

System Level Parameterizations, Constraints and Synthesis

Yuh-Shyang Wang, Nikolai Matni, and John C. Doyle

Abstract—We introduce the system level approach to controller synthesis, which is composed of three elements: System Level Parameterizations (SLPs), System Level Constraints (SLCs) and System Level Synthesis (SLS) problems. SLPs provide a novel parameterization of all internally stabilizing controllers and the system responses that they achieve. These can be combined with SLCs to provide parameterizations of constrained stabilizing controllers. We provide a catalog of useful SLCs, and show that by using SLPs with SLCs, we can parameterize the largest known class of constrained stabilizing controllers that admit a convex characterization. Finally, we formulate the SLS problem, and show that it defines the broadest known class of constrained optimal control problems that can be solved using convex programming. We end by using the system level approach to computationally explore tradeoffs in controller performance, architecture cost, robustness and synthesis/implementation complexity.

I. INTRODUCTION

The Youla parameterization [1] represented an important shift towards a system level approach to optimal controller synthesis. Youla showed that there exists an isomorphism between a stabilizing controller and the resulting closed loop system response from sensors to actuators – therefore rather than synthesizing the controller itself, this system response (or Youla parameter) could be designed directly. Together with state-space methods, this contribution played a major role in shifting controller synthesis from an ad hoc, loop-at-a-time tuning process to a principled one with well defined notions of optimality. Indeed, this approach proved very powerful, and paved the way for the foundational results of robust and optimal control that would follow [2].

This paper presents an approach that is inspired by the system level thinking pioneered by Youla: rather than directly designing only the feedback loop between sensors and actuators, we propose directly designing the *entire closed loop response of the system*, as captured by the maps from process and measurement disturbances to control actions and states – as such, we call the proposed method a system level approach to controller synthesis. A distinction between our approach and Youla’s is that we explicitly model the internal delay structure of the feedback system, whereas Youla (and contemporary state-space methods) hid the internals of the controller, and focused instead on its input-output behavior. However, modern cyber-physical systems (CPS) are large-scale, physically distributed, and interconnected. Rather than a logically centralized controller, these systems are composed of several sub-controllers, each equipped with

their own sensors and actuators – these sub-controllers then exchange locally available information (such as sensor measurements or applied control actions) via a communication network. It follows that the information exchanged between sub-controllers is constrained by the delay, bandwidth and reliability properties of this communication network, ultimately manifesting as information asymmetry among sub-controllers of the system. It is this information asymmetry, as imposed by the underlying communication network, that lies at the heart of what makes distributed optimal controller synthesis challenging [3]–[7].

A defining feature of CPS is that controllers have internal delays, as specified by the exchange of information between constituent sub-controllers. These delays thus needed to be reintroduced into Youla synthesis methods, which aimed to hide the internals of the controller from the system engineer. There were also indications that introducing information asymmetry into the optimal control problem lead to intractable synthesis tasks [8]. Despite these apparent technical and conceptual challenges, a body of work [4]–[7], [9] that began in the early 2000s, and culminated with the introduction of quadratic invariance (QI) in the seminal paper [5], showed that for a large class of practically relevant systems, such internal structure could be incorporated into the Youla parameterization and still preserve the convexity of the optimal controller synthesis task. Informally, a system is quadratically invariant if sub-controllers are able to exchange information with each other faster than their control actions propagate through the CPS [10]. Further, this condition is tight, in the sense that QI is a necessary [11] and sufficient [5] condition for subspace constraints on the controller to be enforceable via convex constraints on the Youla parameter.

These results then set the stage for the advancements in distributed optimal controller synthesis that would follow (e.g., [12]–[16] and references therein). As impressive as these results were, they also showed that the QI framework does not allow for synthesis methods to scale to very large systems. We show in Section II that the QI framework, which adapts the Youla parameterization to a distributed setting, fails to capture certain constraints that are needed for optimal controller synthesis to scale to arbitrarily large systems. In particular, when the underlying physical system is strongly connected,¹ the QI framework does not allow for localized controllers, in which local sub-controllers only access a subset of system-wide measurements (c.f., Section IV-5), to be synthesized using convex programming – perhaps

The authors are with the department of Control and Dynamical Systems, California Institute of Technology, Pasadena, CA 91125, USA (`{yswang, nmatni, doyle}@caltech.edu`).

¹We say that a plant is strongly connected if the state of any subsystem can eventually alter the state of all other subsystems.

counter-intuitively, this statement holds true even when sub-controllers can exchange information with no delay (c.f., Example 1). Although this may seem surprising, note that implicit to the Youla parameterization is that sub-controllers can only exchange locally collected measurements with each other, and not, for instance, locally applied control actions.

To overcome this limitation, we propose the system level approach to controller synthesis, which is composed of three elements: System Level Parameterizations (SLPs), System Level Constraints (SLCs) and System Level Synthesis (SLS) problems. This paper is organized as follows: in §II, we define the system model considered in this paper, review relevant results from the distributed optimal control and QI literature, provide a motivating example for a system level approach, and present a survey of our main results. We then define and analyze SLPs in §III, which provide a novel parameterization of all internally stabilizing controllers and the closed loop system responses that they achieve. In §IV, we provide a catalog of useful SLCs that can then be imposed on these system responses and the controllers that achieve them – we further show that by using SLPs and SLCs, we can parameterize a set of constrained stabilizing controllers that are a strict superset of those that can be parameterized using quadratic invariance, hence generalizing the QI framework. Finally, we use SLPs and SLCs to formulate the SLS problem, and show that it defines the broadest known class of constrained optimal control problems that can be solved using convex programming.

Proofs: Complete proofs for all results, as well as additional results, discussion and examples, can be found in [17]–[19].

II. PROBLEM STATEMENT AND MAIN RESULTS

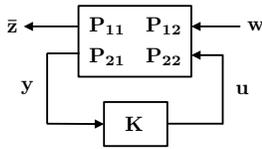


Fig. 1. Interconnection of the plant \mathbf{P} and controller \mathbf{K} .

