

Regularization for Design

Nikolai Matni and Venkat Chandrasekaran *

July 29, 2015

Abstract

When designing controllers for large-scale systems, the architectural aspects of the controller such as the placement of actuators, sensors, and the communication links between them can no longer be taken as given. The task of designing this architecture is now as important as the design of the control laws themselves. By interpreting controller synthesis (in a model matching setup) as the solution of a particular linear inverse problem, we view the challenge of obtaining a controller with a desired architecture as one of finding a structured solution to an inverse problem. Building on this conceptual connection, we formulate and analyze a framework called *Regularization for Design (RFD)*, in which we augment the variational formulations of controller synthesis problems with convex penalty functions that induce a desired controller architecture. The resulting regularized formulations are convex optimization problems that can be solved efficiently; these convex programs provide a unified computationally tractable approach for the simultaneous co-design of a structured optimal controller and the actuation, sensing and communication architecture required to implement it. Further, these problems are natural control-theoretic analogs of prominent approaches such as the Lasso, the Group Lasso, the Elastic Net, and others that are employed in statistical modeling. In analogy to that literature, we show that our approach identifies optimally structured controllers under a suitable condition on a “signal-to-noise” type ratio.

1 Introduction

As we move into the era of large-scale systems such as the smart-grid, software defined networking and automated highways, the design of control systems is becoming increasingly more challenging. Designing the controller *architecture* – placing sensors and actuators as

*N. Matni is with the Department of Control and Dynamical Systems, California Institute of Technology, Pasadena, CA. nmatni@caltech.edu. V. Chandrasekaran is with the Departments of Computational and Mathematical Sciences and of Electrical Engineering, California Institute of Technology, Pasadena, CA. venkatc@caltech.edu. N.M. was supported in part by the NSF, AFOSR, ARPA-E, and the Institute for Collaborative Biotechnologies through grant W911NF-09-0001 from the U.S. Army Research Office. V.C. was supported in part by NSF Career award CCF-1350590 and by Air Force Office of Scientific Research grant FA9550-14-1-0098. The content does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred. A preliminary version of this work [1] has appeared at the 53rd Annual Conference on Decision and Control in December, 2014.

well as the communication links between them – is now as important as the traditional design of the control laws themselves.

A conceptually useful viewpoint in the design of controller architectures is to consider complicated systems as being composed of multiple simpler *atomic* subsystems. For example, if the task is to design the actuation architecture of a controller, a natural atomic element is a controller with a single actuator – it is then clear that a general architecture can be built out of such atoms. In general, controllers with a dense actuation, sensing and communication architecture (i.e., systems that consist of many atomic subsystems) achieve better closed loop performance in comparison with those with sparse architectures (i.e., systems composed of a small number of atomic subsystems). However, as these architectural resources translate into actual hardware requirements, it is desirable from both a maintenance and a cost perspective that we minimize the total number of atomic elements used. Hence, the problem of controller architecture/control law co-design is one of jointly optimizing an appropriately defined structural measure of the controller and its closed loop performance by trading off between these two competing metrics in a principled manner. In other words, we seek an approximation of a given optimal controller by one that utilizes fewer atomic elements without a significant loss in performance. This goal has parallels with the approximation theory literature in which one seeks approximations of complicated functions as combinations of elements from a simpler class of functions such as the Fourier basis or a wavelet basis [2].

In an appropriate parameterization, pure controller synthesis methods in a model matching framework can be interpreted as techniques for solving a particular linear inverse problem in which one is given an open loop response of a system Y and a linear map \mathfrak{L} from the controller to the closed loop response, and one seeks a controller U such that $Y - \mathfrak{L}(U) \approx 0$ (as measured in a suitable performance metric). From this perspective, our objective in joint controller architecture/control law co-design is to obtain *structured* solutions to the linear inverse problem underlying controller synthesis. Such *structured linear inverse problems* (SLIP) are of interest in diverse applications across applied mathematics – for instance, computing sparse solutions to linear inverse problems or computing low-rank solutions to systems of linear matrix equations arise prominently in many contexts in signal processing and in statistics [3–6].

In these problem domains, minimizing the ℓ_1 norm subject to constraints described by the specified equations is useful for obtaining sparse solutions [3, 4], and similarly, nuclear norm minimization is useful for obtaining low-rank solutions to linear matrix equations [5, 6]. These ideas were extended in [7], where the authors describe a generic convex programming approach – based on minimizing an appropriate *atomic norm* [2]– for inducing a desired type of structure in solutions to linear inverse problems. Motivated by these developments, our approach to the problem of joint architecture/control law co-design is to augment variational formulations of controller synthesis methods with suitable convex regularization functions. The role of these regularizers is to penalize controllers with more complex architectures in favor of those with less complex ones, thus inducing controllers with a simpler architecture. We call this framework *Regularization for Design* (RFD).

Related Work: Regularization techniques based on ℓ_1 norms and, more generally, atomic norms have already been employed extensively in system identification, e.g., to identify systems of small Hankel order (cf. [8–10]), and in linear regression based methods [11]. Although the resulting solutions yield structured systems, they typically do not have a di-

rect interpretation in terms of the architecture of a control system (i.e., actuators, sensors and the communication links between them). The use of regularization explicitly for the purpose of designing the architecture of a controller can also be found in the literature. Examples include the use of ℓ_1 regularization to design sparse structures in \mathcal{H}_2 static state feedback gains [12], treatment therapies [13], and synchronization topologies [14]; the use of group norm penalties to design actuation/sensing schemes [1, 15]; and the use of a specialized atomic norm to design communication delay constraints that are well-suited to \mathcal{H}_2 distributed optimal control [16, 17]. Although these methods provide an algorithmic approach for designing controllers with a desired architecture in certain specialized settings, they do not enjoy the same theoretical support that regularization techniques for structured inverse problems enjoy in other settings [2–6]. Methods based on greedy algorithms have also been developed to identify minimal architectures that guarantee the structural observability and controllability of a system [18].

Our Contributions: This paper presents novel computational and theoretical contributions to the area of optimal controller architecture/control law co-design. From a computational perspective, we propose a general RFD framework that is applicable in a much broader range of settings than the previous approaches mentioned above. We restrict ourselves to problems for which the linear optimal structured controller is specified as the solution to a convex optimization problem [19–21]. As a result, RFD optimization problems with convex regularization functions for inducing a desired architecture are convex programs. Specifically, (i) we provide a catalog of atomic norms useful for control architecture design. In particular, in addition to known penalties for actuator, sensor and communication design, we provide novel penalties for *simultaneous* actuator, sensor and/or communication design; (ii) we describe a unifying framework for RFD that encompasses state and output feedback problems in centralized and distributed settings, and in which any subset of actuation, sensing, and/or communication architectures are co-designed; and (iii) we present a two-step algorithm that first identifies the controller architecture via a finite-dimensional convex RFD optimization problem, and then solves for the potentially infinite dimensional linear optimal controller restricted to the designed architecture using methods from classical and distributed optimal control [22–28].

To provide theoretical support for our proposed computational framework, we make explicit links between RFD optimization problems and the use of convex optimization based approaches for structured inference problems. We build on these links to analyze the properties of the structured controllers generated by RFD synthesis methods, which leads to conditions under which RFD methods succeed in identifying optimally-structured controllers. Our analysis and results are natural control-theoretic analogs of similar results in the statistical inference literature. Specifically, (i) we show that finite-horizon finite-order convex approximations of an RFD optimization problem can recover the structure of an underlying infinite dimensional optimal controller; (ii) we define control-theoretic analogs of identifiability conditions and signal-to-noise ratios (SNRs), and we provide sufficient conditions based on these for a controller architecture to be identified by RFD. In particular, we show that controllers that maximize this SNR-like quantity are more easily recovered via RFD than those that do not, and (iii) we provide a concrete example of a system satisfying the above identifiability and SNR conditions. As far as we are aware, this is the first example in the literature of a system for which it is shown that convex optimization provably recovers the

actuation architecture of an underlying optimally structured controller.

Paper Organization: In §2, we define notation and discuss the relevant concepts from operator theory. This paper is then organized in a modular fashion: §3-5 focus on the computational aspects of controller architecture design, whereas §6 and §7 focus on conditions for optimal architecture recovery. Specifically, in §3, we introduce the RFD framework as a natural blend of controller synthesis methods and regularization techniques. In §4, we focus on RFD problems with an \mathcal{H}_2 performance metric and an atomic norm penalty; we present a catalog of atomic norms that are useful for controller architecture design, and make connections to the structured inference literature. The computational component of the paper concludes in §5, in which we formally describe the two-step RFD procedure, and we apply the RFD framework to a simultaneous actuator, sensor and communication design problem. In §6, we shift our focus to analyzing the theoretical properties of the RFD procedure: we make connections between structured controller design and structured inference problems by framing both tasks as finding structured solutions to linear inverse problems, and we leverage these connections to describe interpretable sufficient conditions for the success of a finite-dimensional RFD optimization problem. In §7, we provide a case study to further illustrate the applicability of these results.

2 Preliminaries & Notation

We use standard definitions of the Hardy spaces \mathcal{H}_2 and \mathcal{H}_∞ . As is standard, we denote the restrictions of \mathcal{H}_2 and \mathcal{H}_∞ to the space of real proper transfer matrices \mathcal{R}_p by \mathcal{RH}_2 and \mathcal{RH}_∞ . As we work in discrete time, the two spaces are equal, and as a matter of convention we refer to this space as \mathcal{RH}_∞ . We refer the reader to [22] for a review of this standard material. Let $\mathcal{RH}_\infty^{\leq t}$ denote the subspace of \mathcal{RH}_∞ composed of finite impulse response (FIR) transfer matrices of length t , i.e., $\mathcal{RH}_\infty^{\leq t} := \{G \in \mathcal{RH}_\infty \mid G = \sum_{i=0}^t \frac{1}{z^i} G^{(i)}\}$. We denote the projection of an element $G \in \mathcal{R}_p$ onto the subspace $\mathcal{RH}_\infty^{\leq t}$ by $G^{\leq t}$. Unless required for the discussion, we do not explicitly denote dimensions and we assume that all vectors, operators and spaces are of compatible dimension throughout.

We denote elements of \mathcal{R}_p with upper case Latin letters, and temporal indices and horizons by lower case Latin letters. Linear maps from \mathcal{R}_p to \mathcal{R}_p are denoted by upper case Fraktur letters such as \mathfrak{L} . For such a linear map, we denote the i^{th} impulse response element of \mathfrak{L} by $\mathfrak{L}^{(i)}$. We further use $\mathfrak{L}^{\leq t}$ to denote the restriction of the range of \mathfrak{L} to $\mathcal{RH}_\infty^{\leq t}$, and $\mathfrak{L}^{\leq t, v}$ to denote the restriction of $\mathfrak{L}^{\leq t}$ to the domain $\mathcal{RH}_\infty^{\leq v}$. Thus $\mathfrak{L}^{\leq t, v}$ is a map from $\mathcal{RH}_\infty^{\leq v}$ to $\mathcal{RH}_\infty^{\leq t}$. In particular, if \mathfrak{L} is represented as a semi-infinite lower block triangular matrix, then $\mathfrak{L}^{\leq t, v}$ corresponds to the t by v block row by block column sub matrix $(\mathfrak{L})_{ij}$, $i = 1, \dots, t$, $j = 1, \dots, v$.

Sets are denoted by upper case script letters, such as \mathcal{S} , whereas subspaces of an inner product space are denoted by upper case calligraphic letters, such as \mathcal{S} . The restriction of a linear map \mathfrak{L} to a subspace $\mathcal{S} \in \mathcal{RH}_\infty$ is denoted by $\mathfrak{L}_{\mathcal{S}}$; similarly, the projection of an operator $G \in \mathcal{R}_p$ onto a subspace $\mathcal{A} \subset \mathcal{R}_p$ is denoted by $G_{\mathcal{A}}$. We denote the adjoint of a linear map \mathfrak{L} by \mathfrak{L}^\dagger . The most complicated expression that we use is of the form $[\mathfrak{L}_{\mathcal{S}}^{\leq t, v}]^\dagger$: this denotes the adjoint of the map $\mathfrak{L}^{\leq t}$ restricted to $\mathcal{RH}_\infty^{\leq v} \cap \mathcal{S}$.

We denote the n -dimensional identity matrix and down-shift matrices by I_n and Z_n , respectively. In particular, Z_n is a matrix with all ones along its first sub-diagonal and zero elsewhere. We use e_i to denote a standard basis element in \mathbb{R}^n , and E_{ij} to denote the matrix with $(i, j)^{\text{th}}$ element set to 1 and all others set to 0.

