2011 American Control Conference
on O'Farrell Street, San Francisco, CA, USA
June 29 - July 01, 2011

Bacterial Persistence: Mathematical Modeling and Optimal Treatment Strategy

Nicholas G. Cooper and A. Agung Julius
Department of Electrical, Computer and Systems Engineering
Rensselaer Polytechnic Institute
Troy, NY 12180, USA.
Email:coopen3,juliua2@rpi.edu

Abstract— **Bacterial persistence is an epigenetic phenomenon in which some bacteria cells become immune to antibiotic treatment without undergoing genetic mutation. In this paper, we develop a population dynamic model that captures both short term and long term persistence in bacteria. We subsequently pose the problem of designing an optimal treatment strategy, in terms of minimizing the number of persister cells that transition into long term dormancy. We find that the infinite time horizon optimal control strategy is not unique, and it can be expressed as a feedback law using the information about the population sizes of normal and persister cells. We also show the existence of a theoretical lower bound for the optimal cost value.**

I. Introduction

Bacterial persistence is an epigenetic phenomenon, in which some bacteria cells become immune to antibiotic treatment without undergoing genetic mutation[1] [1]. Biologists observed bacterial persistence through a small fraction of cells in a colony that survive the antibiotic attack. Upon removal of the antibiotics, the colony regrows. However, this colony of survivor cells is found to exhibit the same vulnerability to antibiotics as the previous one. This effectively demonstrates the fact that persistence is not inheritable, and hence is an epigenetic trait [1], [2].

Bacterial persistence poses a serious global health problem. In tuberculosis, which is a disease caused by bacterial infection, it is known that persistent bacteria can infect a patient asymptomatically for decades. Approximately eight million people develop active tuberculosis every year, with two millions dying from the disease [3].

A recent finding by Balaban et al [4], [5] suggests that persistence is related to a phenotypic state of slow growth (dormancy), in which the effect of antibiotics on the bacteria cells is minimized. More interestingly, this state of slow growth does not seem to be induced by any external stressor or stimulus. Rather, some cells (termed *type II persisters*) spontaneously and stochastically become dormant. The mechanism behind this stochastic transition to dormancy is not well understood, although some biologists suggest that it might be caused by fluctuations in the expression of toxin-anti toxin proteins in the cells [6], [7].

A. Mathematical Model for Persistence

Some earlier publications in this field put forward a mathematical model for this phenomenon [4], [8], which is based on empirical observation of the population dynamics. The underlying assumptions are: (i) the transitions between dormancy and normal state can be modeled as a two-state continuous time Markov chain (cf. [9]), and (ii) the population growth can be modeled as a linear term (thus, exponential in time). Mathematically, this can be expressed as

$$\dot{n} = \mu_n n - an + bp, \tag{1}$$

$$\dot{p} = \mu_p p + an - bp. \tag{2}$$

Here, n and p denote the population sizes of the normal cells and persister/dormant cells, respectively. The growth rates at both states are given by μ_n and μ_p, while a and b represent the transition rates into and out of dormancy. For wild type *E. coli*, these transition rates are reported to be [4], [8]:

$$a = 1.2 \times 10^{-6} \text{ hour}^{-1},$$
$$b = 0.1 \text{ hour}^{-1}.$$

The growth rates μ_n and μ_p are variables of the growth condition. For example, in favorable growth condition (i.e. in the absence of antibiotics), these rates are reported to be [4], [8]:

$$\mu_n = 2 \text{ hour}^{-1},\ \mu_p = 0 \text{ hour}^{-1}.$$

With antibiotics present, the growth rates become

$$\mu_n = -4 \text{ hour}^{-1},\ \mu_p = -0.4 \text{ hour}^{-1}.$$

From the model, it is obvious that in favorable growth condition, having dormancy is a hindrance to the proliferation of bacteria. However, the model and empirical observation also suggest that having dormancy is beneficial in adverse conditions. In a constantly changing environment, researchers have hypothesized that the optimal growth policy is to have a certain level of dormancy. Kussell et al [8] evaluated the benefit of having dormancy in a periodically changing environment, under the model given in (1) - (2). Acar et al [10] did the same investigation experimentally by putting cells with different dormancy levels in a periodically changing environment. Gardner et al [11] also performed a mathematical analysis of individual fitness from the evolutionary perspective, in a periodically changing environment. The general outcome of these investigations is that depending on the periodicity of the environmental changes, a population with a given level of dormancy has a competitive advantage over other population with lower level of dormancy.

