Bézier Control Parameterization for Evolutionary Optimization in Disease Models

Tim Rogalsky
Department of Mathematics
Canadian Mennonite University
500 Shaftesbury Blvd
Winnipeg, MB. R3P 2N2
204.487.3300
trogalsky@cmu.ca

ABSTRACT
In many disease models, the dynamics are described by a system of differential equations. When the spread of the disease is controlled by a treatment strategy, an obvious challenge is to find the best treatment possible. Mathematically, this problem is known as optimal control, or dynamic optimization. To solve these problems, researchers are increasingly turning to evolutionary optimization methods. Evolutionary computation, however, operates on discrete, n-dimensional vectors, not on continuous functions, and becomes computationally unmanageable for large n. Thus a parameterization technique is required, that can represent arbitrary functions with a small number of parameters. The typical approach to parameterization in epidemiological and biomedical models is to approximate the control functions as piecewise constant. We show the limitations of this approach, and demonstrate a recently developed method, Bézier Control Parameterization (BCP). With relatively few parameters, BCP can represent continuous control functions, and provides an efficient and effective parameterization method for evolutionary control of disease models.

Categories and Subject Descriptors
J.3 [**Computer Applications**]: Life and Medical Sciences – *biology and genetics, health.*

General Terms
Algorithms.

Keywords
Control vector parameterization, Differential equations, Evolutionary algorithms, Mathematical biology, Optimal control.

1. INTRODUCTION
The spread of a disease can often be modeled by a dynamical system of ordinary differential equations for a set of dependent functions, $\mathbf{x}(t)$. Treatment strategies can then be modeled as independent control functions, $\mathbf{u}(t)$, that control the spread of the disease. These might be public health strategies in an epidemiological model, or drug intervention rates in a biomedical model. With each control comes an associated cost, which might be financial, but might also be a side effect of the treatment. Practitioners are thus highly interested in optimal treatment strategies – ones that maximize healing, while minimizing cost.

This type of problem is known as optimal control, or sometimes, dynamic optimization. Mathematically, the problem can be stated as follows:

$$\min_{\mathbf{u}} F(\mathbf{u}) = \int_{t_0}^{t_f} f(t, \mathbf{x}(t), \mathbf{u}(t))dt, \quad \text{subject to } \begin{cases} \mathbf{x}'(t) = g(t, \mathbf{x}(t), \mathbf{u}(t)) \\ \mathbf{x}(t_0) = \mathbf{x}_0 \end{cases}, \tag{1}$$

where $F(\mathbf{u})$ is the objective function, t_0 and t_f are the initial and final times, and f and g depend on the particular model. The dependent functions $\mathbf{x}(t)$ are known as state functions, and the independent $\mathbf{u}(t)$ as control functions.

In a classic SIR model of an epidemic, for example the state functions are the populations, at time t, of those who are susceptible to the disease, those who are immune from the disease, and those who are recovered from the disease. Control functions might include a vaccination rate and a quarantine rate, both functions of time. An optimal control objective would then be to find a public health policy, in terms of vaccination and quarantine rates, that minimizes both the number of infectious persons, and the cost of implementing the policy.

In this paper, we use an evolutionary direct method to solve optimal control problems arising from disease models. Section 2 of the paper introduces evolutionary direct methods, and describes a recently developed approach, Bézier Control Parameterization (BCP) [1]. Some related work is described and compared to the BCP approach. Part 3 demonstrates the effectiveness of the method for various disease models, including optimal control of a cancerous tumour by chemotherapy, optimal control of an epidemic through vaccination, and optimal control of HIV through reverse transcription inhibitor therapy. In part 4 we draw conclusions and look ahead to future implementations and applications.

GECCO'12 Companion, July 7–11, 2012, Philadelphia, PA, USA.

2. EVOLUTIONARY DIRECT METHODS

2.1 Direct methods for optimal control

There are two general approaches to optimal control. These are often labeled as direct and indirect methods. An indirect method, using calculus of variations techniques, transforms a given optimal control problem into a boundary value problem, which can then be solved analytically or numerically using well-known techniques for differential equations. An excellent introduction to this method, for biological models, can be found in a text by Suzanne Lenhart and John Workman [2].

While the indirect method is a powerful tool mathematically, with existence and uniqueness results, exact solutions, or error estimates, it is also limited to solving only a certain class of problems. Increasingly, then, researchers are turning to direct methods, which is a more universal approach.

In a direct method, optimal control is seen as a standard optimization problem: perform a search for the control function $\mathbf{u}(t)$ that optimizes the objective functional. A solution requires several subroutines: a global optimizer, a parameterization strategy to represent control functions, and a numerical technique for solving the dynamical system.

2.2 Evolutionary optimization

Evolutionary Algorithms (EAs) are powerful, global optimizers that treat optimization from the perspective of natural evolution: an initial population of feasible solutions evolves into a population of globally near-optimal solutions. There are typically two mechanisms by which new feasible solutions are formed: mutation (small perturbations in a binary- or real-valued individual) and recombination (combining the characteristics of two different individuals). Some form of natural selection is used to decide which population members "survive" to the next generation, and after many generations the population converges, to one or several near-optimal solutions.

