2015 American Control Conference
Palmer House Hilton
July 1-3, 2015. Chicago, IL, USA

Convex synthesis of symmetric modifications to linear systems

Neil K. Dhingra and Mihailo R. Jovanović

***Abstract*— We develop a method for designing symmetric modifications to linear dynamical systems for the purpose of optimizing $\mathcal{H}_2$ performance. For systems with symmetric dynamic matrices this problem is convex. While in the absence of symmetry the design problem is not convex in general, we show that the $\mathcal{H}_2$ norm of the symmetric part of the system provides an upper bound on the $\mathcal{H}_2$ norm of the original system. We then study the particular case where the modifications are given by a weighted sum of diagonal matrices and develop an efficient customized algorithm for computing the optimal solution. Finally, we illustrate the efficacy of our approach on a combination drug therapy example for HIV treatment.**

***Index Terms*— combination drug therapy, networks, sparse controller synthesis, structured design, symmetric systems.**

I. Introduction

Structured feedback control design problems are challenging and, in general, nonconvex. Consequently, significant effort has been devoted to identifying classes of problems which admit a convex formulation. These include funnel causal and quadratically invariant systems [1], [2], positive systems [3], structured and sparse consensus and synchronization networks [4]–[6], and optimal sensor/actuator selection [7], [8]. In this paper, we identify another class of optimal control problems that admit a convex characterization. We also examine how design problems for these systems can be used to inform design problems for more general systems.

We are interested in minimizing an $\mathcal{H}_2$ performance objective via symmetric modifications to the dynamical generator of a linear time-invariant system. The symmetric modification is a linear function of a design variable which is subject to additional penalties. Following much recent work on sparse feedback synthesis [9], [10], we penalize the ℓ_1 norm of the design variable to promote sparsity.

An interesting problem instance arises in the design of combination drug therapy strategies for the treatment of the Human Immunodefficiency Virus (HIV). Recent advances is the modeling of disease [11] motivate the design of drug dosages in a manner which accounts for the evolutionary dynamics of the disease. While several recent papers focus on designing $\mathcal{H}_\infty$ or suboptimal $\mathcal{L}_1$ controllers [12]–[14], the $\mathcal{H}_2$ problem for these systems still has not been addressed.

The proposed formulation also encompasses several problems that arise in networked control systems; see Section II-B. For example, the leader selection problem seeks to identify nodes which have the most profound influence on the network performance [15]. Although analytical expressions exist for the selection of one or two leaders [16], the selection of a higher numbers of leaders amounts to solving a combinatorial optimization problem. Moreover, the problem of adding undirected edges to an existing consensus network also fits into this framework [17]. In this case, the objective is to add a sparse set of undirected edges in order to improve network performance [18].

It is worth noting that even with symmetric modifications to the system dynamics, the $\mathcal{H}_2$ problem is in general nonconvex. Our approach is to instead examine the symmetric part of the dynamics. We show that this formulation provides an upper bound on the performance of the original problem, and when the underlying system is not strongly non-normal, it is a good approximation. We also demonstrate that the restriction to symmetric systems makes the design problem a semidefinite program (SDP). We develop a scalable algorithm for designing diagonal modifications to the dynamics and apply it to the HIV combination therapy problem.

The rest of this paper is structured as follows. Section II states the general form of the problem we consider. Section III presents results relating systems to their symmetric parts and justifies our use of these symmetric systems for design. In section IV, we present a relevant application, form the primal and dual optimization problems, and develop a customized algorithm to solve it. Section V applies our algorithm to real world data. Finally, section VI offers concluding remarks and outlines future research directions.

II. Problem formulation

A. General form

We consider a class of systems,

$$\dot{x} \;=\; (A \;-\; K(u))\, x \;+\; d \tag{1}$$

where $K(u) \in \mathbb{R}^{n\times n}$ is a symmetric matrix that is a linear function of $u \in \mathbb{R}^m$, $x(t) \in \mathbb{R}^n$ is the state vector, $d(t) \in \mathbb{R}^p$ is a white stochastic disturbance with $\mathbf{E}\,(d(t_1)\,d^T(t_2)) = I\delta(t_1 - t_2)$, and $\mathbf{E}$ is the expectation operator. Our objective is to design a stabilizing $K(u)$ that minimizes the steady-state variance of the state x,

$$\lim_{t\,\to\,\infty} \mathbf{E}\left(x^T(t)\,x(t)\right).$$

We also impose a quadratic penalty, $u^T R u$ with $R \succ 0$, to limit the magnitude of u, and an ℓ_1 penalty, $\|u\|_1 := \sum_i |u_i|$,

Financial support from the National Science Foundation under award ECCS-1407958, the University of Minnesota Informatics Institute Transdisciplinary Faculty Fellowship, the NASA Harriet G. Jenkins Predoctoral Fellowship Program, and the MnDRIVE Graduate Scholars Program is gratefully acknowledged.

