

A Scalable Formulation for Engineering Combination Therapies for Evolutionary Dynamics of Disease

Vanessa Jonsson, Anders Rantzer and Richard M. Murray

Abstract—It has been shown that optimal controller synthesis for positive systems can be formulated as a linear program. Leveraging these results, we propose a scalable iterative algorithm for the systematic design of sparse, small gain feedback strategies that stabilize the evolutionary dynamics of a generic disease model. We achieve the desired feedback structure by augmenting the optimization problems with ℓ_1 and ℓ_2 regularization terms, and illustrate our method on an example inspired by an experimental study aimed at finding appropriate HIV neutralizing antibody therapy combinations in the presence of escape mutants.

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. 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, leading to so called 'treatment-escape'.

One of the solutions that has been proposed [1], [2] is the rational design of combination therapy, an approach that requires the identification of targets that are known to play a key role in disease progression. An example of a multi target therapy currently used for the treatment of chronic HIV infection is highly active antiretroviral therapy (HAART), which is comprised of a combination of antiretroviral drugs that target specific enzymes involved during different points of the infection cycle.

Recent results by Rosenbloom, et al. [1] have been more quantitative in nature, modeling the evolutionary dynamics of HIV and showing through simulations how the effect of antiretroviral dynamics can determine HIV evolution and therapy outcome. The Michor lab [2] recently showed the effects of different *erlotinib* dosing strategies in the presence of pharmacokinetic fluctuations on the evolution of resistance

of non small cell lung cancer through simulations of a stochastic evolutionary dynamics model.

Although these methods have provided some insight into the problem, the challenge of designing treatment protocols that prevent escape is one that has been more recently addressed by control theoretic methods. For cancer therapy, results in this spirit apply methods from optimal and receding horizon control [3], [4], as well as gain scheduling [5], to synthesize treatment protocols that are robust to parameter uncertainty, an inherent issue in all biological systems. In the context of HIV and antiretroviral therapy, the authors in [6] propose a discrete time formulation that allows for the design of switching therapy strategies to delay the emergence of highly resistant mutant viruses.

In [7], we introduced a general algorithm that used an \mathcal{H}_∞ approach for the principled design of targeted combination therapy concentrations that explicitly account for the evolutionary dynamics of a generic disease model. This algorithm was effective in generating robustly stabilizing controllers, however it suffered from an inherent lack of scalability symptomatic of semidefinite programming formulations. Here, we propose a scalable solution to the combination therapy problem by reformulating it as a second order cone program (SOCP).

We also address the requirement that the synthesized controller be not only robust to unmodeled dynamics but also exhibit sparse structure and small feedback gains. This is motivated by the fact that the number of therapies commonly used in combination to treat a disease is often small while the number of potential usable therapies are often very large [8]. Targeted therapies such as small molecule drugs or antibodies exhibit a maximum effective concentration beyond which side effects are likely to worsen and no additional drug benefits are seen.

In particular, through ℓ_1 and ℓ_2 regularization, we induce sparse structure in the feedback controller while bounding the magnitude of the feedback gains. This leads to a SOCP formulation of the combination therapy synthesis problem. The main contribution of this paper is a *scalable* algorithm for the systematic design of *sparse, small gain* feedback strategies to stabilize the evolutionary dynamics of a generic disease model.

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. In Section III, we present our \mathcal{L}_1 combination therapy synthesis algorithm. Section IV illustrates our al-

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

A. Rantzer is with Automatic Control LTH, Lund University, Box 118, SE-221 00 Lund, Sweden. rantzer@control.lth.se.

gorithm in the context of an HIV antibody therapy design problem previously studied in an experimental setting [9]. In this section we also compare performance and scalability properties of our \mathcal{L}_1 algorithm and the previously developed \mathcal{H}_∞ algorithm [7]. 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 $g * w \in \mathcal{L}_r^p[0, \infty)$ 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. Problem formulation

The quasispecies model [10] 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 [7] we incorporated the effects of potential therapies into the basic quasispecies model, by defining a drug binding reaction, $\ell + x \xrightarrow{K_A} \ell \cdot x$ — drug ℓ binds to self replicating macromolecule x with association rate K_A , giving a neutralized complex $\ell \cdot x$. The extended quasispecies model for n mutants and m drugs, is written as:

$$\dot{x}_i = (r_i q_{ii} - d_i)x_i + \sum_{k \neq i} r_i q_{ik} x_k - \sum_{k=1}^m \psi_{ki} \ell_k x_i \quad (1)$$

where $x_i \in \mathbb{R}_+$ is the concentration of mutant i , $\ell_k \in \mathbb{R}_+$ 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 . Finally, $\psi_{ki} = f(K_{ki})$ is a function of the association constant K_{ki} for each neutralization reaction representing the rate at

which a neutralizing macromolecule ℓ_k neutralizes mutant i . The rates r_i and ψ_{ki} can be viewed as replication and neutralization fitnesses of mutant i . When $\ell_k = 0, \forall k \in \{1, \dots, m\}$, the quasispecies dynamics are unstable.

