

Towards Personalized Cancer Therapy Using Delta-Reachability Analysis*

Bing Liu
School of Medicine
University of Pittsburgh
liubing@pitt.edu

Soonho Kong
Computer Science Dept.
Carnegie Mellon University
soonhok@cs.cmu.edu

Sicun Gao
CSAIL
MIT
sicung@csail.mit.edu

Paolo Zuliani
School of Computer Science
Newcastle University
paolo.zuliani@ncl.ac.uk

Edmund M. Clarke
Computer Science Dept.
Carnegie Mellon University
emc@cs.cmu.edu

ABSTRACT

Recent clinical studies suggest that the efficacy of hormone therapy for prostate cancer depends on the characteristics of individual patients. In this paper, we develop a computational framework for identifying patient-specific androgen ablation therapy schedules for postponing the potential cancer relapse. We model the population dynamics of heterogeneous prostate cancer cells in response to androgen suppression as a nonlinear hybrid automaton. We estimate personalized kinetic parameters to characterize patients and employ δ -reachability analysis to predict patient-specific therapeutic strategies. The results show that our methods are promising and may lead to a prognostic tool for personalized cancer therapy.

Categories and Subject Descriptors

D.2.4 [Software Engineering]: Software/Program Verification—*Model checking*; J.3 [Life and Medical Sciences]: Biology and genetics

General Terms

Theory, Verification

Keywords

hybrid systems, delta-reachability, systems biology, prostate cancer, personalized therapy

1. INTRODUCTION

Prostate cancer is the second leading cause of cancer-related deaths among men in United States [24]. Hormone ther-

apy in the form of androgen deprivation has been a cornerstone of the management of advanced prostate cancer for several decades. However, controversy remains regarding its optimal application [5]. Continuous androgen suppression (CAS) therapy has many side effects including anemia, osteoporosis, impotence, etc. Further, most patients experience a relapse after a median duration of 18-24 months of CAS treatment, due to the proliferation of castration resistant cancer cells (CRCs).

In order to reduce side effects of CAS and to delay the time to relapse, intermittent androgen suppression (IAS) was proposed to limit the duration of androgen-poor conditions and avoid emergence of CRCs [2]. In particular, IAS therapy switches between on-treatment and off-treatment modes by monitoring the serum level of a tumor marker called prostate-specific antigen (PSA):

- When the PSA level decreases and reaches a lower threshold value r_0 , androgen suppression is suspended.

- When the PSA level increases and reaches an upper threshold value r_1 , androgen suppression is resumed by the administration of medical agents.

Recent clinical phase II and III trials confirmed that IAS has significant advantages in terms of quality of life and cost [3, 4]. However, with respect to time to relapse and cancer-specific survival, the clinical trials suggested that to what extent IAS is superior to CAS depends on the individual patient and the on- and off-treatment scheme [3, 4, 10]. Thus, a crucial unsolved problem is how to design a personalized treatment scheme for each individual to achieve maximum therapeutic efficacy.

To answer this question, mathematical models have been developed to study the dynamics of prostate cancer under androgen suppression [16, 15, 14, 12, 17, 23]. Recently, attempts have been made to computationally classify patients and obtain the optimal treatment scheme [13, 25]. However, these results relied on simplifying nonlinear hybrid dynamical systems to more manageable versions such as piecewise linear models [13] and piecewise affine systems [25], which compromises the validity of the models. In this

*This work has been partially supported by award N00014-13-1-0090 of the US Office of Naval Research and award CNS0926181 of the National Science foundation (NSF).

paper, we construct a nonlinear hybrid model to describe the prostate cancer progression dynamics under IAS therapy. Our model extends the models previously proposed in [16, 15, 14]. We use δ -reachability analysis to obtain the following results:

- First, we show that our model is in good agreement with the published clinical data in literature [3, 4]. It can depict the dynamical changes of proliferation rates induced by perturbing androgen levels that are difficult for previous models (e.g. [14]) to capture. It also addresses the variability in individual patients and is able to accurately reproduce the datasets of different patients.

- Second, we obtain interesting insights on CRC proliferation dynamics through analysis of the nonlinear model. Our results support the hypothesis that the physiological level of androgen reduce CRCs [14], while rule out other hypotheses, for instance, CRCs proliferate at a constant rate [23].

- Third, we propose a computational framework for identifying patient-specific IAS schedules for postponing the potential cancer relapse. Specifically, we obtain personalized model parameters by fitting to the clinical data in order to characterize individual patients. We then use δ -decision produces and bounded model checking to predict therapeutic strategies.

Through this case study, we aim to highlight the opportunity for solving realistic biomedical problems using formal methods. In particular, methods based on δ -reachability analysis suggest a very promising direction to proceed.

The rest of the paper is organized as follows. We describe our model in Section 2 and present preliminaries on δ -reachability analysis in Section 3. In Section 4, we present the biological insights we gained through this case study, as well as the model-predicted treatment schemes for individual patients. In the final section, we summarize the paper and discuss future work.

