Receding Horizon Control of HIV

John David [*†] Hien Tran [†] H.T. Banks [†]

December 22, 2009

Abstract

This paper describes a model of the immunologic response of the human immunodeficiency virus (HIV) in individuals. It then illustrates how a Receding Horizon Control (RHC) methodology can be used to drive the system to a stable equilibrium in which a strong immune response controls the viral load in the absence of drug treatment. We also illustrate how this feedback methodology can overcome unplanned treatment interruptions, inaccurate or incomplete data and imperfect model specification. We consider how ideas from stochastic estimation can be used in conjunction with RHC to create a robust treatment methodology.

1 Introduction

There have been various studies illustrating the application of optimal control methodology to HIV drug treatment. A large number of methodologies for the control of HIV are based on open-loop techniques. This involves pre-computing a control over the entire treatment time interval for a given model and initial conditions, without later revising the treatment based on newly-available information. An example of these works can be seen in [19, 16, 3, 10, 13, 14]. This technique may be inadequate in many cases for a number of reasons. First, some unmodelled effects can disturb the system, thus rendering the treatment schedule ineffective or, even worse, detrimental. Secondly, if a patient misses a dose, the open loop control is no longer optimal or even necessarily beneficial. This open-loop methodology also does not take advantage of the measurements collected from an individual. For these reasons, we consider in this paper feedback methodology, where the control depends on the current state of the system, in the context of the control of an HIV infection.

Some feedback techniques may be inappropriate or non-applicable to the problem of HIV control. For example, many feedback techniques are based on linear plants or models. It is difficult to accurately model HIV dynamics with a linear model, especially as the time scale expands and the number of modelled biological mechanisms increases. See [22] for an excellent review of how the complexity of these models has increased as

[*]Department of Mathematics and Computer Science, The College of Wooster, Wooster OH 44691, jdavid@wooster.edu

[†]Department of Mathematics and Center for Research in Scientific Computation, Box 8205, North Carolina State University, Raleigh North Carolina 27695-8205, tran@ncsu.edu, htbanks@ncsu.edu

understanding of disease dynamics has progressed. Techniques based around solutions of Ricatti equations may be augmented to solve nonlinear control problems; however, these techniques require that the model be of a certain form and can require frequent state knowledge or estimation. For an example of this sort of technique applied to HIV control see [6]. A potential solution involves using Bellman's principle of optimality, however this involves solving a nonlinear partial differential equation which, in general, is computationally challenging.

Receding Horizon Control (RHC) seeks to gain the benefits of a feedback control while utilizing the computational simplicity of a calculus of variations approach. RHC involves solving a finite horizon open-loop optimal control problem on-line at each sampling instant. There are several design considerations in using this methodology including the process model, the cost function to be minimized, the sampling period, the control horizon, and the method by which the state is obtained at each sampling instant. The method by which the state is obtained at each sampling instant is of special concern as it generally involves some combination of model prediction and noisy measurements. For this problem, we will employ stochastic estimation; in particular, the Extended Kalman Filter (EKF). An excellent survey paper related to the theory behind RHC is [21]. Another source for a description of the multitude of industrial applications of this technique is [23]. This technique has been used several times in the context of HIV control [24, 26].

The contribution of the present work is that it considers a complex predictive model of HIV dynamics, which has been validated with clinical data (see [1]), that includes multi-drug therapy, multiple target cells and a compartment accounting for the virus specific immune response. We use RHC in conjunction with stochastic estimation to control HIV infection under poor drug adherence, noisy measurements, and inaccurate model parameters.

The organization of this paper is as follows. In Section 2 we discuss the model of HIV infection studied here. We then describe the specific RHC methodology that we have considered in Section 3. Numerical results for the RHC implementation are presented in Section 4. Situations in which treatments are missed and a state estimator is employed are treated in Sections 5 and 6. The numerical solutions to these two combined problems are illustrated in Section 7. In Section 8 we demonstrate how RHC methodology can be used in situations in which model parameters are allowed to vary in some random manner.

2 HIV model

A great deal of effort has gone into modeling the physiologic and immunologic response of HIV in individuals. For excellent reviews of the various types of modeling attempts see, e.g. [8, 9, 22]. In our attempt to model the physiologic and immunologic response of the HIV in individuals, we will consider a variation of the model proposed in [8]. It accounts for a variation of drug efficacy based on target cell type. Note that the differentiation between target cells is based on cell type as opposed to cell location or longevity. A second chief component of this model is the modeling of the bodies HIV specific immune response. This aspect of the model is based on a Michaelis-Menten nonlinearity saturation as proposed by [7]. This specific model has been studied in [3, 4, 1, 6].

This model captures some of the chief characteristics of HIV infection that we would like to consider. In addition to modeling the body's HIV specific immune response and differentiating between types of target cells, this model also accounts for multi-drug therapy, both reverse transcriptase inhibitors (RTIs) and protease inhibitors (PIs). RTIs interfere with the viral RNA-to-DNA synthesis thus blocking the introduction of viral DNA into the cells genetic information. Without this step the cells will not produce new virus. PIs attack the virus at the stage at which infected cells assemble new viral cells thus producing non-infectious viral cells. An important physiological aspect of this model is its ability to reproduce a low non-zero viral load in the presence of multi-drug therapy. The model also possesses multiple stable equilibria including a viral dominant equilibrium and an immune response dominant equilibrium.

The dynamics of our HIV model are described by the system of nonlinear differential equations

$$\begin{aligned}
\dot{T}_1 &= \lambda_1 - d_1 T_1 - (1-\epsilon_1) k_1 V_I T_1 \\
\dot{T}_2 &= \lambda_2 - d_2 T_2 - (1-f\epsilon_1) k_2 V_I T_2 \\
\dot{T}_1^* &= (1-\epsilon_1) k_1 V_I T_1 - \delta T_1^* - m_1 E T_1^* \\
\dot{T}_2^* &= (1-f\epsilon_1) k_2 V_I T_2 - \delta T_2^* - m_2 E T_2^* \\
\dot{V}_I &= (1-\epsilon_2) N_T \delta (T_1^* + T_2^*) - [c + (1-\epsilon_1)\rho_1 k_1 T_1 + (1-f\epsilon_1)\rho_2 k_2 T_2] V_I \\
\dot{V}_{NI} &= \epsilon_2 N_T \delta (T_1^* + T_2^*) - c V_N I \\
\dot{E} &= \lambda_E + b_E \frac{T_1^* + T_2^*}{T_1^* + T_2^* + K_b} E - d_E \frac{T_1^* + T_2^*}{T_1^* + T_2^* + K_d} E - \delta_E E.
\end{aligned} \tag{1}$$

The state variables are T_1, the uninfected CD4+ T-cells; T_2, the uninfected target cells of a second kind; T_1^*, the infected T-cells; T_2^*, the infected target cells of a second kind; V_I, the infectious virus; V_{NI}, the non-infectious virus; and E, the immune effectors. We do not give specific definitions of target cells of second kind and immune effectors. They could, for example be related to macrophages and cytotoxic T-lymphocytes, respectively. Table 1 contains values for model parameters. The predictive capabilities of this model were demonstrated in [1] with clinical data from over 100 individuals from the period between 1996-2004. See [2, 3] for a further description of the model and model parameters.

3 RHC Methodology

As stated before, RHC seeks to gain the benefits of a feedback control while utilizing the computational simplicity of a calculus of variations approach. Thus, there are several basic components to the methodology. First, there are the model equations which we use to predict system behavior and compute the optimal control. There is the estimation of the current state, which can be done with a combination of measured data and predicted output. We will use the methodology described in [12], however there are a variety of techniques to complete this step. There is also the specific optimal control problem involving a cost function one seeks to minimize. Finally, there is the methodology by which the control horizon, the period over which each optimal control

parameter	value	unit	parameter	value	unit
λ_1	10.0	$\frac{cells}{mm^3 \cdot day}$	λ_2	31.98e-3	$\frac{cells}{mm^3 \cdot day}$
d_1	0.01	$\frac{1}{day}$	d_2	0.01	$\frac{1}{day}$
k_1	8e-4	$\frac{mm^3}{virions \cdot day}$	k_2	0.1	$\frac{mm^3}{virions \cdot day}$
m_1	0.01	$\frac{mm^3}{cells \cdot day}$	m_2	0.01	$\frac{mm^3}{cells \cdot day}$
ρ_1	1	$\frac{virions}{cells}$	ρ_2	1	$\frac{virions}{cells}$
δ	0.7	$\frac{1}{day}$	c	13	$\frac{1}{day}$
f	0.34	-	N_T	100	$\frac{virions}{cells}$
λ_E	1e-3	$\frac{cells}{mm^3 \cdot day}$	δ_E	0.1	$\frac{1}{day}$
b_E	0.3	$\frac{1}{day}$	d_E	0.25	$\frac{1}{day}$
K_b	0.1	$\frac{cells}{mm^3}$	K_d	0.5	$\frac{cells}{mm^3}$

Table 1: Values of parameters in the HIV model.

problem is solved, and the period over which this control is used, are selected. Figure 1 is a schematic that illustrates this general process.

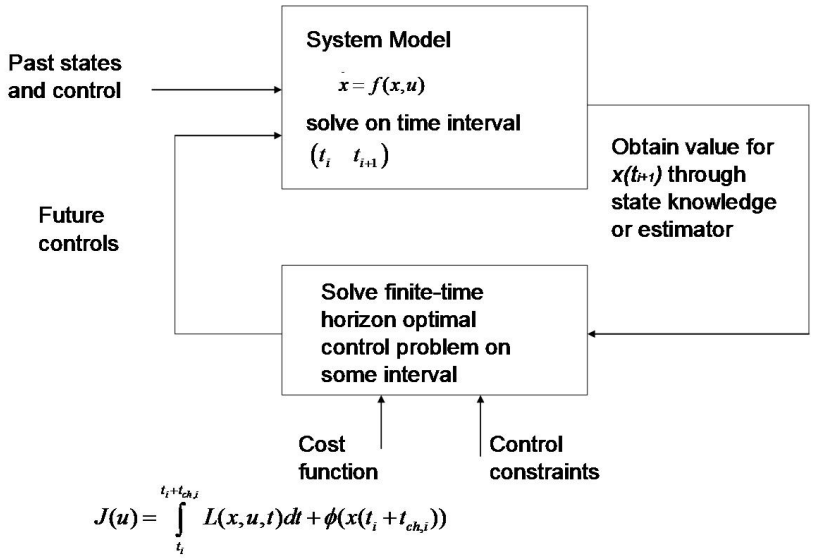

Figure 1: Schematic diagram of RHC methodology.

To formalize this methodology we will present the following mathematical formulation. We will use notation for the RHC methodology that is similar to that presented in [15]. Let $[t_i, t_{i+1}]$ be a sequence of time intervals. Let $t_{ch,i}$ such that $t_{ch,i} \geq t_{i+1} - t_i$ be the *control horizon*. Consider the sequence of problems, O_i,

$$\min_u J(u) = \int_{t_i}^{t_i+t_{ch,i}} L(x,u,t)dt + \phi(x(t_i + t_{ch,i})), \tag{2}$$

subject to

$$\dot{x} = f(x,u),\ x(t_i) = x_i, \tag{3}$$

on the interval $[t_i, t_{i+1}]$. In order to be biologically realistic we will also enforce bounds on the control, $u_{min} \leq u(t) \leq u_{max}$. Note that for this problem the control is the efficacy of the drug dose and thus we will write $u = (\epsilon_1, \epsilon_2)$.

The solution to each optimal control problem O_i can be analyzed and solved in the manner presented in detail in [12] and summarized here. We used a quasi-Newton method BFGS. The code used in this problem can be obtained from the web site `http://www.math.ncsu.edu/~ctk`. For more information on this optimization algorithm see [18].

The RHC methodology can be summarized as follows:

1. Given initial condition $x(t_i)$ solve the control problem O_i on the time interval $[t_i, t_i + t_{ch,i}]$.
2. Use the control as it is defined on the interval $[t_i, t_{i+1}]$ to determine the trajectory on the same time interval.
3. Use observations or an estimator to determine $x(t_{i+1})$.
4. Repeat the process over starting at the next time interval t_{i+1}.

