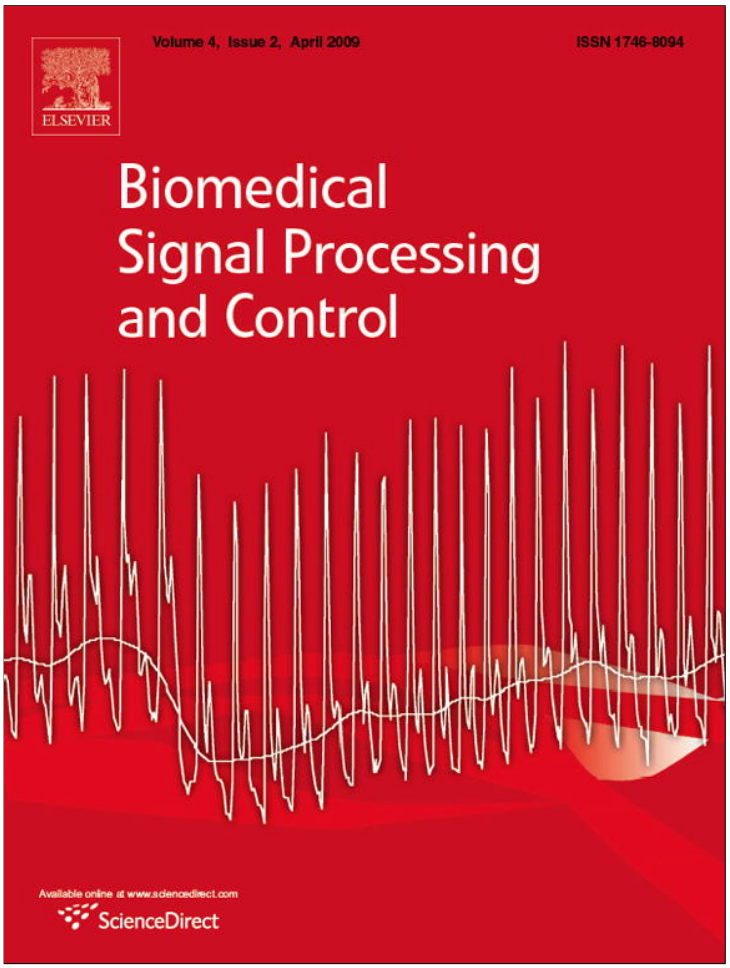

Biomedical Signal Processing and Control 4 (2009) 139–148

Contents lists available at ScienceDirect

Biomedical Signal Processing and Control

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

A closed-loop approach to antiretroviral therapies for HIV infection

V. Costanza [c,a,*], P.S. Rivadeneira [c], F.L. Biafore [a,b], C.E. D'Attellis [a,b]

[a] Centro de Matemática Aplicada, Escuela de Ciencia y Tecnología, UNSAM, M. de Irigoyen 3100, 1650 San Martín, Pcia. de Buenos Aires, Argentina
[b] Departamento de Matemática, Facultad de Ingeniería, Universidad Favaloro, Avda. Belgrano 1723, 1093 Buenos Aires, Argentina
[c] Grupo de Sistemas No Lineales, INTEC(UNL-CONICET), Güemes 3450, 3000 Santa Fe, Argentina

ARTICLE INFO

Article history:
Received 30 October 2008
Received in revised form 16 February 2009
Accepted 17 February 2009
Available online 27 March 2009

Keywords:
HIV
Antiretroviral drugs
Optimal control
Nonlinear dynamics
Dynamic programming

ABSTRACT

Antiretroviral therapies allowing for regular medical intervention in an optimal feedback role are substantiated. Individual therapies are evaluated from a multi-objective cost perspective, under different time and history constrains. The dynamics of the control system is governed by three state variables: healthy T-cells, infected T-cells, and viral particles. The drug dose administrated to the patient is taken as the manipulated or control variable. Illustrations of the methodology employed are provided for a fixed time-horizon of 180 days and variable prescription intervals, inverse discount factors, state discretization, and thresholds. Two cases are discussed: (i) beginning of the infection and (ii) near endemic equilibrium. Both open-loop and closed-loop results are presented. Dynamic Programming has been employed in the numerical treatment of the problem. Safety recommendations are also designed to cope with the chaotic behavior of the free dynamics' flow in critical regions of the states' domain. A software package is envisioned to assist physicians in assessing the patient "state", estimating antiretroviral dose prescriptions for the whole treatment period, simulating the patient's evolution, eventually correcting the original sequence in presence of incoming data, and evaluating combined costs of alternative treatment strategies.

1. Introduction

The treatment of HIV/AIDS has evolved in the last 20 years since the upsurge of its epidemic quality, from no treatment to highly active antiretroviral therapy (HAART), which has radically changed the face of the disease [14,15]. HAART has clearly shown to decelerate AIDS progression and to extend the life of the patient. Suppressing viral replication with HAART allows the body time to rebuild its immune system. However, despite the clinical improvement associated with HAART, current antiviral drug regimens are not able to eradicate HIV. Note that a curative treatment of HIV/AIDS is not possible at the present time, and that whenever HAART is stopped, then after some time HIV becomes detectable in the blood once again [14,15,28].

The conventional treatment consists of drugs that have to be taken every day for the rest of someone's life. People taking antiretroviral drugs may have low adherence to complicated drug regimens. Current recommended regimens involve taking several antiretroviral drugs each day from at least two different classes, some of which may cause unpleasant side effects such as nausea and vomiting, making long-term, continuous therapy impractical for many HIV-infected individuals [28]. About 25% of patients stop therapy within the first year on HAART because of side effects. As a result, treatment of HIV infection has become a complicated balancing act between the benefits of durable HIV suppression and the risks of drug toxicity [14,15,28]. Therefore, the research in more effective drug regimens has become an important challenge.

Mathematical modeling has made a substantial impact on our thinking and understanding of HIV-1 infection. A great number of deterministic models have been developed to describe the immune system and its interaction with HIV-1 as well as the effects of drug therapy [20–22,24,17,18,25,8,22,1,4,16,11]. In most cases, their mathematical expression are based in relatively complex systems of non-linear differential equations. Population models are most commonly used.

Using these models, therapeutic strategies can be simulated with the purpose of reducing the viral load. Also, different control problems can be posed and solved when the dynamics is known. For instance, state feedback and optimal control of HIV infection have been explored in several other works [18,1,4,16,11,13,3]. The dynamic contents of these theoretical problems usually imply the appearance of variable instead of constant doses administration. The variable drug dosage has been questioned by some authors due to the possibility of generating drug resistant mutants. However, recent empirical studies indicate that the relationship between adherence (i.e. keeping rigidly constant the dose all the time by

* Corresponding author.
E-mail address: tsinoli@santafe-conicet.gov.ar (V. Costanza).

1746-8094/$ – see front matter © 2009 Elsevier Ltd. All rights reserved.
doi:10.1016/j.bspc.2009.02.005

prescription) and the appearance of HIV drug resistance, is more complicated than assumed initially. For some regimens, resistance may be more likely to occur in those patients taking more rather than less of their medications. For other therapies, the opposite may be true. Recent data indicate that each antiretroviral therapeutic class has a unique adherence–resistance relationship [2]. Cases of interrupted medication (significant periods with zero drug dose) have been reported [12], not implying the prescription of zero control during part of a recommended protocol. In short, no conclusive results concerning these problems have reached generalized consensus. Therefore, due to the adverse effects and cost of the antiretroviral therapy, the optimal administration of drugs considering multiple objectives (besides the reduction of viral load), becomes particularly constructive, and it will be pursued here.

The continuous-time Pontryagin approach to optimal control is frequently called upon in this context [18,11]. But these problems are often quite difficult to solve numerically due to the two point boundary value problem resulting in the adjoint variables. Also, the eventual finding of a continuous optimal control strategy is not quite suitable for practical implementation: neither blood analyses can be continuously made nor changes in the dose consequently prescribed, the amount of the drug is not freely divisible, and provisions for departures from the optimal trajectory are not implicit in the solution. Discrete-time versions of this approach (see for instance [5]), offer some positive and negative features, as will be illustrated below.

Another optimization tool promoted in recent times is 'model predictive control' (MPC, see for instance [7]), based on the so called receding horizon philosophy: a sequence of future control actions is chosen according to a prediction of the future evolution of the model and applied to the system until new measurements are available. At that time, based on the new measurement, a new sequence is established which replaces the previous one. This method is essentially numerical, usually implemented on-line, which requires significant computer capacity and speed. Off-line methods include 'Dynamic Programming' (adopted here), in its continuous or discrete-time versions, which allow for state-feedback control, a decisive capacity when preserving the 'medical control' over the 'automatic pilot' point of view.

