Synthesizing Combination Therapies for Evolutionary Dynamics of Disease for Nonlinear Pharmacodynamics

Vanessa Jonsson, Nikolai Matni and Richard M. Murray

Abstract—Our previous results proposed an iterative scalable algorithm for the systematic design of sparse, small gain feedback strategies that stabilize the evolutionary dynamics of a generic disease model with linear pharmacodynamics. In this manuscript, we use piecewise linear approximations to model nonlinear drug effects. We leverage results from optimal controller synthesis for positive systems to formulate the feedback synthesis problem as an optimization problem that sequentially explores piecewise linear subsystems corresponding to higher and higher treatment dosages.

I. INTRODUCTION AND MOTIVATION

A challenge inherent to the treatment of certain infectious and non-infectious diseases, such as HIV or cancer, is the risk that the pathogen or tumor will evolve away and become resistant to treatment methods that comprise the standard of care [1], [2], [3], [4]. Especially vulnerable to this phenomenon are treatment methods that involve exposing the disease population (such as viruses or cancer cells) to therapies targeting specific molecules involved in disease progression for an extended period of time. While these targeted therapies have the benefit of allowing physicians to tailor treatments to a patient's tumor cell population, they nonetheless establish an environment in which the occurrence of mildly drug resistant pathogens or tumor cells can develop an evolutionary advantage over those for which the therapy is targeted [5], [6], [7], [8], leading to so called 'treatment-escape'.

Pharmacodynamic models are used to quantitatively describe nonlinear drug dose-responses and model drug-receptor relationships. These models are often described using combinations of nonlinear Hill functions, and allow for the modeling of drug saturation effects [9], [10], [11]. These pharmacodynamic nonlinearities are further increased with the fact that drugs administered in combination can have independent additive effects or can otherwise exhibit synergistic or antagonistic dose dependent behavior, further complicating the design of suitable combination therapies [12].

The challenge of designing treatment protocols that prevent escape is one that has been addressed by control theoretic methods. For cancer therapy, results in this spirit apply methods from optimal and receding horizon control [13], [14], as well as gain scheduling [15], to synthesize treatment protocols that are robust to parameter uncertainty, an inherent issue in all biological systems. The authors in [16], present

a static multi-objective optimization formulation to solve the combination therapy problem for different initial tumor populations, when the drugs under consideration have additive, linear effects on cell viability. Proposed combination treatments were confirmed experimentally for different tumor initial conditions in a murine lymphoma model [17]. In the context of HIV and antiretroviral therapy, the authors in [18] propose a discrete time formulation that allows for the design of switching therapy strategies to delay the emergence of highly resistant mutant viruses. Recent results in [19] and [20], consider a simplified bilinear model and the optimal control problem is shown to be convex over a finite horizon for a predefined set of initial states.

There have been several attempts to deal with nonlinear HIV infection dynamics using model predictive control (MPC) to design optimal antiretroviral drug dosing strategies [21], [22]. An issue with the use of MPC is that an essential feature of the system, that of nonlinearities associated with drug binding, is linearized away. In particular, such an approach can lead to a model that underestimates the efficacy of a drug at lower concentrations, and over estimates its efficacy at high dosages *unless* a sufficiently small update time is taken. This latter restriction may then lead to strategies that are no longer realistic in a biomedical application, where it may be difficult to update the administered therapies at the frequency dictated by the MPC controller. In this paper, we take an alternative approach and attempt to design constant drug therapies by taking input non-linearities into account more explicitly.

In [23], we introduced a scalable, iterative algorithm for the principled design of targeted combination therapy concentrations that explicitly accounts for the evolutionary dynamics of a generic disease model. This algorithm was effective in generating robustly stabilizing controllers, while simultaneously allowing the designer to explore explicit trade offs between closed loop performance, sparsity in controller structure and gain minimization. Leveraging recent results from positive systems [24], we formulated this algorithm as a second order cone program (SOCP), which made the controller synthesis scalable. However it could not take into account a) the pharmacodynamics of the input, potentially suffering from over or underestimating gains, and b) the effects of synergistic or antagonistic drug combinations that can be modeled with additional nonlinear pharmacodynamic terms [12].

Here we propose a new algorithm that solves the combination therapy problem subject to the same design constraints (sparsity of the drug combination, maximum dosage

V. Jonsson, N. Matni and R. M. Murray are with the Department of Control and Dynamical Systems, California Institute of Technology, Pasadena, CA. vjonsson@caltech.edu.

and robustness constraints) formulated as an SOCP while addressing the non linear dynamics of individual drugs and of their combinations. In particular, through the piecewise linearization of individual and combination drug pharmacodynamics, in conjunction with a branch and bound like algorithm for the effective search through these linear pharmacodynamic modes, we reduce the combination therapy problem to applying the SOCP formulation from [23] to a set of pharmacodynamic modes.

The main contribution of this paper is an algorithm for the systematic design of *sparse, small gain* feedback strategies to stabilize the evolutionary dynamics of a generic disease model and general nonlinear pharmacodynamics models, which support synthesis of feedback strategies in light of highly nonlinear drug dynamics.

The article is structured as follows: In Section II, we recall the extended quasispecies evolutionary dynamics model that encodes replication, mutation and neutralization and summarize relevant results in controller design of positive systems and introduce the pharmacodynamics of a non competitive additive drug binding model. In Section III, we present our $\mathcal{L}_1$ combination therapy synthesis algorithm as well as an algorithm designed to increase the scalability of the formulation. Section IV illustrates our algorithm in the context of an additive drug interaction example, that of the HIV antibody therapy design problem previously studied in an experimental setting [25], where the drugs act independently and additively. Section V ends with concluding remarks and directions for future work.