3 RFD as Structured Approximation

Under standard assumptions [22], traditional controller synthesis methods within the framework of *model matching* can be framed as linear inverse problems of the form

$$\underset{U \in \mathcal{RH}_\infty}{\text{minimize}} \quad \Psi(U; Y, \mathfrak{L}) \quad (1)$$

where Y is the open loop response of the system, U is the Youla parameter, \mathfrak{L} is a suitably defined linear map from the Youla parameter U to the closed loop response, and $\Psi(\cdot; Y, \mathfrak{L})$ is a performance metric that measures the size of the closed loop response (i.e., the size of the deviation between Y and $\mathfrak{L}(U)$), such as the \mathcal{H}_2 or \mathcal{H}_∞ norm. We make this connection clear in the following subsection and we also recall how to incorporate *quadratically invariant* [19] distributed constraints on the controller into this framework.

Remark 1 *We use this non-standard notation to facilitate comparisons with the structured inference literature. In particular, this notation emphasizes that the optimal linear controller synthesis task can be viewed as one of solving a linear inverse problem with “data” specified by the open loop response Y of the system and the map \mathfrak{L} from the Youla parameter to the closed loop response.*

3.1 Convex Model Matching




Figure 1: A diagram of the generalized plant defined in (2).

In order to discuss a broad range of model matching problems, we introduce the *generalized plant*, a standard tool in robust and optimal control [22]. In particular, consider the system described by

$$P = \begin{bmatrix} P_{11} & P_{12} \\ P_{21} & P_{22} \end{bmatrix} = \left[\begin{array}{c|cc} A & B_1 & B_2 \\ \hline C_1 & 0 & D_{12} \\ C_2 & D_{21} & 0 \end{array} \right] \quad (2)$$

where $P_{ij} = C_i(zI - A)^{-1}B_j + D_{ij}$. As illustrated in Figure 1, this system describes the four transfer matrices from the disturbance and control inputs w and u , respectively, to the controlled and measured outputs y and m , respectively. We make the (slightly modified) standard orthogonality assumptions that

$$\begin{aligned} D_{12} \begin{bmatrix} C_1^\top & D_{12}^\top \end{bmatrix} &= \begin{bmatrix} 0 & \rho_u I \end{bmatrix} \\ D_{21}^\top \begin{bmatrix} B_1 & D_{21} \end{bmatrix} &= \begin{bmatrix} 0 & \rho_w I \end{bmatrix} \end{aligned} \quad (3)$$

for some $\rho_u, \rho_w \geq 0$. At times we separate the state component of the open and closed loop responses from the components of these transfer matrices that measure control effort. To that end, we define the *state component* of an element X to be the projection of X onto the range of C_1 .

Letting $u(z) = K(z)m(z)$ for a causal linear controller $K \in \mathcal{R}_p$, the closed loop map from the disturbance w to the controlled output y is given by the linear fractional transform $P_{11} - P_{12}K(I - P_{22}K)^{-1}P_{21}$. A typical optimal control problem in this framework is then formulated as

$$\begin{aligned} \underset{K \in \mathcal{R}_p}{\text{minimize}} \quad & \|P_{11} - P_{12}K(I - P_{22}K)^{-1}P_{21}\| \\ \text{s.t.} \quad & K(I - P_{22}K)^{-1} \in \mathcal{RH}_\infty \end{aligned} \quad (4)$$

where $\|\cdot\|$ is a suitable norm, and the constraint ensures internal stability of the closed loop system [22]. Notice however that the optimal control problem (4) is non-convex as stated.

When the open-loop plant is stable (we remark on the unstable case at the end of the subsection), a standard and general approach to solving the optimal control problem (4) is to convert it into a *model matching problem* through the Youla change of variables $U := K(I - P_{22}K)^{-1}$; the optimal controller K can then be recovered via $K = (I + UP_{22})^{-1}U$. The resulting convex optimization problem is then given by

$$\underset{U \in \mathcal{RH}_\infty}{\text{minimize}} \quad \|P_{11} - P_{12}UP_{21}\|, \quad (5)$$

which is of the form of the linear inverse problem (1) if we take $Y := P_{11}$, $\mathfrak{L} = P_{12} \otimes P_{21}$ (where $(P_{12} \otimes P_{21})(U) := P_{12}UP_{21}$) and $\Psi(U; Y, \mathfrak{L}) := \|Y - \mathfrak{L}(U)\|$.

We also often want to impose a *distributed constraint* on the controller K by requiring K to lie in some subspace \mathcal{S} , which specifies information exchange constraints between the sensors and actuators of the controller. It is known that a necessary and sufficient condition for such a distributed constraint to be invariant under the Youla change of variables is that it be *quadratically invariant* with respect to P_{22} [19–21].

Definition 1 (Quadratic Invariance) *A subspace \mathcal{S} is quadratically invariant (QI) with respect to P_{22} if $KP_{22}K \in \mathcal{S} \forall K \in \mathcal{S}$.*

In particular, when a subspace \mathcal{S} is QI with respect to P_{22} , we have that $K \in \mathcal{S}$ if and only if $K(I - P_{22}K)^{-1} \in \mathcal{S}$, allowing us to convert the general optimal control problem (1) with the additional constraint that $K \in \mathcal{S}$ to the following convex optimization problem

$$\underset{U \in \mathcal{RH}_\infty}{\text{minimize}} \quad \Psi(U; Y, \mathfrak{L}) \quad \text{s.t.} \quad U \in \mathcal{S}. \quad (6)$$

This optimization problem is again precisely of the form of the linear inverse problem (1) save for the addition of the subspace constraint $U \in \mathcal{S}$. This framework is fairly general in that it allows for a unified treatment of all structured optimal control problems in which the linear optimal structured controller can be computed via convex optimization [19, 21]. These include state and output feedback problems in either centralized or QI distributed settings. Further, if the optimal control problem is centralized with respect to the \mathcal{H}_2 , \mathcal{H}_∞ or \mathcal{L}_1 metrics, or is QI distributed with a finite horizon \mathcal{H}_2 cost, the linear optimal control is globally optimal [21, 22, 29, 30].

Remark 2 (Extension to Unstable Plants) *The above discussion extends to unstable plants through the use of an appropriate structure preserving Youla-Kucera parameterization built around arbitrary coprime factorizations, which are always available. See [27, 31, 32] for examples of such parameterizations, and [17] for an example of using such a parameterization with a structure inducing penalty.*

3.2 Architecture Design through Structured Solutions

We seek a modification of the optimal controller synthesis procedure to design the controller's architecture. We reiterate that by the architecture of a controller, we mean the actuators, sensors, and communication links between them. In particular, we view the controller K as a map from all potential sensors to all potential actuators, using all potential communication links between these actuators and sensors. The architectural design task is that of selecting which actuators, sensors and communication links need to be used to achieve a certain performance level. This task is naturally viewed as one of finding a structured approximation of the optimal controller (4) that utilizes all of the available architectural choices.

The components of the controller architecture being designed determine the type of structured approximation that we attempt to identify. In particular, each nonzero row of $K(z)$ corresponds to an actuator used by the controller, and likewise, each nonzero column corresponds to a sensor employed by the controller. Further sparsity patterns present within rows/columns of the power series elements $K^{(t)}$ of $K(z)$ can be interpreted as information exchange constraints imposed by an underlying communication network between the sensors and actuators. With these observations, it is clear that specific sparsity patterns in K have direct interpretations in terms of the architectural components of the controller: nonzero rows correspond to actuators, nonzero columns correspond to sensors, and additional sparsity structure corresponds to communication constraints.

Although we seek to identify a suitably structured controller K , for the computational reasons described in §3.1 it is preferable to solve a problem in terms of the Youla parameter U as this parameterization leads to the convex optimization problem (6). Therefore, in the following definition RFD problems are defined as a regularized version of the model matching problem (6) with a penalty function added to the objective to induce suitable structure, rather than as a modification of the controller synthesis problem (4). In the sequel, we justify that for architectural design problems of interest, the structure underlying the controller K is equivalent to the structure underlying the Youla parameter U ; to that end, we show that in the case of actuator, sensor, and/or QI communication topology de-

sign,¹ the structure of the Youla parameter U corresponds to the structure of the controller K .

Definition 2 Let $U, Y \in \mathcal{RH}_\infty$, and $\mathfrak{L} : \mathcal{RH}_\infty \rightarrow \mathcal{RH}_\infty$ be of compatible dimension. The optimization

$$\underset{U \in \mathcal{RH}_\infty}{\text{minimize}} \quad \Psi(U; Y, \mathfrak{L}) + 2\lambda\Omega(U) \quad \text{s.t. } U \in \mathcal{S} \quad (7)$$

is called a RFD optimization problem with cost function $\Psi(\cdot; Y, \mathfrak{L})$ and penalty $\Omega(\cdot)$.

Remark 3 (Static or Dynamic) With the exception of the centralized state-feedback setting, it is known that the optimal linear controller is dynamic [22], and therefore restricting our analysis to dynamic controllers is natural. In the centralized setting (given the equivalence between static and dynamic state-feedback), once an appropriate actuation architecture² is identified, traditional methods can then be used to solve for a static state-feedback controller restricted to that architecture. Further as we show in §6, the dynamic controller synthesis approach is amenable to analysis that guarantees optimal structure recovery.

We discuss natural costs $\Psi(\cdot; Y, \mathfrak{L})$ and penalties $\Omega(\cdot)$ in §4, and we focus now on justifying why we can perform the structural design on the Youla parameter U rather than the controller K .

Actuator/Sensor Design Recall that the actuators (sensors) that a controller K uses are identified by the nonzero rows (columns) of K : the actuator (sensor) design problem therefore corresponds to finding a controller K that achieves a good closed loop response and that is sparse row-wise (column-wise). This corresponds exactly to finding a row (column) sparse solution U to the RFD optimization problem (7). This is true because any subspace \mathcal{D} that is defined solely in terms of row (column) sparsity is QI with respect to any P_{22} . In particular, it is easily verified that if $K \in \mathcal{D}$, then right (left) multiplication leaves \mathcal{D} invariant, i.e., $KX \in \mathcal{D}$ ($XK \in \mathcal{D}$) for all compatible X . It then follows from Definition 1 that \mathcal{D} is QI with respect to any plant P_{22} . We can extend this analysis to incorporate additional QI distributed constraints \mathcal{S} by leveraging the results in [20]: in particular, if \mathcal{S} is QI with respect to P_{22} , then so is $\mathcal{D} \cap \mathcal{S}$.³

Joint Actuator and Sensor Design By virtue of the previous discussion joint actuator and sensor design corresponds to finding a controller K that is simultaneously sparse row-wise and column-wise. It follows immediately from the previous discussion that this corresponds exactly to finding a simultaneously row and column sparse solution U to the

¹We restrict ourselves to communications delays that satisfy the triangle inequality defined in [20]. This assumption implies that information exchanged between sensors and actuators is transmitted along shortest delay paths in the underlying communication graph.

²For a system to be centralized, there can be no communication constraints, and for it to be state-feedback, there can be no sensing constraints, leaving only actuation architecture design as a possible RFD task.

³In particular, since removing actuators does not change the communication delays between the remaining actuators and sensors, if the delay based conditions in [20] hold when all actuators (sensors) are present they also hold with any subset of them being present.

RFD optimization problem (7). In particular, any subspace \mathcal{D} defined solely in terms of row and column sparsity is QI with respect to any plant P_{22} , we can incorporate additional QI distributed constraints \mathcal{S} by leveraging the results in [20].

Communication Design In an analogous manner to the above, one can also associate subspaces to structures corresponding to suitable information exchange constraints that a distributed controller must satisfy. Recall in particular that the information exchange constraints between the sensors and actuators of a controller K are identified by the sparsity structure found within the nonzero rows and columns of K . In [17], the first author showed that a specific type of sparsity structure in K corresponds exactly to sensors and actuators exchanging information according to an underlying communication graph. In particular, given a communication graph between sensors and actuators with adjacency matrix Γ , a distributed controller K can be implemented using the graph defined by Γ if the power series elements $K^{(t)}$ of the controller satisfy $\text{supp}(K^{(t)}) \subseteq \text{supp}(\Gamma^{t-1})$.⁴ The interpretation of the support nesting condition is that the delay from sensor j to actuator i is given by the length of the shortest path from node j to node i in the graph defined by Γ . This support nesting condition thus defines the distributed subspace constraint \mathcal{S} in which K must lie – based on the discussion in §3.1, one can pose the distributed controller synthesis problem as a distributed model matching problem (6) if and only if \mathcal{S} is QI with respect to P_{22} . In light of this, we consider the communication design task proposed in [17]: given an initial graph with adjacency matrix Γ_{QI} that induces a QI distributed subspace constraint \mathcal{S} , what minimal set of additional edges should be added to the graph to achieve a desired performance level.⁵ It is additionally shown in [17] that any communication graph constructed in this manner results in a subspace constraint \mathcal{S} that is QI with respect to P_{22} . Therefore the structure imposed on the controller K by an underlying QI communication graph corresponds exactly to the structure imposed on the Youla parameter U .

