Utilizing alternate target cells in treating HIV infection through scheduled treatment interruptions

Ryan Zurakowski Andrew R. Teel
Department of Electrical and Computer Engineering
University of California, Santa Barbara
Santa Barbara, CA 93106
ryanz@ece.ucsb.edu teel@ece.ucsb.edu
Dominik Wodarz
Ecology & Evolutionary Biology
University of California, Irvine
Irvine, CA 92697-2525
dwodarz@uci.edu

Abstract— **The dynamics of the human immune response to infection are nonlinear and complex, and analysis of these models can yield surprising, counterintuitive insights. In this paper, we explore one such insight concerning the treatment of HIV infection. In a previous paper, we introduced a model predictive control (MPC) based method of determining treatment schedules that would induce a transition to a state in which the patient's immune system controlled the viral infection without the need for further treatment. In this paper, we show how introducing additional, non HIV-specific target cells (cells which act as hosts for the HIV virus) can yield faster convergence with less transient damage to the patient's immune system.**

I. Introduction

Human Immunodeficiency Virus (HIV) is a virus that infects helper-T cells, a component of the immune response which responds to infection by recruiting the various other cell types that mediate viral clearance. In HIV disease, the attrition of these cells is responsible for the patient's inability to mount a successful immune response and leaves the patient open to infection with a variety of other pathogens. At any given time most helper-T cells are inactive, in a "resting" state that effectively shields them from HIV infection. These cells become "active", and consequently susceptible to HIV infection, when they are stimulated by the presence of the specific pathogen to which they respond. In early HIV disease, the HIV-specific helper-T cells are activated by the HIV stimulus, resulting in a situation in which HIV selectively infects and eliminates those helper-T cells capable of responding to it [1]. However, throughout all stages of infection, HIV-infected patients have abnormally high levels of helper-T activation in the non HIV-specific compartment as well [8], [4]. These additional targets for HIV infection result in higher levels of viremia, quicker progression of the disease, and lower levels of helper-T recovery even while undergoing treatment. It has been suggested, therefore, that peripheral, non HIV-specific helper-T cell activation in HIV patients should be suppressed, thus limiting the number of potential targets for HIV infection and the corresponding increase in viremia [5],[4].

In the absence of an effective immune response to HIV, suppression of peripheral helper-T cell activation makes perfect sense. Suppression of HIV replication by anti-retroviral drugs is not directly dependent on the level of viremia, and increased helper-T cell activation can only lead to higher viremia and quicker depletion of the helper-T population. Higher levels of viremia also increase the risk of emergence of anti-retroviral resistant mutant strains. However, in the case where immune-mediated control is considered, intuition does not lead to an easy conclusion. Immune-mediated control, as described by the Wodarz Nowak model ([13], [11]), is dependent on the level of viremia. Additional targets increase both the rate of new infections and the antigenic stimulus which in turn increases the rate at which the immune response grows. Furthermore, the model suggests the possibility of inducing long-term immune-mediated control of HIV-infection in the absence of continued anti-retroviral drug use through the application of a short-term schedule of treatment interruptions. This process may be helped or hindered by the effects of adding additional non HIV-specific activated helper-T cells as targets for viral infection. In this paper we address the question of whether it is ever useful to introduce additional, non HIV-specific target cells in the context of interruption-based strategies for inducing effective immune control.

The paper is organized as follows: in Section 2, we introduce the Wodarz-Nowak model of HIV infection modified to allow for additional target cells. In Section 3,

Research supported in part by NIH grant number R21 AI057071

we review the MPC-based treatment scheduling algorithm of [14], modified to allow for additional target cells, and we demonstrate the possibility of improved performance utilizing additional target cells. In Section 4, we investigate optimal pairings of interruption length and additional target cell density, and compare best-case results with and without additional target cells. In Section 5, we interpret these results and discuss their implications for HIV therapy.