N. K. Dhingra and M. R. Jovanović are with the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455. E-mails: dhin0008@umn.edu, mihailo@umn.edu.

to promote sparsity in u. The positive parameter γ specifies the importance of sparsity relative to system performance.

Our design problem is formulated as

$$\begin{array}{ll} \underset{u}{\text{minimize}} & J(u) \;+\; u^T R u \;+\; \gamma \|u\|_1 \\ \text{subject to} & A \;-\; K(u) \;\; \text{Hurwitz.} \end{array} \tag{2}$$

Here, $J(u)$ is the $\mathcal{H}_2$ norm of system (1) from d to x,

$$J(u) \;:=\; \text{trace}(X)$$

and X is the controllability gramian,

$$(A \;-\; K(u))\, X \;+\; X\,(A \;-\; K(u))^T \;+\; I \;=\; 0.$$

Since in (2) we do not impose penalty on the L_2 norm of the control effort $-K(u)x$, for $\gamma = 0$ the problem (2) is different from the Linear Quadratic Regulator problem.

In general, the optimal control problem (2) is non convex. However, it is instructive to examine

$$\dot{x} \;=\; (A_s \;-\; K(u))\, x \;+\; d \tag{3}$$

where $A_s := (A + A^T)/2$ is the symmetric part of A. As we show below, designing a symmetric $K(u)$ for system (3) is a convex optimization problem. The $\mathcal{H}_2$ norm of system (3) is given by $\text{trace}\,(X_s)$ where,

$$\big(A_s \;-\; K(u)\big)X_s \;+\; X_s\big(A_s \;-\; K(u)\big) \;+\; I \;=\; 0.$$

Since both A_s and $K(u)$ are symmetric,

$$X_s \;=\; -\frac{1}{2}\,(A_s \;-\; K(u))^{-1}$$

and the design problem (3) becomes

$$\begin{array}{ll} \text{minimize} & \dfrac{1}{2}\,\text{trace}\left((K(u) - A_s)^{-1}\right) \;+\; u^T R u \;+\; \gamma\,\|u\|_1 \\ & K(u) \;-\; A_s \;\succ\; 0. \end{array} \tag{4}$$

Taking the Schur complement casts (4) as an SDP,

$$\begin{array}{ll} \underset{\Theta,\,u}{\text{minimize}} & \text{trace}\,(\Theta) \;+\; u^T R u \;+\; \gamma\,\|u\|_1 \\ \text{subject to} & \left[\begin{array}{cc} \Theta & I \\ I & K(u) - A_s \end{array}\right] \succ 0 \end{array} \tag{5}$$

where we drop the constant factor $1/2$ for compactness.

In Section III, we show that stability of the symmetric system (3) implies stability of the corresponding original system (1). Furthermore, the $\mathcal{H}_2$ norm of the symmetric system is an upper bound on the $\mathcal{H}_2$ norm of the original system. Finally, when the difference between A and A_s is small (of the order ϵ, $O(\epsilon)$), the $\mathcal{H}_2$ norms of systems (1) and (3) differ only by $O(\epsilon^2)$.

If $K(u)$ is not symmetric, $K_s(u) := \frac{1}{2}(K(u) + K(u)^T)$ can be used in (4) to find its solution $u^\star$. The closed-loop systems (1), with $K(u^\star)$, and (3), with $K_s(u^\star)$, will have the same relationship as when $K(u)$ is symmetric. However, the neglected effect of the asymmetric component of $K(u)$ makes the degree of conservatism unpredictable.

B. Applications

We next provide examples to illustrate that the considered problem structure arises in several applications.

1) Design of edges in consensus networks: The problem of adding undirected edges to an existing consensus network can be cast in this problem form. The dynamics are,

$$\dot{x} \;=\; -\big(L \;+\; E D(u) E^T\big)\; x \;+\; d$$

where L is a graph Laplacian which contains information about which nodes are connected to each other, E contains information about which edges may be added to the network, and $D(u)$ is a diagonal matrix of added edge weights [17].

2) Leader selection in consensus networks: In this setup, it is desired to identify influential nodes in networks. These special nodes, so-called leaders, can be equipped with additional information in order to influence the network behavior in a beneficial way. One application is in vehicular formations, where the objective is for the vehicles to gather at a certain point. The 'leaders' are equipped with absolute measurements (e.g., from GPS units) and the other nodes have only relative measurements (e.g., their distance from certain neighbors). The dynamics are given by

$$\dot{x} \;=\; -\,(L \;+\; D(u))\; x \;+\; d$$

where L is a graph Laplacian and $D(u)$ is a nonnegative diagonal matrix whose nonzero entries identify the leaders. In contrast to earlier work [15] which treated this problem directly as a combinatorial problem, leader selection with our formulation would amount to a convex relaxation.