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

$$\begin{aligned} \dot{x} &= (A - \Psi L)x + w \\ z &= Cx \end{aligned} \quad (2)$$

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) $\Psi \in \mathbb{R}^{n \times nm}$ a block diagonal matrix that describes the fitness of n mutants with respect to m different drugs, with diagonal elements $\Psi_i = (\psi_{ki}) \in \mathbb{R}_+^{1 \times m}$; (iii) $L = (I \otimes \ell) \in \mathbb{R}_+^{mn \times n}$, with $\ell = (\ell_k) \in \mathbb{R}^m$, a block diagonal matrix that encodes the concentrations of drugs for all n mutants; (iv) $C = [\mathbf{1}_n \ L^T]^T \in \mathbb{R}^{(mn+1) \times n}$; and (v) $w \in \mathbb{R}_+^n$ an arbitrary positive disturbance. Note that $\Psi L \in \mathbb{R}_+^{n \times n}$ is by construction a diagonal matrix.

We set the regulated output $z = [z_1 \ z_2]^T = Cx$ where $z_1 = x_1 + \dots + x_n$ minimizes the total virus population and $z_2 = Lx$ serves as a proxy to minimizing the concentration of drugs needed to robustly stabilize the system.

Remark 1: 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 $10^{-5} - 10^{-9}$ mutations per base pair per replication cycle for reverse transcriptase to DNA replication.

Letting G denote the closed loop system (2), the control task then becomes to design drug concentrations ℓ by finding a controller $L = (I \otimes \ell)$ that leads to a stable G satisfying $\|G\|_{\infty\text{-ind}} < \gamma$, for some robustness level $\gamma > 0$.

C. Linear programming controller synthesis for positive systems

Our previous work [7] provided a semi definite program (SDP) formulation to synthesize \mathcal{H}_∞ controllers that stabilize the evolutionary dynamics of a generic disease model as described by the extended quasispecies model. One feature of our system was its internal positivity, which allowed us to restrict the storage function matrix used in the bounded real lemma to be strictly diagonal [11] [12]. Although this reduced the number of decision variables in the SDP, it was not enough to offset the inherent lack of scalability of semidefinite programming. Scalability is an important consideration when designing combination therapies where the number of possible mutants can be large, such as for the treatment of chronic HIV infection. Recent results [12],[13] on the synthesis of controllers for positive systems show that the design of structured static state feedback controllers for internally positive systems can be reformulated as a convex problem with methods that scale *linearly* with the number of non zero elements in the feedback matrix. In this section we provide a brief survey of the relevant definitions and results from [13]:

Theorem 1: For the system:

$$\begin{aligned}\dot{x} &= (A + ELF)x + Bw \\ z &= (C + GLF)x + Dw\end{aligned}\quad (3)$$

let \mathcal{D} be the set of $m \times m$ diagonal matrices with entries in $[0,1]$. Suppose that $A + ELF$ is Metzler and $C + GLF \geq 0$ for all $L \in \mathcal{D}$. Let $g(t)$ be the impulse response of

$$G(s) = (C + GLF)[sI - (A + ELF)]^{-1}B + D$$

If the matrices B, D and F have non-negative coefficients, then the following two conditions are equivalent:

- 1) There exists an $L \in \mathcal{D}$ with $A + ELF$ Hurwitz and $\|g\|_{\infty\text{-ind}} < \gamma$.
- 2) There exists a $\xi \in \mathbb{R}_+^n, \mu \in \mathbb{R}_+^m$ with

$$\begin{aligned}A\xi + E\mu + B\mathbf{1} &< 0 \\ C\xi + G\mu + D\mathbf{1} &< \gamma\mathbf{1} \\ \mu &\leq F\xi\end{aligned}$$

If ξ, μ satisfy the linear constraints 2) then the stability and norm guarantees of 1) hold for every L such that $\mu = LF\xi$.

