Optimal Scheduling of Drug Treatment for HIV Infection: Continuous Dose Control and Receding Horizon Control

H. Shim*, S.J. Han*, I.S. Jeong**, Y.H. Huh**, C.C. Chung**, S.W. Nam** and J.H. Seo*

* School of Electrical Engineering, Seoul National University, 151-744, Korea
(Email: h.shim@ieee.org, sjhan75@hanmail.net, jhseo@snu.ac.kr)
** Division of Electrical and Computer Engineering, Hanyang University, Seoul, 133-791, Korea
(Email: insuk2000@hanmail.net, huyh7@hanmail.net, cchung@hanyang.ac.kr, swnam@hanyang.ac.kr)

Abstract: It is known that HIV (Human Immunodeficiency Virus) infection, which causes AIDS after some latent period, is a dynamic process that can be modeled mathematically. Effects of available anti-viral drugs, which prevent HIV from infecting healthy cells, can also be included in the model. In this paper we illustrate control theory can be applied to a model of HIV infection. In particular, the drug dose is regarded as control input and the goal is to excite an immune response so that the symptom of infected patient should not be developed into AIDS. Finite horizon optimal control is employed to obtain the optimal schedule of drug dose since the model is highly nonlinear and we want maximum performance for enhancing the immune response. From the simulation studies, we find that gradual reduction of drug dose is important for the optimality. We also demonstrate the obtained open-loop optimal control is vulnerable to parameter variation of the model and measurement noise. To overcome this difficulty, we finally present nonlinear receding horizon control to incorporate feedback in the drug treatment.

Keywords: Chemotherapy, HIV, Optimal control, Receding horizon control

1. Introduction

Human Immunodeficiency Virus (HIV) is the virus that causes Acquired Immune Deficiency Syndrome (AIDS). Infection with HIV can weaken the immune system in human body to a level at which it has difficulty fighting off certain infections. This is because a major target of HIV is the CD4 T-helper cell, which is a key component of the adaptive immune system in human body, and HIV infection reduces the number of CD4 cells in the body. AIDS is usually judged by counting CD4 cells; when the count of CD4 cells is below $200/mm^3$ in blood, the HIV infected person is regarded as an AIDS patient [12]. With very few CD4 cells the immune system of a patient cannot function normally, and thus, the patient is very vulnerable to other infections. These types of infections are known as 'opportunistic infections' because they take the opportunity a weakened immune system gives to cause illness or to result in death.

Although HIV weakens the immune system, it still works inside the patient. For example, APC[1] of a HIV infected patient still signals precursor CTL (Cytotoxic T Lymphocytes) cells to differentiate into killer T cells (effector CTL). In addition, killer T cells still destroy infected CD4 cells in which new virus is born. On the other hand, there are several developed drugs which can inhibit HIV from infecting CD4 cells. This means that HIV infection process is a quite complicated interaction among many different cells, virus and drug. Therefore, some quantitative analysis (as well as a qualitative one) might be important.

In order to gain much insight about this complicated phenomenon, mathematical models of HIV infection (including the effects of drugs) have been developed in, for example, [1, 7, 11, 13, 14]. Based on these models, control engineers have also studied optimal drug dose control problems (see, e.g., [2,4,6,8,9,16,17,21]). In particular, Wodarz and Nowak [18–20] have recently presented a model in which both the memory CTL precursor and the memory CTL effector are appropriately described, and have shown that the medication can be stopped while the viral load and the number of uninfected CD4 cells remain at a low and high level, respectively, so that the HIV infected patient would not progress to AIDS. Zurakowski and Teel [21] have bestowed some control engineering concept upon the result of Wodarz and Nowak by applying Model Predictive Control method to the model of [19].

This paper further investigates the model of [7, 20] in the sense that continuous variation of drug dose is allowed. Our result shows that a *gradual reduction* of dose leads to maximally excited CTL response, which has not been addressed in the previous works of [18–21] where either full dose or no medication is allowed. In Section 2, we formulate our problem and give some analysis of the model considered in this paper. Optimal control result is illustrated in Section 3 while feedback control via receding horizon control technique is presented in Section 4. Conclusions are given in Section 5.

