Heuristic Design of Cancer Chemotherapies

Minaya Villasana and Gabriela Ochoa

***Abstract*—A methodology using heuristic search methods is proposed for optimizing cancer chemotherapies with drugs acting on a specific phase of the cell cycle. Specifically, two evolutionary algorithms, and a simulated annealing method are considered. The methodology relies on an underlying mathematical model for tumor growth that includes cycle phase specificity, and multiple applications of a single cytotoxic agent. The goal is to determine effective protocols for administering the agent, so that the tumor is eradicated, while the immune system remains above a given threshold. Results confirm that modern heuristic methods are a good choice for optimizing complex systems. The three algorithms considered produced effective solutions, and provided drug schedules suitable for practice, although some methods excelled others in performance. A discussion of comparative results is presented.**

***Index Terms*—Cancer model, cycle-specific chemotherapy, evolution strategies (ESs), evolutionary algorithms (EAs), genetic algorithms (GAs), simulated annealing (SA), singular optimal control.**

I. Introduction

TUMOR CELLS can be divided into proliferating or cycling cells and nonproliferating or quiescent cells [3]. A cell is considered cancerous when it has lost its ability to regulate cell growth and division (mitosis). Thus, cancer is a disease of rapid uncontrolled growth of malignant cells. A successful treatment should target rapidly cycling or proliferating cells, as are tumor cells (also hair follicle cells, cells in the digestive tract, etc., where chemotherapy has common side effects). Cycle-phase-specific drugs are those drugs that act on a specific phase of cycling cells. One such example is Taxol® (paclitaxel), which is hailed as a promising antineoplastic agent for treating breast cancer. It was first isolated in 1971 from the bark of the pacific yew, *Taxus brevifolia*, and has shown efficacy in the treatment of cancer. The cytotoxic action of this drug is carried through different mechanisms: it inhibits mitosis, induces apoptosis (programmed cell death), and enhances tumor radiosensitivity. Today paclitaxel is used for treating breast cancer, ovarian cancer, head, and neck cancer, sometimes as a single agent or often accompanied by agents such as 5-FU or Doxorubicin. The optimal scheduling and possible drug interactions for paclitaxel are not totally understood [12] yet. When a new drug is discovered, there are many costly trial and error procedures to determine the best way of administering it.

Manuscript received May 18, 2003; revised March 22, 2004 and April 1, 2004.

M. Villasana is with the Departamento de Cómputo Científico y Estadística, Universidad Simón Bolívar, AP 89000, Caracas 1081-A, Venezuela (e-mail: mvillasa@usb.ve).

G. Ochoa is with the Departamento de Computación, Universidad Simón Bolívar, Caracas 1081-A, Venezuela (e-mail: gabro@ldc.usb.ve).

Digital Object Identifier 10.1109/TEVC.2004.834154

Many authors have dealt with the drug scheduling problem, for example, Cojocaru and Agur [6] studied drug scheduling for cycle-phase-specific drugs based on probabilistic modeling and found a relationship between scheduling times and cell cycle times. Also, Panetta and Adams [21] studied the effect of cycle-specific chemotherapy and determined best periods of application on simple models. De Pillis and Radunskaya [8] posed an optimal control problem (OCP) to schedule chemotherapy. Their implementation resulted in a set of ordinary differential equations with initial and boundary values, which they solved using an iterative numerical method. In [18], McCall and Petrovski present an interactive tool (OWCH) that computes, through the use of genetic algorithms (GAs), possible treatment schedules with different drugs. The authors assumed gompertz dynamic for the patient with no other cell interaction. They accounted for the effect of the treatment on cancer and added constraints on total drug administration and maximum allowed tumor levels.

The aim of our study is to design drug schedules with paclitaxel as the only chemotherapeutic agent. The patient dynamic is modeled following the formulation by Villasana and Radunskaya [28], who proposed a model considering tumor growth, interaction with immune cells, and the application of a cycle-phase-specific drug. For this purpose, an OCP is formulated and heuristic search methods are used to solve it. The goal is to drive the tumor population below a level (given by the basin of attraction of the zero tumor level fixed point) after which the patient's immune system removes the remaining tumor cells.

The present study enhances McCall's and Petrovski's since the patient model is designed to consider cycle-phase-specific drugs, and it also includes the effect of certain immune cell lines that are crucial in the fight against cancer. Our formulation specifies an objective function that not only ensures that the tumor level is low at the end of the treatment but also minimizes tumor levels during the course of treatment. Moreover, three heuristic search methods are tested and compared. Our approach, however, does not include drug cocktails, leaving this as future work.

The OCP we propose admits bang-bang solutions, meaning that the control (the drug injected at time t) attains its maximum or minimum value for bounded control functions. Under this setting it suffices to determine the best switching times of the solution, i.e., the times in which the drug begins or ceases to be administered. Since an analytical solution to this problem is not feasible given the system's complexity, alternative heuristic approaches were adopted. Specifically, we explored three methodologies: two evolutionary algorithms (EAs): 1) GAs; and 2) evolutionary strategies (ESs); and 3) simulated annealing (SA).

This paper is organized as follows. Section II-A briefly recalls the model presented in [28] for a single application of the drug and expands to multiple applications. Section II-B describes the

corresponding OCP and shows the admittance of bang-bang solutions. Section III describes the three heuristic methods considered, while Section IV provides the algorithms' parameter settings used. Section V presents our empirical results and compares the three algorithms' performance. Finally, Section VI summarizes and discusses our findings.

II. Problem Formulation

A. Mathematical Model

The tumor growth model used [28], considers the effects and interactions between tumor cells and immune cells, it also differentiates between cell phases for subsequent treatment with a cycle-phase-specific drug (i.e., a drug that acts on a specific phase of the cell cycle).

The cell cycle is the process between two cell divisions or mitosis. It encompasses four stages: G_1, S, G_2, and M. G_1 is a resting phase (or Gap period) called presynthetic phase. S is the synthetic period, where the replication of deoxyribonucleic acid (DNA) occurs. G_2 is another gap period called postsynthetic phase, and finally, M or mitosis in which cells segregate the duplicated DNA material between daughter cells. Generally, most cycle-phase-specific drugs interfere with mitosis, stopping cell proliferation and allowing their natural death. To properly account for cycle-specificity, and model how the drug acts, the tumor population is subdivided into its different stages.