II. PRELIMINARIES

A. Notation

$\mathbb{R}_+$ denotes the set of nonnegative real numbers. The inequality $X > 0, (X \geq 0)$ means that all elements of the matrix or vector X are positive (nonnegative). $X \succ 0$ means that X is a symmetric and positive definite matrix. The matrix $A \in \mathbb{R}^{n \times n}$ is said to be Hurwitz if all eigenvalues have negative real part. Finally, the matrix is said to be Metzler if all off-diagonal elements are nonnegative. Define $\mathbf{1}_n$ to be the vector of all ones of dimension n . The induced matrix norm for a matrix $M \in \mathbb{R}^{r \times m}$ is $\|M\|_{p-\text{ind}} = \sup_{w \in \mathbb{R}^m} \frac{\|Mw\|_p}{\|w\|_p}$ where $\|w\|_p = (|w_1|^p + \dots + |w_m|^p)^{1/p}$. Let $G(s) = C(sI - A)^{-1}B + D$ be a $r \times m$ matrix transfer function. The induced norms of the corresponding impulse response $g(t) = Ce^{At}B + D\delta(t)$ are $\|g\|_{p-\text{ind}} = \sup_w \frac{\|g*w\|_p}{\|w\|_p}$ for $w \in \mathcal{L}_m^p[0, \infty)$, given that $g*w \in \mathcal{L}_p^r[0, \infty)$ is the convolution of g and w . Finally we refer to the ∞ -induced robust controller as the $\mathcal{L}_1$ controller as is customary in the robust control literature.

B. Evolutionary dynamics model

The quasispecies model [26] was originally developed to describe the dynamics of populations of self replicating macromolecules undergoing mutation and selection. We choose this model for its relative simplicity and its ability to capture the salient features of the evolutionary dynamics of

a simplified generic disease model. In [27] we incorporated the effects of potential therapies into the basic quasi species model, by defining a drug binding reaction, $\ell + \rho \xrightarrow{K_A} \ell \cdot \rho$ — drug ℓ binds to self replicating macromolecule ρ with association rate K_A , giving a neutralized complex $\ell \cdot \rho$. The extended quasispecies model for n mutants and m drugs, is written as:

$$\dot{x}_i = (r_i q_{ii} - d_i)x_i + \sum_{k \neq i}^n r_i q_{ik} x_k - \Psi_i(\ell)x_i \quad (1)$$

where $x_i \in \mathbb{R}_+$ is the concentration of mutant i , $\ell = (\ell_k) \in \mathbb{R}_+^m$ is the drug concentration (assumed to remain at constant concentrations throughout), r_i and d_i are the replication and degradation rates, respectively, of mutant i , and q_{ik} is the probability that mutant k mutates to mutant i (note that q_{ii} is the probability of no mutation occurring). The rates r_i can be viewed as the replication fitnesses of mutant i without the effect of the drug. When $\ell_k = 0, \forall k \in \{1, \dots, m\}$, the quasispecies dynamics are unstable. Finally, the function $\Psi_i(\ell)$ represents the pharmacodynamics of individual drugs ℓ_k and their combinations with respect to the i -th mutant species, namely the sum of nonlinear drug effect functions as represented by Hill equations. In the following section, we present a brief justification for the Hill equation models of these binding reactions.

C. Pharmacodynamic models

When combined, drugs can have additive, synergistic or antagonistic effects that need to be explicitly accounted for when designing combination therapies. When the presence of one drug modulates the effect of another, the combined drug effects are no longer additive, and in particular, the pharmacodynamics of the drug interaction now incorporate additional terms that represent this interaction.

1) *The Hill equation*: The Hill equation has been used in pharmacology to model nonlinear drug dose-responses, for example the effects of drugs on cell viability or virus neutralization. More generally it serves to quantify drug-receptor relationships, the fraction of bound receptors ρ (e.g. cell receptors, virus proteins) as a function of ligand ℓ_k (e.g. drug, antibody) concentrations for the binding reaction

$$\psi_\rho(\ell_k) = \frac{[\ell_k]^{n_k}}{[\ell_k]^{n_k} + K_k^{n_k}} \quad (3)$$

where $\psi_\rho(\ell_k) \in [0, 1]$ is the fraction of bound receptors, $\ell_k \in \mathbb{R}_+$ is the concentration of ligands, $K_k = \frac{1}{K_A} \in \mathbb{R}_+$ is the dissociation constant associated with the binding reaction (4), $n_k \in \mathbb{R}_+$ is the Hill coefficient that represents the degree of cooperativity, i.e. the degree to which binding of a ligand molecule modulates the probability of another ligand molecule binding.

We notice that the Hill function is a biological analog to actuator saturation, in that there is a law of diminishing

returns in terms of the effect of ever increasing drug concentrations on the system. In fact, the Hill functions $\psi_\rho(\ell_k)$ that appear in HIV applications, for example, look approximately linear for small to moderate ℓ_k , and nearly constant for large ℓ_k – it is this observation that motivates our piecewise linear approximation approach.