Preliminaries & Notation: We use lower and upper case boldface Latin letters such as \mathbf{x} and \mathbf{G} to denote signals and transfer matrices, respectively, and calligraphic letters such as \mathcal{S} to denote sets. We work with discrete time linear time invariant systems, but unless stated otherwise, all results extend naturally to the continuous time setting. We use standard definitions of the Hardy spaces \mathcal{H}_2 and \mathcal{H}_∞ , and denote their restriction to the set of real-rational proper transfer matrices by \mathcal{RH}_2 and \mathcal{RH}_∞ . We use $G[i]$ to denote the i th spectral component of a transfer function \mathbf{G} , i.e., $\mathbf{G}(z) = \sum_{i=0}^{\infty} \frac{1}{z^i} G[i]$ for $|z| > 1$. We use \mathcal{F}_T to denote the space of finite impulse response (FIR) transfer matrices with horizon T , i.e., $\mathcal{F}_T := \{\mathbf{G} \in \mathcal{RH}_\infty \mid \mathbf{G} = \sum_{i=0}^T \frac{1}{z^i} G[i]\}$.

1) *System Model:* We consider discrete time linear time invariant (LTI) systems of the form

$$x[t+1] = Ax[t] + B_1w[t] + B_2u[t] \quad (1a)$$

$$\bar{z}[t] = C_1x[t] + D_{11}w[t] + D_{12}u[t] \quad (1b)$$

$$y[t] = C_2x[t] + D_{21}w[t] + D_{22}u[t] \quad (1c)$$

where x , u , w , y , \bar{z} are the state vector, control action, external disturbance, measurement, and regulated output, respectively. Equation (1) can be written in state space form as

$$\mathbf{P} = \left[\begin{array}{c|cc} A & B_1 & B_2 \\ \hline C_1 & D_{11} & D_{12} \\ C_2 & D_{21} & D_{22} \end{array} \right] = \begin{bmatrix} \mathbf{P}_{11} & \mathbf{P}_{12} \\ \mathbf{P}_{21} & \mathbf{P}_{22} \end{bmatrix}$$

where $\mathbf{P}_{ij} = C_i(zI - A)^{-1}B_j + D_{ij}$. We refer to \mathbf{P} as the open loop plant model.

Consider a dynamic output feedback control law $\mathbf{u} = \mathbf{K}\mathbf{y}$. The controller \mathbf{K} is assumed to have the state space realization

$$\xi[t+1] = A_k\xi[t] + B_ky[t] \quad (2a)$$

$$u[t] = C_k\xi[t] + D_ky[t], \quad (2b)$$

where ξ is the internal state of the controller. We have $\mathbf{K} = C_k(zI - A_k)^{-1}B_k + D_k$. A schematic diagram of the interconnection of the plant \mathbf{P} and the controller \mathbf{K} is shown in Figure 1.

The following assumptions are made throughout the paper.

Assumption 1: The interconnection in Figure 1 is well-posed – the matrix $(I - D_{22}D_k)$ is invertible.

Assumption 2: Both the plant and the controller realizations are stabilizable and detectable; i.e., (A, B_2, C_2) and (A_k, B_k, C_k) are stabilizable and detectable.

2) *Structured Controller Synthesis, Youla, and QI:* We follow the paradigm adopted in [5], [12]–[16], and focus on information asymmetry introduced by delays in the communication network. In the references cited above, locally acquired measurements are exchanged between sub-controllers subject to delays imposed by the communication network,² which manifest as subspace constraints on the controller itself.

We consider the following optimal structured controller synthesis task, as defined in [5], [11], [20], [21]:

$$\begin{aligned} & \underset{\mathbf{K}}{\text{minimize}} && \|\mathbf{P}_{11} + \mathbf{P}_{12}\mathbf{K}(I - \mathbf{P}_{22}\mathbf{K})^{-1}\mathbf{P}_{21}\| \\ & \text{subject to} && \mathbf{K} \text{ internally stabilizes } \mathbf{P}, \mathbf{K} \in \mathcal{C}, \end{aligned} \quad (3)$$

for \mathcal{C} a subspace. This subspace can enforce, for instance, the information sharing constraints imposed on the controller \mathbf{K} by the underlying communication network.

A synthesis of the main results of these papers can be expressed as follows: if the subspace \mathcal{C} is quadratically invariant [5] with respect to \mathbf{P}_{22} , then the set of all stabilizing controllers lying in subspace \mathcal{C} can be parameterized by those stable transfer matrices $\mathbf{Q} \in \mathcal{RH}_\infty$ satisfying $\mathfrak{M}(\mathbf{Q}) \in \mathcal{C}$,

²Note that this delay may range from 0, modeling instantaneous communication between sub-controllers, to infinite, modeling no communication between sub-controllers.

Let \mathcal{S} denote such a SLC, and assume that it admits a convex representation. Further, let $g(\cdot)$ be a convex functional. Then the solution to the SLS problem:

$$\underset{\{\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L}\}}{\text{minimize}} \quad g(\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L}) \quad (11a)$$

$$\text{subject to} \quad \text{equations (7) - (9),} \quad \begin{bmatrix} \mathbf{R} & \mathbf{N} \\ \mathbf{M} & \mathbf{L} \end{bmatrix} \in \mathcal{S} \quad (11b)$$

can be found via convex programming. Further, as we briefly discuss in Section V-3, if the SLC and the objective $g(\cdot)$ are localized and separable, then a structured optimal controller can be synthesized and implemented in a scalable manner, i.e., with $O(1)$ parallel computational and implementation complexity relative to the size of the full system.