Let $T_I(t)$ and $T_M(t)$ denote the population of tumor cells during interphase (period comprising G_1 through G_2) and mitosis at time t, and $I(t)$ the immune system population at time t. Let $u(t)$ be the concentration of drug present at time t and τ be the resident time of cells in interphase. The governing equations for the system with a multiple applications of the drug are

$$\begin{aligned}
T_I' &= 2a_4T_M - (c_1I + d_2)T_I - a_1T_I(t-\tau)\\
T_M' &= a_1T_I(t-\tau) - d_3T_M - a_4T_M - c_3T_MI\\
&\quad - k_1\left(1 - e^{-k_2u}\right)T_M\\
I' &= k + \frac{\rho I(T_I+T_M)^3}{\alpha+(T_I+T_M)^3} - c_2IT_I - c_4T_MI\\
&\quad - d_1I - k_3(1-e^{-k_4u})I\\
u_1' &= -\lambda_1u_1 + \mathbf{c}(t)\\
u_2' &= -\lambda_2u_2 + \mathbf{c}(t)
\end{aligned} \tag{1}$$

where $'$ denotes derivatives with respect to time and with initial data given by

$$\begin{aligned}
T_I(t) &= \phi_1(t), \quad \text{for } t \in [-\tau, 0]\\
T_M(t) &= \phi_2(t), \quad \text{for } t \in [-\tau, 0]\\
I(t) &= \phi_3(t), \quad \text{for } t \in [-\tau, 0]\\
u_1(0) &= 0\\
u_2(0) &= 0.
\end{aligned}$$

The immune system is comprised of many different types of cells, [28] focused on a subset of these cells, namely, the population of cytotoxic T cells (CTL), since they play a fundamental role in combating cancer.

The drug free system corresponding to system (1) can have up to five different fixed points depending on the parameter values. However, there is one fixed point that is always present in this system, namely, $(0, 0, k/d_1)$. This fixed point represents a tumor-free environment with positive immune population, which is a desirable scenario. The stability of this fixed point was studied in [28]. The effect of the drug is represented by the term $-k_1(1 - e^{-k_2u})T_M$, which expresses the removal of the cells from the cell cycle due to the action of the drug. A similar term with different parameters values is used for the effect of the drug on the immune system.

The terms d_2T_I, d_3T_M, and d_1I in the model equations represent proportions of natural cell death or apoptosis, a_1 and a_4 represent the different rates at which cells cycle or reproduce, and together with τ regulate the pace of cell division. The terms T_MI and T_II are standard competition terms that represent losses due to encounters among the different cell types. The coefficients that accompany these terms c_i provide the fraction of the losses from these encounters. The term $(\rho I(T_I + T_M)^n)/(\alpha + (T_I + T_M)^n)$ represents the nonlinear growth of the immune population due to stimulus by the tumor cells. With this choice, the recruitment function is zero when there are no tumor cells, and increases monotonically toward a horizontal asymptote: this rational form reflects these characteristics in a simple, smooth function. The parameters ρ, α, and n depend on the type of tumor being considered and the health of the immune system. The model assumes that in the absence of tumor cells, immune cells are produced at a constant rate k (bone marrow production). The tumor cells reside in interphase for a certain period of time τ, before continuing in the cycle to mitosis. Assuming that cells reside in interphase τ units of time, then the cells that enter mitosis at time t are those cells that entered interphase τ units of time before. This explains the two terms $T_I(t-\tau)$ appearing in system (1).

Paclitaxel, like many other drugs, has a decay that can be modeled with two separate elimination rates (bi-exponential curve). A fast rate of decay while the drug is distributed through the blood to the tissues and a second, slower rate of decay in the peripheral compartment or tissue [25]. Thus, the decay function is expressed as

$$\text{decay}(t) = r_1e^{-\lambda_1t} + r_2e^{-\lambda_2t} \tag{2}$$

with r_1, and r_2 real adimensional constants.

Letting u_1 and u_2 be such that the concentration of drug at any given time is a linear convex combination $u(t) = r_1u_1(t) + r_2u_2(t)$. The last two equations of system (1) model this situation with multiple drug applications in time, and are identified with the function $\mathbf{c}(t)$, which is the concentration of paclitaxel that goes in the system at time t. With this choice and initial conditions, we get

$$u(t) = \mathbf{c}(t) * \text{decay}(t)$$

where $*$ denotes convolution.

Parameter estimation was performed on the drug free system (for more details see [28]), where the estimation was done for the drug terms using the information available for paclitaxel in [5], [12], and [29]. We took these reports as guidelines to obtain the effect of the drug on the different cell lines.

Before beginning the analysis, we nondimensionalized and scaled our variables so that the model constants are as close to

TABLE I
PARAMETER VALUES

$\tau = 0.92$	$a_1 = 0.98$	$a_4 = 0.8$	$d_2 = 0.11$
$d_3 = 0.4$	$d_1 = 0.29$	$c_1 = 9 \cdot 10^{-1}$	$c_3 = 9 \cdot 10^{-1}$
$c_2 = 85 \cdot 10^{-3}$	$c_4 = 85 \cdot 10^{-3}$	$\rho = 0.1$	$\alpha = 0.2$
$k = 0.036$	$n = 3$	$\lambda_1 = 126.12$	$\lambda_2 = 0.85$
$k_1 = 0.47$	$k_2 = 0.57$	$k_3 = 0.49$	$k_4 = 0.061$
$r_1 = 0.73$		$r_2 = 0.27$	

unity as possible and in the same range of values. This improves the stability of the numerical method for solving the equation. However, the parameters will vary between tumor types and from person to person. Therefore, there is no unique set of parameters values for any given model. The nondimensional parameter values used throughout this report are summarized in Table I. This particular set represents a patient with a rapidly growing tumor and an immune system not able to control the tumor progression, resulting in his/her eventual death.

The general goal is to use the drug to drive the system inside the estimated basin of attraction of the tumor-free fixed point, while maintaining a minimum level of CTL.