3) Combination drug therapy design for HIV treatment: The problem of designing drug dosages for treating HIV can be cast as [13], [14],

$$\dot{x} \;=\; \left(A \;-\; \sum_{k\,=\,1}^{m} u_k D_k\right) x \;+\; d.$$

Here, the elements of x represent populations of HIV mutants. The diagonal elements of A represent each mutant's replication rate and the off diagonal elements of A represent the probability of mutation from one mutant to another. The components of the vector u are dosages of different drugs, where D_k is a diagonal matrix containing information about how efficiently drug k kills each HIV mutant.

III. Symmetric system design

We justify the use of the symmetric synthesis problem (4) for the design of controllers for the original system (1). First we show that the stability of the symmetric system implies stability of the full system.

Lemma 1: Let the symmetric part of A be Hurwitz. Then, A is Hurwitz.

Proof: We show this by contradiction. Since the symmetric part of A, $A_s := (A + A^T)/2$ is symmetric and Hurwitz, it is negative definite,

$$v^* A_s\, v \;<\; 0 \quad \text{for all } v \neq 0.$$

Assume that A is not Hurwitz. Then there is a v such that $Av = \lambda v$ with $\mathrm{Re}(\lambda) \geq 0$. Furthermore, $v^*Av = \lambda v^*v$. However,

$$\begin{aligned} v^*Av &= v^*A_sv \ + \ \tfrac{1}{2}\, v^*\left(A \ - \ A^T\right) v \\ &= v^*A_sv \ + \ j\,\mathrm{Im}(\lambda)\, \|v\|_2^2. \end{aligned}$$

Since $A_s \prec 0$, v^*A_sv cannot have a nonnegative real part. ∎

Note that Lemma 1 only provides a sufficient condition. It is standard that A may be Hurwitz even if A_s is not.

We next show that the $\mathcal{H}_2$ norm of the symmetric system (3) is an upper bound on the $\mathcal{H}_2$ norm of the system (1). First, we present a useful theorem from linear algebra [19].

Lemma 2: Let A be any matrix and let $A_s = \frac{1}{2}(A + A^T)$ be the symmetric part of A. Then

$$\|\mathrm{e}^A\| \ \leq \ \|\mathrm{e}^{A_s}\|$$

for every unitarily invariant norm.

Proof: See Theorem IX.3.1 in [19]. ∎

The statement about the $\mathcal{H}_2$ norms of systems (1) and (3) is a simple corollary of Lemma 2.

Corollary 3: When systems (1) and (3) are stable, the $\mathcal{H}_2$ norm of (1) is bounded from above by the $\mathcal{H}_2$ norm of (3).

Proof: Recall that the $\mathcal{H}_2$ norm of a system,

$$\dot{x} \ = \ A\,x \ + \ d$$

with A Hurwitz is given by $\mathrm{trace}(X)$ where

$$A\,X \ + \ XA^T \ + \ I \ = \ 0$$

and

$$X \ = \ \int_0^\infty \mathrm{e}^{At}\,\mathrm{e}^{A^Tt}\,\mathrm{d}t. \tag{6}$$

Using the linearity of the trace and of integration, we can rewrite the expression for the $\mathcal{H}_2$ norm as,

$$\mathrm{trace}\left(\int_0^\infty \mathrm{e}^{At}\,\mathrm{e}^{A^Tt}\mathrm{d}t\right) \ = \ \int_0^\infty \|\mathrm{e}^{At}\|_F^2\,\mathrm{d}t.$$

Since the Frobenius norm is unitarily invariant, by Lemma 2 $\|\mathrm{e}^{At}\|_F^2 \leq \|\mathrm{e}^{A_st}\|_F^2$ for any t and therefore,

$$\int_0^\infty \|\mathrm{e}^{At}\|_F^2\ \mathrm{d}t \ \leq \ \int_0^\infty \|\mathrm{e}^{A_st}\|_F^2\ \mathrm{d}t.$$

Since the right-hand-side is the $\mathcal{H}_2$ norm of system (3), this completes the proof. ∎

Remark 1: Lemma 2 relies on the fact that the sum of the k largest eigenvalues of X_s is larger than the sum of the k largest eigenvalues of X for any integer k. As a result, Corollary 3 does not extend to a general state penalty matrix Q where the $\mathcal{H}_2$ norm is given by $\mathrm{trace}\,(QX)$. For the same reason, the result requires $\mathbf{E}\,(d(t_1)\,d^T(t_2)) = I\delta(t_1 - t_2)$.