2) *Non-competitive additive drug binding model*: Under this model, we consider the pharmacodynamics of non-interacting drugs that bind to different receptors, in which the effect of each drug on the system is additive. Consider the system of m drug binding reactions to different receptors on a particular cell or virus x :

$$\ell_1 + \rho_1 \xrightarrow{K_1} \ell_1 \cdot \rho_1, \dots, \ell_m + \rho_m \xrightarrow{K_m} \ell_m \cdot \rho_m \quad (4)$$

where ℓ_k is the drug k and ρ_k is a receptor. If these receptors comprise different drug binding targets on a cell or virus x , then we can describe the total effect of these independently acting drugs $\ell_1, \dots, \ell_m$ on x as

$$\Psi_x(\ell) = \sum_{k=1}^m \psi_{\rho_k}(\ell_k). \quad (5)$$

Remark 1: Although the remainder of the paper focusses on this drug interaction model, we note that our approach applies nearly verbatim to a synergistic/antagonistic drug binding model – it suffices to use a suitable expression for $\Psi_x(\ell)$ to take these interactions into account. We do note however that this may lead to a more involved piecewise linear approximation procedure.

III. A SUB-OPTIMAL $\mathcal{L}_1$ COMBINATION THERAPY CONTROLLER FOR NON LINEAR PHARMACODYNAMICS

In this section, we formulate the task of designing suitable combination therapies as an optimal control problem. The inherent non-linearities of the system make this a challenging task – in [23], [27], we worked with a simplified problem in which we assumed the $\Psi_x(\ell)$ were linear functions – in this paper we relax that assumption, and show that at the expense of some additional modeling complexity, we are able to reduce the problem to that considered in [23]. Our main result is based around the use of piecewise linear approximations, and relating the robustness levels of the approximate system to that of the true underlying system.

A. Problem formulation

The following state space representation of equation (1) emphasizes the inherent feedback structure that arises from drug binding reactions:

$$\begin{aligned} \dot{x} &= (A - \Psi(\ell))x + w \\ z &= Cx \end{aligned} \quad (6)$$

with (i) $A \in \mathbb{R}^{n \times n}$, with $A_{ij} = r_i q_{ij} \geq 0 \forall i \neq j$ and $A_{ii} = r_i q_{ii} - d_i$, that encodes the mutation and replication dynamics; (ii) $\Psi(\ell) \in \mathbb{R}^{n \times n}$, a diagonal matrix, with diagonal elements $\Psi_{x_i}(\ell)$, that maps the effects of the therapies ℓ to the population dynamics; (iii) $\ell = (\ell_k) \in \mathbb{R}^m$, a vector of the concentrations of neutralizing macromolecules; (iv)

$C = \mathbf{1}_n^T \in \mathbb{R}^{1 \times n}$; and (v) $w \in \mathbb{R}_+^n$ an arbitrary positive disturbance.

We set the regulated output $z = \mathbf{1}_n x$ to be the total virus population, so as to ensure that the resulting treatment plan is one that robustly drives the total mutant population to zero.

Remark 2: A is a Metzler matrix with off-diagonal entries that are several orders of magnitude smaller than the diagonal entries. This is due to the biological fact that mutation rates range from $10^{-5} - 10^{-8}$ mutations per base pair per replication cycle for reverse transcriptase [28] to DNA replication [29].

Letting G denote the closed loop system (6), the control task then becomes to reverse engineer neutralizing macromolecule concentrations ℓ by finding a ‘controller’ $\Psi(\ell)$ that leads to a stable G satisfying $\|G\|_{\infty-\text{ind}} < \gamma$, for some robustness level $\gamma > 0$.

B. Piecewise linear mode approximations and mode reduction

In the following, we assume that each $(\Psi(\ell))_{ii} = \Psi_{x_i}(\ell)$ has the form given by (5). To take into account non-linear pharmacodynamics, we propose a piecewise linear approximation algorithm and a mode reduction algorithm for problems where there is a large number of non interacting or synergistic drug combinations.

We assume that the pharmacodynamics for every individual drug and combination are defined over the same drug concentration domain $\mathcal{D} \in \mathbb{R}$. Let $\mathcal{P} = \{p_1, \dots, p_k\}$ be a partitioning of this domain into k intervals.

Definition 1: A pharmacodynamic mode $\omega = (\omega_1, \dots, \omega_m)$ is an m -tuple in $\mathcal{P}^m$. The total number of pharmacodynamic modes is $|\mathcal{P}|^m$ where m is the number of drugs under consideration. For $\ell \in \mathbb{R}^m$, we define $\ell \in \omega \iff \ell_i \in \omega_i, \forall i = \{1, 2, \dots, m\}$.

We let $\psi_{i\omega} : \mathbb{R}^m \rightarrow \mathbb{R}$ be the affine approximation of the pharmacodynamics for each mutant i for $\ell \in \omega$, i.e. the sum of the individual and combination drug effects on mutant i while operating within mode ω . We can then construct a linear approximation to $\Psi(\ell)$ via an appropriately defined block diagonal matrix Ψ_ω , constructed from the $\psi_{i\omega}$ (c.f. [23]), and write $\Psi(\ell) = \Psi_\omega L_\omega$, where $L_\omega = (I \otimes \ell_\omega) \in \mathbb{R}_+^{mn \times n}$ is the block diagonal matrix, with identical block diagonal elements given by the drug concentrations $\ell \in \omega$. The resulting dynamics, for a fixed concentration $\ell \in \omega$, are then described by the transfer function

$$G_{\ell\omega}(s) = C(sI - (A - \Psi_\omega L_\omega)^{-1}B + D). \quad (7)$$

We consider the problem of finding a suitably sparse therapy combination that achieves a certain closed loop performance level γ . As such, our initial goal becomes to reduce the search space to the set of sparse modes that achieve the desired level of robustness, where a sparse mode ω is one that allows at least one drug concentration to be zero, in other words, $0 \in \omega_i$ for at least one $\omega_i \in \omega$.

