Optimal control of innate immune response

Robert F. Stengel*, Raffaele Ghigliazza, Nilesh Kulkarni and Olivier Laplace

School of Engineering and Applied Science, Princeton University, Princeton, NJ 08544, U.S.A

SUMMARY

Treatment of a pathogenic disease process is interpreted as the optimal control of a dynamic system. Evolution of the disease is characterized by a non-linear, fourth-order ordinary differential equation that describes concentrations of pathogens, plasma cells, and antibodies, as well as a numerical indication of patient health. Without control, the dynamic model evidences sub-clinical or clinical decay, chronic stabilization, or unrestrained lethal growth of the pathogen, depending on the initial conditions for the infection. The dynamic equations are controlled by therapeutic agents that affect the rate of change of system variables. Control histories that minimize a quadratic cost function are generated by numerical optimization over a fixed time interval, given otherwise lethal initial conditions. Tradeoffs between cost function weighting of pathogens, organ health, and use of therapeutics are evaluated. Optimal control solutions that defeat the pathogen and preserve organ health are demonstrated for four different approaches to therapy. It is shown that control theory can point the way toward new protocols for treatment and remediation of human diseases. Copyright © 2002 John Wiley & Sons, Ltd.

KEY WORDS: optimal control; biological modelling; bioinformatics; optimization

INTRODUCTION

Immune-system response to invasion by infectious microbes is a dynamic process in which potentially uncontrolled growth of the invader (or *pathogen*) is countered by various protective mechanisms. There are both tactical and strategic responses once the outer perimeter of defense — the surface *epithelial layers* of the body, including the *epidermal cells* of the skin and the *mucosal cells* that line the respiratory, gastrointestinal, and genitourinary tracts — has been broached [1–3]. The *innate immune system* provides a *humoral* (or tactical) response, signalling the presence of 'non-self' organisms and activating *B cells* to produce *antibodies* that bind to the intruders' *antigens*. The antibodies identify targets for scavenging cells, such as *neutrophils* and *macrophages*, which engulf and consume the microbes, reducing them to non-functioning units. They also stimulate the production of molecules (e.g. *cytokines*, *complement*, and *acute-phase proteins*) that either damage an intruder's plasma membrane directly (a process called *lysis*) or help trigger the second phase of immune response. The innate immune system protects against

---
*Correspondence to: Robert F. Stengel, School of Engineering and Applied Science, Princeton University, D-202 Engineering Quadrangle, Princeton, NJ 08544, USA.
E-mail: stengel@princeton.edu

Copyright © 2002 John Wiley & Sons, Ltd.

many *extracellular bacteria* or *free viruses* that are found in blood plasma, lymph, tissue fluid, or interstitial space between cells (and, therefore, that are not hidden within the cells of the body). However, innate immunity normally is not adequate for defeating microbes that burrow into cells, such as *viruses*, *intracellular bacteria*, and *protozoa*.

The *adaptive immune system* provides the *cellular* (or strategic) response to intracellular microbial assault, producing protective cells that remember specific antigens, that produce antibodies 'customized' to the pathogens, and that seek out *epitopes* (or defining regions) of antigens that are expressed on the surfaces of infected cells. Adaptive immune mechanisms are largely based on the actions of *B* and *T lymphocyte cells* that become dedicated to a single antibody type through *clonal selection*. In particular, *killer T cells* (or *cytotoxic T lymphocytes*) bind to infected cells and kill them by initiating *apoptosis* (programmed cell death). *Helper T cells* assist naive *B* cells in maturing into *memory cells* and *plasma cells* that produce the needed antibody type. The process of adaptation is slow, but it generates a cadre of immune cells with narrowly focused memory, cells that are ready to respond rapidly if invading microbes of the same type (i.e. with the same antigen epitopes) are encountered again. As inferred by this brief discussion, elements of the innate and adaptive immune systems are shared, and response mechanisms are coupled, even though separate modes of operation can be identified.

We distinguish between pre-exposure vaccination (or immunization) that is intended to prevent clinical disease from ever occurring and post-exposure therapy for a clinically diagnosed condition. The options available for clinical treatment of infection once it has been recognized focus on killing the invading microbes directly, neutralizing their deleterious effects, enhancing the efficacy of immune response, and providing palliative or healing care to other organs of the body. Few biological or chemical agents have just a single effect; for example, an agent that kills a pathogen also may damage healthy 'self' cells. Nevertheless, such agents normally are introduced with one particular goal in mind, and it is a critical function of drug discovery and development to identify new compounds that have maximum intended efficacy with minimal side effects in the general population. Within the pharmacopoeia, we find examples of *antibiotics* as microbe killers; *interferons* as microbe neutralizers; *interleukins, antigens* from *killed* (i.e. *non-toxic*) *pathogens*, and *pre-formed* and *monoclonal antibodies* as immunity enhancers (each of very different nature); and *anti-inflammatory* and *anti-histamine compounds* as palliative drugs.

Numerous models of immune response to infection have been postulated (e.g. [4–7]), with recent emphasis on the human-immunodeficiency pathogen (HIV) [8–13], but there have been few applications of control theory to the healing process. Norbert Wiener [14] and Richard Bellman [15] appreciated and anticipated that possibility, and References [16–19] advance optimization as a model for immune system dynamics. References [20, 21] survey early optimal control applications to biomedical problems, while References [22–26] investigate the role of optimization in cancer therapy. References [27, 28] offer an optimal control approach to HIV treatment, and intuitive control approaches are presented in References [29–32].

A goal of the present analysis is to show how regimens for applying drugs in a manner that maximizes efficacy while minimizing side effects and cost can be designed. In the current context, that means killing, neutralizing, or limiting growth of the pathogen with small impact on the patient's health and pocketbook. Optimal control theory can help discriminate between options for attacking the infection directly, stimulating the immune system, suppressing or amplifying ancillary biological processes, prescribing drug dosage and timing, and combining therapies for a 'cocktail' of effects.