A. Small asymmetric perturbations

We next show that in addition to being an upper bound, the $\mathcal{H}_2$ norm of the symmetric and full systems are close when A is nearly normal. In what follows, we show that when a normal system is subject to an anti-symmetric perturbation of $O(\epsilon)$, the first order correction to the $\mathcal{H}_2$ norm is zero. A similar result appeared in [20] for the design of an interconnection graph for synchronizing oscillator networks. We present a result for systems with normal dynamical generators A_n. Since a normal matrix A_n commutes with A_n^T, this class includes symmetric systems considered in [20].

Proposition 4: Let A_n be a normal Hurwitz matrix. The $O(\epsilon)$ correction to the $\mathcal{H}_2$ norm of the system

$$\dot{x} \ = \ A_n\,x \ + \ d$$

resulting from an $O(\epsilon)$ anti-symmetric perturbation A_a to A_n is zero.

Proof: The $\mathcal{H}_2$ norm of the above system is given by $\mathrm{trace}(X_n)$ where,

$$A_nX_n \ + \ X_nA_n^T \ + \ I \ = \ 0. \tag{7}$$

From Lemma 1 in [21], $X_n = -(A_n + A_n)^{-1}$. Perturbing A_n by a small anti-symmetric matrix ϵA_a yields a small correction term $\epsilon\tilde{X}$ in the controllability gramian. Collecting the $O(\epsilon)$ terms from the Lyapunov equation,

$$(A_n + \epsilon A_a)(X_n + \epsilon\tilde{X}) + (X_n + \epsilon\tilde{X})(A_n + \epsilon A_a)^T + I = 0.$$

recovers the linear equation,

$$A_n\tilde{X} \ + \ \tilde{X}A_n \ + \ A_aX_n \ + \ X_nA_a^T \ = \ 0.$$

Since $A_a = -A_a^T$, the $O(\epsilon)$ correction to the $\mathcal{H}_2$ norm vanishes,

$$\begin{aligned} \mathrm{trace}\,(\tilde{X}) &= \mathrm{trace}\left(X_n\left(A_aX_n \ + \ X_nA_a^T\right)\right) \\ &= \mathrm{trace}\left(\left(A_a \ - \ A_a\right)X_n^2\right) \qquad = \ 0 \end{aligned}$$

∎

IV. Combination drug therapy for HIV treatment

A. Background

In this section, we focus on the problem described in Section II-B.3,

$$\dot{x} \ = \ \left(A \ - \ \sum_{k=1}^{m} u_kD_k\right) x \ + \ d.$$

where x is a vector of HIV mutant populations, A contains their evolutionary dynamics, and D_k are diagonal matrices containing information about drug treatment. The entry A_{ii} determines how fast mutant i replicates, and the entry A_{ij} quantifies the probability of mutation from mutant j into mutant i.

Each element u_k of the control input u represents the amount of drug k administered to the patient. The ith element of the diagonal of D_k specifies how quickly drug k destroys mutant i.

Combination drug therapy is desirable because using only one drug often leads to the mutant population adapting to the weaknesses of that particular drug [22]. However, because of potential side effects and drug-drug interactions, it is not de-

sirable to use a large number of different drugs. Furthermore, large doses can have additional side effects [23].

Since the probability of mutation is often orders of magnitude less than the rate of replication [11], using a symmetric model here is justified. We are therefore interested in minimizing the $\mathcal{H}_2$ norm of the system,

$$\dot{x} \;=\; (A_s \;-\; \mathrm{diag}\{Du\})\,x \;+\; d.$$

where $D = [d_1 \cdots d_r]$, $d_i = \mathrm{diag}\{D_i\}$ and $A_s = \frac{1}{2}(A \;+\; A^T)$. We use the $\mathrm{diag}\{\cdot\}$ operator to denote either the diagonal entries of a matrix or a diagonal matrix with elements of the vector on its diagonal. Since this system is now symmetric, it fits into the framework of problem (4).

