid dynamical system approach in [7], and numerical bifurcation analyses are applied to study the effect of different therapy protocols on tumor growth and time to relapse. Various works that extend [3] and [7] have recently been developed, and we briefly review some of them. In [8] a nonlinear model is proposed to explain the competition between different cancer cell subpopulations, while switched ordinary differential equations are used in the model from [9]. In [10] the problem of personalized prostate cancer treatment is formulated as an optimal control problem using a piecewise affine system model. In [4] a feedback control system is used to model the prostate tumor under IAS for patient classification; an extension of this work, in which conditions for patient relapse are derived, is presented in [11].

Although the majority of existing models provide insights into the dynamics of prostate cancer evolution under androgen deprivation therapy, they fail to address the issue of therapy design. Moreover, previous works that suggest optimal treatment schemes by classifying patients into groups have been based on more manageable, albeit less accurate, approaches to nonlinear hybrid dynamical systems. Addressing this limitation, a nonlinear hybrid automaton model was recently developed and δ -reachability analysis was performed to identify patient-specific treatment schemes in [12]. In spite of being in good agreement with published clinical data, this model does not account for noise and fluctuations inherently associated with cell population dynamics and monitoring of clinical data. In contrast, a hybrid model of tumor growth under IAS therapy that incorporates stochastic effects is proposed in [13], but is not used for personalized therapy design.

A first attempt to define optimal personalized IAS therapy schemes using stochastic models of prostate cancer evolution was reported in [14]. Building upon the Infinitesimal Perturbation Analysis (IPA) framework established in [14], here we implement an IPA-driven gradient-based optimization algorithm capable of adaptively adjusting controllable therapy settings so as to improve IAS therapy outcomes. From a practical perspective, the goal of this paper is to set the stage for the use of basic IPA techniques for optimal personalized cancer therapy design. Therefore, we introduce the relevant concepts related to an IPA-based system of optimal cancer therapy design and illustrate its application to the case of advanced prostate cancer.

The remainder of this paper is organized as follows. In Section 2, we present a Stochastic Hybrid Automaton (SHA) model of prostate cancer evolution, based on which we formulate the problem of optimal IAS therapy

design. Section 3 details the derivation of IPA estimators for therapy evaluation and optimization. The IPA estimators are then incorporated into a gradient-based optimization algorithm and simulation results providing interesting insights are thus obtained. Sample results are given in Section 4, and we include final remarks, along with a discussion of the proposed methodology in light of the application performed for IAS therapy of prostate cancer, in Section 5.

2. Problem Formulation

2.1. Stochastic Model of Prostate Cancer Evolution

We consider a system composed of a prostate tumor under IAS therapy, which is modeled as a Stochastic Hybrid Automaton (SHA). Details of the problem formulation are given in [14], but for completeness we include here a condensed description of the SHA modeling framework. We adopt a standard SHA definition [15]:

$$G_h = (Q, X, E, U, f, \phi, Inv, guard, \rho, q_0, x_0) \quad (1)$$

where Q is a set of discrete states; X is a continuous state space; E is a finite set of events; U is a set of admissible controls; f is a vector field, $f : Q \times X \times U \rightarrow X$; ϕ is a discrete state transition function, $\phi : Q \times X \times E \rightarrow Q$; Inv is a set defining an invariant condition (when this condition is violated at some $q \in Q$, a transition must occur); $guard$ is a set defining a guard condition, $guard \subseteq Q \times Q \times X$ (when this condition is satisfied at some $q \in Q$, a transition is allowed to occur); ρ is a reset function, $\rho : Q \times Q \times X \times E \rightarrow X$; q_0 is an initial discrete state; x_0 is an initial continuous state.

In this context, a SHA model of prostate cancer progression can be defined in terms of the following:

1. A discrete state set $Q = \{q^{ON}, q^{OFF}\}$, where q^{ON} (q^{OFF} , respectively) is the on-treatment (off-treatment, respectively) operational mode of the system. Patients undergoing IAS therapy will temporarily stop being medicated once the size of their cancer cell populations decreases by a predetermined desirable amount. Since population sizes are not directly observable, this reduction is estimated in terms of the patient's Prostate-Specific Antigen (PSA) level, a biomarker commonly used for monitoring the outcome of hormone therapy. In this context, therapy is suspended when a patient's PSA level reaches a lower threshold

value, and reinstated once the size of cancer cell populations has increased considerably, i.e., once the patient's PSA level reaches an upper threshold value.

2. A state space $X = \{x_1(t), x_2(t), x_3(t), z_1(t), z_2(t)\}$, defined in terms of the biomarkers commonly monitored during IAS therapy, as well as “clock” state variables that measure the time spent by the system in each discrete state. We assume that prostate tumors are composed of two coexisting subpopulations of cancer cells, Hormone Sensitive Cells (HSCs) and Castration Resistant Cells (CRCs), and thus define a state vector $x(t) = [x_1(t), x_2(t), x_3(t)]$ with $x_i(t) \in \mathbb{R}^+$, such that $x_1(t)$ is the total population of HSCs, $x_2(t)$ is the total population of CRCs, and $x_3(t)$ is the concentration of androgen in the serum. Prostate cancer cells secrete high levels of PSA, hence a common assumption is that the serum PSA concentration can be modeled as a linear combination of the cancer cell subpopulations. It is also frequently assumed that both HSCs and CRCs secrete PSA equivalently [7], and in this work we adopt these assumptions. Finally, we define variable $z_i(t) \in \mathbb{R}^+$, $i = 1, 2$, where $z_1(t)$ ($z_2(t)$, respectively) is the “clock” state variable corresponding to the time when the system is in state q^{ON} (q^{OFF} , respectively), and is reset to zero every time a state transition occurs. Setting $z(t) = [z_1(t), z_2(t)]$, the complete state vector is $[x(t), z(t)]$.
3. An event set $E = \{e_1, e_2\}$, where e_1 corresponds to the condition $[x_1(t) + x_2(t) = \theta_1 \text{ from above}]$ (i.e., $x_1(t^-) + x_2(t^-) > \theta_1$) and e_2 corresponds to the condition $[x_1(t) + x_2(t) = \theta_2 \text{ from below}]$ (i.e., $x_1(t^-) + x_2(t^-) < \theta_2$), where the notation t^- indicates the time instant immediately preceding time t .
4. An admissible control set $U = \{0, 1\}$, such that the control is defined, at any time t , as:

$$u(x(t), z(t)) \equiv \begin{cases} 0 & \text{if } x_1(t) + x_2(t) < \theta_2, q(t) = q^{OFF} \\ 1 & \text{if } x_1(t) + x_2(t) > \theta_1, q(t) = q^{ON} \end{cases} \quad (2)$$

This is a simple form of hysteresis control to ensure that androgen deprivation will be suspended whenever a patient's PSA level drops below a minimum threshold value, and that treatment will resume once the patient's PSA level reaches a maximum threshold value. To this end, IAS therapy is viewed as a controlled process characterized by two parameters: $\theta = [\theta_1, \theta_2] \in \Theta$, with $\theta_1 < \theta_2$, and where $\theta_1 \in$

$[\theta_1^{\min}, \theta_1^{\max}]$ is the lower threshold value of the patient's PSA level, and $\theta_2 \in [\theta_2^{\min}, \theta_2^{\max}]$ is the upper threshold value of the patient's PSA level, with $\theta_1^{\max} < \theta_2^{\min}$. An illustrative representation of such threshold-based IAS therapy scheme is depicted in Figure 1. Simulation driven by clinical data [5],[6] was performed to generate the plot in Figure 1, which shows a typical profile of PSA level variations along several treatment cycles.

Figure 1: Schematic representation of Intermittent Androgen Suppression (IAS) therapy

5. System dynamics. The *continuous (time-driven) dynamics* capture the prostate cancer cell population dynamics, which are defined in terms of their proliferation, apoptosis, and conversion rates. As in [14], we incorporate stochastic effects into the deterministic model from [12] as follows:

$$\begin{aligned} \dot{x}_1(t) = & \left[\frac{\alpha_1}{1+e^{-(x_3(t)-k_1)k_2}} - \frac{\beta_1}{1+e^{-(x_3(t)-k_3)k_4}} - m_1 \left(1 - \frac{x_3(t)}{x_{3,0}} \right) - \lambda_1 \right] \cdot x_1(t) \\ & + \mu_1 + \zeta_1(t) \end{aligned} \quad (3)$$

$$\begin{aligned} \dot{x}_2(t) = & \left[\alpha_2 \left(1 - d \frac{x_3(t)}{x_{3,0}} \right) - \beta_2 \right] x_2(t) \\ & + m_1 \left(1 - \frac{x_3(t)}{x_{3,0}} \right) x_1(t) + \zeta_2(t) \end{aligned} \quad (4)$$

$$\dot{x}_3(t) = \begin{cases} -\frac{x_3(t)}{\sigma} + \mu_3 + \zeta_3(t) & \text{if } x_1(t) + x_2(t) > \theta_1 \\ & \text{and } q(t) = q^{ON} \\ \frac{x_{3,0} - x_3(t)}{\sigma} + \mu_3 + \zeta_3(t) & \text{if } x_1(t) + x_2(t) < \theta_2 \\ & \text{and } q(t) = q^{OFF} \end{cases} \quad (5)$$

$$\begin{aligned} \dot{z}_1(t) &= \begin{cases} 1 & \text{if } q(t) = q^{ON} \\ 0 & \text{otherwise} \end{cases} \\ z_1(t^+) &= 0 \quad \begin{cases} \text{if } x_1(t) + x_2(t) = \theta_1 \\ \text{and } q(t) = q^{ON} \end{cases} \end{aligned} \quad (6)$$

$$\begin{aligned} \dot{z}_2(t) &= \begin{cases} 1 & \text{if } q(t) = q^{OFF} \\ 0 & \text{otherwise} \end{cases} \\ z_2(t^+) &= 0 \quad \begin{cases} \text{if } x_1(t) + x_2(t) = \theta_2 \\ \text{and } q(t) = q^{OFF} \end{cases} \end{aligned} \quad (7)$$

where α_1 and α_2 are the HSC proliferation constant and CRC proliferation constant, respectively; β_1 and β_2 are the HSC apoptosis constant and CRC apoptosis constant, respectively; k_1 through k_4 are HSC proliferation and apoptosis exponential constants; m_1 is the HSC to CRC conversion constant; $x_{3,0}$ corresponds to the patient-specific androgen constant; σ is the androgen degradation constant; λ_1 is the HSC basal degradation rate; μ_1 and μ_3 are the HSC basal production rate and androgen basal production rate, respectively. Finally, $\{\zeta_i(t)\}$, $i = 1, 2, 3$, are stochastic processes which we allow to have arbitrary characteristics and only assume them to be piecewise continuous w.p. 1. The processes $\{\zeta_i(t)\}$, $i = 1, 2$, represent noise and fluctuations inherently associated with cell population dynamics, while $\{\zeta_3(t)\}$ reflects randomness associated with monitoring clinical data, more specifically, with monitoring the patient's androgen level.