Example 2 (Example 1 cont'd): We now return to the motivating example introduced above. In the case of a full control ($B_2 = I$) state-feedback ($C_2 = I, D_{21} = 0$) problem, the conditions (7)-(9) simplify to $(zI - A)\mathbf{R} - \mathbf{M} = I$, $\mathbf{R}, \mathbf{M} \in \frac{1}{z}\mathcal{RH}_\infty$, and a controller achieving the desired response is given by $\mathbf{K} = \mathbf{M}\mathbf{R}^{-1}$ [22]. Further this controller can be implemented as

$$\hat{\mathbf{w}} = \mathbf{x} - \hat{\mathbf{x}}, \quad \mathbf{u} = z\mathbf{M}\hat{\mathbf{w}}, \quad \hat{\mathbf{x}} = (z\mathbf{R} - I)\hat{\mathbf{w}}. \quad (12)$$

Suppose that we aim to synthesize an optimal controller that has a communication topology given by the support of A – from the above implementation, it suffices to constrain the support of transfer matrices \mathbf{R} and \mathbf{M} to be a subset of that of A . It can be checked that $\mathbf{R} = \frac{1}{z}I$, and $\mathbf{M} = -\frac{1}{z}A$ satisfy the above constraints, and recover the globally optimal controller $\mathbf{K} = -A$. This controller cannot be computed using quadratic invariance and the Youla parameterization if the system is strongly connected and sparse.

III. SYSTEM LEVEL PARAMETERIZATION

In this section, we show that the affine subspace defined by the constraints (7) - (9) parameterizes all stable achievable system responses $\{\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L}\}$, and that the controller $\mathbf{K} = \mathbf{L} - \mathbf{M}\mathbf{R}^{-1}\mathbf{N}$, which admits a realization as described in (10), parameterizes all internally stabilizing controllers for a strictly proper plant \mathbf{P}_{22} .⁴ Consider a strictly proper plant

$$\mathbf{P} = \begin{bmatrix} A & B_1 & B_2 \\ \hline C_1 & D_{11} & D_{12} \\ C_2 & D_{21} & 0 \end{bmatrix}. \quad (13)$$

Letting $\delta_x[t] = B_1 w[t]$ denote the disturbance on the state, and $\delta_y[t] = D_{21} w[t]$ denote the disturbance on the measurement, the dynamics defined by plant (13) can be written as

$$\begin{aligned} x[t+1] &= Ax[t] + B_2 u[t] + \delta_x[t] \\ y[t] &= C_2 x[t] + \delta_y[t]. \end{aligned} \quad (14)$$

We define a system response $\{\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L}\}$ from perturbations (δ_x, δ_y) to state and control inputs (\mathbf{x}, \mathbf{u}) via the following relation:

$$\begin{bmatrix} \mathbf{x} \\ \mathbf{u} \end{bmatrix} = \begin{bmatrix} \mathbf{R} & \mathbf{N} \\ \mathbf{M} & \mathbf{L} \end{bmatrix} \begin{bmatrix} \delta_x \\ \delta_y \end{bmatrix}. \quad (15)$$

⁴The non-strictly proper case follows from standard arguments [23], c.f., Section III.C in [17].

Substituting the output feedback control law $\mathbf{u} = \mathbf{K}\mathbf{y}$ into the z -transform of system equation (14), we obtain

$$(zI - A - B_2 \mathbf{K} C_2) \mathbf{x} = \delta_x + B_2 \mathbf{K} \delta_y.$$

For a proper controller \mathbf{K} , the transfer matrix $(zI - A - B_2 \mathbf{K} C_2)$ is always invertible, hence we obtain the following expressions for the system response (15) in terms of an output feedback controller \mathbf{K} :

$$\begin{aligned} \mathbf{R} &= (zI - A - B_2 \mathbf{K} C_2)^{-1}, & \mathbf{M} &= \mathbf{K} C_2 \mathbf{R}, \\ \mathbf{N} &= \mathbf{R} B_2 \mathbf{K}, & \mathbf{L} &= \mathbf{K} + \mathbf{K} C_2 \mathbf{R} B_2 \mathbf{K}. \end{aligned} \quad (16)$$

We now present one of the main results of the paper: an algebraic characterization of the set $\{\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L}\}$ of output-feedback system responses that are achievable by an internally stabilizing controller \mathbf{K} .

Theorem 1: For the output feedback system (13), the following are true:

- (a) The affine subspace described by:

$$[zI - A \quad -B_2] \begin{bmatrix} \mathbf{R} & \mathbf{N} \\ \mathbf{M} & \mathbf{L} \end{bmatrix} = [I \quad 0] \quad (17a)$$

$$\begin{bmatrix} \mathbf{R} & \mathbf{N} \\ \mathbf{M} & \mathbf{L} \end{bmatrix} \begin{bmatrix} zI - A \\ -C_2 \end{bmatrix} = \begin{bmatrix} I \\ 0 \end{bmatrix} \quad (17b)$$

$$\mathbf{R}, \mathbf{M}, \mathbf{N} \in \frac{1}{z}\mathcal{RH}_\infty, \quad \mathbf{L} \in \mathcal{RH}_\infty \quad (17c)$$

parameterizes all system responses (16) achievable by an internally stabilizing controller \mathbf{K} .

- (b) For any transfer matrices $\{\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L}\}$ satisfying (17), the controller $\mathbf{K} = \mathbf{L} - \mathbf{M}\mathbf{R}^{-1}\mathbf{N}$ is internally stabilizing and achieves the desired response (16).