II. HIV Model with Additional Targets

In a previous paper [14], we applied an MPC-based method to a models of HIV infection developed in [13] and [11] to determine optimal schedules of anti-retroviral therapy that would lead to drug-free immune control of the HIV infection. To test whether it is best to suppress (as much as possible) the activation of non HIV-specific helper-T cells as suggested in [5], we modified the model, introducing a second control term $\mathbf{u_2}$. This term represents the activation of helper-T cells not specifically involved in the immune response against HIV, instead serving only as targets for the virus. In this formulation, ν is a positive constant which scales this effect. This formulation is admittedly simplified and does not address the means by which the population of additional targets would be modulated. Nor does it incorporate the dynamics of this population. It is further simplified by the fact that, in order to use existing code, we restrict the value of $\mathbf{u_2}$ to be either 0 or 1 for each control interval. Unlike our restriction on the values of $\mathbf{u_1}$, this is not well motivated by any restrictions inherent in the system. However, since our goal is simply to determine whether such modulation would ever be useful, the simplified model is appropriate. The model, as used in this paper, is

$$\begin{aligned}
\dot{\mathbf{x}} &= \lambda - d\mathbf{x} - \beta(1-\eta\mathbf{u_1})\mathbf{xy} \\
\dot{\mathbf{y}} &= \beta(1-\eta\mathbf{u_1})(\mathbf{x}+\nu\mathbf{u_2})\mathbf{y} \\
&\quad -a\mathbf{y} - p_1\mathbf{z_1y} - p_2\mathbf{z_2y} \\
\dot{\mathbf{z}}_\mathbf{1} &= c_1\mathbf{z_1y} - b_1\mathbf{z_1} \\
\dot{\mathbf{w}} &= c_2\mathbf{xyw} - c_2q\mathbf{yw} - b_2\mathbf{w} \\
\dot{\mathbf{z}}_\mathbf{2} &= c_2q\mathbf{yw} - h\mathbf{z_2}
\end{aligned} \tag{1}$$

and the states used describe concentrations of: $\mathbf{x}$, healthy, active HIV-specific helper-T lymphocytes (susceptible to HIV infection), $\mathbf{y}$, HIV-infected helper-T lymphocytes, $\mathbf{z}_1$, helper-independent CTL (Cytotoxic Lymphocytes,killer-T cells), $\mathbf{w}$, CTL precursors (memory CTL), and $\mathbf{z}_2$, helper-dependent CTL. The variable $\mathbf{u}_1$ represents the application of HAART therapy, and η is the therapy's effectiveness. We consider the region where all states are positive, since only this region has physical meaning. This region is forward invariant. The model recognizes the dependence of the CTL immune response on the helper-T system, and distinguishes between the helper-T mediated CTL response, which persists even at low antigen levels, and the helper-T independent response, which dies out at low antigen levels. The control input $\mathbf{u_1}$ enters in a manner which recognizes the effect of HAART, which shuts down viral replication and prevents new infection. In order to avoid increased risk of the emergence of drug resistant viral strains, we restrict $\mathbf{u_1}(t)$ to be either 0 (no treatment) or 1 (full treatment). For a more complete description of the states and their interaction, see [13] and [11].

The steady-state behavior of this model has many possible bifurcations due to parameter changes, which are discussed in [11]. In this paper, however, we consider only the case where in the absence of treatment ($\mathbf{u_1} = 0$, $\mathbf{u_2} = 0$) the model has two stable steady states: one describing a progressive infection leading to AIDS and one describing the establishment of a successful immune response.

This model is normalized, i.e., the values of the states have not been adjusted to correspond to measured data. The basic behavior of the model has been observed in experiments on Simian Immunodeficiency Virus (SIV) infection in apes [12], and treatment interruptions in HIV patients have been associated with CTL control of the virus as well [9],[7],[10],[2]. The successful use of the model to induce immune control in SIV and the corresponding anecdotes in human patients lend hope to the possibility of similar success in HIV-infected patients.