2. A HYBRID MODEL OF PROSTATE CANCER PROGRESSION

In this section, we propose a hybrid automata based model in order to reproduce the clinical observations [3, 4] of prostate cancer cell dynamics in response to the IAS therapy. It is known that the proliferation and survival of prostate cancer cells depend on the levels of androgens, specifically testosterone and 5 α -dihydrotestosterone (DHT). Here we consider two distinct subpopulations of prostate cancer cells: hormone sensitive cells (HSCs) and castration resistant cells (CRCs). Androgen deprivation can lead to remarkable decreases of the proliferation and survival rates of HSCs, but also up-regulates the conversion from HSCs to CRCs, which will keep proliferating under low androgen level. Figure 1 illustrates this cancer progression process and the corresponding hybrid automata model is shown in Figure 2.

Our model is based on previous models developed by [16, 15, 14]. It takes into account the population of HSCs, the population of CRCs, as well as the serum androgen concentration, represented as $x(t)$, $y(t)$, and $z(t)$, respectively. In addition, it also includes the serum prostate-specific antigen




Figure 1: Prostate cancer progression in response to (a) CAS and (b) IAS treatments.

(PSA) level $v(t)$, which is a commonly used biomarker for assessing the total population of prostate cancer cells. The model has two modes: *on-treatment* mode and *off-treatment* mode. Following [14], in the off-treatment mode (mode 2), the androgen concentration is maintained at the normal level z_0 by homeostasis. In the on-treatment (mode 1), the androgen is cleared at a rate $\frac{1}{\tau}$. Further, we also introduce a basal androgen production rate μ_z , in order to reproduce the measured basal testosterone levels in response to androgen suppression [3, 4].

The net growth rate of $x(t)$ equals to $(prolif_x - apopt_x - conv_x) \cdot x(t)$, where $prolif_x$, $apopt_x$ and $conv_x$ denote the proliferation, apoptosis and conversion rates, respectively. In previous studies such as [16, 15, 14], the $prolif_x$ and $apopt_x$ were modeled using Michaelis-Menten-like (MML) functions, in the form of $V_{max} + (1 - V_{max}) \frac{z(t)}{z(t) + K_m}$, where V_{max} and K_m are kinetic parameters. This approach will result in androgen response curves as shown in Figure 3(a). In particular, when one decreases the androgen level starting from the normal level, $prolif_x$ (or $apopt_x$) begins to decrease (or increase) first slowly and then fast until a sufficiently low level of androgen is reached. However, this is inconsistent with the clinical observations presented in [3, 4]. The data show that for most of the patients, androgen suppression around normal level will induce an immediate decrease of the PSA level, which implies an fast decrease (or increase) of $prolif_x$ (or $apopt_x$). Therefore, instead of the MML functions, we adopt sigmoid functions, in the form of $\frac{1}{1 + \exp(-(z(t) - k_1) \cdot k_2)}$, to model $prolif_x$ and $apopt_x$. The corresponding androgen response curves are shown in Figure 3(b). Following [14], we model the conversion rate, proliferation rate and the apoptosis rate of $y(t)$ as $m_1(1 - \frac{z(t)}{z_0})$, $\alpha_y(1 - d \frac{z(t)}{z_0})$ and β_y , respectively. The PSA level v (ng ml $^{-1}$) is defined as $v(t) = c_1 \cdot x(t) + c_2 \cdot y(t)$.

The transitions between two modes depends on the values of v , dv/dt and an auxiliary variable w , which measures the time taken in a mode. Specifically, for each patient we starts with mode 1 to apply the treatment. When the PSA level drops to certain threshold r_0 or w hits time out threshold t_{max} , the treatment will be suspended. When the PSA level is back to threshold r_1 , the treatment will be resumed. Note that w is associated with a dummy differential equation $\frac{dw}{dt} = 1$ (not shown in Figure 2). Its value will be reset to 0 when the jump takes place.




Figure 2: A hybrid automaton model for prostate cancer hormone therapy.

We obtained the parameter values by fitting to patient PSA data reported in [3, 4]. Note that the patient-to-patient variability in terms of parameter values is significant. For example, Figure 4 shows that the proliferation rate of Patient#22 is much lower than the Patient#1. The descriptions and a set of typical values (i.e. estimated from Patient#1 data) of model parameters are listed in Table 1.

3. DELTA-REACHABILITY ANALYSIS

Hybrid automata are difficult to analyze. It has been shown that even simple reachability questions for hybrid systems with linear differential dynamics are undecidable [11]. Therefore, in order to analyze our hybrid model of prostate cancer progression, we employed a δ -reachability based framework [21] which can sidestep undecidability and allows the parameter synthesis problem to be relaxed in a sound manner and solved algorithmically.

3.1 Delta-Decisions

The framework of δ -complete decision procedures [6] aims to solve first-order logic formula with arbitrary computable real functions, such as elementary functions and solutions of Lipschitz-continuous ODEs [7]. The answers returned by such procedures are either *unsat* or δ -sat. Here, *unsat* means the corresponding formula is verifiably false, while δ -sat means a δ -weakening version of the formula is true. In other words, δ -decision procedures overcome undecidability issues by returning answers with one-sided δ -bounded errors.




Figure 3: Androgen response curves of (a) Ideta's model and (b) our model.




Figure 4: The clinical data for PSA time series.