In order to do so, we require two lemmas. The first provides sufficient conditions on the linear approximation terms $\psi_{i\omega}$ that guarantee that the robustness of the piecewise

linear approximation is an upper bound on that of the true system. The second is the simple observation that for non-interacting or synergistic additive drug interactions, the robustness of the closed loop dynamics increases as drug concentrations are uniformly increased (this statement will be made precise). This result allows us to subsequently develop a branch and bound like method that significantly reduces the search space of the algorithm.

We begin with a result on the input-output performance of a positive system, taken from [24].

Lemma 1: Let $G(s) = C(sI - A)^{-1}B + D$ be a positive system. Then $\|G\|_{\infty-\text{ind}} \leq \gamma$ if and only if there exists $x \geq 0$ such that

$$\begin{aligned} Ax + B\mathbf{1}_n &< 0 \\ Cx &\leq \gamma \end{aligned} \quad (8)$$

Lemma 2: Let $\Psi(\ell)$ be the non-linear pharmacodynamics function, $\psi(\ell)$ the vector of its diagonal elements, and denote by $\Psi_\omega(\ell)$ its piecewise linear approximation within mode ω , and by $\psi_\omega(\ell)$ the vector of its diagonal elements. If for every mode ω we have $\psi_\omega(\ell) \leq \psi(\ell) \forall \ell \in \omega$ then the $\mathcal{L}_1$ norm γ of the piecewise linear approximation (7) is an upper bound on that of the true system (6).

Proof: Note that for a fixed ℓ , both the full and piecewise linear approximation systems are linear in x . Therefore, by Lemma 1, the $\mathcal{L}_1$ norm of (7) is upper bounded by γ if and only if there exists an $x > 0$ such that

$$\begin{aligned} (A - \Psi_\omega(\ell))x + \mathbf{1}_n &< 0 \\ \mathbf{1}_n^T x &\leq \gamma \end{aligned} \quad (9)$$

Thus it suffices to show that this same x yields $(A - \Psi(\ell))x + \mathbf{1}_n \leq 0$ and the desired conclusion follows immediately. To that end, rewrite the first inequality of (9) as

$$(A - \Psi(\ell))x + \mathbf{1}_n < (\Psi_\omega(\ell) - \Psi(\ell))x, \quad (10)$$

and notice that the right hand side is less than or equal to 0 by the assumptions of the lemma. ■

This lemma essentially states that if our piecewise linear approximation is conservative, then the norm of the approximate system serves as a certificate for the norm of the true system.

Next we formalize the observation that increasing the concentrations of non or synergistically interacting therapies present in the system will improve robustness.

Lemma 3: Let ℓ^1 and ℓ^2 be therapy combinations such that $\ell^1 \geq \ell^2$. Then if the piecewise linear approximation $\Psi_\omega(\ell)$ is non-decreasing, $\|G_1\|_{\infty-\text{ind}} \leq \|G_2\|_{\infty-\text{ind}}$, where

$$G_i = C(sI - (A - \Psi_\omega(\ell^i))^{-1}I.$$

Proof: By assumption, the piecewise linear function $\Psi_\omega(\ell)$ is non-decreasing. Thus if $\ell^1 \geq \ell^2$, then $\Psi_\omega(\ell^1) \geq \Psi_\omega(\ell^2)$. Let $\gamma_i = \|G_i\|_{\infty-\text{ind}}$. By Lemma 1, γ_i is the solution to the following optimization:

$$\begin{aligned} &\underset{\gamma, x \geq 0}{\text{minimize}} \quad \gamma \\ &\text{subject to} \quad (A - \Psi_\omega(\ell))x + \mathbf{1} < 0 \\ &\quad \quad \quad Cx \leq \gamma \end{aligned} \quad (11)$$

Let γ_2 and x be the optimal solutions of the above program for $i = 2$. Then we have that

$$(A - \Psi_\omega(\ell^1))x + \mathbf{1} \leq (A - \Psi_\omega(\ell^2))x + \mathbf{1} \leq 0. \quad (12)$$

Hence (γ_2, x) is a feasible solution for optimization (11) with $i = 1$, implying that $\gamma_1 \leq \gamma_2$. ■

Thus we see that by constraining the piecewise linear approximations to be under approximations of the effects of the drugs, we are able to bound the performance of the true system.

Remark 3: We note that these results typically will not hold for antagonistic drug interactions – however, many settings in which such models are required (such as cancer therapy design) do not have a large set of therapies or mutants, thus mitigating the computational cost of the sequential search across modes.

We exploit this result to reduce the modes that need to be searched over – in particular, we use the partial order implied by the previous two results to upper and lower bound the uniform concentration treatments required to achieve a prescribed performance level γ . We also exploit the fact that we are searching for *sparse* treatment strategies to further eliminate modes.

This approach is formalized in the following algorithm. Let $\ell_{\omega_{max}} \in \mathbb{R}^m$, be the maximum possible drug concentrations.

