

Multiobjective Optimization Applied to the Eradication of Persistent Pathogens

Ole Steuernagel and Daniel Polani

Abstract—In scenarios such as therapeutic modeling or pest control, one aims to suppress infective agents or maximize crop yields while minimizing the side-effects of interventions, such as cost, environmental impact, and toxicity. Here, we consider the eradication of persistent microbes (e.g., *Escherichia coli*, multiply resistant *Staphylococcus aureus* (MRSA-“superbug”), *Mycobacterium tuberculosis*, *Pseudomonas aeruginosa*) through medication. Such microbe populations consist of metabolically active and metabolically inactive (persistent) subpopulations. It turns out that, for efficient medication strategies, the two goals, eradication of active bacteria on one hand and eradication of inactive bacteria on the other, are in conflict. Using multiobjective optimization, we obtain a survey of the full spectrum of best solutions. We find that, if treatment time is limited and the total medication dose is constant, the application of the medication should be concentrated both at the beginning and end of the treatment. If the treatment time is increased, the medication should become increasingly spread out over the treatment period until it is uniformly spread over the entire period. The transition between short and long overall treatment times sees optimal medication strategies clustered into groups.

Index Terms—Biochemistry, biological systems, biology, biomedical, chemistry, computational bioinformatics, computational intelligence, ecology, evolutionary biology, evolvable hardware, game theory, mathematics.

I. INTRODUCTION

A. Problem of Bacterial Persistence

NOT ONLY hibernating mammals or sporing fungi reduce or stop their metabolic activities, also some microbial organisms are known to randomly slip into and out of “hibernation.” This is essentially characterized by reduced metabolic activity and reduced or suspended reproduction. The disadvantage of reduced population growth, goes hand-in-hand with the advantage of reduced vulnerability to drugs, rendering “hibernating” bacteria persistent in the face of medication treatments [2], [8], [12], [17], [19]. Bacterial population can therefore consist of genetically identical active and persisters subpopulations.

From a human point of view, be it medical or pest control, the presence of persisters can have serious consequences.

Manuscript received April 24, 2009; revised September 15, 2009 and November 25, 2009. Date of publication March 29, 2010; date of current version October 1, 2010.

O. Steuernagel is with the School of Physics, Astronomy and Mathematics, University of Hertfordshire, Hatfield AL10 9AB, U.K. (e-mail: o.steuernagel@herts.ac.uk).

D. Polani is with the School of Computer Science, University of Hertfordshire, Hatfield AL10 9AB, U.K. (e-mail: d.polani@herts.ac.uk).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TEVC.2010.2040181

Bacterial persistence was first observed in *Staphylococcus* when, in 1944, Bigger [2], [12] noted that penicillin did not always kill all exposed bacteria although sufficient toxicity was established. Based on the observation that bacteria surviving penicillin treatment were no less susceptible than their ancestors, it was concluded that heritable bacterial resistance was not involved but persistent behavior could explain such a finding; this has recently been reconfirmed [2].

Bacterial persistence can occur irrespective of environmental conditions [2], [17] and is widespread [2], [7], [8], [11], [12], [16], [17]. It also appears in viruses, which can become persistent by integrating into their host’s genome and suspending production of virus particles, as exemplified by human immunodeficiency virus (HIV), herpes, and the bacteriophage lambda.

The tradeoff between the persisters’ growth-underperformance under benign conditions on one hand, and the wipeout of all active organisms in the case of a catastrophe on the other, leads to a small random subpopulation of persisters individually “bet-hedging” to switch into a persistent state [9] thus effectively establishing a “life-insurance” [13], [14]. Persistence can bring a species “back from the brink,” even when sequences of sudden catastrophes occur, because, in all likelihood, a few persisters will have stayed out of harm’s way [14]. It is thus relevant in disease prevention [12] and requires new treatment regimes [19]. *Escherichia coli* (*E. coli*), MRSA-“superbug”, *Mycobacterium tuberculosis*, and *Pseudomonas aeruginosa* show persistence [2], [12], [17], possibly extending to airborne infectants [11]. Persistence also appears to be important in ecological scenarios [7], [16] and latent *HIV-1* infections [8].

In *E. coli*, the conversion rate from the active subpopulation to the persistent form and the reverse rate have been shown to be independent of environmental factors [2], [12]. In other words, no sensorial input about the quality of the environment is used to trigger the conversion from one to the other. In this paper, we only consider this type of persistence. It is an effective strategy for organisms, which face life in environments where sudden devastating degradation and recovery is an acute possibility and moreover dispenses with the need to maintain sensors for surveying the environment—an important advantage for primitive organisms [14].

B. Our Approach

In this paper, we primarily intend to highlight the features of multiobjective optimization and its applicability to

problems in the life sciences. Our model for the eradication of persistent pathogens shows that eradication of persisters and normal pathogens form conflicting objectives which are best approached using multiobjective optimization. Our model is not intended to quantitatively represent a specific biological or clinical system but to investigate the problem in general terms. Multiobjective optimization [4] has been rarely applied to problems in the life sciences [10], [15]. It is, however, becoming clear that the benefits of using multiobjective optimization in the life sciences could be considerable [10].