Here, we focus on the control of a simple mathematical model for the evolution of a generic pathogen in a human. The model contains essential features of disease dynamics, including a pathogen, the innate immune system, and an organ whose function is affected by the pathogen. The model does not represent a specific pathogen, and the control variables represent hypothetical therapeutic agents that have idealized effect on the system.

We describe the disease model, define a treatment cost function to be minimized by application of the therapeutic agents, present results of the optimization for a range of treatment parameters, and draw conclusions about the method and results. The simulated results show not only the progression of the patient's state from an initially life-threatening state to a controlled or cured condition but the optimal history of therapeutic agents that produces that condition. Extension of the method to more explicit models of specific diseases, e.g. HIV and cancer, is the topic of future papers. The present paper demonstrates the power of the optimization approach to treatment using a representative model.

A MODEL OF DISEASE DYNAMICS

We consider a simple model for a pathogenic attack on an organism and the organism's immunological defense. There are four components to the dynamic state [4]:

x_1 = concentration of a pathogen (= concentration of associated antigen)
x_2 = concentration of plasma cells, which are carriers and producers of antibodies
x_3 = concentration of antibodies, which kill the pathogen
x_4 = relative characteristic of a damaged organ [0 = healthy, 1 = dead]

The mathematical model presented in Reference [4] does not allow for the effects of therapy. Consequently, we modify the model, adding the following idealized therapeutic control agents:

u_1 = pathogen killer
u_2 = plasma cell enhancer
u_3 = antibody enhancer
u_4 = organ healing factor

The four scalar, non-linear, ordinary differential equations of the modified dynamic model are:

$$\dot{x}_1 = (a_{11} - a_{12}x_3)x_1 + b_1u_1$$

$$\dot{x}_2 = a_{21}(x_4)a_{22}x_1(t - \tau)x_3(t - \tau) - a_{23}(x_2 - x_2{}^*) + b_2u_2$$

$$\dot{x}_3 = a_{31}x_2 - (a_{32} + a_{33}x_1)x_3 + b_3u_3$$

$$\dot{x}_4 = a_{41}x_1 - a_{42}x_4 + b_4u_4 \tag{1-4}$$

The a_{ij} and b_i are constants, except $a_{21}(x_4)$, which is a non-linear function that describes the immune deficiency triggered by damage to the organ:

$$a_{21}(x_4) = \begin{cases} \cos \pi x_4, & 0 \leqslant x_4 \leqslant 1/2 \\ 0, & x_4 \geqslant 1/2 \end{cases} \tag{5}$$

The equation can be expressed in the vector form,

$$\dot{\mathbf{x}}(t) = \mathbf{f}[\mathbf{x}(t), \mathbf{u}(t)] \tag{6}$$

where the vector $\mathbf{x}$ is called the state of the system and $\mathbf{u}$ is the control vector. The dynamic equation is solved over the time interval (t_o, t_f) subject to assumed initial conditions, $\mathbf{x}(t_o) = \mathbf{x}_o$. The state and control are subject to inequality constraints, $\mathbf{x} \geqslant \mathbf{0}$, $\mathbf{u} \geqslant \mathbf{0}$ because concentrations cannot go below zero, and $x_4 \leqslant 1$, because organ death is described as $x_4 = 1$. The model allows us to compute the history of organ health, pathogen, plasma, and antibody concentrations over time, given the initial state of the system and a specified therapeutic protocol. By itself, however, the model does not tell us what the therapeutic protocol should be.

This model of pathogenic attack contains linear parameters that cause each element of the state to grow or decay, linear and non-linear coupling of state elements, and affine linear control. The most complicated equation describes the evolution of plasma cells as a non-linear function of organ health and antigen and antibody concentrations, with a pure time delay, τ (here assumed to be zero), and a steady-state concentration of plasma cells, x_2^*. The pathogen has a natural tendency to grow exponentially (hence, $a_{11} > 0$) that is mediated by the presence of antibodies, x_3. Conversely, the evolution of antibodies is governed not only by a balance between natural birth and death of cells expressed by a_{32} (> 0) but by the presence of the pathogen, x_1, and by the generative effect of plasma cells, x_2. The organ is considered to be perfectly healthy when x_4 is zero and to be dead when x_4 is one. Linear control effects are entirely justified in an abstract model such as this, as there is no basis for making the control effects more complex. This simple control model facilitates the main presentation (i.e. the method for development of optimal therapeutic policies), and more complex control effects could easily be incorporated in the model and optimization.

The model presented in Equations (1–4) has some important basic features. One such feature describes the intrinsic ability of the immune system to react to pathogen attacks. Under certain severe circumstances, the immune system may be incapable of recovering from any presence of pathogens, no matter how small, and continuing treatment would be required. For our study, we have chosen model parameters that allow the system to recover or succumb to pathogenic attack (without treatment) as a function of initial conditions.

We define the sub-clinical case (Case 1) as one that is not likely to require medical examination and the clinical case (Case 2) as a more severe condition warranting medical consultation. Parameter values and initial conditions are listed in Table I. The initial plasma cell and the antibody populations are fixed, and the organ is initially healthy. The principal difference in the four cases is the initial concentration of pathogen, $x_1(0)$. Prior to pathogen attack, the steady-state value of antibody concentration that corresponds to $x_2^* = 2$ is

$$x_3(0) = (a_{31}/a_{32})x_2^* \tag{7}$$

The natural response of the system over a period of 10 time units and with no application of therapy is shown in Figure 1. In the figure, with $\mathbf{u} = \mathbf{0}$, there is a rapid buildup of plasma cells and a degradation of organ health. The four cases are classified by the pathogen concentration history, which either decays (Cases 1 and 2), approaches a steady value (Case 3), or diverges (Case 4), paralleling the qualitative illustrations of Ref. [4].