2. Model Description and Analysis

There are already several mathematical models available which describe interactions between HIV and immunocytes in human body. (See [14] or [11] for comprehensive exposition of the HIV infection modeling.) Among them, we have chosen a model from [7, 20] because the CTL response is appropriately modeled[2].

This work has been done while the first author is with Hanyang Univ. This work was supported by the Basic Research Program of the Korea Science and Engineering Foundation (Grant No. KOSEF R01-2002-000-00227-0).

[1] Antigen-Presenting Cell, whose abundance is approximated by the number of infected cells in the model (1) of the next section.

[2] It is interesting to see that other models such as used in [2, 4, 6, 9, 16] do not have memory CTL terms, so that termination of drug treatment again

The model considered in this paper is given by

$$\begin{aligned}
\dot{x}(t) &= \lambda - dx(t) - \beta(1-\eta u(t))x(t)v(t) \\
\dot{y}(t) &= \beta(1-\eta u(t))x(t)v(t) - ay(t) - py(t)z(t) \\
\dot{v}(t) &= ky(t) - \mu v(t) \\
\dot{w}(t) &= cx(t)y(t)w(t) - cqy(t)w(t) - bw(t) \\
\dot{z}(t) &= cqy(t)w(t) - hz(t)
\end{aligned} \tag{1}$$

where all the states implies the population of specified cells (or virus) in a unit volume of blood (and therefore they are meaningful only when positive). Each of them is uninfected CD4 T-helper cell (x), infected CD4 T-helper cell (y), memory CTL precursor (w), memory CTL effector (z) and free virus (v). The input (u) represents the drug dose, which may have values between 0 and 1. If $u = 1$, a patient is fully dosed, while zero input implies no medication at all.

An interpretation of the model is not very difficult. For example, the population of healthy (uninfected) CD4 T-helper cell increases at a rate λ (since it is produced from thymus), and decreases at a rate dx (since a cell dies naturally) which is modeled to be proportional to the population of x. The CD4 T-helper cell (x) is also a target of HIV (v) so that its population decreases proportionally to $x(t)$ and $v(t)$. When the cell x is infected, it becomes the infected cell y that generates new virus, which is modeled by the term $ky(t)$ in the third equation. The infected cell y and the virus v also die out at a rate ay and μv, respectively. The model (1) also describes the 'adaptive cell-mediated immune system' equipped in human body. They consist of the CTL precursor w (a cell which provides a long-term memory for a specific antigen; HIV in this paper) and the CTL effector z (a cell which actually kills the infected cell y at a rate pyz; see the second equation). The cell w differentiates into the cell z at a rate of $cqyw$ (that is, the given model implies that the larger population of the infected cell y and the CTL precursor w makes the population of z increase more quickly). Finally, the term $cxyw$ in the fourth equation implies that the CTL precursor w is generated at a rate proportional to the number of x, y and w itself.

The model (1) also includes the effect of the drug known as Reverse Transcriptase Inhibitor (RTI), whose role is to inhibit the virus from infecting new cells by preventing the reverse transcription[3]. Therefore, the control input u, representing RTI, may alter the infection rate β to $\beta(1-\eta)$ when the drug is maximally prescribed (i.e., $u = 1$). Here, η is a model parameter indicating the effect of a drug; $\eta \in [0, 1]$.

We have also taken the parameters in the model from [7], which are $\lambda = 1$, $d = 0.1$, $\beta = 0.02$, $a = 0.2$, $p = 1$, $c = 0.027$, $q = 0.5$, $b = 0.001$, $h = 0.1$, $k = 25$, $\mu = 1$, and $\eta = 0.98$. The time derivative in the model is taken with respect to time of a day, i.e., $x(1)$, $x(2)$, $\cdots$ imply the quantity of x at the first day, the second day and so forth.

leads to the proliferation of viral load. Although those models fit well to the description of the initial phase of HIV infection, it is not suitable for developing a long-term plan of medical treatment for HIV patients.