[1]In contrast, resistance is a similar phenomenon where genetic modification is involved.

Despite the agreement between the model (1) - (2) and the experimental observation for bacteria that do not exhibit long term dormancy (i.e. *E. coli*), it does not capture long term dormancy exhibited in others (e.g. *M. tuberculosis*). Long term dormancy is a phenomenon where the bacterial population consists of high levels of persisters and low level of normal cells, allowing the population to remain dormant for decades [3]. Once the host's immunity system weakens the pathogenic bacteria reactivate. In this paper, we propose a hybrid system model that captures long-term dormancy. The underlying assumption is that during antibiotic attack, the transition rate from dormancy to active cells (the symbol b in (1) - (2)) is zero. Further, once the number of active bacteria cells is reduced below a threshold, we assume that the host's native immunity system is activated, and the remaining dormant bacteria become long-term persisters.

B. Infection Treatment Strategy

In the controls community, the problem of finding an optimal infection treatment strategy for a given mathematical model of infection has been previously studied. For example, Stengel et al developed a detailed optimal control analysis of infection treatment using a model of the human immunity system [12]. However, in this work, persistence effect is not included in the model. Other researchers investigated the issue of optimal control in the treatment of HIV infection, using a model that also captures the development of drug resistance (see e.g. [13], [14], [15], [16]). Jung et al developed an optimal control based treatment strategy for tuberculosis, based on a mathematical model of the host population [17].

In this paper, we use the population model for the bacteria including long-term dormancy to devise an optimal control strategy to minimize long-term persisters. The outcome of this analysis underlines the necessity of patterned or interrupted treatment strategy to deal with persistence. Interestingly, when persistence was first discovered in the 1940s, field observation reported in the literature also suggested the benefit of interrupted treatment strategy [1], [18].

II. Model Development for Long-Term Persistence

We modify the standard model in (1) - (2) to include the effect of finite environmental support. This is done by incorporating the Verhulst model that models the competition between individuals as a bilinear term in the dynamics (see e.g. [19]).

$$\dot{n} = \mu_n n - an + bp - kn^2, \tag{3}$$

$$\dot{p} = \mu_p p + an - bp, \tag{4}$$

where, as before, n and p denote the population sizes of the normal/active cells and dormant/persister cells, respectively. The parameter k represents the competition rate among the active cells. In our subsequent analysis, we use the assumption that $\mu_p \approx 0$[4], which means that the persisters do not multiply nor get killed in antibiotic attacks. Equations (3) - (4) represent a planar dynamical system, whose phase-plane can be sketched in Figure 1.

Qualitatively, we can easily observe that the first quadrant is invariant in (3) - (4). Further, from Figure 1, we see that the system (generally) admits two equilibria, one at the origin,

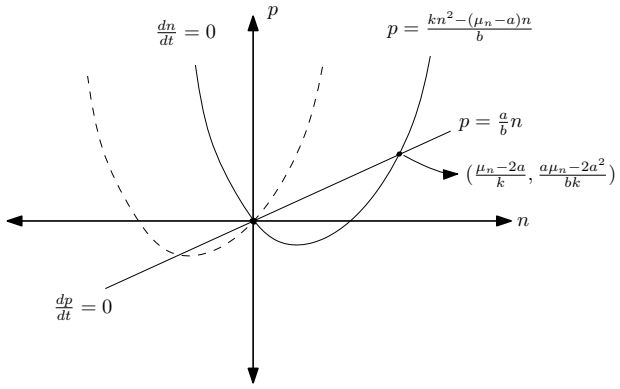

Fig. 1. The phase-plane of (3) - (4). The solid lines represent the nullclines of the dynamics in favorable growth condition ($\mu_n > 0$). During antibiotic attack, $\mu_n < 0$. The nullcline $\dot{n} = 0$ in this case is represented by the dashed line.

and another one that depends on the system parameters. Given that typically $\mu_n \gg a$, the second equilibrium always lies in the first quadrant. The Jacobian of the dynamics in (3) - (4) is given by

$$J(n,p) = \begin{bmatrix} \mu_n - a - 2kn & b \\ a & -b \end{bmatrix}.$$

Analyzing the eigenvalues of $J(n,p)$ at the two equilibria, we find that the origin is an unstable equilibrium, whereas $\left(\frac{\mu_n - 2a}{k}, \frac{a\mu_n - 2a^2}{bk}\right)$ is stable. This means that any positive initial condition will result in a dynamics that converge to the maximum carrying capacity of the environment. Our model is thus able to capture both the exponential growth phase and the stationary growth phase of bacterial population.