It is clear from (3)-(5) that $x_1(t)$ and $x_2(t)$ are dependent on $x_3(t)$, whose dynamics are affected by mode transitions. To make explicit the dependence of $x_1(t)$ and $x_2(t)$ on the discrete state (mode) $q(t)$, we let $\tau_k(\theta)$ be the time of occurrence of the k th event (of any type), and denote the state dynamics over any interevent interval $[\tau_k(\theta), \tau_{k+1}(\theta))$ as

$$\dot{x}_n(t) = f_k^{x_n}(t), \quad \dot{z}_i(t) = f_k^{z_i}(t), \quad n = 1, 2, 3, \quad i = 1, 2$$

We include θ as an argument to stress the dependence of the event times on the controllable parameters, but we will subsequently drop this for ease of notation as long as no confusion arises.

We thus start by assuming $q(t) = q^{ON}$ for $t \in [\tau_k, \tau_{k+1})$. Solving (5) yields, for $t \in [\tau_k, \tau_{k+1})$,

$$\begin{aligned} x_3(t) &= x_3(\tau_k^+) e^{-(t-\tau_k)/\sigma} \\ &\quad + e^{-t/\sigma} \cdot \int_{\tau_k}^t e^{\varepsilon/\sigma} [\mu_3 + \zeta_3(\varepsilon)] d\varepsilon \end{aligned}$$

It is then possible to define, for $t \in [\tau_k, \tau_{k+1})$,

$$\begin{aligned} h^{ON}(t, \tilde{\zeta}_3(t)) &\equiv x_3(\tau_k^+) e^{-(t-\tau_k)/\sigma} \\ &\quad + \mu_3 \sigma [1 - e^{-(t-\tau_k)/\sigma}] + \tilde{\zeta}_3(t) \end{aligned} \tag{8}$$

where, for notational simplicity, we let

$$\tilde{\zeta}_3(t) = \int_{\tau_k}^t e^{-(t-\varepsilon)/\sigma} \zeta_3(\varepsilon) d\varepsilon \tag{9}$$

Next, let $q(t) = q^{OFF}$ for $t \in [\tau_k, \tau_{k+1})$, so that (5) implies that, for $t \in [\tau_k, \tau_{k+1})$,

$$\begin{aligned} x_3(t) &= x_3(\tau_k^+) e^{-(t-\tau_k)/\sigma} \\ &\quad + (\mu_3 \sigma + x_{3,0}) [1 - e^{-(t-\tau_k)/\sigma}] + \tilde{\zeta}_3(t) \end{aligned}$$

Similarly as above, we define, for $t \in [\tau_k, \tau_{k+1})$,

$$\begin{aligned} h^{OFF}(t, \tilde{\zeta}_3(t)) &\equiv x_3(\tau_k^+) e^{-(t-\tau_k)/\sigma} \\ &\quad + (\mu_3 \sigma + x_{3,0}) [1 - e^{-(t-\tau_k)/\sigma}] + \tilde{\zeta}_3(t) \end{aligned} \tag{10}$$

It is then possible to rewrite (5) as follows:

$$x_3(t) = \begin{cases} h^{ON}(t, \tilde{\zeta}_3(t)) & \text{if } q(t) = q^{ON} \\ h^{OFF}(t, \tilde{\zeta}_3(t)) & \text{if } q(t) = q^{OFF} \end{cases}$$

Although we include $\tilde{\zeta}_3(t)$ as an argument in (8) and (10) to stress the dependence on the stochastic process, we will subsequently drop this

for ease of notation as long as no confusion arises. Hence, substituting (8) and (10) into (3)-(4), yields

$$\dot{x}_1(t) = \begin{cases} \left\{ \alpha_1 [1 + \phi_\alpha^{ON}(t)]^{-1} - \beta_1 [1 + \phi_\beta^{ON}(t)]^{-1} \right. \\ \left. + m_1 \left(\frac{h^{ON}(t)}{x_{3,0}} \right) - (m_1 + \lambda_1) \right\} \cdot x_1(t) \\ + \mu_1 + \zeta_1(t) & \text{if } q(t) = q^{ON} \\ \left\{ \alpha_1 [1 + \phi_\alpha^{OFF}(t)]^{-1} - \beta_1 [1 + \phi_\beta^{OFF}(t)]^{-1} \right. \\ \left. + m_1 \left(\frac{h^{OFF}(t)}{x_{3,0}} \right) - (m_1 + \lambda_1) \right\} \cdot x_1(t) \\ + \mu_1 + \zeta_1(t) & \text{if } q(t) = q^{OFF} \end{cases} \quad (11)$$

$$\dot{x}_2(t) = \begin{cases} \left[\alpha_2 \left(1 - d \frac{h^{ON}(t)}{x_{3,0}} \right) - \beta_2 \right] x_2(t) \\ + m_1 \left(1 - \frac{h^{ON}(t)}{x_{3,0}} \right) x_1(t) + \zeta_2(t) & \text{if } q(t) = q^{ON} \\ \left[\alpha_2 \left(1 - d \frac{h^{OFF}(t)}{x_{3,0}} \right) - \beta_2 \right] x_2(t) \\ + m_1 \left(1 - \frac{h^{OFF}(t)}{x_{3,0}} \right) x_1(t) + \zeta_2(t) & \text{if } q(t) = q^{OFF} \end{cases} \quad (12)$$

with

$$\begin{aligned} \phi_\alpha^{ON}(t) &= e^{-(h^{ON}(t)-k_1)k_2} \\ \phi_\beta^{ON}(t) &= e^{-(h^{ON}(t)-k_3)k_4} \\ \phi_\alpha^{OFF}(t) &= e^{-(h^{OFF}(t)-k_1)k_2} \\ \phi_\beta^{OFF}(t) &= e^{-(h^{OFF}(t)-k_3)k_4} \end{aligned}$$

The *discrete (event-driven) dynamics* are dictated by the occurrence of events that cause state transitions. Based on the event set $E = \{e_1, e_2\}$ we have defined, the occurrence of e_1 results in a transition from q^{ON} to q^{OFF} and the occurrence of e_2 results in a transition from q^{OFF} to q^{ON} . The corresponding SHA model of prostate cancer evolution under IAS therapy is shown in Figure 2.

2.2. IAS Therapy Evaluation and Optimization

Within the SHA framework we propose, the problem of personalizing an IAS treatment scheme can be cast as the search for the optimal IAS therapy that satisfies some performance criterion. In this sense, an IAS therapy

can be viewed as a controlled process $u(\theta, t)$ characterized by the parameter vector θ , as in (2), whose effect can be quantified in terms of performance metrics of the form $J[u(\theta, t)]$. Although it is clearly infeasible to evaluate $J[u(\theta, t)]$ over all possible values of θ , there exist very efficient ways to perform the search for an optimal IAS therapy within a stochastic hybrid system framework. In particular, Perturbation Analysis (PA) is a methodology to efficiently estimate the sensitivity of the system’s performance with respect to θ (note that when θ is a real-valued scalar, this amounts to estimating the derivative $dJ/d\theta$). This is accomplished by extracting data from a sample path (simulated or actual) of the observed system based on which an unbiased estimate of $dJ/d\theta$ can indeed be obtained. The attractive feature of PA is that the resulting estimates are extracted from a *single* sample path in a non-intrusive manner and the computational cost of doing so is, in most cases of interest, minimal [15]. This is in contrast to the conventional finite difference estimate of $dJ/d\theta$ obtained through $[J(\theta + \Delta) - J(\theta)] / \Delta$. Thus, for a vector θ of dimension N , estimating the gradient $\nabla J(\theta)$ requires a single sample path (with some overhead) instead of $N + 1$ sample paths. The simplest family of PA estimators is Infinitesimal Perturbation Analysis (IPA), which has been shown to provide *unbiased* gradient estimates [16] for virtually arbitrary stochastic hybrid systems.

We further emphasize that alternative techniques, such as genetic algorithms (GA) or other stochastic optimization methods that seek global optimality (e.g., [17],[18],[19],[20],[21]), require either that the value of $J(\theta)$ be known or that it be estimated through repeated simulation. Moreover, we caution that applying such approaches to cancer therapy design calls for a strictly off-line and extremely time-consuming implementation. We elaborate on this caveat in what follows: first, $J(\theta)$ would have to be estimated, which requires accurate knowledge of the stochastic processes; this is something IPA does not need, as explained above. Second, executing a GA is a very time consuming task in itself, e.g., in selecting the proper parameter values to execute such an algorithm [17]. Given that simple gradient-based algorithms work very well, as our results show (see Section 4), applying sophisticated stochastic global optimization methods such as GAs to solve the specific optimization problem we describe adds unnecessary complexity, even if the stochastic processes involved were known. Regarding the precise stochastic process models used, to the best of our knowledge, all other approaches require full knowledge of such models. However, we emphasize that IPA is robust to the stochastic process models used. It is important to clarify that

this does not imply that IPA is independent of these processes, only that the development of our IPA-based methodology is unaffected by the choice of such processes because their information is contained in the event time data (which are directly observable given that IPA is a data-driven method).