Note that δ is an arbitrarily small positive rational chosen by the user. The algorithms for solving δ -decision problems were described in our previous work [7, 9] and were implemented in the dReal toolset [8].

3.2 Parameter identification

Further, we have also proposed an encoding scheme which aimed to answer bounded reachability problems of hybrid automata with nontrivial invariants [21]. This encoding enabled us to tackle the parameter identification problem by answering a k -step reachability question: “Is there a parameter combination for which the model reaches the goal region in k steps?” Essentially, we describe the set of states of interest (goal region) as a first-order logic formula and perform bounded model checking [1] to determine reachability of these states. We then adapt an interval constraints propagation based algorithm to explore the parameter space and identify the sets of resulting parameters. If none exists, then the model is *unfeasible*. Otherwise, a witness (i.e., a value for each parameter) is returned. We have developed the dReach tool (<http://dreal.cs.cmu.edu/dreach.html>) that automatically builds reachability formulas from a hybrid model and a goal description. Such formulas are then solved by the δ -complete solver dReal [8].

For the interested readers, we refer to Appendix and [21] for more details on δ -decisions and δ -reachability analysis based parameter identification.

4. RESULTS

We have implemented our prostate cancer progression model in the dReach's modeling language. The model files are available at <http://www.cs.cmu.edu/~liubing/hsccl5/>. All the experiments reported below were done using a machine with two Intel Xeon E5-2650 2.00GHz processors and 32GB RAM. The precision δ was set to 10^{-3} .

Table 1: Prostate cancer model parameter values

Parameter	Value	Remark
α_x	0.0204 d^{-1}	HSC proliferation
α_y	0.0242 d^{-1}	CRC proliferation
β_x	0.0201 d^{-1}	HSC apoptosis
β_y	0.0168 d^{-1}	CRC apoptosis
k_1	10.0 nM	HSC proliferation
k_2	1.0	HSC proliferation
k_3	10.0 nM	HSC apoptosis
k_4	2	HSC apoptosis
m_1	0.00005 d^{-1}	HSC to CRC conversion
z_0	12.0 nM	normal androgen level
τ	12.5 d	androgen degradation
λ_x	0.01 d^{-1}	HSC basal degradation
μ_x	0.05 d^{-1}	HSC basal production
μ_z	0.02 d^{-1}	Androgen basal production

4.1 CRC proliferation dynamics

Due to the lack of biomarkers distinguishing HSCs and CRCs *in vivo*, the proliferation kinetics of CRCs in response to androgen is far from known. Three hypotheses, denoted as H_1 , H_2 and H_3 have been proposed to describe the androgen-dependent CRC growth [14], which are discriminated by the value of d in the model, i.e.:


- $H_1 : d = 0$, the grow of CRCs is independent of $z(t)$;
- $H_2 : d = 1 - \frac{\beta_y}{\alpha_y}$, CRCs cease growing when $z(t) = z_0$;
- $H_3 : d = 1$, CRCs decrease when $z(t) = z_0$.

The Patient#1 data presented in Figure 4 shows that with proper treatment schedules, it is possible to avoid his cancer relapse in years. We now show that only H_3 agrees with this observation. As the PSA level $v(t)$ reflects the total number of cancer cells and CRCs are responsible for recurrent cancer, we use two invariants: $0 \leq v(t) \leq 30$ and $0 \leq y(t) \leq 1$ to specify the property of “no cancer relapse”. We then carried out δ -reachability analysis to verify whether the invariants hold for each of the model candidates within a bounded time of 365 days. Here the treatment schedule threshold parameters were provided as ranges: $r_0 \in [0, 7.99]$ (ng ml⁻¹) and $r_1 \in [8, 15]$ (ng ml⁻¹).

The **unsat** answers were returned for H_1 and H_2 (Run#1 and Run#2, Table 2), indicating that they will always lead to cancer relapse no matter which treatment schedule was chosen. In contrast, δ -sat was returned for H_3 (Run#3, Table 2). Witness trajectories are shown in Figure 5), demonstrating that the cancer relapse can be avoided in a bound time as observed experimentally [3, 4]. The rest of the results in this paper were generated using H_3 .

4.2 Androgen-dependent HSC dynamics

As mentioned in Section 2, previous studies [16, 15, 14] modeled the androgen-dependent proliferation and apoptosis of HSCs using MML functions, while we use sigmoid functions. Here we show that the MML based approach is unable to reproduce an important dynamical property, but our model could. The patients’ data in [3, 4] show that the *half-time* $t_{1/2}$ (i.e. the amount of time required for a quantity to fall to one half of its initial value) of PSA level under androgen suppression is often less than 60 days. To specify this property, we introduced an auxiliary mode (Mode 3). If $v(t) = v(0)/2$,


Figure 5: The simulated witness trajectories of the H_3 model.

the system will jump from Mode 1 to Mode 3. Starting with Mode 1 and $20 \leq x(0) \leq 30$, we checked the reachability of a goal state with $0 \leq w \leq 60$ for both Ideta’s model [14] and our model. The results show that δ -sat was returned for our model (Run#4, Table 2), while **unsat** was returned for Ideta’s model (Run#5, Table 2), suggesting the superiority of sigmoid functions over MML functions in capturing HSC dynamics.