B. Primal and dual optimization problems

We first state the problem in a form which is convenient for implementation of the alternating direction method of multipliers (ADMM), a technique well-suited to large-scale problems. ADMM has been recently successfully applied to sparse control synthesis problems [8], [9]. Using the auxiliary variable $G := -A_s + \mathrm{diag}\{Du\}$, problem (4) becomes,

$$\begin{array}{ll}\text{minimize} & \mathrm{trace}\,(G^{-1}) \;+\; \frac{1}{2}\,u^T R u \;+\; \gamma \mathbb{1}^T u \\ \text{subject to} & G \;+\; A_s \;-\; \mathrm{diag}\{Du\} \;=\; 0 \\ & G \succ 0, \quad u \;\geq\; 0.\end{array} \tag{P}$$

Since a negative drug dosage is not possible, the ℓ_1 norm of u is determined by $\mathbb{1}^T u$. The Lagrangian is,

$$\begin{aligned}\mathcal{L}(G,u,Y,\lambda) \;=\; & \mathrm{trace}\,(G^{-1}) \;+\; \tfrac{1}{2}\,u^T R u \;+\; \gamma \mathbb{1}^T u \;- \\ & \lambda^T u \;+\; \langle Y, G + A_s - \mathrm{diag}\{Du\}\rangle\end{aligned}$$

where $\langle\cdot,\cdot\rangle$ denotes the standard inner product between two matrices, $Y = Y^T$, and $\lambda \geq 0$. We omit the Lagrange multiplier associated with the positive definiteness of G for brevity; this omitted dual variable will become a slack variable and it will result in a requirement on the positive definiteness of Y. We substitute the equivalent expression,

$$\mathrm{trace}\,(Y\ \mathrm{diag}\{Du\}) \;=\; y^T D u$$

where $y = \mathrm{diag}\{Y\}$ into the Lagrangian, differentiate it with respect to G and u and set the resulting gradients to zero,

$$\begin{aligned} 0 \;&=\; -G^{-2} + Y \\ 0 \;&=\; R\,u \;+\; \gamma\,\mathbb{1} \;-\; \lambda \;-\; D^T y.\end{aligned}$$

The optimal G and u are thus given by

$$\begin{aligned} G \;&=\; Y^{-1/2} \\ u \;&=\; R^{-1}(D^T y - \gamma v).\end{aligned}$$

Substituting these expressions into the Lagrangian yields

$$\begin{aligned}& 2\,\mathrm{trace}\,(Y^{\frac{1}{2}}) \;+\; \mathrm{trace}\,(A_s\,Y) \\ & -\; \tfrac{1}{2}\,(D^T y - \gamma\mathbb{1} + \lambda)^T R^{-1}(D^T y - \gamma\mathbb{1} + \lambda).\end{aligned}$$

The dual problem requires maximization of the above expression over $\lambda \geq 0$ and $Y \succ 0$. However, since λ must be positive, it is only nonnegative when $D^T y - \gamma\mathbb{1}$ is negative. Therefore, λ can be eliminated and the dual problem can be written as,

$$\begin{array}{ll}\text{maximize} & 2\,\mathrm{trace}\,(Y^{\frac{1}{2}}) \;+\; \mathrm{trace}\,(A\,Y) \\ & -\frac{1}{2}\ \max(D^T y - \gamma\mathbb{1}, 0)^T R^{-1}(D^T y - \gamma\mathbb{1}) \\ \text{subject to} & Y \succeq 0.\end{array} \tag{D}$$

C. Alternating direction method of multipliers

To apply ADMM to (P), we first form the corresponding augmented Lagrangian,

$$\begin{aligned}\mathcal{L}_\rho(G,u,Y) \;:=\; & \mathrm{trace}\,(G^{-1}) \;+\; u^T R u \;+\; \gamma\,\mathbb{1}^T u \\ & +\; \langle Y, G + A_s - \mathrm{diag}\{Du\}\rangle \\ & +\; \frac{\rho}{2}\,\|G + A_s - \mathrm{diag}\{Du\}\|_F^2.\end{aligned}$$

Relative to the standard Lagrangian, $\mathcal{L}_\rho$ contains an additional quadratic penalty on the violation of the linear constraint. The positive parameter ρ specifies the magnitude on the constraint violation penalty at each iteration.

The ADMM iteration uses the update sequence [24]

$$\begin{aligned} G_{k+1} \;&=\; \underset{G}{\mathrm{argmin}}\ \ \mathcal{L}_\rho(G,\,u_k,\,Y_k) \\ u_{k+1} \;&=\; \underset{u}{\mathrm{argmin}}\ \ \mathcal{L}_\rho(G_{k+1},\,u,\,Y_k) \\ Y_{k+1} \;&=\; Y_k \;+\; \rho\ (G_{k+1} + A - \mathrm{diag}\{Du_{k+1}\})\end{aligned}$$

to find the optimal solution to the original problem. The stopping criteria depend on the primal residual, which quantifies how well G_k and u_k satisfy the linear constraint, and the dual residual, which quantifies the difference between u_k and u_{k-1}. We refer the reader to [24] for details.