[3]Reverse transcription is a process of HIV infection that the genetic code of HIV to join the DNA of host cell (y in the model) in order to force the host cell to produce another HIV.

With these parameters at hand, the equilibrium points are easily found by making the right-hand side of (1) zero. We illustrate four equilibrium points that the model has when there's no medication ($u = 0$).

Point 1:

$$\bar{x} = \frac{\lambda}{d}, \quad \bar{y} = \bar{v} = \bar{w} = \bar{z} = 0.$$

This point is for a person not having HIV. Stability analysis by the local linearization of the model with given parameters shows that this point is an unstable equilibrium. The model, therefore, asserts that it is impossible to revert a patient, once infected, back to the normal status before infection with a retraction of medication.

Point 2:

$$\bar{x} = \frac{a\mu}{\beta k}, \quad \bar{y} = \frac{\mu\bar{v}}{k}, \quad \bar{v} = \frac{\lambda - d\bar{x}}{\beta\bar{x}}, \quad \bar{w} = \bar{z} = 0.$$

With the given parameters, it has the values of $\bar{x} = 0.4$, $\bar{y} = 4.8$ and $\bar{v} = 120$, and therefore, this is the status of a patient, for whom HIV dominates. Stability analysis shows it is a stable equilibrium.

Point 3:

$$\bar{x} = \frac{c\mu(\lambda+dq) - kb\beta + \sqrt{[c\mu(\lambda+dq) - kb\beta]^2 - 4c^2dq\lambda\mu^2}}{2cd\mu},$$

$$\bar{y} = \frac{b}{c(\bar{x}-q)}, \quad \bar{v} = \frac{k\bar{y}}{\mu}, \quad \bar{w} = \frac{h\bar{z}}{cq\bar{y}}, \quad \bar{z} = \frac{\beta\bar{x}\bar{v} - a\bar{y}}{p\bar{y}}.$$

With our paramters, it has

$$[\bar{x}, \bar{y}, \bar{v}, \bar{w}, \bar{z}]^T = \Big[9.8, 0.004, 0.1, 8751, 4.7\Big]^T =: X_{eq}, \tag{2}$$

which is locally stable. It is seen that the numbers of viral load and infected cells are maintained small while the CTL precursor has a large number, which is desired. Also, because it is locally stable, our control goal will be to drive an initial state X_0 to a neighborhood of X_{eq}. Then, in spite of treatment suspension, a patient is switched into the status of long-term non-progression to AIDS.

Point 4:

$$\bar{x} = \frac{c\mu(\lambda+dq) - kb\beta - \sqrt{[c\mu(\lambda+dq) - kb\beta]^2 - 4c^2dq\lambda\mu^2}}{2cd\mu},$$

$$\bar{y} = \frac{b}{c(\bar{x}-q)}, \quad \bar{v} = \frac{k\bar{y}}{\mu}, \quad \bar{w} = \frac{h\bar{z}}{cq\bar{y}}, \quad \bar{z} = \frac{\beta\bar{x}\bar{v} - a\bar{y}}{p\bar{y}}.$$

With our parameters, it has

$$[\bar{x}, \bar{y}, \bar{v}, \bar{w}, \bar{z}]^T = [0.51, 3.72, 93.05, 0.11, 0.055]^T,$$

which is unstable. This point is not of interest for our treatment.

For the presentation, two sets of initial conditions are chosen by

$$\begin{aligned}
X_a :&= [x(0), y(0), v(0), w(0), z(0)]^T \\
&= [0.4,\ 4.8,\ 119.9,\ 0.0001,\ 0.0001]^T
\end{aligned}$$

and

$$X_b := [x(0), y(0), v(0), w(0), z(0)]^T = [9.94,\ 0.0069,\ 0.189,\ 0.0026,\ 8.43 \times 10^{-6}]^T.$$

The first one (X_a) represents a patient who was infected by HIV quite a long time ago and thus virus population is dominant while the second one (X_b) is for a patient who has been treated by RTI for a long time after infection[4].