This process is illustrated in Figure 2. This means that O_0 was solved by minimizing the integral from t_0 to T. However the control was only used from t_0 to t_1. At this point a new control was computed by solving O_1 on the interval t_1 to $t_1 + T$. This control was then used from t_1 to t_2. Note that neither the time intervals nor the control horizon need to be constant.

Our numerical routine to solve each control problem O_i, is an iterative method thus requiring an initial iterate. The initial iterate for O_0 was $u_0 = 0.5(\epsilon_{1,max}, \epsilon_{2,max})$. The initial iterate for each subsequent O_i was the optimal solution to O_{i-1}. Two heuristic approaches were used to facilitate numerical computations. One approach was how to choose the initial iterate when the immune response E was stimulated above 200 cells/mm^3. In this case, the initial iterate was decreased by a factor of 10 as this facilitated convergence to a low dose control that stimulated further immune response. The other was in the case that the ode solver did not converge for a candidate control vector. In this case u_p was replaced with $0.99u_p$ and the algorithm proceeded. The optimal control estimation is discussed further in [12] and the full optimality conditions are given in [11].

4 Numerical Results: Perfect State Knowledge

For our study we employed the following cost functional

$$J(\epsilon_1, \epsilon_2) = \int_{t_i}^{t_i + t_{ch,i}} \left[Q_e V(t) + R_1 \epsilon_1^2(t) + R_2 \epsilon_2^2(t) - SE(t) \right] dt + \theta (E(t_i + t_{ch,i}) - 353.108)^2, \quad (4)$$

where Q_e, R_1, R_2, S and θ are weighting coefficients. The maximum values for the drug efficacies for the following simulations are $\epsilon_{1,max} = 0.7$ and $\epsilon_{2,max} = 0.3$. The results of [11] suggest that each optimal control problem O_i has a unique solution with the addition of one condition, that ϕ in (2), the weighting on the final state, is continuous. As ϕ in (4) is a quadratic function of E this condition holds.

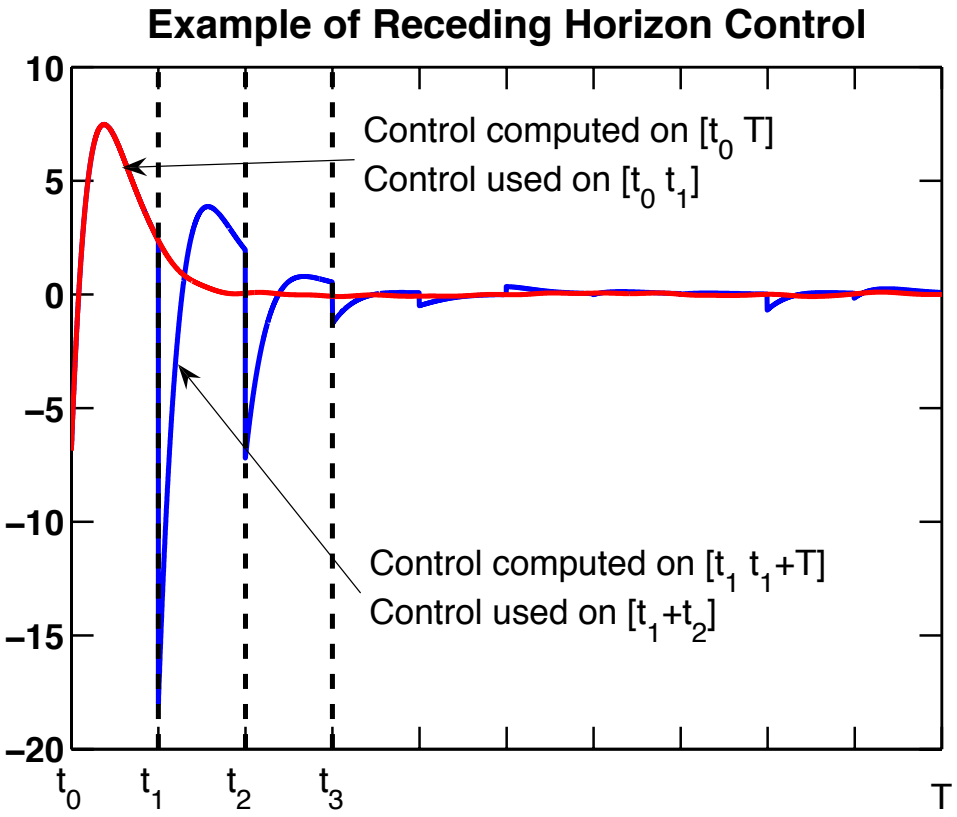

Figure 2: Example of RHC.

Parameter	Value
t_i	$50i$ days
$t_{ch,i}$	50 days
Q_e	10^{-9}
R_1	10^{-9}
R_2	10^{-9}
S	10^{-5}
θ	5^{-4}

Table 2: Parameters for RHC problem.

Table 2 contains the values for the optimal control problem. The small values of weighting coefficients are used so that the gradient values are of approximately the same magnitude as the control bounds thus improving the efficiency of the optimization.

In this section we will assume that we have perfect knowledge of the state at each point in which a new control is computed. This is not physically realistic, but we will use the EKF to address the problem of lack of state knowledge in Section 6. The initial condition for this problem is a healthy individual with a small amount of free virus introduced to their system. This can be termed acute infection. The specific values for this initial condition are

$$\begin{aligned} T_1 &= 1000\, cells/mm^3,\ T_2 = 3.198\, cell/mm^3, \\ T_1^* &= 10^{-5}\, cells/mm^3,\ T_2^* = 10^{-5}\, cell/mm^3, \\ V_I &= 0.001\, virions/mm^3,\ V_{NI} = 10^{-5}\, virions/mm^3, \\ E &= 0.01\, cells/mm^3. \end{aligned} \tag{5}$$

Our goal is to transfer this infected individual to a *healthy immune response dominated equilibrium*. The specific values for this state are

$$\begin{aligned} T_1 &= 967.839\, cells/mm^3,\ T_2 = 0.621\, cell/mm^3, \\ T_1^* &= 0.076\, cells/mm^3,\ T_2^* = 0.006\, cell/mm^3, \\ V_I &= 0.415\, virions/mm^3,\ V_{NI} = 0\, virions/mm^3, \\ E &= 353.108\, cells/mm^3. \end{aligned} \tag{6}$$

This equilibrium has been studied in [2, 3]. We will merely note here that this is a stable equilibrium in the presence of no treatment. We have investigated three cases for how often the control is allowed to change. The treatments are allowed to vary daily, every five days, and every ten days. Note that as the model and control are continuous representations (as opposed to a discrete model) the control values used in between specified treatment values are a linear interpolation.

Figures 3-6 depict control and model response for each of these cases. The non-infectious virus V_{NI} has been omitted in the figures as that compartment has no direct effect on any of the other compartments. Figure 3 depicts each one of the state variables as a function of time in days on the left and the corresponding treatment protocol on the right. Note that the healthy cells decline during the treatment interruptions and

the free virus and infected cells rebound. The state variable E is the best indicator of the progression toward immune system dominance and can be seen in the lower right box of all the state variables. In Figure 4 we plot in the phase plane of the log of the immune effector versus the log of the virus. Note that in this sort of plot time is an implicit variable with only the two state variables being shown. This plot begins at acute infection and thus relatively low levels of E and V. We can see that the viral load undergoes one increase during infection and four increases during treatment interruptions. It is interesting to see that E decreases during most of the treatment interruption only to increase toward the end of the interruption. Observe that each one of these treatment protocols has treatment interruptions under which the viral load and infected cells rebound, while the number of healthy cells drops. Also note that the frequency with which the control is allowed to vary has a large effect on how long it takes the immune response to be stimulated and for the individual to be transferred to the healthy steady state. When the control is allowed to vary daily the control is largely turned off after 400 days, whereas it is not turned off until approximately 600 days when the control varies every 5 days and 900 days are required when the control is allowed to vary every 10 days. The phase plots in Figure 4 also illustrate the difference in system response when the control is only allowed to vary every 10 days. Note how more interruptions are required and there is a good deal more time where the immune response is declining.

We would also like to determine whether this methodology can be used to transfer a patient from the unhealthy equilibrium to the healthy equilibrium. By unhealthy equilibrium we mean the state

$$\begin{aligned} T_1 &= 163.573\, cells/mm^3,\ T_2 = 0.005\, cell/mm^3, \\ T_1^* &= 11.945\, cells/mm^3,\ T_2^* = 0.046\, cell/mm^3, \\ V_I &= 63.919\, virions/mm^3,\ V_{NI} = 0\, virions/mm^3, \\ E &= 0.024\, cells/mm^3. \end{aligned} \tag{7}$$

This equilibrium represents a physical state in which the immune response has largely been destroyed and a large quantity of infected cells and free virus is present in an individual. Note that this equilibrium is shown to be stable in [2, 3]. Figures 7, 9 and 10 illustrate the model simulations and controls when the initial conditions are given by the unhealthy steady state. Figure 8 depicts the phase plot of the log of the immune effector as a function of the log of the virus. In each one of these situations, even when the control was only allowed to vary every 10 days, the RHC methodology was able to transfer the system to the immune effector dominated equilibrium.

5 Unscheduled Treatment Interruption

A natural question to pose is what happens if (or when) a patient goes off medication? This can occur because of new infections, drug side effects, lack of adherence, or a number of other reasons. Open loop methodology has no way to take this into account and will most likely prove insufficient. Feedback controllers can take this treatment lapse into account and adjust the controller. We will examine how RHC performs when a patient disregards the proposed treatment schedule and goes off treatment. For the following simulations a patient went off treatment for 10 days after 300 days

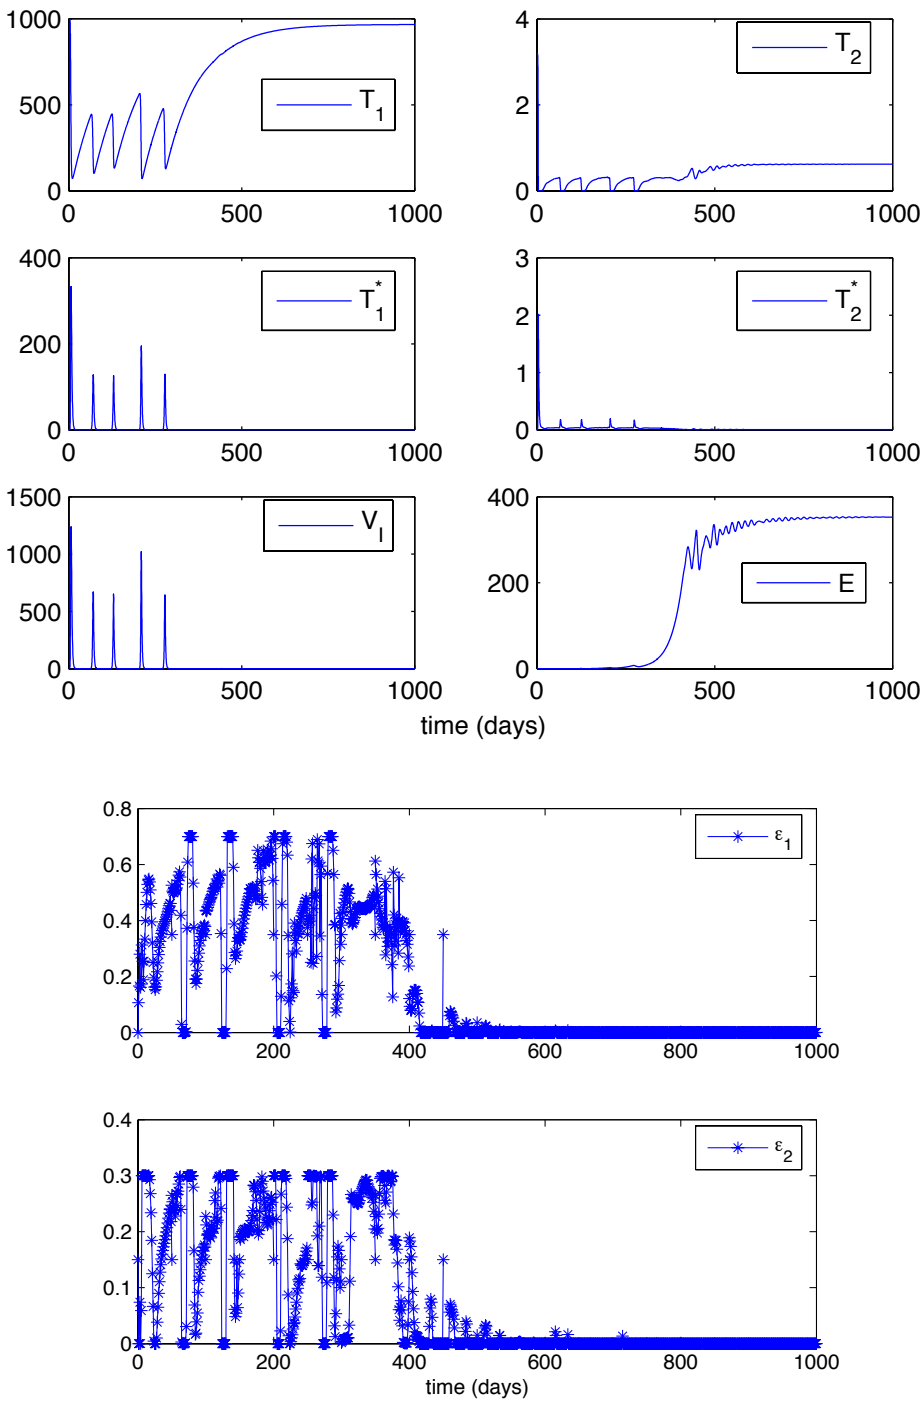

Figure 3: Model response and control when the control is allowed to vary daily. Initial conditions given by (5).