This algorithm is advantageous because the subproblems are much simpler than the original problem. The G-minimization step has an explicit solution, the u-minimization step takes a standard form for which there are efficient algorithms, and the Y-update step is algebraic.

1) G-minimization: The G-minimization step amounts to solving,

$$\begin{array}{ll}\underset{G}{\text{minimize}} & \mathrm{trace}\,(G^{-1}) \;+\; \dfrac{\rho}{2}\,\|G - V_k\|_F^2 \\ \text{subject to} & G \succ 0\end{array}$$

where $V_k := -A_s + \mathrm{diag}\{Du_k\} - \frac{1}{\rho}Y_k$. Seting the gradient,

$$-G^{-2} \;+\; \rho\,G \;-\; V_k$$

to zero is an explicit exercise with the positive definiteness constraint. Since the powers of G appear with no coefficients, the optimal G has the same eigenstructure as V_k. Its eigenvalues are determined by the real positive solution to the cubic equation

$$\rho\,\lambda_i^3 \;+\; \sigma_i\,\lambda_i^2 \;-\; 1$$

where σ_i is the corresponding eigenvalue of V_k. By the convexity of the G-minimization problem, there can be only one real positive solution to the above equation.

2) u-minimization: The u-minimization step amounts to solving,

$$\underset{u}{\text{minimize}} \ \gamma \mathbb{1}^T u + u^T R u + \frac{\rho}{2} \|Du - w_k\|_2^2.$$

where $w_k := g_{k+1} + a + (1/\rho)y_k$, $g_k = \text{diag}\{G_k\}$, $a = \text{diag}\{A\}$, and $y_k = \text{diag}\{Y_k\}$. The objective function is the sum of a quadratic term and an ℓ_1 norm: a problem form is commonly referred to as LASSO. This problem has attracted lots of attention in recent years and there are many efficient methods for computing its solution.

We employ a proximal gradient method known as Iterative Soft-Thresholding (ISTA) [25]. At each point $\bar{u}$, the smooth part of the objective function is linearized and a proximal term is added to obtain,

$$\text{minimize} \ \ \frac{\beta}{2} \|u - \bar{u}\|_F^2 + \langle f(\bar{u}), u \rangle + \gamma \|u\|_1$$

where, $f(\bar{u})$ is the gradient of the smooth part of (8),

$$f(\bar{u}) = (\rho D^T D + R)\bar{u} - D^T(y_k - \rho(g_k + a))$$

This approximation has an explicit minimizer; the update of the ith element of u is given by

$$u_i = \max(\bar{u}_i - f/\beta - \gamma/\beta, 0).$$

3) Computational complexity: The worst-case complexity of generic SDP solvers is $O(n^6)$, where n is the dimension of the positive definite constraint. In contrast, the G-minimization takes $O(n^3)$ operations because it requires an eigenvalue decomposition, the u-minimization step takes $O(nr)$ operations, and the Y-update step takes $O(n^2)$ operations.

V. An example

A. HIV example

Following the example given in [13], [14] based on [26], we examine a system with 35 mutants and 5 potential types of drugs. The diagonal entries are 0.5 and the off diagonal elements range from $O(10^{-8})$ to $O(10^{-6})$. The structure of the matrix A is shown in Fig. 1; clearly A is not symmetric.

We next use our algorithm to design control inputs u for the symmetric system with varying levels of the sparsity promoting parameter γ. As γ is increased, sparsity is prioritized over $\mathcal{H}_2$ performance and therefore the $\mathcal{H}_2$ performance degrades. In Fig. 2, we show the difference in $\mathcal{H}_2$ performance between the symmetric and original systems as a function of γ. In this problem, the symmetric model is a very good approximation of the original system, even up to extremely large levels of γ.

Since the off-diagonal entries of A are small, we artificially increase them by a constant factor to study our approach for systems with larger degrees of asymmetry. We take,

$$A_c = c(A - I \circ A) + 0.5(I \circ A)$$

where c is a constant factor and $\circ$ is the Hadamard (element wise) product. This modification means that A_c does not have physically relevant implications for the drug synthesis problem, but it illustrates the utility of our approach using a realistic problem structure. Figure 3 compares the $\mathcal{H}_2$ performance for $c = 10^5, 1.4\times10^6$, and 1.9×10^7. Compared to the diagonal entries of A_c, the maximum off-diagonal element is of the same order, one order of magnitude higher, and two orders of magnitude higher, respectively.