The paper is organized as follows: in Section 2 the continuos-time model is described and some of the features of the free-dynamics behavior are discussed and illustrated, and in Section 3 the introduction of an appropriate control variable is substantiated. Next Section 4 is dedicated to designing a cost objective functional taking into consideration not only the burden of drugs application but also the patient's health evolution during treatment, the opportunity of its improvement, and the practical restrictions in measuring and applying controls that lead to a discretized version of the problem. Numerical results are showed and discussed in Section 5, especially for treating a 'recent' patient (comparing the open and closed-loop approaches) and for an 'endemic' situation. The last Section 6 presents the conclusions.

2. Description of the 3rd-order model of HIV dynamics

HIV is a very complex disease involving multiple interactions between the virus and the host immune system. However, the main characteristics of the infection can be ascertained from relatively simple equations [22]. This basic model has three variables: healthy (CD4+) T-cells (x), infected (CD4) T-cells (y) and free virus copies or virions (z). It is a one-compartment model (the compartment is the blood), where healthy cells encounter free virions and become infected cells and the rate of production of new infected cells is proportional to the product of the density of uninfected cells times the density of virions. The dynamics consists of a system of three nonlinear ordinary differential equations (ODEs):

$$\begin{aligned} \dot{x} &= \lambda - \delta x - \beta xz \\ \dot{y} &= \beta xz - \mu y \\ \dot{z} &= \kappa y - \gamma z \end{aligned} \tag{1}$$

Sometimes a concise form for these equations is used, namely

$$\begin{aligned} X &\triangleq (x, y, z)' \\ \dot{X} &= f(X), \end{aligned} \tag{2}$$

where f is the vector field given by the right-hand-side of Eq. (1).

The 'population dynamics' point-of view is implicit in this model structure. The parameter λ represents the constant rate of production of the uninfected CD4+ T-cells from the thymus. The infection rate is represented by the weighted product of 'predator-prey' populations βxz, the weight β interpreted as a 'constant rate' of infection. The term δx is the death rate of healthy T-cells. From the uninfected T-cells, the productively infected T-cells y are produced at the same rate βxz. The respective death rate of y-cells is μy. It is also assumed that each productively infected cell y produces κ viral copies per unit of time, and that the death rate of virions is of the form γz.

The numerical simulation of the model illustrates the evolution of infection in an individual human without the intervention of antiretroviral drugs. It is found that the system has two equilibrium points. The infection-free (or 'healthy') equilibrium point $(x_h, y_h, z_h) = (\lambda/\delta, 0, 0)$ and the unique 'endemic' equilibrium $(x_e, y_e, z_e) = (\mu\gamma/\beta\kappa, \lambda - \delta x_e/\mu, \kappa y_e/\gamma)$. Initial conditions near the healthy (unstable) equilibrium evolve towards the endemic (stable) equilibrium. The therapeutic (or control) process consists basically in reversing this propensity of the flow, i.e. in driving the organism as close as possible to the (unstable) healthy equilibrium.

Several simulated trajectories are shown in Figs. 1 and 2, where the time parameterizes each curve. The rays in the plot of Fig. 1 join points of equal time in trajectories starting from different initial conditions and $x_0 \geq 500$. Fig. 2 shows only a few individual trajectories, including some with $x_0 \leq 500$ and no synchronizing rays. The successive stages of HIV-1 infection can be distinguished in the flow. The first stage, called the acute infection, is characterized by an initial overshoot in the productively infected cells y and the viral load z. This was also better depicted through the $x(t), y(t)$, and $z(t)$ trajectories individually [16,4]. The healthy CD4+ cells x decline from 1000 cells per cubic millimeter to 400 in this stage. After the viral peak has been reached, the uninfected cell population increases but finally it is stabilized in a lower value. A

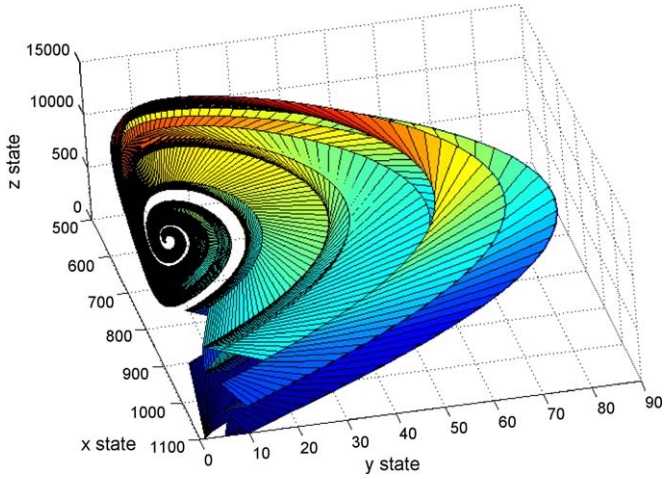

Fig. 1. Flow of free dynamics ($u \equiv 0$), with rays connecting synchronous points in different trajectories.

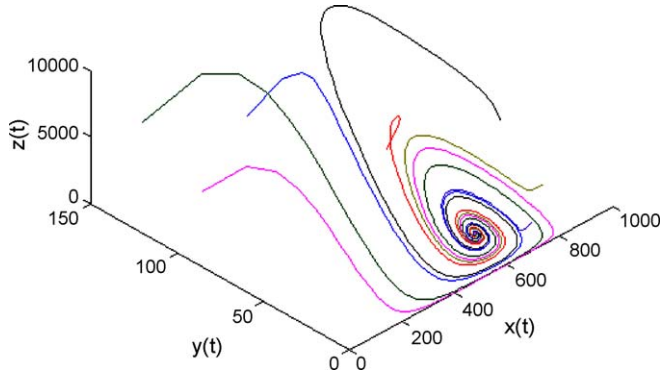

Fig. 2. Individual trajectories of the free dynamics in phase space.

set point is virtually reached for all variables after about 400 days. This stage is known as the asymptomatic period of AIDS disease. See Ref. [22] for the complete analysis of this model, where it is pointed out that this type of predator-prey interaction, reducing the abundance of predators (virions), may cause an increase in the number of prey (uninfected CD4+ T-cells), which in turn causes predators to rise again, leading to the observed oscillations.

The parameters of the model have been estimated by different authors (see [23,25,27], etc.). The nominal values range around those adopted here

$\lambda = 9$
$\delta = 0.009$
$\beta = 4 \times 10^{-6}$
$\mu = 0.3$
$\kappa = 80$
$\gamma = 0.6$

The values of these parameters should be updated by using further experimental data. One of the tools helping to design these experiments for better accuracy is the calculations of sensitivities. The 'sensitivity' S_β of the model with respect to the parameter β is usually defined by the n-column vector

$$S_\beta \triangleq \frac{\partial \phi}{\partial \beta} \tag{3}$$

where $\phi(t, X, P)$ is the 'flow' of the vector ODE (2), where the presence of parameters is made explicit through the vector $P = (\lambda, \delta, \beta, \mu, \kappa, \gamma)$ in $f(X, P)$. The solutions of the ODEs are condensed in ϕ, i.e.

$$\frac{\partial \phi}{\partial t}(t, X, P) = f(\phi(t, X, P)), \tag{4}$$

which allows to find [10] an ODE for the matrix $S_P = (S_\lambda, S_\delta, S_\beta, S_\mu, S_\kappa, S_\gamma)$, namely

$$\dot{S}_P = \frac{\partial f}{\partial X}(\phi(t, X, P), P) \cdot S_P + \frac{\partial f}{\partial P}(\phi(t, X, P), P) \tag{5}$$
$$S_P(0) = 0. \tag{6}$$

The qualitative behavior of the solutions to (5) is most valuable to determine the best time to record experimental measurements, as will be seen in the next section.

3. Curtailing the infection. The therapy process as a control system

Basically, the antiretroviral drugs can be grouped into the following three categories [15]:

(i) Inhibitors of the reverse transcriptase enzyme (RTIs): If RT is inhibited, HIV can enter a cell but will not successfully infect it; a DNA copy of the viral genome will not be made and the cell will not make viral proteins.
(ii) Protease inhibitors (PIs): If HIV protease is inhibited, cleavage of the viral polyprotein will not occur, and viral particles that lack functional enzymes will be made. The net effect of blocking HIV protease is that noninfectious viral particles are made.
(iii) Fusion Inhibitors (FIs): These work by inhibiting the binding of HIV to uninfected CD4+T-cells (used in patients with multi-drug HIV resistance).