Joint Communication, Actuator and/or Sensor Design By virtue of the previous discussion and the results of [20], combining QI communication design with actuator and/or sensor design still leads to the underlying structure of the controller K corresponding to the underlying structure of the Youla parameter U .

Thus for architecture design problems the RFD task can be performed via a model matching problem.

Example 1 *Suppose that different RFD optimization problems are solved for a system with three possible actuators and three possible sensors, resulting in the various sparsity patterns in $U(z)$ shown on the far right of Figure 2. It is easily seen by inspection that the resulting sparsity patterns are QI. In particular Figures 2a) through 2c) correspond to centralized RFD optimization problems (this can be seen from the full matrices in the center of the left hand*

⁴We assume that K is square for simplicity; cf. [17] for the general case.

⁵It is shown in [17] that under mild assumptions on the plant P_{22} , the propagation delays of P_{22} can be used to define an adjacency matrix Γ_{QI} that induces a distributed subspace constraint that is QI with respect to P_{22} .

side), and 2d) to a RFD optimization problem subject to lower triangular constraints, a special case of a nested information constraint.⁶

$$\begin{array}{l}
 \text{a)} \\
 \text{b)} \\
 \text{c)} \\
 \text{d)}
 \end{array}
 \begin{array}{l}
 \begin{bmatrix} 0 & 0 & 0 \\ * & * & * \\ * & * & * \end{bmatrix} \\
 \begin{bmatrix} 0 & * & * \\ 0 & * & * \\ 0 & * & * \end{bmatrix} \\
 \begin{bmatrix} 0 & 0 & 0 \\ 0 & * & * \\ 0 & * & * \end{bmatrix} \\
 \begin{bmatrix} * & 0 & 0 \\ 0 & 0 & 0 \\ * & * & * \end{bmatrix}
 \end{array}
 \begin{array}{l}
 \begin{bmatrix} * & * & * \\ * & * & * \\ * & * & * \end{bmatrix} \\
 \begin{bmatrix} * & * & * \\ * & * & * \\ * & * & * \end{bmatrix} \\
 \begin{bmatrix} * & * & * \\ * & * & * \\ * & * & * \end{bmatrix} \\
 \begin{bmatrix} * & * & * \\ * & * & * \\ * & * & * \end{bmatrix}
 \end{array}
 \begin{array}{l}
 \begin{bmatrix} 0 & 0 & 0 \\ * & * & * \\ * & * & * \end{bmatrix} \\
 \begin{bmatrix} 0 & * & * \\ 0 & * & * \\ 0 & * & * \end{bmatrix} \\
 \begin{bmatrix} 0 & 0 & 0 \\ 0 & * & * \\ 0 & * & * \end{bmatrix} \\
 \begin{bmatrix} * & 0 & 0 \\ 0 & 0 & 0 \\ * & * & * \end{bmatrix}
 \end{array}
 \begin{array}{l}
 \supseteq \\
 \supseteq \\
 \supseteq \\
 \supseteq
 \end{array}
 \begin{array}{l}
 \begin{bmatrix} 0 & 0 & 0 \\ * & * & * \\ * & * & * \end{bmatrix} \\
 \begin{bmatrix} 0 & * & * \\ 0 & * & * \\ 0 & * & * \end{bmatrix} \\
 \begin{bmatrix} 0 & 0 & 0 \\ 0 & * & * \\ 0 & * & * \end{bmatrix} \\
 \begin{bmatrix} * & 0 & 0 \\ 0 & 0 & 0 \\ * & * & * \end{bmatrix}
 \end{array}$$

Figure 2: Examples of QI sparsity patterns generated via a) actuator, b) sensor, and c) actuator/sensor RFD procedures without any distributed constraints, and d) actuator RFD subject to nested information constraints.

4 RFD Cost Functions and Regularizers

In this section we examine convex formulations of the RFD optimization problem (7) by restricting our attention to convex cost functions $\Psi(\cdot; Y, \mathfrak{L})$ and convex penalty functions $\Omega(\cdot)$.

4.1 Convex Cost Functions

Any suitable convex cost function $\Psi(\cdot; Y, \mathfrak{L})$ can be used in (7): traditional examples from robust and optimal control include the \mathcal{H}_2 , \mathcal{H}_∞ [22] and \mathcal{L}_1 norms [30]. We focus on the \mathcal{H}_2 norm as a performance metric because it allows us to make direct connections between the RFD optimization problem (7) and well-established methods employed in statistical inference such as ordinary least squares, Ridge Regression [33], Group Lasso [34] and Group Elastic Net [35].

We begin by introducing a specialized form of the model matching problem (6), and show how state-feedback problems with \mathcal{H}_2 performance metrics can be put into this form.

Definition 3 *Let $U, Y \in \mathcal{RH}_\infty$, and $\mathfrak{L} : \mathcal{RH}_\infty \rightarrow \mathcal{RH}_\infty$ be of compatible dimension. The optimization problem*

$$\underset{U \in \mathcal{RH}_\infty}{\text{minimize}} \quad \|Y - \mathfrak{L}(U)\|_{\mathcal{H}_2}^2 + \rho_u \|U\|_{\mathcal{H}_2}^2 \quad \text{s.t. } U \in \mathcal{S} \quad (8)$$

is the \mathcal{H}_2 optimal control problem for a suitable control penalty weight ρ_u and a distributed constraint \mathcal{S} .

⁶Nested information constraints are a well studied class of QI distributed constraints, cf. [23–25] for examples.

In this definition, Y is the state component of the open loop response, and \mathfrak{L} is the map from the Youla parameter to state component of the closed loop response. Explicitly separating the cost of the state component $\|Y - \mathfrak{L}(U)\|_{\mathcal{H}_2}^2$ of the closed loop response from the control cost $\rho_u \|U\|_{\mathcal{H}_2}^2$ of the closed loop response allows us to connect the \mathcal{H}_2 RFD optimization to several well-established methods in the inference literature. Before elaborating on some of these connections, we provide two examples of standard control problems that can be put into this form.

Example 2 (Basic LQR) Consider the basic LQR problem given by

$$\begin{aligned} & \underset{u \in \ell_2}{\text{minimize}} \quad \sum_{t=0}^{\infty} \|Cx_t\|_{\ell_2}^2 + \|Du_t\|_{\ell_2}^2 \\ & \text{s.t.} \quad x_{t+1} = Ax_t + Bu_t, \quad x_0 = \xi, \end{aligned} \quad (9)$$

and assume that $D^\top D = \rho_u I$, for some $\rho_u \geq 0$. Define $\rho = \rho_u$, $X^{(t)} = CA^t \xi$ for $t \geq 0$, $U^{(t)} = u_t$, and $\mathfrak{L}(U) = -H * U$, where $H \in \mathcal{R}_p$ with $H^{(0)} = 0$, and $H^{(t)} = CA^{t-1}B$ for $t \geq 1$. We can then rewrite the basic LQR problem in the form of optimization problem (8) (with no distributed constraint \mathcal{S}).

Example 3 (\mathcal{H}_2 State Feedback) Assume either that the generalized plant (2) is open-loop stable or that the control problem is over a finite horizon. Let $C_2 = I$ and $D_{21} = 0$ in the generalized plant (2) such that the problem is one of synthesizing an optimal state-feedback controller, and for clarity of exposition, assume that B_1 is invertible.⁷ Define the Youla parameterization for the controller synthesis problem (4) as follows [22]:

$$\tilde{P}_{12} = \frac{1}{z} P_{12}, \quad \tilde{P}_{21} = AP_{21} + B_1, \quad \tilde{U} = K(I - P_{22}K)^{-1} \tilde{P}_{21},$$

with all other parameters remaining the same. Under this parameterization, the optimal control problem (4) (with additional QI distributed constraint \mathcal{S}) with performance metric $\|\cdot\|_{\mathcal{H}_2}^2$ can be written as

$$\underset{\tilde{U} \in \mathcal{RH}_\infty}{\text{minimize}} \|P_{11} - \mathfrak{L}(\tilde{U})\|_{\mathcal{H}_2}^2 + \rho_u \|\tilde{U}\|_{\mathcal{H}_2}^2 \quad \text{s.t.} \quad \tilde{U} \tilde{P}_{21}^{-1} \in \mathcal{S}, \quad (10)$$

The optimal controller K is then recovered from the solution to (10) as $K = (I + \tilde{U} \tilde{P}_{21}^{-1} P_{22})^{-1} \tilde{U} \tilde{P}_{21}^{-1}$. The state-feedback assumption and the choice of Youla parameterization ensure that \tilde{P}_{21} is invertible in \mathcal{RH}_∞ and that $K \in \mathcal{S}$.

Remark 4 A dual argument applies to \mathcal{H}_2 filter design by considering the “full-control” setting, cf. [22] for more details.

As illustrated by Example 3, the \mathcal{H}_2 optimal control problem (8) is simply a more general way of writing the \mathcal{H}_2 state feedback model matching problem – in Remark 7 we show how \mathcal{H}_2 output feedback model matching problems can also be put in a similar form. Writing the \mathcal{H}_2 problem as a linear inverse problem with a least squares like state cost and an explicitly separated control cost already allows us to make connections to classical techniques from statistical inference. These connections (along with others we make later in this section) are summarized in Table 1. In order to keep the discussion as streamlined as possible, we make these connections in the context of the Basic LQR problem presented in Example 2.

⁷The assumption that B_1 is invertible simply implies that no component of the state is deterministic, and can be relaxed at the expense of more complicated formulas.

4.1.1 $\rho_u = 0$

In an inferential context, this is simply ordinary least squares, and is commonly used when U_* is not known *a priori* to have any structure. Further, the resulting estimate of U_* is unbiased, but often suffers from high error variance. Moving now to a control context, It is easy to see that this setting corresponds to “cheap control” LQR, in which there is no cost on u_t – under suitable controllability and observability assumptions, the resulting state trajectory is deadbeat, but the optimal control law is not necessarily unique.

4.1.2 $\rho_u > 0$

This corresponds to Ridge Regression or Tikhonov Regularization [33]. In an inferential context, this regularizer has the effect of *shrinking* estimates towards 0 – this introduces bias into the estimator, but reduces its variance, and is often a favorable tradeoff from a statistical perspective. From a linear algebraic perspective, this is a commonly used technique to improve the numerical conditioning of an inverse problem. Once again, the interpretation in RFD is clear: this corresponds to standard LQR control with $R = \rho I$; the parameter ρ allows for a tradeoff between control effort and state deviation. The optimal control action is then unique and the resulting state trajectory is generally not deadbeat.

4.2 The \mathcal{H}_2 RFD Optimization Problem with an Atomic Norm Penalty

Recall that our strategy for designing controller architectures is to augment the traditional model matching problem (6) with a structure inducing penalty, resulting in the RFD optimization problem (7). In light of the discussion of the previous subsection, we further specialize the RFD optimization problem to have an \mathcal{H}_2 performance metric and an atomic norm penalty $\|\cdot\|_{\mathcal{A}}$.

Definition 4 Let $U, Y \in \mathcal{RH}_\infty$ and $\mathfrak{L} : \mathcal{RH}_\infty \rightarrow \mathcal{RH}_\infty$ be of compatible dimension. The optimization problem

$$\underset{U \in \mathcal{RH}_\infty}{\text{minimize}} \|Y - \mathfrak{L}(U)\|_{\mathcal{H}_2}^2 + \rho \|U\|_{\mathcal{H}_2}^2 + 2\lambda \|U\|_{\mathcal{A}} \quad \text{s.t. } U \in \mathcal{S} \quad (11)$$

is called the \mathcal{H}_2 RFD optimization problem with parameters (ρ, λ) , distributed constraint \mathcal{S} , and atomic norm penalty $\|\cdot\|_{\mathcal{A}}$.

There are two components of note in this definition. The first is that ρ need not be equal to ρ_u , the control cost parameter of the original non-penalized control problem (8); the reasons why a different choice of ρ may be desirable are explained in §6. The second is the use of an atomic norm penalty function to induce structure. Indeed, if one seeks a solution U_* that can be composed as a linear combination of a small number of “atoms” \mathcal{A} , then a useful approach, as described in [3, 4], [5, 6], [7], to induce such structure in the solution of an inverse problem is to employ a convex penalty function that is given by the atomic norm induced by the atoms \mathcal{A} [2]. Examples of the types of structured solutions one may desire in linear inverse problems include sparse, group sparse and signed vectors, and low-rank, permutation and orthogonal matrices (see [7] for a more extensive list).