3. Optimal Scheduling of Drug Treatment

In order to solve the problem of enhancing the CTL response, formulated in the previous section, we formally pose an input-constrained finite-horizon optimal control problem given as follows.

Problem: Find an optimal control $u^*(\cdot) : \mathbb{R}_{[0,T_f]} \to [0,1]$ which minimizes the performance index

$$J(X_0, u(\cdot)) = (X(T_f) - X_{eq})^T Q_f (X(T_f) - X_{eq}) + \int_0^{T_f} u^2(t)dt \quad (3)$$

where $X(t) = [x(t), y(t), v(t), w(t), z(t)]^T$ is the solution trajectory of system (1) initiated at time $t = 0$ by X_0.

For this problem to be meaningful, the target equilibrium X_{eq} is taken by (2) and the initial condition X_0 of system (1) will be taken either by X_a or X_b in the previous section. We also take $T_f = 420$(days) so that the medication schedule of more than one year would be obtained. In the terminal cost, we select $Q_f = \mathrm{diag}(1, 1, 1, 0.001, 1)$ so that the cost for the final state $w(T_f)$ is less weighted. Small weighting to the state w is due to the fact that the dynamics for the CTL precursor (w) is quite slow and with the finite-horizon of 420 days it is not fully developed, and thus, at the end of optimization horizon, the difference between $w(T_f)$ and $\bar{w} = 8751$ (target point) would be of the order of several thousands.

The problem to be solved is a continuous-time optimal control problem, which means the decision vector $u(\cdot)$ belongs to the infinite-dimensional vector space. However, to solve it by the digital computer, it is necessary to approximate $u(\cdot)$ by some finite-dimensional vector and discretize the nonlinear dynamics (1). For our problem, the optimization horizon of 420 days are divided by 420 knot points so that sampling period becomes one day and the continuous signal $u(\cdot)$ has 420 knot points. Although the closed form of discrete time model is not available, we are still able to handle our problem by integrating the continuous time model (1) using Runge-Kutta integration (see, e.g., [3]). To increase the accuracy, we use variable step-size Runge-Kutta algorithm with piecewise linear interpolation of $u(t)$ during the sampling period (i.e., between the knot points). For these purposes, a package RIOTS[5] running in MATLAB environment is used in this paper.

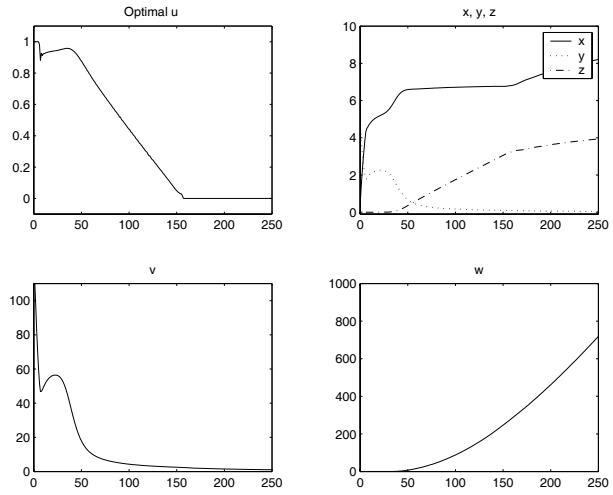

Fig. 1. Simulation result for the patient A. Horizontal and vertical axes indicate time (order of a day) and the value of states, respectively. Plot of optimal control u: In the early stage of approximately 1 week, the drug is fully dosed for the reduction of viral load. However, some amount of virus is necessary in order to increase the CTL response (the state w) that changes slowly. After about 40 days, the optimal dose decreases gradually, which is a compromise between two goals; small number of virus and large number of memory CTL. Plot of states x, y, z: The number of uninfected cell (x) goes back to normal, and infected cell (y) decays, which is caused by the increased CTL z. Plot of virus v: It decays. Plot of CTL w: It goes toward the desired equilibrium point 3.