In this work we will only consider the action of RTIs. Although the medication process is clearly viewed as a control action exerted over the patient's organism (the 'plant'), the physical identification between manipulated and control variable has not always been clear. Most scientists work with drug efficacy (between 0 and 1) as the input (control), equivalent to a coefficient multiplying the parameter β [18,3,16,25,11,13]. But in absence of an adequate pharmaco-dynamical model relating efficacy to real dose, the optimization results obtained this way apply to abstract efficacies, entailing only vague guidelines for drug prescription. For this reason, in the mathematical model used here, the control u represents the amount of drug (i.e. the dose) manipulated by clinicians. Consequently, the parameter β in Eqs. (1) and (2) should be thought as a function of the doses u (without loss of generality, a power series). Here an approximation to this form is adopted, namely:

$$\beta \approx \beta_0 - \alpha_1 u - \alpha_2 u^2 \tag{7}$$

where the values of the parameters were estimated from experimental data, through standard least-squares regression techniques, and under the following assumptions:

(i) HIV-1 viral load reduction curves for patients under mono-therapy (a daily dose of $u = 1.2$ g of aprecitabine, an RTI in clinical development) can be approximated from clinical data [6] and the observations quoted in [25], namely:
(ii) For a short period after therapy has begun, x remains approximately constant (say $x(t) \approx \tilde{x} \triangleq x(0)$).

Then, the equations for $\dot{y}$ and $\dot{z}$ in the model (1) become linear and their solution imply:

$$z(t) \approx (k_1 e^{a_1 t} + k_2 e^{a_2 t}) z(0), \tag{8}$$

with coefficients depending on $x(0), y(0), \beta_0, \alpha_1, \alpha_2$. Data extracted from the literature [6,25] correspond to treatments of recently discovered infections, resulting in the following estimates

$$\beta_0 = 4 \times 10^{-6} \tag{9}$$

$$\alpha_1 = 0.88 \times 10^{-6} \tag{10}$$

$$\alpha_2 = 0.3 \times 10^{-6} \tag{11}$$

Fig. 3 depicts the sensitivity of the system with respect to parameters β_0, α_0, and α_1, calculated from Eq. (5) with a naturally augmented set of parameters P and for a constant u-value of 1.2 g. Some observations and comments from these numerical results:

(i) There exists a 'peak' in all sensitivities at around 20 days. In the modeling context (see [10]) this means that measurements of all variables around $t = 20$ should be preferred in order to assess reliable values for β_0, α_0 and α_1.
(ii) Similar sensitivity values are obtained for α_0 and α_1, from which it can be inferred that both parameters are equally necessary in the expansion of the original parameter β.

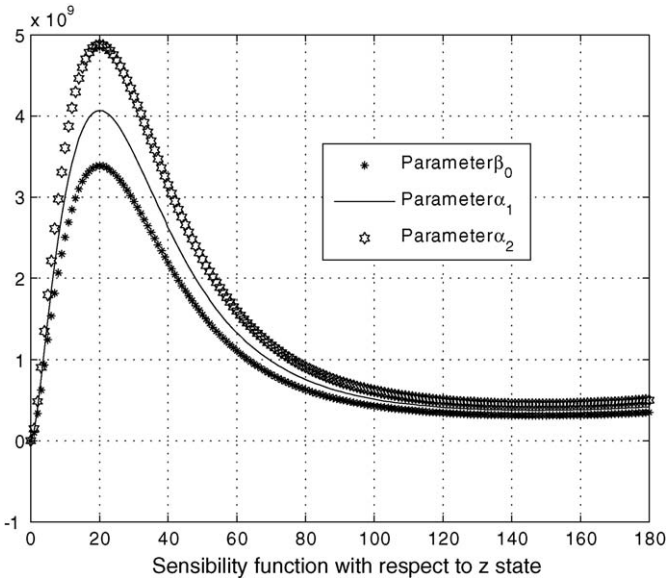

Fig. 3. Sensibility of state z with respect to parameters $\beta_0, \alpha_1, \alpha_2$.

(iii) The high value of sensitivity with respect to β_0 means that β is an appropriate parameter to be 'controlled'. In other words, the administration of drugs (in amounts u) that directly affect the value of β, will in turn influence the dynamics in a degree related to S_β.

4. The total cost associated to a therapeutic strategy

A typical objective functional, representing the "total" cost to minimize amongst all acceptable therapies may be designed as follows

$$J(u) = \int_{t_o}^{T} \left\{ \left[a_1 z(t) + a_2 (x(t) - \bar{x})^2 \right] e^{\alpha t} + a_3 u^2(t) \right\} dt + a_4 z^2(T), \tag{12}$$

consisting of a "trajectory cost" $Q(t_0, T, x_0, y_0, z_0, u)$ (usually expressed as the integral of the Lagrangian function $L(t, x, y, z, u)$ modeling the cost differentials occurring during treatment), and a "final penalty" $K(x, y, z)$ associated to the departure from desired (target) states at the end of the therapy. The Lagrangian in this case includes nontrivial terms $[a_1 z(t) + a_2 (x(t) - \bar{x})^2]$ penalizing the departure of the variables from their 'healthy' equilibrium values. This is a difference with other dynamic programming attempts to this problem (see for instance [19] and some of the references therein), that only take into account the drug cost $u^2(t)$ during treatment. Adding state-trajectory costs will prevent from total drug interruption at any intermediate time within optimal therapies. The choices made here, namely

$$L(t, x, y, z, u) = [a_1 z + a_2 (x - \bar{x})^2] e^{\alpha t} + a_3 u^2 \tag{13}$$

$$K(x, y, z) = a_4 z^2 \tag{14}$$

reflect the emphasis on abating virions (or 'viral copies', namely in terms $a_1 z(t)$ and $a_4 z^2(T)$) and the deviation of x-cells from the equilibrium $\bar{x} = 1000$ cells/mm^3. The $a_3 u^2(t)$ term represents the effective cost of drug consumption. The squares in the previous terms are only used for convenience in determining optimal controls analytically, as will be discussed later.

The (opposite) discount factor term $e^{\alpha t}$, i.e. with $\alpha \geq 0$, punishes the "permanence" of the disease. The penalty on departures of the state variables from their target values ($\bar{x} = 1000, \bar{z} = 0$) is greater as time increases. It is assumed that the immune response of the patient deteriorates with time, and then the same infectious situation causes more harm as time goes on, so its "cost" should increase with time.

The values of the coefficients have been tuned so as to obtain partial costs (each summand in (12)) proportional to 25, 10, 45, and 20% of the total cost associated to a typical trajectory, which is normalized to unity, i.e.

$a_1 = 0.25 c_1$
$a_2 = 0.10 c_2$
$a_3 = 0.45 c_3$
$a_4 = 0.20 c_4$, with
$c_1 = 5.970 \times 10^{-6}$
$c_2 = 1.2962 \times 10^{-7}$
$c_3 = 0.003858$
$c_4 = 0.000164$
$\alpha = 0$ (no inverse discounting); 0.005 (inverse discounting case)
$t_0 = 0$
$T = 180$ days

The initial and final times (t_0 and T respectively) deserve special comment. The $t_0 = 0$ adoption is appropriate for time-constant models and Lagrangian cost functions, although the insertion of a discount factor may change the perspective. Regarding the optimization horizon T, its adopted value is based on empirical evidence: viral load reduction to below limits of assay detection (50 copies/ml) in a treatment-naïve patient usually occurs within the first 12–24 weeks of therapy [29].

The usual theoretical approach to the solution to the optimal control problem just posed, when there are no special restrictions, evolves through the following steps:

(i) Define the 'Hamiltonian' of the problem (see [5,26]), i.e.

$$H(t, X, \ell, u) \triangleq L(t, X, u) + \ell' f(X, u), \tag{15}$$

where ℓ is a new vector-valued variable, of the same dimensions than X, called the 'costate' or 'adjoint variable', related to the Lagrange multipliers appearing in static constrained optimization, and whose precise role in the present context is described below.

(ii) Asses whether the problem is 'regular', i.e. if there exists a continuous function:

$$u^0(t, X, \ell) \triangleq \arg\min_u H(t, X, \ell, u), \tag{16}$$

called the 'H-minimal control'. If the answer is negative, the problem needs to be reworked or treated along other lines.

(iii) The u-minimal Hamiltonian H^0 can then be defined:

$$H^0(t, X, \ell) \triangleq H(t, X, \ell, u^0(t, X, \ell)), \tag{17}$$

(iv) and the Hamiltonian equations of the problem result:

$$\dot{X} = \left(\frac{\partial H^0}{\partial \ell} \right)'; \qquad X(0) = X_0 \tag{18}$$

$$\dot{\ell} = -\left(\frac{\partial H^0}{\partial X} \right)'; \qquad \ell(T) = \frac{\partial K}{\partial X}(X(T)) \tag{19}$$

which poses a two-point boundary value problem, usually very difficult to solve, even numerically.