III. Improvement at 1-Week Intervals

In [14], we introduced an MPC-based method for determining treatment interruption schedules that would lead to immune-mediated control of HIV infection. For a discrete system of the form

$$\mathbf{X_{k+1}} = f(\mathbf{X_k}, \mathbf{u_k}), \tag{2}$$

with current state $\mathbf{X_k}$, we find a length N sequence $\mathbf{U} = \{\mathbf{u_k}, \mathbf{u_{k+1}}, ..., \mathbf{u_{k+N-1}}\}$ which minimizes a cost function of the form

$$V(\mathbf{X_k}, \mathbf{U}) = \sum_{i=k}^{k+N-1} l(\mathbf{X_i}, \mathbf{u_i}), \tag{3}$$

where l is the stage cost. In this paper, we use the stage cost

$$\begin{aligned}
l(\mathbf{X_i}, \mathbf{u_i}) = \alpha_1(\mathbf{x_i} - x_o)^2 + \alpha_2(\mathbf{w_i} - w_o)^2 \\
+\alpha_3|\mathbf{u_{1i}}| + \alpha_4|\mathbf{u_{2i}}|
\end{aligned} \tag{4}$$

where α_j are positive weighting constants and x_o, w_o, X_o are the steady-state values of their respective states at the desired equilibrium. The resulting optimal control sequence is applied for one sampling period , and at the next sampling period a new optimal control is calculated. This formulation satisfies conditions in [3] and [6] guaranteeing asymptotic stability and robustness (for further details, see [14]). We use a numerical simulator to approximate the discretization of the system.

Using the expanded model described in section 2 and restricting, at one week intervals ($T = 7$), both the application of treatment and the addition of targets ($\mathbf{u}_1$ and $\mathbf{u}_2$) to binary values (i.e. either on or off), we applied the MPC-based algorithm described above to determine whether it

was ever useful to have additional target cells in HIV treatment. The optimization showed that, for certain values of ν, the judicious addition of target cells significantly improved the results of treatment, allowing quicker establishment of immune-mediated virus control and reducing the overall damage to the immune system. An example of these results can be seen in Figure 1. The dotted-line plot shows the growth of the anti-HIV CTL response responding to treatment optimized without allowing for the possibility of additional targets. The solid-line plot shows the growth of the anti-HIV CTL response optimized while allowing the possibility of additional targets at each one-week interval. By adding additional target cells, the treatment scheduler is able to significantly increase the rate at which the immune response to HIV grows, resulting in quicker suppression of the virus. In Figure 2 we see the same comparison, but this time showing the behavior of the healthy HIV-specific helper-T cells. Again, we see that using additional target cells allows the treatment interruption period to be more effective, allowing quicker immune system recovery after treatment ceases. Also, the use of control was significantly reduced, with a total application of HAART of only two weeks, compared to five weeks in the case without using additional targets.

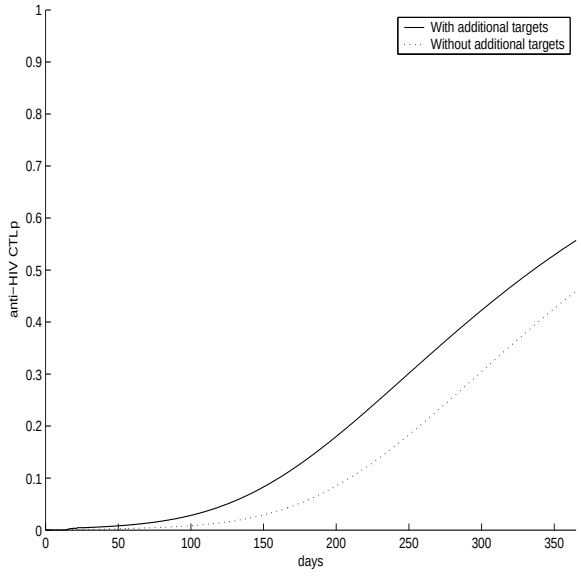

Fig. 1. **Comparison of CTL response** This graph shows the improved growth in the Cytotoxic Lymphocyte response to HIV in the case where additional target cells were introduced. Optimization was carried out over a ten week interval. For the case with additional targets (solid line), $\nu = 0.3636$. For the case without additional targets (dotted line), $\nu = 0$. Initial condition is $\mathbf{x} = 9.9647$, $\mathbf{y} = 0.0176$, $\mathbf{z_1} = 0.00001$, $\mathbf{w} = 0.0551$, $\mathbf{z_2} = 0.0003$.