The system's natural response highlights a feature particular to our choice of the a_{21} function. The function represents the effect of health on the capacity to generate plasma cells, which decreases as damage to the organ increases. When organ health reaches a certain point, namely

Table I. Parameters and initial conditions of the dynamic model.

$a_{11} = 1$	$a_{12} = 1$		$b_1 = -1$	$x_1(0)$ see below
a_{21} see Equation (5)	$a_{22} = 3$	$a_{23} = 1$	$b_2 = 1$	$x_2(0) = 2$
$a_{31} = 1$	$a_{32} = 1.5$	$a_{33} = 0.5$	$b_3 = 1$	$x_3(0) = 2x_2^*/3$
$a_{41} = 1$	$a_{42} = 1$		$b_4 = -1$	$x_4(0) = 0$

Case 1: $x_1(0) = 1.5$ Case 3: $x_1(0) = 2.57$
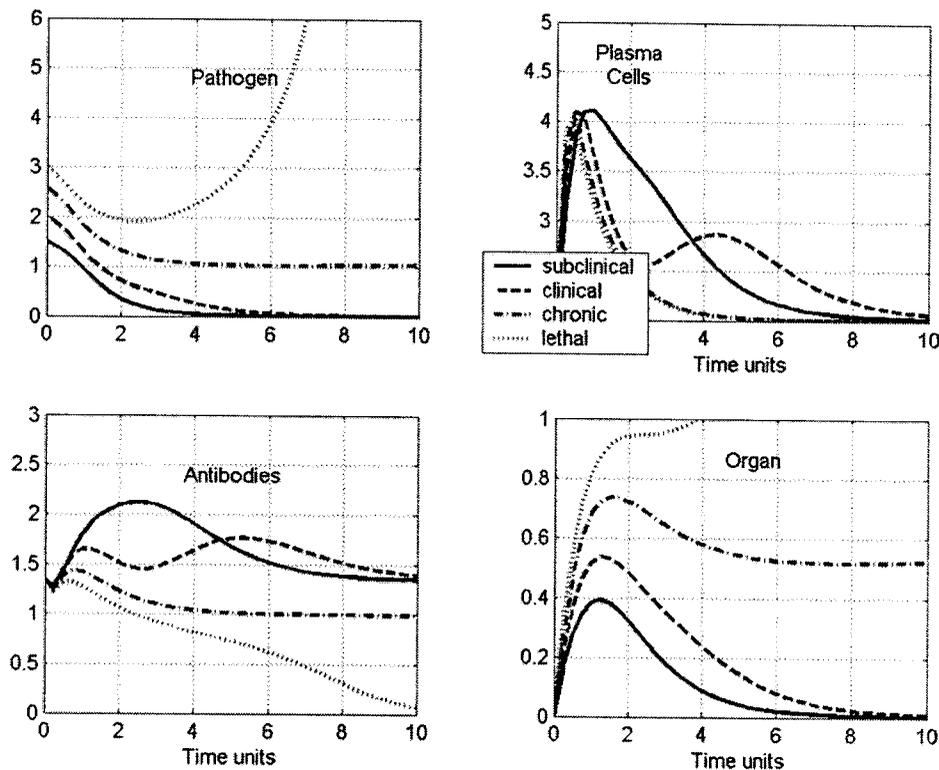Case 2: $x_1(0) = 2$ Case 4: $x_1(0) = 3$

Figure 1. Natural response to pathogenic attack.

$x_4 \leqslant 0.5$, the production of plasma cells stops. This function produces the four initial-condition-dependent histories depicted in Figure 1. In the sub-clinical case, the immune system integrity is never threatened, and pathogens are successfully destroyed. The patient's health is never in danger. If the initial infectious dose is increased, the pathogen compromises immune response. This occurs for the clinical, chronic, and lethal cases after approximately 0.5 time units. Subsequent histories for x_2 show exponential decrease of the plasma cell concentration toward the limit value of $x_2 = 2$ due to aging. In the meantime, antibodies fight off the pathogen. If the antibody concentration stays high relative to pathogen concentration, it ultimately decreases the

latter enough for the organ to heal. The result is an increased immune system efficiency and final destruction of pathogens. The clinical case (#2) reflects momentary threat to the patient's health and thus justifies close monitoring. If the antibody concentration is not high enough, the pathogen growth remains uncurbed, developing into a lethal case. This is what happens for high initial pathogen concentrations. The antibodies are not sufficient to overcome the infection (Case 4), leading to fatal damage of the organ. We define the chronic case (#3) as the limit between the two previous ones, where the pathogen reaches a steady-state value without the patient being completely cured.

These open-loop simulations show that as soon as $x_3 > a_{11}/a_{12}$, $x_1 < x_4$, a_{42}/a_{41}, and $x_4 < 1$, the disease evolves towards complete recovery of the patient. These conditions assure that the pathogen concentration keeps decreasing while organ health improves. Furthermore, if $x_4 < 0.5$, the latter assures steadily growing concentration of antibodies. Given this final state, patient health evolves toward complete recovery after the treatment stops, even if the pathogen concentration is not zero at that time.

TREATMENT COST FUNCTION AND DEFINITION OF THE OPTIMAL CONTROL POLICY

The optimal therapeutic protocol is derived by minimizing a positive-definite treatment cost function, J, that penalizes large values of pathogen concentration, poor organ health, and excessive application of therapeutic agents over the fixed time interval, (t_o, t_f):

$$J = \frac{1}{2}\left(p_{11}x_{1_f}^2 + p_{44}x_{4_f}^2\right) + \frac{1}{2}\int_{t_o}^{t_f}\left(q_{11}x_1^2 + q_{44}x_4^2 + r_{11}u_1^2 + r_{22}u_2^2 + r_{33}u_3^2 + r_{44}u_4^2\right)\mathrm{d}t \qquad (8)$$

The scalar cost function takes the general form,

$$J = \phi\left[\mathbf{x}(t_f)\right] + \int_{t_o}^{t_f} L[\mathbf{x}(t), \mathbf{u}(t)]\,\mathrm{d}t \qquad (9)$$

which includes the *terminal cost*, $\phi[\mathbf{x}(t_f)]$, associated with values of pathogen and organ health at the end of the treatment period, as well as an integral cost of state and control variations from zero during the period. The integrand, $L[.]$, is called the *Lagrangian* of the cost function.