Furthermore, since our goal is to design personalized therapies, it is important to drive our threshold adaptation based on *patient-specific data which are directly observable* (though we only emulate this in this work, since it is currently very difficult to access individual patient data). However, we envision this happening in the near future, so that with every new therapy cycle we would have more personalized data, thus allowing us to properly adapt thresholds to a specific patient. Finally, we highlight that an additional benefit deriving from the use of IPA is the ability to obtain sensitivity estimates with respect to various system parameters; we further discuss this fact in another paper [22].

Returning to our problem of personalized cancer therapy design, for the SHA model of prostate cancer evolution we consider here, our goal is to estimate the effects of different therapies $u(\theta, t)$ by adapting IPA estimators of the form $dJ[u(\theta, t)]/d\theta$, and to ultimately design optimal therapy schemes by solving problems of the form $\min_{\theta \in \Theta} J[u(\theta, t)]$. As in [14], here we make use of a sample function defined in terms of complementary measures of therapy success. In particular, we consider the most adequate IAS treatment schemes to be those that (i) ensure PSA levels are kept as low as possible; (ii) reduce the frequency of on and off-treatment cycles. From a practical perspective, (i) translates into the ability to successfully keep the size of cancer cell populations under control, which is directly influenced by the duration of the on and off-treatment periods. On the other hand, (ii) aims at reducing the duration of on-treatment periods, thus decreasing the exposure of patients to medication and their side effects, and consequently improving the patients' quality of life throughout the treatment. Clearly there is a trade-off between keeping tumor growth under control and the cost associated with the corresponding IAS therapy. The latter is related to the duration of the therapy and could potentially include fixed set up costs incurred when therapy is reinstated. For simplicity, we disconsider fixed set up costs and take (ii) to be linearly proportional to the length of the on-treatment cycles. Hence, we define our sample function as the sum of the average PSA level and the average duration of an on-treatment cycle over a fixed time interval $[0, T]$. We also take into account that it may be desirable to design a therapy scheme which favors (i) over (ii) (or vice-versa) and thus associate weight

W with (i) and $1 - W$ with (ii), where $0 \leq W \leq 1$. Finally, to ensure that the trade-off between (i) and (ii) is captured appropriately, we normalize our sample function: we divide (i) by the value of the patient's PSA level at the start of the first on-treatment cycle (PSA_{init}), and normalize (ii) by T .

Recall that the total population size of prostate cancer cells is assumed to reflect the serum PSA concentration, and that we have defined clock variables which measure the time elapsed in each of the treatment modes, so that our sample function can be written as

$$L(\theta, x(0), z(0), T) = \frac{1}{T} \left[W \int_0^T \left[\frac{x_1(\theta, t) + x_2(\theta, t)}{PSA_{init}} \right] dt + (1 - W) \int_0^T \frac{z_1(t)}{T} dt \right] \quad (13)$$

where $x(0)$ and $z(0)$ are given initial conditions. We can then define the overall performance metric as

$$J(\theta, x(0), z(0), T) = E[L(\theta, x(0), z(0), T)] \quad (14)$$

Hence, the problem of determining the optimal IAS therapy can be formulated as

$$\min_{\theta \in \Theta} E[L(\theta, x(0), z(0), T)] \quad (15)$$

We note that it is not possible to derive a closed-form expression of $J(\theta, x(0), z(0), T)$ without imposing limitations on the processes $\{\zeta_i(t)\}$, $i = 1, 2, 3$. Nevertheless, by assuming only that $\zeta_i(t)$, $i = 1, 2, 3$, are piecewise continuous w.p. 1, we can successfully apply the IPA methodology developed for general SHS in [16] and obtain an estimate of $\nabla J(\theta)$ by evaluating the sample gradient $\nabla L(\theta)$. The knowledge of such gradient values can then be used to improve current operating conditions or to compute an optimal θ^* through an iterative optimization algorithm of the form

$$\theta_{i,k+1} = \theta_{i,k} - \rho_k H_{i,k}(\theta_k, x(0), T, \omega_k) \quad (16)$$

where θ_i , $i = 1, 2$ are the therapy thresholds, ρ_k is the step size at the k th iteration, $k = 1, \dots$, and ω_k denotes a sample path from which data are extracted and used to compute $H_{i,k}(\theta_k, x(0), T, \omega_k)$, which is an estimate of $dJ(\theta)/d\theta_i$. We will assume that the derivatives $dL(\theta)/d\theta_i$ exist w.p. 1 for all $\theta_i \in \mathbb{R}^+$. It is also simple to verify that $L(\theta)$ is Lipschitz continuous for $\theta_i \in \mathbb{R}^+$. We will further assume that $\{\zeta_i(t)\}$, $i = 1, 2, 3$, are stationary random processes over $[0, T]$ and that no two events can occur at the same

time w.p. 1. Under these conditions, it has been shown in [16] that $dL(\theta)/d\theta_i$ is an *unbiased* estimator of $dJ(\theta)/d\theta_i$, $i = 1, 2$. Hence, our goal is to compute the sample gradient $\nabla L(\theta)$ using data extracted from a sample path of the system (e.g., by simulating a sample path of our SHA model using clinical data), and use this value as an estimate of $\nabla J(\theta)$.

3. Infinitesimal Perturbation Analysis

For the sake of completeness, we provide here a brief overview of the IPA framework developed for stochastic hybrid systems in [16]. Consider a sample path of the system over $[0, T]$ and denote the time of occurrence of the k th event (of any type) by $\tau_k(\theta)$, where θ corresponds to the control parameter of interest. Although we use the notation $\tau_k(\theta)$ to stress the dependency of the event time on the control parameter, we will subsequently use τ_k to indicate the time of occurrence of the k th event where no confusion arises. In order to further simplify notation, we shall denote the state and event time derivatives with respect to parameter θ as $x'(t) \equiv \frac{\partial x(\theta, t)}{\partial \theta}$ and $\tau'_k \equiv \frac{\partial \tau_k}{\partial \theta}$, respectively, for $k = 1, \dots, N$. Additionally, considering that the system is at some discrete mode during an interval $[\tau_k, \tau_{k+1})$, we will denote its time-driven dynamics over such interval as $f_k(x, \theta, t)$. It is shown in [16] that the state derivative satisfies

$$\frac{d}{dt}x'(t) = \frac{\partial f_k(t)}{\partial x}x'(t) + \frac{\partial f_k(t)}{\partial \theta} \quad (17)$$

with the following boundary condition:

$$x'(\tau_k^+) = x'(\tau_k^-) + [f_{k-1}(\tau_k^-) - f_k(\tau_k^+)] \cdot \tau'_k \quad (18)$$

when $x(\theta, t)$ is continuous in t at $t = \tau_k$. The notation τ_k^+ (τ_k^- , respectively) indicates the time instant immediately succeeding (preceding, respectively) event time τ_k . Otherwise,

$$x'(\tau_k^+) = \frac{d\rho(q, q', x, e)}{d\theta} \quad (19)$$

where $\rho(q, q', x, e)$ is the reset function defined in (1).

Knowledge of τ'_k is, therefore, needed in order to evaluate (18). Following the framework in [16], the expressions of the event time derivative depend on the type of event that takes place at τ_k and are given as follows:

- (i) *Exogenous event.* This type of event causes a discrete state transition which is independent of parameter θ and, as a result, $\tau'_k = 0$.
- (ii) *Endogenous event.* In this case, there exists a continuously differentiable function $g_k : \mathbb{R}^n \times \Theta \rightarrow \mathbb{R}$ such that $\tau_k = \min \{t > \tau_{k-1} : g_k(x(\theta, t), \theta) = 0\}$, which leads to

$$\tau'_k = - \left[\frac{\partial g_k}{\partial x} \cdot f_{k-1}(\tau_k^-) \right]^{-1} \cdot \left(\frac{\partial g_k}{\partial \phi} + \frac{\partial g_k}{\partial x} \cdot x'(\tau_k^-) \right) \quad (20)$$

where $\frac{\partial g_k}{\partial x} \cdot f_{k-1}(\tau_k^-) \neq 0$.

- (iii) *Induced event.* Such an event is triggered by the occurrence of another event at time $\tau_m \leq \tau_k$ and the expression of τ'_k depends on the event time derivative of the triggering event (τ'_m) (details can be found in [16]).

Thus, IPA captures how changes in θ affect the event times and the state of the system. Since interesting performance metrics are usually expressed in terms of τ_k and $x(t)$, IPA can ultimately be used to infer the effect that a perturbation in θ will have on such metrics. We end this overview by returning to our problem of personalized prostate cancer therapy design and thus defining the derivatives of the states $x_n(\theta, t)$ and $z_j(\theta, t)$ and event times $\tau_k(\theta)$ with respect to θ_i , $i, j = 1, 2$, $n = 1, 2, 3$, as follows:

$$x'_{n,i}(t) \equiv \frac{\partial x_n(\theta, t)}{\partial \theta_i}, \quad z'_{j,i}(t) \equiv \frac{\partial z_j(\theta, t)}{\partial \theta_i}, \quad \tau'_{k,i} \equiv \frac{\partial \tau_k(\theta)}{\partial \theta_i} \quad (21)$$

In what follows, we derive the IPA state and event time derivatives for the events identified in our SHA model of prostate cancer progression.

