2011 50th IEEE Conference on Decision and Control and
European Control Conference (CDC-ECC)
Orlando, FL, USA, December 12-15, 2011

Nonlinear MPC of HIV-1 infection with periodic inputs

João V. Pinheiro, João M. Lemos and Susana Vinga

Abstract— **Although many papers address the issue of therapy design for HIV-1 infection based on control methods, the results available in the literature do not in general consider the fact that the manipulated variable is not continuous but a train of pulses. In addition, they are usually concerned only with the relation between drug effect (assumed to be manipulated) and virus dynamics. In order to improve the direct relation with actual clinical practice, the present work takes as manipulated variable a train of impulses that represent the pills taken by the patient and includes pharmacokinetics (PK) and pharmacodynamic (PD) drug models. In addition, patient adherence to treatment and their impact on virus drug resistance is also modeled. The problem of driving the viral load to a low specified value while minimizing the amount of drugs administered to the patient is thus addressed by nonlinear model predictive control (NMPC) with periodic inputs. Therefore, the paper contributions consist in the characterization of the results obtained with this type of control strategy in a HIV-1 infection model comprising drug PK and PD, development of virus resistance to drugs and virus dynamics. Furthermore, it is shown that various amounts of reverse transcriptase inhibitor (RTI) and protease inhibitor (PI) drugs can be given depending on the weights of the cost function minimized by periodic NMPC, while attaining the same control objective. It is proposed that these weights can be adjusted to minimize the toxicity of the drug cocktail administered.**

Index Terms— **Nonlinear Model Predictive Control, Periodic control, HIV-1 infection control, immunology, biomedical systems.**

I. Introduction

A. Motivation

In recent years, an increasing number of papers have addressed the issue of therapy design for HIV-1 infection based on control methods. Nevertheless, in general, the available literature does not consider the fact that the manipulated variable is a train of pulses and is usually concerned only with the relation between drug effect (assumed to be manipulated) and virus dynamics. In order to approach clinical practice, control algorithms must take into account that manipulated variables are drug doses taken periodically, usually once or twice a day, and resort to models that incorporate pharmacokinetics (PK) and pharmacodynamics (PD) of the drugs considered. Furthermore, modeling of patient adherence to prescribed therapy and the development of virus resistance and the impact of these issues on control performance is also important.

This work was supported by FCT under project *HIVControl – Control Based on dynamic modeling of HIV-1 infection for therapy design*, contract PTDC/EEA-CRO/100128/2008, and INESC-ID multiannual funding through the PIDDAC Program funds.

J. V. Pinheiro is with INESC-ID jvpinheiro@gmail.com

J. M. Lemos is with INESC-ID and IST/UTL, R. Alves Redol 9, 1000-029 Lisboa, Portugal jlml@inesc-id.pt

S. Vinga is with INESC-ID, R. Alves Redol 9, 1000-029 Lisboa, Portugal svinga@kdbio.inesc-id.pt

B. Literature review

A recent paper [1] addresses the application of periodic control to drug delivery. Since the seminal paper [2] that pointed the attention to the importance of phenomena that take place in time scales of days and weeks, that are fast when compared with the time scale of AIDS, many works have been reported in the literature concerning both model development and analysis [3] and therapy design using control algorithms [4], [5]. Other examples of control techniques used include nonlinear control [6], [7], [8] and optimal control [9], [10].

Predictive control is currently receiving an increased attention in relation to HIV-1. In [11] MPC is used to schedule interruptions in highly active anti-retroviral therapy (HAART) used to simulate a therapeutic vacine. Treatment schedules based on robust multirate MPC are proposed in [12]. In order to minimize drug consumption, [13] proposes a MPC based algorithm in which the dose (given with a sampling time of one week) is restricted either to be zero or the maximum acceptable value. In [14] a model of pharmacodynamics of antiretroviral drugs in HIV-1 infected patients that incorporates drug susceptibility and patient adherence is discussed.