Choice of the cost function structure is critical to the development of feasible strategies for therapy. The most important element is that central variables like pathogen concentration and organ health be *observable in the cost* [33], as these variables cannot be allowed to reach unacceptable values. It also is important to establish a framework for trading off the therapy dosage level against deviations of system variables from desirable levels. The cost function variables are squared to amplify the effects of large variations and to de-emphasize contributions of small variations. Each squared element is multiplied by a coefficient (p_{ii}, q_{ii}, or r_{ii}) that establishes the relative importance of the factor in the treatment cost. These coefficients could reflect financial cost associated with each variation, or they could represent physiological 'cost', such as side effects, toxicity, or discomfort. The weighting coefficients manage the maximum values of state and control variations and the speed of response in bringing the pathogenic attack under control.

The necessary conditions for optimality are expressed by the three Euler–Lagrange equations [33]. These equations include a linear, ordinary-differential equation whose integral is $\lambda(t)$, a

terminal boundary condition that specifies $\boldsymbol{\lambda}(t_f)$, and a stationarity condition on the control:

$$\dot{\boldsymbol{\lambda}}(t) = -\left\{\frac{\partial H[\mathbf{x}(t), \mathbf{u}(t), \boldsymbol{\lambda}(t), t]}{\partial \mathbf{x}}\right\}^{\mathrm{T}}$$

$$\boldsymbol{\lambda}(t_f) = \left\{\frac{\partial \phi[x(t_f)]}{\partial \mathbf{x}}\right\}^{\mathrm{T}}$$

$$0 = \frac{\partial H[\mathbf{x}(t), \mathbf{u}(t), \boldsymbol{\lambda}(t), t]}{\partial \mathbf{u}} \tag{10-12}$$

They are expressed using the *Hamiltonian* of the system,

$$H(\mathbf{x}, \mathbf{u}, \boldsymbol{\lambda}, t) = L(\mathbf{x}, \mathbf{u}, t) + \boldsymbol{\lambda}^{\mathrm{T}}\mathbf{f} \tag{13}$$

which adjoins the system dynamics to the cost function Lagrangian through the transpose of the adjoint vector, $\boldsymbol{\lambda}$.

Because the mathematical model for disease dynamics is non-linear, there is no closed-form solution for the Euler–Lagrange equations. Together with the system dynamics (Equation (6)), they must be solved numerically. The solution must be iterative because this is a two-point boundary-value problem: initial conditions are specified for the state, $\mathbf{x}(t_o)$, while terminal conditions are given for the adjoint vector, $\boldsymbol{\lambda}(t_f)$.

There are numerous ways of solving this problem [33]. Here, we have used the 'steepest-descent' gradient method. The method is based on successive approximations of the control history $\mathbf{u}(t)$. Starting with an initial guess for a control history $\mathbf{u}_o(t)$, the state dynamics (Equation (6)) are integrated forward. The adjoint vector $\boldsymbol{\lambda}(t)$ is then integrated backwards with boundary condition expressed at the final time $\boldsymbol{\lambda}(t_f)$ by Equation (11). In general, because $\mathbf{u}_{k-1}$ is only an approximation of the optimal control history, the necessary optimality condition (Equation (12)) is not satisfied, so that $\partial H/\partial \mathbf{u}|_{k-1} \neq 0$. In order to obtain a better estimate, the control history is updated with

$$\mathbf{u}_k = \mathbf{u}_{k-1} - \varepsilon \frac{\partial H}{\partial \mathbf{u}}\bigg|_{k-1} \tag{14}$$

where ε is a small positive constant. For k sufficiently large, $\partial H/\partial \mathbf{u}$ tends to zero, and $\mathbf{u}(t)$ converges to the optimal control history. As shown in Figure 2, this method provides a minimal cost. There are more technical details related to the presence of hard constraints ($\mathbf{x} \geqslant \mathbf{0}$, $\mathbf{u} \geqslant \mathbf{0}$, and $x_4 \leqslant 1$) which we omit for conciseness. In this particular case, the results are similar even without explicitly integrating the constraints. This is partly due to the fact that the uncontrolled dynamics are well defined, i.e. the state is never less than zero on its own. We note that steepest-descent algorithms and quadratic cost functions for systems with smooth non-linearities generally have good convergence properties, as is demonstrated in what follows.