The basin of attraction of a stable equilibrium can be defined as the set of initial (history) functions ϕ for which $\lim_{t\to\infty} x(t, t_0, \phi) = 0$. In other words, it is the set of functions ϕ, for which the orbits go toward an equilibrium. One way to estimate a basin of attraction is by the use of a Lyapunov function, which has been done in [27], where Villasana was able to obtain a subset of the basin for the stable fixed point that represents a tumor-free environment $(0, 0, k/d_1)$. In this paper, we computed numerically the basin of attraction when a constant history function is considered. This computation was done in the following way: For each point in the discretization of three-space (each point representing a history function equal to the constant value), system (1) was integrated using delay differential equation (DDE) for MATLAB, [26]. If at the final time (for sufficiently large time) the tumor level is less than the original level, then this point (function) is regarded as inside the basin of attraction. Fig. 1 shows the estimated basin of attraction when starting with a constant initial function for the tumor-free fixed point.

We note that this estimated basin is a subset of the actual basin, which may or may not depend on the immune system (I). When the drug is added to the system, (a single bolus injection) it can be seen that the region is enlarged, thus increasing the estimated basin of attraction for the tumor-free fixed point.

The basin of attraction for different types of history functions has been computed in [27], and its effect is seen in the enlargement or reduction of the basin of attraction that will extend or shorten the total length of treatment according to the type of history. A constant history function is a simplification of a past history behavior (which is usually unknown). One may argue that the course of the tumor in the past is an increasing function, however, in the time frame where the history is needed (≈1 day),

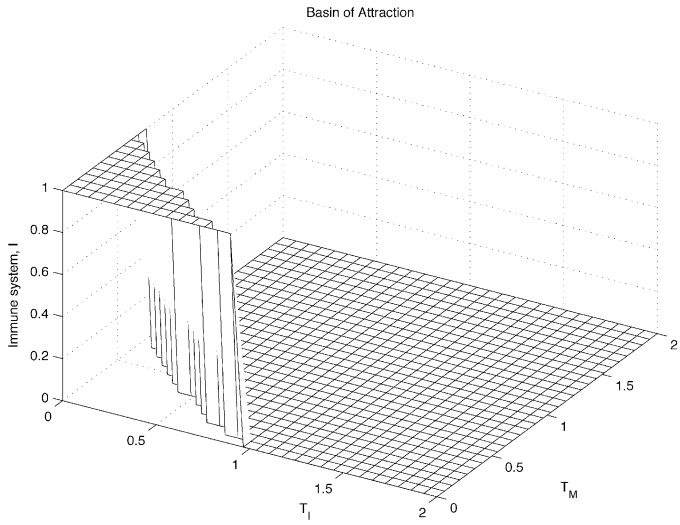

Fig. 1. Basin of attraction of tumor-free fixed point $(0, 0, k/d_1)$ when the history is a constant function. The x axis represents tumor in interphase (T_I), the y axis represents the tumor in mitosis (T_M), and the z axis is the immune system (I).

we find that assuming a constant history function is an acceptable simplification of the model.

Starting with a constant initial function outside the basin of attraction $(T_I(0), T_M(0), I(0)) = (1.3, 1.2, 0.9)$ and apply the drug to drive the tumor population inside its basin of attraction.

B. Optimal Control Problem (OCP)

What is the best course of treatment with the single agent paclitaxel on the model described (1), so that the tumor is eradicated, while the immune system remains above a given threshold?

Previous research shows [2], [4], [8], [24], that optimal control theory can be used to determine the optimal therapeutic regime, whether it is applied to cancer, *H. influenzae*, or any other viral infection. Even though optimal control for systems of DDE seems to be an adequate tool in the biosciences, there are relatively few studies in this area.

The general mathematical formulation is as follows:

$$\text{Min. } \Phi_0(Y(T)) + \int_0^T F_0(t, Y(t), Y(t-\tau_y), \mathbf{c}(t), \mathbf{c}(t-\tau_u))\, dt$$
$$\text{s.t. } Y'(t) = f(t, Y(t), Y(t-\tau_y), \mathbf{c}(t), \mathbf{c}(t-\tau_u)), \quad o \le t \le T$$

and

$$k(t, Y, u) \le 0. \tag{3}$$

The control function $\mathbf{c}(t)$ is the amount of drug introduced into the system as a function of time, determining the scheduling and dosing of the drug. Our goal is to minimize the average and final tumor size, thus, the problem is stated as

$$\text{Min} \quad T_I(t_f) + T_M(t_f) + \frac{1}{t_f}\int_0^{t_f} T_I(t) + T_M(t)\, dt$$
$$\text{s.t.} \quad \text{Equations in system (1)} \tag{4}$$

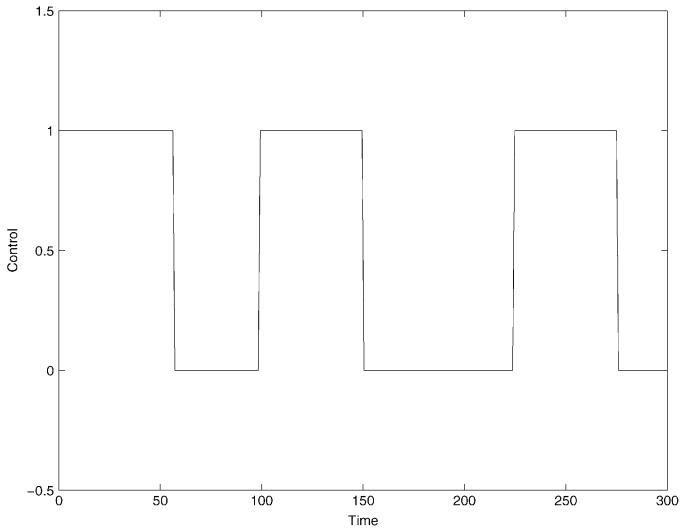

Fig. 2. Schematic view of bang-bang solutions.

along with the added restriction

$$I - \gamma I_{\max} \geq 0.$$

All the parameters involved in (4) are known, but nothing has been said about the threshold imposed over the immune system. There is no methodological way to determine this value. In practice, we want the patient to be as healthy as possible. In our experiments, we required that the immune system does not fall below its initial state.

Pontryagin's Maximum Principle is used to obtain necessary conditions for an analytical solution to this problem (see [15]). For sufficiency conditions, we would have to verify that f is linear in the control. Also, that functions F and f in (3) are concave in the state variables and the control.