Input-output performance is characterized using induced norms which are determined by the closed loop system's static gain:

Theorem 2: For an $r \times m$ transfer matrix $G(s) = C(sI - A)^{-1}B + D$, let $g(t) = Ce^{At}B + D\delta(t)$ be the corresponding impulse response, where $Ce^{At}B \geq 0$ for $t \geq 0$ and $D \geq 0$, with A Hurwitz. Then $\|g\|_{p\text{-ind}} = \|G(0)\|_{p\text{-ind}}$ for $p = 1, p = 2$ and $p = \infty$.

The positive nature of the system allows us to restrict ourselves to linear storage functions, which in turn allows for sparse structure to be imposed on the feedback gain [12]. Our feedback gain $L = I \otimes \ell$ is not only structurally constrained to be block diagonal, but algebraically constrained as well, in that all block diagonal components must be equal. Unfortunately, there is no known convex reformulation for this additional constraint.

D. Regularization for structured controller synthesis

The biomedical justification for wanting a simple controller structure with small gains is twofold: first, the number of therapies that can be used simultaneously to treat a disease is often limited, and second to minimize side effects, it is desirable to keep the magnitude of drug concentrations small while being robust to evolution. These design specifications can be expressed with the use of regularization, a common technique used in machine learning and inverse problems for *model identification* [14], [15], [16], [17] and increasingly used for controller *design* [18], [19], [20], [21]. As such, we introduce ℓ_1 and ℓ_2 penalties in our design objective to promote controller sparsity and minimize controller gains.

We combine these regularization techniques with controller synthesis results for positive systems and present an iterative algorithm that yields a suboptimal \mathcal{L}_1 controller. This formulation of the combination therapy problem allows the designer to explore explicit trade offs between closed

loop performance, sparsity in controller structure and gain minimization.

III. A SUB-OPTIMAL \mathcal{L}_1 COMBINATION THERAPY CONTROLLER

In this section, we address the aforementioned non-convexity of the optimal control problem by formulating an iterative algorithm for finding effective drug concentrations. Our main result addresses the issue of synthesizing a stabilizing controller subject to the constraints imposed by the quasispecies model (2), with acceptable robustness properties characterized in terms of its ∞ -induced closed loop norm.

A. Initializing stabilizing controller

We begin by synthesizing a stabilizing controller to use as an initial controller in our iterative algorithm, as noted in Remark 5. We recall a simple algorithm developed in [7] for the synthesis of a stabilizing controller, which admits a particularly simple formulation in light of the Metzler nature of A and the diagonal structure of ΨL .

Lemma 1: There exists $\epsilon > 0$ such that the solution to the convex program:

$$\begin{aligned}\text{minimize } & \|\ell\|_{\infty} \\ \text{subject to } & \\ & A_d + \epsilon I - \Psi L \prec 0 \\ & L = I \otimes \ell\end{aligned}\quad (4)$$

is a stabilizing controller for system (2), where $A_d \in \mathbb{R}_+^{n \times n}$ is a diagonal matrix such that $(A_d)_{ii} = (A)_{ii}$ where A is the replication and mutation matrix in equation (2).

Remark 2: The resulting LMI is strictly diagonal, and so reduces to an LP.

Proof: Rewrite $A = A_d + M$ where A_d is diagonal and $M = \{m_{ij}\} \in \mathbb{R}^{n \times n}$, $m_{ij} = 0$ for $i = j$ and $m_{ij} > 0$ for $i \neq j$. By the Perron Frobenius theorem, there exists $r > 0$ such that the spectral radius $\rho(M) = r \leq \max_i \sum m_{ij}$. Let $\epsilon = \max_i \sum m_{ij}$ and rewrite $M = \epsilon I - (\epsilon I - M)$. We note that $-(\epsilon I - M) \prec 0$. The closed loop dynamics are then given by $A - \Psi L = A_d + \epsilon I - (\epsilon I - M) - \Psi L \prec A_d + \epsilon I - \Psi L \prec 0$, yielding the desired stability. ■

B. A Suboptimal \mathcal{L}_1 combination therapy controller

Observe that through a straightforward application of Theorem 2, with $B = I$, $C = [\mathbf{1}_n \ L^T]^T$, $D = 0$, $E = -\Psi$, $F = I$ to system (1), solving the following non-convex program:

$$\begin{aligned}\text{minimize } & \|Cx\|_{\infty} + \lambda_1 \|\ell\|_1 + \lambda_2 \|\ell\|_2 \\ \text{subject to } & \\ & Ax + Kx + \mathbf{1} \leq 0 \\ & K = \Psi L \\ & L = I \otimes \ell \\ & x \geq 0\end{aligned}\quad (5)$$

will yield a sparse combination of drug concentrations ℓ , yielding an optimal ∞ -induced closed loop norm for appropriately chosen regularizers $\lambda_1 \geq 0, \lambda_2 \geq 0 \in \mathbb{R}$.