4.1. The need for discretization

The posing of a continuous-time optimal control problem is not practical for HIV treatment. Notwithstanding the patient's health

undergoes a continuos deterioration, possibly following a model like Eq. (2), the assessment of the situation can only be made through periodic blood analyses and doctor inspections, and the present administration of drugs can only be made through discrete amounts and changed every some time according to prescriptions. This means that, even when the system evolves in continuous time and the cost objective can theoretically be posed in the same context, the discrete nature of:

(i) measurements' availability for the state variables,
(ii) the discrete nature of admissible control values,
(iii) the existence of restrictions in the admissible control values, which hinders the continuity of the H-minimal control function u^0 of Eq. (16), and
(iv) the usual delay between physician interventions and control decisions, forces to consider a mixed continuous/discrete approach. In what follows, the values of the states and control variables will be then discretized according to the following scheme

$$\begin{array}{l}\mathcal{X} \triangleq \{x_L, x_L + \Delta x, x_L + 2\Delta x, \ldots, x_U\} \\ \mathcal{Y} \triangleq \{y_L, y_L + \Delta y, y_L + 2\Delta y, \ldots, y_U\} \\ \mathcal{Z} \triangleq \{z_L, z_L + \Delta z, z_L + 2\Delta z, \ldots, z_U\} \\ \mathcal{U} \triangleq \{u_L, u_L + \Delta u, u_L + 2\Delta u, \ldots, u_U\}\end{array} \tag{20}$$

where the values of the lower (L) and upper (U) bounds and the grid size (Δ) for each variable should depend on real constrains on the appreciation of measurement devices, the possibilities of dose subdivision, calculation capabilities, and expectations. However it is important to note that the adopted values for y_L and z_L should be strictly greater than zero to avoid an unrealistic stagnation of the variables y, z (see Eq. (2)). This will be discussed further, together with the numerical trials and results. The cost takes now a slightly different form

$$\begin{array}{l}\mathcal{J}(u) \triangleq \displaystyle\sum_{k=0}^{T/h} \int_{t_k}^{t_{k+1}} \left[(a_1 z(t) + a_2 (x(t) - \bar{x})^2) e^{\alpha t} + a_3 u_k^2\right] dt + a_4 z^2(T) \\ t_k \triangleq t_0 + hk\end{array} \tag{21}$$

where $x(t), y(t), z(t)$ must be understood, in each interval $[t_k, t_{k+1})$, as the rounded result of the state-transition function $\phi(t, t_0, x, y, z, u(.))$ associated with the continuous-time model (1), namely

$$(x(t), y(t), z(t))' = \phi(t, t_k, x_k, y_k, z_k, \tilde{u}_k) \tag{22}$$

$$(x_{k+1}, y_{k+1}, z_{k+1})' = \text{round}(\phi(t_{k+1}, t_k, x_k, y_k, z_k, \tilde{u}_k)) \tag{23}$$

$$\tilde{u}_k(t) \equiv u_k \tag{24}$$

and where 'round' acts over the values $(x(t_{k+1}), y(t_{k+1}), z(t_{k+1}))' = \phi(t_{k+1}, t_k, x_k, y_k, z_k, \tilde{u}_k)$ in a 'safe' way, precisely

$$\begin{array}{l}x_{k+1} \triangleq \text{closest smaller value next to } x(t_{k+1}) \text{ in } \mathcal{X} \\ y_{k+1} \triangleq \text{closest bigger value next to } y(t_{k+1}) \text{ in } \mathcal{Y} \\ z_{k+1} \triangleq \text{closest bigger value next to } z(t_{k+1}) \text{ in } \mathcal{Z}\end{array}$$

The adopted value of $h = 15$ days takes into consideration the observed 'peak time' (of approximately 20 days) occurring in the state variables and sensitivities (see Fig. 3). The possibility of a hiding acute infection period is discarded this way, since at least the results of a blood analysis reflecting the situation will come at some intermediate point.

The Hamiltonian formulation of the problem above reads (see [5] for details)

$$\mathcal{J}(u) = \sum_{k=0}^{N} \mathcal{L}_k(x_k, y_k, z_k, u_k) + a_4 z^2(T), \tag{25}$$

where $N \triangleq T/h$, and the discrete Lagrangian $\mathcal{L}_k(x_k, y_k, z_k, u_k) = \mathcal{L}_k(X_k, u_k)$ is now

$$\mathcal{L}_k(X_k, u_k) \triangleq \int_{t_k}^{t_{k+1}} [(a_1 z(t) + a_2 (x(t) - \bar{x})^2) e^{\alpha t} + a_3 u_k^2] dt, \tag{26}$$

for $k = 0, 1, \ldots, N-1$. The discrete Hamiltonian of the problem can be defined in a convenient form

$$\mathcal{H}_k \triangleq \mathcal{L}_k + \ell'_{k+1} f_k, \tag{27}$$

where the costate (column) vector $\ell_k \in \mathbb{R}^3$ plays the role of the gradient of the value (or Bellman) function as in the continuous time setup (see for instance [5,26]). After taking variations in the cost functional the usual Hamiltonian dynamics get condensed in the following scheme

$$\ell_k = \left(\frac{\partial f_k}{\partial X}\right)' \ell_{k+1} + \left(\frac{\partial \mathcal{L}_k}{\partial X}\right)', \tag{28}$$

$$\frac{\partial \mathcal{H}_k}{\partial u} = \frac{\partial \mathcal{L}_k}{\partial u} + \ell'_{k+1} \frac{\partial f_k}{\partial u} = 0, \tag{29}$$

subject to the final condition

$$\ell_N = \left(\frac{\partial K}{\partial X}\right)'\Big|_N = 2a_4(0, 0, z(T))', \tag{30}$$

where the following interpretation of the discrete state-transition function is adopted

$$X_{k+1} = f_k(X_k, u_k) \triangleq \phi(t_{k+1}, t_k, X_k, u_k). \tag{31}$$

Notice that it is possible to work backwards from the final condition (30). The local optimal control value u_{N-1} can in principle be extracted as the u-solution to Eq. (29). Then this u_{N-1} is replaced in (28) with $k = N-1$, and then the ℓ_{N-1} becomes available. And then the process follows for $k = N-2, \ldots, 0$. But no restrictions in the values of u can be easily imposed in this formulation, nor any prescription to transform the eventual results to a state-feedback form. These inconveniences have influenced the decision to use dynamic programming as the actual numerical method of solution, to reserve the Hamiltonian formalism for checking the validity of results, and to record the values of validated costates as useful additional information (each ℓ_k becomes the vector of 'marginal costs' from stage k on).

5. Numerical results

Due to the nonlinearity of the disease represented by the population models such as the one used in this work, the identification and characterization of the patient state may be non-trivial. That is, a similar set of values in viral load and immune cells population can correspond to different stages of the infection and evolution perspectives, as can be observed in the flow of the system in Fig. 1. This way, the physician's intervention plays a major role in the therapy's efficiency since by means of the clinical history of the patient, medical interrogation, and complementary tests, the real stage of the infection can be determined and eventual corrections over existing protocols can be made. In this section two typical situations will be treated numerically, discussed, and illustrated: (i) initial conditions describing a patient having viral load and immunological cells populations near the beginning of the infection (acute infection), and (ii) a patient near the endemic equilibrium point (asymptomatic phase). A nontrivial ($\alpha \neq 0$) inverse discount factor in the objective cost will be only applied in Section 5.1.3.

5.1. Treatment of a recently discovered infection

In this subsection the optimal treatment of a patient with a recently discovered infection, for instance with initial conditions near (see Fig. 1)

$$x_0 = 800\,\text{cells/mm3}, \qquad y_0 = 30\,\text{cells/mm3}, \qquad z_0 = 3500\,\text{copies/ml}, \tag{32}$$

is illustrated. Dynamic programming was implemented for a range of states around these initial conditions, covering expected behaviors of the patient under different medication strategies. The adopted discretization or spacing (Δ) of the variables and their corresponding lower (L) and upper (U) bounds are listed bellow:

$$\Delta x = 50, \qquad \Delta y = 10, \qquad \Delta z = 50, \qquad \Delta u = 0.2, \tag{33}$$

$$x_L = 600, \qquad y_L = 1, \qquad z_L = 10, \qquad u_L = 0, \tag{34}$$

$$x_U = 1000, \qquad y_U = 91, \qquad z_U = 5010, \qquad u_U = 1.6. \tag{35}$$

It must be remarked that the lower thresholds y_L and z_L are given positive values. This is to avoid that, after an eventual rounding of their values, the discretized trajectories reach any point with $y = z = 0$, since in that case the optimal strategy would continue with $u \equiv 0$ till the end, which is certainly erroneous (and dangerous). Indeed, the real system never reaches $y = z = 0$ from an initial condition different from the unstable equilibrium, which means that any remaining infection $(y > 0, z > 0)$ will grow if $u = 0$, and this growth should eventually be controlled with some $u > 0$.