Our goal is to develop this model into one that can capture the phenomenon of long-term persistence in infections such as tuberculosis. Primary tuberculosis develops within 1 or 2 years after initial infection [3]. Upon treatment, the pathogen can remain in a state of asymptomatic infection. Post-primary tuberculosis typically develops much later in the host's life. This can be caused by reactivation of the remaining bacteria from the initial infection [3]. Currently, it is estimated that 2 billion people are infected with *M. tuberculosis*, the pathogen that causes tuberculosis [20]. In developing our model, we assume that the normal portion of the bacteria population is virulent, while the dormant one corresponds to asymptomatic infection.

An antibiotic attack can be modeled as a change in the value of μ_n to a negative value. This is in accordance with the experimental observation and the model in (1) - (2) (cf. [4], [8]). In this case, there is no equilibrium in the first quadrant (see Figure 1), and the origin is a stable equilibrium. Physically, this corresponds to the prediction that all bacteria cells will eventually be killed by the attack. However, this model is not able to capture long-term dormancy, where persisters can survive for decades. To see this, observe that although persister cells are not susceptible to the attack, they are converted to normal cells in a process with time constant $\frac{1}{b}$. According to the reported parameter value, this time constant is in the order of 10 hours, which is several orders of magnitude lower than the time scale of long-term

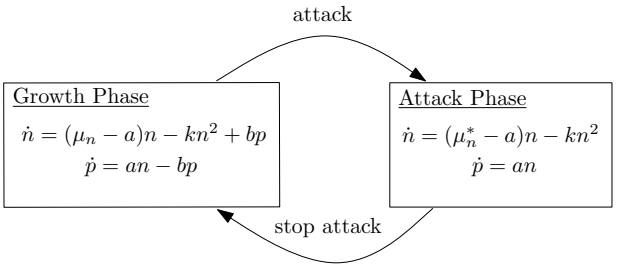

Fig. 2. A hybrid system model allows survival of persister cells. This model is obtained by modifying (3) - (4). The two events `attack` and `stop attack` corresponds to the beginning and the end of the antibiotic treatment period.

persistence.

In order to allow the persisters to survive antibiotic attacks, we further modify (3)-(4) into a hybrid system model [21]. This is shown in Figure 2. The underlying assumption is that during antibiotic attack, both μ_n and b assume new values. The new growth rate is denoted as μ_n^*, which is a negative number. The new value for the transition rate b is effectively 0.

The planar dynamical system during the attack phase has a continuum of equilibria, given by $n = 0$. Linearizing the dynamics around this line, we find that the equilibria are stable. Thus, any initial condition in the first quadrant will converge to $n = 0$. However, this convergence is asymptotic, which means that any dose of antibiotics (represented by arbitrarily negative but finite μ_n^*) given during a treatment period of any finite length will not reduce n to 0. Further, the model also predicts that upon cessation of this treatment period, the infection will return to its full scale. This corresponds to the dynamics in the growth phase, where the trajectory converges to the stable equilibrium in the first quadrant. Therefore, the model given in Figure 2 cannot capture the fact that the normal cell population remains insignificantly low during long-term persistence.

We further enhance the model by including a third discrete location, whose dynamics represents the condition when n is very low. In this condition, we assume that the amount of virulent bacteria is low enough that host's native immunity system can take over the attack against the bacteria. This is shown in Figure 3. We introduce a third mode of dynamics that represents long-term dormancy/persistence. We assume that the transition to this mode automatically happens when the guard condition $n \leq n_{\text{immune}}$ is satisfied. During long-term dormancy, normal cells will continue to decrease and transition to persistence. However, persister cells do not transition back to normal cells. Stopping the treatment in this mode does not have any effect on the dynamics, since the host's native immunity system is mounting an attack on the bacteria. The system can transition from long-term dormancy to growth phase again when the host's immunity system is weak. This is modeled by the event `weak immunity`. The `weak immunity` event is external to the system and assumed uncontrollable. During long term dormancy the number of active and persister cells is relatively constant; therefore, the number of persisters entering the Growth Phase through the `weak immunity` event will be the number of