Specifically, if one assumes that

$$U = \sum_{i=1}^r c_i A_i, \quad A_i \in \mathcal{A}, \quad c_i \geq 0 \quad (12)$$

for a set of appropriately scaled and centered atoms \mathcal{A} , and a small number r relative to the ambient dimension, then solving

$$\underset{U}{\text{minimize}} \quad \Psi(U; Y, \mathfrak{L}) + 2\lambda \|U\|_{\mathcal{A}} \quad (13)$$

with the atomic norm $\|\cdot\|_{\mathcal{A}}$ given by the gauge function⁸

$$\begin{aligned} \|U\|_{\mathcal{A}} : &= \inf\{t \geq 0 \mid U \in t\text{conv}(\mathcal{A})\} \\ &= \inf\{\sum_{A \in \mathcal{A}} c_A \mid U = \sum_{A \in \mathcal{A}} c_A A, c_A \geq 0\} \end{aligned} \quad (14)$$

results in solutions that are both consistent with the data as measured in terms of the cost function $\Psi(\cdot; Y, \mathfrak{L})$, and that are sparse in terms of their atomic descriptions, i.e., are a nonnegative combination of a small number of elements from \mathcal{A} . Note that the unit ball of the atomic norm is given by the convex hull of its constituent atoms – as such, atomic norms are convex functions.

Our discussion in §3.2 on designing controller architecture by finding structured solutions to the model matching problem (6) suggests natural atomic sets for constructing suitable penalty functions for RFD. We make this point precise by showing that actuator, sensor, and/or communication delay design can all be performed through the use of a purposefully constructed atomic norm. We introduce several novel penalty functions for controller architecture design, most notably for the simultaneous design of actuator, sensor and communication delays. Further, all regularizers that have been considered for control architecture design in the literature (cf. [1, 12–17], among others) may be viewed as special instances of the atomic norms described below.

In what follows, the atomic sets that we define are of the form

$$\mathcal{A} = \bigcup_{\mathcal{A} \in \mathcal{M}} \mathcal{A} \cap k_{\mathcal{A}} \mathcal{B}_{\mathcal{H}_2}, \quad (15)$$

for \mathcal{M} an appropriate set of subspaces, $\{k_{\mathcal{A}}\}$ a set of normalization constants indexed by the subspaces $\mathcal{A} \in \mathcal{M}$, and $\mathcal{B}_{\mathcal{H}_2}$ the \mathcal{H}_2 unit norm ball; see the concrete examples below. Note that we normalize our atoms relative to the \mathcal{H}_2 norm as this norm is isotropic; hence this normalization ensures that no atom is preferred over another within a given family of atoms \mathcal{A} . We use n_s and n_a to denote the total number of sensors and actuators, respectively, available for the RFD task.

4.2.1 Actuator/Sensor Norm

For the *Actuator Norm*, we choose the atomic set to be transfer functions in $\mathcal{RH}_{\infty}^{n_a \times n_s}$ that have exactly one nonzero row with unit \mathcal{H}_2 norm, i.e., suitably normalized Youla parameters

⁸If no such t exists, then $\|X\|_{\mathcal{A}} = \infty$.

that use only one actuator. Specifically, the set of subspaces (15) in this context is

$$\mathcal{M}_{\text{act}} := \{ \mathcal{A} \subset \mathcal{RH}_{\infty}^{n_a \times n_s} \mid \mathcal{A} \text{ has one nonzero row} \}, \quad (16)$$

leading to the atomic set

$$\mathcal{A}_{\text{act}} := \{ e_i V \mid V \in \mathcal{RH}_{\infty}^{1 \times n_s}, \|V\|_{\mathcal{H}_2} = 1 \}. \quad (17)$$

The resulting atomic norm is then given by

$$\|U\|_{\text{act}} = \sum_{i=1}^{n_a} \|e_i^{\top} U\|_{\mathcal{H}_2}. \quad (18)$$

In particular, each ‘‘group’’ corresponds to a row of the Youla parameter. For the *Sensor Norm*, we similarly choose transfer functions with exactly one nonzero column with unit \mathcal{H}_2 norm, leading to the atomic norm

$$\|U\|_{\text{sns}} = \sum_{i=1}^{n_s} \|U e_i\|_{\mathcal{H}_2}. \quad (19)$$

Both of these norms are akin to a Group Lasso penalty [34].

4.2.2 Joint Actuator and Sensor Norm

Conceptually, each atom corresponds to a controller that uses only a small subset of actuators and sensors. As each row of the Youla parameter U corresponds to an actuator and each column to a sensor, the atomic transfer matrices have support defined by a submatrix of $U(z)$. Specifically, we choose atoms with at most k_a actuators and k_s sensors:

$$\mathcal{M}_{\text{act+sns}} := \{ \mathcal{A} \subset \mathcal{RH}_{\infty}^{n_s \times n_a} \mid \text{supp}(\mathcal{A}) \text{ is a submatrix with at most } k_a \text{ nonzero rows and } k_s \text{ nonzero columns} \} \quad (20)$$

The scaling terms $k_{\mathcal{A}}$ in the definition of the atomic set (15) are given by $k_{\mathcal{A}} = (\text{card}(\mathcal{A}) + .1)^{-\frac{1}{2}}$, and are necessary as some of the atoms are nested within others – the additional .1 can be any positive constant, and controls how much an atom of larger cardinality is preferred over several atoms of lower cardinality. The resulting *Actuator+Sensor Norm* is then constructed according to (15) and is akin to the latent Group Lasso [36].

4.2.3 Communication Link Norm

As described in §3.2, the communication design task is to select which additional links to introduce into an existing base communication graph. An atom in $\mathcal{A}_{\text{comm}}$ corresponds to such an additional link. We provide an example of such an atomic set for a simple system, and refer the reader to [17] for a more general construction. In particular, consider a three player chain system, with physical topology illustrated in Figure 3, such that P_{22} lies in the subspace

$$\mathcal{S} := \frac{1}{z} \begin{bmatrix} * & 0 & 0 \\ 0 & * & 0 \\ 0 & 0 & * \end{bmatrix} \oplus \frac{1}{z^2} \begin{bmatrix} * & * & 0 \\ * & * & * \\ 0 & * & * \end{bmatrix} \oplus \frac{1}{z^3} \mathcal{R}_p,$$




Figure 3: Three player chain system

where $*$ is used to denote \mathbb{R} to reduce notational clutter. We consider an existing communication graph matching the physical topology illustrated in Figure 3 so that the induced distributed subspace constraint, as described in §3.2, is given precisely by \mathcal{S} . It can be checked that \mathcal{S} is then QI with respect to P_{22} . We consider choosing from two additional links to augment the communication graph: a directed link from node 1 to node 3, and a directed link from node 3 to node 1. Then $\mathcal{M}_{\text{comm}} = \{\mathcal{A}_{13}, \mathcal{A}_{31}\}$, where its component subspaces are given by

$$\mathcal{A}_{13} = \frac{1}{z^2} \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ * & 0 & 0 \end{bmatrix}, \mathcal{A}_{31} = \frac{1}{z^2} \begin{bmatrix} 0 & 0 & * \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

In particular, each subspace \mathcal{A}_{ij} corresponds to the additional information available to the controller *uniquely* due to the added link from sensor j to actuator i . The resulting *Communication Link Norm* $\|\cdot\|_{\text{comm}}$ is then constructed according to (15) with all normalization constants $k_{\mathcal{A}} = 1$. We note that this penalty is also akin to the latent Group Lasso [36].

4.2.4 Joint Actuator (and/or Sensor) and Communication Link Norm

This penalty can be viewed as simultaneously inducing sparsity at the communication link level, while further inducing row sparsity as well. The general strategy is to combine the actuator and communication link penalties in a convex manner. We suggest two such approaches, one based on taking their weighted sums and the other based on taking their “weighted maximum.” In particular, we define the joint actuator *plus* communication link penalty to be:

$$\|U\|_{\text{act+comm}} = (1 - \theta) \|U\|_{\text{comm}} + \theta \|U\|_{\text{act}}, \quad (21)$$

for some $\theta \in [0, 1]$, and the *max* actuator/communication link penalty to be

$$\|U\|_{\max\{\text{act,comm}\}} = \max \{(1 - \theta) \|U\|_{\text{comm}}, \theta \|U\|_{\text{act}}\}, \quad (22)$$

for some $\theta \in [0, 1]$. The analogous Sensor and Communication Link penalties, as well as Sensor+Actuator and Communication link penalties can be derived by replacing \mathcal{A}_{act} with either \mathcal{A}_{sns} or $\mathcal{A}_{\text{act+sns}}$.

4.3 Further Connections with Structured Inference

As already noted in §4.1, by choosing different values of ρ for $\lambda = 0$ we are able to recover control-theoretic analogs to Ordinary Least Squares and Ridge Regression [33]. Noting that the actuator norm penalty (18) is akin to the Group Lasso [34], we now discuss how control theoretic analogs of the Group Lasso and Group Elastic Net [35] can be obtained by setting

Parameters	Stats Problem	Stats Structure	Stats Tradeoff	RFD Problem	RFD Structure	RFD Tradeoff
$\lambda = 0, \rho = 0$	Ordinary Least Squares	None	N/A	Cheap Control LQR	Deadbeat response	N/A
$\lambda = 0, \rho > 0$	Ridge Regression	Small Euclidean norm	Bias, Variance	LQR	Small control action	State deviation, Control effort
$\lambda > 0, \rho = 0$	Group LASSO	Group sparsity	Bias, Variance, Model complexity	RFD LQR	Sparse actuation	State deviation, Control effort, Actuation complexity
$\lambda > 0, \rho > 0$	Group Elastic Net	Correlated group sparsity	Bias, Variance, Model complexity	RFD LQR	Correlated sparse actuation	State deviation, Control effort, Actuation complexity

Table 1: A dictionary relating various SLIP problems in statistical inference and Actuator RFD problems.

$\lambda > 0$ in (11) and using the Actuator Norm (18) penalty – these connections are summarized in Table 1. To simplify the discussion, we once again consider these connections in the context of the basic LQR problem introduced in Example 2, now augmented with the actuator norm penalty (18).

In structured inference problems, the setting $\lambda > 0, \rho > 0$ corresponds to Group Elastic-Net regression. If the groups are single elements, this becomes the traditional Elastic Net and Lasso. The singleton group setting with $\lambda > 0, \rho = 0$ corresponds to Lasso regression, and this inference method is employed when the underlying model is known to be sparse – in particular, the Lasso penalty is used to select which elements U_i^* of the model are non-zero [3, 4]. Continuing with the singleton group setting, if both $\lambda > 0$ and $\rho > 0$, then the corresponding inferential approach is called the Elastic Net. In addition to the sparsity-inducing properties of the Lasso, the Elastic Net also encourages automatic clustering of the elements [35] – in particular, $\rho > 0$ encourages the simultaneous selection of highly correlated elements (two elements U_i^* and U_j^* are said to be highly correlated if $\mathfrak{L}(U_i^*) \approx \mathfrak{L}(U_j^*)$). Thus ρ can be seen as a parameter that can be adjusted to leverage a prior of *correlation* in the underlying measurement operator \mathfrak{L} . These interpretations carry over to more general groups in a natural way.

In RFD, this setting corresponds to our motivating Example 2 augmented with the actuator norm penalty, in which we design the controller’s actuation architecture. As each atom corresponds to an actuator, this RFD procedure then selects a small number of actuators. Porting the clustering effect interpretation from the structured inference setting, we see that ρ promotes the selection of actuators that have similar effects on the closed loop response. In particular, this suggests that for systems in which no such similarities are expected, ρ should be chosen to be small (or 0) during the RFD process, even if the original LQR problem had non-zero control cost.

5 The RFD Procedure

5.1 The Two-Step Algorithm

We now introduce the convex optimization based RFD procedure for the co-design of an optimal controller and the architecture on which it is implemented. The remaining computational challenge is the possibly infinite dimensional nature of the RFD optimization problem (7). To address this issue, we propose a two step procedure: first, a finite dimensional approximation of optimization problem (7) is solved to identify a potential controller architecture and its defining subspace constraint \mathcal{D} . Once this architecture has been identified, a traditional (and possibly infinite dimensional) optimal control problem (1) with Youla parameter restricted to lie in $\mathcal{D} \cap \mathcal{S}$ is then solved – in particular, in many interesting settings the resulting optimal controller restricted to the designed architecture can then be computed exactly leveraging results from the optimal controller synthesis literature [22–28].