Figs. 4 and 5 show projections (individual components) of the optimal state-trajectories $\{(x^*(t), y^*(t), z^*(t)), 0 \le t \le T\}$ over the (t,x) and (t,z)-planes, respectively (to save room the y-component is not depicted since it is qualitatively very similar to the z-component). In a 3-dimensional phase-space, with t as a parameter, these trajectories resemble a contracting flow, and once the flow reaches a point, from there on the trajectory, being optimal, is unique. This aspect is partially reflected in the projections: eventual bifurcations at some point $(t, x^*(t))$ indicate that there exist $M > 1$ different values of $(y^*(t), z^*(t))$ giving rise to $m \le M$ trajectories after t. But in general it can be observed that the variables are driven towards the desired values $(\bar{x}, \bar{y}, \bar{z})$.

Fig. 6 show optimal control strategies corresponding to different initial state conditions. It is observed that the final stage requires an optimal control of $u = 1.6$ g which reflects the influence of the strong final penalty $K(x,y,z) = a_4 z^2$ (the virions z have values in the order of 500 near the end of the treatment (see Figs. 7 and 11), which forces high values of the control to abate $z(T)$ near zero). Also it can be observed that most optimal strategies

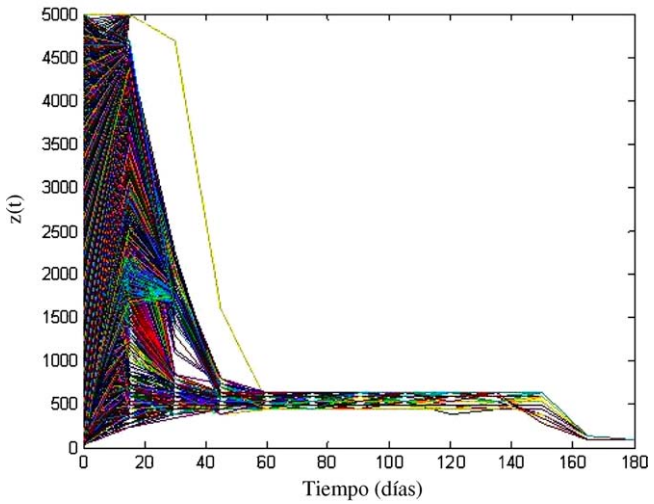

Fig. 5. $z^*(t)$ optimal trajectories, each for a fixed $z^*(0)$ and different $x^*(0), y^*(0)$.

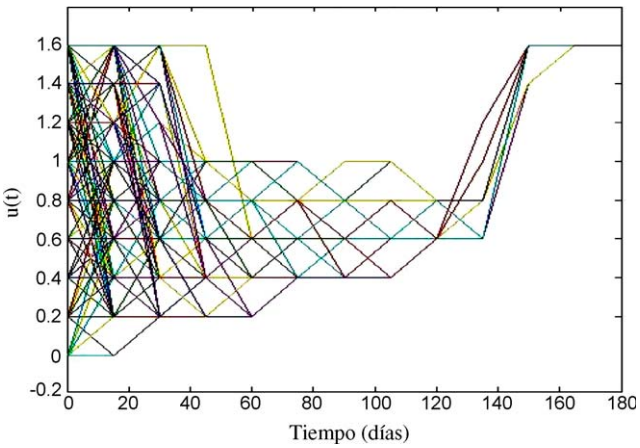

Fig. 6. Typical optimal $u^*(t)$trajectories.

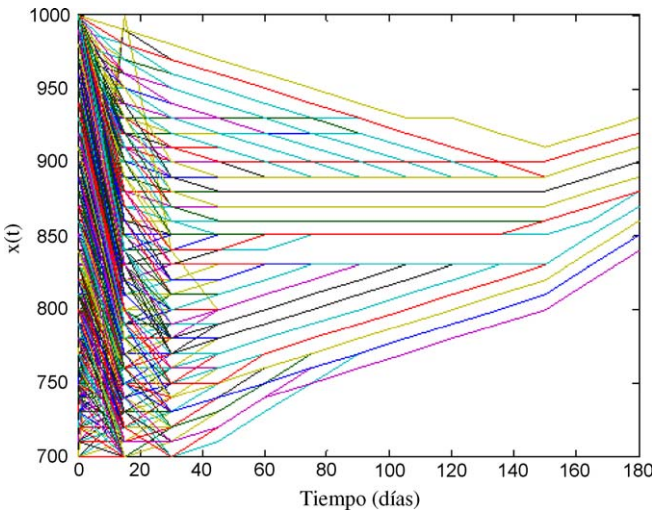

Fig. 4. $x^*(t)$ optimal trajectories, each for a fixed $x^*(0)$ and different $y^*(0), z^*(0)$.

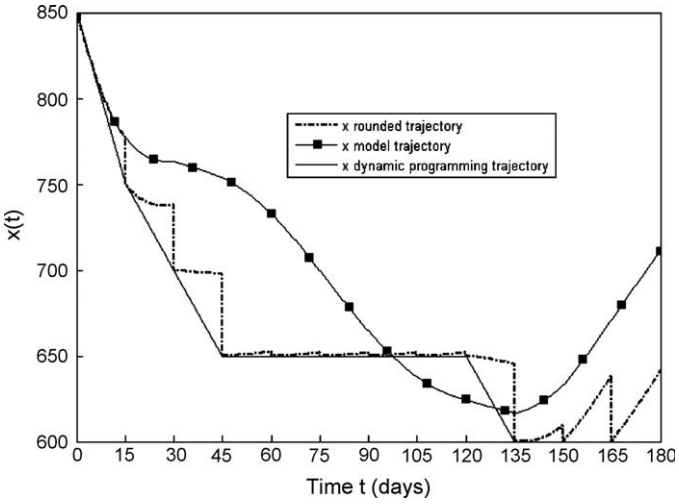

Fig. 7. Optimal trajectory $x^*(t)$ constructed under "worst-situation" rounding, and continuous-time evolution corresponding to the same controls.

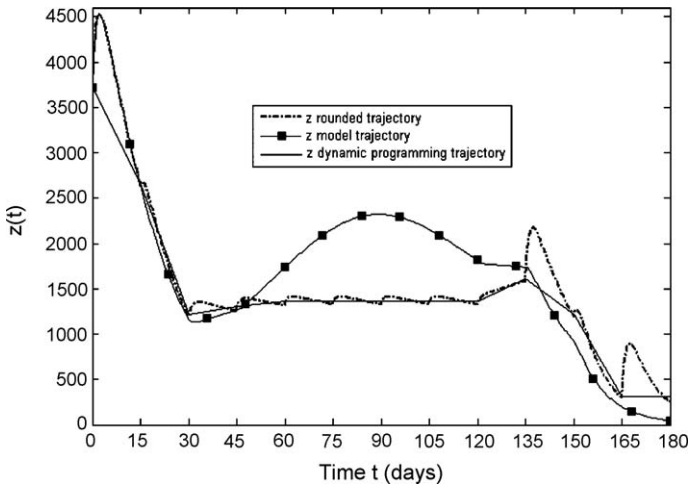

Fig. 8. Optimal trajectory $z^*(t)$ constructed under "worst-situation" rounding, and continuous-time evolution corresponding to the same controls.

require mild values of the control $(0.4 \leq u \leq 1)$ in the central stages of the treatment, all below the usually recommended $u \equiv 1.2$ g [6], reinforcing the possibility of variable doses with smaller impact.

5.1.1. Open-loop trajectories for a particular initial state of the patient

Typical initial conditions associated with a recent infection situation are

$$x_0 = 850\,\text{cells/mm3}, \qquad y_0 = 41\,\text{cells/mm3}, \qquad z_0 = 3710\,\text{copies/ml}. \tag{36}$$

The optimal control strategy was recovered from the dynamic programming solution to the problem set by Eqs. (21), (22), (33)–(35). It is shown in Fig. 12. The corresponding optimal state-trajectory was also stored during the solution process, and it is depicted as broken lines in Figs. 7 and 8. At each sampling time there also starts a curve, solution to the continuous-time flow of Eq. (2), showing the expected evolution of the system during the next sampling period. That curve is truncated at the end of this period, according with the adopted discretization of the states, and in the figures can be observed that this rounding is made along the 'worst situation' rule, i.e. going down when looking for the next admissible x-value, and up for y and z-values. Finally, in the figures are also depicted the evolution of the patient subject to the same control strategy but without any rounding of the states, denoted as 'model trajectories'.