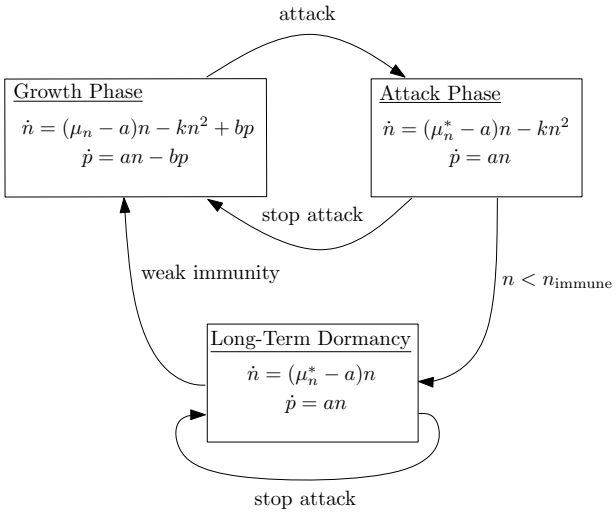

Fig. 3. A hybrid system model that captures long-term persistence and post-primary infection.

persisters entering the Long Term Dormancy Phase through the $n < n_{immune}$ event.

III. Optimal Treatment Strategy

In designing the optimal treatment strategy, the goal is to *minimize the number of persisters upon transition into long-term dormancy*. Given the hybrid system model shown in Figure 3, a treatment strategy is defined as the timing of `attack` and `stop attack` events, which corresponds to the scheduling of the treatment. Therefore, if we define $t_{a,1} < t_{a,2} < \cdots$ as the times when the *attack* event happens, and $t_{s,1} < t_{s,2} < \cdots$ as the times when the `stop attack` event happens, where

$$\forall i \in \{1, 2, \cdots\},\ t_{a,i} < t_{s,i}, \tag{5}$$

then the design of optimal treatment strategy can be formulated as follows:

Problem: Given an initial condition $n(0) > n_{\text{immune}}$ and $p(0)$ in the growth mode. Minimize $p(t_{\text{dorm}})$, where

$$t_{\text{dorm}} = \inf\{t > 0 \mid n(t) < n_{\text{immune}}\}, \tag{6}$$

with $\{t_{a,i}\}_{i\in\{1,2,\cdots\}}$ and $\{t_{s,i}\}_{i\in\{1,2,\cdots\}}$ as the optimization variables.

To illustrate the affect of treatment scheduling, we perform a few numerical simulations. Consider the model in Figure 3, with parameters as follows:

$$\mu_n = 2,\ a = 10^{-6},\ b = 0.1,$$
$$k = 2,\ \mu_n^* = -4,\ n_{\text{immune}} = 10^{-6}.$$

These parameters are the same as the ones reported in [4], except for k and n_{immune}[2]. We pick $k = 2$ to normalize the amount of normal cells at equilibrium during growth phase at 1. For this numerical simulation, an arbitrarily small value of n_{immune} is chosen. We assume that the initial conditions are $n(0) = 1$, $p(0) = 5 \times 10^{-6}$ in growth phase.

[2]Parameters can be scaled up uniformly to simulate a large environment, i.e. host

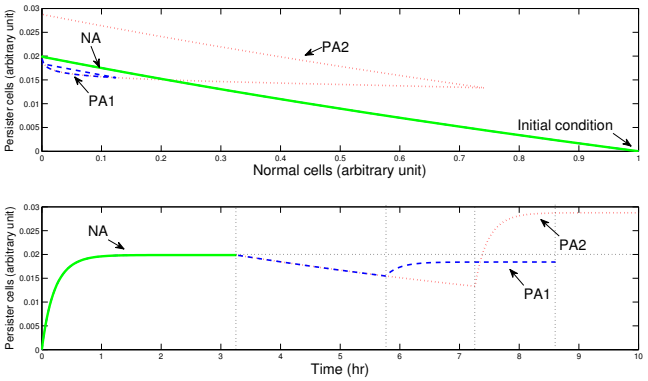

Fig. 4. Top: The phase-plane plots of the outcomes of the three attack patterns. Bottom: The plots of the amount of persister cells vs time for each of the three scenarios.

Consider the following three scenarios:

Naive Attack (NA): Apply antibiotics from time $t = 0$ until long-term dormancy is reached, then stop treatment.

Patterned Attack 1 (PA1): Apply antibiotics until $n = n_{\text{immune}}$ (just before transition to long-term dormancy), then stop the treatment for 6 hours, and then reapply antibiotics until long-term dormancy is reached.