C. Paper contributions and structure

The problem of driving the viral load in HIV-1 infection to a low specified value while minimizing the amount of drugs administered to the patient is addressed by a nonlinear model predictive control (NMPC) with periodic inputs. The paper contributions consist in the characterization of the results obtained with this type of control strategy in a HIV-1 infection model that includes drug PK and PD, virus resistance depending on patient adherence to therapy and virus dynamics. Furthermore, it is shown that various amounts of reverse transcriptase inhibitor (RTI) and protease inhibitor (PI) drugs can be given resulting in the same final virus load, depending on the weights of the cost function minimized by periodic NMPC, a result that may be used to minimize toxicity of the drug cocktail administered.

The paper is organized as follows: After the Introduction (this section) that motivates the problem, performs a concise literature review and describes the paper contributions and structure, the HIV-1 infection model used is presented in section II. This model is exploited in section III to address nonlinear state estimation, required for the control algorithm used, and in section IV to develop a nonlinear model predictive control (NMPC) algorithm with periodic inputs.

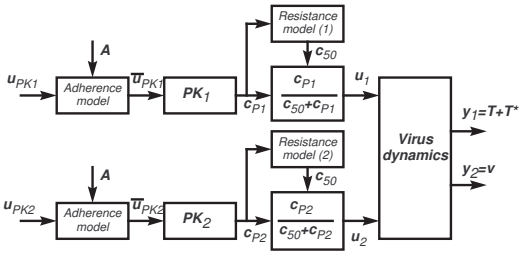

Fig. 1: Pharmacokinetic/pharmacodynamic model of HIV-1 infection incorporating adherence and resistance models.

This section also includes a discussion, based on simulation results, of the effect of controller parameters on the type of response obtained. Finally, section V draws conclusions.

II. PK/PD MODEL FOR HIV-1 INFECTION

Figure 1 shows the model of HIV-1 infection and drug interaction considered in this work, including a model of patient adherence to the therapeutics, drug pharmacokinetics and pharmacodynamic model. The pharmacodynamic model includes a model of the development of virus resistance to drugs, for patients without perfect adherence that is a variant of the one presented in [14]. The pharmacodynamic model includes a model of the development of resistances in response to non-perfect adherence.

It is assumed that two different types of drugs are being administered to the patient: A reverse transcriptase inhibitor (RTI, drug number 1) and a protease inhibitor (PI, drug number 2). Hence, corresponding to each drug, there is one pharmacokinetics and one pharmacodynamic model, corresponding to indexes 1 and 2 in the block diagram of figure 1.

The inputs to the the model are the drug doses of each type, u_{PK1} for the RTI dose and u_{PK2} for the PI dose and the patient adherence A. The drug doses are assumed to be trains of square pulses of very small duration and with an height proportional to the dose taken. The adherence A is defined as the probability that the patient will actually take the dose prescribed by the controller. Hence, at each discrete time k, the adherence model generates a random number α_A with a uniform distribution between 0 and 100. If $\alpha_A \leq A$ then the dose effectively taken of drug i, at time k, $\bar{u}_{PKi}(k)$ is made equal to $u_{PKi}(k)$; otherwise, it is made equal to 0.

The pharmacokinetics (PK) model of drug i (block PK_i in figure 1) is a linear, time invariant, model, with 3 poles, 3 zeros and unit delay. These PK models relate the amounts of drug administered to the patient with the drug plasma concentration c_{p1} and c_{p2}. Figure 2 shows the impulse response of the PK model used for the RTI drug. The impulse response used for PI is similar. These are typical curves, adapted from published drug PK data with modifications and taken for exemplificative purposes. Of course the results depend on the specific drugs considered.