IV. Optimal Combination of Additional Targets and Interruption Length

The improvement seen in the case of one-week intervals was relatively modest, and the range of additional target

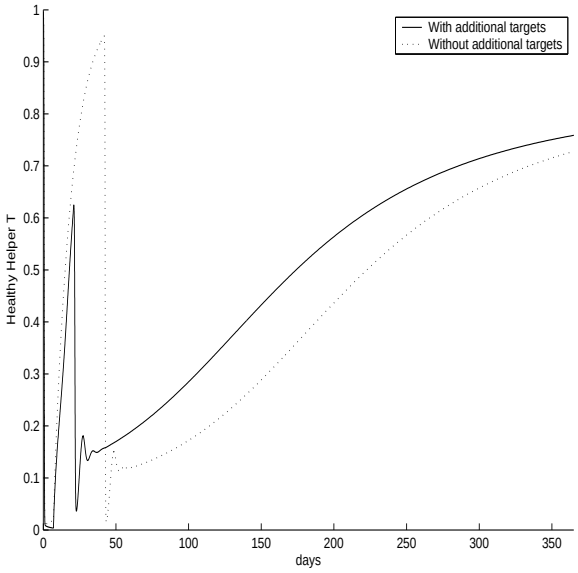

Fig. 2. **Comparison of healthy helper-T populations** This graph shows the improved recovery of the HIV-specific helper-T cell population in the case where additional target cells were introduced. Optimization was carried out over a ten week interval. For the case with additional targets (solid line), $\nu = 0.3636$. For the case without additional targets (dotted line), $\nu = 0$. Initial condition is $\mathbf{x} = 9.9647$, $\mathbf{y} = 0.0176$, $\mathbf{z_1} = 0.00001$, $\mathbf{w} = 0.0551$, $\mathbf{z_2} = 0.0003$.

cells we could add without seeing worse performance was very restricted (approximately up to 5% of the steady-state value for the HIV-specific active helper-T compartment). However, since the addition of target cells effectively accelerates the dynamics, we considered that it may be possible to realize better performance with shorter interruption times. With this in mind, we searched for combinations of treatment interval length and additional target cell density which would maximize the effectiveness of a treatment interruption. We determined the effectiveness of a treatment interruption by the maximum value of $\mathbf{w}$ attained after reinitiating treatment. For an initial condition $\mathbf{X}_{init}$, we found T and ν satisfying

$$\max_{\nu\in[0,10],T\in[0.1,7],t\in[0,50]} f(\phi_\nu(\phi_\nu(\mathbf{X}_{init},0,1,T),1,0,t)) \quad (5)$$

where $\phi_\nu(\mathbf{X},u_1(\cdot),u_2(\cdot),\cdot)$ is the solution of Equation (1) for initial condition $\mathbf{X}$, and $f(\phi_\nu(\mathbf{X},u_1(\cdot),u_2(\cdot),\cdot) = \mathbf{w}(\cdot)$. Because of its computational complexity, we evaluated this maximum only about a single initial condition close to the steady-state of the system when $\mathbf{u_1} = 1, \mathbf{u_2} = 0$. A plot of f as a function of T can be seen in Figure 3. The maximum value of f was achieved when $T = 0.4545$, $\nu = 3.6364$. To do a fair comparison between the case with and without additional target cells, we evaluated the maximum where $\nu = 0$, that is

$$\max_{T\in[0.1,7],t\in[0,50]} f(\phi_0(\phi_0(\mathbf{X}_{init},0,0,T),1,0,t)) \quad (6)$$

A plot of f as a function of T where $\nu = 0$ can be seen in Figure 4. The maximum value of f was achieved when

$T = 0.5909$. .