Patterned Attack 2 (PA2): Apply antibiotics until $n = n_{\text{immune}}$ (just before transition to long-term dormancy), then stop the treatment for 12 hours, and then reapply antibiotics until long-term dormancy is reached.

The outcomes of these treatment strategies are shown in Figure 4. We can see that at the transition to long-term dormancy, the amount of persister cells in the three attack scenarios differ. For the naive attack (NA), the value is at around 0.02. For PA1 and PA2, the values are at around 0.018 and 0.028, respectively. From this observation, we can conclude that the naive attack strategy (simply apply antibiotics until the primary infection subsides) is generally not the best option, if we are to suppress the amount of persister cells that enter long-term dormancy. However, by inspecting the performance of PA1 and PA2 we can also conclude that some patterns of attack can be worse than the naive attack strategy. Similar observations were also made in a recent publication [5], which stated that some periodic treatment patterns are not effective. Therefore, an optimal treatment strategy is needed.

We formulate a solution to this optimization problem that is based on the principles of dynamic programming[22],[23]. As is the case with most optimal control problems, our optimal treatment strategy will then be defined as a feedback policy, rather than an open loop timing sequence.

First, we observe that any treatment strategy will end with an attack that leads to the transition to long-term dormancy. This is implicitly expressed in (6), and the problem statement. Our goal is to minimize the number of persisters as the system enters into long term dormancy, where long term dormancy can only be reached during an attack. We calculate the cost function incurred (i.e. the number of persisters at the transition) as a function of the state when this final attack begins, $[n_a,\, p_a]$ as follows. The orbit of the state trajectory in the attack phase can be obtained by solving

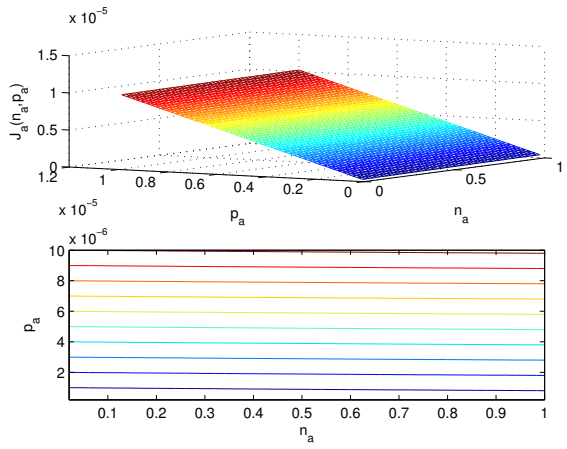

Fig. 5. The plot of the final cost $J_a(n_a, p_a)$ as a function of the state $[n_a, p_a]^T$, at which the final attack is initiated. Top: Surface plot. Bottom: Contour plot

$$\frac{dn}{dp} = \frac{\dot{n}}{\dot{p}} = \frac{(\mu_n^* - a)}{a} - \frac{k}{a}n \tag{7}$$

using $[n_a,\, p_a]$ as a boundary condition. From here, we obtain

$$p = p_a + \frac{a}{k} ln \left(\frac{\mu_n^* - a - kn_a}{\mu_n^* - a - kn(t)} \right) \tag{8}$$

The incurred cost is then obtained by substituting $n = n_{\text{immune}}$ (the start at transition to long-term dormancy).

$$J_a(n_a, p_a) = p_a + \frac{a}{k} ln \left(\frac{\mu_n^* - a - kn_a}{\mu_n^* - a - kn_{\text{immune}}} \right) \tag{9}$$

For the numerical example above, the plot of this cost function is shown in Figure 5. By definition of the cost, the contour lines of $J_a(\cdot)$ are the end of the orbits of the system trajectories in attack mode.

Prior to the start of the final attack, the system must be in the growth phase (the hybrid system contains only two modes). By the principle of dynamic programming, the transition from the final growth phase to the final attack phase (i.e. the start of the final attack) needs to be made when $J_a(\cdot)$ is minimized. Under this assumption, we can compute the final cost as a function of the state when the last growth phase is started, $[n_g,\, p_g]^T$, as follows

$$J_g(n_g, p_g) = \min_{t \geq 0} J_a(n(t), p(t)) \tag{10}$$

where $(n(t), p(t))$ is the state trajectory in growth phase with $[n_g,\, p_g]^T$ as the initial conditions. Furthermore, in the optimal strategy, the final attack should start when the minimum of the right hand side of (10) is attained. The function $J_g(n_g, p_g)$ can be computed and shown in Figure 6.