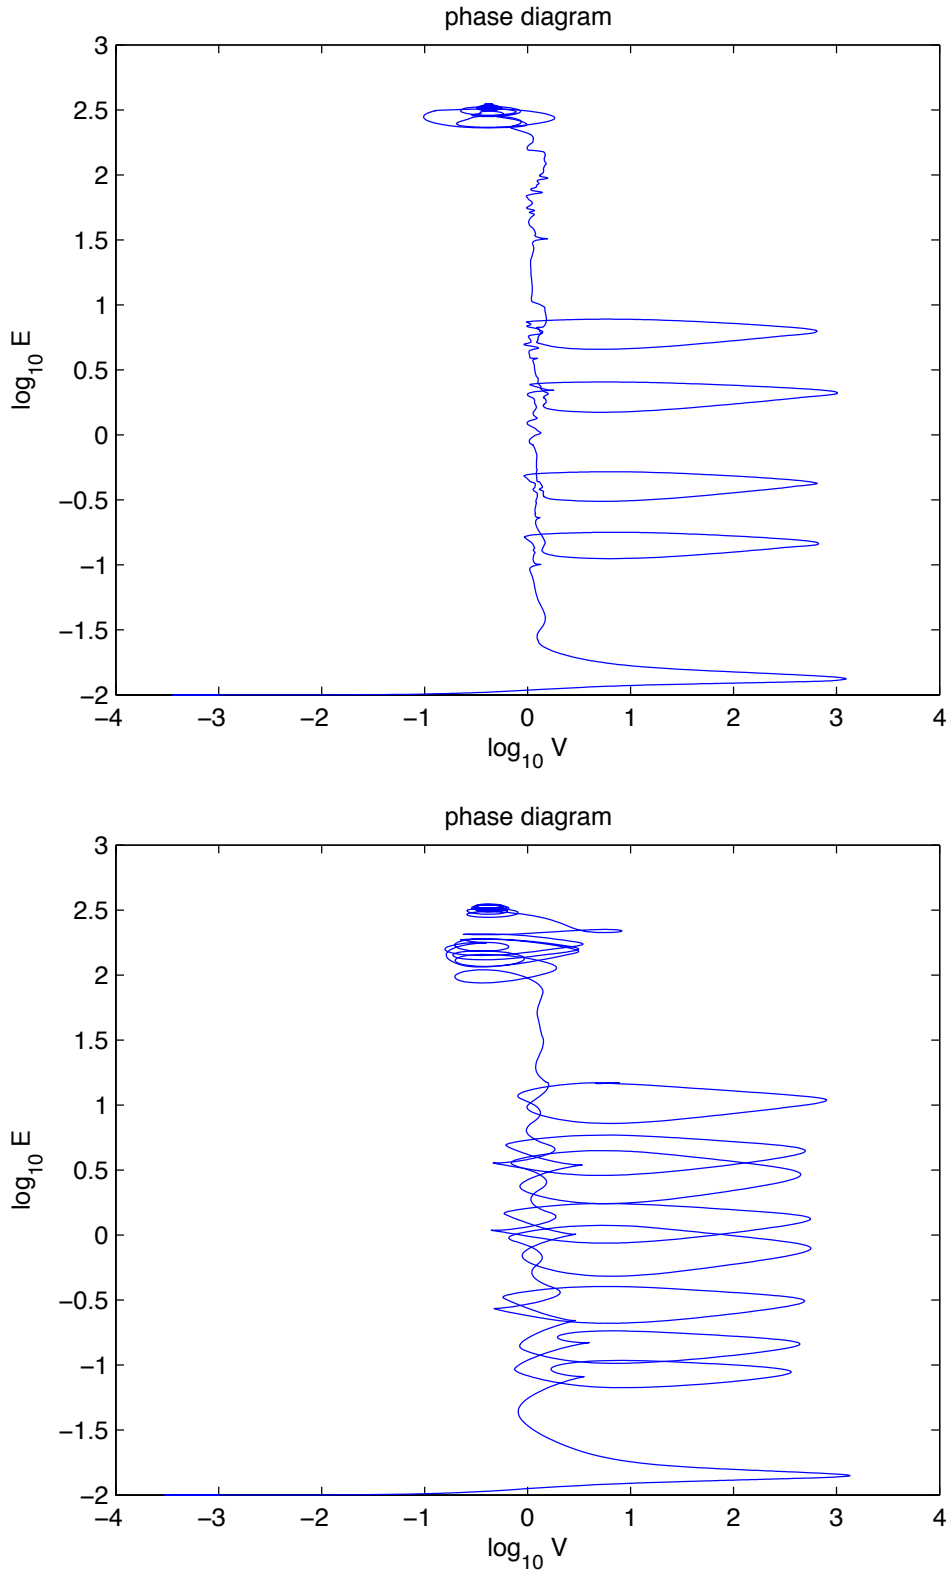

Figure 4: Phase plot of log of immune effectors vs. log of virus when control is allowed to vary daily on the top. The results for when control was allowed to vary every 10 days are on the bottom.

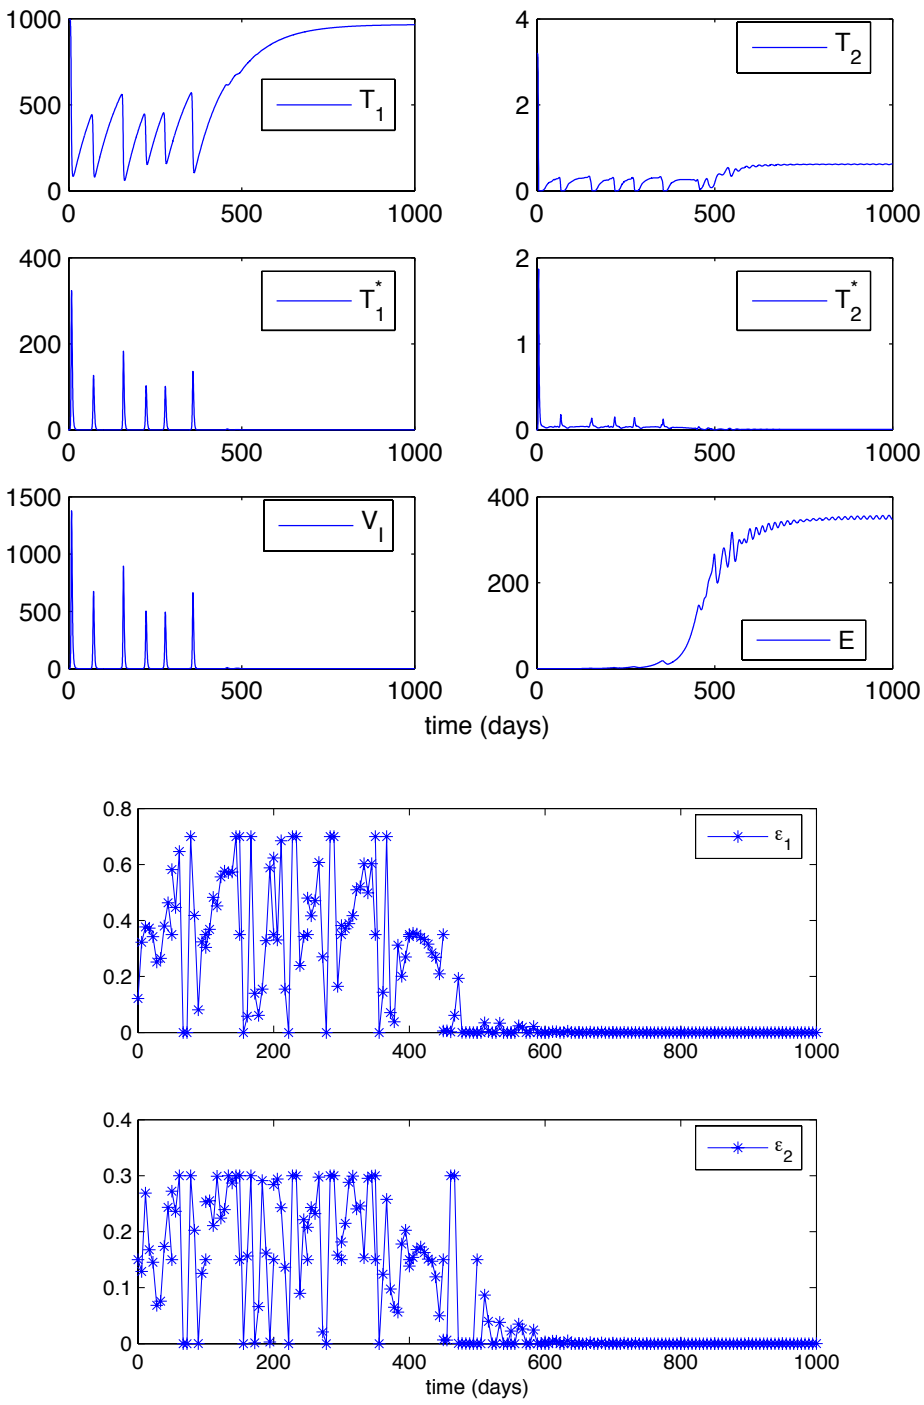

Figure 5: Model response and control when control is allowed to vary every 5 days. Initial conditions given by (5).