The open-loop optimal solutions succeed in reducing the number of viral copies below the 50 copies/ml in the desired time of 180 days. The virions decrease strongly at the beginning and the end of treatment, remaining around 1200 copies/ml during most of the period, which is not too bad as long as no deleterious 'peaks' occur. At the extreme stages the optimal drug doses is high (1.4 g and 1.6 g respectively), but during most of the treatment the values stay around 0.6 g, half of the tested protocol value (see Fig. 12). Also, during the application of the control strategy the number of x-cells is always higher that 600 cell/mm^3, considered clinically acceptable [15].

Numerical checking for optimality of these solutions were performed, by comparing costs with neighboring trajectories, which resulted greater as expected. But also a partially theoretical check was attempted. Fig. 9 shows a comparison between the norm of the costate calculated from discrete Hamilton equations and from dynamic programming. The coincidence is good enough for our purposes, since the theoretical approach does not take into account the discretization nor the upper bound in the u-values (see Eq. (29)). Notice that with unbounded controls, and to cope with the strong final penalty on z, the optimal drug doses for the last stage would have most probably been greater than 1.6 g. And this difference is reflected in the discrepancy between ℓ-values in the last sampling period, which in turn propagates through earlier stages until reaching a complete agreement.

5.1.2. The need for closed-loop strategies

The open-loop trajectories (Figs. 7 and 8) show that the difference between the evolutions of the real patient and its expected discretized behavior is not negligible, specially between the 45 and 105 days. The discretized trajectory being optimal can not be improved, but it is possible to correct the situation after a significant departure from the expected state is noticed. The actual state X_k of the patient, assessed through reliable measurements, might be considered as a 'perturbation' from the expected one X_k^*. And then, the 'closed-loop' methodology in control theory suggests to consider a new optimization problem, starting at X_k and for the remaining stages from k to N.

After the appropriate admissible grid approximation $\widetilde{X_k}$ to X_k is selected, the new optimization problem does not need additional calculation if dynamic programming has been used to determine the open-loop optimal solution, and if the valuable information has

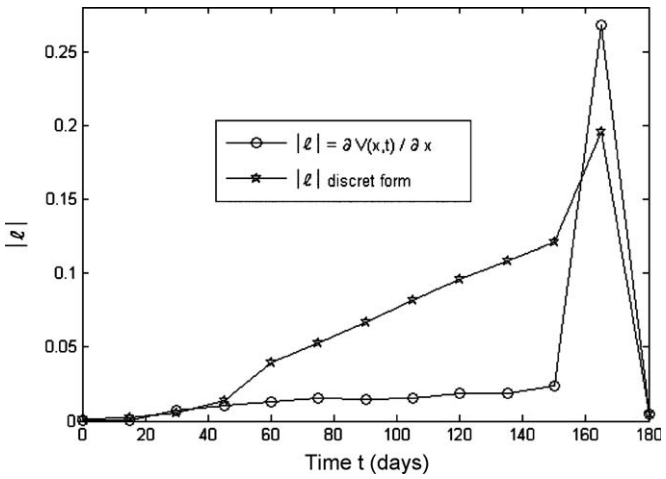

Fig. 9. Comparison between the norm of the costate calculated from discrete Hamilton equations and from dynamic programming.

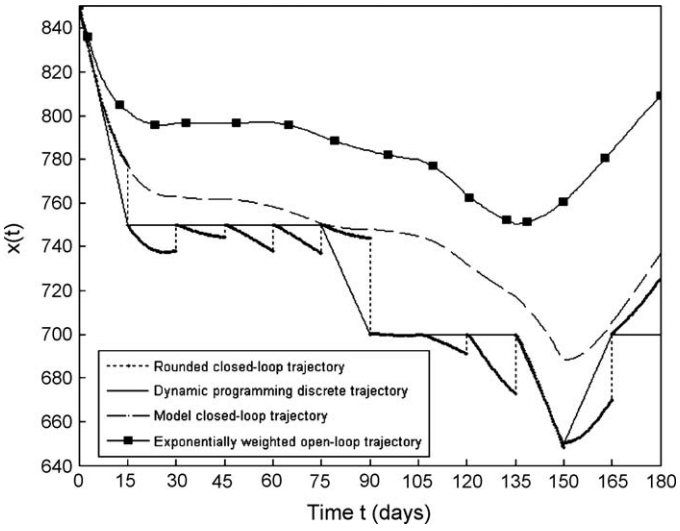

Fig. 10. Comparison of closed-loop and inverse-discount optimal x-trajectoties.

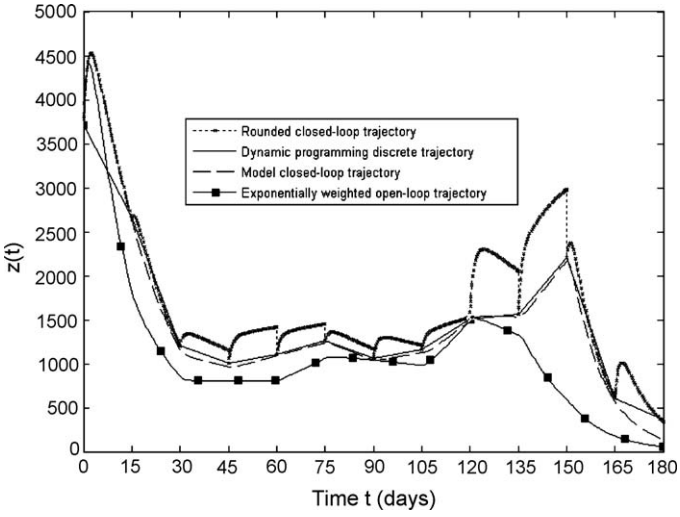

Fig. 11. Comparison of closed-loop and inverse-discount optimal z-trajectoties.

been stored when generated. At each point in the admissible state grid and at each stage k dynamic programming usually stores:

(i) the optimal control $u_k^*(\widetilde{X_k})$ to be applied during the stage k,
(ii) the value of the optimal cost for trajectories starting at $\widetilde{X_k}$ and evolving during stages k to N, and
(iii) the state X_{k+1}^* reached after applying $u_k^*(\widetilde{X_k})$ during the stage k.

Therefore, by tracking the steps from $\widetilde{X_k}, u_k^*(\widetilde{X_k})$ through the end, the new optimal control strategy becomes available. Actually, in practical terms just the information $u_k^*(\widetilde{X_k})$ should be used with certainty, i.e. the first control value of the new strategy. Because, if incoming data X_{k+1} at the next stage if found to differ from X_{k+1}^* in a grid appreciable amount, then a new closed-loop correction will be needed.

Figs. 10 and 11 show the application of this methodology to the same problem stated, but considering the measured (actually the continuous-time $X_{k+1} = \phi(t_{k+1}, t_k, X_k, u_k^*)$) state of the patient as 'perturbations' from the expected (discretized optimal) state $X_{k+1}^* = [\phi(t_{k+1}, \widetilde{t_k}, \widetilde{X}_k, u_k^*)]$. Notice that, in general, $X_{k+1}^* \neq \widetilde{X_{k+1}}$.

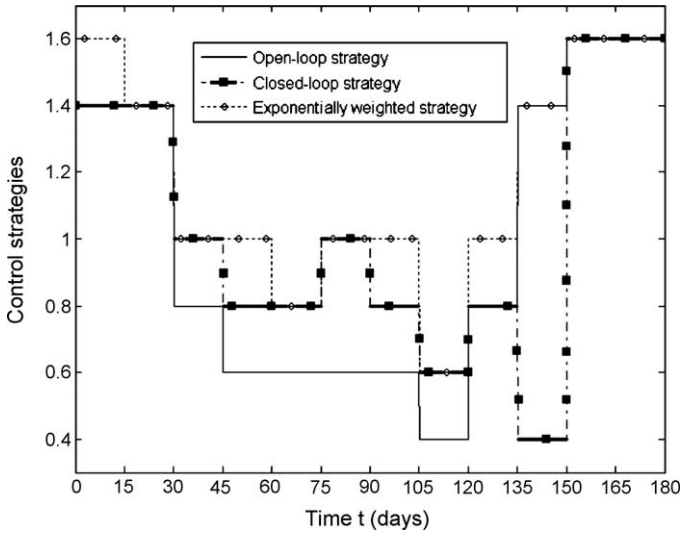

Fig. 12. Comparison of closed-loop and inverse-discount optimal control trajectoties.