3.1. State and Event Time Derivatives

We proceed by analyzing the state evolution of our SHA model of prostate cancer progression considering each of the states (q^{ON} and q^{OFF}) and events (e_1 and e_2) therein defined.

1. *The system is in state q^{ON} over interevent time interval $[\tau_k, \tau_{k+1})$.* Using (17) for $x_1(t)$, we obtain, for $i = 1, 2$,

$$\begin{aligned} \frac{d}{dt} x'_{1,i}(t) &= \frac{\partial f_k^{x_1}(t)}{\partial x_1} x'_1(t) + \frac{\partial f_k^{x_1}(t)}{\partial x_2} x'_2(t) \\ &\quad + \frac{\partial f_k^{x_1}(t)}{\partial z_1} z'_1(t) + \frac{\partial f_k^{x_1}(t)}{\partial z_2} z'_2(t) + \frac{\partial f_k^{x_1}(t)}{\partial \theta_i} \end{aligned}$$

From (11), we have $\frac{\partial f_k^{x_1}(t)}{\partial x_2} = \frac{\partial f_k^{x_1}(t)}{\partial z_i} = \frac{\partial f_k^{x_1}(t)}{\partial \theta_i} = 0$, $i = 1, 2$, and

$$\begin{aligned} \frac{\partial f_k^{x_1}(t)}{\partial x_1} &= \alpha_1 [1 + \phi_\alpha^{ON}(t)]^{-1} - \beta_1 [1 + \phi_\beta^{ON}(t)]^{-1} \\ &\quad - m_1 \left(1 - \frac{h^{ON}(t)}{x_{3,0}}\right) - \lambda_1 \end{aligned}$$

It is thus simple to verify that solving (17) for $x'_{1,i}(t)$ yields, for $i = 1, 2$,

$$x'_{1,i}(t) = x'_{1,i}(\tau_k^+) e^{A(t)}, \quad t \in [\tau_k, \tau_{k+1}) \quad (22)$$

with

$$\begin{aligned} A(t) &\equiv \int_{\tau_k}^t \left[\frac{\alpha_1}{1 + \phi_\alpha^{ON}(\varepsilon)} - \frac{\beta_1}{1 + \phi_\beta^{ON}(\varepsilon)} \right] d\varepsilon \\ &\quad - \int_{\tau_k}^t \frac{m_1}{x_{3,0}} h^{ON}(\varepsilon) d\varepsilon - (m_1 + \lambda_1)(t - \tau_k) \end{aligned} \quad (23)$$

In particular, at τ_{k+1}^- :

$$x'_{1,i}(\tau_{k+1}^-) = x'_{1,i}(\tau_k^+) e^{A(\tau_k)} \quad (24)$$

where $A(\tau_k)$ is given from (23).

Similarly for $x_2(t)$, we have from (12) that $\frac{\partial f_k^{x_2}(t)}{\partial z_i} = \frac{\partial f_k^{x_2}(t)}{\partial \theta_i} = 0$, $i = 1, 2$, and

$$\begin{aligned} \frac{\partial f_k^{x_2}(t)}{\partial x_1} &= m_1 \left(1 - \frac{h^{ON}(t)}{x_{3,0}}\right) \\ \frac{\partial f_k^{x_2}(t)}{\partial x_2} &= \alpha_2 \left(1 - d \frac{h^{ON}(t)}{x_{3,0}}\right) - \beta_2 \end{aligned}$$

Combining the last two equations and solving for $x'_{2,i}(t)$ yields, for $i = 1, 2$ and $t \in [\tau_k, \tau_{k+1})$,

$$x'_{2,i}(t) = x'_{2,i}(\tau_k^+) e^{B_1(t)} + B_2(t, x'_{1,i}(\tau_k^+), A(t)) \quad (25)$$

with

$$\begin{aligned} B_1(t) &\equiv \int_{\tau_k}^t \left[\alpha_2 \left(1 - d \frac{h^{ON}(\varepsilon)}{x_{3,0}}\right) - \beta_2 \right] d\varepsilon \\ B_2(\cdot) &\equiv e^{B_1(t)} \int_{\tau_k}^t G_1(\varepsilon, \tau_k) e^{-B_1(\varepsilon)} d\varepsilon \end{aligned} \quad (26)$$

where $G_1(t, \tau_k) = m_1 \left(1 - \frac{h^{ON}(t)}{x_{3,0}}\right) x'_{1,i}(\tau_k^+) e^{A(t)}$, $t \in [\tau_k, \tau_{k+1})$.

In particular, at τ_{k+1}^- :

$$x'_{2,i}(\tau_{k+1}^-) = x'_{2,i}(\tau_k^+) e^{B_1(\tau_k)} + B_2(\tau_k, x'_{1,i}(\tau_k^+), A(\tau_k)) \quad (27)$$

where $B_1(\tau_k)$ and $B_2(\tau_k, x'_{1,i}(\tau_k^+), A(\tau_k))$ are given from (26).

Finally, for the “clock” state variable, from (6)-(7) we have $\frac{\partial f_k^{z_i}(t)}{\partial x_n} = \frac{\partial f_k^{z_i}(t)}{\partial z_i} = \frac{\partial f_k^{z_i}(t)}{\partial \theta_i} = 0$, $n, i = 1, 2$, so that $\frac{d}{dt} z'_{j,i}(t) = 0$, $j, i = 1, 2$, for $t \in [\tau_k, \tau_{k+1})$. Hence, $z'_{j,i}(t) = z'_{j,i}(\tau_k^+)$, $j, i = 1, 2$ and $t \in [\tau_k, \tau_{k+1})$.

2. *The system is in state q^{OFF} over interevent time interval $[\tau_k, \tau_{k+1})$.*

Starting with $x_1(t)$, based on (11) we once again have $\frac{\partial f_k^{x_1}(t)}{\partial x_2} = \frac{\partial f_k^{x_1}(t)}{\partial z_i} = \frac{\partial f_k^{x_1}(t)}{\partial \theta_i} = 0$, $i = 1, 2$, but now

$$\begin{aligned} \frac{\partial f_k^{x_1}(t)}{\partial x_1} &= \alpha_1 [1 + \phi_\alpha^{OFF}(t)]^{-1} - \beta_1 [1 + \phi_\beta^{OFF}(t)]^{-1} \\ &\quad - m_1 \left(1 - \frac{h^{OFF}(t)}{x_{3,0}} \right) - \lambda_1 \end{aligned}$$

Therefore, (17) implies that, for $i = 1, 2$:

$$x'_{1,i}(t) = x'_{1,i}(\tau_k^+) e^{C(t)}, \quad t \in [\tau_k, \tau_{k+1}) \quad (28)$$

with

$$\begin{aligned} C(t) &\equiv \int_{\tau_k}^t \left[\frac{\alpha_1}{1 + \phi_\alpha^{OFF}(\varepsilon)} - \frac{\beta_1}{1 + \phi_\beta^{OFF}(\varepsilon)} \right] d\varepsilon \\ &\quad - \int_{\tau_k}^t \frac{m_1}{x_{3,0}} h^{OFF}(\varepsilon) d\varepsilon - (m_1 + \lambda_1)(t - \tau_k) \end{aligned} \quad (29)$$

In particular, at τ_{k+1}^- :

$$x'_{1,i}(\tau_{k+1}^-) = x'_{1,i}(\tau_k^+) e^{C(\tau_k)} \quad (30)$$

where $C(\tau_k)$ is given from (29).

Similarly for $x_2(t)$, we have

$$\begin{aligned} \frac{\partial f_k^{x_2}(t)}{\partial x_1} &= m_1 \left(1 - \frac{h^{OFF}(t)}{x_{3,0}} \right) \\ \frac{\partial f_k^{x_2}(t)}{\partial x_2} &= \alpha_2 \left(1 - d \frac{h^{OFF}(t)}{x_{3,0}} \right) - \beta_2 \end{aligned}$$

It is thus straightforward to verify that (17) yields, for $i = 1, 2$ and $t \in [\tau_k, \tau_{k+1})$,

$$x'_{2,i}(t) = x'_{2,i}(\tau_k^+) e^{D_1(t)} + D_2(t, x'_{1,i}(\tau_k^+), C(t)) \quad (31)$$

with

$$\begin{aligned} D_1(t) &\equiv \int_{\tau_k}^t \left[\alpha_2 \left(1 - d \frac{h^{OFF}(\varepsilon)}{x_{3,0}} \right) - \beta_2 \right] d\varepsilon \\ D_2(\cdot) &\equiv e^{D_1(t)} \int_{\tau_k}^t G_2(\varepsilon, \tau_k) e^{-D_1(\varepsilon)} d\varepsilon \end{aligned} \quad (32)$$

where $G_2(t, \tau_k) = m_1 \left(1 - \frac{h^{OFF}(t)}{x_{3,0}}\right) x'_{1,i}(\tau_k^+) e^{C(t)}$, $t \in [\tau_k, \tau_{k+1})$.

In particular, at τ_{k+1}^- :

$$x'_{2,i}(\tau_{k+1}^-) = x'_{2,i}(\tau_k^+) e^{D_1(\tau_k)} + D_2(\tau_k, x'_{1,i}(\tau_k^+), C(\tau_k)) \quad (33)$$

where $D_1(\tau_k)$ and $D_2(\tau_k, x'_{1,i}(\tau_k^+), C(\tau_k))$ are given from (32).