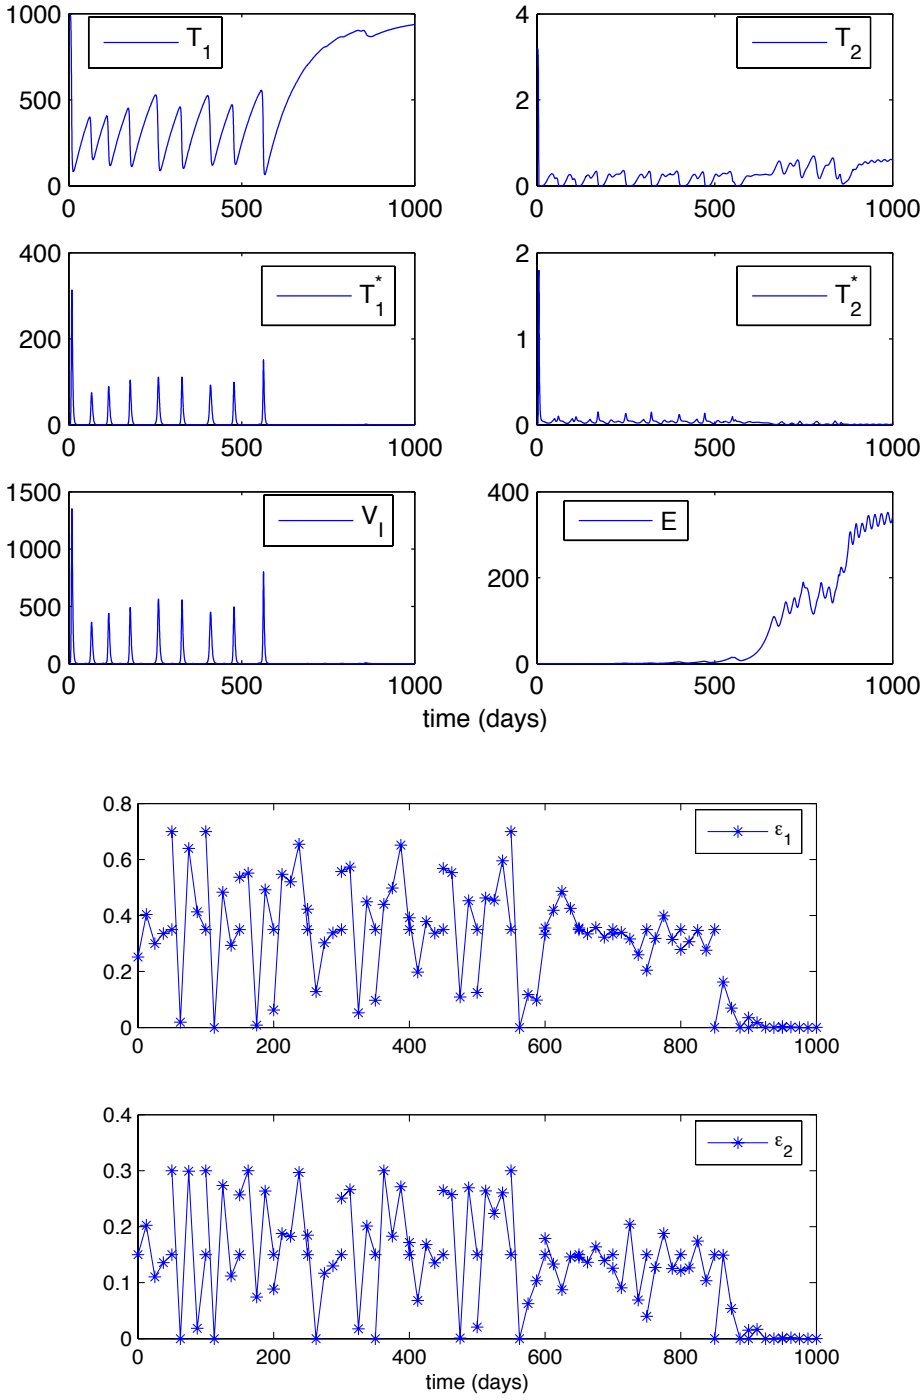

Figure 6: Model response and control when control is allowed to vary every 10 days. Initial conditions given by (5).

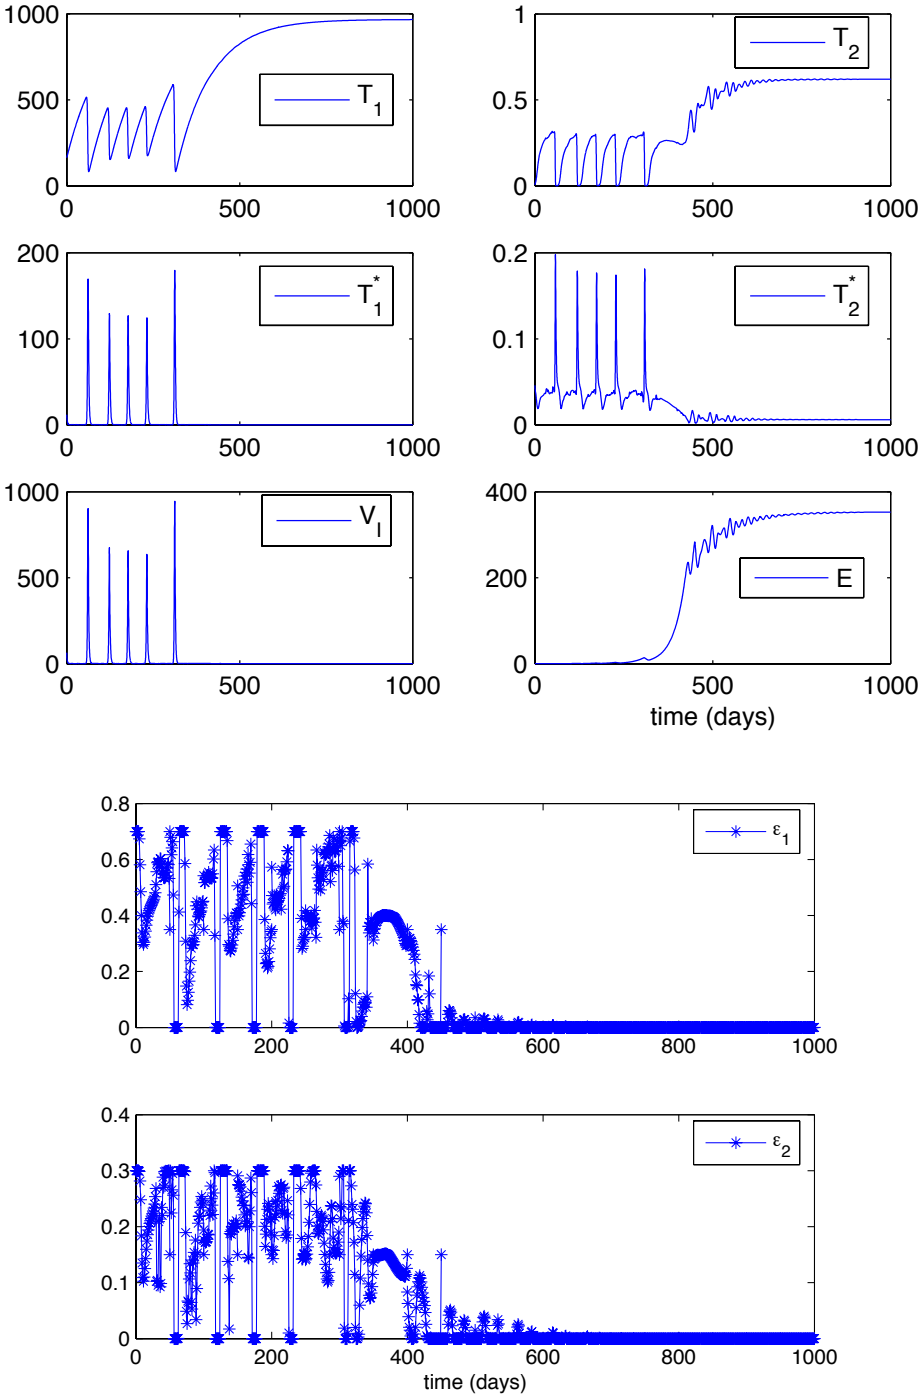

Figure 7: Model response and control when control is allowed to vary daily. Initial conditions given by unhealthy steady state (7).

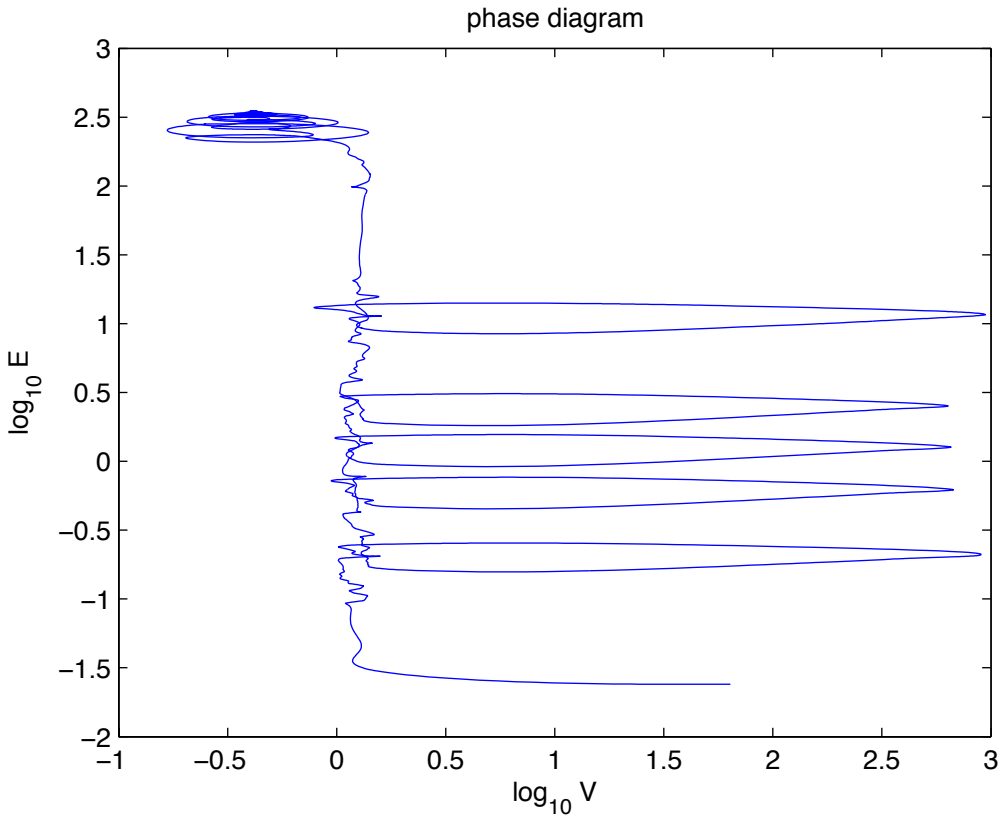

Figure 8: Phase plot of log of immune effectors vs. log of virus when control is allowed to vary daily and the initial condition is the unhealthy equilibrium.

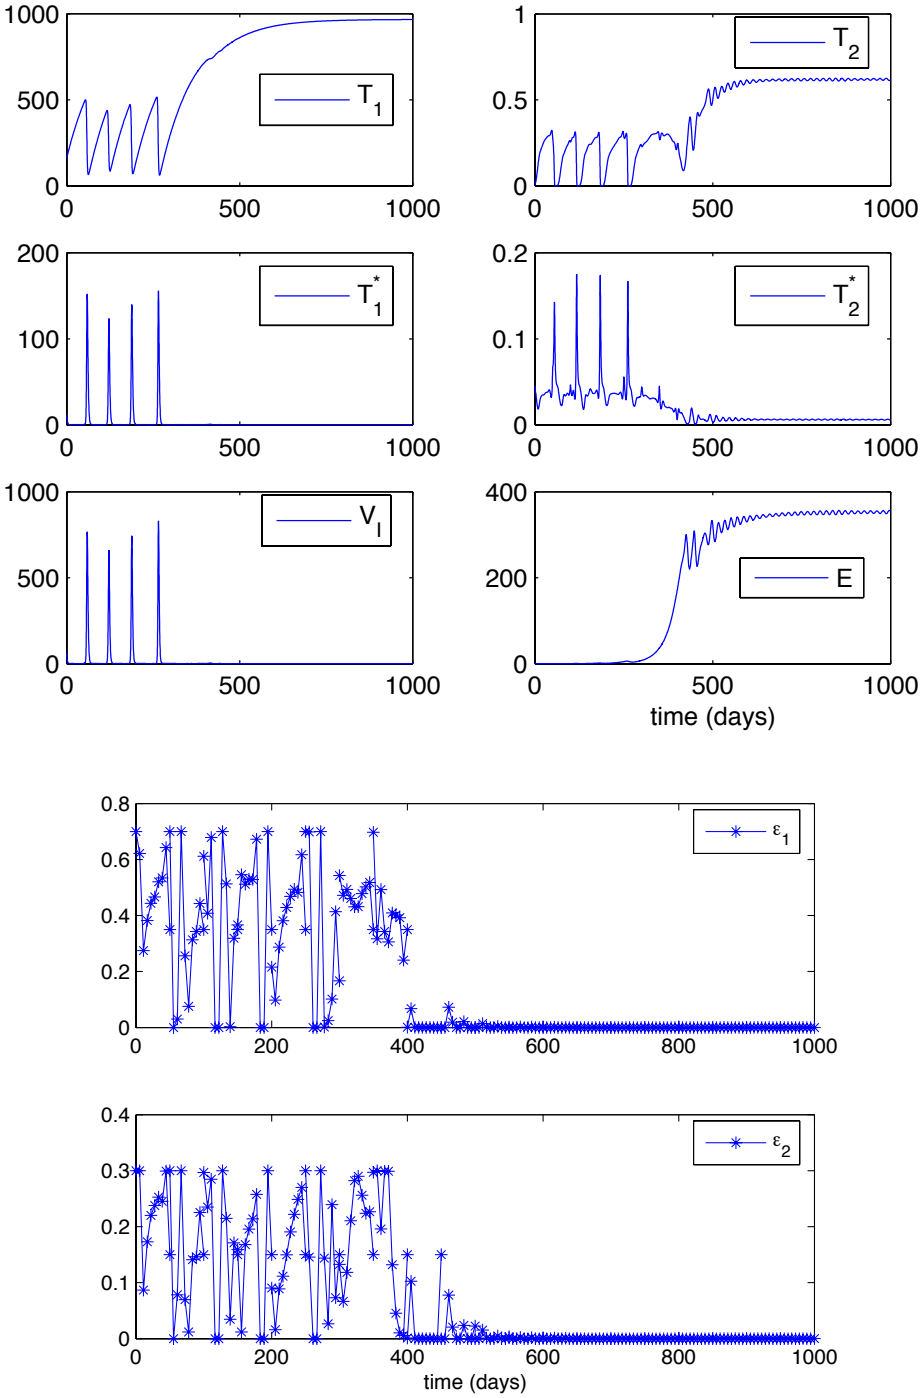

Figure 9: Model response and control when control is allowed to vary every 5 days. Initial conditions given by unhealthy steady state (7).

