Scalable system level synthesis for virtually localizable systems

Nikolai Matni, Yuh-Shyang Wang and James Anderson

Abstract— **In previous work, we developed the system level approach to controller synthesis, and showed that under suitable assumptions, this framework allowed for the synthesis of localized controllers. We further showed that such localized controllers enjoy $O(1)$ synthesis and implementation complexity relative to the dimension of the global system, making them particularly well suited for the control of large-scale cyber-physical systems. However, the assumptions under which a system is localizable are stringent: roughly, a system is localizable if the controller has the necessary actuation, sensing and communication resources to "get out ahead" of the propagation of a disturbance and neutralize it, thus containing its effect to a localized spatiotemporal region. In this paper, we relax the assumption of exact localizability, and develop a controller synthesis methodology that is applicable to arbitrary systems that are in an appropriate sense "easy to control." We focus on the state-feedback setting and develop a simple necessary and sufficient condition for robust stability using the system level approach. We then leverage this condition, along with the introduction of virtual actuation, communication and system responses into the synthesis process, to design stabilizing controllers that have (i) $O(1)$ synthesis and implementation complexity and (ii) guaranteed performance bounds. We end with a power-inspired example demonstrating the usefulness of these techniques, wherein we synthesize a near globally optimal controller for a system that is neither localizable nor quadratically invariant.**

PRELIMINARIES & NOTATION

We use lower and upper case Latin letters such as x and A to denote vectors and matrices, respectively, and lower and upper case boldface Latin letters such as $\mathbf{x}$ and $\mathbf{G}$ to denote signals and transfer matrices, respectively. Calligraphic letters such as $\mathcal{S}$ denote sets. We work with discrete time, linear time invariant systems. We use standard definitions of the Hardy spaces $\mathcal{H}_2$ and $\mathcal{H}_\infty$, and denote their restriction to the set of real-rational proper transfer matrices by $\mathcal{RH}_2$ and $\mathcal{RH}_\infty$. We use $G[i]$ to denote the ith spectral component of a transfer function $\mathbf{G}$, i.e., $\mathbf{G}(z) = \sum_{i=0}^{\infty} \frac{1}{z^i} G[i]$ for $|z| > 1$. Finally, $\mathcal{F}_T$ denotes the space of finite impulse response (FIR) transfer matrices with horizon T, i.e., $\mathcal{F}_T := \{\mathbf{G} \in \mathcal{RH}_\infty \,|\, \mathbf{G} = \sum_{i=0}^{T} \frac{1}{z^i} G[i]\}$. Let $\mathbb{Z}^+$ be the set of all positive integers. We use calligraphic lower case letters such as $\mathcal{r}$ and $\mathcal{c}$ to denote subsets of $\mathbb{Z}^+$. Consider a transfer matrix $\mathbf{\Phi}$ with n_r rows and n_c columns. Let $\mathcal{r}$ be a subset of $\{1, \ldots, n_r\}$ and $\mathcal{c}$ a subset of $\{1, \ldots, n_c\}$. We use $\mathbf{\Phi}(\mathcal{r}, \mathcal{c})$ to denote the submatrix of $\mathbf{\Phi}$ by selecting the rows according to the set $\mathcal{r}$ and columns according to the set $\mathcal{c}$. Finally, the symbol : is used to denote the set of all rows or all columns, i.e., we have $\mathbf{\Phi} = \mathbf{\Phi}(:,:)$. Let $\{\mathcal{c}_1, \ldots \mathcal{c}_p\}$ be a partition of the set $\{1, \ldots, n_c\}$. Then $\{\mathbf{\Phi}(:, \mathcal{c}_1), \ldots, \mathbf{\Phi}(:, \mathcal{c}_p)\}$ is a column-wise partition of the transfer matrix $\mathbf{\Phi}$.

N. Matni and J. Anderson are with the Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA. Y.-S. Wang is with GE Research. This research was supported by grants from the AFOSR and NSF, and by gifts from Huawei and Google. nmatni@caltech.edu, yuhshyangwang@gmail.com, james@caltech.edu.

I. INTRODUCTION

Modern cyber-physical systems (CPS) are large-scale, physically distributed and interconnected – these systems are composed of a multitude of sub-controllers, each equipped with their own sensors and actuators. These sub-controllers then exchange their local information with each other via a communication network, subject to the delay, bandwidth and reliability properties of this network. The resulting information asymmetry among the sub-controllers is the fundamental challenge that must be overcome when synthesizing distributed optimal controllers. [1]–[6].

In recent work, we introduced the System Level Approach (SLA) to controller synthesis [7], [8], and showed that it characterized the broadest known class of structured (e.g., distributed) optimal control problems that admit a convex characterization, generalizing quadratic invariance (QI) [3]. A driving motivation behind this body of work was to address the issue of scalable distributed controller synthesis and implementation – in particular, the SLA allows for the design of localized controllers [9], [10] using convex programming that enjoy $O(1)$ synthesis and implementation complexity relative to the size of the full system.

Roughly speaking, a system is localizable if the effect of each disturbance can be isolated to a localized region of the global system and eliminated in finite time – it can thus be thought of as a generalization of deadbeat control to the spatiotemporal domain. For the effect of a disturbance to be localized, there must be sufficient communication, sensing and actuation resources so that the controller can detect and "get out ahead" of a disturbance as its effects propagate through the physical system. In practice this condition can be quite restrictive, expensive to achieve, and is in general stricter than the conditions imposed by QI[1] as it requires communication between sub-controllers to be faster than the propagation of disturbances through the physical system.

In this paper we address this issue by extending our previous results to systems that are in an appropriate sense "nearly localizable." We focus on the state-feedback setting, and begin by developing necessary and sufficient conditions under which a controller synthesized using the SLA robustly stabilizes a family of physical plants. We then introduce

[1] Under mild assumptions, a system is QI if sub-controllers can communicate with each other faster than their control actions can be felt by other sub-controllers. [11].

virtual actuation, communication and system responses into the synthesis procedure to ensure that the augmented system is indeed localizable, allowing us to design a controller that enjoys the computational benefits of locality, i.e., that has $O(1)$ synthesis and implementation complexity. By viewing these virtual components as perturbations to a nominal system, we then appeal to our robustness result to guarantee that the resulting controller is globally stabilizing, and further provide sub-optimality bounds with respect to the performance achievable by a centralized optimal controller.

Related work: There exists a rich body of work addressing the identification of tractable information exchange constraints that lead to convex distributed optimal control problems, cf. [1]–[7]. These results, and in particular QI, were then used as a starting point to pose and solve distributed optimal controller synthesis problems under sparsity and delay constraints for $\mathcal{H}_2$ (e.g., [12]–[17]) and $\mathcal{H}_\infty$ (e.g., [18]–[20]) performance metrics. However, as all of these results required QI (or QI-like) properties to hold, their complexity did not scale well with the size of the global system. This lack of scalability did not go unnoticed by the community, and techniques based on regularization [21], convex approximation [22], [23], and spatial truncation [24] were used in hopes of finding a near optimal distributed static feedback controller that are scalable to implement. These methods were successful in extending the size of systems for which a distributed controller could be computed, but they are still limited in their scalability as they rely on an underlying centralized synthesis procedure.