Run	Model	Initial State	Result	Time
1	H_1	$r_0 \in [0.0, 7.99], r_1 \in [8.0, 15.0]$	unsat	3.94
2	H_2	$r_0 \in [0.0, 7.99], r_1 \in [8.0, 15.0]$	unsat	5.26
3	H_3	$r_0 \in [0.0, 7.99], r_1 \in [8.0, 15.0]$	δ -sat	472
4	H_3	$x(0) \in [20.0, 30.0]$	δ -sat	10.1
5	Ideta	$x(0) \in [20.0, 30.0]$	unsat	0.5
6	H_3	$r_0 \in [0.0, 7.99], r_1 \in [8.0, 15.0]$	δ -sat	526
7	H_3	$r_0 \in [0.0, 7.99], r_1 \in [8.0, 15.0]$	unsat	0.3
8	H_3	$r_0 \in [0.0, 7.99], r_1 \in [8.0, 15.0]$	δ -sat	28
9	H_3	$r_0 \in [0.0, 7.99], r_1 \in [8.0, 15.0]$	δ -sat	203

Table 2: Experimental results. Result = bounded model checking result, Time = CPU time (s), $\delta = 10^{-3}$.

4.3 Personalized therapy design

We next apply δ -reachability analysis to design treatment schemes for individual patients. The parameter values shown in Table 1 were estimated by fitting the data of Patient#1. Since the IAS response of Patient#1 is typical, we treated its parameter values as the baseline values. As we demonstrated in Figure 4, the values of some parameters vary among patients. Such variability may significantly affect the hormone therapy responses. For example, Figure 6(a-c) illustrates the PSA dynamics of 3 mock patients with different personalized parameters under the same IAS treatment scheme ($r_0 = 4$, $r_1 = 10$). IAS prevents the relapse for Patient A and delays the relapse for Patient B, but does not help Patient C. Figure 6(d) shows that, by modifying the IAS scheduling parameters r_0 and r_1 , the relapse of Patient C can be avoided or delayed.

Table 3: Estimated personalized parameters and suggested treatment schemes

Parameter	Patient#1	Patent#11	Patient # 15	Patient#26
α_x	0.0204 d ⁻¹	0.0204 d ⁻¹	0.0213 d ⁻¹	0.0197 d ⁻¹
α_y	0.0242 d ⁻¹	0.0242 d ⁻¹	0.0242 d ⁻¹	0.0242 d ⁻¹
β_x	0.0201 d ⁻¹	0.02 d ⁻¹	0.01 d ⁻¹	0.0175 d ⁻¹
β_y	0.0168 d ⁻¹	0.0158 d ⁻¹	0.0168 d ⁻¹	0.0168 d ⁻¹
k_1	10.0 nM	7.0 nM	7.0 nM	10.0 nM
k_2	1.0	1.0	1.0	1.0
k_3	10.0 nM	7.0 nM	7.4 nM	10.0 nM
k_4	2	2	2	2
m_1	0.00005 d ⁻¹	0.00005 d ⁻¹	0.00005 d ⁻¹	0.00005 d ⁻¹
z_0	12.0 nM	9.0 nM	8.0 nM	12.0 nM
τ	12.5 d	12.5 d	12.5 d	12.5 d
λ_x	0.01 d ⁻¹	0.0121 d ⁻¹	0.01 d ⁻¹	0.01 d ⁻¹
μ_x	0.05 d ⁻¹	0.06 d ⁻¹	0.02 d ⁻¹	0.03 d ⁻¹
μ_z	0.02 d ⁻¹	0.02 d ⁻¹	0.02 d ⁻¹	0.02 d ⁻¹
Scheme	$r_0 = 5.2, r_1 = 10.8$	N.A	$r_0 = 1.9, r_1 = 8.0$	$r_0 = 4.6, r_1 = 10.7$




Figure 6: Simulated PSA profiles of patients with different parameters. (a) **Patient A:** $\alpha_y = 0.0242$, $\beta_y = 0.0168$, $m_1 = 0.00005$, $z(0) = 12$, $r_0 = 4$, $r_1 = 10$ (b) **Patient B:** $\alpha_y = 0.0328$, $\beta_y = 0.013$, $z(0) = 13$, $m_1 = 0.0001$, $r_0 = 4$, $r_1 = 10$ (c) **Patient C:** $\alpha_y = 0.0426$, $\beta_y = 0.189$, $m_1 = 0.00005$, $z(0) = 15$, $r_0 = 4$, $r_1 = 10$ (d) **Patient C:** $\alpha_y = 0.0426$, $\beta_y = 0.189$, $m_1 = 0.00005$, $z(0) = 15$, $r_0 = 4$, $r_1 = 10.6$.