In the remainder of the paper, we present the results of several system optimizations. Our principal goal is to reveal the qualitative effects of cost function weighting, the choice of therapeutic agents, and the period of treatment on the etiology of the disease and the therapeutic protocol. As a starting point for all calculations, we assume initial conditions that would prove fatal without treatment (Case 4). In the process, we demonstrate the substantial power and flexibility of optimal control theory for suggesting improved procedures for treating a disease.
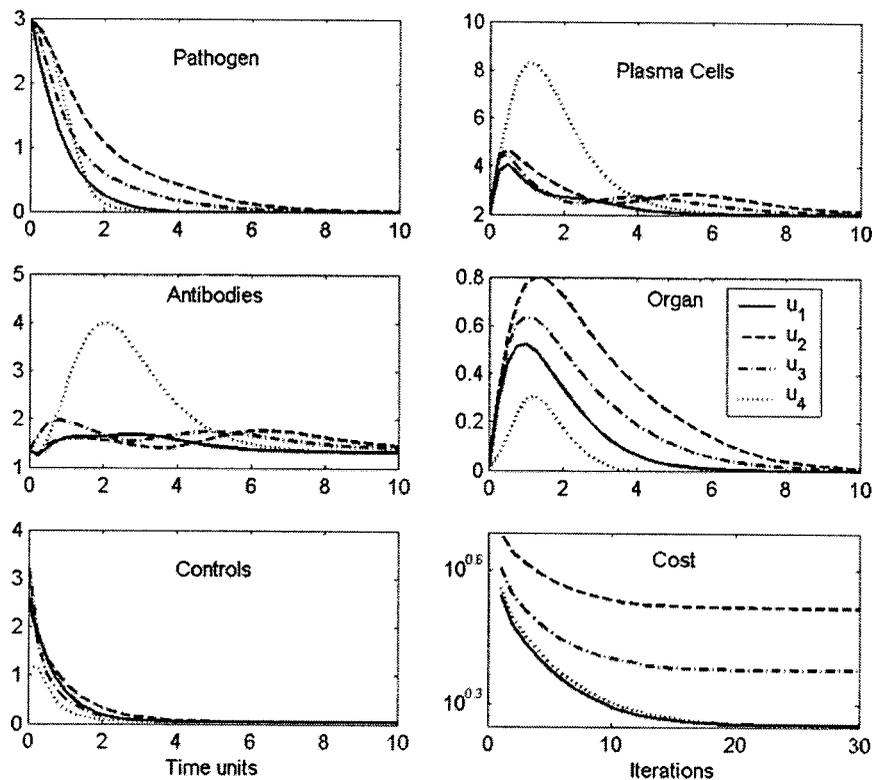
Figure 2. Optimal therapies with unit cost-function weights and scalar controls.

OPTIMAL THERAPIES AND DISEASE ETIOLOGIES

Baseline cases

We begin by considering the effects of each control variable applied and optimized separately, using a reference set of costs. All p_{ii} and q_{ii} are one; r_{jj} is one for the active control and zero for the inactive controls. The treatment interval is 10 time units, which is on the order of the transient time associated with the uncontrolled pathogenic attack (Figure 1). Numerical optimizations are stopped when little further improvement is apparent. The first baseline case assumes that only the pathogen killer (u_1) is applied, the second that only the plasma-cell enhancer (u_2) is applied, and so on.

For the disease dynamics model used here, the only way to decrease the pathogen concentration is to attack it by means of pathogen killer drugs or antibodies; however, the single-control optimizations show that there are two qualitatively very different ways to achieve this result (Figure 2). The direct way consists in relying only on the pathogen killer to do the job, thereby also allowing the organ to heal naturally. On the other hand, histories of u_2, u_3, and u_4 demonstrate indirect ways to fight the pathogen. Using an organ enhancer provides the least direct action. It helps the organ stay healthy, thus enabling the immune system to do its job
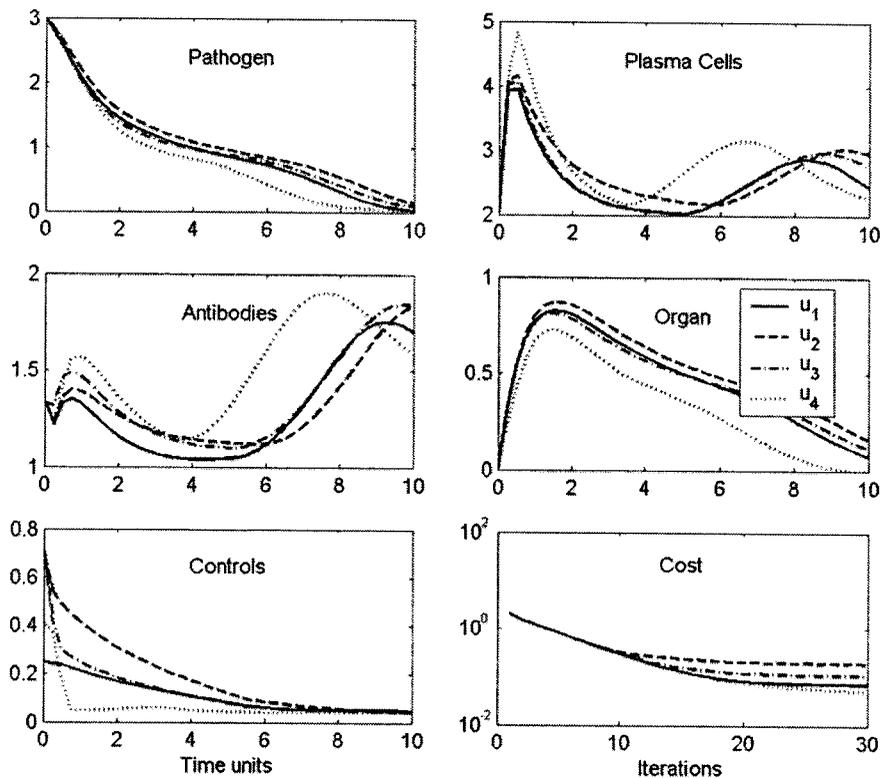
Figure 3. Optimal therapies with integral state weights = 0.01 and scalar controls.

properly. As a consequence, very high concentrations of plasma cells and antibodies are observed.

The pathogen killer, u_1, generates the smallest cost. The control history that produces this result has maximum value at the initial time (3 units), with an exponential decay until the pathogen is eliminated, at which point the control becomes zero. Plasma cell and antibody enhancers (u_2 and u_3) have similar effects, with u_2 being the least efficient because temporary organ damage is as high as 80%. Neither control completely restores the steady-state concentration of plasma cells and antibodies after 10 time units. The treatment histories of u_1, u_2, and u_3 are very similar, presenting an exponential decay. The organ enhancer, u_4, has a longer, initial dosage (1 unit) followed by an exponential decay.

