Optimal duration and planning of switching treatments taking drug toxicity into account a convex optimisation approach

Devia Pinzon, C.A.; Giordano, G.

DOI
10.1109/CDC40024.2019.9028881

Publication date
2019

Document Version
Accepted author manuscript

Published in
Proceedings of the 58th IEEE Conference on Decision and Control

Optimal duration and planning of switching treatments taking drug toxicity into account: a convex optimisation approach

Carlos Andrés Devia and Giulia Giordano

***Abstract*— We consider a multi-compartment evolutionary model representing growth, mutation and migration of cancer cells, as well as the effect of drugs, and we design optimal switching targeted cancer therapies where a single drug, or suitable drug combination, is given at each time so as to minimise not only the overall tumor size over a finite horizon, but also drug-provoked side effects. The strong diagonally-dominant structure of the model allows to solve the problem via convex optimisation. We provide an algorithm that yields optimality throughout the whole treatment duration by solving the convex optimisation problem with different horizons, and show how dwell time can be enforced via heuristics. Also the optimal treatment duration can be computed via convex optimisation. The proposed approaches are applied to a model of ALK-rearranged lung carcinoma.**

I. Introduction and Motivation

Designing optimised treatment protocols for HIV or cancer is an open challenge. Switching among different therapies allows to counteract resistance phenomena occurring in the presence of mutations, and reduce (or mitigate the increase of) the viral load or the tumor mass over a finite horizon. Treatment design can be cast as an optimal control problem associated with the evolutionary dynamics of the disease, where the state components represent the number of viruses or cancer cells of each mutant type, the inputs are the drug doses given at each time, while the function to be minimised is the total viral load or the total number of cancer cells.

Switching strategies to minimise the total HIV load over a finite horizon, in the presence of mutations, were designed in [13]; the associated cost functional is convex [3], [5]. Infinitely fast switching may be needed [10], but is clinically infeasible: each drug must be given for a minimum time before switching, which requires a dwell-time constraint on the switching rule [11]. The optimal control problem was also solved in a receding horizon fashion [12] and its robustness to uncertainty was investigated [7]. Treatment protocols that prevent or delay cancer progression and are robust to parameter uncertainty were planned via optimal and receding horizon control [2], [4] and gain scheduling [1]. For the multi-compartment evolutionary model in [9], [14], [15], describing replication, mutation and migration of cancer cells, and drug response, optimal switching treatments were designed in [9] via convex optimisation, exploiting convexity and monotonicity [3], [5], [16]. Since growth phenomena are dominant with respect to mutation and migration, the state matrix is practically diagonal; hence, the tumor size at the end of the treatment horizon is invariant with respect to the order in which the drugs (or suitable drug combinations) are given, and the variables of the convex optimisation problem can be the total times for which each drug is used [9].

In this paper, we consider the same multi-compartment evolutionary model as in [9], but we include in our novel cost not only the total amount of cancer cells over a finite horizon, but also an integral cost, to prevent the tumor mass from increasing too much during the transient, and a cost that represents drug-provoked side effects. Moreover, we address the problem of choosing an optimal treatment horizon. We consider both the general case, and the special diagonal case when only the final tumor size and the drug-provoked side effects are minimised. We show that:

- convexity of the optimisation problem for the design of the optimal switching treatment is preserved, in general, even when the more complex cost is considered;
- in the special diagonal case, the cost value depends exclusively on the overall treatment duration with each drug, regardless of the time instants when the different drugs are given: we can compute the optimal administration times for each drug via convex optimisation;
- in the special diagonal case, by solving the convex optimisation problem with different horizons, we can plan the treatment to ensure that, if it must be prematurely interrupted at time $\hat{t}$, the achieved cost is practically the optimum for time horizon $\hat{t}$: even though the state-dependent integral cost is not explicitly considered, we achieve optimality along the whole transient;
- we can compute the optimal treatment duration by iteratively solving a convex optimisation problem; in the special diagonal case, we can compute the optimal treatment duration and the optimal intervals for each drug by solving a single convex optimisation problem.

For the model of metastatic ALK rearranged lung adenocarcinoma considered in [9], our approach yields clinically feasible switching therapies that limit the overall disease progression as well as drug-provoked side effects.

II. Problem Formulation

We consider the multi-compartment evolutionary model [9], [14], [15]

$$\dot{x}(t) = \left[A - \sum_{s=1}^{d} D_s \ell_s(t)\right] x(t), \tag{1}$$

Research supported by the DTF Grant at TU Delft and by the Dutch Research Council through the NWO Talent Scheme grant VI.Veni.192.035. The authors are grateful to Anders Rantzer, Vanessa D. Jonsson, and Franco Blanchini for valuable discussions.

The authors are with the Delft Center for Systems and Control, Delft University of Technology, The Netherlands; {c.a.deviapinzon,g.giordano}@tudelft.nl

where x_i, $i = 1, \dots, n$, is the number of cells of line i, a certain mutant type located in a given body compartment, $A_{ii} \sim 10^{-1}$ is its growth rate, while the off-diagonal entries of A represent either mutation (if cell lines j and k are different mutants within the same compartment, $A_{jk} \sim 10^{-19}$ is the probability that a cell of line k mutates into a cell of line j), or migration (if cell lines j and k belong to the same mutant type but are in different body compartments, $A_{jk} \sim 10^{-27}$ is the probability that a cell of line k migrates to the same compartment where the cells of line j are located). Hence, $A \in \mathbb{R}^{n \times n}$ is a Metzler matrix representing tumor growth, mutation and metastasis in the absence of therapy. Matrices D_s are diagonal and $[D_s]_{ii} \geq 0$ represents the effect of drug s on the ith cell line, while

$$\ell_s \in \{0, 1\} \tag{2}$$

is the normalised concentration of drug s: either $\ell_s = 1$, if drug s is used at its maximum tolerated dose, or $\ell_s = 0$, if it is not used at all. A possible choice s could also be *no drug*, $D_s = 0$. At each time, the total maximum tolerated dose is

$$\sum_{s=1}^{d} \ell_s(t) = 1, \tag{3}$$

hence at most one drug can be given at a time. Another clinical requirement is that each ℓ_s is piecewise-constant in time, and is kept constant for a minimum time interval T_{min}.