Let v_i denote the co-state or adjoint variables. To solve the problem described by the set of equations in (4), we form the Hamiltonian H and solve the following additional set of differential equations for the co-state variables. These equations can be derived from the Hamiltonian by taking its derivative with respect to the state variables.

The Hamiltonian in our case is

$$H = \frac{1}{t_f}[T_I(t) + T_M(t)] + v_1(T_I') + v_2(T_M') + v_3(I') + v_4(u_1') + v_5(u_2') + \zeta(I - \gamma I_{\max}). \quad (5)$$

It is not hard to find the system of forward equations that the co-state variables must satisfy (see [15] and [27]), along with the condition:

$$\frac{\partial H}{\partial \mathbf{c}} = v_4 + v_5 = 0. \quad (6)$$

The fact that we have a set of differential equations forwarded in time (for the co-state variables) and differential equations delayed in time (state variables) make finding an analytical solution very difficult. A numerical solution for the augmented system (delayed plus forwarded differential equations) entitles the resolution of a two point boundary value problem since

application time 1	resting time 1	application time 2	resting time 2	...	application time n	resting time n

Fig. 3. Schematic view of the control variable.

the co-state variables have end point conditions. This numerical problem is difficult in itself from a numerical point of view, therefore, we resort to heuristic approaches. From (6), we have a singular problem, that is a problem were the Hamiltonian's gradient does not provide information about the control when it is zero. This occurs when the controls appear linearly in the state equations [17]. Since the amount of drug is bounded above, then candidates for solution are bang-bang, i.e., they attain the maximum value when $v_4 + v_5 < 0$, and the lowest value when $v_4 + v_5 > 0$ (see Fig. 2).

Since the control variable is the amount of drug administered, and solutions are bang-bang, the problem reduces to finding the times where the solution $\mathbf{c}(t)$ switches from "ON" to "OFF." That is, the times at which we begin and cease administering the drug. Each ON-OFF switching constitutes a chemotherapeutic cycle.

C. Problem Encoding and Objective Function

In order to admit variable time intervals, the switching times were set as the control variables, which can be encoded naturally as real numbers. Two types of control variables are distinguished: administration-time lengths and resting-time lengths, these variables are intercalated and concatenated to encode a potential solution to the problem (configuration, see Fig. 3). The range of admitted values is different for each type of variable, being [0.2, 5] days for administration-time lengths and [0, 30] days for resting-time lengths. A parameter P indicates the number of switching times. We found empirically that nine (application/resting) cycles were enough to drive the tumor into the basin of attraction, therefore we report results for treatments consisting of nine cycles (that is $P = 18$).

The objective function was the minimization of average and final tumor size, subject to the state equations, and a penalty term for the restriction over immune system (4).

III. Heuristic Methods

Conventional search techniques are often incapable of optimizing nonlinear multimodal functions. In such cases, an heuristic search method might be required. Heuristic methods do not use much knowledge of the problem to be optimized, nor they depend on special properties of the objective function. This section describes the heuristic algorithms included in our study. Details of implementation, and choice of algorithm's components and parameters will be discussed in Section IV.

A. Evolutionary Algorithms (EAs)

Over the last few decades, several EAs have been proposed and studied, they share the conceptual framework of simulating natural evolution. The fundamental components of an EA are the following:

- a representation of candidate solutions to the problem at hand;
- a population of these candidate solutions;

- mechanisms for generating new solutions from members of the current population (operators such mutation and recombination);
- an evaluation or fitness function to assess the quality (fitness) of a given solution;
- a selection method which gives better chances of survival to good solutions.

Historically, there have been three well-defined approaches to evolutionary computation: evolutionary programming (EP) [9], evolution strategies [23], and GAs [13]. Although similar at the highest level, these approaches differ in the way they implement an EA. The differences touch all the aspects of EAs, including the choices of representation for the individual structures, types of selection mechanisms, forms of variation operations, and measures of performance. The major differences are, however, in the choice of representations, and emphasis and use of variation operators. EP uses representations that are tailored to the problem domain. Similarly, ESs, due to initial interest in hydrodynamic optimization problems, use real-valued vector representations. On the other hand, GAs have traditionally used a more domain independent representation, namely, binary strings. Regarding variation operations, both EP and ES use mutation as the main operator, and propose a form of self-adaptive mutation; whereas GAs emphasize recombination as the main search operator, and use mutation as a secondary operator applied with a small constant probability.

Mutation strength adaptation is a distinctive component of ESs. An ill adjusted mutation strength can considerably slow down progress if it is too low, or lead to divergence if it is too high. Rather than using fixed schedules, ESs employ dynamic schemes that adjust to the local characteristic of the problem landscape. Several mutation strength adaptation methods have been proposed (see [1] for a summarized review). A state of the art scheme: cumulative mutation strength adaptation, has been proposed by Hansen and Ostermeier [10], [11]. This scheme attempts to "derandomize" the process of mutation strength adjusting. Unlike previous methods, cumulative adaptation is deterministic rather than based on variation and selection. It works by accumulating and analyzing information over a number of time steps.

For this paper, we selected the two mainstream EAs, namely, GAs and ESs.

B. Simulated Annealing (SA)

SA [16] builds on an analogy with thermodynamics, specifically with the way that liquids freeze and crystalize, or metals cool and anneal. If the liquid or metal is cooled slowly, the system may reach a state of minimum energy. The analogy with optimization arises when we consider the optimal solution as the state of minimum energy, and the optimization process as the process of cooling. The method works as follows: a solution to a problem is somehow perturbed, and a neighbor solution of less quality is accepted with probability according to a Boltzmann distribution $e^{-(\Delta E/T)}$, where ΔE represents the difference in quality between the current solution and the perturbed one, and T the temperature of the process. The higher the temperature the greater the probability of accepting a perturbed solution whose quality is worse than the current state. By decreasing T using an annealing schedule, it is possible to simulate the cooling or annealing processes.