Formally, we begin by fixing an optimization horizon t and a controller order v . We suggest initially choosing t and v to be small (i.e., 2 or 3), and then gradually increasing these parameters until a suitable controller architecture/control law pair is found. Our motivations for this approach are twofold: (i) first, selecting a small horizon t and small controller order v leads to a smaller optimization problem that is computationally easier to solve; and (ii) as we show in the next section, a smaller horizon t and smaller controller order v can actually aid in the identification of optimal controller architectures. For a given performance metric $\Psi(\cdot; Y, \mathfrak{L})$ and atomic norm penalty $\|\cdot\|_{\mathcal{A}}$, the two step RFD procedure consists of an architecture design step and an optimal control law design step:

1) Architecture design: Select the regularization weight λ and solve the *finite dimensional* RFD optimization problem

$$\underset{U \in \mathcal{RH}_{\infty}^{\leq v} \cap \mathcal{S}}{\text{minimize}} \quad \Psi(U; Y^{\leq t}, \mathfrak{L}^{\leq t, v}) + 2\lambda \|U\|_{\mathcal{A}} \quad (23)$$

The actuators, sensors and communication links defining the designed architecture are specified by the non-zero atoms that constitute the solution \hat{U} to optimization problem (23). The architectural components employed in \hat{U} in turn define a subspace $\mathcal{D}(\hat{U})$ which corresponds to all controllers (within the Youla parameterization) that have the same architecture as \hat{U} .⁹

2) Optimal control law design: Solve the infinite dimensional optimal control problem with Youla parameter additionally constrained to lie in the designed subspace $\mathcal{D}(\hat{U}) \cap \mathcal{S}$:

$$\underset{U \in \mathcal{RH}_{\infty}}{\text{minimize}} \quad \Psi(U; Y, \mathfrak{L}) \quad \text{s.t.} \quad U \in \mathcal{D}(\hat{U}) \cap \mathcal{S}. \quad (24)$$

If the resulting controller architecture and controller performance are acceptable, the RFD procedure terminates. Otherwise, adjust λ accordingly to vary the tradeoff between architectural complexity and closed loop performance. If no suitable controller architecture/controller can be found, increase t and v and repeat the procedure.

Remark 5 (Removing Bias) *The method of solving a regularized optimization problem to identify the architectural structure of a controller and then solving a standard model matching*

⁹As described in §3, the subspace $\mathcal{D}(\hat{U}) \cap \mathcal{S}$ is QI by construction, and hence this subspace also corresponds to all controllers with the same architecture as $\hat{K} = (I + \hat{U}P_{22})^{-1}\hat{U}$.

problem restricted to the identified architecture is analogous to a procedure that is commonly employed in structured inference. In structured inference problems, a regularized problem is solved first to identify a subspace corresponding to the structure of an underlying model U_* . Subsequently, a non-regularized optimization problem with solution restricted to that identified subspace is solved to obtain an unbiased estimator of the underlying model U_* .

5.2 Simultaneous Actuator, Sensor and Communication RFD

In this subsection we demonstrate the full power and flexibility of the RFD framework in designing a distributed controller architecture, jointly incorporating actuator, sensor and communication link design. In particular we consider a plant with eleven subsystems with topology as illustrated in Figure 4. The solid lines correspond to the physical interconnection between subsystems. Choosing $C_2 = B_2 = I$, the adjacency matrix of this graph then defines the support of the A matrix in the state space realization of the generalized plant (2), as well as the required communication links between nodes such that the distributed constraint is QI under P_{22} [17, 20].




Figure 4: Topology of system considered for RFD example. Solid lines indicate both physical interconnections and existing communication links between controllers. Dashed lines correspond to possible additional edges to be added.

The non-zero entries of A were generated randomly and normalized such that $|\lambda_{\max}(A)| = .999$. The remaining state space parameters of the generalized plant (2) satisfy $D_{12} \begin{bmatrix} C_1^\top & D_{12}^\top \end{bmatrix} = \begin{bmatrix} 0 & 25I \end{bmatrix}$, and $D_{21}^\top \begin{bmatrix} B_1 & D_{21} \end{bmatrix} = \begin{bmatrix} 0 & .01I \end{bmatrix}$, with $C_1 C_1^\top = 100I$ and $B_1^\top B_1 = I$.

For the RFD task, we choose an \mathcal{H}_2 norm performance metric; we allow each node to be equipped with an actuator and/or a sensor (for a total of 11 possible actuators and sensors), and we allow the communication graph to be augmented with any subset of the interconnections denoted by the dashed lines in Figure 4, in addition to the already present links given by the solid lines. This leads to 536,870,911 different possible controller architectures.

We solved the RFD optimization (43) with atomic norm $\|\cdot\|_{\text{act+sns+comm}}$ as defined in §4.2, with weighting parameter $\theta = .75$ and with $k_s = k_a = 1$. We performed the RFD procedure for two different horizon/order pairs: $t = 4$ and $v = 2$, as well as $t = 6$ and $v = 3$; for these latter horizon/order values acceptable tradeoffs between architecture complexity and closed loop performance were identified, and hence the RFD procedure terminated. For each horizon/order pair (t, v) , we vary λ , and for each resulting optimal solution \hat{U} , we identified the designed architecture and corresponding subspace $\mathcal{D}(\hat{U})$. We then used the method from [27] to exactly solve the resulting non-regularized distributed \mathcal{H}_2 model matching problem with subspace constraint $\mathcal{D}(\hat{U})$. Note that the Youla parameter solving this non-regularized problem is not restricted to have a finite impulse response. In particular, we can compute the optimal Youla parameter and the corresponding optimal controller

restricted to the architecture underlying \hat{U} in a computationally tractable fashion because we guarantee that the subspace corresponding to the designed architecture is QI, as per the discussion in Section 4.

For horizon $t = 6$ and order $v = 3$, the resulting architectural complexity is plotted against the closed loop norm of the system in Figure 5. As λ is increased, the architectural complexity (i.e. the number of actuators, sensors and communication links) decreases, but at the expense of deviations from the performance achieved by the controller that uses all of the available architectural resources. We also show the resulting architecture for $\lambda = 500$ in Figure 6: as can be seen, a non-obvious combination of eight actuators, eight sensors and five additional communication links are chosen, resulting in only a 0.71% degradation in performance over the distributed controller using all eleven actuators, eleven sensors and seven additional communication links.

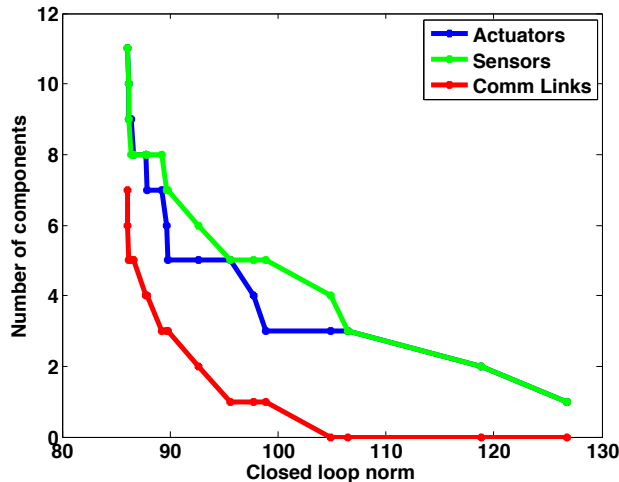


Figure 5: A small degradation in closed loop performance allows for a significant decrease in architectural complexity.




Figure 6: Resulting architecture for $\lambda = 500$: despite only using eight actuators (orange squares), eight sensors (blue triangles) and five additional communication links (green arrows), the performance only degraded by 0.71% relative to the distributed controller using all eleven actuators, eleven sensors and seven communication links.

As this example shows, the RFD procedure is effective at identifying simple controller architectures that approximate the performance of a controller that maximally utilizes the available architectural resources. In the next section, we offer some theoretical justification for the success of our procedure by suitably interpreting the RFD optimization problem (7) in the context of approximation theory and by making connections to analogous problems in structured inference.

6 Recovery of Optimal Actuation Structure

This section is dedicated to the analysis of the \mathcal{H}_2 RFD optimization problem (11) with no distributed constraint, actuator norm penalty (18), and Y and \mathfrak{L} as given in Example 3 – a nearly identical argument applies to a sensor norm regularized problem. We discuss how to extend the analysis to output feedback and distributed problems at the end of the section.

Viewing the model matching problem (6) as a linear inverse problem makes it clear that designing a structured controller is akin to obtaining a structured solution to a linear inverse problem. The problem of obtaining structured solutions to linear inverse problems arises prominently in many contexts, most notably in statistical estimation. In that setting, one posits a linear measurement model¹⁰

$$Y = \mathfrak{L}(U_*) + W, \quad (25)$$

where Y is the vector of observations, \mathfrak{L} is the measurement map, U_* parametrizes an underlying model and W is the measurement error. The linear model (25) also has an appealing interpretation from a control-theoretic perspective. In particular, letting $Y \in \mathcal{RH}_\infty$ be the state component of the open loop response of a LTI system, $U_* \in \mathcal{RH}_\infty$ be a suitably defined Youla parameter, and $\mathfrak{L} : \mathcal{RH}_\infty \rightarrow \mathcal{RH}_\infty$ be the map from Youla parameter to the state component of the closed loop response, it is then immediate that

$$W := Y - \mathfrak{L}(U_*) \quad (26)$$

represents the state component of the closed loop response achieved by the controller U_* . Table 2 summarizes the correspondence between these two perspectives.

Parameter	Structured Controller Design	Structured Inference
Y	Open loop system	Observations
\mathfrak{L}	Map to closed loop	Measurement map
U_*	Desired controller	Underlying model
W	Closed loop response	Measurement noise

Table 2: Interpretation of parameters in Structured Controller Design and Structured Inference.

This conceptual connection suggests a novel interpretation of the role of the closed loop response W achieved by a controller U_* . In an inferential context, since W corresponds to measurement noise, a smaller W makes the task of identifying the structure of the underlying model U_* much easier, as the measurements are more accurate. In a similar spirit, we demonstrate that structured controller design is easier (via the solution to an RFD optimization problem (7)) if the corresponding state component of the closed loop response is small. Thus the state component of the closed loop response of the system plays the role of noise when trying to identify the structure of a suitably defined controller U_* . In the sequel, we describe an appropriate notion of smallness for the state component of the closed loop response in the context of designing structured controllers.

The remainder of the discussion in this section builds on prominent results from the structured inference literature [3–6]. The flavor of these results is somewhat non-standard in the

¹⁰We purposefully use non-standard notation to facilitate comparisons between RFD and SLIP.

controls literature, and we therefore pause briefly to frame the setup in this section appropriately and to discuss how the results of this section should be interpreted. The main result of this section proceeds by assuming that there exists an architecturally simple controller U_* (i.e., one with a small number of actuators) that achieves a good closed loop response, (i.e., that achieves a small W as defined in (26)). Under suitable conditions, Theorems 1 and 2 state that the architectural structure of U_* can be recovered via tractable convex optimization using the RFD procedure. These conditions are phrased in terms of quantities associated to U_* which are typically unknown in advance – however, these conditions are not meant to be checked prior to solving a RFD optimization problem. Although the results are stated in terms of a nominal controller U_* , they should be interpreted as describing the properties satisfied by controller architectures identified via the RFD procedure of §5. In particular, the RFD procedure requires solving RFD optimization problems across a range of controller orders v , optimization horizons t and regularization weights λ : this process leads to a set of controller architectures being identified. Our results allow a practitioner to be confident that all controller architectures satisfying the conditions of our theorems – i.e., those that have a small number of actuators and that achieve a small closed loop state response – are included in this set of identified controller architectures. In this way, the RFD procedure provably identifies good controller architectures, should they exist.

We study finite dimensional variants of the \mathcal{H}_2 RFD optimization problem (11) with the actuator norm penalty (18), and show that such finite dimensional approximations are sufficient to identify the structure of a desired controller U_* . In particular, we truncate the optimization problem (11) to a finite horizon t by restricting $Y - \mathfrak{L}(U) \in \mathcal{RH}_\infty$ to the first t elements of its impulse response, and to a finite controller order v by restricting U to lie in $\mathcal{RH}_\infty^{\leq v}$. The resulting optimization problem is thus finite dimensional, and corresponds to the first step of the RFD procedure defined in the previous section. At this point, it is convenient to introduce the temporally truncated version of (26) for a fixed optimization horizon t and controller order v :

$$\begin{aligned} W^{\leq t} &= Y^{\leq t} - \mathfrak{L}^{\leq t,t}(U_*^{\leq t}) \\ &= Y^{\leq t} - \mathfrak{L}^{\leq t,v}(U_*^{\leq v}) - T^{\leq t,v} \end{aligned} \quad (27)$$

with

$$T^{\leq t,v} := \mathfrak{L}^{\leq t,t}(U_*^{\leq t} - U_*^{\leq v}) \quad (28)$$

corresponding to the effect of the “tail” of U_* on the state component $W^{\leq t}$ of the truncated closed loop response.