Mutation and migration are strongly dominated by cancer growth rates: the off-diagonal entries of A are several orders of magnitude smaller than its diagonal entries. Therefore, in the sequel we *sometimes* consider a simplifying assumption.

Assumption 1: Matrix A is diagonal. ◇

Our first goal is to design an optimal switching therapy that, over a *given* finite horizon t_F, minimises:

- the final tumor size $h^\top x(t_F)$, with $h \in \mathbb{R}^n_+$ ($h = \mathbf{1}_n$);
- the side effects $\int_0^{t_F} g^\top \ell(\zeta) d\zeta$ arising due to drug toxicity where vector $g \in \mathbb{R}^n_+$ quantifies drug toxicity levels;
- a state-dependent integral cost $\int_0^{t_F} f^\top x(\zeta) d\zeta$, with $f \in \mathbb{R}^n_+$, which prevents the components of x from excessively increasing during the transient.

The resulting cost to be minimised is

$$J(\ell) = h^\top x(t_F) + c_g \int_0^{t_F} g^\top \ell(\zeta) d\zeta + c_f \int_0^{t_F} f^\top x(\zeta) d\zeta, \tag{4}$$

where c_g and c_f are positive scalar weights.

Remark 1: Healing does not require $\mathbf{1}^\top x(t_F) = 0$: when $\mathbf{1}^\top x(t_F)$ goes below a certain threshold, the therapy can be successfully interrupted and the immune system will destroy a small enough number of cancer cells. ◇

Given the integer constraint (2) and the sum constraint (3), we can equivalently rewrite system (1) as

$$\dot{x}(t) = A_{\sigma(t)} x(t) = [A - D_{\sigma(t)}] x(t), \tag{5}$$

where $\sigma(t) \in \mathcal{I} \doteq \{1, \dots, d\}$ for all t, with $\sigma(\tilde{t}) = m$ when drug m is used at time $\tilde{t}$, and a dwell time T_{min} is enforced.

Considering system (5) with dwell time corresponds to the exact clinical problem (C). When dwell time is not enforced, infinite-frequency switching is possible (S). Moreover, we can relax the integer constraint and consider $A(\ell) = A - \sum_{s=1}^{d} D_s \ell_s$ with continuous variables $\ell_s \geq 0$ (F). By definition of Filippov solutions [8], the set of the trajectories of (F) includes all the trajectories of (S), hence of (C), while the set of the possible trajectories of (C), with an arbitrarily small dwell time, is dense in the set of the trajectories of (F), hence all the trajectories of (F) can be approximated in the limit of infinite frequency switching (sliding modes). Therefore, the formulations (S) and (F) are equivalent, since $x(t) - x(0) = \int_0^t A_{\sigma(\zeta)} x(\zeta) d\zeta = \int_0^t A(\ell(\zeta)) x(\zeta) d\zeta$, and the optimal cost is also the same [5]. Although the clinical problem is (C), the solution of problem (F) achieved with $\ell_s(t)$ arbitrarily time-varying (corresponding to $\sigma(t)$ allowed to switch at infinite frequency: sliding modes, no dwell time) is interesting because it provides a lower bound for the cost.

First, in Section III, we solve problem (F), equivalent to (S), which turns out to be a convex optimisation problem; under Assumption 1 and when $f = 0$, we provide an explicit solution that already enforces the requirement not to mix different drugs. We also address heuristics to enforce the dwell time, thus solving the clinical problem 1, and provide a scheme to optimise the cost over the whole horizon $[0, t_F]$. Then, in Section IV, we consider the cost

$$J(t_F, \ell(t)) = h^\top x(t_F) + c_g \int_0^{t_F} g^\top \ell(\zeta) d\zeta + c_f \int_0^{t_F} f^\top x(\zeta) d\zeta \tag{6}$$

where the treatment horizon t_F is a *free variable*, and show that also the optimal treatment horizon can be found via convex optimisation (or analytically, in a special case).

III. Optimal Switching Treatments

We first neglect the dwell time and consider the relaxed problem with $\ell_s \geq 0$. Then, we need to solve the optimal control problem:

$$\min_{\ell(t) \in \mathcal{L}} \quad h^\top x(t_F) + c_g \int_0^{t_F} g^\top \ell(\zeta) d\zeta + c_f \int_0^{t_F} f^\top x(\zeta) d\zeta \tag{7}$$

$$\text{s.t.} \qquad \dot{x}(t) = A(\ell(t)) x(t), \quad x(0) = x_0, \tag{8}$$

where $\ell = [\ell_1 \dots \ell_d]^\top$, $A(\ell(t)) = A - \sum_{s=1}^{d} D_s \ell_s(t)$ and $\mathcal{L} = \{\ell : \mathbf{1}^\top \ell = 1 \text{ and } \ell_s \geq 0,\ s = 1, \dots, d\}$.