We adopt population dynamical models based on coupled differential equations [2], [13], [14], [17] for the numerical study of the behavior of persistent pathogens exposed to different medication strategies (we use the terms “medication” and “drug” interchangeably, they stand for the presence of any hazardous entity killing the pathogens, such as radiation, chemicals, antibodies, etc.).

We assume that only two subpopulations are present, active (normal) pathogens that grow at a normal rate, and are susceptible to the medication, and persisters (such as type II in [2]) that grow more slowly and are less susceptible. We assume that the subpopulations are so large that discreteness of population sizes can be neglected. We denote the sizes of the normal and the persistent subpopulations by the time-dependent functions, $n(t)$ and $p(t)$, respectively. This continuous description allows us to employ continuous differential equations (in the “deterministic limit” [13]) which are readily integrated using a computer.

Initially, we will confirm mathematically that the best approach to the eradication of nonpersistent multiplying pathogens is their immediate extermination by as strong a medication dose as possible. Whereas this case is intuitively easy to understand, matters become much more complicated when persistence is taken into account.

The slowdown or shutdown of the persisters’ metabolism protects them from medication. One therefore has to retain some medication to be administered some time after the first dose of medication was applied. This helps to exterminate persistent pathogens that have bypassed the biocidal effects of the initially given medication and subsequently revert back to their active state.

For such followup action neither very long waiting times are allowed, because the surviving active pathogens multiply and thus hurt the host and replenish the persisters’ reservoir, nor is immediate followup medication advised; otherwise the persisters have not had enough time to come out of the persistent state and so the medication hurts the host more than the pathogens.

Neither intuitive nor analytical solutions for this problem are available, we therefore choose a model in which a course of treatment consists of the administration of N equal units of the drug (we choose $N = 10$). The total amount of drug applied during a course of treatment is fixed. The course of treatment extends over a fixed time interval, spanning from the initial time $t = 0$, of the pathogens’ detection, to the time $t = T$ when the final outcome of the treatment is evaluated. Within this interval, times for the individual drug administrations t_k , with ($k = 1, \dots, N$), are chosen freely. Different medication

scenarios, i.e., the effect of different distributions of the N administration times $\{t_k\}$ are compared for their effectiveness. The objective of the treatment is the minimization of the sizes of the normal, $n(T)$, and the persistent subpopulation, $p(T)$, at the end of the course of treatment.

Although we perform multiobjective optimization we fix this treatment time, T , beforehand. One can, of course, generalize our approach to include variable treatment times as well, thus having to consider three objective variables, namely, $n(T)$, $p(T)$, and T . Then, our problem space would be 3-D and the set of best solutions would form a complicated 2-D hypersurface embedded in it: too rich a system for an introductory treatment of our method. We thus consider a 2-D problem space $[n(T), p(T)]$ and the family of best solutions that form a 1-D hypersurface within.

We will, toward the end of this paper consider the general trends of our model system’s behavior when the total treatment time T is varied as well.

Aside from this simplicity issue, there are two more good reasons to fix the total treatment time T beforehand.

We assume that cumulative toxicity of the medication is the major constraint regarding its application (this is reasonable for scenarios, such as radiation therapy, many types of medication treatments and for agricultural and other such environmental scenarios). With a cumulative dosage constraint, medication strategies must not be drawn out too much in time since the medication becomes overdiluted, see below. We therefore arrive at a natural upper limit for the total treatment time T .

If there is no time constraint, and if one makes sure that the medication does not become overdiluted, drawn out medication regimes where the medication is administered at a constant rate throughout the treatment T show the greatest suppression of $n(T)$ and $p(T)$, see Fig. 3. But this kind of treatment regime can become unstable due to the danger of overdilution (see caption of Fig. 3) and is also harder to adhere to than treatment regimes of fixed shorter time.

In light of the fact that imperfect patient adherence to medication strategies is of considerable concern [1], we thus conclude that there are the following reasons to consider fixed total treatment times T : simplicity, safety, and practicality.

C. Problem Space and the Pareto Front

Different medication strategies lead to different final results. When the times at which a medication dose is administered is continuously changed the outcome changes continuously as well. Therefore, the problem space consists of a connected area of feasible solutions outside of which lies the region which cannot be reached by any feasible solution; because, say, perfect or near perfect suppression of the pathogens’ numbers is beyond the eradication power of the medication. An example in our model would be the origin $[n(T), p(T)] = [0, 0]$ and its immediate neighborhood. This area cannot be reached because the differential equations used in our model only allow for exponential suppression of the population, not complete eradication (on the issue of complete extermination due to fluctuations see [13]).

The most interesting area is the boundary that lies between feasible and unfeasible solutions for small numbers of $n(T)$

and $p(T)$, because it contains the optimal cases of what is feasible. The boundary can have a complicated shape, see, for instance, Figs. 10 and 13 of [18].