Paper organization: We begin by reviewing the relevant SLA results from [7] for state-feedback synthesis in Section II. In Section III we develop necessary and sufficient conditions under which a controller synthesized using the SLA robustly stabilizes a family of physical systems. In Section IV, we show how the introduction of virtual actuation and communication can be combined with the robustness result developed in the previous section, as well as the notion of column-wise separable System Level Synthesis (SLS) problems [8] to parameterize a family of stabilizing controllers that can be synthesized and implemented in a localized manner. This parameterization is then used to pose a corresponding optimal control problem in Section V, and show that it can further be used to bound the sub-optimality of the computed controller relative to the performance achievable by a centralized optimal controller. We then demonstrate the usefulness of these ideas on numerical examples in Section VI, and end with conclusions and directions for future work in Section VII.

II. The System Level Approach

We consider linear time invariant (LTI) systems in discrete time of the form

$$x[t+1] = Ax[t] + B_1 w[t] + B_2 u[t] \tag{1a}$$
$$\bar{z}[t] = C_1 x[t] + D_{12} u[t] \tag{1b}$$

where x, u, w, and $\bar{z}$ are the state vector, control action, external disturbance, and regulated output, respectively. In this paper, we focus on the state-feedback setting wherein $\mathbf{u} = \mathbf{K}\mathbf{x}$ for a possibly dynamic feedback gain $\mathbf{K}$.

In recent work [7], [8], [25], we defined and analyzed the System Level Approach (SLA) to controller synthesis, which provides a novel parameterization of structured stabilizing controllers. We recall the relevant state-feedback results that we build upon here.

The z-transform of the state dynamics (1a) is given by

$$(zI - A)\mathbf{x} = B_2\mathbf{u} + \boldsymbol{\delta_x}, \tag{2}$$

where $\boldsymbol{\delta_x} := B_1\mathbf{w}$ denotes the disturbance affecting the state.

We then define $\mathbf{R}$ to be the system response mapping the disturbance $\boldsymbol{\delta_x}$ to the state $\mathbf{x}$, and $\mathbf{M}$ to be the system response mapping the disturbance $\boldsymbol{\delta_x}$ to the control action $\mathbf{u}$. By substituting a dynamic state feedback control rule $\mathbf{u} = \mathbf{K}\mathbf{x}$ into (2), the system response $\{\mathbf{R}, \mathbf{M}\}$ as a function of the controller $\mathbf{K}$ as can be written as

$$\begin{array}{rcl} \mathbf{R} & = & (zI - A - B_2\mathbf{K})^{-1} \\ \mathbf{M} & = & \mathbf{K}(zI - A - B_2\mathbf{K})^{-1}. \end{array} \tag{3}$$

The following key theorem from [7] provides an algebraic characterization of the set $\{\mathbf{R}, \mathbf{M}\}$ of state-feedback system responses that are achievable by an internally stabilizing controller $\mathbf{K}$.

Theorem 1: For the dynamics (1) with state-feedback control rule $\mathbf{u} = \mathbf{K}\mathbf{x}$, the following are true:

(a) The affine subspace defined by

$$\begin{bmatrix} zI - A & -B_2 \end{bmatrix} \begin{bmatrix} \mathbf{R} \\ \mathbf{M} \end{bmatrix} = I \tag{4a}$$
$$\mathbf{R}, \mathbf{M} \in \frac{1}{z}\mathcal{RH}_\infty \tag{4b}$$

parameterizes all system responses from $\boldsymbol{\delta_x}$ to $(\mathbf{x}, \mathbf{u})$, as defined in (3), achievable by an internally stabilizing state feedback controller $\mathbf{K}$.

(b) For any transfer matrices $\{\mathbf{R}, \mathbf{M}\}$ satisfying (4), the controller $\mathbf{K} = \mathbf{M}\mathbf{R}^{-1}$ is internally stabilizing and achieves the desired system response (3).

We further showed that the controller $\mathbf{K} = \mathbf{M}\mathbf{R}^{-1}$ can be implemented as,

$$\begin{array}{rcl} \hat{\boldsymbol{\delta}}_{\boldsymbol{x}} & = & \mathbf{x} - \hat{\mathbf{x}} \\ \mathbf{u} & = & z\mathbf{M}\hat{\boldsymbol{\delta}}_{\boldsymbol{x}} \\ \hat{\mathbf{x}} & = & (z\mathbf{R} - I)\hat{\boldsymbol{\delta}}_{\boldsymbol{x}}. \end{array} \tag{5}$$

as illustrated in Figure 1. An important feature of this controller parameterization and implementation is that if the system responses $\{\mathbf{R}, \mathbf{M}\}$ are structured, then so is the controller implementation defined in terms of the transfer matrices $\tilde{\mathbf{R}}$ and $\tilde{\mathbf{M}}$ (see Figure 1 caption for definitions).

Theorem 1 and the corresponding controller implementation shown Figure 1 allow us to pose the controller synthesis task as finding the solution to the following convex System

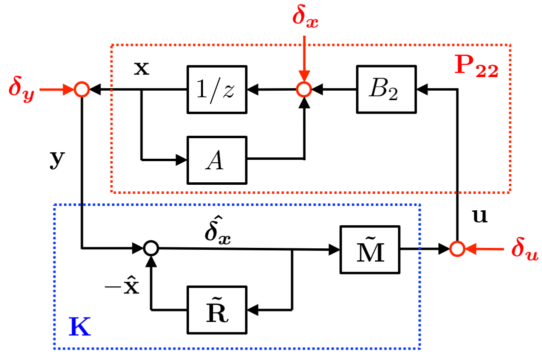

Fig. 1: The proposed state feedback controller structure, with $\tilde{\mathbf{R}} = I - z\mathbf{R}$ and $\tilde{\mathbf{M}} = z\mathbf{M}$.

Level Synthesis (SLS) problem:

$$\begin{aligned} \underset{\{\mathbf{R},\mathbf{M}\}}{\text{minimize}} \quad & g(\mathbf{R}, \mathbf{M}) \\ \text{subject to} \quad & (4a) - (4b) \\ & \begin{bmatrix} \mathbf{R} \\ \mathbf{M} \end{bmatrix} \in \mathcal{S}. \end{aligned} \tag{6}$$

where g is a suitably chosen convex cost functional, and $\mathcal{S}$ is a convex set encoding System Level Constraints (SLCs) that can be used to enforce various properties on the system response $\{\mathbf{R}, \mathbf{M}\}$ and the corresponding controller implementation (cf., §IV [7] for a catalog of useful SLCs).