The necessity of a stable and achievable system response $\{\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L}\}$ lying in the affine subspace (17) follows from rote calculation (c.f. Appendix B in [17]), and hence we focus on proving that for any system response $\{\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L}\}$ lying in the affine subspace (17), there exists an internally stabilizing controller \mathbf{K} that leads to the desired system response (16). From the relations in (16), we notice the identity $\mathbf{K} = \mathbf{L} - \mathbf{K} C_2 \mathbf{R} B_2 \mathbf{K} = \mathbf{L} - \mathbf{M}\mathbf{R}^{-1}\mathbf{N}$. This relation leads to the controller structure given in Figure 2, with $\tilde{\mathbf{R}}^+ = z\tilde{\mathbf{R}} = z(I - z\mathbf{R})$, $\tilde{\mathbf{M}} = z\mathbf{M}$, and $\tilde{\mathbf{N}} = -z\mathbf{N}$. As was the case for the state feedback setting, it can be verified that $\tilde{\mathbf{R}}^+$, $\tilde{\mathbf{M}}$, and $\tilde{\mathbf{N}}$ are all in \mathcal{RH}_∞ . Therefore, the structure given in Figure 2 is well defined. The controller implementation of Figure 2 is governed by the equations:

$$z\beta = \tilde{\mathbf{R}}^+ \beta + \tilde{\mathbf{N}} \mathbf{y}, \quad \mathbf{u} = \tilde{\mathbf{M}} \beta + \mathbf{L} \mathbf{y}. \quad (18)$$

The control implementation equations (18) can be interpreted as an extension of the state-space realization (2) of a controller \mathbf{K} . The benefit of this implementation is that arbitrary convex constraints imposed on the transfer matrices $\tilde{\mathbf{R}}^+, \tilde{\mathbf{M}}, \tilde{\mathbf{N}}, \mathbf{L}$ carry over directly to the controller implementation. It remains to be shown is that the controller implementation (18) is internally stabilizing and achieves the desired system response (16). Notice that all of the blocks in Figure 2 are stable filters – thus, as long as the origin

$(x, \beta) = (0, 0)$ is asymptotically stable, all signals internal to the block diagram will decay to zero. We introduce external perturbations $\delta_x, \delta_y, \delta_u$, and δ_β to the system and note that the perturbations appearing on other links of the block diagram can all be expressed as a combination of the perturbations $(\delta_x, \delta_y, \delta_u, \delta_\beta)$ being acted upon by some stable transfer matrices. Thus it suffices to check the input-output stability of the closed loop transfer matrices from perturbations $(\delta_x, \delta_y, \delta_u, \delta_\beta)$ to controller signals $(\mathbf{x}, \mathbf{u}, \mathbf{y}, \beta)$ to determine the internal stability of the structure [23].

Lemma 1: Consider the output feedback system (13). For any system response $\{\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L}\}$ lying in the affine subspace defined by (17), the controller $\mathbf{K} = \mathbf{L} - \mathbf{M}\mathbf{R}^{-1}\mathbf{N}$ (with structure shown in Figure 2) internally stabilizes the plant. In addition, the desired system response, as specified by $\mathbf{x} = \mathbf{R}\delta_x + \mathbf{N}\delta_y$ and $\mathbf{u} = \mathbf{M}\delta_x + \mathbf{L}\delta_y$, is achieved.

Proof: (sketch) For any system response $\{\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L}\}$ lying in the affine subspace defined by (17), we construct a controller using the structure given in Figure 2. We now check the stability of the closed loop transfer matrices from the perturbations $(\delta_x, \delta_y, \delta_u, \delta_\beta)$ to the internal variables $(\mathbf{x}, \mathbf{u}, \mathbf{y}, \beta)$. It can be checked that the closed loop transfer matrices from $(\delta_x, \delta_y, \delta_u, \delta_\beta)$ to $(\mathbf{x}, \mathbf{u}, \mathbf{y}, \beta)$ are given as in Table I. Equation (17c) implies that all sixteen transfer

TABLE I

CLOSED LOOP MAPS FROM PERTURBATIONS TO INTERNAL VARIABLES

	δ_x	δ_y	δ_u	δ_β
\mathbf{x}	\mathbf{R}	\mathbf{N}	$\mathbf{R}\mathbf{B}_2$	$\frac{1}{2}\mathbf{N}\mathbf{C}_2$
\mathbf{u}	\mathbf{M}	\mathbf{L}	$\mathbf{I} + \mathbf{M}\mathbf{B}_2$	$\frac{1}{2}\mathbf{L}\mathbf{C}_2$
\mathbf{y}	$\mathbf{C}_2\mathbf{R}$	$\mathbf{I} + \mathbf{C}_2\mathbf{N}$	$\mathbf{C}_2\mathbf{R}\mathbf{B}_2$	$\frac{1}{2}\mathbf{C}_2\mathbf{N}\mathbf{C}_2$
β	$-\frac{1}{2}\mathbf{B}_2\mathbf{M}$	$-\frac{1}{2}\mathbf{B}_2\mathbf{L}$	$-\frac{1}{2}\mathbf{B}_2\mathbf{M}\mathbf{B}_2$	$\frac{1}{2}\mathbf{I} - \frac{1}{2}\mathbf{A} + \mathbf{B}_2\mathbf{L}\mathbf{C}_2$

matrices in Table I are stable, so the implementation in Figure 2 is internally stable. Furthermore, the desired system response from (δ_x, δ_y) to (\mathbf{x}, \mathbf{u}) is achieved. ■

IV. SYSTEM LEVEL CONSTRAINTS

In this section, we provide a catalog of useful SLCs that can be naturally imposed on the SLPs described in the previous section. In particular, we show that QI subspace constraints are a special case of SLCs, and as such, we provide here a description of the largest known class of constrained stabilizing controllers that admit a convex parameterization. In the interest of space, we omit descriptions of controller robustness, controller architecture and positivity SLCs, which can be found in [17].

1) *Constraints on the Youla parameter:* We show that any constraint imposed on the Youla parameter can be translated into a SLC, and vice versa. In particular, if this constraint is convex, then so is the corresponding SLC. Consider the following modification of the standard Youla parameterization, which characterizes a set of constrained internally stabilizing controllers \mathbf{K} for a plant (13):

$$\mathbf{K} = \mathfrak{Y}(\mathbf{Q})[\mathfrak{X}(\mathbf{Q})]^{-1}, \quad \mathbf{Q} \in \mathcal{Q} \cap \mathcal{RH}_\infty. \quad (19)$$

Here \mathfrak{Y} and \mathfrak{X} are affine maps defined in terms of the doubly co-prime factorization of the plant (13) mentioned in §II, and \mathcal{Q} is an arbitrary set – if we take $\mathcal{Q} = \mathcal{RH}_\infty$, we recover the standard Youla parameterization (c.f., [23]). By appropriately varying the set \mathcal{Q} , one can then characterize all possible constrained internally stabilizing controllers,⁵ and hence this formulation is as general as possible. We now show that an equivalent parameterization can be given in terms of a SLC.