The *Pareto front* contains all those points of the boundary for which there are no other points which allow for solutions that are *simultaneously* better or equal with respect to all optimization objectives; it only contains “nondominated” solutions. Since the boundary can have a complicated form, the subset of Pareto-optimal points can be discontinuous, see Figs. 11(c) and 14(b) of [18]. Typically, a continuous Pareto front in 2-D has a shape like the curve shown in Fig. 2(b), also, compare Figs. 9, 10, and 13 in [18].

In principle, using conventional single-objective optimization and changing all available relative weight factors that were used to combine several objectives artificially into a single one allows us to find the Pareto front as well [18]. However, for practical reasons this modification of single-objective optimization is unfeasible because many solutions are being entirely missed, see Fig. 14(c) of [18]. Using multiobjective optimization allows us to gain the advantage of being able to explore the entire set of optimal solutions [18].

D. Our Model

For transparency we also employ the following simplifications.

The active subpopulation, $n(t)$, grows at a constant rate μ_n leading to exponential growth, whereas the persisting subpopulation, $p(t)$, grows at a substantially lower rate μ_p [13] which we set to zero for simplicity (without affecting our basic conclusions). We, similarly, neglect the (greatly reduced) kill rate of persisters in the presence of medication [2], [3].

The subpopulations convert into each other at constant rates a and b [2], [13], although these rates may depend on environmental conditions [9]. We assume that only the active subpopulation is being decimated by the medication: we assume its power to kill to be proportional to the drug concentration, $c(t)$, [3] (although nonlinear threshold behavior has been observed as well [3]—in which case other assumptions such as zero growth of the persisters may have to be reviewed). We therefore arrive at the following system of coupled ordinary differential equations for the behavior of the subpopulations as functions of time

$$\frac{dn(t)}{dt} = (\mu_n - c(t) - a) \cdot n(t) + b \cdot p(t) \quad (1)$$

$$\frac{dp(t)}{dt} = a \cdot n(t) - b \cdot p(t). \quad (2)$$

E. Our Assumptions

In what follows we will assume that the total administered medication dose is fixed. This assumption is motivated by the cumulative toxicity of medical treatments. Our approach can be adapted accordingly, if avoidance of peak values of the drug concentration is the primary concern.

In general, the concentration of the drug, $c(t)$, could be given by any nonnegative function. In accord with our approach, we model each administered drug dose by the same Gaussian peaks (bell-shaped curves) with equal strength D_0 , centered

on the respective administration times t_k , a treatment course, $D(t)$, is thus described by the sum

$$D(t) = \sum_{k=1}^N D_0 \frac{\exp\left[\frac{-(t-t_k)^2}{\sigma^2}\right]}{\sqrt{\pi}\sigma}. \quad (3)$$

Here σ scales the widths of the Gaussians (compare Fig. 1) and the normalization factor $1/(\sqrt{\pi}\sigma)$ assures that each peak is of unit strength ($\int_{-\infty}^{\infty} dt \exp\left[\frac{-(t-t_k)^2}{\sigma^2}\right] / (\sqrt{\pi}\sigma) = 1$).

We assume that the drug is cleared out of the system at a constant rate R (in units of h^{-1}), its concentration, $c(t)$, thus obeys the differential equation

$$\frac{dc(t)}{dt} = D(t) - R c(t). \quad (4)$$

For the drug concentration this yields

$$c(t) = c_0 + \int_0^t d\tau D(\tau) e^{-R(t-\tau)} \quad (5)$$

with the assumed initial value $c_0 = 0$, i.e., no medication is present before the treatment starts at time zero. Note that small values of the drug-clearance rate R lead to prolonged presence of the medication and thus to a greater cumulative effect since the accumulated medication dose

$$C(T) = \int_0^T dt c(t) \quad (6)$$

scales with R^{-1} , just like the total integrated dose

$$C_\infty \doteq \int_{-\infty}^{\infty} dt \int_{-\infty}^t d\tau D(\tau) e^{-R(t-\tau)} = N D_0 / R. \quad (7)$$

This expression C_∞ ignores the initial value assumption $c_0 = 0$ and therefore has a simple and transparent form. Because it includes the small tails of the medication distribution extending beyond of the treatment time interval $[0, T]$ it slightly overestimates the value of the total cumulative administered dose, $C(T)$.

Scenarios that can be described by (1) and (2) include a bacterial infection by a persistent species which is being fought with drugs [term “ $-c(t) \cdot n(t)$ ” in (4)] where the drug degrades over time [say, by excretion, aging, or evaporation: term “ $e^{-R(t-\tau)}$ ” in (7)].

Note that our assumption regarding a finite number of administered doses N , in (3), is the generic way in which medication is released (in radiation treatment or pest control with agricultural aircraft, continuous administration may be altogether unfeasible). Continuous medication (drip-feed) can easily be emulated with our model using a large number, N , of doses shots.