A. Locality SLCs and scalable synthesis

Of particular interest are SLCs that impose subspace constraints that define transfer matrices of sparse spatiotemporal support, i.e., SLCs of the form $\mathcal{S} = \mathcal{L} \cap \mathcal{F}_T$, where $\mathcal{L}$ enforces sparse spatial support, and $\mathcal{F}_T$ enforces sparse temporal support (FIR of horizon T). We say that a system is localizable if there exist system responses $\{\mathbf{R}, \mathbf{M}\}$ that simultaneously satisfy constraints (4a)-(4b) and a SLC $\mathcal{S}$ that enforces sparse spatiotemporal support.

An immediate benefit of enforcing such sparsity constraints on the system responses $\{\mathbf{R}, \mathbf{M}\}$ is that implementing the resulting controller (5) can be done in a *localized way* using FIR filter banks, i.e., each controller state $\hat{x}_i, (\hat{\delta}_x)_i$ and control action u_i can be computed using a local subset (as defined by the support of the system response and the FIR horizon T) of the global controller disturbance estimate history $\hat{\delta}_x$. For this reason, we refer to such constraints as a *localized* SLC when it defines a subspace with sparse support. As we discuss in detail in [8] and briefly recall in §IV-A, when such localized constraints are combined with objective functions that satisfy certain separability properties, this also allows for the resulting system response and controller to be synthesized in a localized way, i.e., the global computation decomposes naturally into decoupled subproblems that depend only on local decision variables and local sub-matrices of the state-space representation (1).

As we argued in [7], the feasibility of such locality constraints can be viewed as a generalization of controllability to localized spatiotemporal regions of the system. This condition can place rather stringent demands on the controller architecture: in particular this implies that for each disturbance perturbing the system, there must be sufficient locally available communication and actuation such that the effect of a disturbance on the system can be contained to a localized spatial region, and can further be completely eliminated with T time-steps *within* this region – this clearly is not be possible for sparsely actuated systems or for controllers with communication speed that is no faster than that of the speed of disturbance propagation through the physical plant.

This paper addresses the problem of synthesizing localized controllers, i.e., controllers that have $O(1)$ synthesis and implementation complexity relative to the size of the full system, *for systems that are not exactly localizable*. What we show is that if the system is appropriately easy to control (to be formalized in §IV), then the benefits of locality can still be exploited, allowing for the scalable synthesis of stabilizing controllers with performance guarantees.

III. A ROBUSTNESS RESULT

We begin with a robustness result that provides necessary and sufficient conditions under which a controller (5), implemented using transfer matrices that only approximately satisfy the achievability constraint (4), is still stabilizing. In §IV, we leverage this result and the use of virtual actuation and communication resources at design time to define a localized controller synthesis algorithm for systems that are not localizable.

Theorem 2: Let $(\mathbf{R}_c, \mathbf{M}_c, \mathbf{\Delta})$ be a solution to

$$\begin{bmatrix} zI - A & -B_2 \end{bmatrix} \begin{bmatrix} \mathbf{R}_c \\ \mathbf{M}_c \end{bmatrix} = I + \mathbf{\Delta} \tag{7}$$

Then, the controller implementation

$$\hat{\boldsymbol{\delta}}_{\boldsymbol{x}} = \mathbf{x} - \hat{\mathbf{x}} \tag{8a}$$
$$\mathbf{u} = z\mathbf{M}_c \hat{\boldsymbol{\delta}}_{\boldsymbol{x}} \tag{8b}$$
$$\hat{\mathbf{x}} = (z\mathbf{R}_c - I)\hat{\boldsymbol{\delta}}_{\boldsymbol{x}}. \tag{8c}$$

internally stabilizes the system (A, B_2) if and only if $(I + \mathbf{\Delta})^{-1}$ is stable.

Proof: For a state-feedback system (1) with controller implementation (5) using transfer matrices $\{\mathbf{R}_c, \mathbf{M}_c\}$, the following holds

$$(zI - A)\mathbf{x} = B_2\mathbf{u} + \boldsymbol{\delta_x} \tag{9a}$$
$$\mathbf{u} = z\mathbf{M}_c\hat{\boldsymbol{\delta}}_{\boldsymbol{x}} + \boldsymbol{\delta_u} \tag{9b}$$
$$\mathbf{x} = z\mathbf{R}_c\hat{\boldsymbol{\delta}}_{\boldsymbol{x}} - \boldsymbol{\delta_y}., \tag{9c}$$

where $\boldsymbol{\delta_y}$ and $\boldsymbol{\delta_u}$ are respectively perturbations on the measurement and control action (as illustrated in Figure 1) introduced to verify the internal stability of the resulting closed loop system.

Substituting (9b) and (9c) into (9a), we have

$$z(zI - A)\mathbf{R}_c\hat{\boldsymbol{\delta}}_{\boldsymbol{x}} - (zI - A)\boldsymbol{\delta_y} = zB_2\mathbf{M}_c\hat{\mathbf{w}} + B_2\boldsymbol{\delta_u} + \boldsymbol{\delta_x}.$$

Moving $\hat{\mathbf{w}}$ to the left-hand-side and using the relation (7) from the assumption, we obtain

$$z(I+\mathbf{\Delta})\hat{\boldsymbol{\delta}}_{\boldsymbol{x}} = (zI-A)\boldsymbol{\delta_y} + B_2\boldsymbol{\delta_u} + \boldsymbol{\delta_x}.$$

Denote $\boldsymbol{I_\Delta} = (I+\mathbf{\Delta})^{-1}$. The closed loop transfer matrices from $(\boldsymbol{\delta_x},\ \boldsymbol{\delta_y},\ \boldsymbol{\delta_u})$ to $\hat{\boldsymbol{\delta}}_{\boldsymbol{x}}$ are now given by

$$\hat{\boldsymbol{\delta}}_{\boldsymbol{x}} = \frac{1}{z}\boldsymbol{I_\Delta}\boldsymbol{\delta_x} + \boldsymbol{I_\Delta}(I-\frac{1}{z}A)\boldsymbol{\delta_y} + \frac{1}{z}\boldsymbol{I_\Delta}B_2\boldsymbol{\delta_u}. \tag{10}$$

Substituting (10) into (9b) and (9c), we have the closed loop transfer matrices from $(\boldsymbol{\delta_x},\ \boldsymbol{\delta_y},\ \boldsymbol{\delta_u})$ to $(\mathbf{x},\ \mathbf{u},\ \hat{\boldsymbol{\delta}}_{\boldsymbol{x}})$ summarized in Table I. Clearly, if $\boldsymbol{I_\Delta}$ is stable, then all the transfer matrices in Table I are stable. If $\boldsymbol{I_\Delta}$ is unstable, then the closed loop maps from $\boldsymbol{\delta_x}$ to $\hat{\boldsymbol{\delta}}_{\boldsymbol{x}}$ will be unstable, and the controller does not internally stabilize the system. Therefore, the stability of $\boldsymbol{I_\Delta} = (I+\mathbf{\Delta})^{-1}$ is necessary and sufficient condition for the controller implementation (8) to internally stabilize the system (A, B_2). ∎