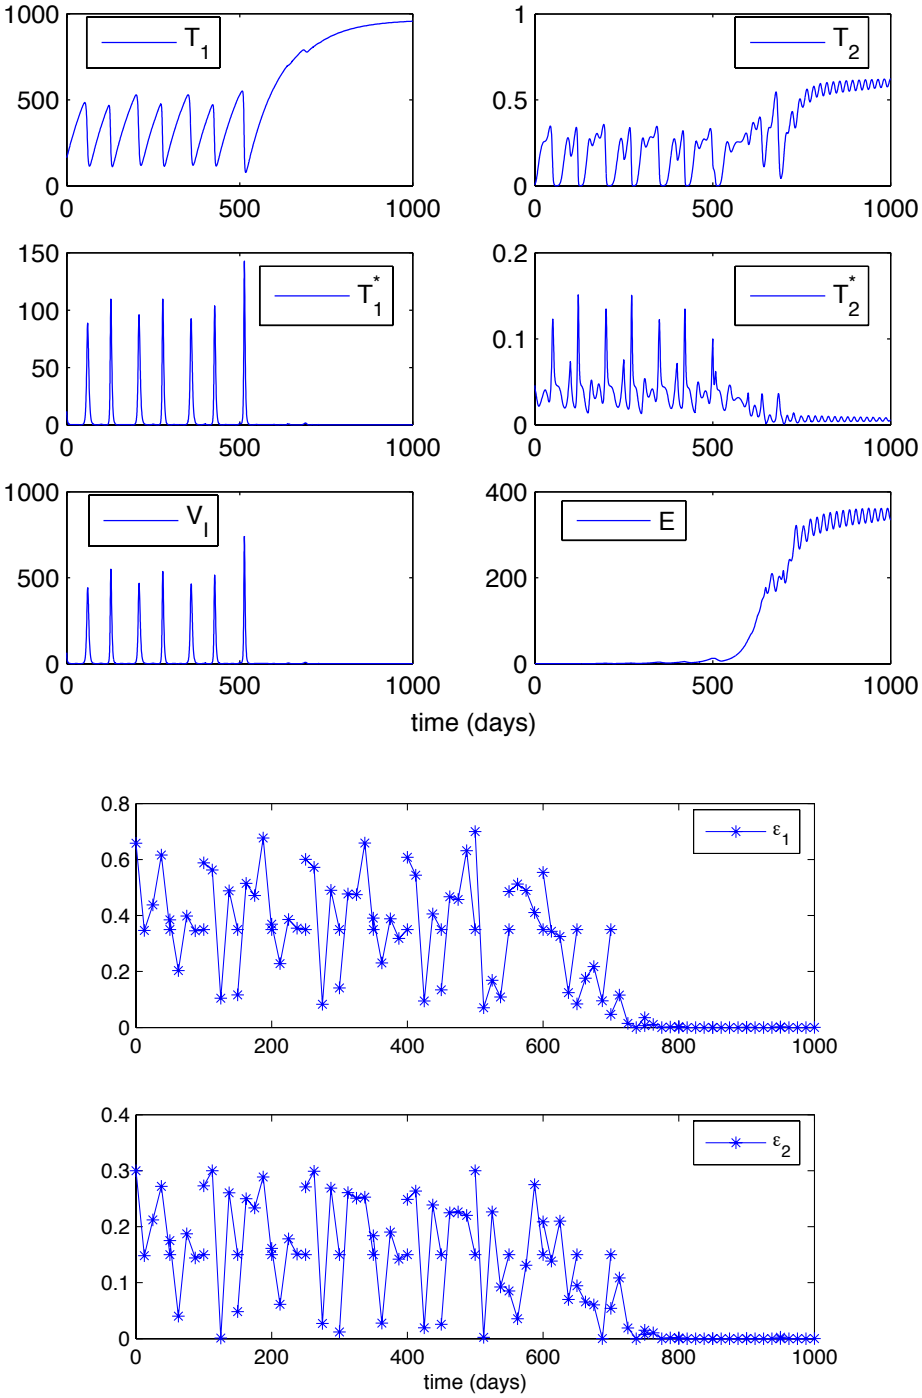

Figure 10: Model response and control when control is allowed to vary every 10 days. Initial conditions given by unhealthy steady state (7).

of RHC based treatment. Other than this cessation of treatment the experimental design is identical to that presented in Section 4. Figure 11 contains the results from this simulation. Note that this situation requires 6 periods of treatment interruptions whereas the identical situation in which the treatment regiment was followed required only 4 treatment interruptions.

As the immune response is key to the healthy steady state we should examine the effect this treatment interruption has on the immune response. Figure 12 has a comparison of the immune response when the control is allowed to vary daily but in one case the patient adhered to the treatment protocol whereas in the other case the patient went off treatment. First it delays the stimulation of the immune response by approximately 120 days. However it should also be noted that the RHC based treatment strategy was able to overcome the failure to take medication and still transferred the patient to healthy steady state. Figure 13 repeats this situation when the control is allowed to vary every 5 days. Note that when the control can only vary every 5 days the stimulation of the immune response is delayed by approximately 180 days, a longer delay than when the control can vary daily.

6 Numerical Results: State Estimation

In this section we will address the problem of implementing RHC methodology while relying on the EKF to estimate the state each time a new control is implemented. (In some cases - see ([5]) - a more efficient filter might be found, but we use EKF here for simplicity.) RHC is a full-state feedback control technique, thus it requires state knowledge. It does not require state knowledge continually as a Ricatti equation technique does, but it does require an estimate of the state each time a new control is computed. As we have a nonlinear model we choose to employ the EKF. In particular we will consider a formulation based on a continuous process model and discrete time measurements. The general problem we want to consider is given a mathematical model for a physical system, x, corrupted by some process noise, w, and a model for data measurements, z, also corrupted by some noise, v, what is the best, in some sense, estimate for the true physical system.

Mathematically the partial feedback control problem can be stated as

$$\begin{aligned} \dot{x} &= f(x,t) + g(t)w(t) \\ z_k &= h(x(t_k),k) + v_k, \end{aligned} \tag{8}$$

where $x(t_0)$ is normally distributed with mean $\bar{x}_0$ and covariance P_0, $w(t)$ and v_k are white noise processes uncorrelated with $x(t_0)$ and with each other and with means 0 and covariances Q and R, respectively. The formulation of this problem in this manner is presented in [20]. In the case where f and h are linear, the problem reduces to the famed Kalman filter (KF) [17]. Note in this type of state estimation we are trying to estimate statistics which describe a probability distribution on the true state of the process. In the KF the optimal estimate is defined by a normal distribution and thus only the mean, $\hat{x}$, and the covariance, P, are needed to completely define the conditional distribution of the state dependent on the data, $p[x(t)|(z_k)]$. See [25] for an excellent introduction to the topic.

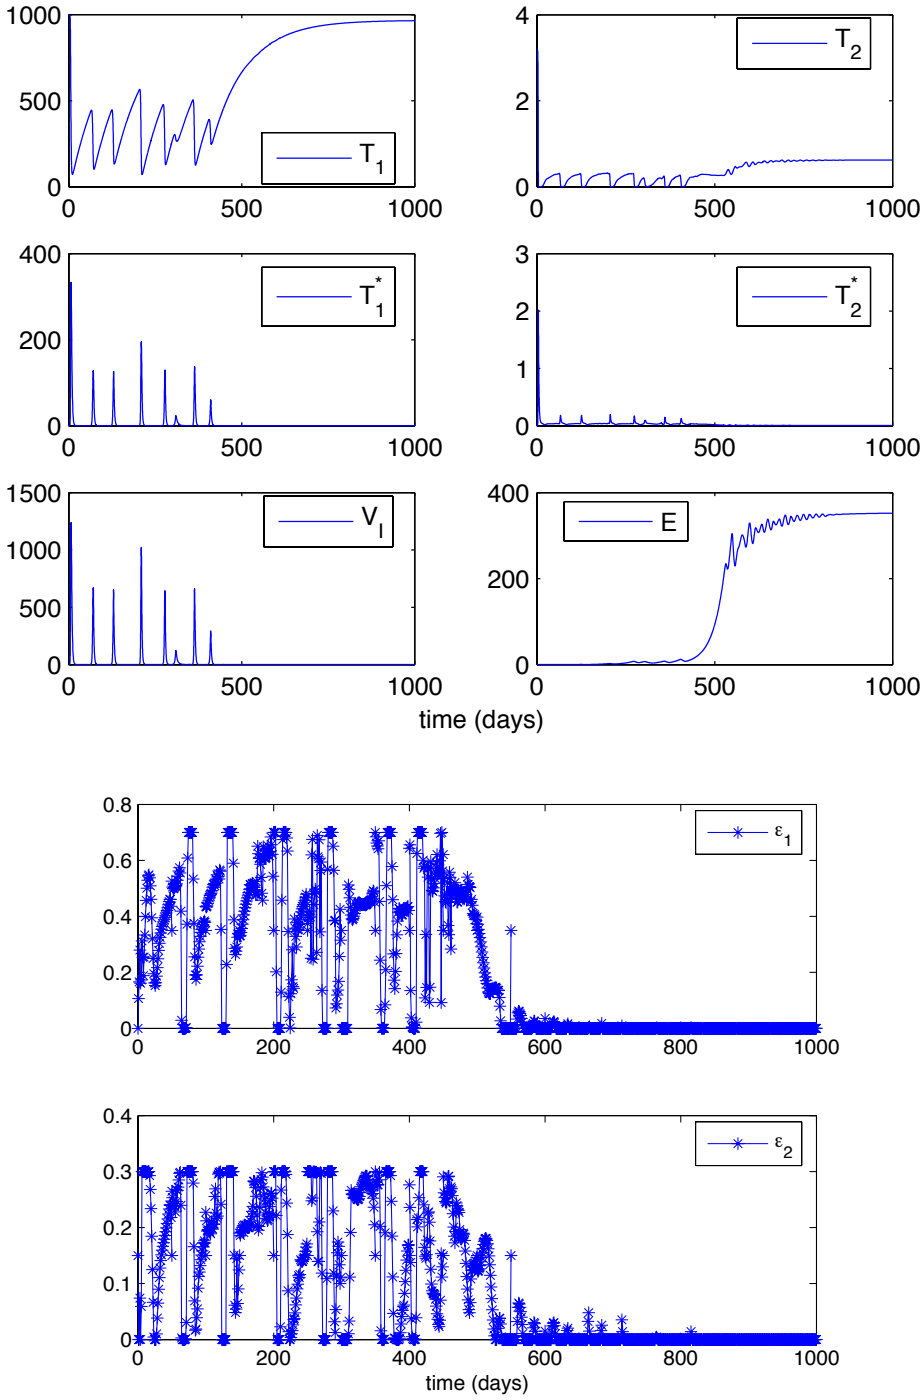

Figure 11: Model response and control when control is allowed to vary daily and control was interrupted after 300 days for 10 days.