Table 1
Cost values for open and closed-loop control, and for inverse-discount evaluation.

	Trajectory cost	Final penalty	Total cost
Open-loop trajectory	1.0344	2.9520	3.9864
Closed-loop trajectory	1.0128	4.0180	5.0308
Discounted trajectory	1.0453	2.9520	3.9973

The 'exponentially weighted open-loop trajectories' are obtained from applying the Dynamic Programming optimal recipe to the model. The expected behavior, however, would be the 'Dynamic Programming discrete trajectories', which clearly depart from the model in open-loop. The first significant detected departure motivates a correction to the open-loop prescription, which results in the 'model closed-loop trajectory' when is applied to the patient.

The closed-loop control reduces the differences between the states of the patient and the expected one, and the abatement of virions z is better than the obtained in open-loop all over the period of treatment, while the healthy cells x recuperate also better under feedback control conditions. Changes in the closed-loop control doses during treatment are not trivial, as can be observed in Fig. 12, which strengthen the merit of close monitoring by the medical doctor during treatment.

In Table 1 the costs of open and closed-loop trajectories are given. As expected, closed-loop costs more since it is not optimal with respect to the original set-up. The decision to consider the real behavior of the patient as 'perturbations' from the expected discretized solution, and the corrections made at the t_k's in the values of the 'new initial' states X_k (for the remaining stages $k, k+1, \ldots, N$), change the optimal control strategy for the remaining horizon, and then it should have a greater cost. When corrections are frequent, as in the example shown, these differences in cost may be significant.

5.1.3. The effect of an inverse discount factor

There exists empirical support (and the model reflects it) of 'silent' or relatively stable periods occurring after infection peaks. But these should not be assumed as healing phases. It has been noted [9] that chronic inflammation is taking place after peaks, affecting blood vessels in the whole organism, and increasing cardiovascular risk, together with related sequels in brain

irrigation and the functioning of kidneys, liver and lungs. This chronic inflammation seems in turn to lower the effectiveness or to generate intolerance to subsequent therapies. Also it is known that treatments are less tolerated when the patient does not begin taking drugs at an early stage. Therefore, even when the model is kept having nearly time-constant coefficients, the need to promote early treatment (or to penalize the postponing of control actions) can be made explicit through the addition of the 'inverse discount' term $e^{\alpha t}$ in the Lagrangian function.

In Figs. 10 and 11 the optimal solution to an exponentially weighted trajectory cost is added to the close-loop behavior corresponding to $\alpha = 0$. In the inversely discounted case a value of $\alpha = 0.005$ was used, which more than duplicates the state-departure differential costs near the horizon: $e^{\alpha T} = \exp(0.005 \cdot 180) \approx 2.46$. Since the control-weight r was kept the same for $\alpha \geq 0$, then the exponential factor worked as an extra-penalization of state departures, which in optimal control always means a sub-penalization of the control effort. The result was that, with $\alpha > 0$, trajectories improved when comparing to the case $\alpha = 0$, due to the application of bigger control values. Therefore the total optimal costs for the discounted and neutral cases (both open-loop) were very similar (see Table 1) since their components compensated each other.

5.2. Medication of an endemic situation

The 'endemic' equilibrium of the uncontrolled system represents a patient that has been infected long ago, and that has reached an apparent steady state (while subsequent blood analyses may suggest a regular behavior, there will obviously exist a global deterioration due to the high discrepancy from 'healthy' equilibrium values). Therefore a strategy to pull out the patient from the endemic toward a healthy condition should be adopted. From the control point of view the problem looks as a 'change of set-point', i.e. to conduct the system from an equilibrium to another. In this case it is required to go from a stable to an unstable steady state, which is likely to imply a significant control effort. An illustration of the dynamic programming approach to this problem is given in Figs. 13–15 . An initial condition near the endemic equilibrium was chosen:

$$x_0 = 600\,\text{cells/mm3}, \qquad y_0 = 11\text{cells/mm3}, \qquad z_0 = 1760\,\text{copies/ml}, \tag{37}$$

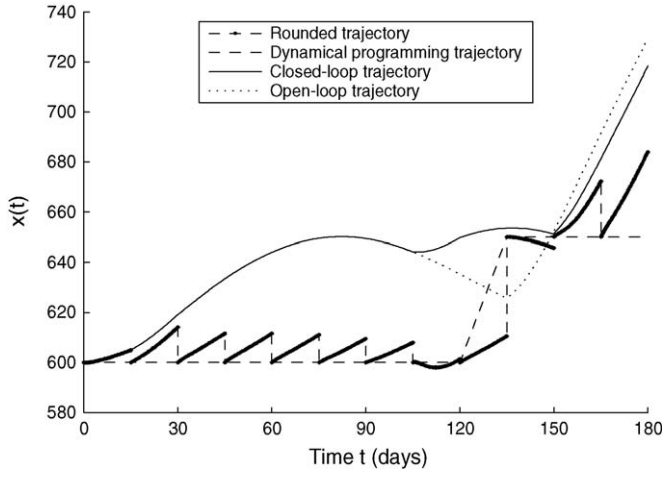

Fig. 13. Comparison of open and closed-loop optimal x-trajectoties for the endemic case.

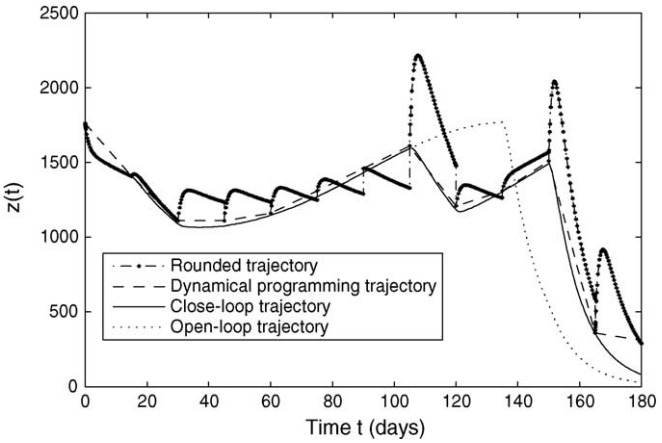

Fig. 14. Comparison of open and closed-loop optimal z-trajectoties for the endemic case.

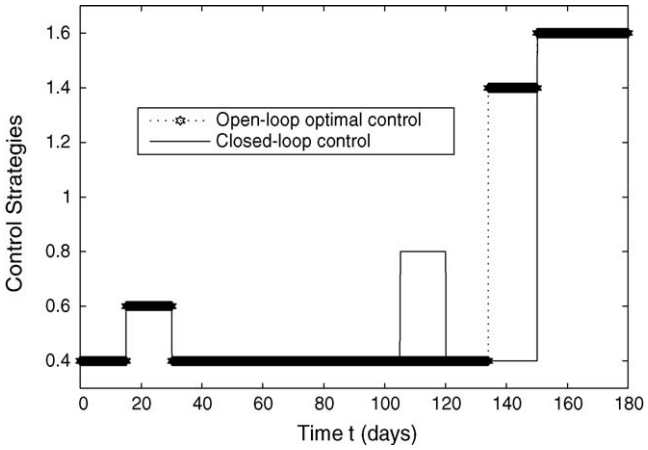

Fig. 15. Comparison of open and closed-loop optimal control trajectoties for the endemic case.

and the following discretization parameters

$$\Delta x = 50, \qquad \Delta y = 10, \qquad \Delta z = 50, \qquad \Delta u = 0.2, \tag{38}$$

$$x_L = 500, \qquad y_L = 1, \qquad z_L = 10, \qquad u_L = 0, \tag{39}$$

$$x_U = 1000, \qquad y_U = 91, \qquad z_U = 2510, \qquad u_U = 1.6. \tag{40}$$

The closed-loop trajectory shows a better tracking of the real patient behavior to the cost of abandoning the optimality of the dynamic programing solution (which actually produces the curves with the 'open loop' legend). The qualitative evolution of the x- and z-trajectories reflect their resistance to abandon the endemic region, which insinuates only after approximately 100 days. But, on the other hand, only mild doses of drugs are necessary during most of the treatment, and the high doses are concentrated at both extremes of the optimization period.