The pharmacodynamic (PD) model consists of two cascade parts (figure 1). The first part relates the plasma concentration of each drug with the corresponding drug effect (u_1 for RTI and u_2 for PI). This is assumed to be the static nonlinear relation known as the "Hill equation":

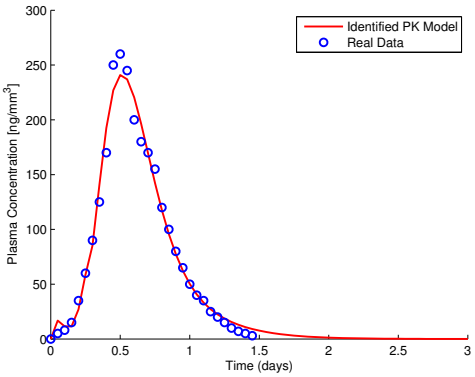

Fig. 2: RTI pharmacokinetics model impulse response.

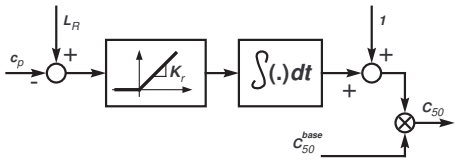

Fig. 3: Resistance model of HIV-1.

$$u_i = \frac{c_{pi}}{c_{50}^i + c_{pi}} \tag{1}$$

where c_{50}^i is a parameter for each i. For drug i, when the plasma concentration is $c_{pi} = 0$ the corresponding effect u_i is zero. When the plasma concentration c_{pi} grows, the effect u_i approaches 1. For $c_{pi} = c_{50}^i$, the effect is $u_i = 0.5$.

The resistance model tries to represent the effect that low plasma drug concentrations induce virus mutations that render the virus more resistant to the drug used. The increase of virus resistance means that the same effect can only be achieved with an higher plasma concentration. A way of representing this effect consists in increasing the c_{50}^i parameter depending on the periods in which the plasma concentration is below the threshold L_R that allows virus mutation [14]. This is the resistance model that is detailed in figure 3 (for simplicity the index of the drug has been dropped in this figure).

Figure 4 is a sketch of the time dependency of plasma concentration, used to explain the resistance model. It is assumed that the parameter C_{50} increases proportionally

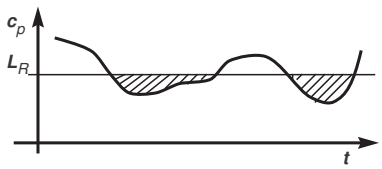

Fig. 4: Explaining the resistance model of HIV-1.

to the area of the shadowed zones. Whenever the plasma concentration decreases below L_R this area will increase, causing the factor f and, consequently, C_{50} to increase. In compact terms, this reads

$$C_{50}(t) = f(t)C_{50}^{base} \tag{2}$$

where C_{50}^{base} is a constant giving the initial value of C_{50} and $f(t)$ is given by

$$f(t) = 1 + K_r \int_0^t \max[0,\, L_r - c_p(\tau)]d\tau \tag{3}$$

where K_r is a parameter that measures the capacity of the virus to develop resistance in the presence of a drop of drug plasma concentration. For simplicity of notation the index of the drug has been omitted.

The second block of the PD model represents the interaction between drug effects u_1 and u_2 and virus dynamics. A simple model is used that considers only the following three state variables [2]:

- x_1, plasma concentration of healthy T-CD4+ cells;
- x_2, plasma concentration of infected T-CD4+ cells;
- x_3, plasma concentration of free virus particles (virions).

A balance, together with simple assumptions on interaction rates between the different elements involved yields the following nonlinear state space model:

$$\begin{cases} \dot{x}_1 &= s - dx_1 - (1-u_1)\beta x_1 x_3 \\ \dot{x}_2 &= (1-u_1)\beta x_1 x_3 - \mu x_2 \\ \dot{x}_3 &= (1-u_2)k x_2 - c x_3 \end{cases}$$
$$y = \begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} \tag{4}$$

where β, μ, k and c are patient dependent parameters.