Algorithm 1 Sparse mode reduction algorithm

```

Set  $\ell_{\omega_i} \leftarrow \ell_{\omega_{max}}, \gamma > 0$ .
while  $\neg(\ell_{\omega_i} == \mathbf{0}_n \wedge \epsilon)$  :
    if  $\|G_{\ell_{\omega_i}}\| < \gamma$ ,
         $\mathcal{S} = \mathcal{S} \cup \omega_i$ 
    else
         $\mathcal{U} = \mathcal{U} \cup \omega_i$ 
    Set  $\ell_{\omega_{i+1}} = \frac{\ell_{\omega_i}}{2}$ .
    Set  $\epsilon = (\omega_i \in \mathcal{S} \parallel \omega_i \in \mathcal{U})$ 

```

The sparse mode reduction algorithm will generate a set of modes that are guaranteed to be stable and achieve a desired robustness level γ , and “sparse”, i.e. allowing modes such that at least one drug concentration is allowed to be zero, significantly reducing the number of the modes over which to apply the combination therapy algorithm. For the HIV example in Section IV, to synthesize controllers with robustness level $\gamma = 14$, we start with $|\mathcal{P}|^m = 10000$ possible modes to search over, and reduce this number to 397 sparse modes with robustness level $\gamma = 14$, via Algorithm 1.

Algorithm 2 Scalable Combination Therapy

- 1) Set $\ell_0 = \ell_{\omega_{max}}$.
 - 2) Check if $P1_{\ell_0, \omega}(x, s)$ is feasible. If feasible
 - Set $(x', s') = P1_{\ell_0, \omega}(x, s)$.
 - Set $(\ell') = P2_{(x', 0, 0, \omega)}(\ell)$.
 else, move to next mode and return to Step 1.
 - 3) Find $(\lambda'_1, \lambda'_2, \ell_\omega)$ for mode ω :
 - $\forall (\lambda_1, \lambda_2) \in \Lambda_1 \times \Lambda_2, \Lambda_1, \Lambda_2 \in \mathbb{R}_+^k$,
 - Set $s = 1$.
 - while $\neg(s' == s)$:
 - Set $s = s'$.
 - Set $(x', s') = P1_{\ell', \omega}(x, s)$.
 - Set $(\ell') = P2_{(x', \lambda_1, \lambda_2, \omega)}(\ell)$.
-

C. A suboptimal combination therapy algorithm

As discussed in [23], [27] there are no known convex reformulations for the robust combination therapy problem over an infinite horizon due to the additional structure on L . As such we use the iterative approach developed in [23], as formalized in Algorithm 2, based on the convex programs (13) and (14), to find a stabilizing controller, given a desired robustness level γ . For notation, let $Y' = P_{Z'}(x, s)$ denote an optimization problem P in which we optimize over x and s leaving Z' fixed and with solution Y' . These optimization programs, taken from [24], are a synthesis variant of the conditions stated in Lemma 1.

Program 1. $P1_{\ell, \omega}(x, s)$:

$$\begin{aligned}
 & \underset{x \in \mathbb{R}_+^n, s}{\text{minimize}} \\
 & \text{subject to} \\
 & A_\omega x + \Psi_\omega Lx + \mathbf{1} \leq s \\
 & L = I \otimes \ell, \ell \in \omega \\
 & \mathbf{1}_n^T x \leq \gamma \\
 & s < 0, \quad x \geq 0
 \end{aligned} \tag{13}$$

Program 2. $P2_{(x, \lambda_1, \lambda_2, \omega)}(\ell)$

$$\begin{aligned}
 & \underset{\ell \in \mathbb{R}_+^k}{\text{minimize}} \quad \lambda_1 \|\ell\|_1 + \lambda_2 \|\ell\|_2 \\
 & \text{subject to} \\
 & A_\omega x + \Psi_\omega Lx + \mathbf{1} < 0 \\
 & L = I \otimes \ell, \ell \in \omega \\
 & \mathbf{1}_n^T x \leq \gamma
 \end{aligned} \tag{14}$$

Remark 4: Note the introduction of a slack variable s into Program 1, to help prevent immediate convergence to a local minimum. Minimizing s has the effect of maximizing the slack in the first constraint, allowing for more freedom in the design of the concentration vector ℓ in Program 2.

IV. APPLICATIONS TO ADDITIVE PHARMACODYNAMIC BINDING MODELS

A. Additive drug effects: HIV and antibody therapy

Our results provide a principled approach to the design of antibody treatments for chronic infection with human

immunodeficiency virus-1 (HIV-1) in light of the nonlinear pharmacodynamics and saturation concerns associated with antibody neutralization of HIV-1. We illustrate this with an example motivated by experimental results of evolutionary dynamics of HIV-1 in the presence of antibody therapy obtained in [25].

A relatively recent discovery is that a minority of HIV-infected individuals can produce broadly neutralizing antibodies (bNAbs), that is, antibodies that inhibit infection by many strains of HIV [30]. These have been shown to inhibit infection by a broad range of viral isolates *in vitro* but also protect non-human primates against infection [30], [31], [32]. Recent experimental results conducted in the Nussenzweig lab at Rockefeller University have demonstrated that the use of single antibody treatments can exert selective pressure on the virus, but escape mutants due to a single point mutation can emerge within a short period of time [25]. Although antibody monotherapy did not prove effective, it was shown that equal, high concentrations of an antibody pentamix effectively control HIV infection and suppress viral load to levels below detection.