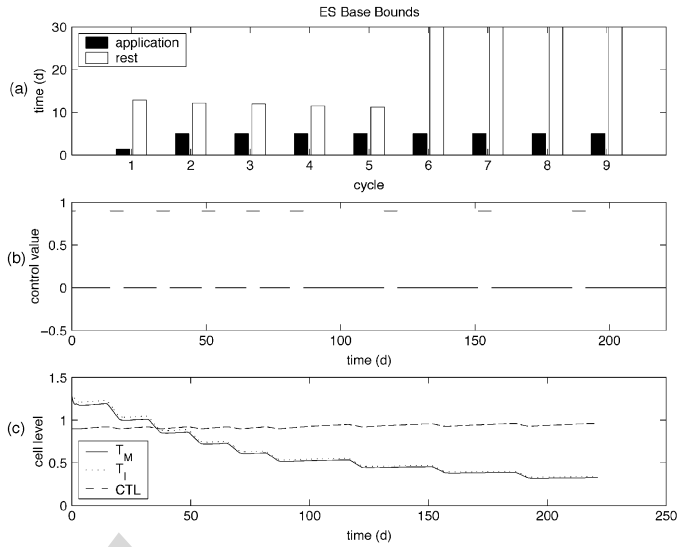

Fig. 4. Treatment obtained running the ES with the first set of bounds: [0.2, 10] days for administration times and [0, 30] days for resting times. (a) Application and resting times for each cycle. (b) Application times (horizontal lines at level 1), and resting times (horizontal lines at level 0). (c) Behavior of the system: T_M = tumor level in mitosis, T_I = tumor level in interphase, and CTL = immune cells level.

To implement a SA algorithm, one must provide the following elements [22]: 1) a description of possible system configurations (or candidate solutions); 2) a generator of random changes or perturbations in the configuration; 3) an objective function (analog of energy) whose minimization is the goal of the procedure; and 4) a control parameter T (analog of temperature) together with an *annealing schedule*, which tells how it is lowered from high to low values.

IV. Experiments

The range of values for administration and resting times were set as follows. According to the literature the maximum tolerated dose for Paclitaxel is five days of infusion at 30 mg/m^2/day, every three weeks [12], which imposes an upper bound for drug administration times. A lower limit of 3 h infusions, a common practice when using paclitaxel was selected. The resting times were set in the interval [0, 30] (zero means that there is no resting period and the treatment continues), where the upper bound (30 days) follows the current practice of standard chemotherapy schedule (i.e., infusions taking up to a week and a resting period of at least three weeks). With these constraints (see Fig. 4), most empirical results reached the maximum of five days for application times on all treatments cycles. Moreover, from the sixth cycle onwards, resting times also saturated to the maximum allowed of 30 days. Following these observations, and with the aim of exploring a wider range of schedules, we decided to run experiments extending the constraints of both the administration and resting times. Specifically, we set the following ranges: [0.2, 15] days for administration time lengths, and [0, 50] days for resting time lengths. In other words, we extended the upper bounds from 5 to 15 (administration times) and from 30 to 50 (resting times).

TABLE II
GENETIC OPERATORS AND FREQUENCY OF APPLICATION (I.E., DISCRETE NUMBER OF TIMES TO CALL THE OPERATOR EVERY GENERATION)

Name	Frequency
Uniform Mutation	4
Non-uniform Mutation	4
Multi-non-uniform Mutation	6
Simple Crossover	4
Arithmetic Crossover	4
Heuristic Crossover	2

The optimization algorithms were implemented in Matlab, since the model for tumor growth was already running in Matlab. For the first set of bounds (called from now onwards base bounds), the three algorithms were run ten times with different random seeds. For the extended bounds, given the longer integration (and, thus running) times, we could only afford five replications per algorithm. For comparison purposes, we set a maximum of 3000 evaluations per run for each algorithm. It is worth mentioning that each function evaluation required the integration of a DDE system for large periods of time. Thus, a single heuristic algorithm run took several days to complete.[1] Performing any extensive parameter tuning was not feasible on our current implementation.

A. Genetic Algorithms (GAs)

We used GAs for Optimization Toolbox (GAOT) a freeware well documented Matlab library by Houck *et al.* [14]. The authors tested the toolbox in a series of nonlinear, multimodal, nonconvex test problems (taken from Corana *et al.* [7]), and compared these results using SA. They found the GA with real-valued representation to be superior to both a binary GA and SA in terms of efficiency and quality of the solution. GAOT implements several genetic operators for real-valued representation (originally introduced by Michalewicz [19]), and suggests parameter settings for them (see Table II).[2] Furthermore, the toolbox provides three selection mechanisms: roulette wheel, tournament, and normalized geometric selection. This last method, the toolbox default, is a ranking scheme. Ranking methods, only require the objective function to map the solutions to a partially ordered set. They assign the probability of selection based on the rank of the solution when all solutions are sorted. Normalized geometric ranking [14] defines P_i for each individual as

$$P_i = q'(1-q)^{r-1}$$

where q' is the probability of selecting the best individual (with a default value of 0.8), r the rank of the individual, where 1 is the best, and

$$q' = \frac{q}{1-(1-q)^{\mathrm{Pop\cdot Size}}}.$$

[1]Up to five days for the base bounds and ten days for the extended bounds, on an up to date PC.

[2]For a detailed description of each operator see [14].

For our experiments, we selected the float representation, genetic operators, parameter settings, and selection method as suggested by Houck *et al.* [14]. The other evolutionary parameters were set as follows: a population of 30 individuals, and a fixed termination criterium of 100 generations.

B. Evolution Strategies

We selected a derandomized ES with covariance matrix adaptation (CMA-ES). This algorithm was shown, empirically, to have convergence velocity improvements over other ESs on a large function optimization test suite [10], [11]. The authors also provide a freeware, modular, and well documented Matlab implementation of their algorithm. Furthermore, CMA-ES provides default values for its strategy parameters: the number of offspring λ, has a value of $\lambda = 4 + \lfloor 3 \ln N \rfloor$ (where N is the problem size), and the number of parents μ is set to $\mu = \lfloor \lambda/2 \rfloor$. In our experiments, $N = 18$, so $\lambda = 12$ and $\mu = 6$. Finally, the weights $(w_i, \ldots, w_\mu)$ for weighted recombination, are given by

$$w_i = \ln \frac{\lambda+1}{2} - \ln i$$

for $i = 1, \ldots, \mu$. We selected these default values and set the number of iterations to 250 in order to reach the selected 3000 evaluations per run.