For the sake of simplicity in mathematical manipulations, model (4) can be written in the following condensed form

$$\begin{aligned} \dot{x} &= \phi(x(t),\, u(t)) \\ y &= h(x,\, u) \end{aligned} \tag{5}$$

with obvious definitions for functions ϕ and h.

Since the control algorithms considered operate in discrete time, the above model is approximated using the Euler method. Let Δ be the sampling interval. With an obvious definition for Φ, the discrete model is written as

$$x(k+1) = \Phi(x(k),\, u(k)) \tag{6}$$

III. Nonlinear State Estimation

The NMPC algorithm used requires knowledge of the state variables. Since not all of these are directly measured, the Extended Kalman Filter (EKF) is used to produce state estimates according to the following equations:

For prediction:

$$\hat{x}_{k|k-1} = \Phi(\hat{x}_{k-1|k-1}, u_{k-1}) \tag{7}$$

$$P_{k|k-1} = F_k P_{k-1|k-1} F_k^T + Q_k \tag{8}$$

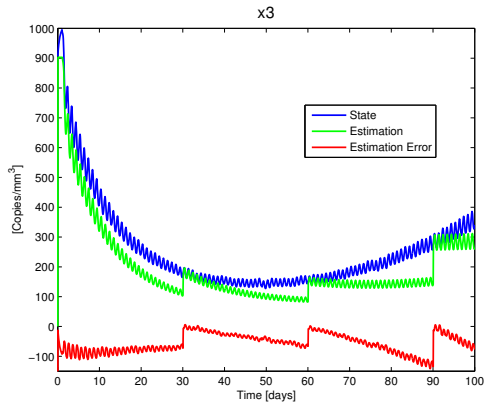

Fig. 5: Multi-rate EKF estimate of the viral load for a sampling interval of $T_m = 30\,day$.

with the jacobian given by

$$F_{k-1} = \frac{\partial \Phi}{\partial x}|_{\hat{x}_{k-1|k-1}, u_{k-1}} =$$
$$= \begin{bmatrix} 1 - \Delta(d+U_1)\beta x_3 & 0 & -\Delta U_1 \beta x_1 \\ \Delta U_1 \beta x_3 & 1 - \Delta\mu & \Delta U_1 \beta x_1 \\ 0 & \Delta U_2 k & 1 - \Delta c \end{bmatrix} \tag{9}$$

where $U_1 = 1 - u_1$, $U_2 = 1 - u_2$ and Δ is the sampling interval chosen to be a submultiple of the one used for control.

For update:

$$\tilde{y}_k = z_k - h(\hat{x}_{k|k-1}) \tag{10}$$

$$S_k = H_k P_{k|k-1} H_k^T + R_k \tag{11}$$

with

$$H_{k-1} = \frac{\delta h}{\delta x}|_{\hat{x}_{k-1|k-1}} = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{12}$$

and

$$\begin{aligned} K_k &= P_{k|k-1} H_k^T S_k^{-1} \\ \hat{x}_{k|k} &= \hat{x}_{k|k-1} + K_k \tilde{y}_k \end{aligned} \tag{13}$$

$$P_{k|k} = (I - K_k H_k) P_{k|k-1} \tag{14}$$

The noise variances are selected as

$$Q_k = \begin{bmatrix} 0.5 & 0 & 0 \\ 0 & 0.001 & 0 \\ 0 & 0 & 50 \end{bmatrix} \quad R_k = \begin{bmatrix} 0.5 & 0 \\ 0 & 0.0005 \end{bmatrix} \tag{15}$$

and the initial error covariance matrix is $P_{0|0} = I$.

A multi-rate version of EKF is used. The prediction stage is used in every iteration while filtering is used in a multiple of this time unit. Figure 5 shows the result for the estimate of the viral load (x_3). In order to test filter robustness, a value of k with a 10% error has been used.