TABLE I: Closed Loop Maps With Non-localizability

	$\boldsymbol{\delta_x}$	$\boldsymbol{\delta_y}$	$\boldsymbol{\delta_u}$
$\mathbf{x}$	$\mathbf{R}_c\boldsymbol{I_\Delta}$	$\mathbf{R}_c\boldsymbol{I_\Delta}(zI-A)-I$	$\mathbf{R}_c\boldsymbol{I_\Delta}B_2$
$\mathbf{u}$	$\mathbf{M}_c\boldsymbol{I_\Delta}$	$\mathbf{M}_c\boldsymbol{I_\Delta}(zI-A)$	$I+\mathbf{M}_c\boldsymbol{I_\Delta}B_2$
$\hat{\boldsymbol{\delta}}_{\boldsymbol{x}}$	$\frac{1}{z}\boldsymbol{I_\Delta}$	$\boldsymbol{I_\Delta}(I-\frac{1}{z}A)$	$\frac{1}{z}\boldsymbol{I_\Delta}B_2$

Theorem 2 can now be combined with small gain theorems to provide simple sufficient conditions for robust stability.

Corollary 1 ($\mathcal{H}_\infty$ robustness): Under the conditions of Theorem 2, the closed loop system is stable if $\|\mathbf{\Delta}\|_{\mathcal{H}_\infty} < 1$.

Proof: Classical, see [26] for example. ∎

Corollary 2 ($\mathcal{L}_1$ robustness): Under the conditions of Theorem 2, the closed loop system is stable if $\|\mathbf{\Delta}\|_{\mathcal{L}_1} < 1$.

Proof: Classical, see [27] for example. ∎

Corollary 3 ($\mathcal{E}_1$ robustness): Under the conditions of Theorem 2, the closed loop system is stable if $\|\mathbf{\Delta}\|_{\mathcal{E}_1} := \|\mathbf{\Delta}^\top\|_{\mathcal{L}_1} < 1$.

Proof: The transfer matrix $(I+\mathbf{\Delta})^{-1}$ is stable if and only if $(I+\mathbf{\Delta})^{-\top} = (I+\mathbf{\Delta}^\top)^{-1}$ is stable. The result then follows from Corollary 2. ∎

Remark 1: Note that $\|\mathbf{\Delta}\|_{\mathcal{L}_1}$ and $\|\mathbf{\Delta}\|_{\mathcal{E}_1}$ are the worst case $\ell_\infty \to \ell_\infty$ and $\ell_1 \to \ell_1$ gains of $\mathbf{\Delta}$, respectively.

It therefore follows that if a set of transfer matrices $\{\mathbf{R}_c, \mathbf{M}_c, \mathbf{\Delta}\}$ satisfy the following constraints:

$$\begin{gathered}\begin{bmatrix} zI-A & -B_2 \end{bmatrix}\begin{bmatrix}\mathbf{R}_c \\ \mathbf{M}_c\end{bmatrix} = I+\mathbf{\Delta} \\ \begin{bmatrix}\mathbf{R}_c \\ \mathbf{M}_c\end{bmatrix} \in \mathcal{L}\cap\mathcal{F}_T\cap\mathcal{X},\ \|\mathbf{\Delta}\|_\bullet < 1,\end{gathered} \tag{11}$$

for $\bullet \in \{\mathcal{H}_\infty, \mathcal{L}_1, \mathcal{E}_1\}$, then a controller (5) implemented using the transfer matrices $\{\mathbf{R}_c, \mathbf{M}_c\}$ is globally stabilizing for the system dynamics (1).

IV. Virtually Localizable Systems

In this section, we leverage the previously developed robustness results to formulate a localized synthesis procedure that generates stabilizing controllers for systems that are not localizable. We begin with a brief review of relevant results from [8] on column-wise separable constraints and objective functions that allow the global synthesis task to be decomposed into local subproblems that can be solved in parallel.

A. Column-wise separable SLS problems

We recall here the notion of a SLS problem (6) being column-wise separable, as defined in [8]. For compactness of notation, we use

$$\mathbf{\Phi} = \begin{bmatrix}\mathbf{R} \\ \mathbf{M}\end{bmatrix}$$

to represent the system response we want to optimize for, and we denote by $\mathbf{Z}_{AB}$ the transfer matrix $\begin{bmatrix} zI-A & -B_2\end{bmatrix}$.

Definition 1: The functional objective $g(\mathbf{\Phi})$ in (6) is *column-wise separable* with respect to the column-wise partition $\{c_1, \ldots, c_p\}$ if

$$g(\mathbf{\Phi}) = \sum_{j=1}^{p} g_j(\mathbf{\Phi}(:, c_j)) \tag{12}$$

for some functionals $g_j(\cdot)$ for $j = 1, \ldots, p$.

Definition 2: The set constraint $\mathcal{S}$ in (6) is *column-wise separable* with respect to the column-wise partition $\{c_1, \ldots, c_p\}$ if the following condition is satisfied:

$$\mathbf{\Phi} \in \mathcal{S} \quad \text{if and only if} \quad \mathbf{\Phi}(:, c_j) \in \mathcal{S}_j \text{ for } j = 1, \ldots, p \tag{13}$$

for some sets $\mathcal{S}_j$ for $j = 1, \ldots, p$.

In [8] we show that locality, FIR SLCs, as well as the achievability constraint (4) are all column-wise separable with respect to arbitrary column-wise partitions. In particular, if $g(\cdot)$ and $\mathcal{X}$ are respectively a column-wise separable objective function and SLC with respect to a column-wise partition, then the resulting SLS problem is also column-wise separable with respect to that same partition.

We can then specialize the state feedback SLS problem (6) to a localized one that can be written as

$$\begin{aligned} \underset{\mathbf{\Phi}}{\text{minimize}} \quad & g(\mathbf{\Phi}) && \text{(14a)} \\ \text{subject to} \quad & \mathbf{Z}_{AB}\mathbf{\Phi} = I && \text{(14b)} \\ & \mathbf{\Phi} \in \mathcal{L}\cap\mathcal{F}_T\cap\mathcal{X}. && \text{(14c)} \end{aligned}$$

Recall that the subspace $\mathcal{L}$ enforces spatial sparsity, the FIR constraint $\mathcal{F}_T$ enforces that the system responses be FIR transfer matrices of horizon T, and that the additional SLC $\mathcal{X}$ can be used to encode other constraints such as information exchange delays between sub-controllers. Note that the affine subspace constraint (14b), the locality SLC $\mathcal{L}$ and the FIR SLC $\mathcal{F}_T$ are column-wise separable with respect to any column-wise partition, and thus the overall problem is column-wise separable with respect to a given column-wise partition if the objective function g and SLC $\mathcal{X}$ are also column-wise separable with respect to that partition.