From the discussion above, we observe that as long as we can traverse down the contour of $J_a(\cdot)$, the number of persisters at the transition to long-term dormancy can be reduced. In order to go to a lower contour (lower orbit), the transitions between growth phase and attack phase are utilized. Specifically, the growth phase initially leads the

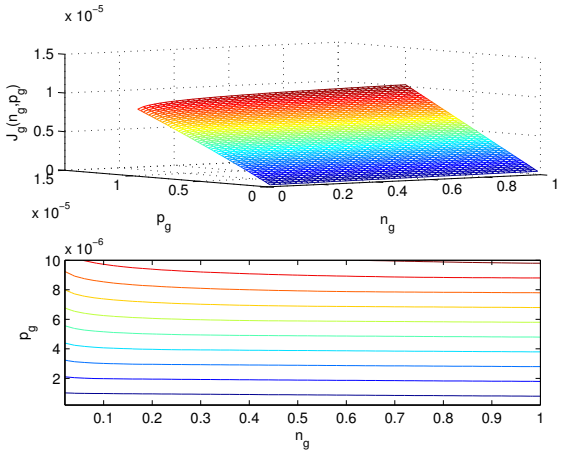

Fig. 6. The plot of the final cost $J_g(n_g, p_g)$ as a function of the state $[n_g, p_g]^T$, at which the final growth phase is started. Top: Surface plot. Bottom: Contour plot

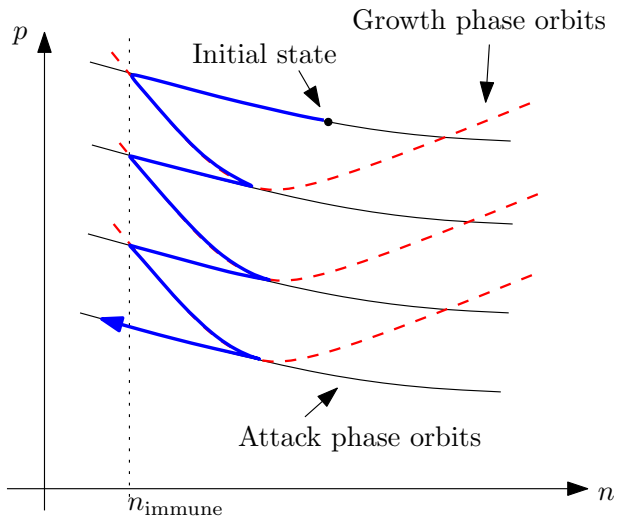

Fig. 7. The sketch of the treatment strategy to reduce the final cost. By switching back and forth between attack and growth phases, we traverse down the contour of $J_a(\cdot)$. Solid lines represent the orbits of the attack phase dynamics, while the dashed line represents the orbits of the growth phase dynamics

system toward lower orbits. A simple sketch of this is shown in Figure 7

From (9), the gradient of $J_a(\cdot)$ is given by

$$\nabla J_a(n,p) = \begin{bmatrix} \frac{-a}{\mu_n^* - a - kn} \\ 1 \end{bmatrix} \tag{11}$$

We can characterize the set of all states at which the growth phase orbit goes down the contour of $J_a(\cdot)$. The set is characterized by

$$\begin{bmatrix} \frac{-a}{\mu_n^* - a - kn} & 1 \end{bmatrix} \begin{bmatrix} (\mu_n - a)n - kn^2 + bp \\ an - bp \end{bmatrix} > 0$$

or

$$p > \frac{-a\,(\mu_n - \mu_n^*)\,n}{b\,(\mu_n^* - kn)} \tag{12}$$

Let us define $Q := \{(n,p) \mid (12) \text{ holds}\}$. For the numerical example above, this set is shown in Figure 8.