Given the parameter values of an particular patient, we can design a treatment scheme, which might help him avoid cancer relapse with bounded time by solving the following parameter identification problem: (i) set the ranges of scheduling parameters as $r_0 \in [0, 7.99]$ (nM) and $r_1 \in [8, 15]$; (ii) check if H_3 can reach the goal state without violating the “no cancer relapse” invariants within 1 year. If `unsat` was returned, it means that androgen suppression therapy is not suitable for the patient. The patient has to resort to other kinds of therapeutic interventions. Otherwise, when the δ -sat answer is returned, a treatment scheme containing feasible values of r_0 and r_1 will also be returned, which could help preventing or delaying the relapse within bounded time. Note that if $r_0 = 0$ is returned, it implies that the CAS scheme, instead of IAS scheme, might be more suitable for the patient.

The personalized parameters of individual patients can be obtained by collectively fitting the available experimental data using global optimization method. We tested our method on real patients data collected by [4]¹. The parameter values for each randomly selected patient were estimated by fitting the model to the PSA time series data under the IAS therapy using evolutionary strategy search. As an example, Figure 7 shows the comparison between model predictions and the experimental data of PSA and androgen levels for Patient#1, Patient#11, Patient#15, and Patient#26. We then predicted the treatment schemes for the future year using δ -reachability analysis (Run#6 for Patient#1, Run#7 for Patient#15, Run#8 for Patient#26 and Run#9 for Patient#11, Table 2). The results are summarized in Table 3. Note that for Patient#11, `unsat` was returned, implying that no suitable treatment schemes were identified. This might be due to the raised population size of CRCs in the late phase of clinical trials.




Figure 7: Model prediction vs. experimental data.

¹Data available at <http://www.nicholasbruchovsky.com/clinicalResearch.html>.

5. CONCLUSION

We have proposed a hybrid model to study the prostate cancer cell dynamics in response to hormone therapy. Using δ -reachability analysis, we obtained interesting biological insights into the prostate cancer heterogeneity. We also developed a δ -decisions based computational framework for predicting patient-specific treatment schedules. We have demonstrated the applicability of our method with the help of real clinical datasets. Our study explored the possibilities of using formal methods to tackle quantitative systems pharmacology problems. Our results also highlighted δ -reachability analysis as a potent technique in this line of research.

Experimental validation of our method might require years of clinical studies, which is beyond the scope of this case study. It is worth noting that our therapy design framework is generic and can be applied to other settings, for example, predicting the radiation dosing schedules for brain cancer [18]. Furthermore, another interesting direction is to extend our model and framework to take into account the stochasticity in cellular environment. In this respect, the probabilistic modeling and statistical analysis techniques in [20, 22, 19] might offer helpful pointers.

6. REFERENCES

- [1] A. Biere, A. Cimatti, E. M. Clarke, and Y. Zhu. Symbolic model checking without BDDs. In *TACAS*, volume 1579 of *LNCS*, pages 193–207, 1999.
- [2] N. Bruchovsky, S. L. Goldenberg, P. S. Rennie, and G. M. E. Theoretical considerations and initial clinical results of intermittent hormone treatment of patients with advanced prostatic carcinoma. *Urology*, 34:389–392, 1995.
- [3] N. Bruchovsky, L. Klotz, J. Crook, S. Malone, C. Ludgte, W. Morris, M. E. Gleave, S. L. Goldenberg, and P. S. 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*, 107:389–395, 2006.
- [4] N. Bruchovsky, L. Klotz, J. Crook, S. Malone, C. Ludgte, W. Morris, M. E. Gleave, S. L. Goldenberg, and P. S. 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*, 109:858–867, 2007.
- [5] N. C. Buchan and S. L. Goldenberg. Intermittent androgen suppression for prostate cancer. *Nat. Rev. Urol.*, 7:552–560, 2010.
- [6] S. Gao, J. Avigad, and E. M. Clarke. Delta-complete decision procedures for satisfiability over the reals. In *IJCAR*, pages 286–300, 2012.
- [7] S. Gao, J. Avigad, and E. M. Clarke. Delta-decidability over the reals. In *LICS*, pages 305–314, 2012.
- [8] S. Gao, S. Kong, and E. M. Clarke. dReal: An SMT solver for nonlinear theories of reals. In *CADE*, pages 208–214, 2013.
- [9] S. Gao, S. Kong, and E. M. Clarke. Satisfiability modulo ODEs. In *FMCAD*, pages 105–112, 2013.
- [10] G. Hamilton and G. Theyer. *Advances in Prostate Cancer*, chapter 13, pages 305–330. INTECH, 2013.
- [11] T. A. Henzinger. The theory of hybrid automata. In *LICS*, pages 278–292, 1996.
- [12] Y. Hirata, N. Bruchovsky, and K. Aihara. Development of a mathematical model that predicts the outcome of hormone therapy for prostate cancer. *J. Theor. Biol.*, 264:517–527, 2010.
- [13] Y. Hirata, M. di Bernordo, N. Bruchovsky, and K. Aihara. Hybrid optimal scheduling for intermittent androgen suppression of prostate cancer. *Chaos*, 20(4):045125, 2010.
- [14] A. M. Ideta, G. Tanaka, T. Takeuchi, and K. Aihara. A mathematical model of intermittent androgen suppression for prostate cancer. *J. Nonlinear Sci.*, 18:593–614, 2008.
- [15] T. L. Jackson. A mathematical investigation of the multiple pathways to recurrent prostate cancer: comparison with experimental data. *Neoplasia*, 6:697–704, 2004.
- [16] T. L. Jackson. A mathematical model of prostate tumor growth and androgen-independent replace. *Discrete Cont. Dyn. Syst. Ser. B*, 4:187–201, 2004.
- [17] H. V. Jain, S. K. Clinton, A. Bhinder, and A. Friedman. Mathematical modeling of prostate cancer progression in response to androgen ablation therapy. *P. Natl. Acad. Sci. USA.*, 108(49):19701–19706, 2011.
- [18] K. Leder, K. Pitter, Q. LaPlant, D. Hambardzumyan, B. D. Ross, T. A. Chan, E. C. Holland, and F. Michor. Mathematical modeling of PDGF-driven glioblastoma reveals optimized radiation dosing schedules. *Cell*, 156:603–616, 2014.
- [19] B. Liu, A. Hagiescu, S. K. Palaniappan, B. Chattopadhyay, Z. Cui, W. Wong, and P. S. Thiagarajan. Approximate probabilistic analysis of biopathway dynamics. *Bioinformatics*, 28(11):1508–1516, 2012.
- [20] B. Liu, D. Hsu, and P. Thiagarajan. Probabilistic approximations of ODEs based bio-pathway dynamics. *Theor. Comput. Sci.*, 412(21):2188–2206, 2011.
- [21] B. Liu, S. Kong, S. Gao, P. Zuliani, and E. M. Clarke. Parameter synthesis for cardiac cell hybrid models using δ -decisions. In *CMSB’14*, pages 99–113, 2014.
- [22] S. K. Palaniappan, B. M. Gyori, B. Liu, D. Hsu, and P. S. Thiagarajan. Statistical model checking based calibration and analysis of bio-pathway models. In *CMSB*, volume 8130 of *LNCS*, pages 120–134, 2013.
- [23] T. Portz, Y. Kuang, and J. D. Nagy. A clinical data validated mathematical model of prostate cancer growth under intermittent androgen suppression therapy. *AIP Advances*, 2:011002, 2012.
- [24] R. Siegel, D. Naishadham, and A. Jemal. Cancer statistics, 2013. *CA. Cancer J. Clin.*, 63:11–30, 2013.
- [25] T. Suzuki, N. Bruchovsky, and A. K. Piecewise affine systems modelling for optimizing hormone therapy of prostate cancer. *Philos. Trans. A Math. Phys. Eng. Sci.*, 368(1930):5045–5059, 2010.