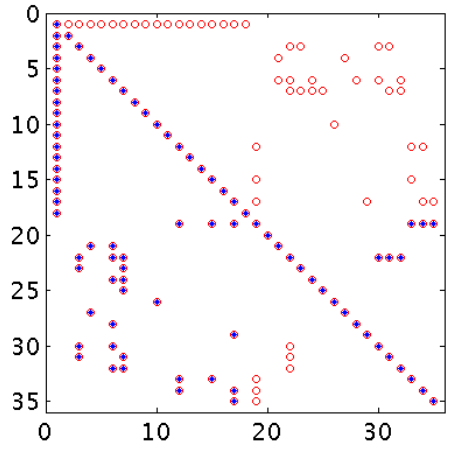

Fig. 1. Sparsity structure of the matrix A and its symmetric counterpart A_s. The elements of A are shown with blue dots, and the elements of A_s are shown with red circles.

When the off-diagonal elements are of the same order of magnitude or smaller than the diagonal elements, there is almost no difference between the symmetric and full models. As the off-diagonal elements get larger, the fidelity of the approximation suffers. Unsurprisingly, as the system becomes more asymmetric, the symmetric approximation becomes more conservative [27].

We note that for a realistic synthesis problem, γ would be varied to find sparsity structures for u. Once a desired sparsity structure is identified, (4) would be performed with $\gamma = 0$ but u constrained to have that particular sparsity structure. This process, known as polishing or de-biasing, yields u with a desired sparsity structure that provides better performance than u, with the same structure, that came from solving (4).

VI. Concluding remarks

We have introduced an approach to designing symmetric additions to dynamical systems for the purpose of optimizing the $\mathcal{H}_2$ norm of the closed-loop system. We showed that the symmetric system inherits many of the properties of the original system; in particular, that it is a convex upper bound on the original design problem. However, this approximation can be very conservative if A is highly non-normal. We show the utility of our approach on a HIV combination drug therapy treatment design example.

There are many exciting avenues for future development. In particular, we are exploring how to establish a bound on

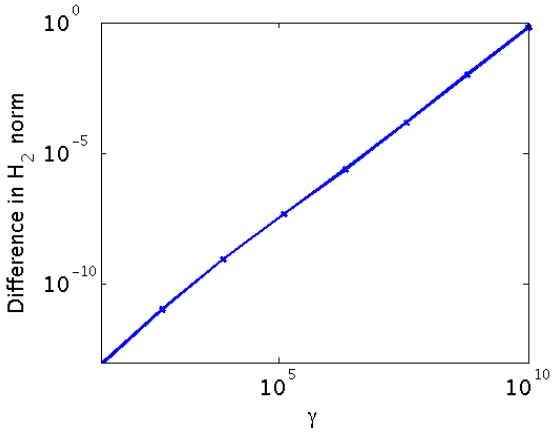

Fig. 2. Difference in $\mathcal{H}_2$ norm between the symmetric and original systems. Different controllers were designed as a function of γ, and the results were normalized using the $\mathcal{H}_2$ norm of the original system.

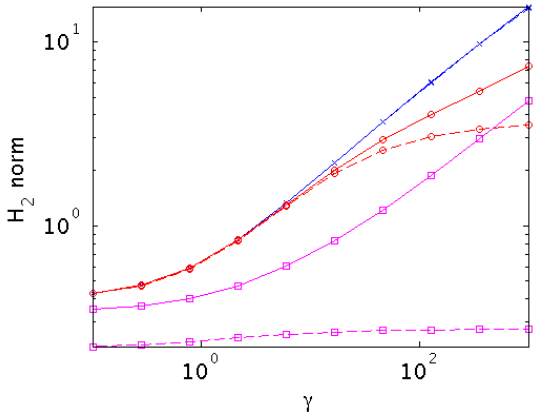

Fig. 3. The solid lines are the $\mathcal{H}_2$ norms of the symmetric systems and the dotted lines are the $\mathcal{H}_2$ norms of the original systems. The blue $\times$, red $\circ$, and magenta $\square$ designate $c = 10^5, 1.4 \times 10^6$, and 1.9×10^7, respectively.

the difference between the original and symmetric systems, especially for positive systems.

VII. Acknowledgments

We would like to thank Vanessa Jonsson for providing the data for the HIV drug therapy problem and Anders Rantzer for useful discussions.

References

[1] B. Bamieh and P. G. Voulgaris, "A convex characterization of distributed control problems in spatially invariant systems with communication constraints," *Systems & Control Letters*, vol. 54, no. 6, pp. 575–583, 2005.

[2] M. Rotkowitz and S. Lall, "A characterization of convex problems in decentralized control," *IEE Trans. Automat. Control*, vol. 51, no. 2, pp. 274–286, Feb 2006.

[3] A. Rantzer, "Distributed control of positive systems," 2012, available at arXiv:1203.0047.