Theorem 2: The set of constrained internally stabilizing controllers described by (19) can be equivalently expressed as $\mathbf{K} = \mathbf{L} - \mathbf{M}\mathbf{R}^{-1}\mathbf{N}$, where the system response $\{\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L}\}$ lies in the set

$$\{\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L} \mid (17a) - (17c) \text{ hold, } \mathbf{L} \in \mathfrak{M}(\mathcal{Q})\}, \quad (20)$$

for \mathfrak{M} an invertible affine map. Further, this parameterization is convex if and only if \mathcal{Q} is convex.

In order to prove this result, we first need to understand the relationship between the controller \mathbf{K} , the Youla parameter \mathbf{Q} , and the system response $\{\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L}\}$. Fortunately, there exists a natural connection between these three objects.

Lemma 2: Let \mathbf{L} be defined as in (16), and an invertible affine map \mathfrak{M} as defined in [21], [24]. We then have that

$$\mathbf{L} = \mathbf{K}(\mathbf{I} - \mathbf{P}_{22}\mathbf{K})^{-1} = \mathfrak{M}(\mathbf{Q}). \quad (21)$$

Proof: (sketch) From [21], [24], we have $\mathbf{K}(\mathbf{I} - \mathbf{P}_{22}\mathbf{K})^{-1} = \mathfrak{M}(\mathbf{Q})$. The rest of the proof follows from algebraic manipulation of the closed loop response of system (13) and the definition of the map \mathbf{L} in (16). ■

Proof: [Proof of Theorem 2] The equivalence between the parameterizations (19) and (20) is readily obtained from Lemma 2. As \mathfrak{M} is an invertible affine mapping between \mathbf{L} and \mathbf{Q} , any convex constraint imposed on the Youla parameter \mathbf{Q} can be equivalently translated into a convex SLC imposed on \mathbf{L} , and vice versa. ■

2) *Quadratically invariant subspace constraints:* Recall that for a subspace \mathcal{C} that is quadratically invariant with respect to a plant \mathbf{P}_{22} , the set of internally stabilizing controllers \mathbf{K} that lie within the subspace \mathcal{C} can be expressed as the set of stable transfer matrices $\mathbf{Q} \in \mathcal{RH}_\infty$ satisfying $\mathfrak{M}(\mathbf{Q}) \in \mathcal{C}$, for \mathfrak{M} an invertible affine map as defined in [21], [24]. The following corollary is then immediate.

Corollary 1: Let \mathcal{C} be a subspace constraint that is quadratically invariant with respect to \mathbf{P}_{22} . Then the set of internally stabilizing controllers satisfying $\mathbf{K} \in \mathcal{C}$ can be parameterized as in Theorem 2.

Proof: Invoking Theorem 14 of [5], we have that $\mathbf{K} \in \mathcal{C}$ if and only if $\mathbf{L} = \mathbf{K}(\mathbf{I} - \mathbf{P}_{22}\mathbf{K})^{-1} \in \mathcal{C}$. The claim then follows immediately from Theorem 2. ■

Thus QI subspace constraints are a special case of SLCs.

Finally, we note that in [11], [20], the authors show that quadratic invariance is necessary for a subspace constraint \mathcal{C} on the controller \mathbf{K} to be enforceable via a convex constraint on the Youla parameter \mathbf{Q} . However, when \mathcal{C} is not a subspace constraint, no general methods exist to determine whether the set of internally stabilizing controllers lying in

⁵In particular, to ensure that $\mathbf{K} \in \mathcal{C}$, it suffices to enforce that $\mathfrak{Y}(\mathbf{Q})[\mathfrak{X}(\mathbf{Q})]^{-1} \in \mathcal{C}$.

\mathcal{C} admits a convex representation. In contrast, determining the convexity of a SLC is trivial.

3) *System performance constraints:* Let $g(\cdot)$ be a functional of the system response – it then follows that all internally stabilizing controllers satisfying a performance level, as specified by a scalar γ , are given by transfer matrices $\{\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L}\}$ satisfying the conditions of Theorem 1 and the SLC

$$g(\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L}) \leq \gamma. \quad (22)$$

Further, recall that the sublevel set of a convex functional is a convex set, and hence if g is convex, then so is the SLC (22). A particularly useful choice of convex functional is

$$g(\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L}) = \left\| \begin{bmatrix} C_1 & D_{12} \end{bmatrix} \begin{bmatrix} \mathbf{R} & \mathbf{N} \\ \mathbf{M} & \mathbf{L} \end{bmatrix} \begin{bmatrix} B_1 \\ D_{21} \end{bmatrix} + D_{11} \right\|, \quad (23)$$

for a system norm $\|\cdot\|$, which is equivalent to the objective function of the decentralized optimal control problem (3). Thus by imposing several performance SLCs (23) with different choices of norm, one can naturally formulate multi-objective optimal control problems.

Remark 1: For a continuous time system with norm $\|\cdot\|$ in (23) chosen to be the \mathcal{H}_2 norm, the transfer matrix (23) must be strictly proper, i.e., $D_{12}L[0]D_{21} + D_{11} = 0$.