Remark 2: The evolutionary model (1) is a convex-monotone system [16] and the objective function is convex. In particular, convexity when the cost is a generic linear combination of the final state was proven in [3] and in Lemma 3, Lemma 4, Theorem 2 in [5]. However, as briefly suggested in [6, Section 2.1], the convexity properties are preserved even when more complex costs, such as (4), are considered. Hence, we can recast the optimal control problem as convex optimisation by augmenting the state. ◇

Define the augmented system

$$\dot{x}(t) = \left[A - \sum_{s=1}^{d} D_s \ell_s(t) \right] x(t) \tag{9}$$

$$\dot{\chi}(t) = c_f f^\top x(t) \tag{10}$$

$$\dot{\xi}(t) = 0 \tag{11}$$

$$\dot{\omega}(t) = g^\top \ell(t) \xi(t) \tag{12}$$

with $x(0)=x_0$, $\chi(0)=0$, $\xi(0)=c_g$ and $\omega(0)=0$, and the new state vector $z=[x^\top\ \chi\ \xi\ \omega]^\top$.

Theorem 1: The optimisation problem (7)–(8) can be equivalently rewritten as

$$\min_{\ell(t)\in\mathcal{L}} \widehat{h}^\top z(t_F) \text{ s.t. } \dot{z}(t)=H(\ell(t))z(t),\ z(0)=[x_0^\top\ 0\ c_g\ 0]^\top \tag{13}$$

where

$$H(\ell(t))=\begin{bmatrix} A(\ell(t)) & 0 & 0 & 0\\ c_f f^\top & 0 & 0 & 0\\ 0 & 0 & 0 & 0\\ 0 & 0 & g^\top\ell(t) & 0\end{bmatrix},\quad \widehat{h}=\begin{bmatrix}h\\1\\0\\1\end{bmatrix}.$$

Proof. Since $\xi(0)=c_g$ and ξ is constant in view of (11), integrating (12) gives $\omega(t_F)=c_g\int_0^{t_F} g^\top\ell(\zeta)d\zeta$. Also, integrating (10) yields $\chi(t_F)=c_f\int_0^{t_F} f^\top x(\zeta)d\zeta$. Therefore $\widehat{h}^\top z(t_F)$ is equal to the cost in (7). ■

Even with this more complex cost, we simply have to solve a convex programming problem, as in the case with the terminal cost only [3], [5], [6], [16]. Since matrices $H(\ell(t))$ do not commute, even when A is diagonal (this is intuitive: if the state transient is taken into account, then applying a drug sooner or later can affect the cost), a simple solution along the lines in [9] cannot be obtained.

Case $f=0$: If only the terminal cost and the therapy side-effects are considered, the optimal control problem is

$$\min_{\ell(t)\in\mathcal{L}} \quad h^\top x(t_F)+c_g\int_0^{t_F} g^\top\ell(\zeta)d\zeta \tag{14}$$

$$\text{s.t.}\quad \dot{x}(t)=A(\ell(t))x(t),\quad x(0)=x_0. \tag{15}$$

Define the augmented system

$$\dot{x}(t)=A(\ell(t))x(t) \tag{16}$$

$$\dot{\xi}(t)=0 \tag{17}$$

$$\dot{\omega}(t)=g^\top\ell(t)\xi(t) \tag{18}$$

with $x(0)=x_0$, $\xi(0)=c_g$, $\omega(0)=0$, and $y=[x^\top\,\xi\,\omega]^\top$.

Proposition 1: When $f=0$, the optimisation problem (14)–(15) can be equivalently rewritten as

$$\min_{\ell(t)\in\mathcal{L}} \tilde{h}^\top y(t_F) \text{ s.t. } \dot{y}(t)=M(\ell(t))y(t),\ y(0)=[x_0^\top\ c_g\ 0]^\top, \tag{19}$$

where

$$M(\ell(t))=\begin{bmatrix} A(\ell(t)) & 0 & 0\\ 0 & 0 & 0\\ 0 & g^\top\ell(t) & 0\end{bmatrix},\quad \tilde{h}=\begin{bmatrix}h\\0\\1\end{bmatrix}.$$

Proof. Since $\xi(0)=c_g$ and ξ is constant in view of (17), integrating (18) gives $\omega(t_F)=c_g\int_0^{t_F} g^\top\ell(\zeta)d\zeta$, hence $\tilde{h}^\top y(t_F)$ is equal to the cost in (14). ■

In general, therefore, we need to solve a convex programming problem [3], [5], [6], [16]. As in the simpler case in which just the final tumor size has to be minimised [9], when A is diagonal, we can exploit matrix commutation. Denote by M_k (respectively A_k) the matrix corresponding to $M(\ell)$ (respectively $A(\ell)$) when $\ell=e_k$, a vector of the Euclidean standard basis. Then we have the following result.

Theorem 2: When $f=0$, under Assumption 1, we can equivalently rewrite the optimisation problem (19) as

$$\min_{\tau_1,\dots,\tau_d} \tilde{h}^\top\prod_{s=1}^{d} e^{M_s\tau_s}y(0) \text{ s.t. } \tau_s\geq 0, s=1,\dots,d,\ \sum_{s=1}^{d}\tau_s=t_F.$$

Proof. It follows from the fact that, as can be easily verified, matrices M_s commute as long as matrices A_s commute, and matrices A_s commute if A is diagonal. ■

In view of the commutation property, the value of $y(t_F)$ does not depend on *when* each drug is given during the treatment horizon, but only on *how long* each drug is given. Thus, if the treatment duration with drug s is τ_s, the cost is $\tilde{h}^\top\prod_{s=1}^{d} e^{M_s\tau_s}y(0)=h^\top\prod_{s=1}^{d} e^{A_s\tau_s}x(0)+c_g\sum_{s=1}^{d} g_s\tau_s$. Of course treatment durations τ_s must be nonnegative and must sum up to the treatment horizon t_F.