Assume that the the objective function g and SLC $\mathcal{X}$ are both column-wise separable with respect to a column-wise partition $\{c_1, \ldots, c_p\}$. In this case, we have that the state feedback localized SLS problem (14) is a *column-wise separable SLS problem*. Specifically, (14) can be partitioned into p parallel subproblems as

$$\begin{aligned} \underset{\mathbf{\Phi}(:,c_j)}{\text{minimize}} \quad & g_j(\mathbf{\Phi}(:,c_j)) && \text{(15a)}\\ \text{subject to} \quad & \mathbf{Z}_{AB}\mathbf{\Phi}(:,c_j) = I(:,c_j) && \text{(15b)}\\ & \mathbf{\Phi}(:,c_j) \in \mathcal{L}(:,c_j) \cap \mathcal{F}_T \cap \mathcal{X}_j && \text{(15c)} \end{aligned}$$

for $j = 1, \ldots, p$.

In this case, not only can each subproblem be solved independently and in parallel, but as we show in [8], the locality SLC $\mathcal{L}$ allows us to reduce the size of each subproblem's optimization variable and the corresponding system model from global to local scale, as defined by the support of the corresponding column-wise constraint $\mathcal{L}(:,c_j)$. It is this decomposition and complexity reduction that we seek to preserve in the following when locality constraints are not exactly feasible.

Remark 2: In [8] we introduce the notion of partially separable SLS problems, which are a substantial generalization of column-wise separable problems. Such problems also enjoy favorable computational properties, in that they can be solved using distributed optimization methods such the alternating direction method of multipliers (ADMM) wherein each iterate update subproblem can be decomposed column(row)-wise and solved with $O(1)$ computational complexity. Although we focus on column-wise separable problems in what follows, the results generalize to partially separable SLS problems in a natural way.

B. Virtually localizable systems

In particular, we seek a computational method with $O(1)$ complexity relative to the size of the global system for finding transfer matrices $\mathbf{R}_c, \mathbf{M}_c, \mathbf{\Delta}$ that satisfy

$$\begin{gathered} \begin{bmatrix} zI - A & -B_2 \end{bmatrix} \begin{bmatrix} \mathbf{R}_c \\ \mathbf{M}_c \end{bmatrix} = I + \mathbf{\Delta} \\ \begin{bmatrix} \mathbf{R}_c \\ \mathbf{M}_c \end{bmatrix} \in \mathcal{L} \cap \mathcal{F}_T \cap \mathcal{X}, \ \|\mathbf{\Delta}\|_{\mathcal{E}_1} < 1, \end{gathered} \tag{16}$$

where the SLCs $\mathcal{L}$, $\mathcal{F}_T$ and $\mathcal{X}$ are as described above. For convenience, we assume that the additional SLC $\mathcal{X}$ is column-wise separable with respect to an arbitrary column-wise partition – the discussion extends to general column-wise partitions at the expense of more cumbersome notation.

It therefore follows that the convex set defined by the constraints (16) is itself column-wise decomposable with respect to an arbitrary column-wise partition – this is because of our choice of the use of the $\|\cdot\|_{\mathcal{E}_1}$ norm to enforce the robust stability condition defined in Theorem 2. What is not apparent yet is if the resulting subproblems enjoy the same dimensionality reduction properties of localized problems, as in general, the error term $\mathbf{\Delta}$ is a dense transfer matrix that captures the non-localizable terms of the closed loop dynamics. In what follows, we suggest two complementary approaches to circumventing this issue.

1) Virtual actuation: We first propose introducing a virtual centralized full-actuation controller during the synthesis procedure. Such a virtual controller ensures that any spatiotemporal constraint of the form $\mathcal{L} \cap \mathcal{F}_T$ that is feasible under centralized full-actuation control is feasible in our framework as well, thus allowing for a localized synthesis procedure and controller implementation – further, if the effects of this virtual controller on the closed loop are suitably bounded, we may then appeal to Theorem 2 to guarantee the closed loop stability of the system.

Specifically, we synthesize a localized controller for a system described by the following augmented dynamics:

$$\begin{aligned} x[t+1] &= Ax[t] + B_1w[t] + B_2u[t] + v[t] && \text{(17a)}\\ \bar{z}[t] &= C_1x[t] + D_{12}u[t], && \text{(17b)} \end{aligned}$$

where all variables and matrices are as in equation (1), and v is the previously described set of virtual actuators. The corresponding SLP of achievable closed loop responses is then given by

$$\begin{bmatrix} zI - A & -B_2 \end{bmatrix} \begin{bmatrix} \mathbf{R} \\ \mathbf{M} \end{bmatrix} = I + \mathbf{V}, \tag{18}$$

where now $\mathbf{V}$ is the closed loop map from the process noise $\boldsymbol{\delta}_x$ to the virtual control action $\mathbf{v}$. If we further impose that $\mathbf{V}$ be localized and FIR, then we can combine the SLP of the augmented system (18), with the parameterization of robustly stabilizing approximate system responses (16) to obtain the following description of a set of robustly stabilizing controllers for a given system (1):

$$\begin{gathered} \begin{bmatrix} zI - A & -B_2 \end{bmatrix} \begin{bmatrix} \mathbf{R}_c \\ \mathbf{M}_c \end{bmatrix} = I + \boldsymbol{V} \\ \begin{bmatrix} \mathbf{R}_c \\ \mathbf{M}_c \end{bmatrix} \in \mathcal{L} \cap \mathcal{F}_T \cap \mathcal{X}, \\ \boldsymbol{V} \in \mathcal{L}' \cap \mathcal{F}_{T'}, \ \|\boldsymbol{V}\|_{\mathcal{E}_1} < 1. \end{gathered} \tag{19}$$

If the norm bound on the virtual system response $\boldsymbol{V}$ is removed, then the set described by the constraints (19) can always be made to be non-empty by selecting the virtual spatiotemporal SLC $\mathcal{L}' \cap \mathcal{F}_{T'}$ to be appropriately large. It therefore follows that if the corresponding augmented dynamics are suitably "easy to control," i.e., can be localized using only a small amount of virtual control authority, then the set of stabilizing localized controllers described by (19) is non-empty as well.