C. Simulated Annealing (SA)

An SA algorithm was implemented in Matlab. As the perturbation operators, we selected the mutation operators for real-valued representation provided by GAOT, specifically, uniform, nonuniform, multinonuniform and boundary mutation (see Section IV-A). For the annealing schedule, we set a initial temperature $T_{\mathrm{Max}} = 10$, and a temperature decreasing factor of 0.85. Each temperature was tried 30 times and the total number of (different) temperatures tested was 100. These parameter values were selected following suggested values taken from the literature [20], [22].

V. RESULTS

Before comparing the relative algorithms' performance, we assessed the quality of the solutions. A good solution should have tumor levels within the basin of attraction at the end of treatment, and should maintain the immune system initial state. All experiments provided these elements and were regarded as feasible solutions.

Fig. 4 shows results of a typical run using the first set of constraints (and in this case, the ES algorithm). Fig. 4(a) illustrates the treatment obtained: for each cycle, we depict the application (black bar) and resting time (white bar). Fig. 4(b) shows another view, where the whole treatment can be appreciated. Horizontal lines at level 0 represent the resting times, whereas lines at level 1 represent the application times. Finally, Fig. 4(c), provides the behavior of system (1) with the multiple drug applications. Specifically, the tumor levels in mitosis (T_M) and interphase (T_I) and the immune cells level (CTL), can be appreciated. Notice in Fig. 4(a) that the first drug application is relatively low, about one and a half days. However, from the second cycle onwards, the administration time reaches the maximum allowed

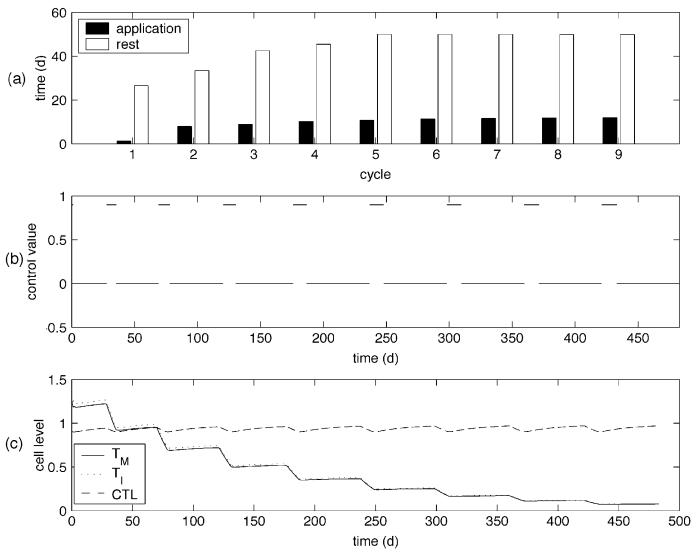

Fig. 5. Treatment obtained running the ES with the extended set of bounds ([0.2, 15] days for administration times and [0, 50] days for resting times). (a) Application (black bar) and resting time (white bar) for each cycle. (b) Application times (horizontal lines at level 1); and resting times (horizontal lines at level 0). (c) Behavior of the system: T_M = tumor level in mitosis, T_M = tumor level in interphase, and CTL = immune cells level.

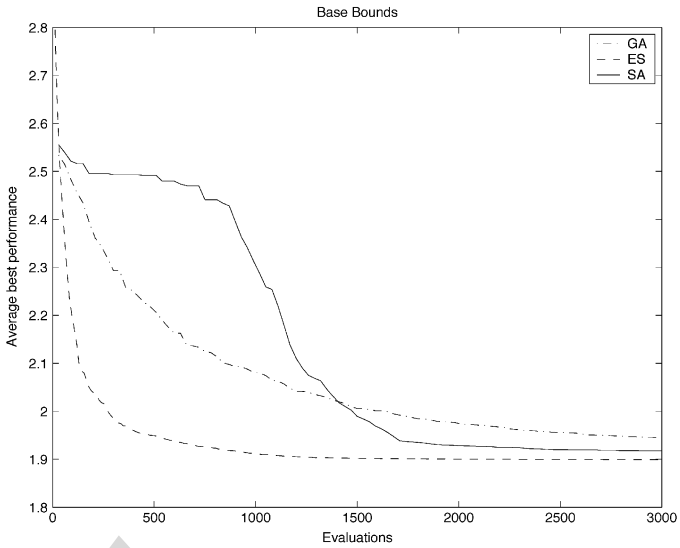

Fig. 6. Average best performance of the three methods using base bounds.

of five days. The resting times are between 11 and 13 days at the beginning of treatment, reaching the maximin allowed of 30 days from the sixth cycle onwards. The whole treatment lasts 221.13 days, and the performance index, or objective function value is $J^* = 1.8965$. Fig. 4(c) illustrates the outcome of the system with this protocol. The immune system (CTL) is always above its initial level. The tumor levels decrease steadily in a stepwise fashion, reaching a value of 0.3 for both T_M and T_I, which lies within the basin of attraction of the tumor-free fixed point.

Fig. 5 shows results of a typical run using the extended constraints (and again the ES algorithm). Notice in Fig. 5(a) that the first drug application is rather low, about 1.2 days, however, from the second cycle onwards, the administration time increases successively to 8, 9, and 10, and finally stabilizes to between 11 and 12 days. The maximum allowed of 15 days is never reached, suggesting that the system is really deciding for optimal doses instead of being limited by the upper bound. Fig. 5(b) shows another view, the whole treatment period of 483 days, or about a year and a half, can be appreciated. Note that this protocol is longer than the one presented in Fig. 4, which was expected since the allowed application and resting time ranges are longer. The resting times increase steadily from 26 days in the first cycle, reaching the maximum of 50 days from the fifth cycle onwards. Finally, Fig. 5(c) illustrates the outcome of the system with this protocol. Notice that the immune system (CTL) is, again, maintained above its initial state, showing a slight decrease after each administration time. So, despite extending the maximum allowed application time of five days, the patient health under our system is not threatened. Again, tumor levels decrease steadily in a stepwise fashion, this time reaching very low tumor levels (0.08 for both T_M and T_I). For this treatment, the performance index is $J^* = 1.0178$.