There are two types of EA, distinguished by the way in which they represent individual feasible solutions. Genetic Algorithms (GAs) [3] use binary representation, and are thus suitable for discrete or integer optimization problems. Evolutionary Strategies (ESs) [4] use real-valued vectors, and are better suited for the kind of continuous parameter optimization required for optimal control.

Differential Evolution (DE) [5] is the evolutionary optimizer used in this paper. DE emerged in the 1990s as one of the most impressive ESs, converging faster and with more certainty than many other acclaimed global optimization methods [6]. In the years since, it has successfully been used in many different applied fields [5]. DE has been shown to be a robust and efficient global optimizer for direct methods in optimal control, [7]-[11].

The crucial difference between DE and other ESs lies in mutation. ESs normally use predetermined probability distribution functions to perturb vectors, leaving them unable to adapt the perturbation magnitude to the topology of the objective function. DE uses the "differential" of two randomly chosen population vectors, $\mathbf{u}_a$ and $\mathbf{u}_b$, to perturb a base vector $\mathbf{u}_c$, $\mathbf{u}_{new} = \mathbf{u}_c + F(\mathbf{u}_a - \mathbf{u}_b)$, where F is the differential weight. The perturbation magnitude is thus automatically appropriate to the given landscape, and the search is less random, being dictated by the shape of the objective function itself. This property of DE is known as self-organization. Ultimately, it results in better convergence properties as the algorithm nears the global minimum.

The specific DE strategy used here is known as DE/local-to-best/1. In this strategy, the base vector $\mathbf{u}_c$ is a combination of one "local" vector chosen at random from the population, and the "best" vector so far – the one with the lowest objective function value. $F = 0.85$ is the recommended differential weight. This strategy tends to balance robustness with fast convergence, and has been demonstrated as one of the more effective DE strategies [12]. Usually a population size of $NP = 10D$ is effective, where D is the dimension of the vector $\mathbf{u}$. Occasionally, when misconvergence occurs, NP needs to be increased.

2.3 Control Vector Parameterization

In a direct approach to optimal control, the optimization algorithm searches for real-valued functions, $\mathbf{u}(t)$ that best meet the objective. These search algorithms, however, typically operate on n-dimensional vectors, not on infinite dimensional function spaces. Thus, a parameterization strategy is required, by which control functions can be represented as $\mathbf{R}^n$ vectors. This is known as Control Vector Parameterization (CVP). A wide variety of CVPs have been used with non-evolutionary optimizers, including piecewise constant [13], Chebyshev polynomials [14], Lagrange polynomials [15], and piecewise Lagrange polynomials [16].

Evolutionary direct methods have been less creative. Most simply discretize the control function space, so that control functions are in fact piecewise constant. The reason for this implementation may be that it is the easiest parameterization to encode, or it may be that current researchers are simply following the path trod by those who first applied EAs to optimal control [17], [18]. In any case, there is room for improvement.

For situations in which control functions are actually continuous, a piecewise constant parameterization is highly inadequate. On the one hand, a very high number of parameters is needed for an accurate approximation of the continuous functions. On the other, EAs are computationally expensive, and require a small number of parameters to converge to a near-optimal solution within a reasonable amount of time. In other words, using a large number of parameters will result in slow convergence to an accurate approximation of the true solution, while a small number will result in quick convergence to a poor approximation. A more creative CVP is desirable for evolutionary direct methods.

2.4 Bézier Control Parameterization

To be effective, the CVP should be able closely to approximate arbitrary, continuous, control functions. To be efficient, it must do so with a relatively small number of parameters. Also, CVPs that increase the nonlinearity of the objective function can lead to epistasis [19] – the nonlinear and interdependent manner in which the objective function relates to the design parameters. Small changes in several variables can result in large changes in the objective function. Epistatic functions can lead to premature convergence, because they provide so few clues as to the location of the global minimum. In general, a reduction of this nonlinear interaction, by having parameters more directly linked to the objective function, will enable the optimizer to converge more quickly.

Bézier Control Parameterization (BCP) [1] is a CVP designed to parameterize continuous functions with minimal parameters and minimal epistasis. Bézier curves are common in engineering

applications, and have been used effectively with DE to optimize turbomachinery airfoils [20].

An nth order Bézier curve, $\mathbf{P}(z)$, is defined parametrically using n+1 two-dimensional control points $\mathbf{P}_i(t_i, u_i)$, as follows:

$$\mathbf{P}(z) = \sum_{i=0}^{n} \mathbf{P}_i \frac{n!}{i!(n-i)!} z^i (1-z)^{n-i}, \ 0 \le z \le 1, \qquad (2)$$

where z is the parameter. Several Bézier curve properties make this parameterization fit naturally within an optimal control strategy. Curves begin at control point $\mathbf{P}_0$, end at control point $\mathbf{P}_n$, have initial slope equal to that of the line segment $\mathbf{P}_0\mathbf{P}_1$, ending slope equal to that of $\mathbf{P}_{n-1}\mathbf{P}_n$, and always lie within the convex hull formed by the control points. The curve is nth order continuous throughout and never oscillates wildly away from its defining control points. Thus Bézier curves can parameterize smooth, non-oscillatory functions, with minimal epistasis, using only a few parameters.