Finally, for the “clock” state variable, based on (6)-(7) we once again have $\frac{\partial f_k^{z_i}(t)}{\partial x_n} = \frac{\partial f_k^{z_i}(t)}{\partial z_i} = \frac{\partial f_k^{z_i}(t)}{\partial \theta_i} = 0$, $n, i = 1, 2$, so that $\frac{d}{dt} z'_{j,i}(t) = 0$, $j, i = 1, 2$, for $t \in [\tau_k, \tau_{k+1})$. As a result, $z'_{j,i}(t) = z'_{j,i}(\tau_k^+)$, $j, i = 1, 2$ and $t \in [\tau_k, \tau_{k+1})$.

3. A state transition from q^{ON} to q^{OFF} occurs at time τ_k . This necessarily implies that event e_1 took place at time τ_k , i.e., $q(t) = q^{ON}$, $t \in [\tau_{k-1}, \tau_k)$ and $q(t) = q^{OFF}$, $t \in [\tau_k, \tau_{k+1})$. From (18) we have, for $i = 1, 2$,

$$x'_{1,i}(\tau_k^+) = x'_{1,i}(\tau_k^-) + [f_k^{x_1}(\tau_k^-) - f_{k+1}^{x_1}(\tau_k^+)] \cdot \tau'_{k,i} \quad (34)$$

and

$$x'_{2,i}(\tau_k^+) = x'_{2,i}(\tau_k^-) + [f_k^{x_2}(\tau_k^-) - f_{k+1}^{x_2}(\tau_k^+)] \cdot \tau'_{k,i} \quad (35)$$

where $f_k^{x_1}(\tau_k^-) - f_{k+1}^{x_1}(\tau_k^+)$ and $f_k^{x_2}(\tau_k^-) - f_{k+1}^{x_2}(\tau_k^+)$ ultimately depend on $h^{ON}(\tau_k^-)$ and $h^{OFF}(\tau_k^+)$. Evaluating $h^{ON}(\tau_k^-)$ from (8) over the appropriate time interval results in

$$\begin{aligned} h^{ON}(\tau_k^-) &= x_3(\tau_{k-1}^+) e^{-(\tau_k - \tau_{k-1})/\sigma} \\ &\quad + \mu_3 \sigma [1 - e^{-(\tau_k - \tau_{k-1})/\sigma}] + \tilde{\zeta}_3(\tau_k) \end{aligned}$$

and it follows directly from (10) that $h^{OFF}(\tau_k^+) = x_3(\tau_k^+)$. Moreover, by continuity of $x_n(t)$ (due to conservation of mass), $x_n(\tau_k^+) = x_n(\tau_k^-)$, $n = 1, 2$. Also, since we have assumed that $\{\zeta_i(t)\}$, $i = 1, 2, 3$, is piecewise continuous w.p.1 and that no two events can occur at the same time w.p.1, $\zeta_i(\tau_k^-) = \zeta_i(\tau_k^+)$, $i = 1, 2, 3$. Hence, for $x_1(t)$, evaluating $\Delta_f^1(\tau_k) \equiv f_k^{x_1}(\tau_k^-) - f_{k+1}^{x_1}(\tau_k^+)$ yields

$$\begin{aligned} \Delta_f^1(\tau_k, \zeta_3(\tau_k)) &= \left\{ \alpha_1 [1 + \phi_\alpha^{ON}(\tau_k^-)]^{-1} \right. \\ &\quad - \alpha_1 [1 + \phi_\alpha^{OFF}(\tau_k^+)]^{-1} - \beta_1 [1 + \phi_\beta^{ON}(\tau_k^-)]^{-1} \\ &\quad + \beta_1 [1 + \phi_\beta^{OFF}(\tau_k^+)]^{-1} \\ &\quad \left. + \frac{m_1}{x_{3,0}} [h^{ON}(\tau_k^-) - x_3(\tau_k)] \right\} \cdot x_1(\tau_k) \end{aligned} \quad (36)$$

Finally, the term $\tau'_{k,i}$, which corresponds to the event time derivative with respect to θ_i at event time τ_k , is determined using (20), as detailed in (41) later.

A similar analysis applies to $x_2(t)$, so that $f_k^{x_2}(\tau_k^-)$ and $f_{k+1}^{x_2}(\tau_k^+)$ ultimately depend on $h^{ON}(\tau_k^-)$ and $h^{OFF}(\tau_k^+)$, respectively. Hence, evaluating $\Delta_f^2(\tau_k) \equiv f_k^{x_2}(\tau_k^-) - f_{k+1}^{x_2}(\tau_k^+)$ from (12) yields

$$\begin{aligned} \Delta_f^2(\tau_k, \zeta_3(\tau_k)) &= \frac{\alpha_2 d}{x_{3,0}} [x_3(\tau_k) - h^{ON}(\tau_k^-)] \cdot x_2(\tau_k) \\ &\quad - \frac{m_1}{x_{3,0}} [h^{ON}(\tau_k^-) - x_3(\tau_k)] \cdot x_1(\tau_k) \end{aligned} \quad (37)$$

In the case of the “clock” state variable, $z_1(t)$ is discontinuous in t at $t = \tau_k$, while $z_2(t)$ is continuous. Hence, based on (19) and (6), we have that $z'_{1,i}(\tau_k^+) = 0$. From (18) and (7), it is straightforward to verify that $z'_{2,i}(\tau_k^+) = z'_{2,i}(\tau_k^-) - \tau'_{k,i}$, $i = 1, 2$.

4. *A state transition from q^{OFF} to q^{ON} occurs at time τ_k .* This necessarily implies that event e_2 took place at time τ_k , i.e., $q(t) = q^{OFF}$, $t \in [\tau_{k-1}, \tau_k)$ and $q(t) = q^{ON}$, $t \in [\tau_k, \tau_{k+1})$. The same reasoning as above holds, so that (34)-(35) also apply. For $x_1(t)$, $f_k^{x_1}(\tau_k^-) - f_{k+1}^{x_1}(\tau_k^+)$ can be evaluated from (11) and ultimately depends on $h^{OFF}(\tau_k^-)$ and $h^{ON}(\tau_k^+)$. Evaluating $h^{OFF}(\tau_k^-)$ from (10) over the appropriate time interval results in

$$\begin{aligned} h^{OFF}(\tau_k^-) &= x_3(\tau_{k-1}^+) e^{-(\tau_k - \tau_{k-1})/\sigma} \\ &\quad + (\mu_3 \sigma + x_{3,0}) [1 - e^{-(\tau_k - \tau_{k-1})/\sigma}] + \tilde{\zeta}_3(\tau_k) \end{aligned}$$

and it follows directly from (8) that $h^{ON}(\tau_k^+) = x_3(\tau_k^+)$.

As in the previous case, continuity due to conservation of mass applies, so that evaluating $\Delta_f^1(\tau_k) \equiv f_k^{x_1}(\tau_k^-) - f_{k+1}^{x_1}(\tau_k^+)$ yields

$$\begin{aligned} \Delta_f^1(\tau_k, \zeta_3(\tau_k)) &= \left\{ \alpha_1 [1 + \phi_\alpha^{OFF}(\tau_k^-)]^{-1} \right. \\ &\quad - \alpha_1 [1 + \phi_\alpha^{ON}(\tau_k^+)]^{-1} - \beta_1 [1 + \phi_\beta^{OFF}(\tau_k^-)]^{-1} \\ &\quad + \beta_1 [1 + \phi_\beta^{ON}(\tau_k^+)]^{-1} \\ &\quad \left. + \frac{m_1}{x_{3,0}} [h^{OFF}(\tau_k^-) - x_3(\tau_k)] \right\} \cdot x_1(\tau_k) \end{aligned} \quad (38)$$

Similarly for $x_2(t)$, by evaluating $\Delta_f^2(\tau_k) \equiv f_k^{x_2}(\tau_k^-) - f_{k+1}^{x_2}(\tau_k^+)$ from (12), and making the appropriate simplifications due to continuity, we obtain

$$\begin{aligned} \Delta_f^2(\tau_k, \zeta_3(\tau_k)) &= \frac{\alpha_2 d}{x_{3,0}} [x_3(\tau_k) - h^{OFF}(\tau_k^-)] \cdot x_2(\tau_k) \\ &\quad - \frac{m_1}{x_{3,0}} [h^{OFF}(\tau_k^-) - x_3(\tau_k)] \cdot x_1(\tau_k) \end{aligned} \quad (39)$$

In the case of the “clock” state variable, $z_1(t)$ is continuous in t at $t = \tau_k$, while $z_2(t)$ is discontinuous. As a result, based on (18) and (6), we have that $z'_{1,i}(\tau_k^+) = z'_{1,i}(\tau_k^-) - \tau'_{k,i}$. From (19) and (7), it is simple to verify that $z'_{2,i}(\tau_k^+) = 0$, $i = 1, 2$.

Note that, since $z'_{j,i}(t) = z'_{j,i}(\tau_k^+)$, $t \in [\tau_k, \tau_{k+1})$, we will have that $z'_{j,i}(\tau_k^-) = z'_{j,i}(\tau_{k-1}^+)$, $j, i = 1, 2$. Moreover, the sample path of our SHA consists of a sequence of alternating e_1 and e_2 events, which implies that $z'_{1,i}(\tau_k^-) = 0$ if event e_1 occurred at τ_{k-1} , while $z'_{2,i}(\tau_k^-) = 0$ if event e_2 occurred at τ_{k-1} . Then, adopting the notation $p, \bar{p} = \{1, 2\}$ such that $p + \bar{p} = 3$, we have:

$$z'_{p,i}(\tau_k^+) = \begin{cases} -\tau'_{k,i} & \text{if event } e_{\bar{p}} \text{ occurs at } \tau_k \\ 0 & \text{otherwise} \end{cases} \quad (40)$$

We now proceed with a general result which applies to all events defined for our SHA model. We denote the time of occurrence of the j th state transition by τ_j , define its derivative with respect to the control parameters as $\tau'_{j,i} \equiv \frac{\partial \tau_j}{\partial \theta_i}$, $i = 1, 2$, and also define $f_j^{x_n}(\tau_j) \equiv \dot{x}_n(\tau_j)$, $n = 1, 2$.