The goal of this example is to demonstrate how our proposed algorithm offers a principled way to design combination antibody therapies that control HIV infection and prevent evolution of any set of known resistant mutants given the nonlinear pharmacodynamics of antibody neutralization of HIV-1. In a realistic setting, the ability to do this relies on the knowledge of what resistant viruses may be selected for with single therapies, and knowledge of antibody pharmacodynamics. This algorithm would be most effective in conjunction with single antibody selection experiments and knowledge of antibody Hill function properties.

1) Model parameters: We consider four potential antibodies to use in combination PG16, 45-46G54W, PGT128 and 10-1074 on an evolutionary dynamics model of twenty one mutants.

The Hill function associated with the fitness of a virus with respect to neutralization by an antibody is described in equation (3). We used Hill function parameters that were experimentally derived in [33] for antibodies 45-46G54W, PGT128 and PG16 and approximated the Hill parameters associated with antibody 10-1074. We consider a set of twenty one mutants that evolved from antibody monotherapy experiments in [25]. Their corresponding half maximal inhibitory antibody concentration (IC50) in $\mu g/ml$, is specified in the Supplementary Figure 8 of [25]. We choose replication rates to be $0.5 \text{ (ml} \cdot \text{day)}^{-1}$ for all mutants. We justify this selection by noting that escape mutants grew to be dominant mutants during selection experiments and assume that replication rate variability due to mutations are negligible. The mutation rate for HIV reverse transcriptase is $u = 3 \times 10^{-5}$ mutations/nucleotide base pair/replication cycle [28], and the HIV replication cycle is approximately 2.6 days [34]. We approximate the rate of mutation for a particular amino acid mutation at a particular location to be $\frac{1}{n_a} u (1 - u)^k = 1.443 \times 10^{-6}$ per replication cycle, where $k \approx 3000$ is the size of the genome in residues and

$n_a = 19$ is the number of amino acids that can be mutated to. Our model supports forward point mutations and two point mutations. We do not consider back mutations, as the probability of mutation is negligible. Units of concentration in number of viruses/ml or number of antibodies/ml are used for states, and time is measured in days. The standard volume is 1 ml.

B. L_1 controller synthesis

We performed ten piecewise linear approximations on each of the Hill functions associated to each of the twenty one mutants and each of the four antibodies we considered as possible candidate therapies. This generated 10000 possible pharmacodynamic modes of which we found 397 to be sparse and stable by Algorithm 1. Figure 1 shows an example of a conservative piecewise linear approximation to a Hill function associated with a virus mutant antibody pairs.

We synthesized a family of robustly stabilizing controllers using our algorithms for a range of desired robustness levels and found that two different trimixes were predominant - one more frequent combination comprised of (45-46G54W, PGT128,10-1074) antibodies was present for smaller robustness levels and another combination, (PG16, 45-46G54W, PGT128) appeared in addition to the first one when the desired robustness was allowed to be larger as shown in Figure 1. According to [25], the mean IC50 for virus mutants is greatest for the PG16 mono therapy than any other single antibody.

In [25], a different antibody trimix containing PG16, (3BC176, PG16, 45-46G54W) was suggested and experimentally shown to produce a decline in the initial viral load. However, a majority of mice in the experimental study had a viral rebound to the trimix pre-treatment levels, suggesting that in these cases, the virus had evolved mutations that were resistant to the trimix treatment. Further study showed that the evolved mutants had mutations found in the PG16 and 45-46G54W monotherapy groups. This could be suggestive that the combination of (PG16 and 45-46G54W) with another antibody although stabilizing, may not be robust enough to the type of perturbations witnessed in a biological setting. In Figure 2, we see that although the (PG16,45-46G54W, PGT128) (red) and (45-46G54W, PGT128,10-1074) (green) and combinations were synthesized to have the same robustness guarantee, the combination containing PG16 and 45-46G54W is less robust to perturbations.

These simulations demonstrate that although many stabilizing solutions to the combination therapy problem exist, the best ones are found when design parameters such as a sparsity, limits on the magnitude of gains, and robustness guarantees are simultaneously considered. Experimentally searching for these combinations is infeasible as the number of potential therapies and possible concentrations to consider is experimentally intractable. We propose to guide these experimental activities with our ability to design and synthesize combination therapy controllers. As such, one could generate a family of controllers based on “design specifications” tailored not only the (viral or cellular) composition of the

Fig. 1. (Left) Piecewise linear approximation (blue dashed) of the Hill function (red) for PG16- wild type clade B YU2 HIV-1. (Right) Robustness versus number of sparse controllers synthesized using the $\mathcal{L}_1$ nonlinear combination therapy algorithm for the first lexicographically ordered twenty robustly stable sparse modes. Two different 3-sparse antibody trimixes are synthesized, with the dominant trimix being the one that excludes PG16.

Fig. 2. Sum of virus populations subject to random time invariant perturbations of 0.13% in the dynamics for 30 different simulations for (left) a stabilizing closed loop controller comprised of antibody trimixes (0,0.96,2,1) (blue), (0,1,0.177,2) (green) and (2,0.35,0.96,0) (red) $\mu\text{g/ml}$ of (PG16,4546-G54W, PGT128, 10-1074) combinations synthesized using our algorithm and (right) for the same antibody trimixes subject to perturbations of 0.15% in the dynamics.

disease, but to explore tradeoffs between number of therapies used (sparsity), therapy concentrations (magnitude of the gain) and ability to support pharmacokinetic fluctuations (robustness to perturbations) and subsequently verify these experimentally.