The Bézier Control Parameterization (BCP) introduced in [1] is designed for a single control function. A fixed, regular mesh is used on the t-axis. This forces the curve to be single-valued, and also reduces the dimension of the optimization vectors to n + 1. That is, the BCP $\mathbf{u} = \left[u_i\right]_{i=0}^{n}$ completely encodes a control function $u(t)$ as the nth order parametric Bézier curve $u(t) = \langle t(z), u(z) \rangle$, as follows:

$$\left\{ \begin{array}{l} t(z) = \sum_{i=0}^{n} \left(t_0 + i\Delta t\right) \frac{n!}{i!(n-i)!} z^i (1-z)^{n-i} \\ u(z) = \sum_{i=0}^{n} u_i \frac{n!}{i!(n-i)!} z^i (1-z)^{n-i} \end{array} \right\}, \ 0 \le z \le 1, \qquad (3)$$

where $\Delta t = (t_f - t_0)/n$, t_0 is the initial time, and t_f is the final time.

The objective function, $F(\mathbf{u})$, is computed as follows. The control function $u(t)$, is found using the Bézier curve parameterization. It is stored as a set of data points, at parameters $z = 0, h, 2h, \ldots, 1$. A step-size of $h = 0.01$ is used here, and can be refined when more accuracy is required. The IVP is then solved numerically for $x(t)$, interpolating the data points to approximate $u(t)$ as necessary. The differential equation solver used is MATLAB's ode45 function, an explicit Runge-Kutta (4,5) formula, with the Dormand-Prince pair. Finally, the objective integral is evaluated, again interpolating to approximate $x(t)$ and $u(t)$, as necessary. The numerical integration routine is MATLAB's quad function, a recursive adaptive Simpson quadrature. The value of the integral is the "cost" F of the vector $\mathbf{u}$. The cost is then minimized by DE.

2.5 Related Work

The BCP method is designed to have advantages over other work in the field. Using only a few parameters, it can approximate arbitrary continuous control functions with excellent precision, while keeping epistatic interactions to a minimum. Contrast this with the following evolutionary approaches to optimal control models.

Chiou and Wang [8] use a hybrid DE algorithm for optimal control of a fed-batch fermentation process for ethanol production. The goal is to find the optimal glucose feed flow rate, as a function of time, such that the ethanol production rate is maximized. The optimal feed flow rate is found by discretizing a 24 hour period into 20 equal time segments of 1.2 h each. In other words, a separate constant flow rate is found for each 1.2 h segment of time. Optimization is then performed on these 20 parameters. Their work is extended by Kapadi and Gudi [11], who use a "non-uniform control vector parameterization." This is still piecewise constant, but permits variable time intervals. This more flexible CVP does in fact result in an improved solution, with a higher ethanol concentration produced in less time. However, to accomplish this they have had to double the number of parameters, to 40. Compare this to the BCP results below, section 3, in which optimal, continuous functions are found using under 10 parameters.

Similarly, Lopez-Cruz et al [10] investigate optimal control of nitrate in lettuce growth using DE as the optimizer. The objective is to minimize the potentially harmful nitrate concentration in lettuce, and also minimize the amount of artificial greenhouse light required, while maintaining a certain lettuce head size over a fixed growing period. The controls are light intensity, carbon dioxide concentration, and temperature. These three potentially continuous functions of time were approximated as piecewise constant, with 20 equally-spaced time intervals, resulting in an optimization problem with 60 parameters. Note that the BCP method developed currently can parameterize only one continuous control function. In future work, BCP will be extended to permit the use of multiple controls.

In medical applications, evolutionary direct methods have been used to control SARS, cancer, and HIV. Yan and Zou use a GA for optimal control of a SARS epidemic [21]. Two control functions are used, a quarantine rate for those exposed to SARS but asymptomatic, and an isolation rate for people displaying symptoms of the disease. The CVP for each control is piecewise constant, in three pieces: One rate for the first 60 days of the disease, one for the next 240 days, and a third for the last 60 days. Yan and Zou also solve the optimal control problem indirectly, permitting them to compare their piecewise constant approach with a much more precise solution. They find that their direct method controls the disease nearly as well as the indirect method, but at the higher cost of quarantining and isolating more people. This higher public health cost can potentially be eliminated by use of BCP, and this will be investigated in future work.

In collaborative work with an oncologist, Liang et al [22] use a GA to find optimal chemotherapy schedules for treating cancer. The control function is piecewise constant, with a fixed drug dosage per day. An 84-day treatment is considered, representing 84 parameters to be optimized. For practical reasons, oncologists may not in fact schedule chemotherapy on a continuous basis over an 84-day period. However, in this case a continuous BCP solution could be used to select daily dosages, with the advantage of improved algorithm efficiency.

It should be noted that not all medical applications have continuous control functions. For example, Neri et al [23] use an EA to design optimal multidrug Structured Treatment Interruption (STI) therapies for HIV. STI therapies are on-off – the patient receives either the maximum dosage or none at all – and can potentially lower the risk of HIV mutating to drug-resistant strains. This type of control function is known as a bang-bang control. Clearly, when the control is bang-bang, piecewise constant CVP is precisely the correct parameterization to use, and we will not consider this situation.

3. RESULTS

The effectiveness of the BCP parameterization is demonstrated by applying it to three optimal control models for diseases: optimal chemotherapy treatment for controlling a cancerous tumour, optimal vaccination schedule for controlling an epidemic, and optimal Reverse Transcription Inhibitor therapy for controlling HIV. These examples have been solved numerically by the indirect method in [2], permitting a comparison with the BCP direct method, for validation of this approach.