The solution for our problem is illustrated in Fig. 1 and Fig. 2, which are for the patients A and B having the initial condition of X_a and X_b, respectively. From these simulation results, it is clear that the optimal drug dose shows the pattern of *gradual reduction* taking about 3 months. While it has been shown in [18–21] that the termination of medication is possible with the stimulation of CTL response, it has not been reported that gradual reduction of dose would enhance the CTL response optimally, to the authors' knowledge.

In Fig. 2, it can be noticed that a patient is not dosed for around 3 days. Since the initial condition X_b is for the patient who has been dosed for a long time, this result implies that an interruption of treatment is necessary[6], which has

[4] In fact, we have identified those points X_a and X_b by simulating the model (1) with an initial condition $[10, 0.0001, 0.01, 0.0001, 0.0001]$ representing newly infected patient (compare the equilibrium point 1). X_a is obtained by integrating (1) for 60 days without medication, and X_b is obtained by integrating (1) again from X_a for another 60 days with full medication $u = 1$.

[5] Recursive Integration Optimal Trajectory Solver: A commercial package to solve continuous time optimal control problem for nonlinear dynamics.

[6] The fourth equation of (1) can also be written by $\dot{w} = (c(x - q)y - b)w$. Since our ultimate goal is to have a large number of w so that the CTL effector z would suppress the infected cell (and thus, the virus), it is necessary to have positive value of $(c(x - q)y - b)$ in the early phase of HIV infection. This can be achieved when both the uninfected cell x and the infected y need to be abundant, but by the drug treatment for a long former days the abundance of infected cell y might be too small. In this case, some interruption of treatment will enhance the population of v and y, which in turn enhances the CTL response w.

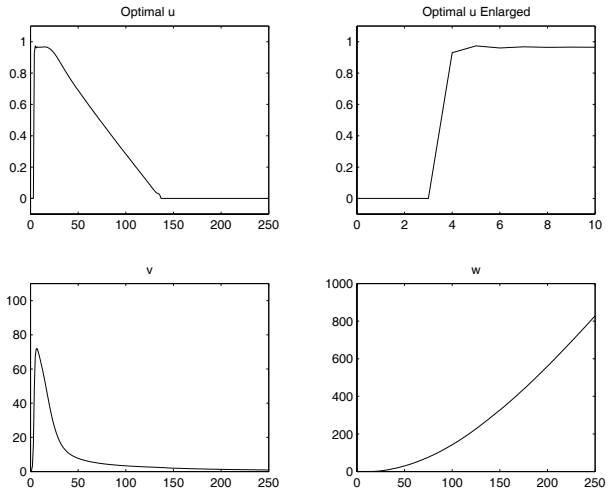

Fig. 2. Simulation result for the patient B. This is the case when the patient has been treated by RTI for a long period so that the initial conditions for viral load v, CTL precursor w and CTL effector z have very small values. Since the CTL response cannot be excited with these values, the drug input initially maintains zero value for short period of time (3 days), after that it is rapidly changed to nearly full dose. Again, gradual reduction of dose is observed after the initial period.