Enforcing no-drug-mixing and dwell time: Solving the convex optimisation problems in Theorem 1 and Proposition 1 yields a minimising function $\ell(t)$ that is not piecewise constant and where, at each time instant, several entries can be nonzero. To enforce the clinical requirements, we can discretise as in [5] and divide the time interval $T=[0,t_F]$ into N intervals where $\ell(t)$ remains constant: $\ell(t)=\ell^{(i)}$ if $t\in T_i=[t_{i-1},t_i)$, for $i=1,\dots,N$, with $t_0=0$ and $t_N=t_F$, where all intervals are chosen much longer than T_{min}. Problem (13) then becomes

$$\min_{\ell^{(i)}\in\mathcal{L}} \quad \widehat{h}^\top\prod_{i=1}^{N} e^{H(\ell^{(i)})T_i}z(0), \tag{20}$$

with $z(0)=[x_0^\top\,0\,c_g\,0]^\top$. Problem (19) when $f=0$ becomes

$$\min_{\ell^{(i)}\in\mathcal{L}} \quad \tilde{h}^\top\prod_{i=1}^{N} e^{M(\ell^{(i)})T_i}y(0), \tag{21}$$

with $y(0)=[x_0^\top\ c_g\ 0]^\top$, or equivalently

$$\min_{\ell^{(i)}\in\mathcal{L}} \left\{h^\top\prod_{i=1}^{N} e^{A(\ell^{(i)})T_i}x(0)+c_g\left(\sum_{i=1}^{N} g^\top\ell^{(i)}T_i\right)\right\}, \tag{22}$$

with $x(0)=x_0$.

All the above are convex optimisation problems.

Each $\ell^{(i)}$ may be any vector with positive components summing up to one. To enforce the no-drug-mixing requirement, we can approximate the "sliding mode" by choosing each single drug for a fraction of the interval T_i corresponding to the closest multiple of T_{min}; this is why it was suggested to pick $T_i\gg T_{min}$.

Under Assumption 1, when $f=0$, Theorem 2 gives a solution that does not mix the drugs. To find a numerical solution in practice, we can solve the convex problem:

$$\min_{\tau_1,\dots,\tau_d}\left\{\sum_{i=1}^{n} h_i e^{K_i}x_i(t_o)+c_g\sum_{s=1}^{d} g_s\tau_s\right\} \tag{23}$$

$$\text{s.t.}\qquad \tau_s\geq 0,\ \forall s\in\{1,\dots,d\},\quad \sum_{s=1}^{d}\tau_s=t_F,\quad c_g>0,$$

where $K_i=A_{ii}t_F-\sum_{s=1}^{d}[D_s]_{ii}\tau_s$.

Then, to practically enforce a dwell time T_{min}, it is enough to set $\tau_s=T_{min}$ whenever the solution of the optimisation problem yields $0<\tau_s<T_{min}$, and decrease some other $\tau_k\gg T_{min}$ accordingly, so that $\sum_{s=1}^{d}\tau_d=t_F$.

Remark 3: The optimal therapy durations $\tau_1,\dots,\tau_d$ do not completely determine the exact therapy scheduling. We can choose the drug switching law to minimise the maximum number of cancer cells over the whole horizon, or the integral

cost. Given the number of intervals ($\approx t_F/T_{min}$), we can choose the best among all combinations that use drug s for a total time $\approx \tau_s$, according to the desired performance index. Another approach is suggested next. $\diamond$

Optimal drug scheduling with A diagonal and $f = 0$: Under Assumption 1 and when $f = 0$ in the cost (4), all the switching rules that use drug s for a total time τ_s, obtained by solving the optimisation problem in Theorem 2, lead to the same (optimal) cost. This gives a degree of freedom worth exploiting to optimally schedule drug administration. Here we aim at minimising the value of the cost function not only at the final time, but also throughout the whole time interval. So, if the treatment is interrupted at time $\hat{t} \in [0, t_F]$, the cost function value is as close as possible the optimal solution of the problem with horizon $\hat{t}$.

We divide the complete time interval $T = [0, t_F]$ into n subintervals $\{[0, t_{F1}], (t_{F1}, t_{F2}], \dots (t_{F(n-1)}, t_{Fn}]\}$, with $t_{Fn} = t_F$, and solve the optimisation problem (23) n times, by taking the final time of each subinterval as the treatment horizon: first we optimise over the horizon $[0, t_{Fn}]$, then over the horizon $[0, t_{F(n-1)}]$, and so forth, up to horizons $[0, t_{F2}]$ and $[0, t_{F1}]$. This yields a sequence of optimal drug delivery times $\tau_1[k], \dots, \tau_d[k]$, for $k = 1, \dots, n$. We go backward in time and we start by considering the n-th subinterval, $(t_{F(n-1)}, t_{Fn}]$. In the i-th subinterval, each drug needs to be given for a duration that can be computed as the difference $\Delta_s[i] = \tau_s[i] - \tau_s[i-1]$. It is not always possible exactly minimise the cost function at the end of each time interval, since it might be $\Delta_s[i] \leq 0$; hence a final step is added where the total drug delivery time is computed and corrected if necessary. The procedure is summarised in Algorithm 1.