3.1 Optimal chemotherapy regimen for cancer control

Cancerous cells are those that can no longer regulate their cell growth. Due to their uncontrolled growth, these cells form malignant tumours that grow quickly and spread throughout the body. Chemotherapeutic drugs are designed to kill rapidly proliferating cells. Unfortunately, they do not typically distinguish between normal and abnormal growth. Healthy rapid-growth cells are also targeted, such as hair follicle cells, or cells in the digestive tract and bone marrow. In fact, an overly aggressive chemotherapy regimen is as capable of causing death as the cancer itself. Treatment is thus restricted to a fixed time period, and has two important goals: to minimize both the tumour size and the total amount of the drug.

For demonstration purposes, we use the following model, developed in [24], that describes the growth rate of a cancerous tumour in the presence of chemotherapy:

$$N'(t) = rN(t)\ln\left(\frac{1}{N(t)}\right) - u(t)\delta N(t), \qquad N(0) = N_0. \tag{4}$$

The state function, $N(t)$, is the tumour density, normalized to $0 \le N \le 1$. In the absence of any treatment, the model assumes Gompertzian tumour growth, which is exponential for small tumours ($N \approx 0$), but decreases as the tumour increases in size. As the tumour approaches its maximum size ($N \approx 1$), the growth rate slows to zero. The parameter r represents the relative growth rate of the particular cancer. Chemotherapy is modelled with Skipper's log-kill hypothesis, which assumes that cells are killed at a rate proportional to the tumour size. The control function $u(t)$ represents the pharmacokinetic effect of the drug over time, and the parameter δ represents the magnitude of the dose.

Since chemotherapy kills both healthy and unhealthy cells, there are two objectives: first to minimize the tumour density, represented by $N(t)$, and second to minimize the overall toxicity, represented by $u(t)$. These are combined into the single objective functional,

$$F(u) = \int_0^T AN(t)^2 + u(t)^2\, dt\,, \tag{5}$$

where A is the relative weight of the first objective over the second, and T is the duration of the treatment. The integrand in the objective functional (5) is quadratic to make the problem possible to solve by the indirect method. The optimal control problem is then to minimize (5) subject to (4).

This problem has been solved by the indirect method, in [2], for parameters $r = 0.3$, $\delta = 0.45$, $A = 3$, $N_0 = 0.975$, and $T = 20$. Using Pontryagin's Maximum Principle, the problem is converted to a boundary value problem, which is solved numerically in MATLAB. This known solution can then be compared with the BCP results.

The direct method solution uses DE to find the optimal BCP variables for a continuous control function $u(t)$. The specific optimization strategy used is DE/local-to-best/1, with $F = 0.85$ and $CR = 1$. BCP solutions are found for a range of Bézier curve orders, from $n = 3$ to $n = 7$. The dimension of the optimization problem is then $D = n + 1$, and the DE populations are of size $NP = 10(n + 1)$. Initial populations are formed by random selection of control parameters, within the bounds [-5, 5]. Optimization is terminated after 200 generations. The BCP approach is successful in solving this problem. See Figure 1 for a sample solution, comparing the BCP control of degree 5 with the indirect solution. With only 6 optimization parameters, the BCP result is an excellent approximation of the indirect result.

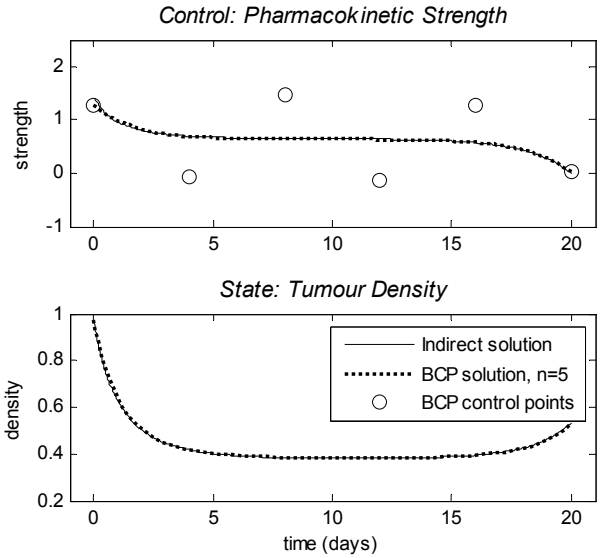

Figure 1. Comparison of solutions for optimal control of cancer by chemotherapy, model (4), (5).

In this particular example, a 20-day round of chemotherapy is prescribed for a very large tumour, 97.5% of its maximum possible size. The size makes tumour reduction more important than the side effects of healthy cell-kill, and a weighting factor of 3 is chosen. Note the pattern of the resulting optimal chemotherapy regimen. The drug strength is high early in the treatment, and drops to zero by day 20. This overall pattern is a fairly standard medical practice. It results in a rapid initial cell-kill for the tumour, without too much damage to the healthy cells, but the drop in chemotherapy strength over the last few days of treatment does allow the tumour to increase slightly. Typically a rest period would be prescribed after this regimen, to permit the body to recover. This would be followed by further rounds of treatment-rest cycles until the tumour is controlled.