Remark 3: We can impose an additional constraint limiting the concentrations of candidate therapies. This is necessary with certain drugs that have maximum tolerated doses dictated by clinical trials.

As mentioned earlier, there are no known convex reformulations of this problem due to the additional structure on L . As such, we suggest the following iterative algorithm, based on the convex programs (6) and (7), to find a stabilizing controller. For notation, let $Y = P_Z(x, \gamma)$ denote an optimization problem P in which we optimize over x and γ leaving Z fixed and with solution Y .

Program 1. $P1_\ell(x, \gamma)$:

$$\begin{aligned} & \underset{\gamma, x \in \mathbb{R}_+^n}{\text{minimize}} \quad \gamma \\ & \text{subject to} \\ & \quad Ax + \Psi Lx + \mathbf{1} \leq 0 \\ & \quad L = I \otimes \ell, \quad C = [\mathbf{1}_n \quad L^T]^T \\ & \quad x \geq 0 \\ & \quad \gamma \geq \|Cx\|_\infty \end{aligned} \quad (6)$$

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

$$\begin{aligned} & \underset{\gamma, \ell \in \mathbb{R}_+^m}{\text{minimize}} \quad \gamma + \lambda_1 \|\ell\|_1 + \lambda_2 \|\ell\|_2 \\ & \text{subject to} \\ & \quad Ax + \Psi Lx + \mathbf{1} < 0 \\ & \quad L = I \otimes \ell, \quad C = [\mathbf{1}_n \quad L^T]^T \\ & \quad \gamma \geq \|Cx\|_\infty \end{aligned} \quad (7)$$

We now present our algorithm:

Algorithm 1 Scalable Combination Therapy

- 1) Set $\epsilon > 0$.
 - 2) *Solve for initial stabilizing controller ℓ' :*
 Solve (4) to obtain controller ℓ_0 .
 Set $(x', \gamma) = P1_{\ell_0}(x, \gamma)$.
 Set $(\ell', \gamma) = P2_{(x', 0, 0)}(\ell, \gamma)$.
 - 3) *Find $(\lambda'_1, \lambda'_2, \ell)$ that minimize γ :*
 $\forall (\lambda_1, \lambda_2) \in \Lambda_1 \times \Lambda_2, \Lambda_1, \Lambda_2 \in \mathbb{R}_+^k$,
 while $\gamma' - \gamma > \epsilon$:
 Set $(x', \gamma) = P1_{\ell'}(x, \gamma)$.
 Set $(\ell', \gamma) = P2_{(x', \lambda_1, \lambda_2)}(\ell, \gamma)$.
 Set $\gamma' = \gamma$.
-

Remark 4: The sequence of γ 's defined by the iterative process in Algorithm 1 is non increasing by construction, and bounded below by 0, thus implying convergence. An initial stabilizing controller can always be found as shown in [7], and thus the algorithm can always be initialized. We therefore have that our algorithm always converges to a feasible, local minimum robustness value γ , generating a stabilizing controller for (2).

Remark 5: In practice, we note that the \mathcal{L}_1 controller suffers from dependence on initial conditions and converges

to local optima quickly, yielding a stabilizing controller with robustness properties that are not significantly different from the nominal controller. A solution to this is to iterate once through P1 and P2, with $\lambda_1 = \lambda_2 = 0$ and initialize the algorithm with the resulting controller.

Remark 6: Due to the presence of the ℓ_2 regularizer, P2 and (5) are SOCPs and not linear programs. As it will be clearly demonstrated in the example in the next section, this is still more efficient than the SDP combination therapy algorithm from [7]. In addition, the second order cone constraint is only on the drug concentrations so it should have minimal effect on performance, given that the number of antibodies is small compared to the number of mutants.

IV. HIV/ANTIBODY THERAPY APPLICATION

Our results provide a principled approach to the design of antibody treatments for chronic infection with human immunodeficiency virus-1 (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 [9].

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 [22]. These have been shown to inhibit infection by a broad range of viral isolates in vitro but also protect non-human primates against infection [22],[23], [24]. 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 [9]. 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. 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 so this algorithm would be most effective in conjunction with single antibody selection experiments.

1) *Model parameters:* We consider a system of eighteen to thirty five HIV mutants with five potential antibodies to use in combination. Figure 5 lists the mutants that evolved from monotherapy experiments with their corresponding half maximal inhibitory antibody concentration (IC50) in $\mu\text{g/ml}$, as measured by the Nussenzweig lab in [9]. Antibodies 3BC176, PG16, 45-46G54W, PGT128 and 10-1074 are potential combination therapy candidates.