The flexibility in the choice of the optimization horizon t and controller order v will be the focus of much of our discussion. In particular, it is of interest to find the smallest t and v for which we can guarantee that the RFD procedure recovers the structure underlying U_* – the smaller the horizon and controller order, the smaller the size of the optimization problem that needs to be solved. Perhaps counter-intuitively, we show that larger t and v do not necessarily help in recovering the structure of an underlying parameter U_* . We make this statement precise in what follows, but again drawing on intuition from the structured inference literature, we note that increasing v in RFD is analogous to increasing the allowed model complexity when solving an inference problem: if the model class is too rich, we risk over-fitting and thus obfuscating the structure of the underlying model U_* .

Our goal is to prove that the solution \tilde{U} to the finite dimensional \mathcal{H}_2 RFD optimization problem

$$\tilde{U} = \underset{U \in \mathcal{RH}_\infty^{\leq v}}{\operatorname{argmin}} \|Y^{\leq t} + \mathfrak{L}^{\leq t, v}(U)\|_{\mathcal{H}_2}^2 + \rho \|U\|_{\mathcal{H}_2}^2 + 2\lambda \|U\|_{\operatorname{act}} \quad (29)$$

has the same architectural structure as U_* for appropriately chosen t and v . To show this, we study the solution \hat{U} to the following *architect* optimization problem:

$$\begin{aligned} \hat{U} = \underset{U \in \mathcal{RH}_\infty^{\leq v}}{\operatorname{argmin}} & \|Y^{\leq t} + \mathfrak{L}^{\leq t, v}(U)\|_{\mathcal{H}_2}^2 + \rho \|U\|_{\mathcal{H}_2}^2 + 2\lambda \|U\|_{\operatorname{act}} \\ \text{s.t. } & U \in \mathcal{A}_* \end{aligned} \quad (30)$$

where \mathcal{A}_* is a subspace of Youla parameters U with the property that a row of U_* being zero implies that the corresponding row of U is zero. In words, \mathcal{A}_* may be viewed as the set of Youla parameters corresponding to actuation schemes matching the actuation scheme of U_* . We also define $\mathcal{M}_* \subset \mathcal{M}_{\operatorname{act}}$, with $\mathcal{M}_{\operatorname{act}}$ defined as in (16), to be

$$\mathcal{M}_* := \{\mathcal{A} \in \mathcal{M}_{\operatorname{act}} \mid (U_*)_{\mathcal{A}} \neq 0\}. \quad (31)$$

In words, the elements of \mathcal{M}_* correspond to actuation schemes that use a single actuator, where these actuators are defined by the nonzero rows of the desired controller U_* .

We show under suitable conditions on t , v , U_* and $\mathfrak{L}^{\leq t, v}$ that $\hat{U} = \tilde{U}$; that is to say that the architect solution \hat{U} is also the unique optimal solution to the RFD optimization problem (29) without the additional constraint $U \in \mathcal{A}_*$. As a result, since \hat{U} is constrained to lie in \mathcal{A}_* , the solution to the RFD optimization problem \tilde{U} also lies in \mathcal{A}_* and hence has the same architectural structure as U_* . We emphasize that at no stage during the RFD procedure described in §5 do we require knowledge of \mathcal{A}_* and \mathcal{M}_* – the investigation of the architect problem (30) is only a theoretical tool used to prove structural recovery results.

6.1 Identifiability Conditions in Control

We begin by introducing two *restricted gains* in terms of the subspace \mathcal{A}_* and its orthogonal complement \mathcal{A}_*^\perp . In order to do so, we introduce the dual norm to $\|\cdot\|_{\operatorname{act}}$, which is given by

$$\|U\|_{\operatorname{act}}^* = \max_{\mathcal{A} \in \mathcal{M}_{\operatorname{act}}} \|U_{\mathcal{A}}\|_{\mathcal{H}_2}. \quad (32)$$

These restricted gains are then

$$\alpha^{\leq t, v} := \min_{\Delta} \left\| \left(\left[\mathfrak{L}_{\mathcal{A}_*}^{\leq t, v} \right]^\dagger \mathfrak{L}_{\mathcal{A}_*}^{\leq t, v} + \rho I \right) (\Delta) \right\|_{\operatorname{act}}^* \quad (33)$$

s.t. $\|\Delta\|_{\operatorname{act}}^* = 1, \Delta \in \mathcal{A}_* \cap \mathcal{RH}_\infty^{\leq v}$

$$\beta^{\leq t, v} := \max \left\| \left[\mathfrak{L}_{\mathcal{A}_*^\perp}^{\leq t, v} \right]^\dagger \mathfrak{L}_{\mathcal{A}_*^\perp}^{\leq t, v} (\Delta) \right\|_{\operatorname{act}}^* \quad (34)$$

s.t. $\|\Delta\|_{\operatorname{act}}^* \leq 1, \Delta \in \mathcal{A}_* \cap \mathcal{RH}_\infty^{\leq v}$.

The minimum gain $\alpha^{\leq t, v}$ is a quantitative measure of the injectivity of the operator $\mathfrak{L}^{\leq t, v} \times \sqrt{\rho}I$ restricted to the subspace \mathcal{A}_* . Intuitively, it characterizes the distinctions

among the effects of the different actuators within \mathcal{A}_* . The maximum gain $\beta^{\leq t, v}$, on the other hand, is a measure of how different the effects of actuators in \mathcal{A}_* are from those of actuators in \mathcal{A}_*^\perp .

We can already see some immediate implications of different choices of the horizon t and controller order v on these quantities. In particular, $\alpha^{\leq t, v}$ is non-increasing in the controller order v . This minimum gain's dependence on the horizon t is more subtle. Define the *mixing time* of \mathcal{M}_* to be

$$\tau_{\mathcal{M}_*} := \max \left\{ t \in \mathbb{Z}_+ \mid [\mathfrak{L}_{\mathcal{A}}^{\leq t}]^\dagger \mathfrak{L}_{\mathcal{B}}^{\leq t} = 0, \forall \mathcal{A} \neq \mathcal{B} \in \mathcal{M}_* \right\}. \quad (35)$$

If no t exists such that the condition within the max $\{\cdot\}$ is satisfied, we set $\tau_{\mathcal{M}_*} = 0$. The mixing time $\tau_{\mathcal{M}_*}$ measures how long it takes for the effects of the distinct actuators used by U_* to overlap, or mix, in the closed loop response. Consequently the minimum gain $\alpha^{\leq t, v}$ is non-decreasing in t so long as $t \leq \tau_{\mathcal{M}_*}$, i.e., so long as t is sufficiently small that the effects of the different actuators used by U_* do not overlap. We then have the following lemma:

Lemma 1 *Let $\tau_{\mathcal{M}_*}$ be as defined in (35). Then*

$$\alpha^{\leq t, v} = \rho + \min_{\mathcal{A} \in \mathcal{M}_*} \sigma_{\min} \left(\left[\mathfrak{L}_{\mathcal{A}}^{\leq t, v} \right]^\dagger \mathfrak{L}_{\mathcal{A}}^{\leq t, v} \right) \quad (36)$$

for all $1 \leq t \leq \tau_{\mathcal{M}_*}$. In particular, $\alpha^{\leq t, v}$ is non-decreasing in t for all $1 \leq t \leq \tau_{\mathcal{M}_*}$.

Proof: It is easily verified that for $\mathcal{A} \neq \mathcal{B} \in \mathcal{M}_*$ and $t \leq \tau_{\mathcal{M}_*}$, we have $\left[\mathfrak{L}_{\mathcal{A}}^{\leq t, v} \right]^\dagger \mathfrak{L}_{\mathcal{B}}^{\leq t, v} = 0$, from which (36) follows. To see that $\alpha^{\leq t, v}$ as given in (36) is non-decreasing in t , it suffices to note that

$$\left[\mathfrak{L}_{\mathcal{A}}^{\leq t+1, v} \right]^\dagger \mathfrak{L}_{\mathcal{A}}^{\leq t, v} = \left[\mathfrak{L}_{\mathcal{A}}^{\leq t, v} \right]^\dagger \mathfrak{L}_{\mathcal{A}}^{\leq t, v}, \quad (37)$$

leading to the conclusion that

$$\left[\mathfrak{L}_{\mathcal{A}}^{\leq t+1, v} \right]^\dagger \mathfrak{L}_{\mathcal{A}}^{\leq t+1, v} = \left[\mathfrak{L}_{\mathcal{A}}^{\leq t, v} \right]^\dagger \mathfrak{L}_{\mathcal{A}}^{\leq t, v} + \left[\mathfrak{L}_{\mathcal{A}}^{\leq t+1, v} - \mathfrak{L}_{\mathcal{A}}^{\leq t, v} \right]^\dagger \left(\mathfrak{L}_{\mathcal{A}}^{\leq t+1, v} - \mathfrak{L}_{\mathcal{A}}^{\leq t, v} \right). \quad (38)$$

The result follows by noting that the final term in this expression is positive semidefinite. ■

In particular, this result suggests that actuation schemes with more evenly distributed actuators (i.e., those with larger mixing times $\tau_{\mathcal{M}_*}$ (35)) are easier to identify.

The maximum gain $\beta^{\leq t, v}$, however, is clearly seen to be non-decreasing both in the controller order v and the horizon t . This is consistent with our interpretation of $\beta^{\leq t, v}$ as a measure of similarity between actuators: as either v or t increase, there is more time for the mixing of the actuators' control actions via the propagation of dynamics in the system, increasing their worst-case "similarity." We now assume that the following *identifiability condition* is satisfied.

Assumption 1 (Identifiability) *There exist $1 \leq v \leq t < \infty$ such that*

$$\frac{\beta^{\leq t, v}}{\alpha^{\leq t, v}} =: \delta \in [0, 1). \quad (39)$$

In light of the previous discussion, it is immediate that a larger controller order v decreases the likelihood of the identifiability condition being satisfied, and should therefore be taken as small as possible. The effect of increasing the horizon t is less clear, but we see that it may help if the minimum gain $\alpha^{\leq t,v}$ increases sufficiently fast with t relative to the increase in the gain $\beta^{\leq t,v}$ with respect to t – further there is no need to increase t beyond the mixing time $\tau_{\mathcal{M}_*}$.

In the inference literature, the analog of these identifiability assumptions are given by conditions known as the *restricted eigenvalue condition* [37] and the *restricted isometry property* [38]. In the sequel, we give an example of deterministic and structured state space matrices that satisfy these identifiability conditions. Specifically, we focus on systems (2) that have block diagonal B_2 and C_1 matrices (i.e., decoupled actuators and state costs), and block banded state matrices A (i.e., locally coupled dynamics).

Remark 6 Notice that if $\mathfrak{L}^{\leq t,v} = I$, then $\alpha^{\leq t,v} \geq 1$, $\beta^{\leq t,v} = 0$ and $\delta = 0$. These conditions are satisfied if $B = C = I$ and $v = t = 1$ in Example 2 (Basic LQR), or if $C_1 = B_2 = I$ and $v = 1$, $t = 2$ in Example 3 (\mathcal{H}_2 State Feedback). Thus sufficiently small values of v and t ensure that condition (39) holds. However, the resulting optimization problem only incorporates low order effects of the dynamics (as encoded in $\mathfrak{L}^{\leq t,v}$) in the RFD optimization problem, suggesting that t and v should be also be chosen large enough to sufficiently capture the dynamics of the system. This observation and Lemma 1 motivate our suggestion in Section 5 to begin with small horizon t and controller order v and to then gradually increase these values until a suitable controller architecture/control law pair is found.

6.2 Sufficient Conditions for Recovery

The following theorem provides sufficient conditions for (i) the architect solution \hat{U} to be the unique optimal solution to the finite dimensional RFD optimization (29), and (ii) an actuator of the desired controller, identified by a subspace $\mathcal{A} \in \mathcal{M}_*$, to be identified by the RFD procedure.

Theorem 1 (Structural Recovery) Fix a horizon $1 \leq t < \infty$, and a controller order $1 \leq v \leq t$ such that Assumption 1 holds. If

$$\lambda > \frac{\delta}{1-\delta} \left(\left\| \left[\mathfrak{L}_{\mathcal{A}_*}^{\leq t,v} \right]^\dagger (W^{\leq t} + T^{\leq t,v}) \right\|_{act}^* + \rho \|U_*^{\leq v}\|_{act}^* \right) + \frac{1}{1-\delta} \left\| \left[\mathfrak{L}_{\mathcal{A}_*^\perp}^{\leq t,v} \right]^\dagger (W^{\leq t} + T^{\leq t,v}) \right\|_{act}^* \quad (40)$$

we have that \hat{U} as defined in (30) is the unique optimal solution to (29), and that the row support of \hat{U} is contained within the row support of U_* . Further if $\mathcal{A} \in \mathcal{M}_*$ and

$$\|(U_*^{\leq v})_{\mathcal{A}}\|_{\mathcal{H}_2} > \frac{1}{\alpha} \left(\lambda + \left\| \left[\mathfrak{L}_{\mathcal{A}_*}^{\leq t,v} \right]^\dagger (W^{\leq t} + T^{\leq t,v}) \right\|_{act}^* + \rho \|U_*^{\leq v}\|_{act}^* \right), \quad (41)$$

then $\hat{U}_{\mathcal{A}} \neq 0$.