Bézier curves with more control points are better able to approximate arbitrary functions. One would thus expect lower order BCP curves to have higher deviations from the indirect solution. This is in fact the case, as shown in Figure 2. The degree 3 curve has the highest deviation, and that deviation decreases monotonically as the degree of the curve is increased to 7. This suggests that DE has indeed closely approximated the

globally optimal solution in each case. That is, for each n, DE has found the degree n Bézier curve that best minimizes the objective. As n is increased, those optimal Bézier curves converge to the actual optimal control.

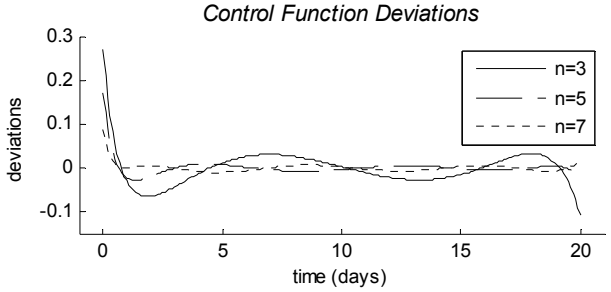

Figure 2. Deviations from the indirect solution diminish as the BCP order is increased.

3.2 Optimal vaccination schedule for epidemic control

Next we consider a model for controlling an epidemic disease. Suppose that immunity can occur either through recovery from the disease, or through vaccination, and that everyone is born susceptible to the disease. Then the dynamics can be modelled with the well-known SEIR system of differential equations:

$$\begin{cases} S'(t) = bN(t) - dS(t) - cS(t)I(t) - u(t)S(t), & S(0) = S_0 \\ E'(t) = cS(t)I(t) - (e+d)E(t), & E(0) = E_0 \\ I'(t) = eE(t) - (g+a+d)I(t), & I(0) = I_0 \\ R'(t) = gI(t) - dR(t) + u(t)S(t), & R(0) = R_0 \end{cases} \tag{6}$$

where the state functions S, E, I, and R represent the numbers of individuals who are Susceptible to, Exposed to, Infected by, and Recovered from (or immune to) the disease, respectively; and N is the total population size, $N = S + E + I + R$, at time t. SEIR models represent diseases with an incubation, or latency, period. Upon being exposed to the disease, and joining the Exposed class, a person remains symptom-free for a time, before becoming infectious and joining the Infected class of individuals.

The real-valued parameters in (6) are as follows: b is the natural birth rate of the population; d the natural death rate; c is the incidence rate of the disease, where the number of people exposed is proportional to the product of those who are susceptible and those who are infected; e is the rate at which those exposed to the disease become symptomatic; and g is the rate at which infectious people recover. The control function $u(t)$ represents the percentage of susceptible individuals being vaccinated per unit time. Since it is unrealistic to expect that all susceptible individuals could be vaccinated, the control function is limited to be less than 90%:

$$0 \le u(t) \le 0.9 \,. \tag{7}$$

When an epidemic disease is spreading through a population, an optimal public health strategy is one that meets two objectives. The primary objective is to minimize the total number of infections, and the secondary is to do so at minimal cost. One way to formulate the two objectives is

$$\min_u F(u) = \int_0^T AI(t) + u(t)^2 \, dt \,, \tag{8}$$

where A is the relative weight of the first objective over the second, and T is the duration of the vaccination program. Again, the quadratic term makes the indirect method transformation possible, mathematically. The optimal control problem is then to minimize $F(u)$ (8), subject to the dynamical system (6), within the bounds (7).

The bounds on the control are imposed as hard constraints. Wherever the control function exceeds the bounds, it is redefined to be the boundary value. That is, the BCP method computes the Bézier curve, $u_{Bez}(t)$, as usual from (3), but the control function itself is defined as

$$u(t) = \begin{cases} u_{Bez}(t), & 0 \le u_{Bez}(t) \le 0.9 \\ 0, & u_{Bez}(t) < 0 \\ 0.9, & u_{Bez}(t) > 0.9 \end{cases} \,. \tag{9}$$

The epidemic control problem is solved for parameters $a = 0.2$, $b = 0.525$, $c = 10^{-4}$, $d = 0.5$, $e = 0.5$, $g = 0.1$, and initial values $S_0 = 1000$, $E_0 = 100$, $I_0 = 50$, $R_0 = 15$. Vaccination duration is $T = 20$ months, and the relative objective weight is $A = 0.1$. For details of the indirect solution, see [2]. The DE strategy used is DE/local-to-best/1, $F = 0.85$, $CR = 1$. BCP solutions were found for curves of order 2, 3, and 4. The dimension of the optimization problem is then $D = n + 1$, and the DE population size used is $NP = 10(n + 1)$. The initial population is formed by random selection of control parameters, within the bounds [-5, 5], and optimization is terminated after 200 generations.