Without medication ($c(t) = 0$) the system (1)–(2) has constant coefficients and is therefore analytically solvable with the general solution $[n(t), p(t)] = w_+ \vec{e}_+ \exp[\lambda_+ t] + w_- \vec{e}_- \exp[\lambda_- t]$. Here, the two eigenvector $\vec{e}_\pm = [\mu_n - a + b \pm \sqrt{(a+b)^2 + \mu_n(\mu_n - 2a + 2b)}, 2a]$ are associated with eigenvalues $\lambda_\pm = [b\mu_n \pm \sqrt{b^2\mu_n^2 + 4(\mu_n - b - a)}]/2$. The component $w_+ \vec{e}_+$ will quickly outgrow its counterpart $w_- \vec{e}_-$ because $\lambda_+ > \lambda_-$. The wild-type without the influence of medication is therefore typically well described by the state $w_+ \vec{e}_+ \exp[\lambda_+ t]$. This implies that the generic initial ratio of

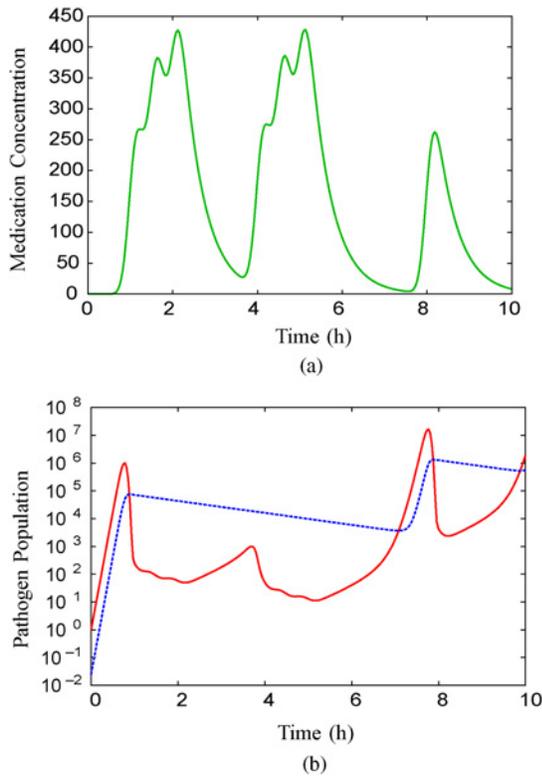


Fig. 1. Response of pathogen population to medication. (a) Drug concentration, $c(t)$, as a function of time. $N = 7$ dose units, of strength $D_0 = 430$ each, are administered at times 1 h, 1.5 h, 2 h, 4 h, 4.5 h, 5 h, 8 h. The widths of the peaks is $\sigma = 0.2$ h and the drug-clearance rate is $R = 2 \text{ h}^{-1}$. (b) Time evolution of active organisms $n(t)$ (red solid line) and persisters $p(t)$ (blue dotted line) with parameters $\mu_n = 20 \text{ h}^{-1}$, $a = 0.5 \text{ h}^{-1}$ and $b = 0.5 \text{ h}^{-1}$, note that the persister responses are delayed in time.

active to persistent pathogens is given by the ratio of the components of the wild-type \bar{z}_+ , namely

$$\frac{n(0)}{p(0)} = \frac{\mu_n + b - a + \sqrt{(b+a)^2 + 2\mu_n(b-a) + \mu_n^2}}{2a}. \quad (8)$$

We therefore choose $n(0) = 1$ and $p(0)$ in accordance with (8) as a *natural initial* condition which models pathogens found in their natural infection habitat. If they are found under very different circumstances, such as bacteria residing in a bacterial biofilm, the initial fraction of persisters can be much higher [12] than assumed here.

II. RESULTS

Equations (1) and (2) do not, in general, allow for an analytical solution. This is why we investigate them numerically. A typical scenario is portrayed in Fig. 1. It illustrates that active organisms, $n(t)$, primarily get killed by the medication whereas the inactive ones, $p(t)$, primarily suffer losses due to conversion into active ones and regain numbers when the active ones recover. The two subpopulations sustain each other.

The goal is to push the entire pathogen population toward its possible extinction (i.e., to such small numbers that action of the host's immune system or random fluctuations can wipe it out [13]).

A. Fighting Nonpersistent Pathogens

Without persistence all pathogens are affected by the medication and should be killed immediately. This can be shown formally: assuming the infection is discovered at time zero, an integration of (1) yields $n(T) = n(0) \exp[\int_0^T (\mu_n - c(\tau)) d\tau] = n(0) \cdot e^{E(T)}$, where the effective exponent

$$E(T) = \mu_n \cdot T - C(T) \quad (9)$$

contains the accumulated medicine dose, $C(t)$, of (6). Maximal suppression of the pathogen population requires the largest achievable negative values of $E(T)$: the positive growth term " $\mu_n \cdot T$ " has to be minimized. This shows the medication has to be given immediately.