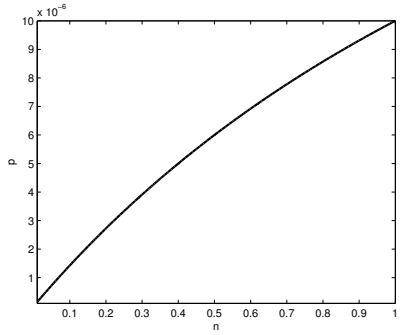

Fig. 8. Q is the area above the curve

Consequently, we can find a lower bound for the number of persisters at the transition to long-term dormancy. This value is given by the intersections of the boundary of Q and $n = n_{\text{immune}}$, or

$$p_{\min} = \frac{-a\,(\mu_n - \mu_n^*)\,n_{\text{immune}}}{b\,(\mu_n^* - kn_{\text{immune}})} \tag{13}$$

For the above numerical example, $p_{\min} = 1.5 \times 10^{-11}$. The optimal treatment strategy is thus not unique, however such strategies can be characterized as in Algorithm 1:

Algorithm 1 Optimal Treatment Strategy

During the attack phase: the attack must stop before $n = n_{\text{immune}}$, but only when $(n,p) \in Q$.

During the growth phase: the attack must begin when $p = \frac{-a(\mu_n - \mu_n^*)n}{b(\mu_n^* - \mu_n)}$

This algorithm does not explicitly define the $t_{1,a}, t_{1,2}, \ldots$ from (5) explicitly, rather the timing of antibiotic dosing is implicit upon the states. A simulation of such strategy is shown in Figure 9. In this figure, we can see the first four waves of attack. Observe that after four waves of antibiotic attack. Observe that after four waves of antibiotic attack, the number of persisters is about 20% of the level that would be attained by a naive attack. We also would like to note that although the attack pattern shown in the bottom panel of Figure 9 seems periodic, it is actually not. Upon continuation of the pattern, which is not shown here, we observe that the intervals between attack waves becomes shorter.

IV. Discussion

We developed a hybrid system model for the dynamics of the populations of normal and persister bacterial cells in favorable growth condition and under antibiotic attacks. Our model is able to capture long-term persistence, a health problem that is common for some infections, such as tuberculosis. Using this model, we posed the problem of designing an optimal infection treatment strategy so as to minimize the number of persister cells that go into long-term dormancy. We subsequently characterized the optimal treatment strategy, which turns out to be non-unique. We also computed the theoretical lower bound on the number of persister cells that transition into long-term dormancy under the optimal treatment scheduling. The non-uniqueness

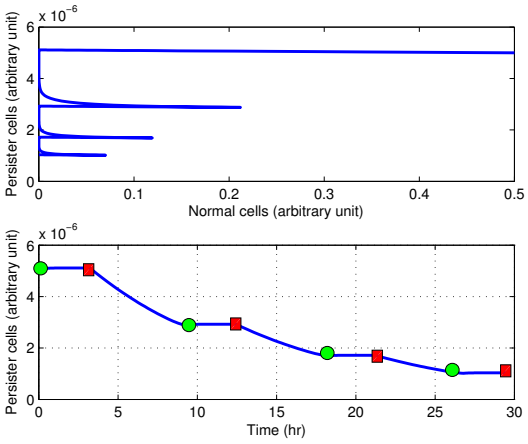

Fig. 9. Top: A simulation of an optimal treatment strategy. Top: Phase-plane plot of the state trajectory. Bottom: A plot of the number of persisters vs time. The circles mark the start of an antibiotic attack, while the squares mark the end of the attack

is due to the approach of an infinite time horizon problem, and the lack of a continual or running cost in the problem statement. If a suboptimal decision is made (application of antibiotics too early or too late) this will only temporarily put the system on a suboptimal path. When looking at a predetermined number of dosing regiments (in Figure 9 there are 4 regiments, or waves of attacks) the optimal solution is unique. The infinite time horizon problem is considered instead of a finite time due to the differences in time scales of active cell growth and dormancy transition rates. Currently, we are pursuing the finite time problem, specifically with the addition of antibiotic resistance to the model. This will be pursued and put forth in future publications.

For future work we also plan to study the problem of optimizing not only the scheduling but also the dosing of the antibiotics. It will be sensible to consider the total administered dose as an additional design parameter (as a cost function or constraint in the optimization). Although not explicit, antibiotic resistance is taken into consideration when developing this control law. Optimizing the beginning of an attack phase but not the end ensures long pulses of antibiotic attack. Prolonged exposure to antibiotics, especially in intermediate concentrations, can lead to the emergence of resistant strains of bacteria. In future work we look to add the resistance state to the model.

Finally, we would like to note that so far our proposed control law is expressed in terms of the population sizes of the normal and persistent cells. This kind of state feedback information might not be available in practice. Therefore, to improve the applicability of control theoretical results in this field, we may need to further develop some notion of state observer, such as the host's physiological state (e.g. body temperature), or the amount of serum or hormone or other signalling molecules in host's bloodstream.