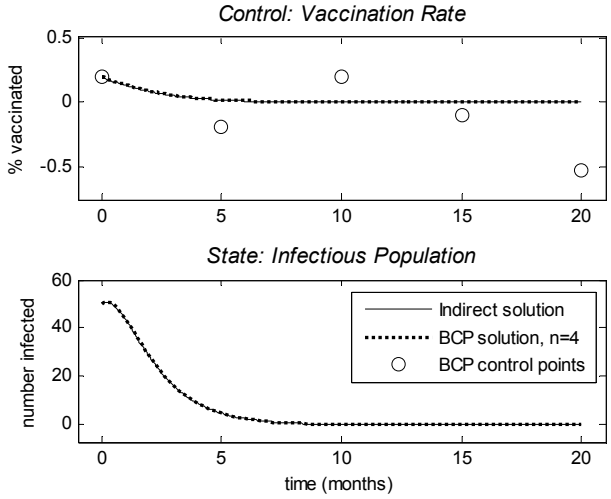

Figure 3. Comparison of solutions for optimal control of an epidemic by vaccination, model (6) - (8).

As shown in Figure 3, BCP again provides an excellent approximation to the numerical inverse solution. In fact, the order 2 curve already agreed reasonably well with the indirect solution, and the degree 4 curve, shown, has only minor deviations. Once again the deviations decreased as the Bézier curve order was increased. Note that the curve shown is the hard constrained control function (9), not the actual Bézier curve.

The particular SEIR model parameters chosen here simulate a disease with a fairly low incidence rate. The optimal vaccination schedule involves an early round of vaccinations, beginning at about 20% in the first month, and then dropping significantly after that. This vaccination strategy protects the susceptible population from the significant numbers who are initially exposed to and infectious with the disease. This results in a rapid control of the epidemic. By month six, the vaccination program has ended, and very few people are infected. By month ten the disease has been eradicated.

3.3 Optimal Reverse Transcription Inhibitor therapy for HIV control

The final case study is a model for controlling the Human Immunodeficiency Virus (HIV) with Reverse Transcription Inhibitor (RTI) therapy. This therapy is based on the mechanism by which HIV attaches to helper T cells, the white blood cells that coordinate our immune system. A free HIV particle contains a single strand of viral RNA, its genetic instructions, and a protein called reverse transcriptase. This free virus particle attaches itself to a healthy helper T cell, injecting its viral RNA and reverse transcriptase. While most cells convert DNA to RNA, in order to communicate with other cells, retroviruses like HIV convert RNA to DNA through reverse transcription. The HIV particle essentially "writes backwards," producing from the RNA instructions a chain of viral DNA. This changes the T cell. Rather than contributing to the immune system, the T cell is recoded to produce more free HIV particles, instead. Thus, if the reverse transcription process can be inhibited, then the spread of the virus can be contained. Reverse transcription inhibitors, including drugs such as AZT, 3TC, d4T, ddc, and ddl, do just that. They block the recoding of viral RNA into DNA.

One model [25] of the interaction is as follows

$$\begin{cases} T'(t) = \dfrac{s}{1+V(t)} - m_1 T(t) + rT(t)\left[1 - \dfrac{T(t)+I(t)}{T_{\max}}\right] \\ \qquad\qquad -u(t)kV(t)T(t) \\ I'(t) = u(t)kV(t)T(t) - m_2 I(t) \\ V'(t) = Nm_2 I(t) - m_3 V(t) \end{cases} \tag{10}$$

The state functions T, I, and V represent the concentrations of healthy helper T-cells, infected helper T-cells, and free HIV particles. The control function $u(t)$ represents the strength of the RTI therapy at time t, normalized to $0 \le u \le 1$, where $u = 0$ represents the maximum drug strength, and $u = 1$ represents no therapy.

The real-valued parameters in (10) are as follows: s is the proportionality constant for the generation of new T cells, whose production rate is inversely proportional to $1 + V$; m_1 is the natural death rate of T cells; m_2 is the natural death rate of infected cells; m_3 is the death rate of free HIV particles; r is the growth rate of T cells, where the growth is logistic; T_{max} is the maximum possible number of T cells; k is the proportionality constant for the infection of T cells, where the rate is proportional to the product VT; and N is the average number of virus particles produced by an infected T cell before it dies. Thus Nm_2I is the growth rate of free HIV particles.

The primary objective here is to maximize the number of healthy helper T cells. Recognizing that HIV may develop drug resistance to the RTI therapy, and that the therapy itself has side effects such as headaches, nausea, numbness, severe fatigue, or even kidney problems, the secondary goal is to minimize the cumulative drug strength. One way to formulate the dual objective is

$$\min_u F(u) = \int_0^{t_f} \left(1 - u(t)\right)^2 - AT(t)dt\,, \tag{11}$$

where t_f is the duration of the RTI therapy, and A is the relative weight of the primary objective over the secondary. Note that minimizing $-\int_0^{t_f} T(t)dt$ in (11) is equivalent to maximizing $\int_0^{t_f} T(t)dt$. Finally, the bounds $0 \le u \le 1$ are imposed as hard constraints, as above.

The HIV control problem is solved for model parameters s = 10, m_1 = 0.02, m_2 = 0.5, m_3 = 4.4, r = 0.03, T_{max} = 1500, k = .000024, and initial values T_0 = 800, I_0 = 0.04, and V_0 = 1.5. The relative objective weight is A = 0.05, and the treatment duration is t_f = 20 days. For details of the indirect solution, see [2]. The DE strategy used is DE/local-to-best/1, F = 0.85, CR = 1. BCP solutions were found for curves of order 3 to 6. The dimension of the optimization problem is then $D = n + 1$, and the DE population size used is $NP = 10(n + 1)$. The initial population is formed by random selection of control parameters, within the bounds [-5, 5], and optimization is terminated after 200 generations.