Lemma 1. *When an event e_p , $p = 1, 2$, occurs, the derivative $\tau'_{j,i}$, $i = 1, 2$, of state transition times τ_j , $j = 1, 2, \dots$ with respect to the control parameters θ_i , $i = 1, 2$, satisfies:*

$$\tau'_{j,i} = \begin{cases} \frac{1 - x'_{1,i}(\tau_j^-) - x'_{2,i}(\tau_j^-)}{f_{j-1}^{x_1}(\tau_j^-) + f_{j-1}^{x_2}(\tau_j^-)} & \text{if event } e_1 \text{ occurs and } i = 1 \\ & \text{or event } e_2 \text{ occurs and } i = 2 \\ \frac{-x'_{1,i}(\tau_j^-) - x'_{2,i}(\tau_j^-)}{f_{j-1}^{x_1}(\tau_j^-) + f_{j-1}^{x_2}(\tau_j^-)} & \text{if event } e_1 \text{ occurs and } i = 2 \\ & \text{or event } e_2 \text{ occurs and } i = 1 \end{cases} \quad (41)$$

PROOF. We begin with an occurrence of event e_1 which causes a transition from state q^{ON} to state q^{OFF} at time τ_j . This implies that $g_j(x, \theta) = x_1 + x_2 - \theta_1 = 0$. As a result, $\frac{\partial g_k}{\partial x_1} = \frac{\partial g_k}{\partial x_2} = 1$, $\frac{\partial g_k}{\partial x_3} = \frac{\partial g_k}{\partial z_i} = \frac{\partial g_k}{\partial \theta_2} = 0$, $i = 1, 2$, and $\frac{\partial g_k}{\partial \theta_1} = -1$, and it is simple to verify that (41) follows from (20).

Next, consider event e_2 at time τ_j , leading to a transition from state q^{OFF} to state q^{ON} . In this case, $g_j(x, \theta) = x_1 + x_2 - \theta_2 = 0$, so that $\frac{\partial g_k}{\partial x_1} = \frac{\partial g_k}{\partial x_2} = 1$, $\frac{\partial g_k}{\partial x_3} = \frac{\partial g_k}{\partial z_i} = \frac{\partial g_k}{\partial \theta_1} = 0$, $i = 1, 2$, and $\frac{\partial g_k}{\partial \theta_2} = -1$. Substituting into (20) once again yields (41).

We note that the numerator in (41) is determined using (24) and (27) if $q(\tau_j^-) = q^{ON}$, or (30) and (33) if $q(\tau_j^-) = q^{OFF}$. Moreover, the denominator in (41) is computed using (11)-(12) and it is simple to verify that, if event e_1 takes place at time τ_j ,

$$\begin{aligned} f_{j-1}^{x_1}(\tau_j^-) + f_{j-1}^{x_2}(\tau_j^-) &= \alpha_1 [1 + \phi_\alpha^{ON}(\tau_j^-)]^{-1} \cdot x_1(\tau_j) \\ &\quad - \left\{ \beta_1 [1 + \phi_\beta^{ON}(\tau_j^-)]^{-1} + \lambda_1 \right\} \cdot x_1(\tau_j) + \mu_1 \\ &\quad + \left[\alpha_2 \left(1 - d \frac{h^{ON}(\tau_j^-)}{x_{3,0}} \right) - \beta_2 \right] \cdot x_2(\tau_j) \\ &\quad + \zeta_1(\tau_j) + \zeta_2(\tau_j) \end{aligned} \quad (42)$$

and, if event e_2 takes place at time τ_j ,

$$\begin{aligned} f_{j-1}^{x_1}(\tau_j^-) + f_{j-1}^{x_2}(\tau_j^-) &= \alpha_1 [1 + \phi_\alpha^{OFF}(\tau_j^-)]^{-1} \cdot x_1(\tau_j) \\ &\quad - \left\{ \beta_1 [1 + \phi_\beta^{OFF}(\tau_j^-)]^{-1} + \lambda_1 \right\} \cdot x_1(\tau_j) + \mu_1 \\ &\quad + \left[\alpha_2 \left(1 - d \frac{h^{OFF}(\tau_j^-)}{x_{3,0}} \right) - \beta_2 \right] \cdot x_2(\tau_j) \\ &\quad + \zeta_1(\tau_j) + \zeta_2(\tau_j) \end{aligned} \quad (43)$$

We now proceed to present the expression of the cost derivative corresponding to the performance metric defined in (13).

3.2. Cost Derivative

Let us denote the total number of on and off-treatment periods (complete or incomplete) in $[0, T]$ by K_T . Also let ξ_k denote the start of the k^{th} period and η_k denote the end of the k^{th} period (of either type). Finally, let $M_T = \lfloor \frac{K_T}{2} \rfloor$ be the total number of complete on-treatment periods, and Δ_m^{ON} denote the duration of the m^{th} complete on-treatment period, where clearly

$$\Delta_m^{ON} \equiv \eta_m - \xi_m, \quad m = 1, 2, \dots$$

Theorem 1. *The derivative of the sample function $L(\theta)$ with respect to the control parameters satisfies:*

$$\begin{aligned} \frac{dL(\theta)}{d\theta_i} &= \frac{W}{T} \sum_{k=1}^{K_T} \int_{\xi_k}^{\eta_k} \left[\frac{x'_{1,i}(\theta, t) + x'_{2,i}(\theta, t)}{PSA_{init}} \right] dt \\ &\quad + \frac{(1-W)}{T} \sum_{m=1}^{M_T} \frac{\Delta_m^{ON}}{T} \cdot (\eta'_{m,i} - \xi'_{m,i}) \\ &\quad - \frac{(1-W)}{T} \mathbf{1}[K_T \text{ is odd}] \cdot \xi'_{M_T+1,i} \cdot \left(\frac{T - \xi_{M_T+1}}{T} \right) \end{aligned} \quad (44)$$

where $\mathbf{1}[\cdot]$ is the usual indicator function and PSA_{init} is the value of the patient's PSA level at the start of the first on-treatment cycle.

PROOF. We assume, without loss of generality, that the start of our sample path will coincide with the start of the first on-treatment period. Note also that we choose to end our sample path at time T , and that this choice is independent of θ_i , $i = 1, 2$. Consequently, we will have $[0, T] \equiv [\xi_1, \eta_{K_T}]$, which implies that $\frac{\partial \xi_1}{\partial \theta_i} = \frac{\partial \eta_{K_T}}{\partial \theta_i} = 0$, $i = 1, 2$. Recall that the sample path of our SHA will consist of alternating on and off-treatment periods.

Since $z_1(t) = 0$ when $q(t) = q^{OFF}$, we can rewrite (13) as

$$\begin{aligned} L(\theta, x(0), z(0), T) &= \frac{W}{T} \sum_{k=1}^{K_T} \int_{\xi_k}^{\eta_k} \left[\frac{x_1(\theta, t) + x_2(\theta, t)}{PSA_{init}} \right] dt \\ &+ \frac{(1-W)}{T} \left[\sum_{m=1}^{M_T} \int_{\xi_m}^{\eta_m} \frac{z_1(t)}{T} dt + \int_{\xi_{M_T+1}}^T \frac{z_1(t)}{T} dt \right] \end{aligned} \quad (45)$$

Note that our sample path can either (a) end with an incomplete on-treatment period, or (b) end with an incomplete off-treatment period. In (45), we assume that (a) holds, since (b) is a special case of (a) for which $\int_0^T \frac{z_1(t)}{T} dt = \sum_{m=1}^{M_T} \int_{\xi_m}^{\eta_m} \frac{z_1(t)}{T} dt$. Observe that the end of an on-treatment period is coupled with the start of the subsequent off-treatment period, i.e., $x_i(\eta_k) = x_i(\xi_{k+1})$, $i = 1, 2$, $k = 1, \dots, K_T - 1$. Using this notation and taking the derivative of (45) yields

$$\begin{aligned} \frac{dL(\theta)}{d\theta_i} &= \frac{W}{T \cdot PSA_{init}} \sum_{k=1}^{K_T-1} \int_{\xi_k}^{\xi_{k+1}} [x'_{1,i}(\theta, t) + x'_{2,i}(\theta, t)] dt \\ &+ \frac{W}{T \cdot PSA_{init}} \sum_{k=1}^{K_T-1} [x_1(\xi_{k+1}) + x_2(\xi_{k+1})] \frac{\partial \xi_{k+1}}{\partial \theta_i} \\ &- \frac{W}{T \cdot PSA_{init}} \sum_{k=1}^{K_T-1} [x_1(\xi_k) + x_2(\xi_k)] \frac{\partial \xi_k}{\partial \theta_i} \\ &+ \frac{W}{T \cdot PSA_{init}} \int_{\xi_{K_T}}^T [x'_{1,i}(\theta, t) + x'_{2,i}(\theta, t)] dt \\ &+ \frac{W}{T \cdot PSA_{init}} [x_1(T) + x_2(T)] \frac{\partial T}{\partial \theta_i} \\ &- \frac{W}{T \cdot PSA_{init}} [x_1(\xi_{K_T}) + x_2(\xi_{K_T})] \frac{\partial \xi_{K_T}}{\partial \theta_i} \\ &+ \frac{(1-W)}{T} \sum_{m=1}^{M_T} \left[\int_{\xi_m}^{\eta_m} \frac{z'_{1,i}(t)}{T} dt + \frac{z_1(\eta_m^-)}{T} \frac{\partial \eta_m}{\partial \theta_i} - \frac{z_1(\xi_m^+)}{T} \frac{\partial \xi_m}{\partial \theta_i} \right] \\ &+ \frac{(1-W)}{T} \int_{\xi_{M_T+1}}^T \frac{z'_{1,i}(t)}{T} dt + \frac{z_1(T^-)}{T} \frac{\partial T}{\partial \theta_i} - \frac{z_1(\xi_{M_T+1}^+)}{T} \frac{\partial \xi_{M_T+1}}{\partial \theta_i} \end{aligned} \quad (46)$$