Appendix

$\mathcal{L}_{\mathbb{R}\mathcal{F}}$ -Formulas and δ -Decidability

We will use a logical language over the real numbers that allows arbitrary *computable real functions*. We write $\mathcal{L}_{\mathbb{R}\mathcal{F}}$ to represent this language. Intuitively, a real function is computable if it can be numerically simulated up to an arbitrary precision. For the purpose of this paper, it suffices to know that almost all the functions that are needed in describing hybrid systems are Type 2 computable, such as polynomials, exponentiation, logarithm, trigonometric functions, and solution functions of Lipschitz-continuous ordinary differential equations.

More formally, $\mathcal{L}_{\mathbb{R}\mathcal{F}} = \langle \mathcal{F}, > \rangle$ represents the first-order signature over the reals with the set $\mathcal{F}$ of computable real functions, which contains all the functions mentioned above. Note that constants are included as 0-ary functions. $\mathcal{L}_{\mathbb{R}\mathcal{F}}$ -formulas are evaluated in the standard way over the structure $\mathbb{R}_{\mathcal{F}} = \langle \mathbb{R}, \mathcal{F}^{\mathbb{R}}, >^{\mathbb{R}} \rangle$. It is not hard to see that we can put any $\mathcal{L}_{\mathbb{R}\mathcal{F}}$ -formula in a normal form, such that its atomic formulas are of the form $t(x_1, \dots, x_n) > 0$ or $t(x_1, \dots, x_n) \geq 0$, with $t(x_1, \dots, x_n)$ composed of functions in $\mathcal{F}$. To avoid extra preprocessing of formulas, we can explicitly define $\mathcal{L}_{\mathcal{F}}$ -formulas as follows.

Definition 1 $\mathcal{L}_{\mathbb{R}\mathcal{F}}$ -Formulas Let $\mathcal{F}$ be a collection of computable real functions. We define:

$$t := x \mid f(t(\vec{x})), \text{ where } f \in \mathcal{F} \text{ (constants are 0-ary functions);}$$

$$\varphi := t(\vec{x}) > 0 \mid t(\vec{x}) \geq 0 \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid \exists x_i \varphi \mid \forall x_i \varphi.$$

In this setting $\neg\varphi$ is regarded as an inductively defined operation which replaces atomic formulas $t > 0$ with $-t \geq 0$, atomic formulas $t \geq 0$ with $-t > 0$, switches $\wedge$ and $\vee$, and switches $\forall$ and $\exists$.