The results for the HIV model correspond to those of the earlier models. The BCP method closely approximates the indirect solution, even for the degree 3 curve, shown in Figure 4. Deviations from the indirect solution again diminish as the degree of the curve increases.

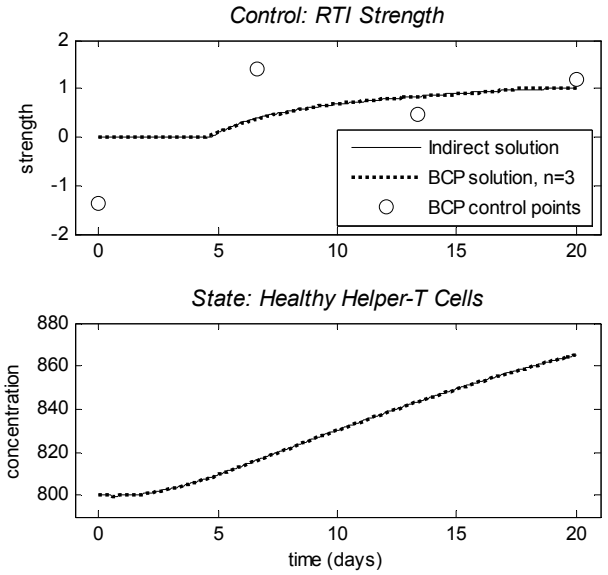

Figure 4. Comparison of solutions for optimal control of HIV by RTI therapy, model (10), (11).

The optimal treatment begins with maximum RTI strength for the first five days, and then decreases to zero strength by day 20. Under this strategy, the concentration of healthy T cells steadily increases over the duration of the therapy. This behaviour is indeed commonly seen in drugs such as AZT and DDT. While not shown in Figure 4, the infected T cells and the free virus particles decrease initially. They do recover slightly as the RTI strength lessens, but never to concentration levels that approach those of the initial infection.

4. CONCLUSIONS

The BCP method proves successful for each of the varied optimal control problems for disease models considered here. It results in an accurate approximation of solutions obtained by indirect methods, but with considerably less effort, as only one algorithm is required to solve three very different problems. A common feature of the BCP solutions is their convergence to the actual optimal solution, as the order of the Bézier curve is increased. In practice, of course, a numerical or exact solution will not be known in advance, and this feature will permit a researcher using the BCP method to begin with a low order solution, then increase the degree of the curve to test whether convergence occurs. If it does, one can be reasonably confident that the method has indeed found the optimal control. The three test cases considered here suggest that a degree three curve is a good place to start, and that convergence is likely to occur by about degree six or seven.

The evolutionary BCP direct method has potential to be a simple, general solution method for any optimal control problem. It can be applied to virtually any disease that can be modeled with a continuous-time control. This can be extremely helpful in the field of epidemiological and biomedical modeling, in which researchers requiring optimal public health policies or optimal treatment schedules may not have the mathematical skills, or the time, to solve the model indirectly.

Future work will focus on enhancements to make the method more general yet. In its current implementation, only one control function can be parameterized. Many disease models, however, have multiple inputs controlling the system. For example, a public health policy might consider three strategies simultaneously: vaccination, quarantine, and isolation rates. The BCP method will be extended to represent multiple control functions. An excellent test case would be the SARS model controlled by quarantine and isolation [21].

Evolutionary optimization also has the ability to solve many control models that indirect methods cannot. Because the indirect approach uses variational calculus to find optimal controls, it is inherently a local optimizer. Many disease models, however, are multimodal. That is, these systems are likely to have many local optima, making it difficult to be confident that the indirect solution is actually the global optimum. Evolutionary optimizers, on the other hand, are known for their abilities to find near-global solutions, even when searching through very complicated landscapes. DE, in particular has been shown is known to be an effective optimizer for multimodal optimal control problems [7], but only for piecewise constant control parameterization. Future work will investigate the efficacy of BCP for these multimodal problems.

Another limitation of indirect methods is that they must be formulated with a single objective functional, cast in integral form, as in equation (1), whose integrand must often be quadratic in the control function. This is rather restrictive. When, for example, disease models have multiple objectives, these must be combined into one functional, and the mathematician must decide in advance its formulation, including the relative weight of the different objectives. This was the approach taken in the three models studied in this paper, in order to compare the direct and indirect results. However, other formulations of the objective may in fact result in treatment strategies that for various reasons might be preferred by health professionals. Thus, true multi-objective optimization is a better approach, in which a range of possible solutions are produced, all of them, for example, Pareto optimal. One of the strengths of EAs is that they can be adapted to do so. Future work will thus incorporate multi-objective search into the algorithm. This will permit a more general formulation of the optimal control problem. More importantly, perhaps, by providing multiple results, all of them in some sense optimal, it will put the decision-making back into the hands of medical practitioners, where it belongs.

5. ACKNOWLEDGMENT

This work has been supported financially by a Canadian Mennonite University Faculty Research Grant. The author also thanks Dr. Abba Gumel, professor of mathematics at University of Manitoba, for suggesting the problem of optimal control, motivated by his work in epidemiological modeling.

6. REFERENCES

[1] Rogalsky, T. 2011. Bézier parameterization for optimal control by differential evolution. In *Proceedings of the 2011 International Conference on Genetic and Evolutionary Methods* (Las Vegas Nevada, July 2011). CSREA Press, 28-34.