A. Comparing Algorithms' Performance

We compared the relative performance of the three heuristics for both the base and extended bounds. Comparisons were done with a fixed set of algorithm parameters for each method. A comparison with just one set of parameters does not yield a general assessment of the algorithms' behavior in this problem. For a more general picture, an extensive parameter study should be considered. This is a topic for future research, but is currently not feasible given the extensive computational times. The results discussed in the following are intended to provide just an impression of the behavioral differences of ESs, GAs, and SA, using "standard" parameter settings.

1) Base Bounds: First, we discuss the results obtained for the base bounds, as shown in Fig. 6, where the actually best objective function value is plotted against the number of function evaluations. ES showed very fast convergence, clearly demonstrating its capabilities of approaching a good solution quickly. GA performance demonstrated a slower convergence and a missing emphasis on local optimization. SA had the slowest convergence, but had better capabilities of local tuning as compared with GA.

Table III provides quantitative measures for the three heuristics after 1000 and 3000 evaluations. Clearly, ES outperformed the other two algorithms providing not only the best solution, but also giving better results on average. After 1000 evaluations, GA slightly outperformed SA on average, indicating its global searching emphasis. Whereas after 3000 evaluations the opposite occurred, illustrating the local search abilities of SA. ES, on the other hand, provided a good balance between global and local search, making this algorithm the best suited for our problem.

2) Extended Bounds: In this case, quantitative measures are summarized in Table IV, while Fig. 7 provides a qualitative view of relative performance. The situation is similar to the scenario with base bounds described above. ES is again the most suitable algorithm providing the fastest convergence, and an appropriate balance between global and local search. GA showed a faster

TABLE III
BASE BOUNDS

	After 1,000 evaluations					After 3,000 evaluations				
	Best	Worst	Median	Mean	Std. Dev.	Best	Worst	Median	Mean	Std. Dev.
GA	2.01	2.15	2.09	2.08	0.0464	1.91	2.02	1.93	1.95	0.0358
ES	1.90	1.92	1.91	1.91	0.0077	1.90	1.91	1.90	1.90	0.0043
SA	2.14	2.43	2.27	2.29	0.0956	1.90	1.96	1.91	1.92	0.0187

TABLE IV
EXTENDED BOUNDS

	After 1,000 evaluations					After 3,000 evaluations				
	Best	Worst	Median	Mean	Std. Dev.	Best	Worst	Median	Mean	Std. Dev.
GA	1.40	1.55	1.47	1.47	0.0561	1.17	1.31	1.27	1.24	0.0631
ES	1.06	1.31	1.11	1.16	0.1019	1.02	1.13	1.02	1.04	0.0514
SA	1.32	1.61	1.39	1.43	0.1153	1.03	1.05	1.04	1.04	0.0085

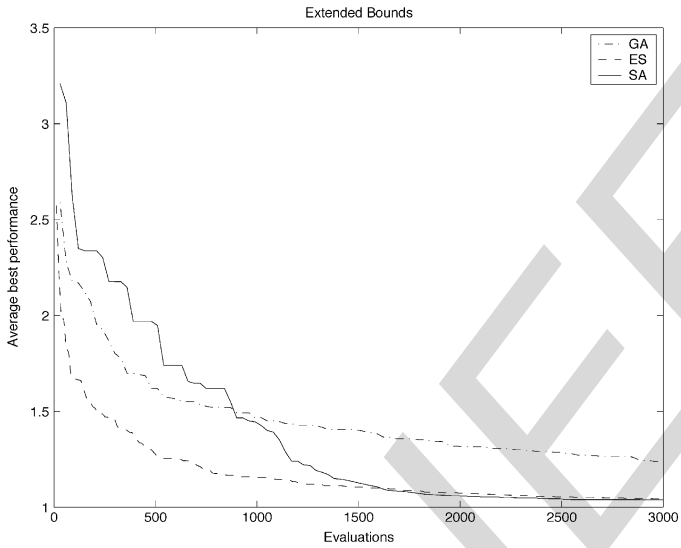

Fig. 7. Average best performance of the three methods using extended bounds.

initial convergence than SA, but stagnated and failed to fine tune the solutions.

VI. DISCUSSION

The design of efficient drug schedules for a cycle-phase-specific agent can be seen, in the present model formulation, as a singular OCP admitting bang-bang solutions, (i.e., solutions that either attain its maximum or minimum value for bounded control function). The problem can then be stated as finding the switching times, that are the times at which we begin or cease administering the anti-neoplastic agent. The heuristic methods proposed here were successful at finding suitable switching times. The algorithms were able to solve the OCP, and produced solutions that drove the system inside the basin of attraction of the tumor-free fixed point. Results confirm that modern heuristic methods are a good choice for optimizing complex systems. Although all the proposed methods produced good solutions, some methods excelled others in performance. ES clearly has the best speed of convergence, and provides an appropriate balance between global and local search. SA proved to be a good contender to ES if enough iterations are afforded, and it is a much simpler algorithm.

Using the proposed approach, we can design protocols for newly discovered drugs saving on costly trial and error procedures. These techniques are not intended to replace the expertise of the medical oncology community, but to aid them in the scheduling of new agents. Furthermore, the slightly more complex problem of scheduling drug cocktails (infusion of various drugs during the same treatment period) can also be addressed using heuristic methods. In such cases, a candidate solution must encode the switching times of all the drugs involved.

Each patient has his or her own set of biological parameters. Good parameter estimation is always a delicate issue given both ethical implications, and lack of reliable methods. The tumor model parameters selected here can be considered as a test set; further studies in this subject must be made. They served, however, for the purpose of illustrating the suitability of the proposed approach for designing efficient drug schedules.