6. Conclusions

The elaboration on a multi-objective cost optimization expounded in previous sections sustains several propositions, such as

- It may be appropriate to adopt strategies of variable drug administration during the period of treatment of an HIV-infected patient. Recent evidences indicate that the generation of

resistance to drugs by the virions is a complex phenomenon, not directly related to keep or quit high doses constant (adherence). Therefore the schemes of treatment proposed in this work, that imply variable doses (see especially Figs. 6, 12, and 15), deserve to be explored in greater depth. The current protocols of antiretroviral drugs administration imply the quantification of immunological and virological variables each 3–4 months. The tools devised here may assist the medical professional in monitoring the patient by means of more frequent blood analysis, which might be taken as an extra burden, but benefits appear through

(i) the reduction of indirect effects over the patient's health coming from a smaller impact of lower doses, and
(ii) the savings in public funds implied by a lower average of drug consumption.

- No interruption of drug treatment seems possible under the adopted model and evaluation scheme. The infected cells and virions can not be abated to zero with reasonable doses of the control variable. Then, any remaining infection will grow and produce a dangerous peak in absence of a new drug doses, however small, prescribed by the dynamic programming solution.
- Optimal solutions can not be thought as enduring recipes. The supervision of their application by the medical doctor is essential during treatment, since most probably there will exist 'perturbations' in state values. The state of a patient may appreciably differ from the expected optimal 'open-loop' trajectory. The 'closed-loop' scheme gives 'better' results: having dynamic programming solutions at hand allows the physician to correct the state as time proceeds and to recover the optimal control strategy from that time on, without need for further calculations.
- The relative weights of the partial costs should also be tuned by the professional. The optimal strategies can, in principle, be simulated and assessed along variable sets of priorities before selecting the definitive cost functional. There exists practically no methodology for the design and tuning of the objective cost, not even for the most frequently used Lagrangian functions (see [5] for the linear-quadratic LQR context).
- Other parameters in the model of the dynamics and the constrains should be ascertained, and eventually updated by the potential user of the proposed method. Sensitivity analysis can help to determine the appropriate values of parameters for different patients and drugs used, and to select the best time instants to take the measurements that substantiate these values.
- There exist several approaches to the optimal control problem at hand. Some alternatives to dynamic programming are less versatile, like the discrete Hamiltonian approach referred to in several points of Section 4; or plainly inadequate to describe the relation patient–doctor, as the continuous-time application of Hamilton equations or the Pontryagin principle. But all of them can provide useful information and eventually validate dynamic programming results. In all, an expert-oriented software package to assist medical doctors in assessing, simulating, and prescribing therapeutic strategies to HIV-infected patients appears as a feasible valuable tool. Variable drug dosage seems to be optimal under certain circumstances, when a multiple-objectives cost evaluation of the HIV patient's evolution is considered. Circumstances include:

(i) that a discretization of the main variables has been adopted, due to the impossibility of continuous variable measurement and medical inspection, the availability of drugs only in fixed doses, and so on.
(ii) there exists an upper restriction in the amount of drug than can be administered during a fixed period of time.
(iii) the initial values of the three main state variables can be precisely estimated, and so can every 'sampling time' (typically a period in the order of 15 days).

The combination of a discretized objective functional and a continuous-time dynamics at the interior of each stage has shown to be practical in implementing an accurate dynamic programming scheme. The results provide optimal strategies starting at each node in the discretized state space, and starting from any intermediate stage through the end. This allows to put into action the closed-loop point of view when applying to a real patient, hence returning the control of the situation to the doctor. A little training of medical personnel on the basics of dynamical optimization (the optimal control way of reasoning) will endorse the generation and use of software tools based on these principles.

References

[1] J. Alvarez-Ramirez, M. Meraz, J.X. Velasco-Hernandez, Feedback control of the chemotherapy of HIV, Int. J. Bifurcation Chaos 10 (2000) 2207–2219.
[2] D. Bangsberg, A. Moss, S. Deek, Paradoxes of adherence and drug resistance to HIV antiretroviral therapy, J. Antimicrob. Chemother. 53 (2004) 696–699.
[3] M. Barão, J.M. Lemos, Nonlinear control of HIV-1 infection with a singular perturbation model, Biomed. Signal Process. Control 2 (2007) 248–257.
[4] M.E. Brandt, G. Chen, Feedback control of a biodynamical model of HIV-1, IEEE Trans. Biomed. Eng. 48 (2001) 754–759.
[5] A.E. Bryson, Y.-C. Ho, Applied Optimal Control, Wiley, Washington, DC, 1975.
[6] P. Cahn, I. Cassetti, R. Wood, P. Phanuphak, L. Shiveley, R.C. Bethell, J. Sawyer, Efficacy and tolerability of 10-day monotherapy with apricitabine in antiretroviral-naive, HIV-1 infected patients, AIDS 20 (9) (2006) 1261–1268.
[7] E.F. Camacho, C. Bordons, Model Predictive Control, 2nd edition, Springer-Verlag, 2003
[8] F.M. Campello de Souza, Modeling the dynamics of HIV-1 and CD4 and CD8 lymphocytes, IEEE Eng. Med. Biol. Mag. 18 (1999) 21–24.
[9] B. Caramelli, Cardiovascular risk and metabolic effects in HIV patients, in: 29th World Congress of Internal Medicine, Buenos Aires, Argentina, September 16–20, 2008.
[10] A. Cheruy, Méthodologie de la Modélisation, Les Cahiers d'EDORA (Reports INRIA-Sophia Antipolis, France) 866 (1988) 23–48.
[11] R.V. Culshaw, S. Ruan, R.J. Spiteri, Optimal HIV treatment by maximizing immune response, J. Math. Biol. 48 (2004) 545–562.
[12] R.T. Davey Jr., HIV-1 and T-cell dynamics after interruption of highly active antiretroviral therapy (HAART) in patients with a history of sustained viral suppression, Proc. Natl. Acad. Sci. U.S.A. 98 (26) (1999) 15109–15114.
[13] S.S. Ge, Z. Tian, T.H. Lee, Nonlinear control of a dynamic model of HIV-1, IEEE Trans. Biomed. Eng. 52 (2005) 353–361.
[14] HIV Treatment Options, October 2008, http://www.hivinfosource.org.
[15] C. Hoffmann, B. Kamps, HIV Medicine 2007, Flying Publisher, October 2008, http://www.HIVmedicine.com.
[16] A.M. Jeffrey, X. Xia, I.K. Craig, When to initiate HIV therapy: a control theoretic approach, IEEE Trans. Biomed. Eng. 50 (11) (2003) 1213–1220.
[17] D. Kirschner, Using mathematics to understand HIV immune dynamics, AMS Notices (1996) 191–202.
[18] D. Kirschner, S. Lenhart, S. Serbin, Optimal control of the chemotherapy of HIV, J. Math. Biol. 35 (1997) 775–792.
[19] J.H. Ko, W.H. Kim, C.C. Chung, Optimized structured treatment interruption for HIV therapy and its performance analysis on controllability, IEEE Trans. Biomed. Eng. 53 (3) (2006) 380–386.
[20] M.A. Nowak, R.M. May, Mathematical biology of HIV infections—antigenic variation and diversity threshold, Math. Biosci. 106 (1991) 1–21.
[21] M.A. Nowak, R.M. May, AIDS pathogenesis—mathematical models of HIV and SIV infections, AIDS 7 (1993) 3–18.
[22] M.A. Nowak, R.M. May, Virus Dynamics: Mathematical Principles of Immunology and Virology, Oxford University Press, New York, 2000.
[23] D.A. Ouattara, C.H. Moog, Modeling of the HIV/AIDS infection: an aid for an early diagnosis of patients, in: I. Queinnec, et al. (Eds.), Biology and Control Theory: Current Challenges, LNCIS 357, Springer, 2007.
[24] A.S. Perelson, D.E. Kirschner, R. De Boer, Dynamics of HIV infection of CD4+ T-cells, Math. Biosci. 114 (1993) 81–125.
[25] A.S. Perelson, P.W. Nelson, Mathematical analysis of HIV-1 dynamics in vivo, SIAM Rev. 41 (1999) 3–44.
[26] E.D. Sontag, Mathematical Control Theory. Deterministic Finite Dimensional Systems, Springer, New York, 1995.
[27] M.A. Stafford, Modeling plasma virus concentration during primary HIV infection, J. Theor. Biol. 203 (2000) 285–301.
[28] Treatment of HIV Infection, October 2008, http://www.niaid.nih.gov/factsheets/treat-hiv.htm.
[29] U.S. Department of Health and Human Services, Guidelines for the Use of Antiretroviral Agents in HIV-1-Infected Adults and Adolescents, October 2008, http://www.aidsinfo.nih.gov/ContentFiles/AdultandAdolescentGL.p.