V. CONCLUSION AND FUTURE WORK

We proposed an iterative algorithm and mode reduction scheme for the systematic design of sparse, small gain feedback strategies that stabilize the evolutionary dynamics of a generic disease model, subject to nonlinear pharmacodynamics. In future work, we plan to explore a principled integration of our methods with recent results on the robust $\mathcal{L}_1$ stability of positive systems [35]. In particular, we hope to explicitly account for unmodeled dynamics due to pharmacokinetic perturbations.

VI. ACKNOWLEDGEMENTS

We would like to thank David Baltimore and Pamela Bjorkman for discussions regarding using antibody therapy for HIV treatment. We appreciate the help of Bjorkman laboratory staff scientist Anthony West for information on antibody neutralization parameters.

REFERENCES

- [1] R. Kantor, R. Shafer, S. Follansbee, J. Taylor, D. Shilane, N. Hurley, Leo, D. Phuong, D. Katzenstein, and W. J. Fesselb, "Evolution of resistance to drugs in HIV-1-infected patients failing antiretroviral therapy," *AIDS*, vol. 18, no. 11, pp. 1503–1511, July 2004.
- [2] W. Pao, V. A. Miller, K. A. Politi, G. J. Riely, R. Somwar, M. F. Zakowski, M. G. Kris, and H. Varmus, "Acquired resistance of lung adenocarcinomas to gefitinib or erlotinib is associated with a second mutation in the egfr kinase domain," *PLoS medicine*, vol. 2, no. 3, p. e73, 2005.
- [3] A. Hochhaus, S. Kreil, A. Corbin, P. La Rosee, M. Muller, T. Lahaye, B. Hanfstein, C. Schoch, N. Cross, U. Berger, *et al.*, "Molecular and chromosomal mechanisms of resistance to imatinib (sti571) therapy," *Leukemia*, vol. 16, no. 11, pp. 2190–2196, 2002.
- [4] S. Kobayashi, T. J. Boggon, T. Dayaram, P. A. Janne, O. Kocher, M. Meyerson, B. E. Johnson, M. J. Eck, D. G. Tenen, and B. Halmos, "Egfr mutation and resistance of non-small-cell lung cancer to gefitinib," *New England Journal of Medicine*, vol. 352, no. 8, pp. 786–792, 2005.
- [5] M. D. Thakur, F. Salangsang, A. S. Landman, W. R. Sellers, N. K. Pryer, M. P. Levesque, R. Dummer, M. McMahon, and D. D. Stuart, "Modelling vemurafenib resistance in melanoma reveals a strategy to forestall drug resistance," *Nature*, vol. 494, no. 7436, pp. 251–255, 2013.
- [6] J. A. Engelman and P. A. Janne, "Mechanisms of acquired resistance to epidermal growth factor receptor tyrosine kinase inhibitors in non-small cell lung cancer," *Clinical Cancer Research*, vol. 14, no. 10, pp. 2895–2899, 2008.
- [7] R. C. Doebele, A. B. Pilling, D. L. Aisner, T. G. Kutateladze, A. T. Le, A. J. Weickhardt, K. L. Kondo, D. J. Linderman, L. E. Heasley, W. A. Franklin, *et al.*, "Mechanisms of resistance to crizotinib in patients with alk gene rearranged non-small cell lung cancer," *Clinical Cancer Research*, vol. 18, no. 5, pp. 1472–1482, 2012.
- [8] B. E. Johnson, T. Mazar, C. Hong, M. Barnes, K. Aihara, C. Y. McLean, S. D. Fouse, S. Yamamoto, H. Ueda, K. Tatsuno, *et al.*, "Mutational analysis reveals the origin and therapy-driven evolution of recurrent glioma," *Science*, vol. 343, no. 6167, pp. 189–193, 2014.
- [9] R. R. Regoes, C. Wiuff, R. M. Zappala, K. N. Garner, F. Baquero, and B. R. Levin, "Pharmacodynamic functions: a multiparameter approach to the design of antibiotic treatment regimens," *Antimicrobial agents and chemotherapy*, vol. 48, no. 10, pp. 3670–3676, 2004.
- [10] M. E. S. Sampah, L. Shen, B. L. Jilek, and R. F. Siliciano, "Dose-response curve slope is a missing dimension in the analysis of hiv-1 drug resistance," *Proceedings of the National Academy of Sciences*, vol. 108, no. 18, pp. 7613–7618, 2011.
- [11] R. Phillips, J. Kondev, J. Theriot, N. Orme, and H. Garcia, *Physical biology of the cell*. Garland Science New York, 2009.
- [12] K. B. Wood, K. C. Wood, S. Nishida, and P. Cluzel, "Uncovering scaling laws to infer multidrug response of resistant microbes and cancer cells," *Cell reports*, vol. 6, no. 6, pp. 1073–1084, 2014.
- [13] M. Alamir and S. Chareyron, "State-constrained optimal control applied to cell-cycle-specific cancer chemotherapy," *Optimal Control Applications and Methods*, vol. 28, no. 3, pp. 175–190, 2007.
- [14] S. Chareyron and M. Alamir, "Model-free feedback design for a mixed cancer therapy," *Biotechnology progress*, vol. 25, no. 3, pp. 690–700, 2009.
- [15] M. Alamir, "Robust feedback design for combined therapy of cancer," *Optimal Control Applications and Methods*, 2012.