Algorithm 1 Optimal Therapy and Drug Scheduling

1: Input: matrices A and D_s for $s = 1, \dots, d$, vectors h and g, real scalars t_F, c_g, and integer n.
2: Divide the whole time interval $T = [0, t_F]$ in n subintervals $\{[0, t_{F1}], (t_{F1}, t_{F2}], \dots (t_{F(n-1)}, t_{Fn}]\}$.
3: Set $k = n$.
4: Solve the optimisation problem (23) with final time t_{Fn}.
5: Set $\tau_s[n] = \tau_s$ for $s = 1, \dots, d$
6: **while** $k > 0$ **do**
7: Compute the optimal times $\tau_s[k-1]$ by solving the convex optimisation problem (23) with final time $t_{F(k-1)}$.
8: Compute for each drug: $\Delta_s[k] = \tau_s[k] - \tau_s[k-1]$
9: Truncate negative values: if $\Delta_s[k] < 0$, then $\Delta_s[k] = 0$
10: Update the complete optimal time for the next iteration: $\tau_s[k-1] = \tau_s[k] - \Delta_s[k]$
11: Choose the usage times for each drug in $(t_{F(k-1)}, t_{Fk}]$
12: Set $k = k - 1$
13: Add the drug usage times for each interval and, if necessary, increase the times in order to meet the total optimal times τ_s.

Of course the cost is exactly optimal only at the end of each interval $(t_{F(k-1)}, t_{Fk}]$, and not at any time $\hat{t} \in [0, t_F]$. To maximise the optimality of the drug scheduling policy over the whole horizon, we need to maximise the number of time intervals: the optimal number of subintervals is therefore $n_{\max}$, i.e. the largest number of intervals in which the treatment can be divided, given the dwell time constraint.

IV. Optimal Treatment Horizon

So far, the treatment horizon t_F was fixed *a priori*. In this section, we wish to find the optimal choice of t_F.

Considering a free-horizon problem makes sense only if each cancer cell line x_i is knocked down by at least one drug. Hence, in this section we make the following assumption.

Assumption 2: For $i = 1, \dots, n$, it is $[A(\ell)]_{ii} < 0$ for at least one choice of $\ell = e_k$. $\diamond$

If drug toxicity is taken into account, even when the treatment is effective in reducing the number of cancer cells, it is not true that the longer the treatment, the better: the longer the treatment lasts, the more the body is harmed by the drug side-effects. Therefore, we expect that an optimal treatment duration T_{opt} exists. We consider the cost (6) and choose a large enough t_L. Then, we wish to compute

$$J^* = \min_{t_F \leq t_L} \min_{\ell(t) \in \mathcal{L}} J(t_F, \ell(t)) \tag{24}$$

$$\text{s.t.} \quad \dot{x}(t) = A(\ell(t))x(t), \quad x(0) = x_0. \tag{25}$$

The value J^*, and the corresponding T_{opt}, can be computed in general by solving the inner convex optimisation problem (minimisation with respect to $\ell(t)$) for m finely sampled values of $0 < t_F \leq t_L$, and choosing the smallest; this requires solving m convex optimisation problems.

When $f = 0$ and A is diagonal, we can compute the optimal treatment horizon by solving a *single* convex optimisation problem; in a particular case, we even have an analytic expression.

The cost can be rewritten as $J(t_F, \ell(t)) = J(t_F, \tau)$, where $\tau = [\tau_1 \dots \tau_d]^\top \in \mathbb{R}^d$ and $\mathbf{1}^\top \tau = t_F$, and

$$J(t_F, \tau) = h^\top \prod_{s=1}^{d} e^{A_s \tau_s} x_0 + c_g \sum_{s=1}^{d} g_s \tau_s. \tag{26}$$

Proposition 2: Under Assumption 1, when $f = 0$ in the cost $J(t_F, \ell(t))$ as in (6), the optimisation problem (24)–(25) can be equivalently rewritten as

$$J^* = \min_{t_F \leq t_L} \min_{\tau_1, \dots, \tau_d} J(t_F, \tau) \tag{27}$$

$$\text{s.t.} \quad \tau_s \geq 0,\ s \in \{1, \dots, d\}, \quad \mathbf{1}^\top \tau = t_F. \tag{28}$$

Proof: It is an immediate consequence of Theorem 2. ■

Then, we can denote

$$J_{opt}(t_F) = \min_{\tau_1, \dots, \tau_d} J(t_F, \tau) \tag{29}$$

$$\text{s.t.} \quad \tau_s \geq 0,\ s \in \{1, \dots, d\}, \quad \mathbf{1}^\top \tau = t_F \tag{30}$$

and $T_{opt} = \arg\min_{t_F \leq t_L} J_{opt}(t_F)$. Since $J(t_F, \tau)$ is convex with respect to τ, we can show convexity of $J_{opt}(t_F)$.

Theorem 3: $J_{opt}(t_F)$ is a convex function of t_F.

Proof: Pick $t_A < t_B < t_C$, with $t_B = \alpha t_A + (1-\alpha) t_C$ for $0 < \alpha < 1$. Then, $J_{opt}(t_A) = J(t_A, \hat{\tau}_A)$ for some $\hat{\tau}_A$ such that $\mathbf{1}^\top \hat{\tau}_A = t_A$ and $J_{opt}(t_C) = J(t_C, \hat{\tau}_C)$ for some $\hat{\tau}_C$ such that $\mathbf{1}^\top \hat{\tau}_C = t_C$. We can choose $\tilde{\tau}_B = \alpha \hat{\tau}_A + (1-\alpha)\hat{\tau}_C$, so that $\mathbf{1}^\top \tilde{\tau}_B = t_B$. If the optimal $J_{opt}(t_B)$ is achieved for $\tau = \hat{\tau}_B$, then $J_{opt}(t_B) = J(t_B, \hat{\tau}_B) \leq J(t_B, \tilde{\tau}_B) \leq \alpha J(t_A, \hat{\tau}_A) + (1-\alpha) J(t_C, \hat{\tau}_C) = \alpha J_{opt}(t_A) + (1-\alpha) J_{opt}(t_C)$, where the second inequality holds since $J(t_F, \tau)$ is convex with respect to τ. ■