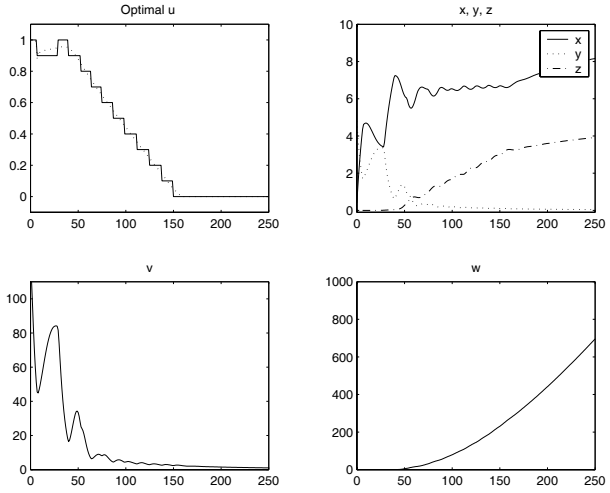

Fig. 3. Simulation result for the patient A. For practical implementation, drug dose is quantized by 10 levels. Nevertheless, it still maintains similar trend of the state trajectories in Fig. 1.

been already observed in [18–20]. However, our result of Fig. 2 further shows that, rapid increase and gradual reduction of dose would be better for the same case.

Finally, we have tried quantized dose treatments of drug in Fig. 3 because continuous variation of dose seems hard to apply in the real treatment of patients. Here, we have divided the dose by 10 levels. The value of performance index J in this case is 50142 which is not very different from the optimal case of continuous drug change ($J = 49796$). To the contrary, an abrupt change of dose from 1 to 0 leads to the cost of 60307, which is quite far from the optimal. We

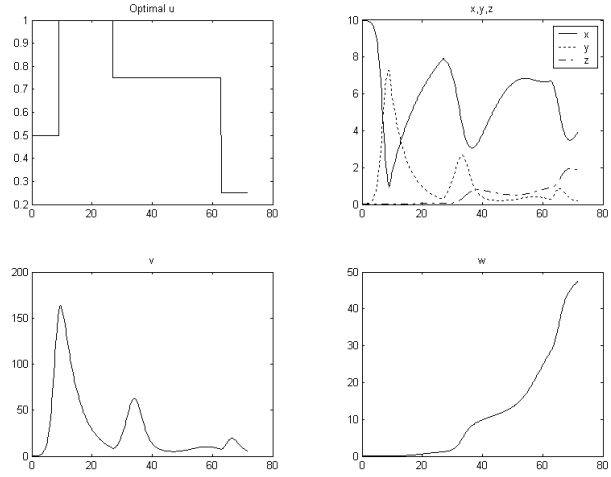

Fig. 4. Simulation result for the patient B. The same optimal control problem is solved with a constraint that the input u can have 0, .25, .5, .75 and 1 only. The result still shows the similar pattern of gradual dose reduction as in Fig. 2.

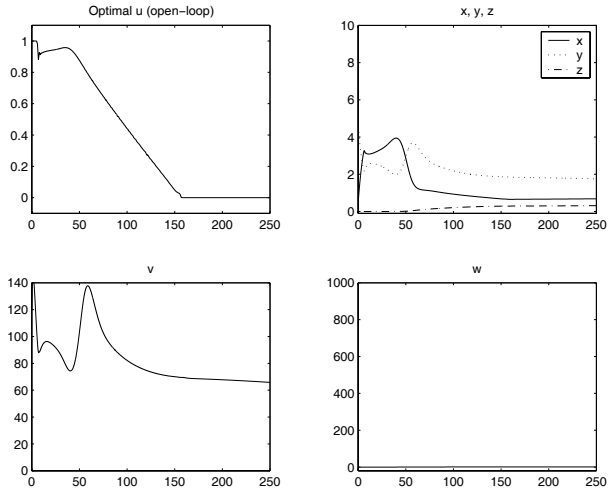

Fig. 5. Blind application of the (open-loop) optimal control u of Fig. 1 to the system having parameter variation of $\pm 20\%$. Viral load v does not decrease and the CTL response w is not excited.

also have found the optimal dose control with the constraint that the input u can only have 0, .25, .5, .75, and 1 with sampling period of 2 weeks. The result is given in Fig. 4 where the similar gradual reduction of dose can be observed.