Although not presented in the paper, we conducted optimizations using all of the therapeutic agents at once. With equal control weighting, the effect of the pathogen killer is dominant, and the computed result is nearly the same as that of the pathogen killer alone. From a clinical perspective, all of these therapies are successful in arresting the explosive growth of the pathogen and preventing organ death, although the costs of treatment and the time to achieve results vary. Actual response characteristics are the result not only of the disease dynamic model but also of the cost function weights. Thus, the cost and time to eliminate the pathogen can be adjusted by an appropriate choice of weights, as illustrated in the subsequent cases.

Decreased integral state weighting

Decreasing the integral state weights, q_{11} and q_{44} to 0.01 without changing the other weights allows less control to be used, which slows the rate of pathogen clearance. The results of the optimization for each of the four scalar controls are plotted in Figure 3. The pathogen killer and organ enhancer continue to be the most effective treatments. In this case, the latter is even more effective but requires higher doses for the first 10% of the treatment, with a peak of 0.4 units and then an almost constant dose of 0.05 units for the remaining 90% of the period. The pathogen killer would require instead slowly decaying doses from 0.3 to 0.05 units. The decrease in the initial concentrations of the administered drug with respect to the baseline cases (Figure 2) range from 50% for u_4 to nearly 90% for u_1.

The weight distribution used for this set of optimizations helps balance possible undesirable drug side effects or high monetary costs against promptness of recovery, penalizing high control values more than high state values during optimization.

Increased integral state weighting

Raising the integral state weights, q_{11} and q_{44} to 100, while keeping the others fixed is appropriate if high drug doses are of smaller concern. The optimization results show that the most effective and less costly treatment in this case is the pathogen killer (Figure 4). There is a
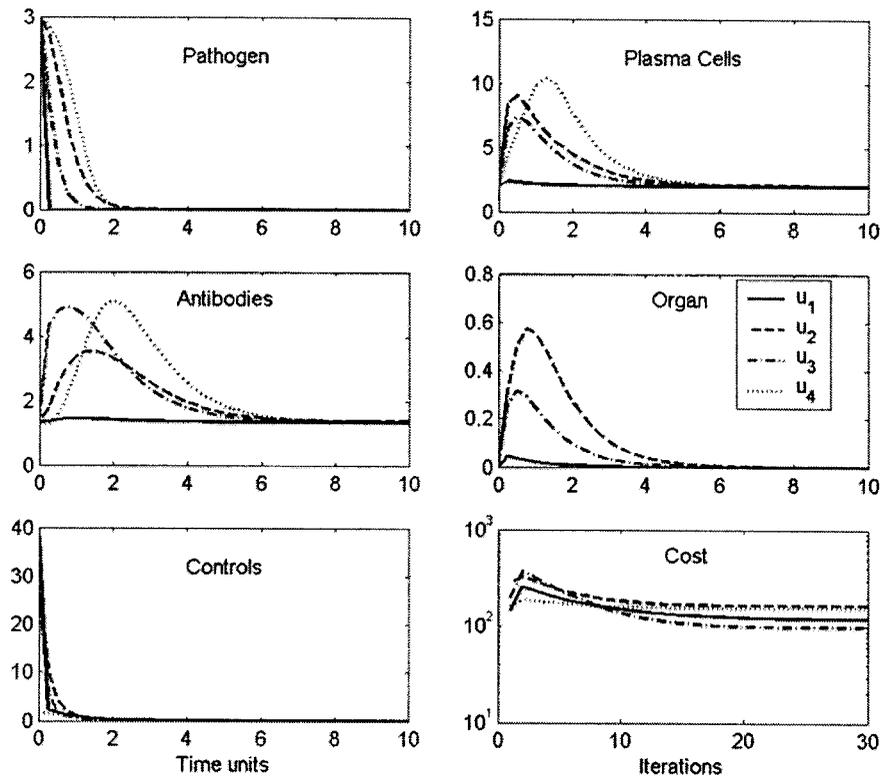

Figure 4. Optimal therapies with integral state weights = 100 and scalar controls.

five-fold increase in pathogen clearance speed at an expense of about 10 times higher dosages of u_1, u_2, and u_3 compared to the baseline cases. The dosage of u_4 used does not increase in the same proportion, but it is strong enough to foster strong response of the natural system while protecting the organ against damage by the pathogen. It is obvious that a further increase in the organ enhancer concentration would only result in higher control costs without therapeutic improvement.

When u_2 or u_3 are used as controls, plasma cell and antibody histories show concentrations comparable to what is obtained by using u_4, though the time to complete recovery is smaller. It illustrates the fact that antibodies are critical when fighting the pathogen through immune-system-generated responses. Thus the drug that generates the earliest peak in antibody concentration obtains the fastest response. On the other hand, when using u_1, the pathogen is eliminated before any significant immune response has time to develop.