[4] L. Xiao, S. Boyd, and S.-J. Kim, "Distributed average consensus with least-mean-square deviation," *J. Parallel Distrib. Comput.*, vol. 67, no. 1, pp. 33–46, 2007.

[5] F. Lin, M. Fardad, and M. R. Jovanović, "Identification of sparse communication graphs in consensus networks," in *Proceedings of the 50th Annual Allerton Conference on Communication, Control, and Computing*, Monticello, IL, 2012, pp. 85–89.

[6] M. Fardad, F. Lin, and M. R. Jovanović, "Design of optimal sparse interconnection graphs for synchronization of oscillator networks," *IEEE Trans. Automat. Control*, vol. 59, no. 9, pp. 2457–2462, 2014.

[7] B. Polyak, M. Khlebnikov, and P. Shcherbakov, "An LMI approach to structured sparse feedback design in linear control systems," in *Proceedings of the 2013 European Control Conference*, 2013, pp. 833–838.

[8] N. K. Dhingra, M. R. Jovanović, and Z. Q. Luo, "An ADMM algorithm for optimal sensor and actuator selection," in *Proceedings of the 53rd IEEE Conference on Decision and Control*, 2014, pp. 4039–4044.

[9] F. Lin, M. Fardad, and M. R. Jovanović, "Design of optimal sparse feedback gains via the alternating direction method of multipliers," *IEEE Trans. Automat. Control*, vol. 58, no. 9, pp. 2426–2431, 2013.

[10] N. Matni, "Communication delay co-design in H_2 decentralized control using atomic norm minimization," in *Proceedings of the 52nd IEEE Conference on Decision and Control*, 2013, pp. 6522–6529.

[11] D. I. Rosenbloom, A. L. Hill, S. A. Rabi, R. F. Siliciano, and M. A. Nowak, "Antiretroviral dynamics determines hiv evolution and predicts therapy outcome," *Nature medicine*, vol. 18, no. 9, pp. 1378–1385, 2012.

[12] A. Rantzer and B. Bernhardsson, "Control of convex monotone systems," in *Proceedings of the 53rd IEE Conference on Decision and Control*, 2014.

[13] V. Jonsson, A. Rantzer, and R. M. Murray, "A scalable formulation for engineering combination therapies for evolutionary dynamics of disease," *Proceedings of the 2014 American Control Conference*, 2014.

[14] V. Jonsson, N. Matni, and R. M. Murray, "Reverse engineering combination therapies for evolutionary dynamics of disease: An H_∞ approach," in *Proceedings of the 52nd IEEE Conference on Decision and Control*, 2013, pp. 2060–2065.

[15] F. Lin, M. Fardad, and M. R. Jovanović, "Algorithms for leader selection in stochastically forced consensus networks," *IEEE Trans. Automat. Control*, vol. 59, no. 7, pp. 1789–1802, July 2014.

[16] K. Fitch and N. E. Leonard, "Information centrality and optimal leader selection in noisy networks," in *Proceedings of the 52nd IEEE Conference on Decision and Control*, 2013, pp. 7510–7515.

[17] M. Fardad, F. Lin, and M. R. Jovanović, "On optimal link creation for facilitation of consensus in social networks," in *Proceedings of the 2014 American Control Conference*, 2014, pp. 3802–3807.

[18] A. Ghosh and S. Boyd, "Growing well-connected graphs," in *Proceedings of the 45th IEEE Conference on Decision and Control*, 2006, pp. 6605–6611.

[19] R. Bhatia, *Matrix analysis*. Springer, 1997, vol. 169.

[20] M. Fardad, F. Lin, and M. R. Jovanović, "Design of optimal sparse interconnection graphs for synchronization of oscillator networks," 2013, arXiv:1302.0449v1.

[21] B. Bamieh and M. Dahleh, "Exact computation of traces and H_2 norms for a class of infinite-dimensional problems," *IEEE Trans. Automat. Control*, vol. 48, no. 4, pp. 646–649, 2003.

[22] F. Michor, Y. Iwasa, and M. A. Nowak, "Dynamics of cancer progression," *Nature Reviews Cancer*, vol. 4, no. 3, pp. 197–205, 2004.

[23] B. Al-Lazikani, U. Banerji, and P. Workman, "Combinatorial drug therapy for cancer in the post-genomic era," *Nature biotechnology*, vol. 30, no. 7, pp. 679–692, 2012.

[24] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1–124, 2011.

[25] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM Journal on Imaging Sciences*, vol. 2, no. 1, pp. 183–202, 2009.

[26] 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.

[27] L. N. Trefethen and M. Embree, *Spectra and pseudospectra: the behavior of nonnormal matrices and operators*. Princeton University Press, 2005.