4. Employing Feedback: Receding Horizon Control

In the real environment, parameter uncertainty in the model and measurement noise are inevitable. Fig. 5, for example, shows the case that the optimal drug control is calculated from the given nominal model (i.e., the control in Fig. 1) and is applied to a model having perturbed parameters. For the simulation, each parameter is perturbed by $\pm 20\%$ from the nominal values; that is, $\lambda = 0.8$, $d = 0.12$, $\beta = 0.016$, $a = 0.16$, $p = 0.8$, $c = 0.0216$, $q = 0.6$, $b = 0.0012$,

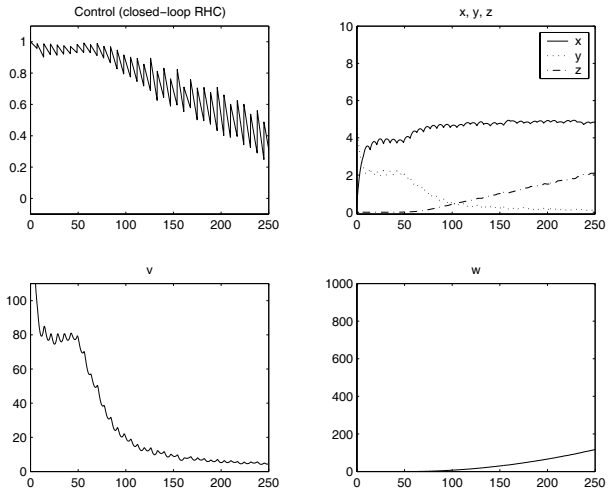

Fig. 6. Simulation result for the patient A. Parameters are perturbed by ±20% and random noise of ±10% is added to each measurement of the state. It is observed that, although the reaction of the system to the drug treatment becomes quite slow, it still excites the CTL response.

$h = 0.08$, $k = 30$ and $\mu = 0.8$ (η is the same as before). The result shows that the CTL response is not excited and viral load remains at high level of abundance. To overcome this problem, some type of feedback control seems necessary.

Receding Horizon Control (or Model Predictive Control which is now alternatively used in the literature) is a feedback scheme that has been widely employed in many practical situations, and is known to be a way to incorporating 'feedback' when the control is obtained from an open-loop optimal control problem [10]. Main idea is to solve an open-loop finite horizon optimization first, and the resulting control trajectory is applied to the system for a fraction of the optimization length. At the end of the fraction of time, the optimization is solved again with new initial condition measured at that time and the resulting control is applied again for the fraction of time. This process is repeated, which leads to a sampled feedback control.

While the idea of using RHC for the HIV infection control has been already presented in [21] where drug dose has either 0 (no medication) or 1 (full dose), we propose, in this paper, the receding horizon control with continuous change of drug dose. For the stability of closed-loop system, we take the optimization horizon large enough as 700 days. (See [5] for a justification that long enough optimization horizon guarantees the closed-loop stability.) The sampling period for the simulation is 1 week; that is, the open-loop control is obtained every 7 days and it is applied to the system for 7 days while the rest of control trajectory is discarded. The result is given in Fig. 6, in which it is observed that the CTL response is excited and viral load decreases, even under ±10% measurement noise (and the same parameter variation as Fig. 5).

5. Conclusions

In this paper, we have shown that a continuous-time nonlinear optimal control scheme can be employed for enhancing the CTL immune response by optimally scheduling the drug treatment for HIV infected patients. Our finding from simulation studies indicates that gradual reduction of drug dose is important for enhancing the CTL immune response. This point has not yet been exploited in the similar results of [7, 18–21]. In addition, to overcome the problem that the obtained open-loop control may not be directly applied to the system having uncertain parameters to some extent, we have applied nonlinear receding horizon control scheme to our system, which results in moderate performance under measurement noise and parameter uncertainty.

References

[1] H.K. Altes, D.A. Price and V.A. Jansen, "Effector cytotoxic T lymphocyte numbers induced by vaccination should exceed levels in chronic infection for protection from HIV," *Vaccine,* vol. 20, pp. 3–6, 2002.