This implies that the minimum T_{opt} is global. We now show that problem (27)–(28) can be equivalently solved as a *single* convex optimisation problem. We consider the system

$$\dot{x}(t) = A(\hat{\ell})x(t) = \left[A - \sum_{s=1}^{d+1} D_s \hat{\ell}_s(t)\right] x(t), \qquad (31)$$

where $D_{d+1} = A$. We augmented the number of choices, by introducing the "dummy" choice $d+1$ that *freezes the time*, since the corresponding state matrix is zero. Given this system, under Assumption 1 and $f = 0$, we can find the optimal time for each choice as per Theorem 2. Consider the augmented vector $\hat{\tau} = [\hat{\tau}_1 \dots \hat{\tau}_d \ \hat{\tau}_{d+1}]^\top \in \mathbb{R}^{d+1}$ and the convex optimisation problem

$$\tilde{J} = \min_{\hat{\tau}_1,\dots,\hat{\tau}_d,\hat{\tau}_{d+1}} J(t_L, \hat{\tau}) \qquad (32)$$

$$\text{s.t.} \quad \hat{\tau}_s \geq 0,\ s \in \{1,\dots,d,d+1\}, \quad \mathbf{1}^\top \hat{\tau} = t_L. \qquad (33)$$

Theorem 4: The solution J^* of the optimisation problem (27)–(28) is equal to the solution $\tilde{J}$ of the optimisation problem (32)–(33): $J^* = \tilde{J}$.

Proof: In problem (32)–(33) the time is freezed for a portion $\hat{\tau}_{d+1}$ of the horizon t_L. If the optimal J^* is obtained for $t_F = t_F^*$, with $\tau = \tau^*$, then $\tilde{J}$ can be rendered at least as small by choosing $\hat{\tau} = [(\tau^*)^\top \ t_L - t_F^*]^\top$; hence, $\tilde{J} \leq J^*$. If the optimal $\tilde{J}$ is obtained with $\hat{\tau} = \hat{\tau}^*$, then J^* can be rendered at least as small by choosing $t_F = t_L - \hat{\tau}_{d+1}^*$ and setting $\tau_i = \hat{\tau}_i^*$ for $i = 1,\dots,d$; hence, $J^* \leq \tilde{J}$. Therefore, the two problems have the same solution. ■

Hence, a single convex optimisation problem (32)–(33) can be solved to find $T_{opt} = t_L - \hat{\tau}_{d+1}^*$.

Under suitable assumptions, we can directly compute the optimal treatment durations for each drug, without introducing the "dummy" mode. First note that the cost $J(\tau_1,\dots,\tau_d) = h^\top \prod_{s=1}^d e^{A_s\tau_s} x_0 + c_g \sum_{s=1}^d g_s\tau_s = \sum_{j=1}^n h_j x_{0,j} e^{\sum_{s=1}^d [A_s]_{jj}\tau_s} + c_g \sum_{s=1}^d g_s\tau_s$, where the sum of the τ_i's is *unconstrained*, is convex. Then:

Lemma 1: The function $J(\tau_1,\dots,\tau_d)$ admits a minimum with $\tau_i > 0$ for all i if and only if $\frac{\partial J(\tau_1,\dots,\tau_d)}{\partial \tau_i} = 0$ for all i.

Even when there is no *internal* minimum (with all $\tau_i > 0$), a constrained gradient method can be used to find the global minimum of the convex function.

If there is an internal minimum, we can give an analytical expression for the optimal τ_i's when the number of drugs is equal to the number of cell lines, $d = n$. The expression is based on the solution of two nested linear systems, associated with matrices $D \in \mathbb{R}^{n\times d}$ and $S \in \mathbb{R}^{d\times n}$, where $[D]_{ij} = [A_j]_{ii}$ and $[S]_{ij} = h_j x_{0,j} [A_i]_{jj}$.

Theorem 5: Assume that matrices D and S, with $n = d$, are nonsingular. Under Assumption 1, when $f = 0$ in (6), if there exists a minimum of the cost $J(\tau_1,\dots,\tau_d)$ with all $\tau_i > 0$, then such an optimal choice of $\tau = [\tau_1 \dots \tau_d]^\top$ is obtained as the solution of two nested algebraic systems: $D\tau = \log v$, where $v \in \mathbb{R}_+^n$ is the positive solution (if any, otherwise no optimal choice of τ exists) of $Sv = b$, with $b \in \mathbb{R}^d$, $b_i = -c_g g_i$; namely, $\tau^* = D^{-1} \log[S^{-1} b]$. Then, $T_{opt} = \mathbf{1}^\top \tau^*$.