- [16] B. Zhao, M. T. Hemann, and D. A. Lauffenburger, "Intratumor heterogeneity alters most effective drugs in designed combinations," *Proceedings of the National Academy of Sciences*, vol. 111, no. 29, pp. 10773–10778, 2014.
- [17] B. Zhao, J. R. Pritchard, D. A. Lauffenburger, and M. T. Hemann, "Addressing genetic tumor heterogeneity through computationally predictive combination therapy," *Cancer discovery*, vol. 4, no. 2, pp. 166–174, 2014.
- [18] R. M. E. HernandezVargas, P. Colaneri and F. Blanchini, "Discrete-time control for switched positive systems with application to mitigating viral escape," *International Journal on Robust and Nonlinear Control*, vol. 21, pp. 1093–1111, 2011.
- [19] A. Rantzer and B. Bernhardsson, "Control of convex monotone systems," in *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*. IEEE, 2014.
- [20] F. Blanchini, P. Colaneri, and R. H. Middleton, "A convexity result for the optimal control of a class of positive nonlinear systems," in *19th IFAC World Congress*, 2014.
- [21] R. Luo, M. J. Piovoso, J. Martinez-Picado, and R. Zurakowski, "Optimal antiviral switching to minimize resistance risk in hiv therapy," *PLoS one*, vol. 6, no. 11, p. e27047, 2011.
- [22] R. Zurakowski, "Nonlinear observer output-feedback mpc treatment scheduling for hiv," *Biomedical engineering online*, vol. 10, no. 1, p. 40, 2011.
- [23] V. Jonsson, A. Rantzer, and R. M. Murray, "A scalable formulation for engineering combination therapies for evolutionary dynamics of disease," in *American Control Conference (ACC), 2014*. IEEE, 2014, pp. 2771–2778.
- [24] A. Rantzer, "Distributed control of positive systems," in *Automatic Control, IEEE Transactions on*. In preparation, 2014.
- [25] F. Klein, A. Halper-Stromberg, J. A. Horwitz, H. Gruell, J. F. Scheid, S. Bournazos, H. Mouquet, L. A. Spatz, R. Diskin, A. Abadir, *et al.*, "Hiv therapy by a combination of broadly neutralizing antibodies in humanized mice," *Nature*, vol. 492, no. 7427, pp. 118–122, 2012.
- [26] M. Eigen, "Viral quasispecies," *Scientific American*, vol. 269, pp. 42–49, 1993.
- [27] V. Jonsson, N. Matni, and R. M. Murray, "Reverse engineering combination therapies for evolutionary dynamics of disease: An ?? approach," in *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*. IEEE, 2013, pp. 2060–2065.
- [28] L. Mansky and H. Temin, "Lower in vivo mutation rate of human immunodeficiency virus type 1 than that predicted from the fidelity of purified reverse transcriptase," *Journal of Virology*, vol. 69, no. 8, pp. 5087–94, 1995.
- [29] M. W. Nachman and S. L. Crowell, "Estimate of the mutation rate per nucleotide in humans," *Genetics*, vol. 156, no. 1, pp. 297–304, 2000.
- [30] M. Li, F. Gao, J. R. Mascola, L. Stamatos, V. R. Polonis, M. Koutsoukos, G. Voss, P. Goepfert, P. Gilbert, K. M. Greene, *et al.*, "Human immunodeficiency virus type 1 env clones from acute and early subtype b infections for standardized assessments of vaccine-elicited neutralizing antibodies," *Journal of virology*, vol. 79, no. 16, pp. 10108–10125, 2005.
- [31] A. B. Balazs, J. Chen, C. M. Hong, D. S. Rao, L. Yang, and D. Baltimore, "Antibody-based protection against hiv infection by vectored immunoprophylaxis," *Nature*, vol. 481, no. 7379, pp. 81–84, 2011.
- [32] R. Diskin, J. F. Scheid, P. M. Marcovecchio, A. P. West Jr, F. Klein, H. Gao, P. N. Gnanapragasam, A. Abadir, M. S. Seaman, M. C. Nussenzweig, *et al.*, "Increasing the potency and breadth of an hiv antibody by using structure-based rational design," *Science*, vol. 334, no. 6060, pp. 1289–1293, 2011.
- [33] J. A. Horwitz, A. Halper-Stromberg, H. Mouquet, A. D. Gitlin, A. Tretiakova, T. R. Eisenreich, M. Malbec, S. Gravemann, E. Billerbeck, M. Dorner, *et al.*, "Hiv-1 suppression and durable control by combining single broadly neutralizing antibodies and antiretroviral drugs in humanized mice," *Proceedings of the National Academy of Sciences*, vol. 110, no. 41, pp. 16538–16543, 2013.
- [34] A. S. Perelson, A. U. Neumann, M. Markowitz, J. M. Leonard, and D. D. Ho, "Hiv-1 dynamics in vivo: virion clearance rate, infected cell life-span, and viral generation time," *Science*, vol. 271, no. 5255, pp. 1582–1586, 1996.
- [35] C. Briat, "Robust stability and stabilization of uncertain linear positive systems via integral linear constraints: L1-gain and l-inf gain characterization," *International Journal of Robust and Nonlinear Control*, vol. 23, no. 17, pp. 1932–1954, 2013.