IV. Periodic Nonlinear MPC of HIV-1 Infection

The control problem consists in bringing down the viral load, that corresponds to the state x_3, to a level of 50 $copies/mm^3$ and keeping it there, in less than 56 days. The manipulated variables are the drug doses u_{PK1} and u_{PK2}

A. NMPC algorithm

The NMPC control is obtained by solving in a receding horizon way, at each current discrete time t, the following constrained minimization problem:

$$\min_{\bar{u}_t^{t+T_c-1}} J(x(t), \bar{u}_t^{T_c}, T_c, T_p) \tag{16}$$

where T_c is the control horizon and T_p is the prediction horizon, verifying $T_p \geq T_c$. The sequence of virtual control $\bar{u}$ is

$$\bar{u}_t^{t+T_c-1} := [\bar{u}(t) \ \ldots \ \bar{u}(t+T_c-1)]^T$$

The receding horizon cost J is defined as

$$J(x(t), \bar{u}_t^{T_c}, T_c, T_p) = \sum_{k=1}^{T_p} L[\bar{x}(t+k), \bar{u}(t+k-1)] \tag{17}$$

where the stage cost is defined by

$$L(x, u) = (x - x_s)^T Q (x - x_s) + u^T R u$$

with

$$Q = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad R = \begin{bmatrix} w_{u_1} & 0 \\ 0 & w_{u_2} \end{bmatrix} \tag{18}$$

and x_s is the reference, given by

$$x_{s,3}(t) = 50 + (902.8 - 50)e^{-\tau_r t} \tag{19}$$

The variable τ_r is the time constant that determines the speed at which the reference signal decreases towards the desired value of 50 $copies/mm^3$, and was selected as $\tau_r = 0.2\ days^{-1}$.

It is assumed that, for $i = 1, \ldots, T_p - T_c$

$$\bar{u}(T_c + i) = \bar{u}(T_c) \tag{20}$$

The virtual state trajectory satisfies the system dynamics with initial condition $\bar{x}(t)$ given by the actual state x at time t. Since the actual state is not available for direct measurement, it is replaced by its EKF estimate, $\hat{x}$. Thus, for $k = 0, \ldots, T_{p-1}$:

$$\bar{x}(t+k+1) = \Phi(\bar{x}(t+k), \bar{u}(t+k)) \tag{21}$$

starting from

$$\bar{x}(t) = \hat{x}(t) \tag{22}$$

According to a receding horizon strategy, once the problem (16) is solved with respect to the virtual control $\bar{u}$, the value of the manipulated variable to actually apply to the plant at time t is $u(t) = \bar{u}(t)$. The same procedure is again repeated at time $t + 1$. Figures 6 and 7 show the viral load and the drug dosage for a patient with an adherence of 95% controlled with NMPC. Figure 8 shows the corresponding plasma concentration of RTI and PI. As seen in figure 7 the doses of both drugs applied grow initially such as to drive the viral load to a low level. Then, the doses are reduced, to keep the viral load where desired. After the fail of a control sample, the controller raises the value of subsequent samples to compensate the drop of drug concentration in plasma.

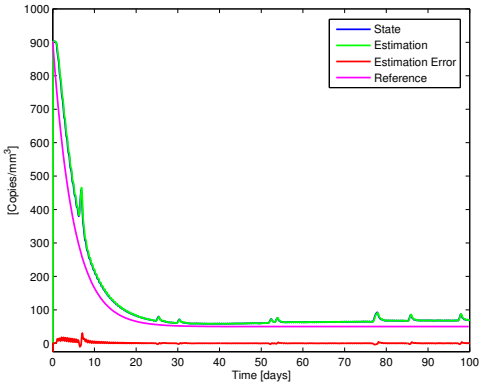

Fig. 6: Viral load response for a patient with adherence of 95% and $K_r = 1 \times 10^{-4}$.