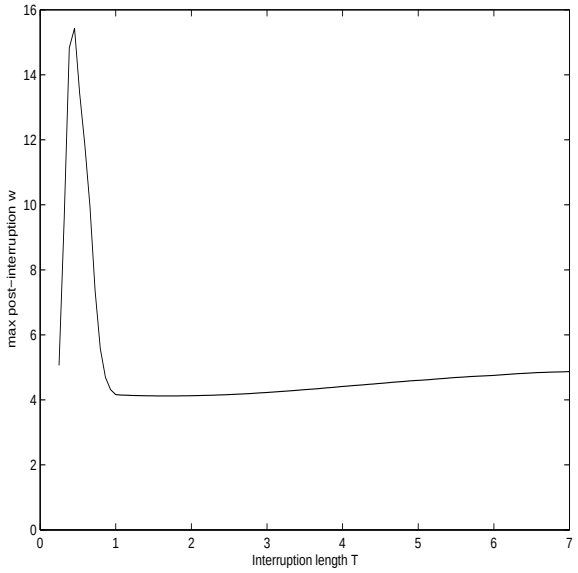

Fig. 3. **Interruption effectiveness as a function of interruption length** Parameters are as in Figure 1 except T and ν. This graph shows the effectiveness of an interruption starting at initial condition $\mathbf{x} = 9.9647$, $\mathbf{y} = 0.0176$, $\mathbf{z_1} = 0.00001$, $\mathbf{w} = 0.0551$, $\mathbf{z_2} = 0.0003$ as a function of interruption length T. For each value of T ν takes the value which maximizes the effectiveness. The maximum occurs at $T = 0.4545$, $\nu = 3.6364$.

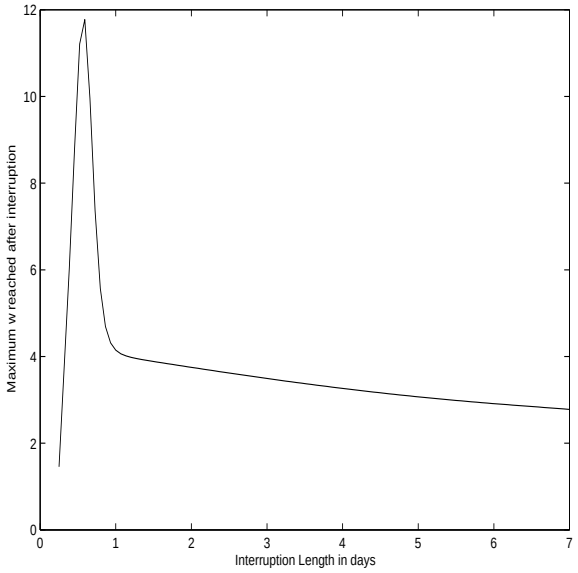

Fig. 4. **Interruption effectiveness as a function of interruption length, no additional targets** Parameters are as in Figure 1 except T and ν. This graph shows the effectiveness of an interruption starting at initial condition $\mathbf{x} = 9.9647$, $\mathbf{y} = 0.0176$, $\mathbf{z_1} = 0.00001$, $\mathbf{w} = 0.0551$, $\mathbf{z_2} = 0.0003$ as a function of interruption length T, $\nu = 0$. The maximum occurs at $T = 0.5909$.

Using the values for T and ν obtained from this optimization, we again applied the MPC algorithm described in Section 3 and compared the results. In both the case with additional targets and without additional targets we saw a significant improvement in performance. In the case where $T = 0.4545$, $\nu = 3.6364$ (Figure 5, solid line), $\mathbf{w}$ was about 50% higher than the value attained at the same time by the algorithm where $T = 7$, $\nu = 0.3636$ (Figure 1, solid line). There was a smaller but still measurable improvement in $\mathbf{w}$ in the $T = 0.5909$, $\nu = 0$ case (Figure 5, dotted line) compared to the $T = 7$, $\nu = 0$ case (Figure 1, dotted line). When comparing the case where $T = 0.4545$, $\nu = 3.6364$ to the case where $T = 0.5909$, $\nu = 0$ as seen in Figure 5, we see the most dramatic difference. Where the use of additional target cells yielded only a modest increase in the growth of $\mathbf{w}$ when $T = 0$, using optimal values for T and ν we see a better-than 50% improvement in the growth of $\mathbf{w}$ after one year.