Proof: The cost function $J(\tau_1,\dots,\tau_d)$ is convex. To find its internal minimum, we can look at the choice of the τ_i's that zeroes all its partial derivatives $\frac{\partial J(\tau_1,\dots,\tau_d)}{\partial \tau_i} = \sum_{j=1}^n h_j x_{0,j} [A_i]_{jj} e^{\sum_{s=1}^d [A_s]_{jj}\tau_s} + c_g g_i$. By setting $v_j = e^{\sum_{s=1}^d [A_s]_{jj}\tau_s}$, $b_i = -c_g g_i$, and $[S]_{ij} = h_j x_{0,j} [A_i]_{jj}$, this amounts to solving the algebraic system $Sv - b = 0$. Once we have found the solution v, we can compute the τ_i's by solving a second algebraic system, $D\tau = \log v$, where $\tau = [\tau_1 \dots \tau_d]$ and $[D]_{ij} = [A_j]_{ii}$. ■

V. Case Study: ALK-rearranged Lung Carcinoma

We consider a model of metastatic ALK rearranged non-small cell lung cancer [9], and show that the proposed convex optimisation approach yields switching therapies that optimally control tumor progression, in spite of imperfect drug penetration in the two different compartments involved (the lung and the brain) and heterogeneity (nine different mutant types in each compartment). It is possible to switch among four targeted therapies: *crizotinib*, *alectinib*, *ceritinib* and *lorlatinib*. Unfortunately, many mutations have shown to be resistant to at least one of these therapies [9]; therefore switching is crucial to achieve the best possible treatment.

In our numerical simulations, we take matrices A and D_s as in [9] and consider the initial condition $\hat{x}_0 = [956, 5, 5, 5, 5, 5, 5, 4, 10, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]^\top \cdot 10^7$. We set $h = \mathbf{1}^\top$, $g = [1 \cdot 10^6, 5 \cdot 10^4, 1 \cdot 10^5, 1 \cdot 10^6]$, $c_g = c_f = 0.7$, and $t_F = 100$ days. To apply Algorithm 1, as well as to compute the solution of (22), the horizon t_F is divided into $n = 10$ subintervals of equal duration (10 days). The enforced dwell time is 1 day.

Figure 1A shows the drug scheduling policy obtained with Algorithm 1. The switching law becomes periodic in the second half of the time interval; this behaviour can be explained by introducing the variables $\delta_s = \frac{\tau_s}{t_F}$. Since each subinterval in Algorithm 1 has the same duration, these variables represent the 'duty cycle' of the control signal. With this change of variables, the optimisation problem (23) becomes $\min_{\delta_1,\dots,\delta_d} \left\{ \sum_{i=1}^n h_i e^{t_F} e^{\hat{K}_i} x_i(t_o) + c_g t_F \sum_{s=1}^d g_s \delta_s \right\}$ s.t. $\delta_s \geq 0$, $\forall s \in \{1,\dots,d\}$, $\sum_{s=1}^d \delta_s = 1$, $c_g > 0$, where $\hat{K}_i = A_{ii} - \sum_{s=1}^d (D_s)_{ii} \delta_s$. For large values of t_F the first term dominates the cost function; the constant e^{t_F} multiplies the cost function and the solutions $\{\delta_s\}_{s=1}^d$ converge to a constant value as $t_F \to \infty$. Figure 1B compares the optimal times τ_s achieved by solving (23) and those obtained with Algorithm 1, for different values of t_F. When t_F increases, the two values become closer, as expected. Also during the complete control interval both values remain close, meaning that the drug scheduling policy given by the algorithm is close to the optimal throughout the whole time interval, as desired. Figure 1C compares the time evolution of the total cost obtained with Algorithm 1 (solid) and with the optimisation problem (22) (dotted). The drug scheduling policy provided by Algorithm 1 minimises the cost function not only for $t_F = 100$ days but also for earlier values of time. This ensures that, if the therapy is prematurely interrupted, the

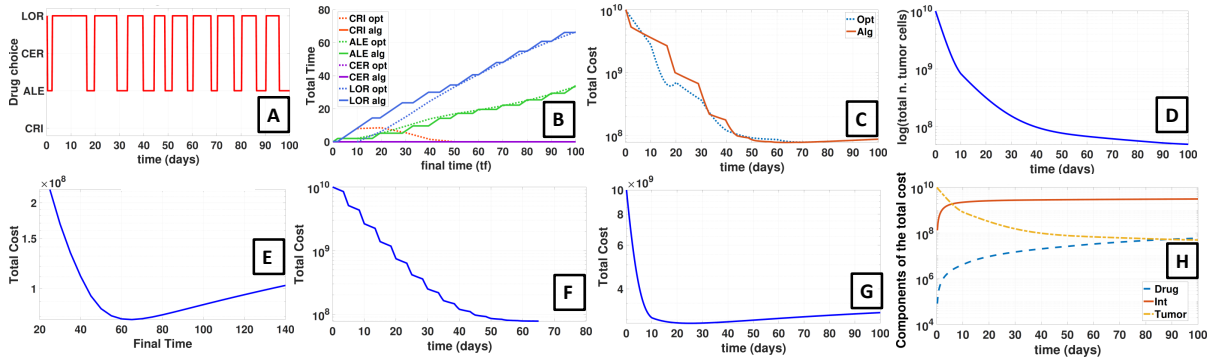

Fig. 1. *Optimal therapy with Algorithm 1 and initial condition* $\hat{x}_0$*:* (A) Drug choice with Algorithm 1 as a function of time. (B) Comparison between the optimal drug times obtained by solving (23) (*opt*) and the times given by Algorithm 1 (*alg*). Red and violet lines are overlapped. (C) Comparison between the cost achieved with Algorithm 1 (*Alg*) and with the optimisation problem (22) (*Opt*). (D) Number of total tumor cells for the initial condition $\hat{x}_0$. *Finding the optimal treatment horizon and the associated therapy:* (E) Optimal cost obtained by solving the optimisation problem (23) for different values of the horizon t_F. (F) Time evolution of the total cost for the optimal treatment horizon $t_F^* = 64.24$, obtained with Algorithm 1 with $n = 13$ equal subintervals. *Optimal therapy with the full cost functional in* (20)*:* (G) Time evolution of the total cost in (20). (H) Time evolution of the unweighted components of the total cost in (20): drug-associated cost (*Drug*), integral cost (*Int*) and total number of tumor cells (*Tumor*).