REFERENCES

[1] D. V. Arnold and H.-G. Beyer, "A comparison of evolution strategies with other direct search methods in the presence of noise," *Comput. Opt. Applicat.*, vol. 1, no. 24, pp. 135–159, 2003.
[2] Y. Biran and B. McInnis, "Optimal control of bilinear systems: Time-varying effects of cancer drugs," *Automatica*, vol. 15, pp. 325–329, 1979.
[3] B. Birkhead, E. Rankin, S. Gallivan, L. Dones, and R. Rubens, "A mathematical model of the development of drug resistance to cancer chemotherapy," *Eur. J. Cancer Clinical Oncology*, vol. 23, no. 9, pp. 1421–1427, 1987.
[4] A. Buldayev, "A numerical method of control optimization in delay systems for immune response modeling," *U.S.S.R. Comput. Maths. Math. Phys.*, vol. 30, no. 5, pp. 18–28, 1990.
[5] L. Chuang, E. Lotzova, K. Cook, P. Cristoforoni, M. Morris, and T. Wharton, "Effect of new investigational drug taxol on oncolytic activity and stimulation of human lymphocytes," *Gynecologic Oncol.*, vol. 49, pp. 291–298, 1993.
[6] L. Cojocaru and Z. Agur, "A theoretical analysis of interval drug dosing for cell-cycle-phase-specific drugs," *Math. Biosci.*, vol. 109, pp. 85–97, 1992.

[7] A. Corana, M. Marchesi, C. Martini, and S. Ridella, "Minimizing multimodal functions of continuous variables with the "simulated annealing" algorithm," in *ACM Trans. Math. Softw.*, vol. 13, 1987, pp. 262–280.
[8] L. de Pillis and A. Radunskaya, "A mathematical tumor model with immune resistance and drug therapy: An optimal control approach," *J. Theor. Med.*, pp. **XXX–XXX**, 2001.
[9] L. J. Fogel, A. J. Owens, and M. J. Walsh, *Artificial Intelligence Through Simulated Evolution*. New York: Wiley, 1966.
[10] N. Hansen and A. Ostermeier, "Adapting arbitrary normal mutation distributions in evolution strategies: The covariance matrix adaptation," in *Proc. 1996 IEEE Int. Conf. Evolutionary Computation*, Nagoya, Japan, 1996, pp. 312–317.
[11] ——, "Completely derandomized self-adaptation in evolution strategies," *Evol. Comput.*, vol. 9, pp. 159–195, 2001.
[12] J. Hardman and L. Linbird, *Goodman and Gilman's the Pharmacological Basis of Therapeutics*, 9th ed. New York: McGraw-Hill, 1996.
[13] J. H. Holland, *Adaptation in Natural-and Artificial Systems*. Ann Arbor, MI: Univ. Michigan Press, 1975.
[14] C. Houck, J. Joines, and M. Kay, "A genetic algorithm for function optimization: A Matlab implementation," North Carolina State Univ., Raleigh, NC, Tech. Rep. NCSU-IE TR 95-09, 1995.
[15] M. Kamien and N. Schwartz, *Dynamic Optimization*, 2nd ed. Amsterdam, The Netherlands: North-Holland, 1991, vol. 31. Advanced textbooks in economics.
[16] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by simulated annealing," *Science*, vol. 220, no. 4598, pp. 671–680, May 13, 1983.
[17] R. Luus, *Iterative Dynamic Programming, Monographs and Surveys in Pure and Applied Mathematics*. London, U.K.: Chapman & Hall, 2000, vol. 110.
[18] J. McCall and A. Petrovski, "A decision support system for cancer chemotherapy using genetic algorithms," in *Proc. Int. Conf. Computational Intelligence for Modeling, Control, Automation*, vol. 1, M. Mouhammadian, Ed., Vienna, Austria, 1999, pp. 65–70.
[19] Z. Michalewicz, *Genetic Algorithms + Data Structures = Evolution Programs*. Berlin, Germany: Springer-Verlag, 1992.
[20] Z. Michalewicz and B. F. Fogel, *How to Solve it: Modern Heuristics*. Berlin, Germany: Springer-Verlag, 2000.
[21] J. C. Panetta and J. Adam, "A mathematical model of cycle-specific chemotherapy," *Math. Comput. Modeling*, vol. 22, no. 2, pp. 67–82, 1995.
[22] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes in C*. Cambridge, U.K.: Cambridge Univ. Press, 1991.
[23] I. Rechenberg, *Evolution Strategy: Optimization of Technical Systems by Means of Biological Evolution*. Stuttgart, Germany: Fromman-Holzboog, 1973.
[24] A. Rundell, R. DeCarlos, V. Balakrishnan, and H. HogenEsch, "Systematic method for determining intravenous drug treatment strategies aiding the humoral immune response," *IEEE Trans. Biomed. Eng.*, vol. 45, pp. 429–439, Apr. 1998.
[25] J. Sear, "Clinics in anaesthesiology," in *Intravenous Anaesthesiology*. Philadelphia, PA: Saunders, 1984, vol. 2, ch. 12, p. 223.
[26] L. F. Shampine and S. Thompson, Solving DDES with Matlab, http://www.runet.edu/~thompson/webddes/index.html.
[27] M. Villasana, "A delay differential equation model for tumor-immune system interactions," Ph.D. dissertation, Claremont Graduate Univ., Claremont, CA, 2001.
[28] M. Villasana and A. Radunskaya, "A delay differential equation model for tumor growth," *J. Math. Bio.*, vol. 47, pp. 270–294, 2003.
[29] W. Zoli, A. Flamigni, G. Frassineti, P. Bajorko, F. De Paola, C. Milandri, D. Amadori, and A. Gasperi-Campani, "In vitro activitity of taxol and toxotere in comparison with doxorubicin and cisplatin on primary cell cultures of human breast cancers," *Breast Cancer Res. Treatment*, vol. 34, pp. 63–69, 1995.

Minaya Villasana received the B.S. degree in applied mathematics from Simón Bolívar University, Caracas, Venezuela, in 1996 and the M.S. and Ph.D. degrees in mathematics from Claremont Graduate University, Claremont, CA, in 1999 and 2001, respectively.

She is currently a Professor at the Scientific Computation and Statistic Department, Simón Bolívar University, Caracas, Venezuela. Her research interests are in the fields of mathematical modeling with a special emphasis for applications in the natural sciences.

Gabriela Ochoa received the B.Sc. degree in computer engineering in 1991 and the M.Sc. degree in computer science from the Simón Bolívar University, Caracas, Venezuela, and the Ph.D. degree from the University of Sussex, Sussex, U.K., in 2004.

In 1996, she was the Computer Science Department, Simón Bolívar University. She is currently a Professor of Computer Science at the Simón Bolívar University. Her research interests are in the fields of evolutionary computation and machine learning, with a special emphasis in applications in medicine and biology.