Although virus replication rates can vary considerably depending on the nature of the mutations a virus may undergo, 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




Fig. 1. The graphs depict the average of thirty simulations subject to random time invariant perturbations of 5.5 % in the plant dynamics found with the \mathcal{H}_∞ and \mathcal{L}_1 combination therapy algorithms for evolutionary dynamics of the first 18 point mutants in Figure (5). (Left) Full support controllers synthesized with the pentamix of antibodies available: 3BC176, PG16, 45-46G54W, PGT128 and 10-1074 and (Right) Sparse controllers synthesized with only two antibodies 45-46G54W and PGT128.

experiments and assume that replication rate variability due to mutations were negligible.

The fitness function associated with the neutralization of a virus i with respect to an antibody j is a Hill function $\psi_{ij} = \frac{\ell_j^n}{\ell_j^n + K_{ij}^n}$ where n is the Hill coefficient, ℓ_j is the concentration of a given antibody j , and $K_{ij} = \frac{k_{on}}{k_{off}} = \frac{[x_i \ell_j]}{[x_i][\ell_j]}$ is the association constant for the virus/antibody binding reaction $\ell_j + x_i \xrightarrow{k_{on}} \ell_j \cdot x_i$, and k_{on} and k_{off} are the on and off reaction rate constants. Note that the association constant represents the fraction bound of antibody/virus complexes in solution and that $K_{ij} = \frac{3 \cdot \text{IC}_{50_{ij}}}{3r_i + \ln(2) - \text{IC}_{50_{ij}}}$, is found by solving Equation (1) for one virus/antibody pair for the duration $[t_0, t_f] = [0, 3]$. We simplify the Hill function by setting the Hill coefficient $n = 1$, as there is evidence that that antibodies do not bind cooperatively. Our algorithm yields antibody concentrations near zero and this yields the linear approximation $\psi_{ij} = \frac{1}{K_{ij}} \ell_j$. In addition, the antibodies we consider in our example do not target the same epitope, in other words, do not bind competitively to the same sites on the virus, thereby reducing any coupling between antibody concentrations.

2) *Mutation process*: The mutation rate for HIV reverse transcriptase is $u = 3 \times 10^{-5}$ mutations/nucleotide base pair/replication cycle, and the HIV replication cycle is approximately 2.6 days. 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.

A. A comparison between \mathcal{H}_∞ and \mathcal{L}_1 controllers

We seek to provide a qualitative comparison between controllers synthesized using the scalable \mathcal{L}_1 and the \mathcal{H}_∞ algorithms. To do this, we adapt the formulation in [7]




Fig. 2. The graph depicts the average \mathcal{H}_∞ (left) norm and ∞ -ind norm as a function of the sparsity of controllers found using either the \mathcal{H}_∞ or \mathcal{L}_1 combination therapy algorithms.

to include ℓ_1 and ℓ_2 regularization terms and solve the following non convex problem using our iterative algorithm.

$$\begin{aligned} & \text{minimize } \gamma + \lambda_1 \|\ell\|_1 + \lambda_2 \|\ell\|_2 \\ & \text{subject to} \\ & \begin{bmatrix} A_{cl}^T X + X A_{cl} + C^T C & X \\ X & -\gamma^2 I \end{bmatrix} \prec 0 \\ & A_{cl} = (A - \Psi L) \\ & C = [\mathbf{1}^T L^T]^T \\ & L = I \otimes \ell \\ & X \succ 0, X \text{ diagonal} \end{aligned} \quad (8) \end{aligned}$$

We synthesized a nominal stabilizing controller using (4), a robust controller that minimizes the \mathcal{H}_∞ closed loop norm using (8), and a robust controller using (7) that minimizes the \mathcal{L}_1 closed loop norm for the evolutionary dynamics of the eighteen HIV point mutants listed in Figure 5. We found similar gains and robustness properties for both sparse and full controllers using either algorithm with the notable difference seen in computational time. Not surprisingly, the \mathcal{L}_1 algorithm has far superior performance, beating the runtime for the \mathcal{H}_∞ synthesis algorithm by four orders of magnitude.