Although the values of $\mathbf{x}$ are transiently better (as seen in Figure 6) in the case where $T = 0.5909$, $\nu = 0$, the faster growth of $\mathbf{w}$ in the case where $T = 0.4545$, $\nu = 3.6364$ causes this situation to reverse around 150 days after the beginning of treatment. In both cases, allowing shorter treatment interruptions has yielded much better response in the $\mathbf{x}$ compartment, allowing us to achieve immune-mediated control with less transient damage to the healthy, HIV-specific helper-T population. It is worth noting that the optimization of T and ν was focused entirely on $\mathbf{w}$, and that an optimization which contained a trade-off between $\mathbf{x}$ and $\mathbf{w}$ could yield values for T and ν which would allow us to see a clearer improvement in the $\mathbf{x}$ response due to the addition of non HIV-specific target cells.

As expected, the control strategies returned by the scheduling algorithms in the cases allowing shorter treatment interruptions were quite complex, involving many short interruptions in the use of anti-retrovirals ($\mathbf{u_1}$). What was surprising was that the use of additional targets was quite simple; in the case where $T = 0.4545$, $\nu = 3.6364$, the MPC scheduling algorithm set $\mathbf{u_2} = 1$ for the duration, apart from one approximately one day interval near the beginning of treatment. Omitting this brief interruption resulted in no discernible change in the output, suggesting that modulation of the concentration of additional target cells may not be necessary to realize their benefits in inducing immune control.

V. Discussion

As modeled in this paper, the addition of non HIV-specific activated helper-T cells as targets for HIV infection can, under certain circumstances, dramatically increase the performance of an algorithm designed to induce effective, immune-mediated control of HIV through a pattern of interruptions in anti-retroviral therapy. From any static understanding of the immune response, this result is nonintuitive. However, an explanation of this behavior can be made by closely analyzing the dynamics of the immune system. The immune response to HIV is dependent on the development of a large population of HIV-specific CTL precursors ($\mathbf{w}$). The growth of this compartment is controlled by a term,

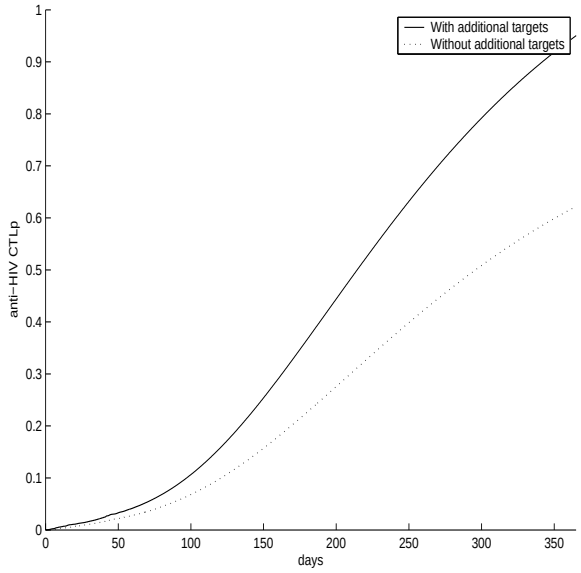

Fig. 5. **Optimal treatment length CTL comparison** Parameters are as in Figure 1 except T and ν. The solid-line plot represents $\mathbf{w}$ when $T = 0.4545$, $\nu = 3.6364$, the dotted-line plot when $T = 0.5909$, $\nu = 0$.