4) *FIR constraints:* Given the parameterization of stabilizing controllers of Theorem 1, it is trivial to enforce that a system response be FIR with horizon T via the SLC $\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L} \in \mathcal{F}_T$. We argue that imposing a FIR SLC is beneficial in the following ways: (a) the closed loop response to an impulse disturbance is FIR of horizon T , where T can be set by the control designer – thus the settling time of the system can be accurately tuned; (b) the controller achieving the desired system response can be implemented using the FIR filter banks $\hat{\mathbf{R}}^+, \hat{\mathbf{M}}, \hat{\mathbf{N}}, \hat{\mathbf{L}} \in \mathcal{F}_T$, as illustrated in Figure 2 – this simplicity of implementation is useful when transitioning to practice; (c) when a FIR SLC is imposed, the resulting set of stable achievable system responses and corresponding controllers admit a finite dimensional representation – specifically, the constraints specified in Theorem 1 only need to be applied to the impulse response elements $\{R[t], M[t], N[t], L[t]\}_{t=0}^T$.⁶

5) *Subspace and sparsity constraints:* Let \mathcal{L} be a subspace of \mathcal{RH}_∞ . We can then parameterize all stable achievable system responses that lie in this subspace by adding the following SLC to the parameterization of Theorem 1:

$$\begin{bmatrix} \mathbf{R} & \mathbf{N} \\ \mathbf{M} & \mathbf{L} \end{bmatrix} \in \mathcal{L}. \quad (24)$$

Of particular interest are subspaces \mathcal{L} that define transfer matrices of sparse support. An immediate benefit of enforcing such sparsity constraints on the system response is that implementing the resulting controller (18) can be done in a *localized way*, i.e., each controller state β_i and control action u_i can be computed using a local subset (as defined by the

⁶The computational benefits claimed above hold only for discrete time systems, as a continuous time FIR transfer matrix is still an infinite dimensional object.

support of the system response) of the global controller state β and sensor measurements y . For this reason, we refer to the constraint (24) as a *localized SLC* when it defines a subspace with sparse support. Further, such localized constraints also allow for the resulting system response to be computed in a localized way, i.e., the global computation decomposes naturally into decoupled subproblems that depend only on local sub-matrices of the state-space representation (1) (c.f., [22], [25]). Clearly, both of these features are extremely desirable when computing controllers for large-scale systems.

Selecting an appropriate and feasible localized SLC, as defined by the subspace \mathcal{L} , is a subtle task: it depends on an interplay between actuator and sensor density, information exchange delay and disturbance propagation delay. Formally defining and analyzing a procedure for designing a feasible localized SLC is beyond the scope of this paper: we refer the reader to our recent paper [26], where we show how to co-design an actuator architecture and corresponding feasible localized SLC. We generalize these methods to the output feedback setting in [18], where we show that actuation and sensing architectures, as well as feasible localized SLCs, can be co-designed using convex programming.

6) *Intersections of SLCs and spatiotemporal constraints:* A benefit of SLCs is that several such constraints can be imposed on the system response at once. Further, as convex sets are closed under intersection, convex SLCs are also closed under intersection. To illustrate the usefulness of this property, consider the intersection of a QI subspace SLC, a FIR SLC and a localized SLC. The resulting SLC can be interpreted as enforcing a spatiotemporal constraint on the system response and its corresponding controller.

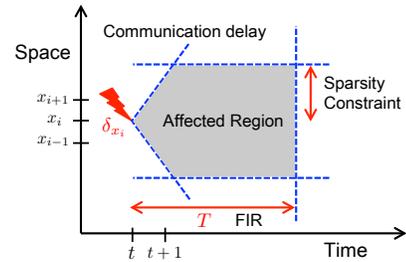


Fig. 3. Space time diagram for a single disturbance striking the chain system described in §IV-6.

Figure 3 shows a diagram of a chain of scalar subsystems responding to a particular disturbance $(\delta_x)_i$. In this figure, the vertical axis denotes the spatial coordinate of a state in the chain, and the horizontal axis denotes time: hence we refer to this figure as a space-time diagram. Depicted are the three components of the spatiotemporal constraint, namely the communication delay imposed on the controller via the QI subspace SLC, the deadbeat response of the system to the disturbance imposed by the FIR SLC, and the localized region affected by the disturbance $(\delta_x)_i$ imposed by the localized SLC. When the effect of each disturbance $(\delta_x)_i$ can be localized within such a spatiotemporal SLC, the system is said to be *localizable* (c.f., [22], [25] and references therein).

V. SYSTEM LEVEL SYNTHESIS

In this section we formulate the SLS problem and show that it can parameterize the largest known class of constrained optimal control problems that admit a convex formulation, with all possible structured optimal control problems of the form (3) that admit a convex representation in the Youla domain as a special case.

1) *General Formulation*: Let $g(\cdot)$ be a functional capturing a desired measure of the performance of the system (as described in §IV-3), and let \mathcal{S} be a SLC. We then pose the SLS problem as

$$\begin{aligned} & \underset{\{\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L}\}}{\text{minimize}} && g(\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L}) \\ & \text{subject to} && (17\text{a}) - (17\text{c}), \begin{bmatrix} \mathbf{R} & \mathbf{N} \\ \mathbf{M} & \mathbf{L} \end{bmatrix} \in \mathcal{S}. \end{aligned} \quad (25)$$

For g a convex functional and \mathcal{S} a convex set, the resulting SLS problem is a convex optimization problem.

2) *Distributed optimal control*: Recall that the objective function in (3) is specified by a suitably chosen system norm measuring the size of the closed loop transfer matrix from the external disturbance \mathbf{w} to the regulated output $\bar{\mathbf{z}}$. Therefore it suffices to select the objective functional g to be as described in equation (23), and to select the SLC constraint set \mathcal{S} as described in equation (20). The resulting SLS problem

$$\begin{aligned} & \underset{\{\mathbf{R}, \mathbf{M}, \mathbf{N}, \mathbf{L}\}}{\text{minimize}} && \left\| \begin{bmatrix} C_1 & D_{12} \end{bmatrix} \begin{bmatrix} \mathbf{R} & \mathbf{N} \\ \mathbf{M} & \mathbf{L} \end{bmatrix} \begin{bmatrix} B_1 \\ D_{21} \end{bmatrix} + D_{11} \right\| \\ & \text{subject to} && (17\text{a}) - (17\text{c}), \mathbf{L} \in \mathcal{C} \end{aligned}$$

is then equivalent to the distributed optimal control problem (3) when the subspace \mathcal{C} is QI with respect to the plant \mathbf{P}_{22} . Thus all distributed optimal control problems of the form (3) that can be formulated as convex optimization problems in the Youla domain are special cases of the SLS problem (25).