Observe that multiple cancelations of the second, third and sixth terms in (46) simplify to

$$\begin{aligned}
& \sum_{k=1}^{K_T-1} [x_1(\xi_{k+1}) + x_2(\xi_{k+1})] \frac{\partial \xi_{k+1}}{\partial \theta_i} \\
& - \sum_{k=1}^{K_T-1} [x_1(\xi_k) + x_2(\xi_k)] \frac{\partial \xi_k}{\partial \theta_i} \\
& = [x_1(\xi_{K_T}) + x_2(\xi_{K_T})] \frac{\partial \xi_{K_T}}{\partial \theta_i}
\end{aligned} \tag{47}$$

Further note that the sixth term in (46) cancels out with the second term on the right hand side of (47). Moreover, it is clear from (6) that $z_1(\xi_{M_T+1}^+) = z_1(\xi_m^+) = 0$ and $z_1(\eta_m^-) = \eta_m - \xi_m$, $m = 1, \dots, M_T$. Since $z'_{j,i}(t) = z'_{j,i}(\tau_k^+)$, $j, i = 1, 2$, over any interevent interval $[\tau_k, \tau_{k+1})$, and recalling that $\frac{\partial T}{\partial \theta_i} = \frac{\partial \xi_1}{\partial \theta_i} = 0$, the last two terms in (46) simplify to

$$\begin{aligned}
& \frac{(1-W)}{T^2} \left[\sum_{m=1}^{M_T} z'_{1,i}(\xi_m^+) (\eta_m - \xi_m) + (\eta_m - \xi_m) \frac{\partial \eta_m}{\partial \theta_i} \right] \\
& + \frac{(1-W)}{T^2} z'_{1,i}(\xi_{M_T+1}^+) (T - \xi_{M_T+1})
\end{aligned}$$

Recall that ξ_m is the start of the m th on-treatment period, which necessarily corresponds to the $(m-1)$ th occurrence of event e_2 . Hence, $z'_{1,i}(\xi_m^+) = -\xi'_{m,i}$, $m = 1, \dots, M_{T+1}$ from (40). As a result, (46) can be further simplified to

$$\begin{aligned}
\frac{dL(\theta)}{d\theta_i} &= \frac{W}{T \cdot PSA_{init}} \sum_{k=1}^{K_T-1} \int_{\xi_k}^{\xi_{k+1}} [x'_{1,i}(\theta, t) + x'_{2,i}(\theta, t)] dt \\
&+ \frac{W}{T \cdot PSA_{init}} \int_{\xi_{K_T}}^T [x'_{1,i}(\theta, t) + x'_{2,i}(\theta, t)] dt \\
&+ \frac{(1-W)}{T^2} \left[\sum_{m=1}^{M_T} -\xi'_{m,i} (\eta_m - \xi_m) + (\eta_m - \xi_m) \eta'_{m,i} \right] \\
&- \frac{(1-W)}{T^2} \xi'_{M_T+1} (T - \xi_{M_T+1})
\end{aligned} \tag{48}$$

The result in (48) is obtained under the assumption that our sample path ends with an incomplete on-treatment period, i.e., K_T is odd. If this is not the case, the last term in (48) can be disregarded. It is then straightforward to verify that (48) can be rewritten as (44).

It is clear that evaluating (44) requires knowledge of: (i) the event times $\xi_{n,m}$ and $\eta_{n,m}$, and (ii) the value of the state derivatives $x'_{1,i}(\theta, t)$ and $x'_{2,i}(\theta, t)$ over all on and off-treatment periods. The quantities in (i) are easily observed using timers whose start and end times are observable events; eventually

knowledge of the noise processes $\zeta_1(t)$ and $\zeta_2(t)$ evaluated at event times *only* is also needed to compute (42)-(43). Information on the noise processes can be extracted from the observed sample path, as explained in Section 4. The state derivatives in (ii) are obtained from (22) and (25) over on-treatment periods, and from (28) and (31) over off-treatment periods. Ultimately, these expressions depend on (8) and (10), so that it is necessary to evaluate the integral of the noise process $\zeta_3(t)$, which can also be accomplished using data extracted from the observed sample path.

As a result, it is straightforward to implement an algorithm for updating the value of $dL(\theta)/d\theta_i$ after each observed event, as outlined in Algorithm 1.

Algorithm 1 IPA Algorithm for Optimal IAS Therapy Design

Whenever an event occurs at time τ_k , $k = 1, 2, \dots$

Step 1 Update event time derivatives using (41)

Step 2 Update state derivatives using (34)-(35) and (40)

Step 3 Update cost derivatives using (44)

End

Repeat

We draw attention to the fact that one of the appealing features of Algorithm 1 is its event-driven simplicity: despite the apparent complexity of the equations involved, the algorithm is simply executed whenever one of the events we have defined occurs and the only action required is an arithmetic computation (e.g., see (41)). In practical terms, running Algorithm 1 requires that some settings be adjusted, as discussed next. For one, the gradient step used in our simulations was of the form $\rho_k = \frac{C}{k^{3/2}}$, where ρ_k denotes the step size at the k^{th} iteration, $k = 1, 2, \dots$, and C is the step constant. The value of constant C was adjusted on a case-by-case basis, i.e., the value used in the simulation shown in Figure 3(a) was different than the one used in the simulations shown in Figure 3(b)-(c) and Figure 4. Additionally, convergence was evaluated based on the absolute difference in values of the average cost over subsequent iterations. More specifically, we considered convergence as having been reached when such difference was at most of the order of 10^{-3} over at least 3 consecutive iterations. Lastly, the time required for our IPA algorithm to achieve convergence was in most cases at least 3 times smaller than the time required to evaluate the response surface using a brute force approach. We end by noting that *Step 1* requires knowledge

of the noise processes $\zeta_1(t)$ and $\zeta_2(t)$ in order to evaluate (42) whenever an event e_1 takes place, and (43) whenever an event e_2 occurs. In what follows, details are given on how to include this type of information using data from the observed sample path.

4. Results

In what follows, we detail how an IPA-driven gradient-based optimization approach can be used for personalized IAS therapy design. The results shown here represent a first attempt at incorporating randomness into a SHA model of prostate cancer evolution in which we consider only noise and fluctuations associated with cell population dynamics, and do not account for noise in the patient's androgen level. Representing randomness as Gaussian white noise, the authors in [13] verified that variable time courses of the PSA levels were produced without losing the tendency of the deterministic system, thus yielding simulation results that were comparable to the statistics of clinical data. For this reason, in this work we take $\{\zeta_i(t)\}$, $i = 1, 2$, to be Gaussian white noise with zero mean and standard deviation of 0.001, similarly to [13], although we remind the reader that our methodology applies independently of the distribution chosen to represent $\{\zeta_i(t)\}$, $i = 1, 2$. We estimate the noise associated with cell population dynamics at event times by randomly sampling from a uniform distribution with zero mean and standard deviation of 0.001. Simulations of the prostate cancer model as a pure DES are thus run to generate sample path data to which the IPA estimator is applied. In all results reported here, we measure the sample path length in between updates of the controllable parameter vector θ in terms of the number of days elapsed since the onset of IAS therapy, which we choose to be $T = 2500$ days.

Two sets of simulations are reported here: one in which we set $W = 0.5$ and determine personalized treatment schemes for two different patients, and another in which we analyze the effect of W on the design of a given patient's optimal therapy. For the former, we make use of the clinical models of Patient #15 and Patient #1 [12].

Figure 3 presents the convergence plots of the average cost and control parameters considering different initial configurations for Patient #15.

Figure 4 presents the convergence plots of the average cost and control parameters considering different initial configurations for Patient #1.

Using a brute force approach, the response surface corresponding to our cost function was generated. In Figure 5, the convergence trajectories from Figure 3 are plotted against the response surface of Patient #15, and it can be seen that all configurations converge to the region of minimum cost.

In Figure 6, the convergence trajectories from Figure 4 are plotted against the response surface of Patient #1, and again it can be seen that all configurations converge to the region of minimum cost.

In our second set of simulations, we use the clinical model of Patient #1 [12] and take $W \in \{0.1, 0.5, 0.9\}$. Recall that our sample function contains two terms, each of which represents complementary measures of therapy success. We take into account the fact that it may be desirable to design a therapy scheme that favors one of these terms over the other by associating weight W with the first term and $1 - W$ with the second term, as detailed in Section 2.2. Table 1 presents the values of optimal lower and upper PSA threshold values (θ_1^* and θ_2^* , respectively) and the corresponding cost of treatment (J^*) using the clinical model Patient #1 for different values of W . In what follows, we adopt the notation $x \approx y$ to indicate that x takes values *approximately equal* to y , and $x \gtrsim y$ to indicate that x takes values *approximately equal or slightly greater* than y .

Table 1: Optimization results for different values of W using the clinical model of Patient 1

W	θ_1^*	θ_2^*	J^*
0.1	≈ 2.5	≈ 8.0	0.017
0.5	≈ 2.5	≈ 8.0	0.086
0.9	≈ 2.5	≈ 8.0	0.154

It can be seen from Table 1 that the ranges of optimal lower and upper PSA threshold values are equivalent for $W \in \{0.1, 0.5, 0.9\}$, i.e., $\theta_1^* \gtrsim \theta_1^{\min}$ and $\theta_2^* \gtrsim \theta_2^{\min}$ for all values of W considered here. In other words, the value of J^* changes with the value of W , but the regions of minimum and maximum cost remain essentially unchanged. This means that a common optimal treatment scheme exists irrespective of the chosen value of W . Hence, it is possible to consistently achieve therapy personalization, as detailed in what follows.