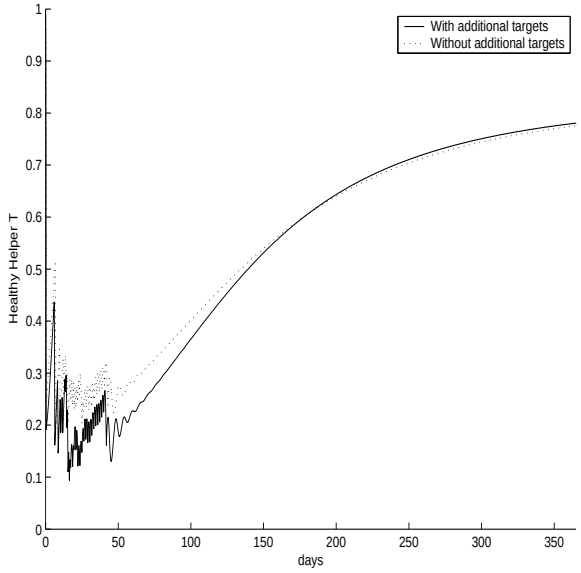

Fig. 6. **Optimal treatment length helper-T comparison** Parameters are as in Figure 1 except T and ν. The solid-line plot represents $\mathbf{x}$ when $T = 0.4545$, $\nu = 3.6364$, the dotted-line plot when $T = 0.5909$, $\nu = 0$.

$c_2\mathbf{xyw}$, that depends on the ratio of healthy HIV-specific helper-T cells and HIV-infected helper-T cells. In order to increase the growth of this compartment, we must increase either $\mathbf{x}$ or $\mathbf{y}$. Adding additional target cells by activating non-HIV specific helper-T cells allowed us to increase $\mathbf{y}$ without decreasing $\mathbf{x}$. Of course, increasing $\mathbf{y}$ also increases the rate $\beta\mathbf{xy}$ at which healthy HIV-specific helper-T cells are infected, but by timing our interruptions and the activation of additional targets correctly, we realized better overall performance. This performance is maximized with treatment intervals of approximately one-half day, and additional target cells added at a concentration approximately 35% of the steady-state level of HIV-specific helper-T cells. In a simulated comparison of "best-case" quasi-optimal parameter sets, the addition of target cells resulted in an increase of better than 50% in the concentration of HIV-specific CTL-precursor cells after an MPC-based treatment algorithm was applied for 1 year. Realizing these benefits did not require complicated modulation of the concentration of additional targets: simply adding the additional targets for the duration of the treatment was sufficient.

This result is interesting, but there are a number of reasons why it would be difficult to implement as a viable treatment option. We don't know how to manipulate the number of non HIV-specific helper-T cells with any degree of accuracy. We didn't need to modulate the level during the course of treatment, and we saw improvement for relatively high levels of additional targets. These two facts suggest that the abnormally high levels of non HIV-specific helper-T cell activation associated with HIV infection may not be a problem for the use of treatment interruptions to induce effective immune control. However, any schedules which depend on modulating the level of non HIV-specific target cells are currently infeasible. Also, the short treatment interruptions for which we saw the most benefit from additional target cells violate certain assumptions used in the model: it is known that there is a short delay between the infection of a cell with HIV and the beginning of virus production, and the clearance of anti-retroviral drugs is not instantaneous upon ceasing treatment. For interruption lengths of one week or greater, it is a reasonable assumption that the aggregate effect of these relatively fast phenomena is negligible, but for interruption lengths of less than a day, this assumption is almost certainly not valid. These phenomena would have to be included in the model to make valid deductions about such short treatment interruptions. Finally, the model we use assumes that any abnormal helper-T cell death is due either to direct viral cytotoxicity or to CTL-mediated killing of infected cells. The authors of [5] and [4] suggest that the depletion of the T-cell population may be due directly to the abnormally high levels of activation, through either accelerated senescence or abnormal apoptosis. If this is the case, it is not addressed in our model and any risks or benefits would have to be evaluated through other means.

It is unlikely that intentionally manipulating the level of nonspecific target cells will be a useful treatment option in the near future. The consequences of T-cell activation are complex enough that exploiting their potential benefits would require further work, both in modeling and experiment. But the possibility of using an increase in non-HIV specific target cells to benefit a patient is one that is highly counterintuitive. This suggests the potential for equally counterintuitive (but more useful) treatment options to be found through the mathematical modeling of HIV infection.

REFERENCES

[1] D. Douek, J. Brenchley, M. Betts, D. Ambrozak, B. Hill, Y. Okamoto, J. Casazza, J. Kuruppu, K. Kunstman, S. Wollnsky, Z. Grossman, M. Dybul, A. Oxenlus, D. Price, M. Conners, and R. Koup. HIV preferentially infects HIV-specific CD4+ T cells. *Nature*, 417, 2002.