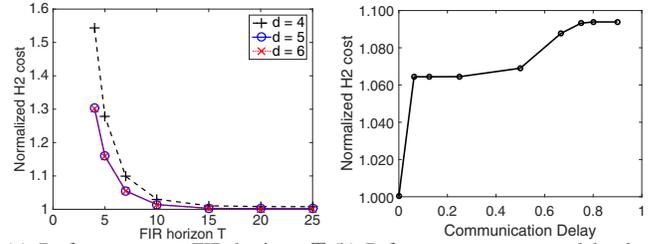
3) *Localized LQG control*: In [22], [25] we posed and solved a localized LQG optimal control problem. It can be recovered as a special case of the SLS problem (25) by selecting the functional g to be of the form (23) (with the system norm $\|\cdot\|$ chosen to be the \mathcal{H}_2 norm), and selecting the constraint set \mathcal{S} to be a spatiotemporal SLC.

In [22], we observe that a state-feedback SLS problem⁷ with a localized SLC can be decomposed into a set of independent sub-problems that solve for the columns \mathbf{R}_i and \mathbf{M}_i of the transfer matrices \mathbf{R} and \mathbf{M} – as these problems are independent, they can be solved in parallel. Further, the sparsity constraint \mathcal{L} restricts each sub-problem to a local subset of the system model and states, as specified by the nonzero components of the corresponding column of the transfer matrices \mathbf{R} and \mathbf{M} (e.g., as was described in Example 1), allowing each of these sub-problems to be expressed in terms of optimization variables (and corresponding sub-matrices of the state-space realization (17)) that are of significantly smaller dimension than the global system response $\{\mathbf{R}, \mathbf{M}\}$. Thus for a given feasible spatiotemporal SLC, the localized SLS problem can be solved for arbitrarily

⁷We extend these results to the output-feedback setting in [25].

large-scale systems, assuming that each sub-controller can solve its corresponding sub-problem in parallel.⁸ In [18], we extend these concepts to the system level approach, and show how appropriate notions of separability for SLCs can be defined that allow for optimal controllers to be synthesized and implemented with order constant complexity (assuming parallel computation is available) relative to the global system size. This approach allows to, for example, compute an optimal controller for a 12800 state power system example (see §VI-2 for a description of the dynamics) on a laptop.

VI. SYSTEM LEVEL TRADEOFFS



(a) Performance vs. FIR horizon T (b) Performance vs. comm. delay between neighboring controllers

Fig. 4. Tradeoffs between performance and comm. speed, localized region size, and FIR horizon (normalized with respect to the optimal centralized \mathcal{H}_2 cost). In Fig 4(b), when the delay relative to the plant is 0, this corresponds to a centralized controller; for any relative delay of 1 or larger, the system cannot be localized.

1) *Performance vs. implementation complexity*: We use a 100 node bi-directional chain with scalar subsystems for this numerical study. We place 40 actuators in the chain network, with actuator location specified by $i = 5j - 4$ and $5j$ for $j = 1, \dots, 20$. We compute optimal controllers with respect to the \mathcal{H}_2 norm, for measured output $\bar{\mathbf{z}}[t] = [x[t]^T u[t]^T]^T$ (see § VI.A-B of [17] for specific parameter values used). Throughout, we impose a spatiotemporal SLC as described in §IV-6 and study the effects of choosing different sized localized regions, parameterized by d -hops of the physical plant topology (c.f., [25]), the length of the FIR horizon and the communication speed of the controller. As shown in Figures 4(a) and 4(b), appropriate choices of these parameters lead to no degradation in performance with respect to a centralized optimal controller, while leading to significant improvements in synthesis and implementation complexity.

2) *Performance vs. architectural complexity*: For this numerical study, we use a power system model with (undirected) topology and subsystem interaction dynamics depicted in Fig. 5. The dynamics of each subsystem are specified by a discretization of the swing equations (see § V.A of [18] for specific parameter values used).

We formulate a \mathcal{H}_2 optimal control problem for which the norm of the optimal system response achieved by a proper centralized controller is 13.32. Here, we assumed that each subsystem in the power network has a phase measurement unit (PMU), a frequency sensor, and a controllable load. In practice, the installation of these sensors and actuators is

⁸We show how to co-design an actuation architecture and feasible spatiotemporal constraint in [26].

expensive: we therefore study the degradation in performance incurred as we remove sensors and actuators from the system.

A challenging problem is to determine the location of these sensors and actuators: to do so, we use techniques developed in the regularization for design framework [27] to penalize the use of sensors and actuators. Thus we impose both spatiotemporal and architectural complexity SLCs on the corresponding SLS problem. We set the localized region to be 4-hops (as defined by the system topology), the FIR horizon to be $T = 30$, and assume sub-controllers can communicate twice as quickly as dynamics propagate through the system. We assume that the relative prices of each frequency sensor, PMU, and controllable load are 1, 100, and 300, respectively.

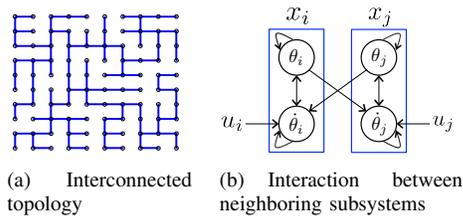


Fig. 5. Power system example interaction graph.

Using the methods described in § IV of [18], we are able to reduce the number of controllable loads and PMUs used by 43% and 46%, respectively (no frequency sensors are removed due to the chosen relative pricing) while still satisfying the spatiotemporal SLC. Further, the \mathcal{H}_2 cost of the optimal system response is given by 17.86, which is a 10% and 34% degradation in performance from that of the corresponding sparsely and fully sensed/actuated centralized controller, respectively.