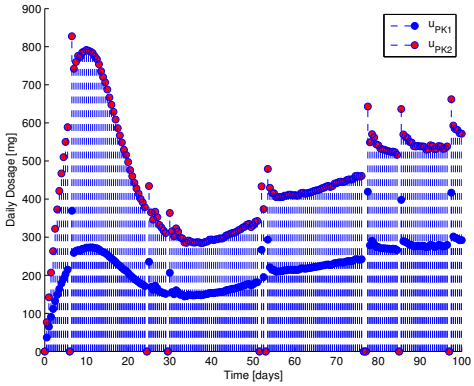

Fig. 7: Dosage for a patient with adherence of 95% and $K_r = 1 \times 10^{-4}$.

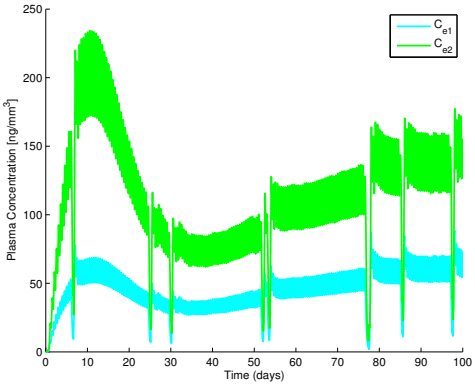

Fig. 8: Drug plasma concentration in the situation of figures 6 and 7.

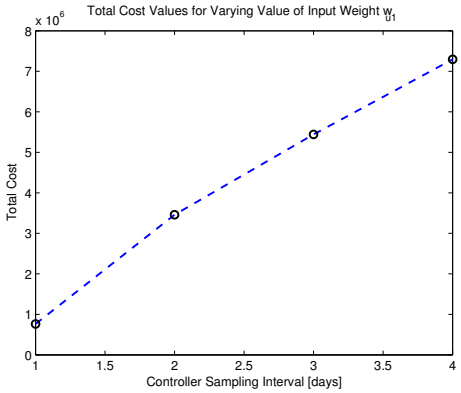

Fig. 9: Cost and Average Error Values for Varying Controller Sampling Interval Values

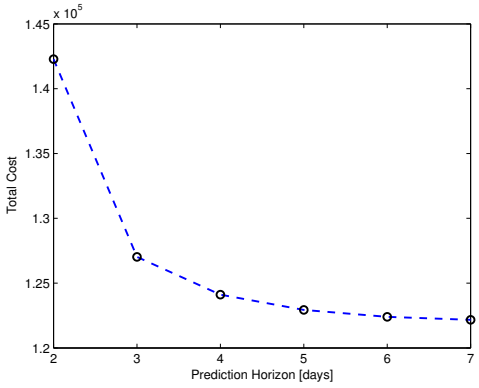

Fig. 10: Cost and Average Error Values for Varying Prediction Horizons

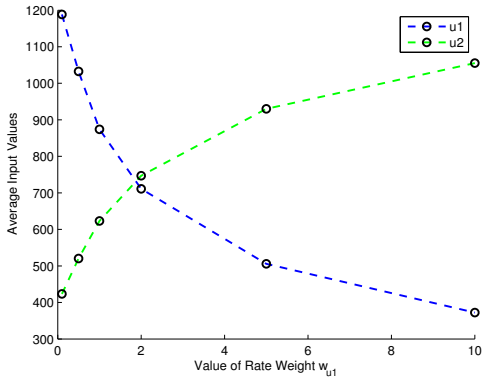

Fig. 11: Average input Values as a function of the input weight w_{u_1}.

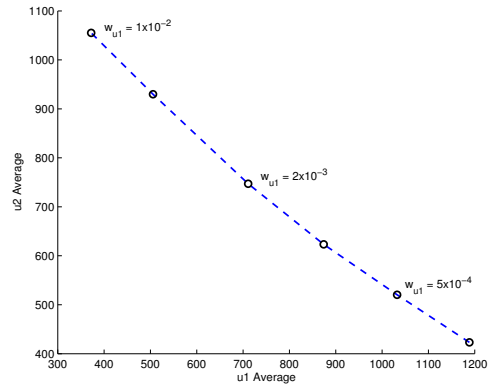

Fig. 12: Selecting the average input values for u_{PK1} and u_{PK2} by varying the value of the input weight w_{u_1}

B. Selecting the controller parameters

1) Sampling interval: Figure 9 shows the dependency of the cost on the sampling interval.