We averaged thirty simulations of closed loop evolutionary dynamics subject to 5.5% random time invariant perturbations in the plant dynamics using both sparse and full support \mathcal{H}_∞ and \mathcal{L}_1 controllers. The sparse controller found by the \mathcal{L}_1 algorithm performed better than the one found by \mathcal{H}_∞ algorithm, whereas the situation was reversed for the respective synthesized full support controllers. As previously mentioned, the motivation for generating sparse controllers for combination therapy is that number of therapies commonly used in combination to treat a disease is often limited for clinical reasons. Therefore, the potential for the \mathcal{L}_1 algorithm to synthesize controllers that are not only sparse but more robustly stable than \mathcal{H}_∞ algorithm is a desirable feature.

Figure 2 shows the relationship between gain sparsity and both \mathcal{H}_∞ and ∞ -induced norms in the synthesis of \mathcal{H}_∞ and \mathcal{L}_1 controllers. Although closed loop \mathcal{H}_∞ norms remain constant with respect to sparsity for both controllers, the closed loop ∞ -induced norm decreases with sparsity with the suggesting that the \mathcal{L}_1 synthesis algorithm, as expected, finds better performing *sparse* controllers.

We note that due to computational limitations, we were

not able to synthesize controllers using the \mathcal{H}_∞ algorithm for the full set of thirty five mutants from Figure 5.

B. L_1 controller synthesis, full HIV example

We synthesized a nominal stabilizing controller using (4) comprised of an antibody pentamix (0.4687, 0.7815, 0.6129, 0.6279, 0.8831) $\mu\text{g/ml}$ of (3BC176, PG16, 45-46G54W, PGT128, 10-1074), and a robust controller using (8) that consisted of antibody trimix (0.6891, 0.6712, 1.0706) $\mu\text{g/ml}$ of (3BC176, 4546-G54W, PGT128). These both were generated for the evolutionary dynamics of the full, thirty five HIV mutants listed in Figure 5.

Both antibody pentamix (stabilizing) and trimix (robustly stabilizing) controllers have similar gains and based on a cursory first glance, one might be tempted to believe these have comparable robustness properties. Indeed, for some simulations of the closed loop dynamics subjected to 5% random time invariant perturbations in plant dynamics, the nominal controller is stabilizing as seen in Figure 3. This is qualitatively consistent with the experimental results done in by the Nussenzweig lab in [9]. It was shown that with weekly injections of equal concentrations of the antibody pentamix holding concentrations constant, had viral loads that remained below the limit of detection during an entire treatment course in mice.

In [9] again, an antibody trimix of equal concentrations of 3BC176, PG16 and 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 pre-treatment levels, suggesting that in these cases, the virus had evolved mutations that were resistant to the trimix treatment. To compare the performance of our \mathcal{L}_1 synthesized controller with gains of (0.6891, 0.6712, 1.0706) $\mu\text{g/ml}$ of (3BC176, 4546-G54W, PGT128) to the experimentally studied trimix, we chose equal concentrations of (3BC176, PG16, 45-46G54W), namely (1, 1, 1) $\mu\text{g/ml}$ for the experimentally derived trimix. We found that even though total antibody concentrations were larger in our version of the experimental trimix, the robustly stabilizing controller synthesized by the \mathcal{L}_1 algorithm nonetheless performed overall better; the closed loop norms were $\|G\|_\infty = 0.2941$ and $\|G\|_{\infty-\text{ind}} = 0.6533$ for the \mathcal{L}_1 controller versus $\|G\|_\infty = 0.26433$ and $\|G\|_{\infty-\text{ind}} = 0.74572$ for the experimental trimix.

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 disease, but to explore tradeoffs between number of therapies




Fig. 3. Sum of virus populations subject to random time invariant perturbations of 5% in the dynamics for 30 different simulations for (left) a stabilizing closed loop controller comprised of antibody pentamix (0.4687, 0.7815, 0.6129, 0.6279, 0.8831) $\mu\text{g/ml}$ of (3BC176, PG16, 45-46G54W, PGT128, 10-1074) synthesized using the convex program (5) and (right) a robustly stabilizing closed loop controller comprised of antibody trimix (0.6891, 0.6712, 1.0706) $\mu\text{g/ml}$ of (3BC176, 4546-G54W, PGT128) synthesized using the \mathcal{L}_1 combination therapy algorithm.

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

Leveraging recent results in positive systems, we proposed a scalable SOCP based iterative algorithm for the systematic design of sparse, small gain feedback strategies that stabilize the evolutionary dynamics of a generic disease model. Through the addition of ℓ_1 and ℓ_2 regularization terms to the objective function, we achieved the desired feedback structure. 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 [25]. In particular, we hope to explicitly account for model error introduced due to model linearization, parametric uncertainty and unmodeled dynamics due to drug interactions.