The effective exponent also yields the condition, $E(T) = 0$, which estimates where the medication just balances pathogen growth. Assuming, as above, that $C(T) \approx C_\infty$, we find that for values of T surpassing

$$T_{\max} = \frac{ND_0}{\mu_n R}. \quad (10)$$

$E(T)$ becomes positive and pathogen growth is no longer kept in check. A fixed total dosage C_∞ thus implies a natural constraint on the total treatment time beyond which drug overdilution renders treatments ineffective.

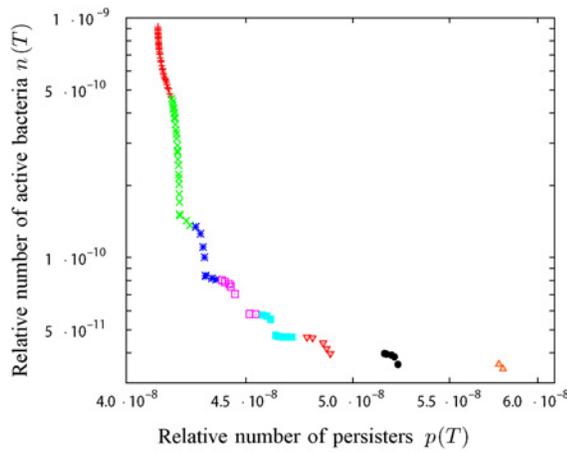
B. Fighting Persistent Pathogens: Pareto Front

Transition into and out of the persistent state ($a, b > 0$) allows pathogens to avoid the effects of medication and shortens the effective maximal stalemate-time considerably, thus our estimate for T_{\max} , derived for the case of nonpersisting pathogens, only establishes an upper bound on the permissible total treatment time for an effective treatment of persisters.

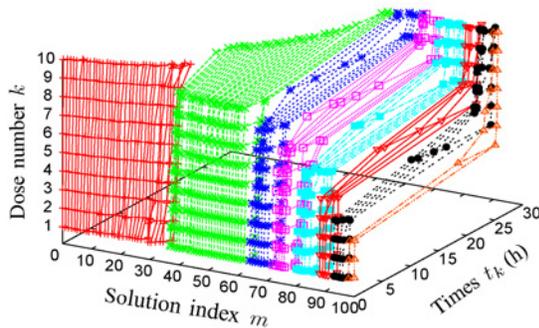
Due to their persistence ($a, b > 0$) pathogens show a delayed response [2] (compare Fig. 1) which complicates their eradication. We now compare different eradication strategies. First, values for D_0 and R (keeping the total effective dose C_∞ constant), and a fixed total treatment time T are chosen. Then we vary the (ten) medication times ($\{t_k\}, k = 1, \dots, 10$) thus modifying the dosage strategies [choice of time-points $t_k \in [0, T]$ in (3)]. Upon integration of (1) and (2), using the fourth order Runge–Kutta method, we determine the number of survivors $n(T)$ and $p(T)$ as our quality criterion.

The delayed response leads to a tradeoff between eradication of active versus persister subpopulations, this complicates the analysis; without further assumptions a best treatment strategy cannot be identified. To map out the solution space, we therefore perform multiobjective optimization [4], using the nondominated sorting genetic algorithm (NSGA)-II [6]. We determine the set of Pareto-optimal strategies: in an $n(T)$ -over- $p(T)$ plot they form a Pareto-optimal front [4] of points corresponding to dosage strategies that lead to simultaneously minimized (nondominated) final values of $n(T)$ and $p(T)$.

For the integration of the differential equation, we use a regular fourth-order Runge–Kutta integrator, with a step length of $\Delta t = 0.01$. The integration was found to provide consistent results for all runs up to a step length of at least $\Delta t = 0.015$, thus ensuring numerical stability of the employed integration routine.



(a)



(b)

Fig. 2. Typical example of optimal treatment strategies. Treatment time $T = 30$ h, bacterial parameters of *E. coli* wild-type [13] (except for our choice of $\mu_p = 0$): $\mu_n = 2 \text{ h}^{-1}$, $a = 1.2 \times 10^{-6} \text{ h}^{-1}$, $b = 0.1 \text{ h}^{-1}$. We choose $n(0) = 1$ and $p(0) \approx 5.714 \times 10^{-7}$ in accord with the natural initial condition (8). Medication parameters $D_0 = 100$, $N = 10$, $\sigma = 10$ h and $R = 0.2 \text{ h}^{-1}$. (a) Pareto front of optimal strategies shows the tradeoff between suppressing active and persister subpopulations. The response margin for suppressing persisters is relatively much narrower than that for suppressing active pathogens (the spread of values on the horizontal coordinate axis is small). (b) Same treatment regimes, plot displays dose number k administered at time t_k over total time t and index m (representing the solution number—out of 100) of solutions found by the NSGA-II algorithm sorted with respect to increasing values of surviving persisters $p(t)$. Note the emergence of distinct steps separating groups of treatment regimes despite a large value of σ (this distinction becomes clearer still for smaller values of σ). The color- and symbol-coding matches subsections of the Pareto front in (a) with the treatment regimes displayed in (b).