Effects of treatment duration and terminal cost weights

The formulation of the optimal treatment as in Equation (8) has no particular restriction on the treatment time interval $(t_f - t_i)$. This means that we can modify the time interval at will to see
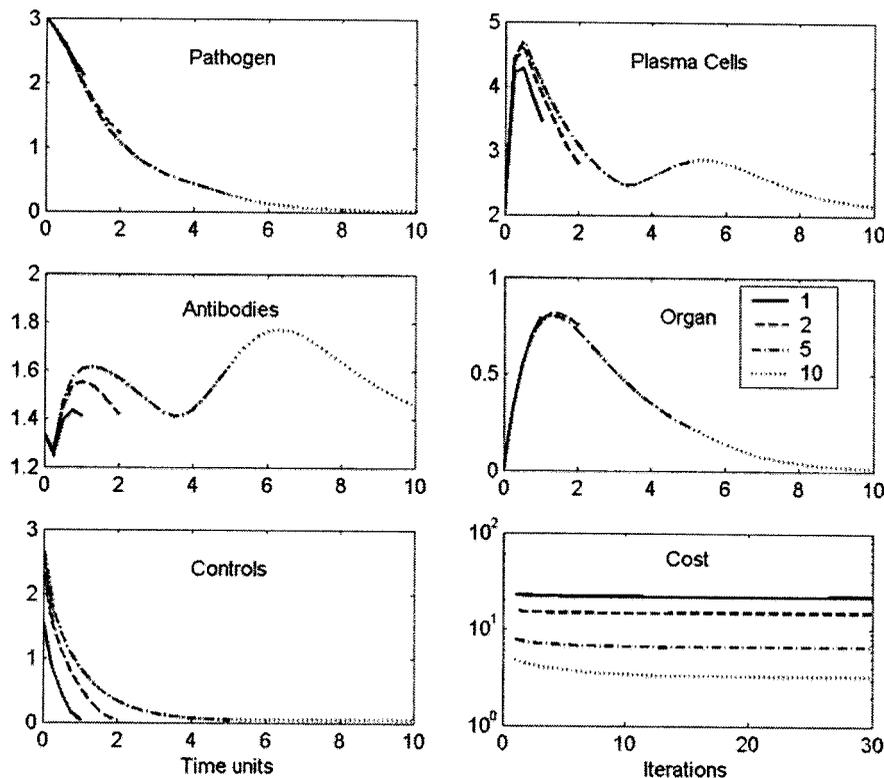

Figure 5. Effects of period of therapy with $p_{ii} = 1$. Legend indicates number of time units from start to finish of therapy.

what effect treatment duration has and how it influences the optimal treatment history. Examples are given in Figures 5 and 6, where we have used the plasma cell enhancer, u_2, generally the least effective control, to treat a hypothetical patient. In the first case (Figure 5), where all cost weights are one, we see that if the treatment is too short (1 or 2 time units), the pathogen concentrations are still high at the end of the treatment, and a relapse would occur. Forcing the final values to be closer to the target values by increasing terminal cost weights ($p_{ii} = 100$) produces higher therapeutic dosage (Figure 6). There seems to exist a natural optimal treatment length; forcing the system to respond more quickly may require significantly higher doses, whereas longer treatments are unnecessary.

There are no differences between both 10-time-unit optimizations using p_{ii} equal to 1 and 100, apart from axis scaling of the figures. Both pathogen concentration and organ health reach zero before the end of the simulation, illustrating that the terminal cost has no significant effect for this model if the treatment time is long enough.
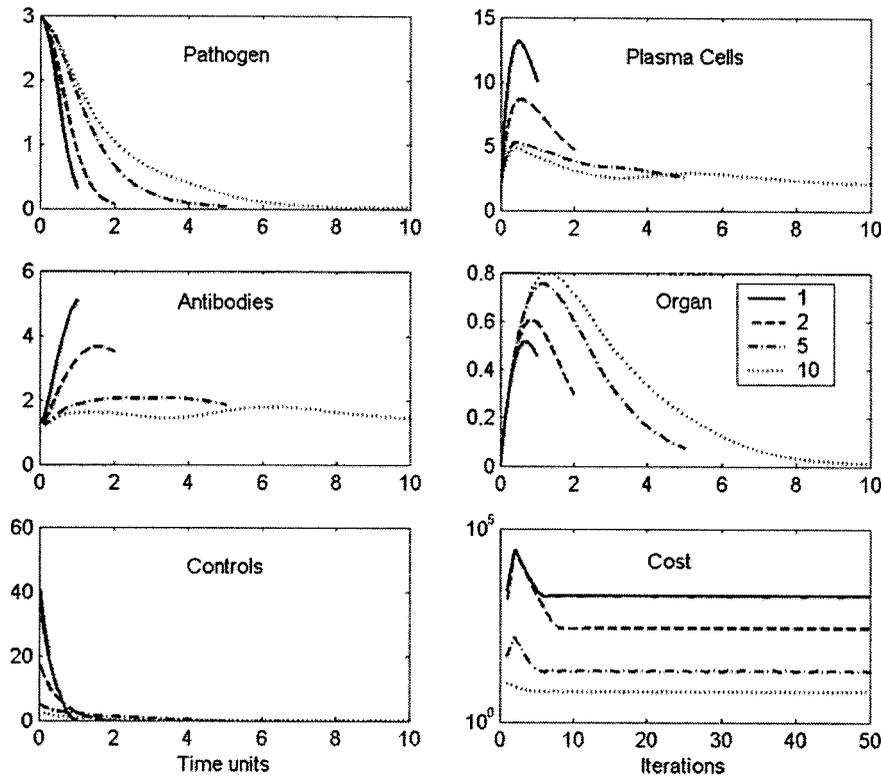

Figure 6. Effect of period of therapy with $p_{ii} = 100$. Legend indicates number of time units from start to finish of therapy.

CONCLUSION

A method for computing optimal therapies to counteract disease is presented, and it is applied to a generic model of pathogenic attack on the innate immune system. The dynamic state of the system describes concentrations of the pathogen, plasma cells, and antibodies, as well as a quantified description of organ health. After showing the natural response of the system to a range of initial pathogen concentrations, the effects of optimal therapies based on four different treatment approaches are demonstrated. Without therapy, the microbial assault would be lethal; with therapy, the assault can be contained and organ health can be preserved. Tradeoffs between treatment dosage and immune system response are examined by altering the weights in a quadratic cost function that penalizes both integral and terminal levels of the dynamic state and control. A critical finding is that successful yet distinctly different approaches to therapy can be identified, either through agents that attack the pathogen directly or by enhancing the natural immune system's defensive capability.

Optimal control theory can help develop clinical insight in monitoring and treating illness. The flexibility and relative simplicity of this technique can lead to improved therapy for individuals, whose unique characteristics can be taken into account when establishing treatment protocols. While this control-theoretic approach to developing novel therapeutic protocols has been illustrated using a simple mathematical model, its application to more precise representations of actual disease dynamics is apparent.