REFERENCES

[1] J. W. Bigger, "Treatment of staphylococcal infections with penicillin," *Lancet*, pp. 497–500, 1944.

[2] H. S. Moyed and K. P. Bertrand, "*hipA*, a newly recognized gene of *Escherichia coli* K-12 that affects frequency of persistence after inhibition of murein synthesis," *J. Bacteriology*, vol. 155, pp. 768–775, 1983.

[3] G. Stewart, B. Robertson, and D. Young, "Tuberculosis: A problem with persistence," *Nature Reviews Microbiology*, vol. 1, pp. 97–105, 2003.

[4] N. Q. Balaban, J. Merrin, R. Chait, L. Kowalik, and S. Leibler, "Bacterial persistence as a phenotypic switch," *Science*, vol. 305, no. 5960, pp. 1622 – 1625, 2004.

[5] O. Gefen and N. Balaban, "The importance of being persistent: heterogeneity of bacterial populations under antibiotic stress," *FEMS Microbiology Review*, no. 33, pp. 704–717, 2009.

[6] N. Dhar and J. D. McKinney, "Microbial phenotypic heterogeneity and antibiotic tolerance," *Current Opinion in Microbiology*, vol. 10, no. 1, pp. 30 – 38, 2007.

[7] R. Jayaraman, "Bacterial persistence: some new insights into an old phenomenon," *J. Biosci*, no. 33, pp. 795–805, 2008.

[8] E. Kussell, R. Kishony, N. Balaban, and S. Leibler, "Bacterial persistence: A model of survival in changing environments," *Genetics*, no. 169, pp. 1807–1814, 2005.

[9] C. G. Cassandras and S. Lafortune, *Introduction to Discrete Event Systems*. Kluwer, 1999.

[10] M. Acar, J. Mettetall, and A. Oudenaarden, "Stochastic switching as a survival strategy in fluctuating environments," *Nature Genetics*, vol. 40, no. 4, pp. 471–475, 2008.

[11] A. Gardner, S. West, and A. Griffin, "Is bacterial persistence a social trait?" *PLoS One*, no. 8, 2007.

[12] R. F. Stengel, R. Ghigliazza, N. Kulkarni, and O. Laplace, "Optimal control of innate immune response," *Optimal Control Applications and Methods*, vol. 23, pp. 91–104, 2002.

[13] D. E. Kirschner, S. Lenhart, and S. Serbin, "Optimal control of the chemotherapy of HIV," *Journal of Mathematical Biology*, vol. 35, pp. 77–96, 1997.

[14] J. A. M. F. de Souza, M. A. L. Caetano, and T. Yoneyama, "Optimal control theory applied to the anti-viral treatment of AIDS," in *Proc. IEEE Conf. Decision and Control*, Sydney, Australia, 2000.

[15] R. Zurakowski and A. R. Teel, "A model predictive control based scheduling method for HIV therapy," *Journal of Theoretical Biology*, vol. 238, pp. 368–382, 2006.

[16] R. Luo, M. Piovoso, and R. Zurakowski, "A generalized multi-strain model of HIV evolution with implications for drug-resistance management," in *Proc. American Control Conference*, St. Louis, Missouri, USA, 2009, pp. 2295–2300.

[17] E. Jung, S. Lenhart, and Z. Feng, "Optimal control of treatments in a two-strain tuberculosis model," *Discrete and Continuous Dynamical Systems - Series B*, vol. 2, pp. 473–482, 2002.

[18] K. Lewis, "Persister cells, dormancy and infectious disease," *Nature Reviews Microbiology*, vol. 5, pp. 48–56, 2007.

[19] E. Renshaw, *Modelling biological populations in space and time*. Cambridge University Press, 1991.

[20] WHO, "Tuberculosis," http://www.who.int/topics/tuberculosis/en/ : Last accessed: 3/27/10.

[21] R. Alur, C. Courcoubetis, N. Halbwachs, T. A. Henzinger, P. H. Ho, X. Nicollin, A. Olivero, J. Sifakis, and S. Yovine, "The algorithmic analysis of hybrid systems," *Theoretical Computer Science*, vol. 138, pp. 3–34, 1995.

[22] D. E. Kirk, *Optimal Control Theory*. Dover Publications Inc., 1970.

[23] D.P.Bertsekas, *Dynamic Programming and Optimal Control*. Athena Scientific, 2007.