2) Leaky boundaries: We start with the SLP (18), and further decompose the state system response $\mathbf{R}$ as the sum of two components: a localized response $\mathbf{R}_c$ and a "leak" term $\mathbf{R}_l$ that can be used to capture the non-localizable component of the state-response $\mathbf{R}$. We can then appropriately modify the parameterization (19) to account for the contribution of this leak term, and write

$$
\begin{aligned}
&\begin{bmatrix} zI - A & -B_2 \end{bmatrix} \begin{bmatrix} \mathbf{R}_c \\ \mathbf{M}_c \end{bmatrix} = I + \boldsymbol{V} - (zI - A)\mathbf{R}_l \\
&\begin{bmatrix} \mathbf{R}_c \\ \mathbf{M}_c \end{bmatrix} \in \mathcal{L} \cap \mathcal{F}_\mathcal{T} \cap \mathcal{X}, \\
&\begin{bmatrix} \mathbf{R}_l \\ \boldsymbol{V} \end{bmatrix} \in \mathcal{L}' \cap \mathcal{F}_{T'}, \ \|\boldsymbol{V} - (zI - A)\mathbf{R}_l\|_{\mathcal{E}_1} < 1.
\end{aligned} \tag{20}
$$

The advantage of introducing this extra element to the synthesis task is that it allows more freedom in selecting the spatiotemporal constraint $\mathcal{L}' \cap \mathcal{F}_{T'}$ in which the virtual system responses $\{\mathbf{R}_l, \mathbf{V}\}$ are constrained to lie.

Remark 3: Similarly to the leaky boundaries setting above, if the state matrices (A, B_2) naturally decompose into localizable and non-localizable components (e.g., $A = A_{\text{loc}} + A_{\text{leak}}$, $B_2 = B_{\text{loc}} + B_{\text{leak}}$), then a similar modification can be made to the parameterization (19).

We use these stabilizing controller parameterizations in the next section to formulate an optimal control problem that produces controllers with provable sub-optimality bounds relative to the centralized optimal controller.

V. Performance Bounds

From Table I we see that the actual closed loop map from the disturbance $\boldsymbol{\delta}_x$ to state $\mathbf{x}$ and control action $\mathbf{u}$ achieved by a controller implemented using the approximate system responses $\{\mathbf{R}_c, \mathbf{M}_c\}$ is given by

$$
\begin{bmatrix} \mathbf{x} \\ \mathbf{u} \end{bmatrix} = \begin{bmatrix} \mathbf{R}_c \\ \mathbf{M}_c \end{bmatrix} (I + \boldsymbol{\Delta})^{-1} \boldsymbol{\delta}_x. \tag{21}
$$

Combining the regulated output description (1b), the actual closed loop maps (21) and either of the augmented SLPs (19) or (20), we can pose the following virtually localizable optimal control problem

$$
\begin{array}{cl}
\underset{\mathbf{R}_c, \mathbf{M}_c, \boldsymbol{\Delta}}{\text{minimize}} & \left\| \begin{bmatrix} C_1 & D_{12} \end{bmatrix} \begin{bmatrix} \mathbf{R}_c \\ \mathbf{M}_c \end{bmatrix} (I + \boldsymbol{\Delta})^{-1} B_1 \right\| \\
\text{s.t.} & \text{constraint (19) or (20),}
\end{array} \tag{22}
$$

where we use $\boldsymbol{\Delta}$ to denote $\mathbf{V}$ if we are using the augmented SLP (19), and $\boldsymbol{\Delta}$ to denote $\mathbf{V} - (zI - A)\mathbf{R}_l$ if we are using the augmented SLP (20).

Optimization problem (22) is non-convex: however, if the norm chosen in the objective function is sub-multiplicative, we can upper bound the objective function by

$$
\frac{\|B_1\|}{1 - \|\boldsymbol{\Delta}\|} \left\| \begin{bmatrix} C_1 & D_{12} \end{bmatrix} \begin{bmatrix} \mathbf{R}_c \\ \mathbf{M}_c \end{bmatrix} \right\|, \tag{23}
$$

where we have used the sub-multiplicative property, that $\|\boldsymbol{\Delta}\| < 1$ and the power-series expansion of the the inverse $(I + \boldsymbol{\Delta})^{-1} = \sum_{k=0}^{\infty} \boldsymbol{\Delta}^k$. As the next lemma shows, this bound is indeed quasi-convex and hence can be effectively optimized.

Lemma 1: For a convex function $f : \mathcal{D} \to \mathbb{R}$, a non-negative convex function $g : \mathcal{D} \to \mathbb{R}_+$, both defined on some domain $\mathcal{D}$, and a convex set $C \subseteq \mathcal{D}$, it holds that

$$
\min_{x \in C} \frac{f(x)}{1 - g(x)} = \min_{\gamma \in [0,1)} \tfrac{1}{1-\gamma} \min_{x \in C} \{ f(x) \mid g(x) \leq \gamma \} \tag{24}
$$

Proof: (Sketch) As $g(x)$ is non-negative, we can write

$$
\begin{aligned}
\min_{x \in C} \frac{f(x)}{1 - g(x)} &= \min_{\gamma, x \in C} \frac{f(x)}{1 - \gamma} \quad \text{s.t.} \quad g(x) \leq \gamma \\
&= \min_{\gamma} \tfrac{1}{1-\gamma} \min_{x \in C} \{ f(x) \mid g(x) \leq \gamma \},
\end{aligned}
$$

which proves the result. ∎

Thus applying the bound (23) and Lemma 1, we can compute an upper bound to optimization problem (22) by solving the following quasi-convex problem:

$$
\begin{array}{cl}
\min_{\gamma \in [0,1)} \frac{1}{1-\gamma} \underset{\mathbf{R}_c, \mathbf{M}_c, \boldsymbol{\Delta}}{\text{minimize}} & \left\| \begin{bmatrix} C_1 & D_{12} \end{bmatrix} \begin{bmatrix} \mathbf{R}_c \\ \mathbf{M}_c \end{bmatrix} \right\| \\
\text{s.t.} & \|\boldsymbol{\Delta}\| \leq \gamma \\
& \text{constraint (19) or (20).}
\end{array} \tag{25}
$$

Corresponding lower bounds on the cost achievable by a centralized optimal controller can be further computed using the methods described in [28], allowing us to bound the sub-optimality gap of the resulting stabilizing controller. We illustrate these ideas in the next section, and explore different trade-offs that may arise in the design of a large-scale cyber-physical system.

VI. Case Studies

We now apply the methods developed to a simple bidirectional chain system and a power-inspired system example; in both cases, we compute control policies subject to communication delays between sub-controllers such that the resulting problem is neither localizable nor quadratically invariant. We note that under our setup, a communication speed of zero leads to a completely decentralized controller implementation, whereas when communication speed tends to ∞ one obtains a fully centralized controller.

We solve the virtual localized optimal control problem (25) using the augmented SLP (19), and use the $\|\cdot\|_{\mathcal{E}_1}$ norm for the objective function. We use this norm for computational convenience, as it leads to a column-wise separable problem, although recent and classical work do support the use of this metric in certain cases [29]–[31] – if a different induced norm, such as the $\mathcal{L}_1$ norm, is chosen then the methods for partially separable problems described [8] can be used to solve the corresponding SLS problem.