VI. ACKNOWLEDGEMENTS

We would like to thank David Baltimore and Pamela Bjorkman for discussions regarding using antibody therapy for HIV treatment, and research scientist Anthony West for information on antibody neutralization parameters. We would like to thank Nikolai Matni for the review of the technical content. A. Rantzer gratefully acknowledges support of the LCCC Linnaeus Center and the eLLIIT Excellence Center at Lund University.

REFERENCES

- [1] D. I. S. Rosenbloom, A. L. Hill, S. A. Rabi, R. F. Siliciano, and M. A. Nowak, "Antiretroviral dynamics determines hiv evolution and predicts therapy outcome," *Nature Biotechnology*, vol. 18, pp. 1378–1385, June 2012.
- [2] F. Michor, Y. Iwasa, and M. A. Nowak, "Dynamics of cancer progression," *Nature Reviews Cancer*, vol. 4, no. 3, pp. 197–205, 2004.
- [3] 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.
- [4] S. Chareyron and M. Alamir, "Model-free feedback design for a mixed cancer therapy," *Biotechnology progress*, vol. 25, no. 3, pp. 690–700, 2009.
- [5] M. Alamir, "Robust feedback design for combined therapy of cancer," *Optimal Control Applications and Methods*, 2012.

Controller	Gains $\mu\text{g/ml}$					$\ G\ _{\infty}$	$\ G\ _{\infty-\text{ind}}$	Time
	3BC176	PG16	45-46G54W	PGT128	10-1074			
Nominal Stabilizing	0.0125	0.0125	0.0125	0.0125	0.0125	2.7543	8.9786	1.53 s
Full \mathcal{H}_{∞}	0.0550	0.0398	0.1587	0.1767	0.0629	0.1723	0.5988	> 8 hours
Full \mathcal{L}_1	0.0424	0.0038	0.1125	0.2635	0.0548	0.1786	0.6722	35.803 s
Sparse \mathcal{H}_{∞}	0	0	0.1485	0.1987	0	0.1917	0.6592	> 8 hours
Sparse \mathcal{L}_1	0	0	0.1175	0.2971	0	0.1766	0.6616	35.803 s
Nominal Stabilizing (35 mutants)	0.4687	0.7815	0.6129	0.6279	0.8831	0.334	0.805	3.76 s
Sparse \mathcal{L}_1 (35 mutants)	0.6891	0	0.6712	1.0706	0	0.2941	0.6533	30.05 s

Fig. 4. Stabilizing gains found for nominal stabilizing controller, a robust controller using (8) that minimizes the \mathcal{H}_{∞} closed loop norm and a robust controller using (6) that minimizes the \mathcal{L}_1 norm of the closed loop system for evolutionary dynamics systems of the first eighteen HIV point mutants listed in Figure 5.

Antibody associated escape mutants	Mutation	3BC176 IC50 $\mu\text{g/ml}$	PG16 IC50 $\mu\text{g/ml}$	45-46G54W IC50 $\mu\text{g/ml}$	PGT128 IC50 $\mu\text{g/ml}$	10-1074 IC50 $\mu\text{g/ml}$
	WT	0.319	0.612	0.024	0.169	0.312
3BC176	G471R	0.159	0.154	0.008	0.02	0.091
PG16	N160K	0.145	50	0.007	0.086	0.155
	T162N	0.154	50	0.013	0.166	0.175
45-46G54W	N279H	0.209	0.294	50	0.064	0.177
	N280Y	0.276	0.145	50	0.031	0.126
PGT128 or 10-1074	N332K	0.232	0.988	0.017	50	50
	N332Y	0.269	0.632	0.01	50	13.596
	S334N	0.218	0.615	0.02	50	7.308
Passenger mutations	Y61H	0.243	0.285	0.015	0.098	0.26
	E102K	0.173	0.341	0.023	0.11	0.207
	N295S	0.347	0.5	0.017	0.145	0.159
	I311M	0.23	2.67	0.013	0.248	0.253
	S365L	0.26	0.273	0.009	0.045	0.153
	G366E	0.187	0.167	0.001	0.021	0.074
	I371M	0.2	0.303	0.013	0.064	0.164
	N413K	0.188	0.557	0.014	0.032	0.109
	E429K	0.146	0.503	0.017	0.082	0.167
N295S-G366E-N413K	0.222	0.131	0.001	0.012	0.021	
tri-mix	T162I-G458D	0.275	50	14.33	0.012	0.047
	T162N-N280Y	0.138	50	50	0.027	0.079
penta-mix	N160K-N280Y-N332K	0.146	50	50	50	50
	N160K-A281T-N332K	0.1	50	50	50	50
	T162I-N280Y-N332K	0.13	50	50	50	50
	T162I-N279K-N332K	0.149	50	50	50	50
Signature + Passenger	T162I-Y61H	0.156	50	0.014	0.088	0.115
	T162N-V430E	0.167	50	0.003	0.037	0.106
	N280Y-A174T	0.064	0.138	50	0.01	0.021
	N332S-N413K	0.181	0.526	0.017	50	50
Estimated Mutations	N160K-N280Y	0.276	50	50	0.086	0.155
	N160K-N332K	0.232	50	0.017	50	50
	N280Y-N332K	0.276	0.988	50	50	50
	N295S-G366E	0.347	0.5	0.017	0.145	0.159
	N295S-N413K	0.347	0.557	0.017	0.145	0.159
	G366E-N413K	0.188	0.557	0.014	0.032	0.109