Definition 2 Bounded $\mathcal{L}_{\mathbb{R}\mathcal{F}}$ -Sentences We define the bounded quantifiers $\exists^{[u,v]}$ and $\forall^{[u,v]}$ as $\exists^{[u,v]} x. \varphi =_{df} \exists x. (u \leq x \wedge x \leq v \wedge \varphi)$ and $\forall^{[u,v]} x. \varphi =_{df} \forall x. ((u \leq x \wedge x \leq v) \rightarrow \varphi)$ where u and v denote $\mathcal{L}_{\mathbb{R}\mathcal{F}}$ terms, whose variables only contain free variables in φ excluding x . A *bounded $\mathcal{L}_{\mathbb{R}\mathcal{F}}$ -sentence* is

$$Q_1^{[u_1, v_1]} x_1 \dots Q_n^{[u_n, v_n]} x_n \psi(x_1, \dots, x_n),$$

where $Q_i^{[u_i, v_i]}$ are bounded quantifiers, and $\psi(x_1, \dots, x_n)$ is quantifier-free.

Definition 3 δ -Variants Let $\delta \in \mathbb{Q}^+ \cup \{0\}$, and φ an $\mathcal{L}_{\mathbb{R}\mathcal{F}}$ -formula

$$\varphi : Q_1^{I_1} x_1 \dots Q_n^{I_n} x_n \psi[t_i(\vec{x}, \vec{y}) > 0; t_j(\vec{x}, \vec{y}) \geq 0],$$

where $i \in \{1, \dots, k\}$ and $j \in \{k+1, \dots, m\}$. The δ -weakening φ^δ of φ is defined as the result of replacing each atom $t_i > 0$ by $t_i > -\delta$ and $t_j \geq 0$ by $t_j \geq -\delta$:

$$\varphi^\delta : Q_1^{I_1} x_1 \dots Q_n^{I_n} x_n \psi[t_i(\vec{x}, \vec{y}) > -\delta; t_j(\vec{x}, \vec{y}) \geq -\delta].$$

It is clear that $\varphi \rightarrow \varphi^\delta$ (see [7]). In [6], we have proved

that the following δ -decision problem is decidable, which is the basis of our framework.

Theorem 1 δ -Decidability [6] Let $\delta \in \mathbb{Q}^+$ be arbitrary. There is an algorithm which, given any bounded $\mathcal{L}_{\mathbb{R}\mathcal{F}}$ -sentence φ , correctly returns one of the following two answers:

- δ -True: φ^δ is true.
- False: φ is false.

When the two cases overlap, either answer is correct.

The following theorem states the (relative) complexity of the δ -decision problem. A bounded Σ_n sentence is a bounded $\mathcal{L}_{\mathbb{R}\mathcal{F}}$ -sentence with n alternating quantifier blocks starting with $\exists$.

Theorem 2 Complexity [7] Let S be a class of $\mathcal{L}_{\mathbb{R}\mathcal{F}}$ -sentences, such that for any φ in S , the terms in φ are in Type 2 complexity class $\mathcal{C}$. Then, for any $\delta \in \mathbb{Q}^+$, the δ -decision problem for bounded Σ_n -sentences in S is in $(\Sigma_n^P)^C$.

Basically, the theorem says that increasing the number of quantifier alternations will in general increase the complexity of the problem, unless $P = NP$ (recall that $\Sigma_0^P = P$ and $\Sigma_1^P = NP$). This result can be specialized for specific families of functions. For example, with polynomially-computable functions, the δ -decision problem for bounded Σ_n -sentences is (Σ_n^P) -complete. For more details and results we again point the interested reader to [7].

Delta-Decisions for Hybrid Models

Now we state the encoding for hybrid models. Recall that hybrid automata generalize finite-state automata by permitting continuous-time evolution (or *flow*) in each discrete state (or *mode*). Also, in each mode an *invariant* must be satisfied by the flow, and mode switches are controlled by *jump* conditions.

Definition 4 $\mathcal{L}_{\mathbb{R}\mathcal{F}}$ -Representations of Hybrid Automata A hybrid automaton in $\mathcal{L}_{\mathbb{R}\mathcal{F}}$ -representation is a tuple

$$H = \langle X, Q, \{\text{flow}_q(\vec{x}, \vec{y}, t) : q \in Q\}, \{\text{inv}_q(\vec{x}) : q \in Q\}, \\ \{\text{jump}_{q \rightarrow q'}(\vec{x}, \vec{y}) : q, q' \in Q\}, \{\text{init}_q(\vec{x}) : q \in Q\} \rangle$$

where $X \subseteq \mathbb{R}^n$ for some $n \in \mathbb{N}$, $Q = \{q_1, \dots, q_m\}$ is a finite set of modes, and the other components are finite sets of quantifier-free $\mathcal{L}_{\mathbb{R}\mathcal{F}}$ -formulas.

We now show the encoding of bounded reachability, which is used for encoding the parameter synthesis problem. We want to decide whether a given hybrid system reaches a particular region of its state space after following a (bounded) number of discrete transitions, *i.e.*, jumps. First, we need to define auxiliary formulas used for ensuring that a particular mode is picked at a certain step.