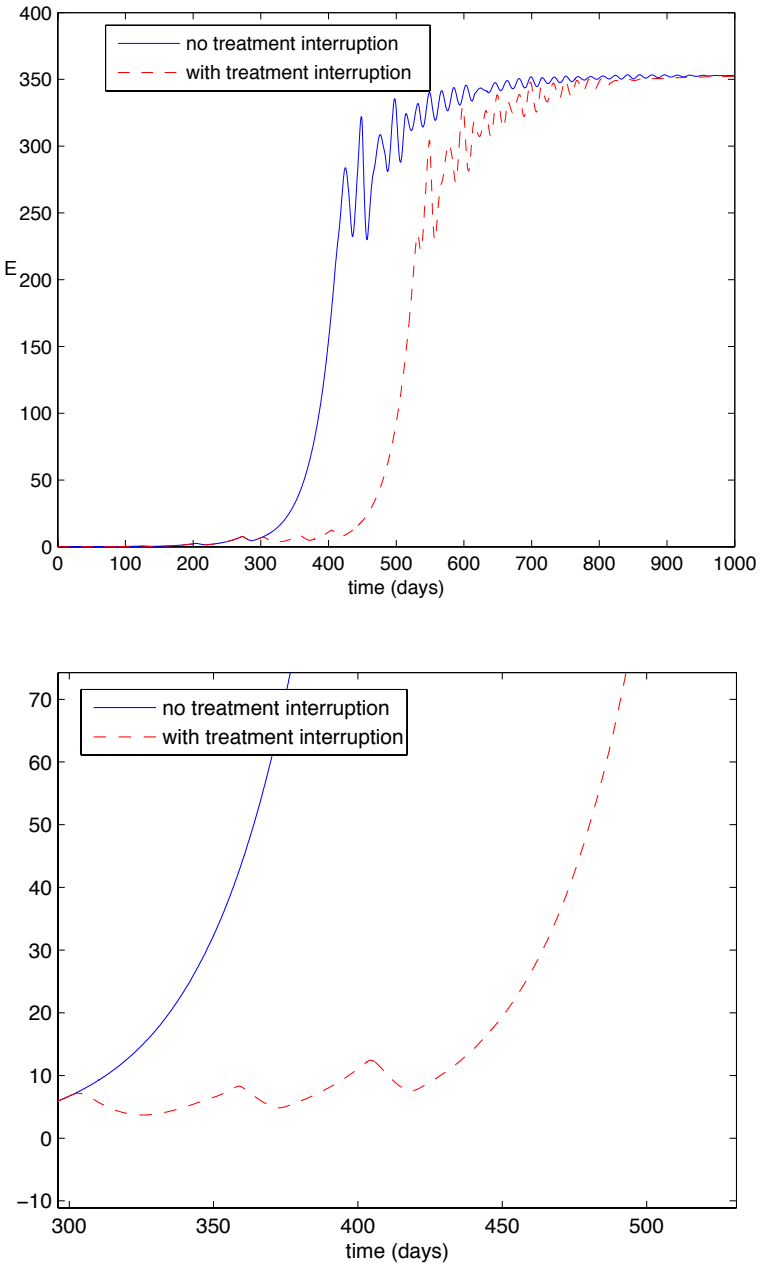

Figure 12: Immune response E when control is allowed to vary daily. Comparison between following treatment and when treatment was interrupted after 300 days for 10 days. Plot on the top contains the full time interval, while the plot on the bottom focuses on the time period directly after the patient went off medication.

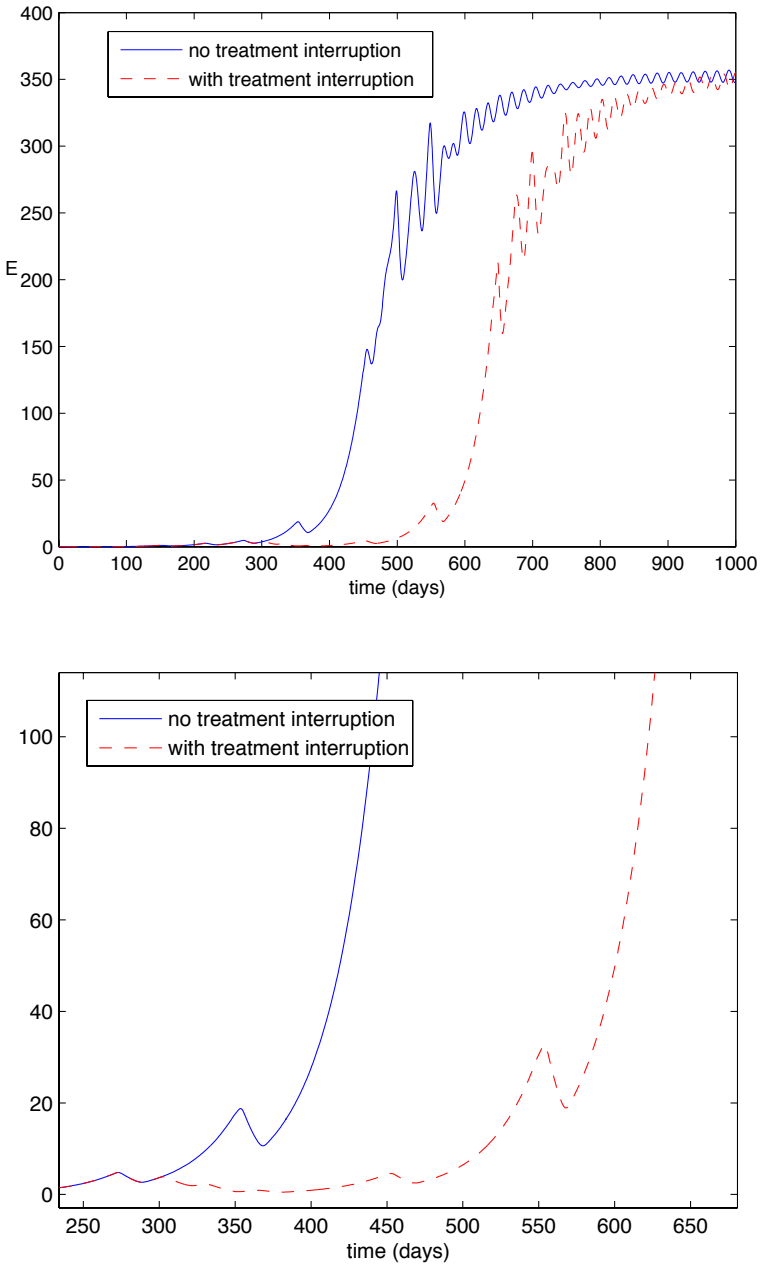

Figure 13: Immune response E when control is allowed to vary every 5 days. Comparison between following treatment and when treatment was interrupted after 300 days for 10 days. Plot on the top contains the full time interval, while the plot on the bottom focuses on the time period directly after the patient went off medication.

The EKF attempts to make an estimate of the true state with a "predictor-corrector" type of implementation. First in the "predictor" stage of the algorithm where no data is available and the dynamics are defined by the model, the time update is employed

$$\begin{aligned} \dot{\hat{x}} &= f(\hat{x}, t) \\ \dot{P} &= P\nabla f(\hat{x})^T + \nabla f(\hat{x})P + gQg^T. \end{aligned} \tag{9}$$

These equations are integrated between time points t_k and t_{k+1}. Then when new data is available the "corrector" stage or measurement update is employed

$$\begin{aligned} K_k &= P^-(t_k)\nabla h(\hat{x})^T[\nabla h(\hat{x})P^-(t_k)\nabla h(\hat{x})^T + R]^{-1} \\ P(t_k) &= [I - K_k\nabla h(\hat{x})]P^-(t_k) \\ \hat{x}_k &= \hat{x}_k^- + K_k[z_k - \nabla h(\hat{x})\hat{x}_k^-]. \end{aligned} \tag{10}$$

Note that this method only gives us the mean and covariance. Since f and h are nonlinear the conditional distribution is generally non-normal and would require perhaps infinitely many statistics to fully describe it. Further discussion of the issues related to implementing this filter on this particular model can be found in [12] and [11].

We will use simulated data created by the model, corrupted by observation noise. The only process noise we will assume is based on the accuracy of the integrator and this is assumed small. Thus we will choose $Q = 10^{-6}I$. The form of the data we are going to use is

$$z_k = \begin{pmatrix} T_1(t_k) + T_1^*(t_k) \\ V_I(t_k) + V_{NI}(t_k) \\ E(t_k) \end{pmatrix} + r_k, \tag{11}$$

where r_k is normally distributed with mean 0 and covariance R. For the following experiments R will be of the form

$$R = \begin{pmatrix} r_1 & 0 & 0 \\ 0 & r_2 & 0 \\ 0 & 0 & r_3 \end{pmatrix}, \tag{12}$$

where $r_1 = 1000$, $r_2 = 1000$, and $r_3 = 10$. The control was allowed to vary every 10 days. Data was collected every 10 days. The initial state for the problem was taken as the unhealthy steady state. The initial condition for the estimator was chosen as $x_0^{EKF} = 0.3\gamma x_0^{true}$ where γ is normally distributed with mean 1 and standard deviation 1. Figures 14 and 15 depict the results from this computational experiment. Note that the EKF is rather effective in filtering the noisy data. This allows for the estimate of the initial condition x_i to the control problem, O_i, to be accurate. This ultimately allows the RHC methodology to be successful in transferring the system to the immune response dominant equilibrium.

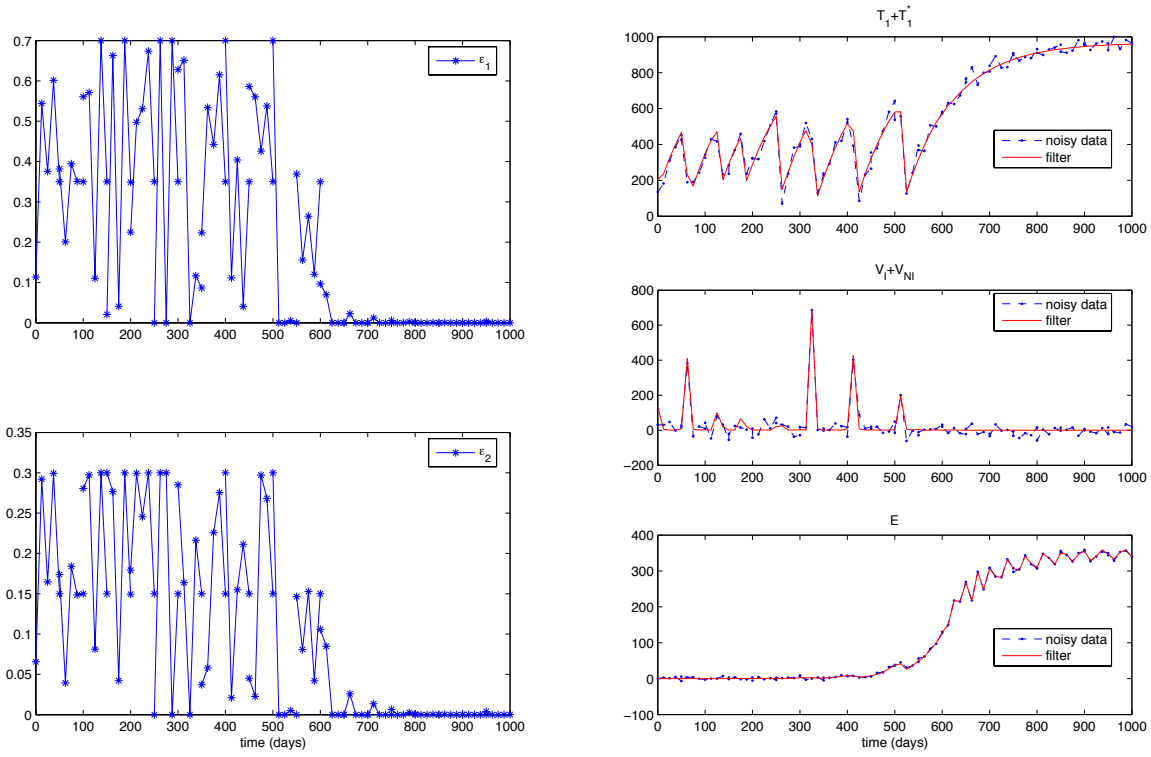

Figure 14: Control found when EKF is used as state estimator on the left. Performance of EKF as a filter of noisy data on the right.