Fig. 5. IC50 values for the indicated antibodies on YU2 mutant viruses found in continuous antibody mono therapy experiments conducted by the Nussenzweig lab at Rockefeller University [9]. The trimix of antibodies is : 3BC176,PG16,45-46G54W, the penta-mix is 3BC176, PG16, 45-46G54W , PGT128 and 10-1074. Estimated two point mutations represent intermediary mutations needed for our model but not included in experimental results in [9]. The IC50 values were taken to be the maximum IC50 of both mutations.

- [6] 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.
- [7] V. Jonsson, N. Matni, and R. M. Murray, “Reverse engineering combination therapy for evolutionary dynamics of disease: an \mathcal{H}_∞ approach,” in *To appear in the IEEE Conf. on Decision and Control, 2013*, 2013.
- [8] B. Al-Lazikani, U. Banerji, and P. Workman, “Combinatorial drug therapy for cancer in the post-genomic era,” *Nature Biotechnology*, vol. 30, pp. 679–692, July 2012.
- [9] 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*, 2012.
- [10] M. Eigen, “Viral quasispecies,” *Scientific American*, vol. 269, pp. 42–49, 1993.
- [11] T. Tanaka and C. Langbort, “The bounded real lemma for internally positive systems and h-infinity structured static state feedback,” *Automatic Control, IEEE Transactions on*, vol. 56, no. 9, pp. 2218–2223, 2011.
- [12] A. Rantzer, “Distributed control of positive systems,” in *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, Dec. 2011, pp. 6608–6611.
- [13] —, “Distributed control of positive systems,” in *Automatic Control, IEEE Transactions on*. In preparation, 2014.
- [14] E. Candès and B. Recht, “Exact matrix completion via convex optimization,” *Commun. ACM*, vol. 55, no. 6, pp. 111–119, June 2012.
- [15] D. L. Donoho, “Compressed sensing,” *IEEE Trans. Inform. Theory*, vol. 52, pp. 1289–1306, 2006.
- [16] B. Recht, M. Fazel, and P. A. Parrilo, “Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization,” *SIAM Rev.*, vol. 52, no. 3, pp. 471–501, Aug. 2010.
- [17] V. Chandrasekaran, B. Recht, P. Parrilo, and A. Willsky, “The convex geometry of linear inverse problems,” *Foundations of Computational Mathematics*, vol. 12, pp. 805–849, 2012.
- [18] F. Lin, M. Fardad, and M. Jovanovic, “Sparse feedback synthesis via the alternating direction method of multipliers,” in *American Control Conference (ACC), 2012*, 2012, pp. 4765–4770.
- [19] N. Dhingra, F. Lin, M. Fardad, and M. R. Jovanovic, “On identifying sparse representations of consensus networks,” in *3rd IFAC Workshop on Distributed Estimation and Control in Networked Systems, Santa Barbara, CA, 2012*, pp. 305–310.
- [20] M. Fardad, F. Lin, and M. R. Jovanović, “Design of optimal sparse interconnection graphs for synchronization of oscillator networks,” *arXiv preprint arXiv:1302.0449*, 2013.
- [21] N. Matni, “Communication delay co-design in \mathcal{H}_2 decentralized control using atomic norm minimization,” in *IEEE Conf. on Decision and Control, 2013*, 2013.
- [22] M. Li, F. Gao, J. R. Mascola, L. Stamatatos, 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. 10 108–10 125, 2005.
- [23] 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.
- [24] 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.
- [25] C. Briat, “Robust stability and stabilization of uncertain linear positive systems via integral linear constraints: L1-and l-infinity-gains characterization,” *arXiv preprint arXiv:1204.3554*, 2012.