A. Chain System

The first example considered is a marginally stable bi-directional chain with $n = 51$ nodes, where each node corresponds to a scalar state. For a coupling value of $\alpha \in [0, .5)$ node dynamics are given by

$$
x_i[t+1] = (1-2\alpha)x_i[t] + \alpha x_{i-1}[t] + \alpha x_{i+1}[t] + b_i u_i[t] + w_i[t],
$$

where $b_i = 0$ if i is odd, else $i = 1$. For the interior nodes $1 < i < 51$, and

$$
x_1[t+1] = (1 - \alpha)x_1[t] + \alpha x_2[t] + u_1[t] + w_1[t],
$$

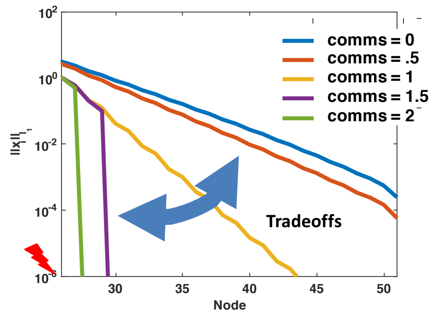

Fig. 2: For an impulse striking node 25, we plot the amplification at each node in the right half of the chain (from node 25 on) for various communication speeds. Here amplification is measured by $\|x_i\|_{\ell_1} = \sum_{t=0}^{\infty} |x_i(t)|$.

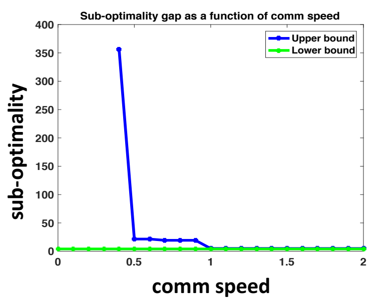

Fig. 3: Sub-optimality of the (approximately) localized controller as a function of communication speed. Note that for a communication speed of 1x (the minimum required for QI), we achieve (approximately) globally optimal behavior via an approximately localized-controller that enjoys $O(1)$ synthesis and implementation complexity.

$$x_{51}[t+1] = (1-\alpha)x_{51}[t] + \alpha x_{50}[t] + u_{51}[t] + w_{51}[t]$$

for the end nodes $i = 1, 51$. The regulated output of the system is given by $\bar{z}^T[t] = [x^T[t], u^T[t]]^T$. For this case study we pick $\alpha = .3$, set $d = 4$ and $T = 15$, and vary communication speed from 2x to 0x that of dynamics propagation – we further impose a sensing/actuation delay of 1. Figure 2 illustrates the effect of an impulse striking node 25 on the right-half (i.e., from node 25 on) of the chain for varying communication speeds – in particular for each node we plot $\|x_i\|_{\ell_1} = \sum_{t=0}^{\infty} |x_i(t)|$. We see that as communication speed decreases, so does the degree of localization of the disturbance and the closed loop performance of the system – note that we are able to compute a completely decentralized controller with communication speed of 0x. Figure 3 plots the ratio of our computed upper-bound to the lower-bound computed using the methods in [28]. We see in particular that for a communication speed of 1x (the minimum needed to satisfy QI), we achieve (approximately) globally optimal behavior via an approximately localized-controller that enjoys $O(1)$ synthesis and implementation complexity. We further note the conservativeness of our sufficient condition, in that we are not able to guarantee (a priori) the stability of the completely decentralized controller.

In the proceeding case study we consider nodes with more complex (non-scalar) dynamics and show illustrate that the sub-optimality gap can be made small even when the open-loop is unstable.

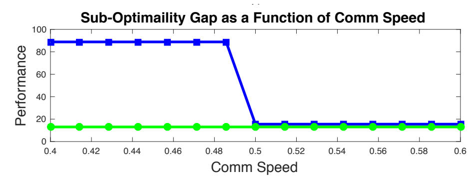

Fig. 4: Upper and lower bounds for the closed loop performance of the power system with 36 busses. With a communication speed of 0.5x the system in not localizable nor QI, but achieves near globally optimal performance.

B. Power System Example

Consider the swing equations for a n-bus power network which evolves according to

$$m_i \ddot{\theta}_i = -d_i \dot{\theta}_i - \sum_{j \in \mathcal{N}_i} k_{ij}(\theta_i - \theta_j) + w_i + u_i, \quad i = 1, \ldots, n, \tag{26}$$

where $\theta_i, d_i, m_i, w_i, u_i$ denote model the phase angle deviation, damping, inertia, disturbance, and control action of controllable load at bus i, finally $j \in \mathcal{N}_k$ implies busses i and j are connected. This can be re-written as a discrete time state-space system with state vector $x_i = [\theta_i, \dot{\theta}_i]^T$ and

$$A_{ii} = \begin{bmatrix} 1 & dt \\ \frac{k_i}{m_i} dt & 1 - \frac{d_i}{m_i} dt \end{bmatrix}, \; A_{ij} = \begin{bmatrix} 0 & 0 \\ \frac{k_{ij}}{m_i} dt & 0 \end{bmatrix},$$
$$B_{ii} = \begin{bmatrix} 0 & 1 \end{bmatrix}^T,$$

where $k_i = \sum_{j \in \mathcal{N}_i} k_{ij}$ and dt is the discretization time step. An equal penalty on state deviation and control action is chosen, i.e. $\begin{bmatrix} C_1 & D_{12} \end{bmatrix} = I$. Further modelling details can be found in [8, §V]. The parameters m_i^{-1}, d_i, k_{ij} are sampled from a uniform distribution over $[0, 2], [1, 1.5], [0.5, 1]$ respectively. The full matrix A is then scaled to make it marginally stable. In this example we consider an $n = 36$ bus model which gives $A \in \mathbb{R}^{72 \times 72}$ and use a discretization step of $dt = 2$. For the case study that follows we use a horizon length of $T = 15$, a 1-hop locaizability region, and communication speed of 0.5. This parameter setting renders the system non-QI and not localizable.

We investigate performance bounds as a function of communication speed, and as before compute a lower-bound using the methods in [28]. Figure 4 shows that as communication speed decreases the sub-optimality gap can rise significantly, and in fact seems to exhibit a "phase transition." We believe that when communication speed decreases below the "effective propagation rate" of disturbances, the closed loop performance of the corresponding controller degrades significantly – future work will look to formalize this idea. As with the previous examples, because the communication speed is slower than the propagation of dynamics in the plant, we are working in the non-QI regime. The observed sharp phase transition is not specific to this example, and has been observed in other numerical experiments.

VII. Conclusions

In this paper we defined and analyzed the notion of virtually localizable systems under state-feedback control.