[2] Lenhart, S. and Workman, J.T. 2007. *Optimal Control Applied to Biological Models*. Chapman & Hall/CRC, Taylor & Francis Group, Boca Raton.

[3] Goldberg, D.E. 1989. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, Reading.

[4] Beyer, H.-G. 2001. *The Theory of Evolution Strategies*. Springer, New York.

[5] Price, K., Storn, R. and Lampinen, J. 2005. *Differential Evolution: A Practical Approach to Global Optimization*. Springer, New York.

[6] Deb, K. 2001. *Multi-Objective Optimization Using Evolutionary Algorithms*. John Wiley & Sons, New York.

[7] Lopez-Cruz, I.L., Van Willigenburg, L.G., and Van Straten, G. 2003. Efficient differential evolution algorithms for multimodal optimal control problems. *Applied Soft Computing*, 3, 2 (Sept. 2003), 97–122.

[8] Chiou, J.P. and Wang, F.S. 1999. Hybrid method of evolutionary algorithms for static and dynamic optimization problems with application to a fed-batch fermentation process. *Computers and Chemical Engineering*, 23, 9 (Nov. 1999), 1277-1291.

[9] Lee, M.H. Han, C., and Chang, K.S. 1999. Dynamic optimization of continuous polymer reactor using a modified differential evolution algorithm. *Ind. Eng. Chem. Res*., 38, 12 (Dec. 1999), 4825–4831.

[10] Lopez-Cruz, I.L., Van Willigenburg, L.G., and Van Straten, G. 2003. Optimal control of nitrate in lettuce by a hybrid approach: differential evolution and adjustable control weight gradient algorithms. *Computers and Electronics in Agriculture*, 40, 1-3 (Oct. 2003), 179-197.

[11] Kapadi, M.D. and Gudi, R.D. 2004. Optimal control of fed-batch fermentation involving multiple feeds using differential evolution. *Process Biochemistry*, 39, 11 (July, 2004), 1709–1721.

[12] Auger, A., Hansen, N. Perez Zerpa, J.M., Ros, R., and Schoenauer, M. 2009. Experimental comparisons of derivative free optimization algorithms. In *Lecture Notes in Computer Science, 5526: Proceedings of the 8th International Symposium on Experimental Algorithms*. Springer-Verlag, New York, 3-15.

[13] Goh, C.J. and Teo, K.L. 1998. Control parametrization: a unified approach to optimal control problems with general constraints. *Automatica*, 24, 1 (Jan. 1988), 3–18.

[14] Vlassenbroeck, J. 1988. A chebyshev polynomial method for optimal control with state constraints. *Automatica*, 24, 4 (July 1988), 499–506.

[15] Biegler, L. 1984. Solution of dynamic optimization problems by successive quadratic programming and orthogonal collocation. *Comp. Chem. Engng*., 8, 3-4 (1984), 243–248.

[16] Vassiliadis, V.S., Sargent, R.W.H., and Pantelides, C.C. 1994. Solution of a class of multistage dynamic optimization problems 1. Problems without path constraints. *I&EC Res*., 33, 9, (Sept. 1994), 2111–2122.

[17] Smith, S. 1995. An evolutionary program for a class of continuous optimal control problems. *Proc. IEEE Conference on Evolutionary Computation*, Piscataway, 418–422.

[18] Dakev, N.V., Chipperfield, A.J., and Flemming, P.J. 1995. A general approach for solving optimal control problems using optimization techniques. *Proc. IEEE Conference on Systems, Man and Cybernetics, part 5* (Vancouver, 1995), 4503–4508.

[19] Bäck, T. 1996. *Evolutionary Algorithms in Theory and Practice: Evolution Strategies, Evolutionary Programming, Genetic Algorithms*. Oxford University Press, New York.

[20] Rogalsky, T., Derksen, R.W., and Kocabiyik, S. 2000. Differential evolution in aerodynamic optimization. *Canadian Aeronautics and Space Journal,* 46, 4 (Dec. 2000), 183-190.

[21] Yan, X. and Zou, Y. 2008. Optimal and sub-optimal quarantine and isolation control in SARS epidemics. *Mathematical and Computer Modelling*, 47, 1-2 (Jan. 2008), 235-245.

[22] Liang, Y., Leung, K-S., and Shu Kam Mok, T. 2006. A novel evolutionary drug scheduling model in cancer chemotherapy. *IEEE Transactions on Information Technology in Biomedicine*, 10, 2 (April 2006), 237-245.

[23] Neri, F., Toivanen, J., and Mäkinen, R. 2007. An adaptive evolutionary algorithm with intelligent mutation local searchers for designing multidrug therapies for HIV. *Applied Intelligence*, 27, 3 (Dec. 2007), 219-235.

[24] Fister, K.R. and Panetta, J.C. 2003. Optimal control applied to competing chemotherapeutic cell-kill strategies. *SIAM Journal of Applied Mathematics*, 63, 6 (2003), 1954-71.

[25] Butler, S., Kirschner, D., and Lenhart, S. 1997. Optimal control of chemotherapy affecting the infectivity of HIV. *Advances in Mathematical Population Dynamics – Molecules, Cells and Man*, 6 (1997), 557-69.