best possible outcome is obtained, and that the total number of tumor cells and drug-provoked side effects are minimised throughout the treatment duration. Comparing these results with those obtained by solving (22) (dotted), which *enables drug-mixing* (and hence provides the best result one could get), the cost value at $t_F = 100$ days is the same. Throughout the whole time interval, the proposed algorithm approaches the optimal solution and also guarantees patient safety, by enforcing no-drug-mixing. The corresponding number of total tumor cells over time is shown in Figure 1D. As for the optimal treatment duration, Figure 1E shows the optimal cost of problem (23) for different values of t_F (time horizon). Solving the augmented optimisation problem (32) yields the optimal treatment duration $t_F^* = 64.24$, in agreement with Figure 1E. The corresponding optimal cost evolution, computed via Algorithm 1, is shown in Figure 1F. All the previously presented results consider the cost with $f = 0$, i.e. without the state-dependent integral cost. The optimal solution computed, when $f = \mathbf{1}^\top \cdot 10^{-1}$, by solving the optimisation problem (20), which includes the full cost functional, is shown in Figure 1G. Figure 1H shows the (unweighted) components of the total cost: drug-associated cost, state-dependent integral cost and total number of tumor cells. The dominant term in the cost function is initially related to the total number of tumor cells; however, as the treatment reduces the number of tumor cells, the dominant term becomes the state-dependent integral. In this particular case (being the components of f relatively large), the drug-associated cost is secondary throughout the whole interval.

VI. Conclusions

We showed convexity of the design of optimal switching cancer therapies that minimise the tumor size, the tumor integral and drug-provoked side effects; convexity enables effective solutions when mutation and migration phenomena are negligible with respect to cancer growth. When the integral cost is neglected, we proposed an explicit algorithm that, by solving n convex optimisation problems, minimises the cost functional throughout the whole treatment horizon, thus implicitly limiting tumor growth in the transient. We showed that the optimal treatment duration can be computed via convex optimisation.

References

[1] M. Alamir, "Robust feedback design for combined therapy of cancer," *Optimal Control Applications and Methods*, 2012.

[2] M. Alamir and S. Chareyron, "State-constrained optimal control applied to cell-cycle-specific cancer chemotherapy," *Optimal Control Applications and Methods*, vol. 28, no. 3, pp. 175–190, 2007.

[3] F. Blanchini, P. Colaneri, and R. H. Middleton, "A convexity result for the optimal control of a class of positive nonlinear systems," in *19th IFAC World Congress*, 2014.

[4] S. Chareyron and M. Alamir, "Model-free feedback design for a mixed cancer therapy," *Biotech. progress*, vol. 25, no. 3, pp. 690–700, 2009.

[5] P. Colaneri, R. H. Middleton, Z. Chen, D. Caporale, and F. Blanchini, "Convexity of the cost functional in an optimal control problem for a class of positive switched systems," *Automatica*, vol. 50, no. 4, pp. 1227–1234, 2014.

[6] P. Colaneri, R. H. Middleton, and F. Blanchini, "Optimal control of a class of positive Markovian bilinear systems," *Nonlinear Analysis: Hybrid Systems*, vol. 21, pp. 155–170, 2016.

[7] M. Colombino, N. Dhingra, M. Jovanovic, and R. Smith, "Convex reformulation of a robust optimal control problem for a class of positive systems," in *IEEE Conf. Dec. Control*, 2016, pp. 5263–5268.

[8] A. F. Filippov, *Differential equations with discontinuous right-hand sides*. Kluwer Academic Publishers Group, 1988.

[9] G. Giordano, A. Rantzer, and V. Jonsson, "A convex optimization approach to cancer treatment to address tumor heterogeneity and imperfect drug penetration in physiological compartments," in *IEEE Conf. Dec. Control*, 2016, pp. 2494–2500.

[10] E. A. Hernandez-Vargas, P. Colaneri, and R. H. Middleton, "Optimal therapy scheduling for a simplified HIV infection model," *Automatica*, vol. 49, no. 9, pp. 2874–2880, 2013.

[11] ——, "Sub-optimal switching with dwell time constraints for control of viral mutation," in *IEEE Conf. Dec. Control*, 2012, pp. 4906–4911.

[12] ——, "Switching strategies to mitigate HIV mutation," *IEEE Trans. Control Systems Technology*, vol. 22, no. 4, pp. 1623–1628, July 2014.

[13] E. Hernandez-Vargas, P. Colaneri, R. Middleton, and F. Blanchini, "Discrete-time control for switched positive systems with application to mitigating viral escape," *International Journal of Robust and Nonlinear Control*, vol. 21, no. 10, pp. 1093–1111, 2011.

[14] V. Jonsson, N. Matni, and R. M. Murray, "Reverse engineering combination therapies for evolutionary dynamics of disease: An H-inf approach," in *IEEE Conf. Dec. Control*, 2013, pp. 2060–2065.

[15] V. Jonsson, A. Rantzer, and R. M. Murray, "A scalable formulation for engineering combination therapies for evolutionary dynamics of disease," in *American Control Conference, (ACC)*, 2014, pp. 2771–2778.

[16] A. Rantzer and B. Bernhardsson, "Control of convex monotone systems," in *IEEE Conf. Dec. Control*, 2014, pp. 2378–2383.