4.1. Achieving Therapy Personalization

Recall that solving problem (15) is equivalent to searching for the IAS treatment scheme that not only keeps PSA levels as low as possible, but that also reduces the frequency of on and off-treatment cycles. In this context, the resulting optimal IAS therapy is one in which both the low and high PSA thresholds take values as small as possible, i.e., $\theta_1^* \gtrsim \theta_1^{\min}$ and $\theta_2^* \gtrsim \theta_2^{\min}$. This tendency is verified in Figures 5 and 6, where it can be seen that the regions of minimum cost are those immediately surrounding $[\theta_1^{\min}, \theta_2^{\min}]$.

Moreover, it is interesting to note that although this tendency is consistent across different patients, successful IAS therapy schemes are only obtained when $\theta_1 \geq 1.5$ for Patient #15 and $\theta_1 \geq 2.5$ for Patient #1. This means that, for Patient #15, a therapy in which the low PSA threshold takes values smaller than 1.5 will eventually lead to uncontrolled cancer cell growth and disease relapse. The same analysis holds for Patient #1, except that in this case, there is a higher lower bound on the value of θ_1 . Clearly this variation across patients is associated with the fact that each patient responds differently to IAS therapy; hence the underlying need for designing *personalized* treatment schemes.

Finally, personalizing IAS therapy based on the cost metric proposed in this work involves, among other things, assessing the smallest value that the low PSA threshold can be allowed to reach, which varies across patients. This can be done using clinical data from patients, when available. In cases where no model exists to predict a patient's response to therapy (e.g., for patients who have never been submitted to IAS therapy), a possible course of action would be to devise an optimal therapy scheme using another patient's model, which could be selected based on clinical indicators of patient similarity and/or insights from the physician. After recording the new patient's response to the first cycle(s) of treatment, it would be possible to iteratively adjust his initial treatment until an improved scheme is found. Such iterative search for an optimal IAS therapy scheme could be successfully performed using the methodology proposed in this work.

We end with a note on the relevance of experimental or clinical supporting data for validation of our methodology. Such validation would ideally rely on randomized clinical trials where, for example, a cohort of patients would be submitted to current treatment protocols for prostate cancer, while another group of patients would undergo personalized treatment as determined by our approach. Data obtained by monitoring all patients throughout the

study would allow us to assess the relative advantages of optimal (personalized) therapies over existing therapy schemes. As already mentioned, the purpose of this paper is to demonstrate the applicability of IPA techniques for personalized therapy design by means of a case study of advanced prostate cancer. While performing clinical validation of our methodology lies outside the scope of this work, we stress that the simulation results shown here support our claim and lay the foundation for such future validation.

5. Conclusion

This work sets the stage for the use of basic IPA techniques for optimal personalized cancer therapy design. We propose a methodology applicable to stochastic models of cancer progression and illustrate our analysis with a case study of optimal IAS therapy design for advanced prostate cancer. We develop a threshold-based policy for optimal IAS therapy design that is parameterized by lower and upper PSA threshold values and is associated with a cost metric that combines clinically relevant measures of therapy success. We use a Stochastic Hybrid Automaton (SHA) model of prostate cancer evolution under IAS and perform Infinitesimal Perturbation Analysis (IPA) to adaptively adjust PSA threshold values so as to improve therapy outcomes.

Results obtained by applying our methodology to clinical data from real prostate cancer patients suggest that optimal IAS treatment schemes are those in which both the low and high PSA thresholds take values as small as possible. In spite of the fact that this tendency is consistent across different patients, lower bounds on PSA threshold values vary from one patient to the next. This variation across patients is associated with the fact that each patient responds differently to IAS therapy; hence the underlying need for designing *personalized* treatment schemes.

It is possible to extend the framework presented here to analyze different potentially interesting controllable parameters, such as different drugs and/or dosages. In this sense, our method would yield information on the effect of, e.g., mixing different medication components or timing therapy periods, on the overall effectiveness of the treatment. Our ongoing work includes performing sensitivity analysis on relevant model parameters other than the PSA threshold values. Lastly, we note that our methodology is general and thus easily applicable not only to other types of cancer, but also to other diseases that are known to progress in stages (e.g., tuberculosis).

- [1] D. Hanahan, R. Weinberg, Hallmarks of cancer: The next generation, *Cell* 144 (2011) 646–674.
- [2] D. Longo, A. Fauci, D. Kasper, S. Hauser, J. Jameson, J. Loscalzo (Eds.), *Harrison’s principles of internal medicine*, 18th Edition, McGraw-Hill, Medical Pub. Division, New York, 2012.
- [3] T. Jackson, A mathematical investigation of the multiple pathways to recurrent prostate cancer: comparison with experimental data, *Neoplasia* 6 (2004) 697–704.
- [4] Y. Hirata, N. Bruchovsky, K. Aihara, Development of a mathematical model that predicts the outcome of hormone therapy for prostate cancer, *J. Theor. Biology* 264 (2010) 517–527.
- [5] N. Bruchovsky, L. Klotz, J. Crook, S. Malone, C. Ludgte, W. Morris, M. Gleave, S. Goldenberg, P. Rennie, Final results of the Canadian prospective phase II trial of intermittent androgen suppression for men in biochemical recurrence after radiotherapy for locally advanced prostate cancer, *Cancer* 107 (2006) 389–395.
- [6] N. Bruchovsky, L. Klotz, J. Crook, S. Malone, C. Ludgte, W. Morris, M. Gleave, S. Goldenberg, P. Rennie, Locally advanced prostate cancer biochemical results from a prospective phase II study of intermittent androgen suppression for men with evidence of prostate-specific antigen recurrence after radiotherapy, *Cancer* 109 (2007) 858–867.
- [7] A. Ideta, G. Tanaka, T. Takeuchi, K. Aihara, A mathematical model for intermittent androgen suppression for prostate cancer, *J. Nonlinear Sci.* 18 (2008) 593–614.
- [8] T. Shimada, K. Aihara, A nonlinear model with competition between prostate tumor cells and its application to intermittent androgen suppression therapy of prostate cancer, *Mathematical Biosciences* 214 (2008) 134–139.
- [9] Y. Tao, Q. Guo, K. Aihara, A mathematical model of prostate tumor growth under hormone therapy with mutation inhibitor, *J. Nonlinear Sci.* 20 (2010) 219–240.

- [10] T. Suzuki, N. Bruchovsky, K. Aihara, Piecewise affine systems modelling for optimizing therapy of prostate cancer, *Philos. Trans. R. Soc.* 368 (2010) 5045–5059.
- [11] Y. Hirata, M. di Bernardo, N. Bruchovsky, K. Aihara, Hybrid optimal scheduling for intermittent androgen suppression of prostate cancer, *Chaos* 20 (2010) 045125.
- [12] B. Liu, S. Kong, S. Gao, P. Zuliani, E. Clarke, Towards personalized cancer therapy using delta-reachability analysis, *HSCC2015*.
- [13] G. Tanaka, Y. Hirata, S. Goldenberg, N. Bruchovsky, K. Aihara, Mathematical modelling of prostate cancer growth and its application to hormone therapy, *Philos. Trans. R. Soc.* 368 (2010) 5029–5044.
- [14] J. L. Fleck, C. G. Cassandras, Infinitesimal perturbation analysis for personalized cancer therapy design, *Proc. of 5th IFAC Conference on Analysis and Design of Hybrid Systems* (2015) 205–210.
- [15] C. Cassandras, S. Lafortune, *Introduction to Discrete Event Systems*, 2nd Edition, Springer, 2008.
- [16] C. Cassandras, Y. Wardi, C. Panayiotou, C. Yao, Perturbation analysis and optimization of stochastic hybrid systems, *European Journal of Control* 6 (6) (2010) 642–664.
- [17] D. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison-Wesley Longman, 1989.
- [18] C. Chen, L. Lee, *Stochastic Simulation Optimization: An Optimal Computing Budget Allocation*, World Scientific Publishing Co., 2011.
- [19] Y. Ho, Q. Zhao, Q. Jia, *Ordinal Optimization: Soft Optimization for Hard Problems*, Springer Science and Business Media, 2008.
- [20] L. Hong, B. Nelson, Discrete optimization via simulation using compass, *Operations Research* 54 (1) (2006) 115–129.
- [21] L. Shi, S. Olafsson, Nested partitions method for global optimization, *Operations Research* 48 (3) (2000) 390–407.
- [22] J. Fleck, C. Cassandras, Personalized cancer therapy design: robustness vs. optimality, Preprint at <http://arxiv.org/abs/1603.00895>.

Figure 2: Stochastic Hybrid Automaton model of prostate cancer evolution under IAS therapy

Figure 3: Convergence plots of average cost and PSA threshold values for Patient #15: (a) initial configuration $[\theta_1^{init}, \theta_2^{init}] = [2.5, 14.0]$; (b) initial configuration $[\theta_1^{init}, \theta_2^{init}] = [7.0, 12.5]$; (c) initial configuration $[\theta_1^{init}, \theta_2^{init}] = [6.5, 9.0]$

Figure 4: Convergence plots of average cost and PSA threshold values for Patient #1: (a) initial configuration $[\theta_1^{init}, \theta_2^{init}] = [5.5, 14.5]$; (b) initial configuration $[\theta_1^{init}, \theta_2^{init}] = [7.0, 15.0]$; (c) initial configuration $[\theta_1^{init}, \theta_2^{init}] = [4.5, 11.5]$

Figure 5: Response surface and convergence trajectories (Patient #15)

Figure 6: Response surface and convergence trajectories (Patient #1)