3) *Mixed Objective Optimal Control*: Finally, we explore the tradeoff between average system performance (the \mathcal{H}_2 norm of $w \rightarrow \bar{z}$) and the robustness of the controller to worst-case signal (the \mathcal{L}_1 norm of $w \rightarrow \bar{z}$) on the mesh example shown in Fig. 5(a). Fig. 6 shows that a tradeoff between these two metrics exist, and that it can be efficiently explored using the system level approach. Such a tradeoff may be relevant in computationally constrained settings such as neural computations and embedded systems.

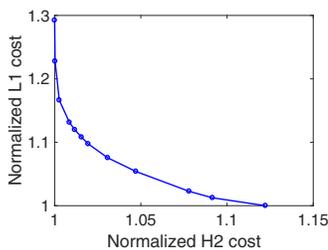


Fig. 6. Tradeoff between \mathcal{H}_2 norm and \mathcal{L}_1 norm

REFERENCES

- [1] D. C. Youla, H. A. Jabr, and J. J. B. Jr., “Modern wiener-hopf design of optimal controllers-part ii: The multivariable case,” *Aut. Control, IEEE Trans. on*, vol. 21, no. 3, pp. 319–338, 1976.
- [2] J. C. Doyle, K. Glover, P. P. Khargonekar, and B. A. Francis, “State-space solutions to standard H_2 and H_∞ control problems,” *IEEE Trans. on Aut. Control*, vol. 34, no. 8, pp. 831–847, Aug 1989.

- [3] Y.-C. Ho and K.-C. Chu, “Team decision theory and information structures in optimal control problems—part i,” *Aut. Control, IEEE Trans. on*, vol. 17, no. 1, pp. 15–22, 1972.
- [4] A. Mahajan, N. Martins, M. Rotkowitz, and S. Yuksel, “Information structures in optimal decentralized control,” in *Decision and Control, 2012 IEEE 51st Annual Conf. on*, 2012, pp. 1291–1306.
- [5] M. Rotkowitz and S. Lall, “A characterization of convex problems in decentralized control,” *Aut. Control, IEEE Trans. on*, vol. 51, no. 2, pp. 274–286, 2006.
- [6] 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.
- [7] A. Nayyar, A. Mahajan, and D. Teneketzis, “Decentralized stochastic control with partial history sharing: A common information approach,” *IEEE Trans. on Aut. Control*, vol. 58, no. 7, pp. 1644–1658, July 2013.
- [8] H. S. Witsenhausen, “A counterexample in stochastic optimum control,” *SIAM Journal of Control*, vol. 6, no. 1, 1968.
- [9] G. E. Dullerud and R. D’Andrea, “Distributed control of heterogeneous systems,” *Aut. Control, IEEE Trans. on*, vol. 49, no. 12, pp. 2113–2128, 2004.
- [10] 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>
- [11] L. Lessard and S. Lall, “Convexity of decentralized controller synthesis,” *IEEE Trans. on Aut. Control, To appear*, 2016. [Online]. Available: <http://arxiv.org/pdf/1305.5859v2.pdf>
- [12] —, “Optimal controller synthesis for the decentralized two-player problem with output feedback,” in *2012 IEEE American Control Conf.*, 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, 2010 49th IEEE Conf. 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 Conf.*, June 2013.
- [15] C. W. Scherer, “Structured \mathcal{H}_∞ -optimal control for nested interconnections: A state-space solution,” *Systems and Control Letters*, vol. 62, pp. 1105–1113, 2013.
- [16] N. Matni, “Distributed control subject to delays satisfying an \mathcal{H}_∞ norm bound,” in *2014 53rd IEEE Conf. on Decision and Control*, 2014. [Online]. Available: <http://arxiv.org/pdf/1402.1559.pdf>
- [17] Y.-S. Wang, N. Matni, and J. C. Doyle, “A system level approach to controller synthesis,” *IEEE Trans. on Aut. Control, submitted*, 2016. [Online]. Available: <https://arxiv.org/abs/1610.04815>
- [18] —, “Separable and localized system level synthesis for large-scale systems,” *IEEE Trans. on Aut. Control, submitted*, 2017. [Online]. Available: <https://arxiv.org/abs/1701.05880>
- [19] Y.-S. Wang, “A system level approach to optimal controller design for large-scale distributed systems,” Ph.D. dissertation, California Institute of Technology, 2016.
- [20] L. Lessard and S. Lall, “Quadratic invariance is necessary and sufficient for convexity,” in *Proceedings of the 2011 American Control Conf.*, June 2011, pp. 5360–5362.
- [21] Ş. Sabāu and N. C. Martins, “Youla-like parametrizations subject to QI subspace constraints,” *Aut. Control, IEEE Trans. on*, vol. 59, no. 6, pp. 1411–1422, 2014.
- [22] Y.-S. Wang, N. Matni, and J. C. Doyle, “Localized LQR optimal control,” in *2014 53rd IEEE Conf. on Decision and Control*, 2014.
- [23] K. Zhou, J. C. Doyle, and K. Glover, *Robust and optimal control*. Prentice Hall New Jersey, 1996.
- [24] A. Lamperski and J. C. Doyle, “The \mathcal{H}_2 control problem for quadratically invariant systems with delays,” *Aut. Control, IEEE Trans. on*, vol. 60, no. 7, pp. 1945–1950, 2015.
- [25] Y.-S. Wang and N. Matni, “Localized LQG optimal control for large-scale systems,” in *2016 IEEE American Control Conf.*, 2016.
- [26] Y.-S. Wang, N. Matni, and J. C. Doyle, “Localized LQR control with actuator regularization,” in *2016 IEEE American Control Conf.*, 2016.
- [27] N. Matni and V. Chandrasekaran, “Regularization for design,” *IEEE Trans. on Aut. Control*, vol. PP, no. 99, pp. 1–1, 2016.

Acknowledgements: Thanks to funding from AFOSR and NSF and gifts from Huawei and Google.