Definition 5 Let $Q = \{q_1, \dots, q_m\}$ be a set of modes. For any $q \in Q$, and $i \in \mathbb{N}$, use b_q^i to represent a Boolean variable. We now define

$$\text{enforce}_Q(q, i) = b_q^i \wedge \bigwedge_{p \in Q \setminus \{q\}} \neg b_p^i$$

$$\text{enforce}_Q(q, q', i) = b_q^i \wedge \neg b_{q'}^{i+1} \wedge \bigwedge_{p \in Q \setminus \{q\}} \neg b_p^i \wedge \bigwedge_{p' \in Q \setminus \{q'\}} \neg b_{p'}^{i+1}$$

We omit the subscript Q when the context is clear.

We can now define the following formula that checks whether a *goal* region of the automaton state space is reachable after exactly k discrete transitions. We first state the simpler case of a hybrid system without invariants.

Definition 6 k -Step Reachability, Invariant-Free Case
Suppose H is an invariant-free hybrid automaton, U a subset of its state space represented by **goal**, and $M > 0$. The formula $\text{Reach}_{H,U}(k, M)$ is defined as:

$$\begin{aligned} & \exists^X \vec{x}_0 \exists^X \vec{x}_0^t \dots \exists^X \vec{x}_k \exists^X \vec{x}_k^t \exists^{[0,M]} t_0 \dots \exists^{[0,M]} t_k. \\ & \bigvee_{q \in Q} \left(\text{init}_q(\vec{x}_0) \wedge \text{flow}_q(\vec{x}_0, \vec{x}_0^t, t_0) \wedge \text{enforce}(q, 0) \right) \\ \wedge & \bigwedge_{i=0}^{k-1} \left(\bigvee_{q, q' \in Q} \left(\text{jump}_{q \rightarrow q'}(\vec{x}_i^t, \vec{x}_{i+1}) \wedge \text{enforce}(q, q', i) \right. \right. \\ & \quad \left. \left. \wedge \text{flow}_{q'}(\vec{x}_{i+1}, \vec{x}_{i+1}^t, t_{i+1}) \wedge \text{enforce}(q', i+1) \right) \right) \\ \wedge & \bigvee_{q \in Q} \left(\text{goal}_q(\vec{x}_k^t) \wedge \text{enforce}(q, k) \right) \end{aligned}$$

where $\exists^X x$ is a shorthand for $\exists x \in X$. Intuitively, the trajectories start with some initial state satisfying $\text{init}_q(\vec{x}_0)$ for some q . Then, in each step the trajectory follows $\text{flow}_q(\vec{x}_i, \vec{x}_i^t, t)$ and makes a continuous flow from $\vec{x}_i$ to $\vec{x}_i^t$ after time t . When the automaton makes a **jump** from mode q' to q , it resets variables following $\text{jump}_{q' \rightarrow q}(\vec{x}_k^t, \vec{x}_{k+1})$. The auxiliary **enforce** formulas ensure that picking $\text{jump}_{q \rightarrow q'}$ in the i -th step enforces picking $\text{flow}_{q'}$ in the $(i+1)$ -th step.

When the invariants are not trivial, we need to ensure that for all the time points along a continuous flow, the invariant condition holds. We need to universally quantify over time, and the encoding is as follows:

Definition 7 k -Step Reachability, Nontrivial Invariant
Suppose H contains invariants, and U is a subset of the state space represented by **goal**. The $\mathcal{L}_{\mathbb{R}\mathcal{F}}$ -formula $\text{Reach}_{H,U}(k, M)$ is defined as:

$$\begin{aligned} & \exists^X \vec{x}_0 \exists^X \vec{x}_0^t \dots \exists^X \vec{x}_k \exists^X \vec{x}_k^t \exists^{[0,M]} t_0 \dots \exists^{[0,M]} t_k. \\ & \bigvee_{q \in Q} \left(\text{init}_q(\vec{x}_0) \wedge \text{flow}_q(\vec{x}_0, \vec{x}_0^t, t_0) \wedge \text{enforce}(q, 0) \right. \\ & \quad \left. \wedge \forall^{[0,t_0]} t \forall^X \vec{x} (\text{flow}_q(\vec{x}_0, \vec{x}, t) \rightarrow \text{inv}_q(\vec{x})) \right) \\ \wedge & \bigwedge_{i=0}^{k-1} \left(\bigvee_{q, q' \in Q} \left(\text{jump}_{q \rightarrow q'}(\vec{x}_i^t, \vec{x}_{i+1}) \wedge \text{flow}_{q'}(\vec{x}_{i+1}, \vec{x}_{i+1}^t, t_{i+1}) \right. \right. \\ & \quad \wedge \text{enforce}(q, q', i) \wedge \text{enforce}(q', i+1) \\ & \quad \left. \left. \wedge \forall^{[0,t_{i+1}]} t \forall^X \vec{x} (\text{flow}_{q'}(\vec{x}_{i+1}, \vec{x}, t) \rightarrow \text{inv}_{q'}(\vec{x})) \right) \right) \\ \wedge & \bigvee_{q \in Q} \left(\text{goal}_q(\vec{x}_k^t) \wedge \text{enforce}(q, k) \right). \end{aligned}$$

The extra universal quantifier for each continuous flow expresses the requirement that for all the time points between the initial and ending time point ($t \in [0, t_i + 1]$) in a flow, the continuous variables $\vec{x}$ must take values that satisfy the invariant conditions $\text{inv}_q(\vec{x})$.