The condition (40) states that, under suitable identifiability assumptions, the regularization weight λ needs to be sufficiently large to guarantee that the architect solution \hat{U} is also the solution to (30). However, this can always be made to hold by choosing λ sufficiently large so that $\hat{U} = \tilde{U} = 0$. The second condition (41) provides an upper bound on the values of λ for which a specific actuator (i.e., a specific component $(U_*^{\leq v})_{\mathcal{A}}$, $\mathcal{A} \in \mathcal{M}_*$) is identified by the architect solution \hat{U} . The following corollary then guarantees the recovery of \mathcal{M}_* .

Corollary 1 *Let $\mu := \min_{\mathcal{A} \in \mathcal{M}_*} \|(U_*^{\leq v})_{\mathcal{A}}\|_{\mathcal{H}_2}$, and suppose that the open interval*

$$\Lambda := \left(\frac{\delta}{1-\delta} \left[\left\| \left[\mathfrak{L}_{\mathcal{A}_*}^{\leq t, v} \right]^\dagger (W^{\leq t} + T^{\leq t, v}) \right\|_{act}^* + \rho \|U_*^{\leq v}\|_{act}^* \right] + \frac{\left\| \left[\mathfrak{L}_{\mathcal{A}_*}^{\leq t, v} \right]^\dagger (W^{\leq t} + T^{\leq t, v}) \right\|_{act}^*}{1-\delta}, \right. \\ \left. \alpha^{\leq t, v} \mu - \left\| \left[\mathfrak{L}_{\mathcal{A}_*}^{\leq t, v} \right]^\dagger (W^{\leq t} + T^{\leq t, v}) \right\|_{act}^* - \rho \|U_*^{\leq v}\|_{act}^* \right) \quad (42)$$

is non-empty. Then the solution \tilde{U} to the RFD optimization (29), with any regularization weight λ chosen within Λ , has row support equal to that of U_ .*

Note that it is useful to have a given architecture be identifiable for a range of regularization weights λ , as prior information about the values needed to specify (42) are typically not available. We exploited this fact when we defined the RFD procedure in Section 5 by suggesting that λ be varied until a suitable architecture/control law pair is identified. This corollary also makes explicit that larger values of ρ shrinks the range of λ for which the RFD procedure is successful. It also shows that larger $T^{\leq t, v}$ tail terms (28) are deleterious to the performance of the RFD procedure as well – therefore although we previously stated that the controller order v should be chosen as small as possible, it should not be so small that the tail term $T^{\leq t, v}$ is too large.

Remark 7 (Extension to Output Feedback) *A similar argument applies to the output feedback problem, but at the expense of more complicated formulas. In particular the \mathcal{H}_2 RFD optimization takes the form*

$$\underset{U \in \mathcal{RH}_\infty}{\text{minimize}} \quad \|Y - \mathfrak{L}(U)\|_{\mathcal{H}_2}^2 + \|\mathfrak{F}(U)\|_{\mathcal{H}_2}^2 + 2\lambda \|U\|_{act}, \quad (43)$$

for $Y = P_{11}$, $\mathfrak{L}(U) = P_{12}UP_{21}$, and $\mathfrak{F}(U) = [P_{12}UD_{21} \quad D_{12}UP_{21} \quad D_{12}UD_{21}]$. Notice that if D_{12} and D_{21} are set to 0 in the RFD optimization problem (43), we recover an optimization problem of exactly the same form as (11) with $\rho = 0$, in which case the analyses of this section and the next section are applicable.

Remark 8 (Extension to Distributed Constraints) *For the purposes of analysis, the additional constraint that $U \in \mathcal{S}$ can be incorporated by considering the restriction of \mathfrak{L} to \mathcal{S} , resulting in a centralized problem (cf. [19] for an example of how this can be done).*

6.3 A RFD Signal to Noise Ratio

Theorem 1 as stated does not yet provide an immediate interpretation of the effect of the choices of the horizon t and the controller order v on the performance of the RFD procedure. In order to better understand the effects of the horizon t and controller order v on the success of the RFD procedure, we describe more interpretable bounds on $\alpha^{\leq t, v}$ and $\beta^{\leq t, v}$.

Lemma 2 *Fix $1 \leq t < \infty$ and $1 \leq v \leq t$. The parameters $\alpha^{\leq t, v}$ and $\beta^{\leq t, v}$, as defined in equations (33) and (34), can be bounded from below and above, respectively, as follows:*

$$\alpha^{\leq t, v} \geq \rho + \gamma^{\leq t, v}, \quad (44)$$

where

$$\gamma^{\leq t, v} := \min_{\mathcal{B} \subseteq \mathcal{M}_*, |\mathcal{B}| \geq 1} \max_{\mathcal{A} \in \mathcal{B}} \left[\sigma_{\min} \left(\left[\mathfrak{L}_{\mathcal{A}}^{\leq t, v} \right]^\dagger \mathfrak{L}_{\mathcal{A}}^{\leq t, v} \right) - \sum_{\mathcal{B} \neq \mathcal{A} \in \mathcal{B}} \sigma_{\max} \left(\left[\mathfrak{L}_{\mathcal{A}}^{\leq t, v} \right]^\dagger \mathfrak{L}_{\mathcal{B}}^{\leq t, v} \right) \right] \quad (45)$$

and

$$\beta^{\leq t, v} \leq \max_{\mathcal{A} \in (\mathcal{M}_{act} \setminus \mathcal{M}_*)} \sum_{\mathcal{B} \in \mathcal{M}_*} \sigma_{\max} \left(\left[\mathfrak{L}_{\mathcal{A}}^{\leq t, v} \right]^\dagger \mathfrak{L}_{\mathcal{B}}^{\leq t, v} \right). \quad (46)$$

Consequently, we can upper bound the ratio (39) as

$$\delta \leq \frac{\beta^{\leq t, v}}{\rho + \gamma^{\leq t, v}}. \quad (47)$$

In particular, it is a straightforward consequence of Lemma 1 that for all $t \leq \tau_{\mathcal{M}_*}$ (where the mixing time $\tau_{\mathcal{M}_*}$ is as in (35)), the intermediate quantity $\gamma^{\leq t, v}$, as introduced in Lemma 2, is given by

$$\gamma^{\leq t, v} = \min_{\mathcal{A} \in \mathcal{M}_*} \sigma_{\min} \left(\left[\mathfrak{L}_{\mathcal{A}}^{\leq t, v} \right]^\dagger \mathfrak{L}_{\mathcal{A}}^{\leq t, v} \right),$$

and is non-decreasing in t . Further it is easily verified that the bounds (44), (46) and (47) can be taken with equality if $v = 1$, as each $\mathfrak{L}_{\mathcal{A}}^{\leq t, 1}$ is isomorphic to a column vector.

The bounds computed in Lemma 2 can be combined with the sufficient conditions of Theorem 1 to describe sufficient conditions for the successful recovery of the architecture of U_* in terms of a signal to noise like quantity – to that end, we introduce the following definitions.

Definition 5 (RFD Noise) *We define the RFD Noise level $\eta_{\mathcal{M}_*}^{\leq t, v}$ for an \mathcal{H}_2 RFD optimization problem (29) to be*

$$\eta_{\mathcal{M}_*}^{\leq t, v} := \left\| \left[\mathfrak{L}_{\mathcal{A}_*}^{\leq t, v} \right]^\dagger (W^{\leq t} + T^{\leq t, v}) \right\|_{act}^* + \left\| \left[\mathfrak{L}_{\mathcal{A}_*^\dagger}^{\leq t, v} \right]^\dagger (W^{\leq t} + T^{\leq t, v}) \right\|_{act}^*. \quad (48)$$

The control theoretic interpretation (26) of the linear model (25) used in inference problems motivates our terminology – recall in particular that in (26) the state component of the closed loop response W is interpreted as measurement noise in the context of identifying a structured controller. Likewise, $T^{\leq t, v}$ can be viewed as additional noise introduced into the architecture identification procedure by the temporal truncation procedure described in (27). We proceed to define a control theoretic analog to the signal in the context of RFD optimization problems.

Definition 6 (RFD SNR) *In the context of architecture recovery via RFD, the magnitude of each atom, $\|(U_*^{\leq v})_{\mathcal{A}}\|_{\mathcal{H}_2}$ plays the role of a signal, and the RFD Noise level $\eta_{\mathcal{M}_*}^{\leq t, v}$ that of noise, leading to the definition of the SNR of a component $(U_*^{\leq v})_{\mathcal{A}}$, $\mathcal{A} \in \mathcal{M}_*$ as*

$$SNR((U_*^{\leq v})_{\mathcal{A}}) := \frac{\|(U_*^{\leq v})_{\mathcal{A}}\|_{\mathcal{H}_2}}{\eta_{\mathcal{M}_*}^{\leq t, v}}. \quad (49)$$

These definitions allow us to state simple conditions in terms of the SNR (49) for the successful recovery of an actuation architecture via the solution to the \mathcal{H}_2 RFD optimization problem (29).

Theorem 2 *Let $\rho = 0$, $\lambda = \lambda' + \kappa$, where λ' is given by the right hand side of (40), and $\kappa > 0$ is an arbitrarily small constant, and assume that $\beta^{\leq t, v} / \gamma^{\leq t, v} < 1$. If*

$$SNR((U_*^{\leq v})_{\mathcal{A}}) > \frac{1}{\gamma^{\leq t, v} - \beta^{\leq t, v}} \quad (50)$$

for all $\mathcal{A} \in \mathcal{M}_$, then for sufficiently small κ , the solution \tilde{U} to the \mathcal{H}_2 RFD optimization problem (29) has the same row support as $U_*^{\leq v}$.*

Proof: Follows from rearranging terms in (41), Definition 6 and letting κ tend to 0 from above. ■

Setting $\rho = 0$ increases the range Λ , as defined in (42), for which the RFD optimization problem is successful in recovering the structure of U_* , and the assumption that $\beta^{\leq t, v} / \gamma^{\leq t, v} < 1$ ensures that Assumption 1 holds. Thus Theorem 2 can be viewed as a slightly stronger, but more interpretable, set of sufficient conditions for the success of the RFD procedure.

Notice in particular that the left hand side of condition (50), i.e., the SNR, is mainly a function of the desired controller $U_*^{\leq v}$ and the closed loop performance $W^{\leq t}$ that it achieves, whereas the right hand side of (50) is mainly a function of the structure of the optimal controller and $\mathfrak{L}^{\leq t, v}$. Thus we expect controllers with sparse and evenly distributed actuation, i.e., controllers that minimize the SNR threshold $(\gamma^{\leq t, v} - \beta^{\leq t, v})^{-1}$, that act quickly and aggressively to achieve a good closed loop norm, i.e., controllers that maximize the SNR (49), to be recovered by the RFD procedure.





Figure 7: A diagram of the Stable Unidirectional Chain System case study.



(a) Fixed horizon $t = 6$. (b) Fixed controller order $v = 1$. (c) Fixed controller order $v = 2$.

Figure 8: Behavior of identifiability parameters $\gamma^{\leq t,v}$ and $\beta^{\leq t,v}$.

7 Case Study

The following case study illustrates the concepts introduced in the previous section on a concrete system that satisfies our sufficient conditions.

We consider a \mathcal{H}_2 RFD optimization with control cost $\rho_u = .1$, and the remaining generalized plant (2) state space parameters set as $B_2 = C_1 = I_{10}$, $A = \frac{1}{2}I_{10} + \frac{1}{2}Z_{10}$, and $B_1 = 1.1(E_{11} + E_{55}) + .7E_{99} + .1I_{10}$. This system is illustrated in Figure 7. This simple example is chosen in order to allow a direct computation of various bounds and parameters, and to easily interpret the propagation of inputs and disturbances.

We consider the task of recovering the optimal actuation schemes that use either 2 actuators or 3 actuators. In particular, we take the desired controller U_s , for $s = 2$ and $s = 3$, to be

$$U_s := \underset{U \in \mathcal{RH}_\infty}{\operatorname{argmin}} \quad \|Y - \mathfrak{L}(U)\|_{\mathcal{H}_2}^2 + .1\|U\|_{\mathcal{H}_2}^2 \quad (51)$$

s.t. U has at most s nonzero rows,

with open loop state response Y and map \mathfrak{L} as defined in Example 3. We solve this optimization problem by enumerating all possible actuation schemes, and we find that the optimal actuation scheme for $s = 2$ is given by actuators at nodes 1 and 5, and for $s = 3$ by actuators at nodes 1, 5 and 9.