The NSGA-II multiobjective optimization algorithm [6] with a population size of 100 was used, running for 500 generations, using the final populations $n(T)$ and $p(T)$ as the two objectives. The objectives were constrained to nonnegative values to prevent the genetic algorithm from being caught in spurious numerical instabilities. The algorithm optimized the medication times $\{t_k\}$ ($k = 1 \dots 10$ in our case), the crossover probability was set to 0.9 and the mutation probability to 0.1. The parameters η_c and η_m for the polynomial distributions used in the SBX crossover and in the mutation operator [5], [6], were both set to 16.

The precise choice of these parameters turned out to be uncritical. We found that the results reported below are robust with respect to variations of crossover and mutation probabil-

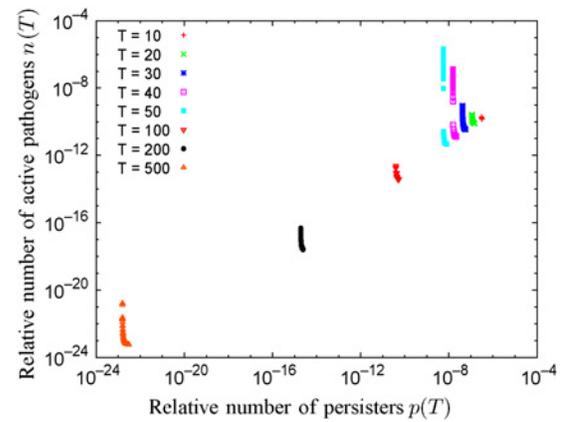


Fig. 3. Collection of several Pareto fronts for various treatment times T shows increased effectiveness of longer treatments. Same parameters as in Fig. 2 except for treatment times ranging from $T = 10 \text{ h}, \dots, 500 \text{ h}$, see legend. We observe breakdown of treatment at $T \approx 600 \text{ h}$ (not shown) because medication becomes too much diluted.

ities from 0.01 to 0.99, similarly the η_c and η_m parameters could be varied between 1.6 and 160 without affecting the position of the Pareto front. Typically, the essential features of the Pareto front emerged reproducibly after approximately 50 generations, whereas the remainder of the run served to fine-tune the precise features of the front and the corresponding solutions. Only for extreme choices of the parameters, namely crossover and mutation probabilities equal to 0.01 and very narrow SBX characteristics, $\eta_c = 160$, $\eta_m = 160$, was the extent of the Pareto front covered significantly more slowly; apart from such extreme choices only insignificant performance differences could be observed.

We now discuss the features of the solutions in detail.

At one end of the Pareto front one finds the strongest suppression of persisters, at the other end the strongest suppression of active bacteria, compare Fig. 2(a).

Fig. 2(b) shows that a strategy aiming to suppress the persistent subpopulation requires early administering of large doses of medication. This is due to the fact that the active population has to be suppressed early on and then some medication has to be used to hold the persister's in check that are “waking up” and become active again. Alternatively, when the primary strategic aim is the suppression of the active subpopulation or a mixed strategy, a later application of the bulk of the medication is advised, although (depending on details) some medication should also be given at the start (as soon as the infection is discovered). Associated optimal strategies may therefore be very different from strategies which aim for uniform constant exposure to medication, or, from the “kill before they multiply” strategy described above. One should note the emergence of discrete “bands” of optimal strategies in Fig. 2(b).

Although the relative population suppression factor due to the medication treatment may be satisfactory in the example sketched in Figs. 2(a) and 3, we are clearly most interested in the critical cases where (because of cost, toxicity, or other reasons) the treatment is in danger of failing. In this context, it should be pointed out that, worryingly, the narrow response

margin of final persisters subpopulations $p(T)$ is a generic feature (see Fig. 2). We are thus led to consider a variation of the total treatment time T as well. Fig. 3 displays several Pareto fronts for different values of T . Each individual front shares the features displayed in Fig. 2. Overall, there is a trend to more effective treatments with lengthened treatment time in which case the medication has to be administered more uniformly over the entire treatment period. It must not be lengthened too much though, because the medication would become too diluted (see discussion leading up to expression (10) above).

III. DISCUSSION

The overall treatment time T has to be sufficiently long to kill persisters which are protected by the time lag in the pathogens' response dynamics but short enough not to dilute the medication concentration too much. If short treatment times can yield sufficient pathogen suppression, the use of such strategies may well be safer, since they lead us away from a possible breakdown due to overdilution.

Also, as one shortens T , the characteristics of optimal strategies change: instead of being uniformly distributed over T , doses are typically increasingly concentrated at the beginning and/or end of the treatment time, compare Fig. 2(b). Such strategies do not only need less time than drawn out therapies, but they are also simpler to administer. We believe that in view of widespread problems with patient adherence to long lasting medication regimens [1] such optimized strategies may offer relevant alternatives that deviate from current clinical practice. In this context, we would like to point out that these treatment regimes appear to be quite stable with respect to small changes in strategy, such as completely concentrating all medication toward beginning and end of the treatment period. In other words, a judiciously dosed two-shot approach, almost the simplest conceivable strategy, can yield nearly optimal results, is shorter than a maximally drawn out therapy and not in danger of failure due to overdilution of the medication.