[2] M.E. Brandt and G. Chen, "Feedback control of a biodynamical model of HIV-1," *IEEE Trans. on Biomedical Engineering,* vol. 48, no. 7, pp. 754–759, 2001.

[3] A.E. Bryson, *Dynamic Optimization,* Addison Wesley Longman, Inc., 1999.

[4] J.A.M. Felippe de Souza, M.A.L. Caetano and T. Yoneyama, "Optimal control theory applied to the antiviral treatment of AIDS," *In proceedings of Conference on Decision and Control,* Sydney, 2000.

[5] G. Grimm, M.J. Messina, A.R. Teel and S. Tuna, "Model predictive control when a local control Lyapunov function is not available," *In proceedings of American Control Conference,* 2003.

[6] D. Kirschner, S. Lenhart and S. Serbin, "Optimal control of the chemotherapy of HIV," *J. of Mathematical Biology,* vol. 35, pp. 775–792, 1997.

[7] S. Kubiak, H. Lehr, R. Levy, T. Moeller, A. Parker and E. Swim, "Modeling control of HIV infection through structured treatment interruptions with recommendations for experimental protocol," *CRSC Technical Report (CRSC-TR01-27),* 2001. (Also at `http://www.math.montana.edu/~parker`.)

[8] J.J. Kutch and P. Gurfil, "Optimal control of HIV infection with a continuously-mutating viral population," *In proceedings of American Control Conference,* pp. 4033–4038, 2002.

[9] U. Ledzewicz and H. Schattler, "On optimal controls for a general mathematical model for chemotherapy of HIV," *In proceedings of American Control Conference,* pp. 3454–3459, 2002.

[10] D. Mayne, J.B. Rawlings, C.V. Rao and P. Scokaert, "Constrained model predictive control: stability and optimality," *Automatica,* vol. 36, pp. 789–814, 2000.

[11] M.A. Nowak and R.M. May, *Virus Dynamics,* Oxford University Press Inc., New York, 2000.

[12] On-line documents at `http://www.cdc.gov/hiv`, Divisions of HIV/AIDS Prevention, Centers for Disease Control and Prevention.

[13] H. Ortega and M. Martin-Landrove, "A model for continuously mutant HIV-1," *In proceedings of 22nd Annual EMBS International Conference,* pp. 1917–1920, Chicago, 2000.

[14] A.S. Perelson and P.W. Nelson, "Mathematical analysis of HIV-1 dynamics in vivo," *SIAM Review,* vol. 41, no. 1, pp. 3–44, 1999.

[15] A. Schwartz, "Theory and implementation of numerical methods based on Runge-Kutta integration for solving optimal control problems," *Ph.D Dissertation,* Univ. of California, Berkeley, 1996.

[16] R.F. Stengel, R. Ghigliazza, N. Kulkarni and O. Laplace, "Optimal control of a viral disease," *In proceedings of American Control Conference,* pp. 3795–3800, 2001.

[17] L.M. Wein, S.A. Zenios and M.A. Nowak, "Dynamic multidrug therapies for HIV: A control theoretic approach," *J. of Theoretical Biology,* vol. 185, pp. 15–29, 1997.

[18] D. Wodarz and M.A. Nowak, "Specific therapy regimes could lead to long-term immunological control of HIV," *Proc. of National Academy Science,* vol. 96, no. 25, pp. 14464–14469, 1999.

[19] D. Wodarz, "Helper-dependent vs. helper-independent CTL responses in HIV infection: implications for drug therapy and resistance," *J. of Theoretical Biology,* vol. 213, pp. 447–459, 2001.

[20] D. Wodarz and M.A. Nowak, "Mathematical models of HIV pathogenesis and treatment," *BioEssays,* vol. 24, pp. 1178–1187, Wiley Periodicals, 2002.

[21] R. Zurakowski and A.R. Teel, "Enhancing immune response to HIV infection using MPC-based treatment scheduling," *In proceedings of American Control Conference,* Denver, 2003.