[2] F. Garcia, M. Plana, G. Ortiz, S. Bonhoeffer, A. Soriano, C. Vidal, A. Cruceta, M. Arnedo, C. Gil, G. Pantaleo, T. Pumarola, T. Gallart, D. Nixon, J. Miro, and J. Gatell. The virological and immunological consequences of structured treatment interruptions in chronic HIV-1 infection. *AIDS*, 15(9), 2001.

[3] G. Grimm, M. J. Messina, A. R. Teel, and S. E. Tuna. Model predictive control when a local control lyapunov function is not available. In *Proceedings of the 2003 American Control Conference*, June 2003.

[4] M. D. Hazenberg, S. A. Otto, B. H. B. van Benthem, M. T. L. Roos, R. A. Coutinho, J. M. A. Lange, D. Hamann, M. Prins, and F. Miedema. Persistent immune activation in HIV-1 infection is associated with progression to AIDS. *AIDS*, 17(13):1881–8, Sep 2003.

[5] P. Hunt, J. Martin, E. Sinclair, B. Bredt, E. Hagos, H. Lampiris, and S. Deeks. T cell activation is associated with lower CD4+ T cell gains in human immunodeficiency virus infected patients with sustained viral suppression during antiretroviral therapy. *Journal of Infectious Diseases*, 187(10), 2003.

[6] C. M. Kellett and A. R. Teel. On robustness of stability and lyapunov functions for discontinuous difference equations. In *Proceedings of the 41st IEEE Control and Decision Conference*, Dec. 2002.

[7] G. Ortiz, D. Nixon, A. Trkola, J. Binley, X. Jin, S. Bonhoeffer, P. Kuebler, S. Donahoe, M.-A. Demoitie, W. Kakimoto, T. Ketas, B. Clas, J. Heymann, L. Zhang, Y. Cao, A. Hurley, J. Moore, D. Ho, and M. Markowitz. HIV-1-specific immune responses in subjects who temporarily contain virus replication after discontinuation of highly active antiretroviral therapy. *Journal of Clinical Immunology*, 104, 1999.

[8] Papagno, Spina, Marchant, Salio, Rufer, Little, Dong, Chesney, Waters, Easterbrook, Dunbar, Shepherd, Cerundolo, Emery, Griffiths, Conlon, McMichael, Richman, Rowland-Jones, and Appay. Immune Activation and CD8(+) T-Cell Differentiation towards Senescence in HIV-1 Infection. *PLoS Biol*, 2(2):E20, Feb 2004.

[9] E. Papasavvas, G. Ortiz, R. Gross, J. Sun, E. C. Moore, J. Heymann, M. Moonis, J. K. Sandberg, L. A. Drohan, B. Gallagher, J. Shull, D. F. Nixon, J. Kostman, and L. Montaner. Enhancement of human immunodeficiency virus type 1-specific CD4 and CD8 T cell responses in chronically infected patients after long-term viral suppression. *Journal of Infectious Diseases*, 182, 2000.

[10] E. Rosenberg, M. Altfeld, S. Poon, M. Phillips, B. Wilkes, R. Eldridge, G. Robbins, R. D'Aquila, P. Goulder, and B. Walker. Immune control of HIV-1 after early treatment of acute infection. *Nature*, 407, 2000.

[11] D. Wodarz. Helper-dependent vs. helper-independent CTL responses in HIV infection. *Journal of Theoretical Biology*, 213, 2001.

[12] D. Wodarz, R. Arnaout, M. Nowak, and J. Lifson. Transient antiretroviral treatment during acute simian immunodeficiency virus infection facilitates long-term control of the virus. *Philosophical Transactions of the Royal Society of London in Biology*, 355, 2000.

[13] D. Wodarz and M. A. Nowak. Specific therapy regimes could lead to long-term immunological control of HIV. *Proceedings of the National Academy of Sciences*, 96(25), 1999.

[14] R. Zurakowski and A. R. Teel. Enhancing immune response to HIV infection using MPC-based treatment scheduling. In *Proceedings of the 2003 American Control Conference*, June 2003.