For sufficiently large medication doses, the qualitative results reported above also apply to the case of amplified persistence (as is the case for the high persistence (*hip*) mutants of *E. coli* analyzed in Balaban *et al.* [2]). When persistence is increased, but a simultaneous increase in medication is impossible, the results can change dramatically. An eradicable disease can become unstoppable. To illustrate this point, we compare the response of the wild-type of *E. coli*, Fig. 3, with its highly persistent *hipQ*-type twin [2], [13] in Fig. 4. We assume the two to be identical except for their different persistence rates a and b [13]. Under identical treatment and initial conditions as for the wild-type, we let the search algorithm find the modified Pareto front. Fig. 4 shows that short, highly concentrated treatments allow us to suppress the active subpopulation but they are too short to affect the persisters. For longer treatments, the persisters' reflux rate b , back to the normal state, is still too low to deplete them sufficiently, the pathogen has become untreatable.

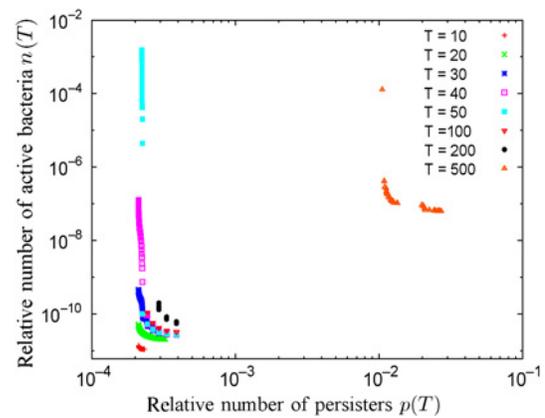


Fig. 4. Mutant pathogen infection is incurable by the approach displayed in Fig. 3. A collection of several Pareto fronts for various treatment times T (h), see inset, using the same parameters as in Figs. 2 and 3. This includes use of the same initial conditions as in the wild-type (which are not the equilibrium conditions from (8) to allow for comparison); only $a = 10^{-3} \text{ h}^{-1}$, $b = 10^{-5} \text{ h}^{-1}$ are altered to describe the *hipQ*-variant of *E. coli* instead of its wild-type [13].

When our simplifying assumption that persisters are entirely resistant to medication is modified, in favor of reduced susceptibility to medication, an extra term of the form “ $-\omega \cdot c(t) \cdot p(t)$ ” has to be added to the right hand side of (2). With a reasonable factor of the order of $\omega \approx \mu_p / \mu_n$ (≈ 0.1 in the case of *E. coli* [13]), our model still displays similar generic features for optimal treatments. Regimes still form groups of distinct strategies, doses for optimal treatment over intermediate lengths T are still administered early and late. The greatest difference is due to the greater vulnerability of the bacteria (i.e., their smaller overall survival rates).

Finally, for other, different conditions and scenarios, such as modified toxicity behavior, nonlinear dose-response [3] and modified quality criteria, one can also use multiobjective optimization to explore and map the pertinent optimality regimes.

ACKNOWLEDGMENT

The authors would like to thank K. Deb for providing them with the code for the NSGA-II algorithm and his helpful comments. They are also grateful to F. van den Berg, T. Aldsworth, and J. T. Kim for their discussions, critical reading of the manuscript, and suggestions for improvement.

REFERENCES

- [1] A. Atreja, N. Bellam, and S. R. Levy, “Strategies to enhance patient adherence: Making it simple,” *Medscape Gen. Med.*, vol. 7, no. 1, p. 4, 2005.
- [2] N. Q. Balaban, J. Merrin, R. Chait, L. Kowalik, and S. Leibler, “Bacterial persistence as a phenotypic switch,” *Science*, vol. 305, pp. 1622–1625, Sep. 2004.
- [3] V. Boonkitticharoen, J. C. Ehrhardt, and P. T. Kirchner, “Quantification of antibiotic drug potency by a two-compartment radioassay of bacterial growth,” *Antimicrob. Agents Chemother.*, vol. 34, no. 6, pp. 1035–1040, Jun. 1990.
- [4] K. Deb, “Evolutionary algorithms for multicriterion optimization in engineering design,” in *Evolutionary Algorithms in Engineering and Computer Science*, K. Miettinen, M. M. Mäkelä, P. Neittaanmäki, and J. Periaux, Eds. Chichester, U.K.: Wiley, 1999, ch. 8, pp. 135–161.