In particular, we first developed a necessary and sufficient condition for a controller synthesized using the SLA to stabilize a family of physical plants. We then showed how this robustness result could be combined with the idea of virtual actuation and system responses, used solely at the time of synthesis, to compute stabilizing localized controllers that enjoy $O(1)$ synthesis and implementation complexity relative to the full size of the system. We then used these ideas to pose a virtually localized optimal control problem, and showed how it could be used to provide sub-optimality guarantees on the resulting stabilizing controllers. We finished with some numerical examples exploring design tradeoffs that arise in the context of the design of large-scale non-localizable systems, and observed an interesting phase transition phenomena as a function of communication speed. Future work will look to extend these results to the output-feedback setting, and to formally understand the empirically observed phase transition.

REFERENCES

[1] Y.-C. Ho and K.-C. Chu, "Team decision theory and information structures in optimal control problems–part i," *Automatic Control, IEEE Transactions on*, vol. 17, no. 1, pp. 15–22, 1972.

[2] A. Mahajan, N. Martins, M. Rotkowitz, and S. Yuksel, "Information structures in optimal decentralized control," in *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, 2012, pp. 1291–1306.

[3] M. Rotkowitz and S. Lall, "A characterization of convex problems in decentralized control," *Automatic Control, IEEE Transactions on*, vol. 51, no. 2, pp. 274–286, 2006.

[4] B. Bamieh, F. Paganini, and M. A. Dahleh, "Distributed control of spatially invariant systems," *Automatic Control, IEEE Transactions on*, vol. 47, no. 7, pp. 1091–1107, 2002.

[5] 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.

[6] A. Nayyar, A. Mahajan, and D. Teneketzis, "Decentralized stochastic control with partial history sharing: A common information approach," *IEEE Transactions on Automatic Control*, vol. 58, no. 7, pp. 1644–1658, July 2013.

[7] Y.-S. Wang, N. Matni, and J. C. Doyle, "A system level approach to controller synthesis," *submitted to IEEE Transactions on Automatic Control*, 2017. [Online]. Available: https://arxiv.org/abs/1610.04815

[8] ——, "Separable and localized system level synthesis for large-scale systems," *submitted to IEEE Transactions on Automatic Control*, 2017. [Online]. Available: https://arxiv.org/abs/1701.05880

[9] ——, "Localized LQR optimal control," in *2014 53rd IEEE Conference on Decision and Control (CDC)*, 2014.

[10] Y.-S. Wang and N. Matni, "Localized LQG optimal control for large-scale systems," in *2016 IEEE American Control Conference (ACC)*, 2016.

[11] M. Rotkowitz, R. Cogill, and S. Lall, "Convexity of optimal control over networks with delays and arbitrary topology," *Int. J. Syst., Control Commun.*, vol. 2, no. 1/2/3, pp. 30–54, Jan. 2010. [Online]. Available: http://dx.doi.org/10.1504/IJSCC.2010.031157

[12] L. Lessard and S. Lall, "Optimal controller synthesis for the decentralized two-player problem with output feedback," in *2012 IEEE American Control Conference (ACC)*, June 2012.

[13] P. Shah and P. A. Parrilo, "$\mathcal{H}_2$-optimal decentralized control over posets: A state space solution for state-feedback," in *Decision and Control (CDC), 2010 49th IEEE Conference on*, 2010.

[14] A. Lamperski and J. C. Doyle, "Output feedback $\mathcal{H}_2$ model matching for decentralized systems with delays," in *2013 IEEE American Control Conference (ACC)*, June 2013.

[15] L. Lessard, M. Kristalny, and A. Rantzer, "On structured realizability and stabilizability of linear systems," in *American Control Conference (ACC), 2013*, June 2013, pp. 5784–5790.

[16] T. Tanaka and P. A. Parrilo, "Optimal output feedback architecture for triangular LQG problems," in *2014 IEEE American Control Conference (ACC)*, June 2014.

[17] A. Lamperski and L. Lessard, "Optimal decentralized state-feedback control with sparsity and delays," *Automatica*, vol. 58, pp. 143–151, 2015.

[18] C. W. Scherer, "Structured $\mathcal{H}_\infty$-optimal control for nested interconnections: A state-space solution," *Systems and Control Letters*, vol. 62, pp. 1105–1113, 2013.

[19] L. Lessard, "State-space solution to a minimum-entropy $\mathcal{H}_\infty$-optimal control problem with a nested information constraint," in *2014 53rd IEEE Conference on Decision and Control (CDC)*, 2014. [Online]. Available: http://arxiv.org/pdf/1403.5020v2.pdf

[20] N. Matni, "Distributed control subject to delays satisfying an $\mathcal{H}_\infty$ norm bound," in *2014 53rd IEEE Conference on Decision and Control (CDC)*, 2014. [Online]. Available: http://arxiv.org/pdf/1402.1559.pdf

[21] M. Fardad, F. Lin, and M. R. Jovanovic, "Sparsity-promoting optimal control for a class of distributed systems," in *2011 IEEE American Control Conference (ACC)*, June 2011.

[22] G. Fazelnia, R. Madani, and J. Lavaei, "Convex relaxation for optimal distributed control problem," in *2014 53rd IEEE Conference on Decision and Control (CDC)*, 2014.

[23] Y. Wang, J. Lopez, and M. Sznaier, "Sparse static output feedback controller design via convex optimization," in *2014 53rd IEEE Conference on Decision and Control (CDC)*, 2014.

[24] N. Motee and A. Jadbabaie, "Approximation methods and spatial interpolation in distributed control systems," in *2009 IEEE American Control Conference (ACC)*, June 2009.

[25] Y.-S. Wang, N. Matni, and J. C. Doyle, "System level parameterizations, constraints and synthesis," *IEEE American Control Conference, to appear*, 2017.

[26] G. E. Dullerud and F. Paganini, *A Course In Robust Control Theory: A Convex Approach*. Springer-Verlag, 2000.

[27] M. A. Dahleh and Y. Ohta, "A necessary and sufficient condition for robust bibo stability," *Syst. Control Lett.*, vol. 11, no. 4, pp. 271–275, Oct. 1988. [Online]. Available: http://dx.doi.org/10.1016/0167-6911(88)90070-9

[28] Y. S. Wang, "Localized lqr with adaptive constraint and performance guarantee," in *2016 IEEE 55th Conference on Decision and Control (CDC)*, Dec 2016, pp. 2769–2776.

[29] J. Yu and A. Sideris, "Optimal induced l1-norm state feedback control," *Automatica*, vol. 35, no. 5, pp. 819 – 827, 1999. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0005109898002143

[30] G. Vossen and H. Maurer, "On l1-minimization in optimal control and applications to robotics," *Optimal Control Applications and Methods*, vol. 27, no. 6, pp. 301–321, 2006. [Online]. Available: http://dx.doi.org/10.1002/oca.781

[31] Y. Ebihara, D. Peaucelle, and D. Arzelier, "L1 gain analysis of linear positive systems and its application," in *2011 50th IEEE Conference on Decision and Control and European Control Conference*, Dec 2011, pp. 4029–4034.