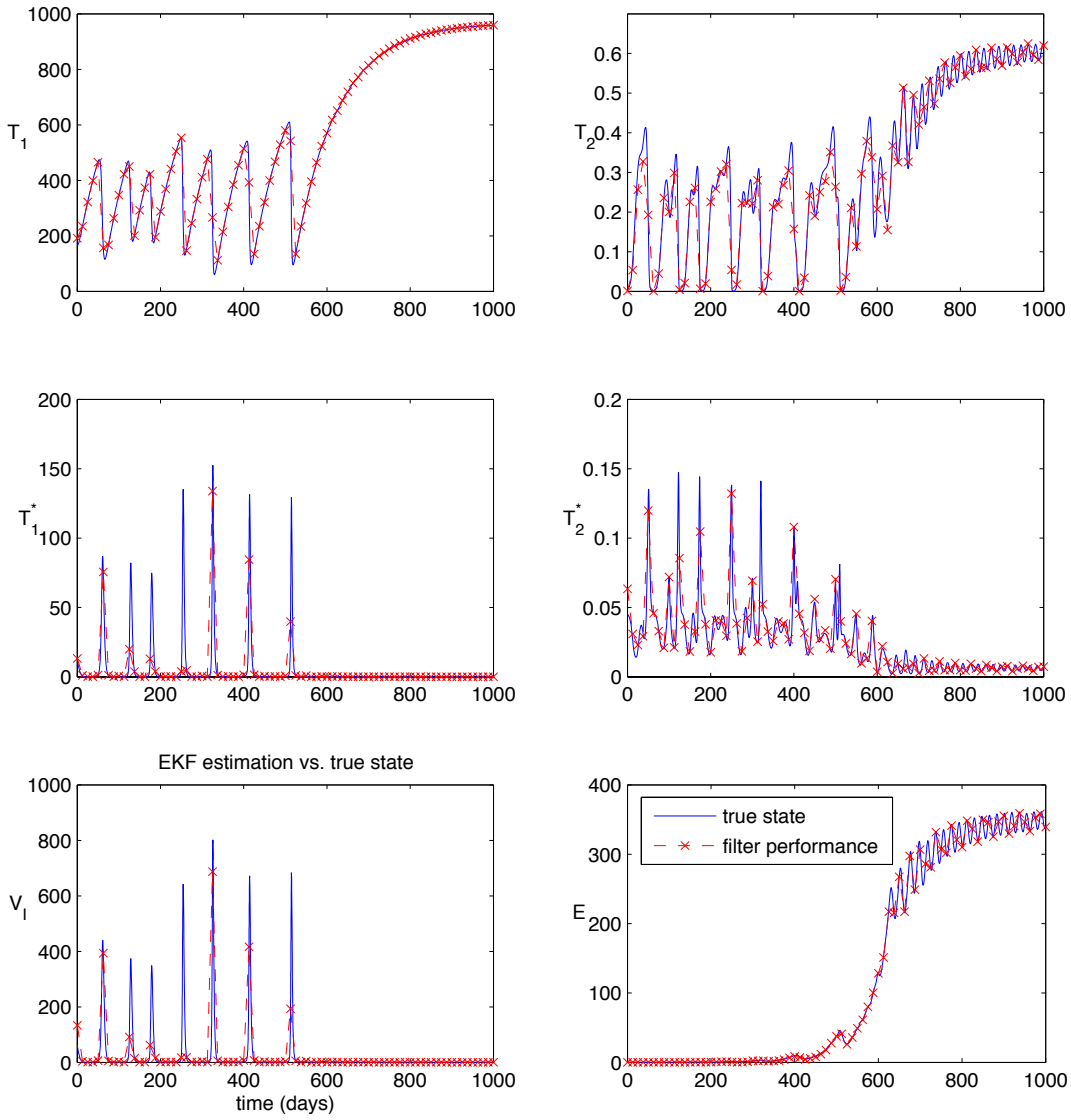

Figure 15: Performance of EKF as a state estimator. Dash line (star) denotes EKF estimation, while solid line represents true state.

7 State Estimation and Unscheduled Treatment Interruptions

An additional test of this methodology that we would like to investigate is the case of using the EKF as a state estimator while interrupting the treatment (not in accordance with the proposed treatment schedule) for 10 days after 300 days . Data is collected every 5 days while the control is allowed to vary every 5 days. The parameters for the observation noise for the following experiment are $r_1 = 1000$, $r_2 = 1000$, and $r_3 = 10$. The initial condition for the estimator is chosen as $x_0^{EKF} = 0.3\gamma x_0^{true}$ where γ is normally distributed with mean 1 and standard deviation 1. The plots for these simulations can be seen in Figures 16 and 17. These results are encouraging. We were able to transfer from an unhealthy steady state to a healthy one in the presence of unscheduled treatment interruption and noisy data collected every five days within approximately 800 days.

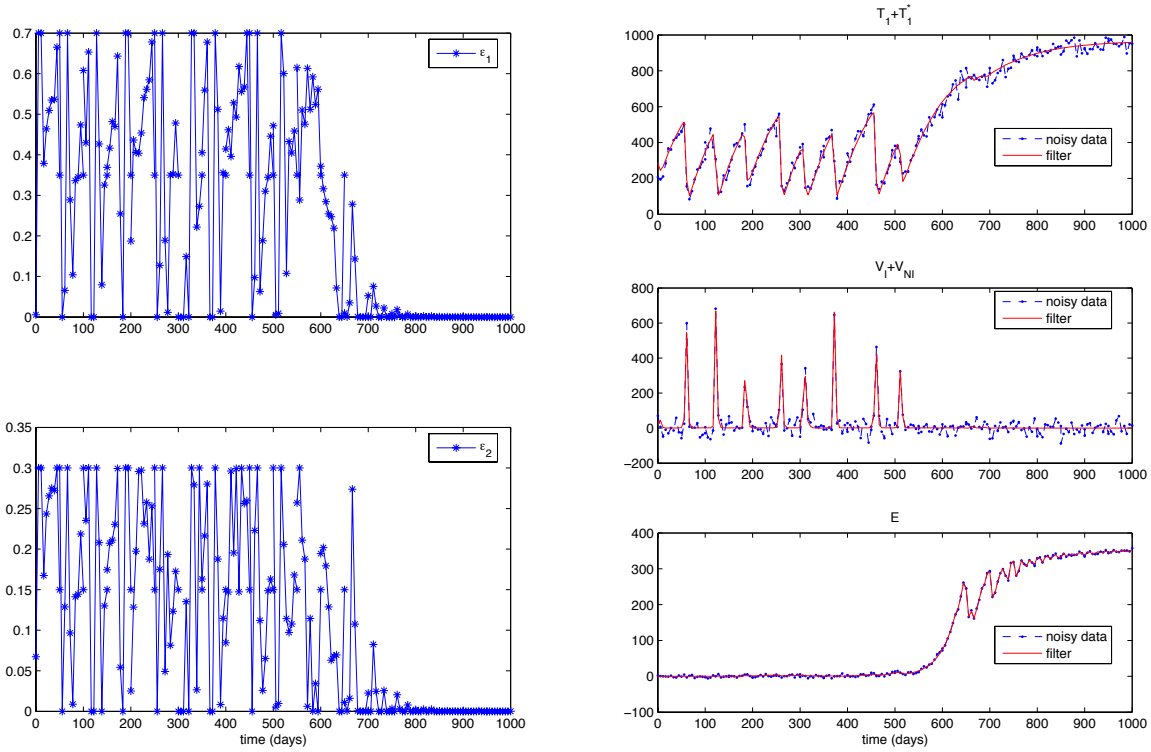

Figure 16: Control found when EKF is used as state estimator on the left. Performance of EKF as a filter of noisy data on the right. Unscheduled 10 day treatment lapse after 300 days.

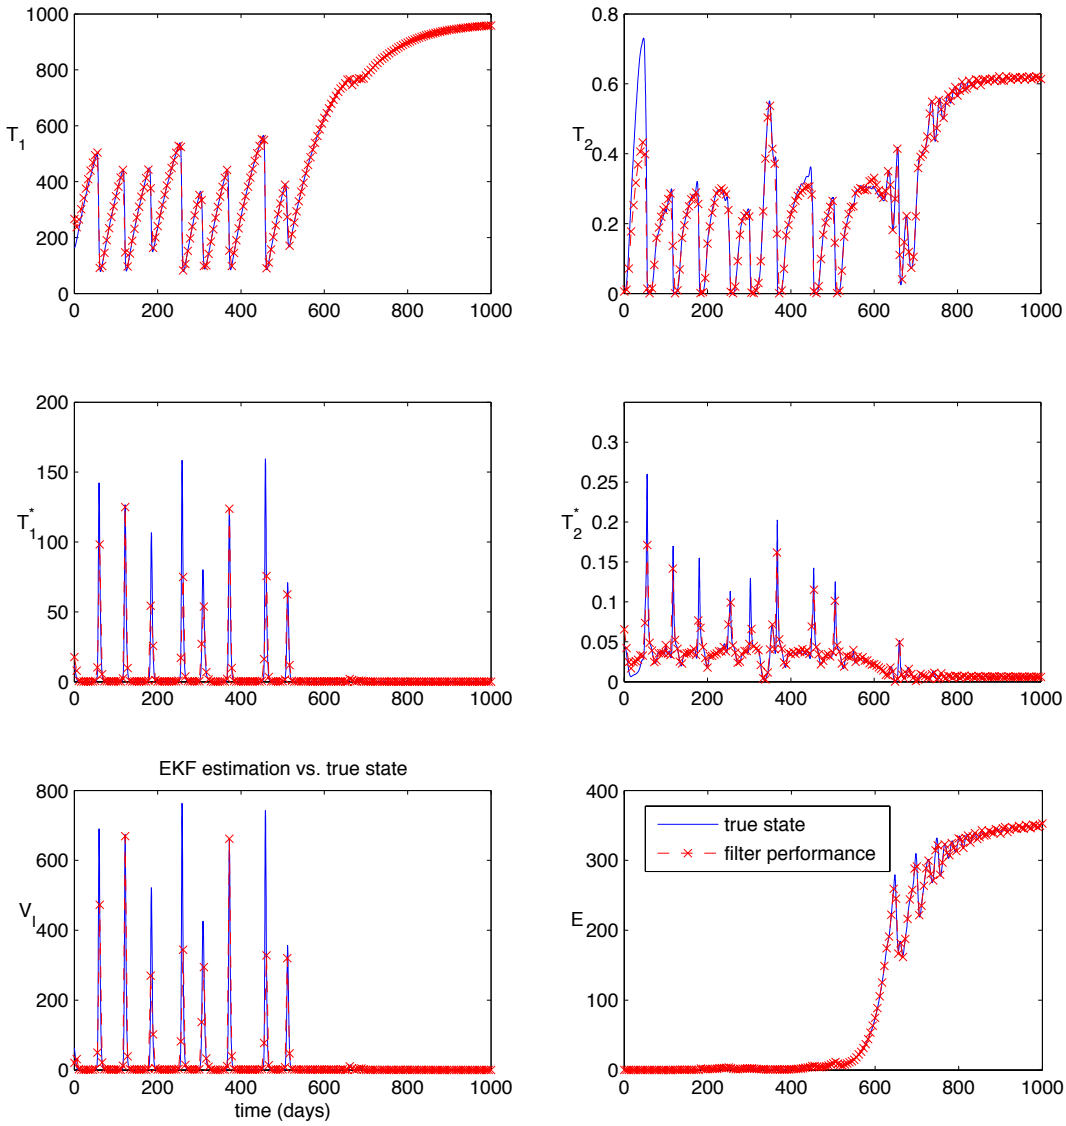

Figure 17: Performance of EKF as a state estimator. Dash line (star) denotes EKF estimation, while solid line represents true state. Data was collected every 5 days. Unscheduled 10 day treatment lapse after 300 days.