- [5] K. Deb and R. B. Agrawal, "Simulated binary crossover for continuous search space," *Complex Syst.*, vol. 9, pp. 115–148, Apr. 1995.
- [6] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," *IEEE Trans. Evol. Comput.*, vol. 6, no. 2, pp. 182–197, Apr. 2002.
- [7] P. A. del Giorgio, Y. T. Prairie, and D. F. Bird, "Coupling between rates of bacterial production and the abundance of metabolically active bacteria in lakes, using CTC reduction and flow cytometry," *Microb. Ecol.*, vol. 34, no. 2, pp. 144–154, Sep. 1997.
- [8] D. Finzi, J. Blankson, J. D. Siliciano, J. B. Margolick, K. Chadwick, T. Pierson, K. Smith, J. Lisziewicz, F. Lori, C. Flexner, T. C. Quinn, R. E. Chaisson, E. Rosenberg, B. Walker, S. Gange, J. Gallant, and R. F. Siliciano, "Latent infection of CD4⁺ T cells provides a mechanism for lifelong persistence of HIV-1, even in patients on effective combination therapy," *Nature Med.*, vol. 5, no. 6, pp. 512–517, Jun. 1999.
- [9] A. Gardner, S. A. West, and A. S. Griffin, "Is bacterial persistence a social trait?" *PLoS One*, vol. 2, no. 1, p. e752, Aug. 2007.
- [10] J. Handl, D. B. Kell, and J. Knowles, "Multiobjective optimization in bioinformatics and computational biology," *IEEE Trans. Comp. Biol. Bioinform.*, vol. 4, no. 2, pp. 279–292, Apr. 2007.
- [11] M. Hernandez, S. L. Miller, D. W. Landfear, and J. M. Macher, "A combined fluorochrome method for quantitation of metabolically active and inactive airborne bacteria," *Aerosol Sci. Technol.*, vol. 30, no. 2, pp. 145–160, Feb. 1999.
- [12] I. Keren, D. Shah, A. Spoering, N. Kaldalu, and K. Lewis, "Specialized persister cells and the mechanism of multidrug tolerance in *Escherichia coli*," *J. Bacteriol.*, vol. 186, no. 24, pp. 8172–8180, Dec. 2004.
- [13] E. Kussell, R. Kishony, N. Q. Balaban, and S. Leibler, "Bacterial persistence: A model of survival in changing environments," *Genetics*, vol. 169, pp. 1807–1814, Apr. 2005.
- [14] E. Kussell and S. Leibler, "Phenotypic diversity, population growth, and information in fluctuating environments," *Science*, vol. 309, pp. 2075–2078, Sep. 2005.
- [15] M. Lahanas, E. Schreibmann, and D. Baltas, "Multiobjective inverse planning for intensity modulated radiotherapy with constraint-free gradient-based optimization algorithms," *Phys. Med. Biol.*, vol. 48, pp. 2843–2871, Aug. 2003.
- [16] P. Lebaron, P. Servais, H. Agogué, C. Courties, and F. Joux, "Does the high nucleic acid content of individual bacterial cells allow us to discriminate between active cells and inactive cells in aquatic systems?" *Appl. Environ. Microbiol.*, vol. 67, no. 4, pp. 1775–1782, Apr. 2001.
- [17] B. R. Levin, "Noninherited resistance to antibiotics," *Science*, vol. 305, pp. 1578–1579, Sep. 2004.
- [18] A. Messac, A. Ismail-Yahaya, and C. A. Mattson, "The normalized normal constraint method for generating the Pareto frontier," *Struct. Multidisc. Optim.*, vol. 25, no. 2, pp. 86–98, 2003.
- [19] C. Wiuff, R. M. Zappala, R. R. Regoes, K. N. Garner, F. Baquero, and B. R. Levin, "Phenotypic tolerance: Antibiotic enrichment of noninherited resistance in bacterial populations," *Antimicrob. Agents Chemother.*, vol. 49, no. 4, pp. 1483–1494, Apr. 2005.



Ole Steuernagel received the Ph.D. degree in physics from Humboldt University, Berlin, Germany, in 1996.

From 1997 to 1998, he was a Research Fellow with the Department of Chemistry, Princeton University, Princeton, NJ. Since 1998, he has been with the School of Physics, Astronomy and Mathematics, University of Hertfordshire, Hatfield, U.K., as a Senior Lecturer in physics. His research interests include the foundations of quantum mechanics and works on classical and quantum optics. He also has a long standing fascination with the life sciences, particularly with questions of mathematical modeling of biological systems.



Daniel Polani received the Doctor of Natural Sciences degree from the Faculty of Mathematics and Computer Science, University of Mainz, Mainz, Germany, in 1996.

From 1996 to 2000, he was a Research Assistant with the Faculty of Mathematics and Computer Science, University of Mainz. In 1997, he was a Visiting Researcher with the Neural Networks Research Group, University of Texas, Austin. From 2000 to 2002, he was a Research Fellow with the Institute for Neuro- and Bioinformatics, University of Lübeck, Lübeck, Germany. From 2002 to 2008, he was a Principal Lecturer with the Adaptive Systems and Algorithms Research Groups, School of Computer Science, University of Hertfordshire (UH), Hatfield, U.K. Since 2008, he has been a Reader in Artificial Life and Member of the Adaptive Systems and Algorithms Research Groups, School of Computer Science, UH. His research interests include the foundations of intelligent behavior in biological as well as artificial agents, especially in terms of optimality principles.