2) Prediction horizon: Figure 10 shows the dependence of the cost on the prediction horizon. It should be remarked that the horizon actually "seen" by the algorithm is the product of T_p by the sampling interval. Smaller sampling intervals require a bigger T_p so that the same performance is achieved. For a given sampling rate, when T_p increases, the performance increases. However, this increase is smaller and smaller as T_p gets bigger. Due to plant model mismatches, it is expected that, when T_p grows, after the initial decrease of the cost this will increase again, for larger values of T_p, due to the degradation of prediction performance over long time horizons.

3) Cost weights and toxicity: The plots in figure 11 show the effect of the cost weight parameter w_{u_1}. Plots for the influence of w_{u_2} are similar, being omitted for lack of space. When w_{u_1} increases there is a decrease in the average of the amount of RTI that is administered to the patient and an increase in PI. Thus, the parameter w_{u_1} regulates a trade-off between these two types of drug (shown in figure 12) that can be explored to minimize toxicity of the overall treatment. Denote by $< u_{PKi} >$ the average value of u_{PKi} and by g_i the function that measures toxicity for drug i, $i = 1, 2$. The average toxicity of both dugs is given by

$$T_x = g_1(< u_{PK1} >) + g_2(< u_{PK2} >)$$

Using NMPC, the average values of both drugs are related by the function plotted in figure 12, that we call Υ. Hence, we can find a value for the weight w_{u_1} that minimizes treatment toxicity, by finding the function u_{PK1} that solves the optimization problem

$$\min_{<u_{PK1}>} \{g_1(< u_{PK1} >) + g_2(\Upsilon(< u_{PK1} >))\} \qquad (23)$$

and then computing the corresponding value for w_{u_1} from figure 12.

C. Modeling errors

Figure 13 shows the dependence of the cost on modeling errors for parameter k. There is a graceful degradation of performance in the presence of a mismatch between the model used by the controller and the "true" model. The behavior is similar with respect to other parameters.

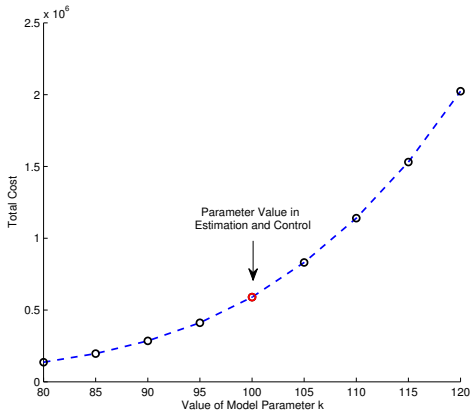

Fig. 13: Cost as a function of model parameter k value.

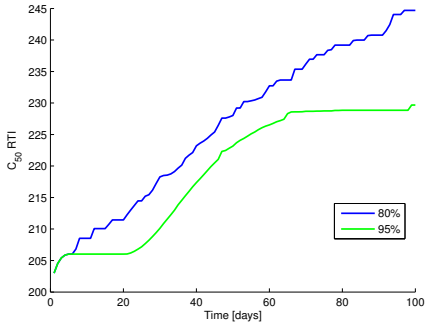

Fig. 14: Time evolution of C_{50} for two patients with different adherence.

D. Patient adherence and virus mutation

An adherence value smaller than 1 has two consequences: First, the drug doses not taken in due time are seen by the controller as disturbances. More important, they may lead to an increase of virus mutation probability that is modeled by an increase in C_50. Figure 14 shows the time increase of C_{50} for two patients with different adherence and figure 15 shows the dependency of optimal cost on the patient adherence parameter. When adherence reduces the performance quickly degrades.