8 Inaccurate Model Parameters

Another natural question to pose especially when considering feedback control is how robust this methodology is with respect to inaccurate models (more precisely, a reasonable model but with inaccurate parameter values). As this or any other model is only an approximation of the true physical/biological process, we will study the effectiveness of a dosing strategy based on a model with incorrect parameters. The experimental set up will remain the same as in Section 4 except for the fact that our treatment schedule will be based on a model with incorrect parameters. What we mean by inaccurate model is that we will assume the true system is governed by equation (1) with model parameters (values in Table 2) that vary randomly in each time interval. In computing the control we will not use these randomly perturbed parameters, but rather the values in Table 2. The following list will further illustrate this experiment. Let q denote the vector of model parameters.

1. Using the model, $\dot{x} = f(x, q_{tab}, u)$, where q_{tab} are the parameter values from the table, solve problem O_i for optimal control u_i.
2. Create the random set of parameters which will define the true state of the patient for the time interval $[t_i\ t_{i+1}]$, using the formula $(q_{pat})_j = (\gamma)_j (q_{tab})_j$ where $(\cdot)_j$ is the j^{th} component of a vector and each component of γ is normally distributed with mean 1 and a given standard deviation.
3. The model, $\dot{x} = f(x, q_{pat}, u)$, is used to simulate the response of the patient on the given time interval.
4. The process is repeated on the following time interval.

For the model responses and controls in Figure 18 the standard deviation of $(\gamma)_j$ is 0.05. The results when using inaccurate model parameters were far less promising than those for unscheduled treatment interruptions of when using noisy data. To obtain the results in Figure 18 we had to frequently change doses (daily) and the level of perturbation of the parameters had to be low. Experiments in which more noise was added to the parameters or a less frequent change in dosing was allowed were generally unsuccessful in the transferring the system to healthy steady state.

This behavior should not entirely be unexpected. The results from the sensitivity analysis in [12] show that this model, especially when the immune response is stimulated, is much more sensitive to variation in parameters than variation in initial conditions. Thus we should expect our methodology to overcome imperfect state knowledge much more readily than inaccurate parameters.

9 Conclusion and Future Directions

In this paper we introduced a RHC control methodology in the context of an HIV mathematical model. To the authors knowledge this is the first work to combine RHC and nonlinear filtering in the context of HIV control. It was shown that these techniques can be used to derive treatment strategies that can effectively control a model of HIV infection under a variety of adverse conditions. These adverse conditions include unscheduled treatment interruption, beginning treatment on a suppressed

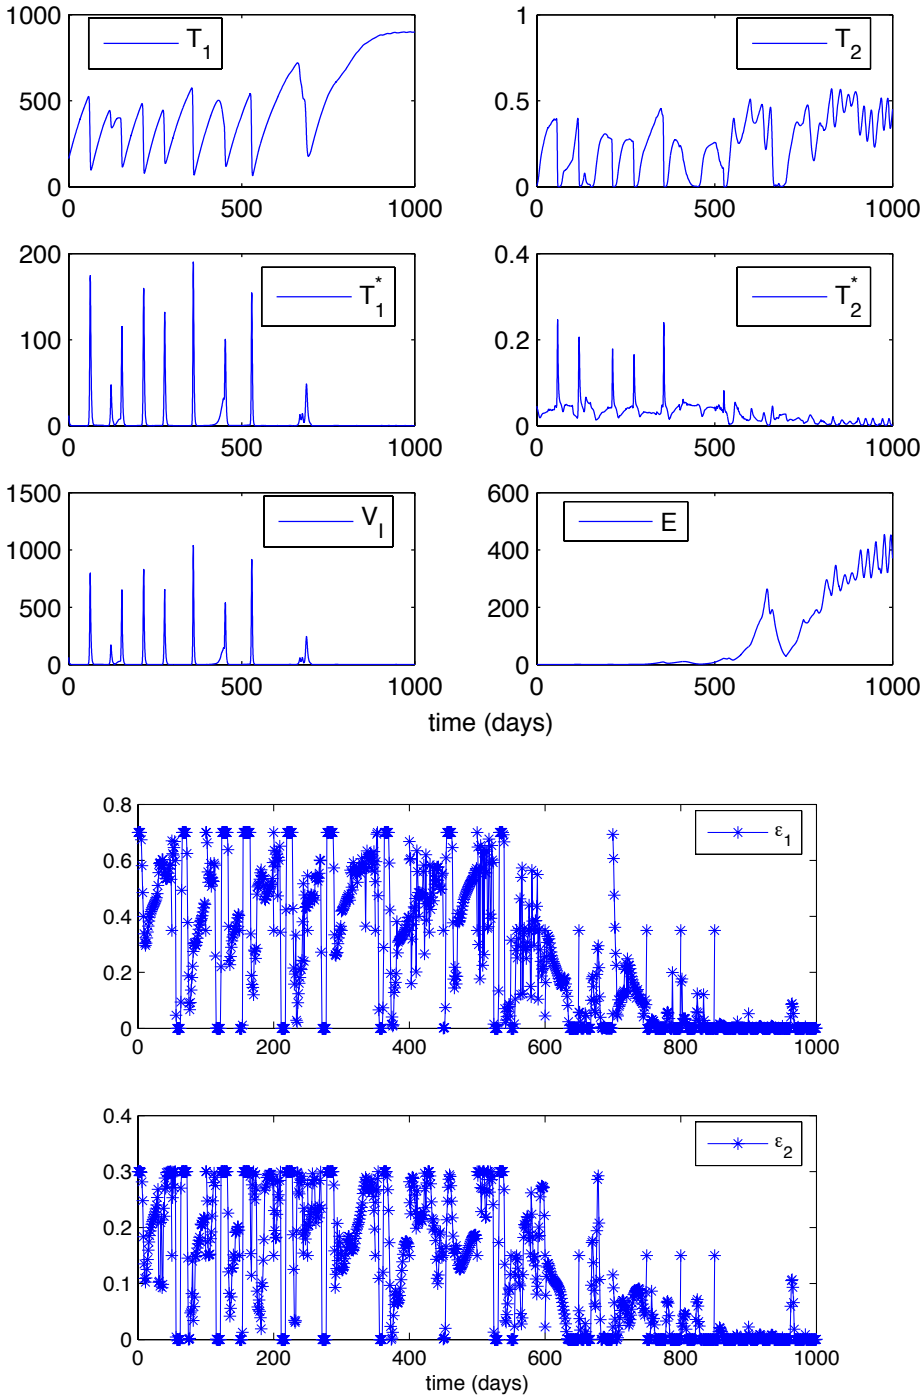

Figure 18: Model response and control when control is allowed to vary daily and model parameters are randomly perturbed.

immune system, and incomplete and noisy data. This work also illustrated the relationship between how often the efficacy of the drugs can be changed and the length of time necessary to reach the immune dominant steady state. We also explored the number of interruptions necessary under several detrimental conditions to control HIV infection.

This work sets the stage for a variety of ways mathematical analysis can be used to improve treatment of HIV patients. We have dealt with several of the issues related to designing control based treatment protocols. However this work also begs the question of several courses of analysis related to better designing these protocols. These improvements include the implementation of a filter that incorporates the censored data produced by many viral assays, a pharmacokinetic model of specific drugs, incorporation of the potential for drug-resistance in model and intervals of data collection and treatment changes that are more in line with a clinical condition.

Acknowledgements

This work in this manuscript was supported in part by grant NIH 1 R01 GM67299-04 and in part by grant NIAID 9R01AI071915-07.

References

[1] B. M. Adams, H.T. Banks, M. Davidian, and E.S. Rosenberg. Model fitting and prediction with HIV treatment interruption data. *Bulletin of Mathematical Biology*, 69:563–584, 2007.

[2] B.M. Adams, H.T. Banks, M. Davidian, H.D. Kwon, H.T. Tran, S.N. Wynne, and E.S. Rosenberg. HIV dynamics: Modelling, data analysis and optimal treatment protocols. *Journal of Computational and Applied Mathematics*, 184:10–49, 2005.

[3] B.M. Adams, H.T. Banks, Hee-Dae Kwon, and Hien T. Tran. Dynamic multidrug therapies for HIV: Optimal and STI control approaches. *Mathematical Biosciences and Engineering*, 1(2):223–241, September 2004.

[4] Brian M. Adams. *Non-Parametric Parameter Estimation and Clinical Data Fitting with a Model of HIV Infection.* PhD thesis, North Carolina State University, 2005.

[5] H.T. Banks, S. Hu, Z.R. Kenz, and Hien T. Tran. A comparison of nonlinear filtering approaches in the context of an HIV model. Technical report, CRSC-TR09-16, Center for Research in Scientific Computation, North Carolina State University, Raleigh N.C., 2009, Mathematical Biosciences and Engineering, to appear.

[6] H.T. Banks, Hee-Dae Kwon, J.A. Toivanen, and Hien Tran. An SDRE base estimator approach for HIV feedback control. *Optimal Control Applications and Methods*, 27(2):93–121, 2006.

[7] S. Bonhoeffer, M. Rembiszewski, G.M. Ortiz, and D.F. Nixon. Risks and benefits of structured antiretroviral drug therapy interruptions in HIV-1 infection. *AIDS*, 14:2313–2322, 2000.

[8] Duncan S. Callaway and Alan S. Perelson. HIV-1 infection and low steady state viral loads. *Bulletin of Mathematical Biology*, 64(1):29–64, 2002.

[9] Rebecca V. Culshaw. Review of HIV models: The role of the natural immune response and implications for treatment. *Journal of Biological Systems*, 12(2):123–135, 2004.

[10] Rebecca V. Culshaw, Shigui Ruan, and Raymond J. Spiteri. Optimal HIV treatment by maximising immune response. *Journal of Mathematical Biology*, 48:545–562, 2004.

[11] John A. David. *Optimal Control, Estimation and Shape Design: Analysis and Applications*. PhD thesis, North Carolina State University, 2007.

[12] John A. David, Hien T. Tran, and H. T. Banks. HIV model analysis and estimation implementation under optimal control based treatment strategies. Technical report, CRSC-TR08-07, Center for Research in Scientific Computation, North Carolina State University, Raleigh N.C., 2008, International Journal of Pure and Applied Mathematics, to appear.

[13] K. Renee Fister, Suzanne Lenhart, and Joseph Scott McNally. Optimizing chemotherapy in an HIV model. *Electronic Journal of Differential Equations*, 1998(32):1–12, 1998.

[14] Winston Garira, Senelani D. Musekwa, and Tinevimbo Shiri. Optimal control of combined therapy in a single strain HIV-1 model. *Electronic Journal of Differential Equations*, 2005(52):1–22, 2005.

[15] Kazufumi Ito and Karl Kunisch. Asymptotic properties of receding horizon optimal control problems. *SIAM J. Control Optim.*, 40(5):1585–1610, 2001.

[16] Hem Raj Joshi. Optimal control of an HIV immunology model. *Optimal Control Applications and Methods*, 23:199–213, 2002.

[17] Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME–Journal of Basic Engineering*, 82(Series D):35–45, 1960.

[18] Carl T. Kelley. *Iterative Methods for Optimization*. SIAM, 1999.

[19] Denise Kirschner, Suzanne Lenhart, and Steve Serbin. Optimal control of the chemotherapy of HIV. *Journal of Mathematical Biology*, 35:775–792, 1997.

[20] Frank Lewis. *Optimal Estimation: With an Introduction to Stochastic Control*. John Wiley and Sons Inc, 1986.

[21] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P.O.M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36:789:814, 2000.

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

[23] S. Qin and T. Badgewell. An overview of industrial model predictive control technology. *Chemical Process Control*, 93(316):232–256, 1997.

[24] H. Shim, S.J. Han, C.C. Chung, S.W. Nam, and J.H. Seo. Optimal scheduling of drug treatment for HIV infection: Continuous dose control and receding horizon control. *International Journal of Control, Automation and Systems*, 1(3):401–407, Sept. 2003.

[25] Greg Welch and Gary Bishop. An introduction to the Kalman filter. Technical Report TR95-041, University of North Carolina, Dept. of Computer Science, 2001.

[26] R. Zurakowski, M.J. Messina, S.E. Tuna, and A.R. Teel. HIV treatment scheduling via robust nonlinear model predictive control. *Proceedings of the 5th Asian Control Conference*, 2004.