REFERENCES

1. Lydyard PM, Whelan A, Fanger MW. *Instant Notes in Immunology*. Springer-Verlag: New York, 2000.
2. Janeway CA, Travers P, Walport M, Shlomchik M. *Immunobiology*. Garland: London, 2001.
3. Thain M, Hickman M. *The Penguin Dictionary of Biology*. Penguin Books: London, 2000.
4. Asachenkov A, Marchuk G, Mohler R, Zuev S. *Disease Dynamics*. Birkhauser: Boston, 1994.
5. Rundell A, HogenEsch H, DeCarlo R. Enhanced modeling of the immune system to incorporate natural killer cells and memory. *Proceedings of the 1995 American Control Conference*, Seattle, June 1995; 255–259.
6. Perelson AS. Immunology for physicists. *Reviews of Modern Physics* 1997; **69**(4):1219–1267.
7. Nowak MA, May RM. *Viral Dynamics: Mathematical Principles of Immunology and Virology*. Oxford University Press: Oxford, UK, 2000.
8. Nowak MA, May RM, Phillips RE, Rowland-Jones S, Lalloo DG, McAdam S, Klenerman P, Köppe B, Sigmund K, Bangham CRM, McMichael AJ. Pathogenic oscillations and shifting immunodominance in HIV-1 infections. *Nature* 1995; **375**:606–611.
9. Perelson AS, Neumann AV, Markowitz M, Leonard JM, Ho DD. HIV-1 dynamics in vivo: virion clearance rate, infected cell lifespan, and pathogenic generation time. *Science* 1996; **271**:1582–1586.
10. Perelson AS, Nelson PW. Mathematical analysis of HIV-1 dynamics in vivo. *SIAM Review* 1999; **41**(1):3–44.
11. Wodarz D, May RM, Nowak MA. The role of pathogen-independent persistence of memory cytotoxic T lymphocytes. *International Immunology* 2000; **12**(4):467–477.
12. Stafford MA, Cao Y, Ho DD, Corey L, Perelson AS. Modeling plasma pathogen concentration and CD4 + T cell kinetics during primary HIV infection. *Journal of Theoretical Biology* 2000; **203**:285–301.
13. Wodarz D, Nowak MA. CD8 memory, immunodominance, and pathogenic escape. *European Journal of Immunology* 2000; **30**:2704–2712.

14. Wiener N. *Cybernetics: or Control and Communication in the Animal and the Machine*. Technology Press: Cambridge, 1948.
15. Bellman RE. *Mathematical Methods in Medicine*. World Scientific: Singapore, 1983.
16. Perelson AS, Mirmirani M, Oster GF. Optimal strategies in immunology, I: B-cell differentiation and proliferation. *Journal of Mathmatical Biology* 1976; **3**:325–367.
17. Perelson AS, Mirmirani M, Oster GF. Optimal strategies in immunology, II: B-memory cell production. *Journal of Mathematical Biology* 1978; **5**:213–256.
18. Perelson AS, Goldstein B, Rocklin S. Optimal strategies in immunology, III: the IgM–IgG switch. *Journal of Mathematical Biology* 1980; **10**:209–256.
19. Perelson AS. Applications of optimal control theory to immunology. In *Recent Developments in Variable Structure Systems, Economics and Biology*, Mohler R, Ruberti A (eds). Springer-Verlag: New York, 1978; 272–287.
20. Swan GW. Optimal control applications in biomedical engineering — a survey. *Optimal Control Applications and Methods* 1981; **2**(4):311–334.
21. Swan GW. *Applications of Optimal Control Theory in Medicine*. Marcel Dekker: New York, 1984.
22. Swan GW. *Optimization of Human Cancer Radiotherapy*. Lecture Notes in Biomathematics, vol. 42. Springer-Verlag: New York, 1981.
23. Murray JM. Optimal control for a cancer chemotherapy problem with general growth and loss functions. *Mathematical Biosciences* 1990; **98**:273–287.
24. Murray JM. Some optimal control problems in cancer chemotherapy with a toxicity limit. *Mathematical Biosciences* 1990; **100**:49–67.
25. Novak A, Feichtinger G. Optimal treatment of cancer diseases. *International Journal of Systems Science* 1993; **24**(7):1253–1263.
26. Fister KR, Panetta JC. Optimal control applied to cell-cycle-specific cancer chemotherapy. *SIAM Journal of Applied Mathematics* 2000; **60**(3):1059–1072.
27. Kirschner D, Lenhart S, Serbin S. Optimal control of the chemotherapy of HIV. *Journal of Mathematical Biology* 1997; **35**:775–792.
28. Wein LM, Zenios SA, Nowak MA. Dynamic multidrug therapies for HIV: a control theoretic approach. *Journal of Theoretical Biology* 1997; **185**:15–29.
29. Bonhoeffer S, May RM, Shaw GM, Nowak MA. Viral dynamics and drug therapy. *Proceedings of the National Academy of Sciences USA* 1997; **94**:6971–6976.
30. Wein LM, D'Amato RM, Perelson AS. Mathematical analysis of antiretropathogenic therapy aimed at HIV-1 eradication or maintenance of low pathogenic loads. *Journal of Theoretical Biology* 1998; **192**:81–98.
31. Wodarz D, Nowak MA. Specific therapy regimes could lead to long-term immunological control of HIV. *Proceedings of the National Academy of Sciences USA* 1999; **96**(25):14,464–14,469.
32. Wodarz D, Page KM, Arnaout RA, Thomsen AR, Lifson JD, Nowak MA. A new theory of cytotoxic T-lymphocyte memory: implications for HIV treatment. *Philosophical Transactions of the Royal Society* B 2000; **355**:329–343.
33. Stengel RF. *Optimal Control and Estimation*. Dover Publications: New York, 1994 (originally published as *Stochastic Optimal Control: Theory and Application*. Wiley: New York, 1986).