V. Conclusions and Discussion

The use of a periodic nonlinear MPC, together with a model incorporating drug PK and PD, virus dynamics and patient adherence provides an adequate way to compute optimal drug dosage for HIV-1 infection therapy. For a patient adherence of 90% the controller could still meet the specifications. The cost dependence on controller parameters as well as on model parameters have been studied, including parameter mismatching. The choice of weights in the cost function minimized by periodic NMPC provides a degree of freedom that can be explored to optimize therapy toxicity.

Future work should remove a number of simplifications in order to make the controller more realistic. These include assuming drug dose quantification, the use of PK models for specific drugs, inclusion of an effect compartment and improved observation noise models.

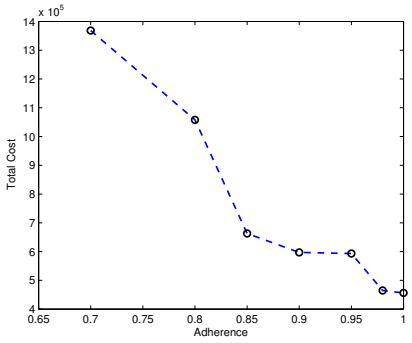

Fig. 15: Cost as a function of adherence.

References

[1] S. Varigonda, T. T. Georgiou, R. A. Siegel and P. Daoutidis (2008). Optimal periodic control of a drug delivery system. *Computers and Chemical Engineering*, 32:2256-2262.

[2] A. S. Perelson and P. W. Nelson (1999). Mathematical analysis of HIV-1 dynamics in vivo. *SIAM Review*, 41(1):3-44.

[3] A. L. Knorr and R. Srivastava (2005). Evaluation of HIV-1 kinetic models using quantitative discrimination analysis. *Bioinformatics*, 21(8): 1668-1677.

[4] H. Chang and A. Astolfi (2009). Control of HIV infection dynamics: Approximating high-order dynamics by adapting reduced order model parameters. *IEEE Control Systems Mag.*, 28(April):28-39.

[5] H. Chang and A. Astolfi (2009). Enhancement of the immune system in HIV dynamics by output feedback. *Automatica* 45:1765-1770.

[6] S. S. Ge, Z. Tian and T. H. Lee (2005). Nonlinear control of a dynamic model of HIV-1. *IEEE Trans. Biomed. Eng.*, 52(3):353-361.

[7] M. Barão and J.M. Lemos (2007). Nonlinear control of HIV-1 infection with a singular perturbation model, *Biomedical Signal Processing and Control* 2:248-257.

[8] M. Mhawej, C. H. Moog, F. Biafore and C. Brunet-Franois (2010). Control of the HIV infection and drug dosage. *Biomedical signal processing and control*, 5:45-52.

[9] J. J. Kutch and P. Gurfil (2002). Optimal control of HIV infection with a continuously-mutating viral population. *Proc. Amer. Control Conf.*, Anchorage, AK May 8-10, 2002. 4033-4038.

[10] R. F. Stengel (2008). Mutation and control of the human immunodeficiency virus. *Mathematical Biosciences*, 213:93-102.

[11] R. Zurakowski and A. R. Teel (2006). A model predictive control based scheduling method for HIV therapy. *J. of Theoretical Biology*, 238:368-382.

[12] A. M. Elaiw and X. Xia (2009). HIV dynamics: Analysis and robust multirate MPC-based treatment schedules, 359:285-301.

[13] G. Pannocchia, M. Laurino and A. Landi (2010). A model predictive control strategy toward optimal structured treatment interruptios in anti-HIV therapy. *IEEE Trans. Biomed. Eng.*, 57(5):1040-1050.

[14] H. Wu, Y. Huang, E. P. Acosta *et al.* (2006). Pharmacodynamics of antiretroviral agents in HIV-1 infected patients: using viral dynamic models that incorporate drug susceptibility and adherence. *J. of Pharmacokinetics and Pharmacodynamics*, 33(4):399-413.