We emphasize that the goal of this case study is to illustrate the concepts introduced in the previous section, and to help the reader understand how the various parameters affect the recovery conditions – in practice, \mathcal{M}_* and \mathcal{A}_* are not available. Further, we note that the case study presented is, as far as we are aware, the first example in the literature of a system for which convex optimization provably identifies an optimal actuation architecture.

With these optimal actuation schemes at our disposal, we vary the parameters t and v to investigate how our recovery conditions are affected. As per the discussion in §6.3, we set $\rho = 0$. We also show that for appropriate fixed controller order v and horizon t ,

increasing λ shifts the identified architecture from actuators at nodes 1, 5 and 9 to actuators at nodes 1 and 5. This is a very desirable property from an architecture design perspective, as it suggests that increasing the regularization weight λ causes the identified architecture to move to a simpler, but still optimal, actuator configuration. As predicted by Corollary 1, each optimal architecture is identified for a range of regularization weights λ . Further, as predicted by Theorem 2, the identified architectures are those for which the optimal control law achieves a small closed loop norm.

We begin by examining how the lower bound parameter $\gamma^{\leq t,v}$ and the maximum gain $\beta^{\leq t,v}$ are affected as we vary the horizon t and the controller order v . In particular, for actuation sparsity $s = 2$, we compute $\gamma^{\leq t,v}$ and $\beta^{\leq t,v}$ for (i) $t = 6$ and $v \in \{1, 2, 3\}$ (shown in Figure 8a), (ii) for $v = 1$ and $t \in \{1, 2, 3, 4\}$ (shown in Figure 8b), and (iii) for $v = 2$ and $t \in \{2, 3, 4\}$ (shown in Figure 8c). As expected, there is a decrease in the lower bound parameter $\gamma^{\leq t,v}$ and an increase in the maximum gain $\beta^{\leq t,v}$ as v increases, while both $\gamma^{\leq t,v}$ and $\beta^{\leq t,v}$ are non-decreasing for a fixed controller order v and increasing horizon t as long as $t \leq \tau_{\mathcal{M}_*}$. For this problem, the mixing time $\tau_{\mathcal{M}_*} = 5$. Further we see that $\gamma^{\leq t,v}$ begins to decrease for horizons $t > 6$ when $v = 2$.

t	$\frac{1}{\gamma^{\leq t,v} - \beta^{\leq t,v}}$	SNR ₁	SNR ₅	λ	$\ \Delta\ _{\text{act}}^*$	Bound
2	1	1.27	1.27	.8	.73	.89
3	.8	.87	.88	1.46	.91	1.03
4	.727	.732	.735	2.01	1.00	1.12
5	.7	.67	.68	2.45	1.05	1.16

Table 3: Summary of relevant values for the controller U_2 with actuators at nodes 1 and 5.

The conditions of Theorem 2 are satisfied for $v = 1$ and several values of t . For example, if we select $t = 4$, $v = 1$, $\rho = 0$ and U_* as defined in (51), we can compute $(\gamma^{\leq 4,1} - \beta^{\leq 4,1})^{-1} = .7273$, $\text{SNR}((U_*^{\leq 1})_1) = .7324$, and $\text{SNR}((U_*^{\leq 1})_5) = .7353$, thus satisfying condition (50) for each of the two actuators. Further, selecting $\lambda = 2.0119 \in \Lambda$, and using this value for λ in the truncated RFD optimization (29) recovers a solution with non-zero first and fifth rows.

Perhaps surprisingly, similar positive recovery results can be verified for all $2 \leq t \leq \tau_{\mathcal{M}_*}$ – the relevant values are summarized in Table 3. In this table, SNR_i corresponds to the SNR achieved by the controller component corresponding to the actuator at node i . Further, $\Delta := \hat{U} - U_*^{\leq v}$ is the approximation error between the architect parameter \hat{U} and underlying parameter $U_*^{\leq v}$, and the values in the “Bound” column are given by equation (52) in the appendix giving upper bounds on $\|\Delta\|_{\text{act}}^*$. It is worth noting that for $t = 5$, we do not satisfy the sufficient conditions of Theorem 2, but nonetheless recover the correct actuation architecture.

We now consider the case of actuation sparsity $s = 3$. Much as in the $s = 2$ case, we can verify that the conditions of Theorem 2 hold for $v = 1$, and $2 \leq t \leq \tau_{\mathcal{M}_*}$, where the mixing time $\tau_{\mathcal{M}_*}$ is still 5. However, since the controller with 3 actuators is able to achieve a much better closed loop norm, the SNRs are significantly larger, while the SNR threshold $(\gamma^{\leq t,v} - \beta^{\leq t,v})^{-1}$ does not change significantly. In particular, for the case of $t = 5$, we have a threshold of $(\gamma^{\leq 5,1} - \beta^{\leq 5,1})^{-1} = .82$, and SNRs of 4.04, 4.04 and 2.67 for the three actuator components.

This is consistent with our original interpretation of the closed loop state response $W^{\leq t}$

playing the role of measurement noise – the better the performance of the controller, the easier it is to identify via RFD. These experiments demonstrate that controllers with sparse and diffuse actuation schemes that achieve a small state response $W^{\leq t}$ are easy to identify as the solutions to RFD optimization problems. In summary, our analysis and case studies demonstrate that: (i) the parameter ρ can be set to 0 in the RFD optimization problem (7), even if the original model matching problem (6) had non-zero control cost ρ_u ; (ii) choosing small controller order v and horizon t can actually lead to a favorable threshold $(\gamma^{\leq t,v} - \beta^{\leq t,v})^{-1}$ (50); (iii) actuation schemes that are more evenly distributed (so that they lead to large mixing times $\tau_{\mathcal{M}_*}$ in (35)) are easier to identify; and (iv) controller components that maximize the RFD analog of a SNR (49) are more likely to satisfy our recovery conditions. These consist of controllers that have a concentration of energy in their early impulse response elements, and that achieve a closed loop with small state response component.

8 Future Work

A priori bounds on incoherence: It is of great interest to derive *a priori* bounds on the gain parameters $\alpha^{\leq t,v}$ (33) and $\beta^{\leq t,v}$ (34) in terms of the state-space parameters of the system and a lower bound on the mixing time $\tau_{\mathcal{M}_*}$ (35). This would allow for broader classes of systems to be classified as being amenable to RFD. We are currently pursuing semidefinite relaxation based methods to obtain bounds on these parameters [39].

Scalability: The scalability of the RFD framework is limited by the underlying quadratic invariance based controller synthesis algorithms upon which it is built. In order to allow the RFD framework, and distributed optimal control theory in general, to scale to large heterogeneous systems, the first author and co-authors have developed the localized optimal control framework (cf. [40] and references therein). The algorithmic component of the RFD framework has already been ported [41]; it is of interest to see if analogous recovery conditions can also be developed.

References

- [1] N. Matni and V. Chandrasekaran, “Regularization for design,” *CoRR*, vol. arXiv:1404.1972, 2014.
- [2] F. Bonsall, “A general atomic decomposition theorem and banach’s closed range theorem,” *The Quarterly Journal of Mathematics*, vol. 42, no. 1, pp. 9–14, 1991.
- [3] E. Candes, J. Romberg, and T. Tao, “Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency info.” *Info. Theory, IEEE Trans. on*, vol. 52, no. 2, pp. 489–509, Feb. 2006.
- [4] D. L. Donoho, “Compressed sensing,” *IEEE Trans. Inform. Theory*, vol. 52, pp. 1289–1306, 2006.
- [5] M. Fazel, “Matrix rank minimization with applications,” Ph.D. dissertation, PhD thesis, Stanford University, 2002.

- [6] B. Recht, M. Fazel, and P. A. Parrilo, “Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization,” *SIAM Rev.*, vol. 52, no. 3, pp. 471–501, Aug. 2010.
- [7] V. Chandrasekaran, B. Recht, P. Parrilo, and A. Willsky, “The convex geometry of linear inverse problems,” *Foundations of Computational Mathematics*, vol. 12, pp. 805–849, 2012.
- [8] P. Shah, B. N. Bhaskar, G. Tang, and B. Recht, “Linear system identification via atomic norm regularization,” in *Decision and Control (CDC), 2012 IEEE 51st Annual Conf. on*, 2012, pp. 6265–6270.
- [9] M. Fazel, H. Hindi, and S. P. Boyd, “A rank minimization heuristic with application to minimum order system approximation,” in *American Control Conference. Proc. of 2001*, vol. 6. IEEE, 2001, pp. 4734–4739.
- [10] N. Matni and A. Rantzer, “Low-rank and low-order decompositions for local system identification,” *CoRR*, vol. arXiv:1403.7175, 2014.
- [11] L. Ljung, “Some classical and some new ideas for identification of linear systems,” *Journal of Control, Automation and Electrical Systems*, vol. 24, no. 1-2, pp. 3–10, 2013.
- [12] F. Lin, M. Fardad, and M. R. Jovanovic, “Design of optimal sparse feedback gains via the alternating direction method of multipliers,” *Automatic Control, IEEE Trans. on*, vol. 58, no. 9, pp. 2426–2431, 2013.
- [13] V. Jonsson, A. Rantzer, and R. M. Murray, “A scalable formulation for engineering combination therapies for evolutionary dynamics of disease,” in *The IEEE American Control Conf. (ACC), 2014.*, 2014.
- [14] M. Fardad, F. Lin, and M. R. Jovanović, “Design of optimal sparse interconnection graphs for synchronization of oscillator networks,” *arXiv preprint arXiv:1302.0449*, 2013.
- [15] N. Dhingra, M. R. Jovanovic, and Z.-Q. Luo, “An admm algorithm for optimal sensor and actuator selection,” in *IEEE Conf. on Decision and Control (CDC)*, 2014.
- [16] N. Matni, “Communication delay co-design in \mathcal{H}_2 decentralized control using atomic norm minimization,” in *Decision and Control (CDC), 2013 IEEE 52nd Annual Conf. on*, Dec 2013, pp. 6522–6529.
- [17] —, “Communication delay co-design in \mathcal{H}_2 distributed control using atomic norm minimization,” *IEEE Transactions on Networked Control Systems*, Submitted to the, vol. arXiv:1404.4911, 2015.
- [18] S. Pequito, S. Kar, and G. J. Pappas, “Minimum cost constrained input-output and control configuration co-design problem: A structural systems approach,” *arXiv preprint arXiv:1503.02764*, 2015.

- [19] M. Rotkowitz and S. Lall, “A characterization of convex problems in decentralized control,” *Automatic Control, IEEE Trans. on*, vol. 51, no. 2, pp. 274–286, 2006.
- [20] 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.
- [21] L. Lessard and S. Lall, “Quadratic invariance is necessary and sufficient for convexity,” in *American Control Conf. 2011*, 2011, pp. 5360–5362.
- [22] K. Zhou, J. C. Doyle, K. Glover, *et al.*, *Robust and optimal control*. Prentice Hall New Jersey, 1996, vol. 40.
- [23] L. Lessard and S. Lall, “A state-space solution to the two-player decentralized optimal control problem,” in *49th Annual Allerton Conf. on Comm., Control, and Computing*. IEEE, 2011, pp. 1559–1564.
- [24] C. W. Scherer, “Structured \mathcal{H}_∞ -optimal control for nested interconnections: A state-space solution,” *arXiv preprint arXiv:1305.1746*, 2013.
- [25] 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 Conf. on*. IEEE, 2010, pp. 6722–6727.
- [26] A. Lamperski and L. Lessard, “Optimal state-feedback control under sparsity and delay constraints,” in *3rd IFAC Workshop on Distributed Estimation and Control in Networked Systems*, 2012, pp. 204–209.
- [27] A. Lamperski and J. C. Doyle, “The \mathcal{H}_2 control problem for decentralized systems with delays,” *CoRR*, vol. abs/1312.7724, 2013.
- [28] N. Matni, “Distributed control subject to delays satisfying an \mathcal{H}_∞ norm bound,” in *Decision and Control (CDC), 2014 IEEE Annual Conf. on*, Dec 2014, pp. 4006–4013.
- [29] Y.-C. Ho and K.-C. Chu, “Team decision theory and information structures in optimal control problems—part i,” *Automatic Control, IEEE Trans. on*, vol. 17, no. 1, pp. 15–22, 1972.
- [30] M. A. Dahleh and I. J. Diaz-Bobillo, *Control of uncertain systems: a linear programming approach*. Prentice-Hall, Inc., 1994.
- [31] S. Sabau and N. C. Martins, “Stabilizability and norm-optimal control design subject to sparsity constraints,” *arXiv:1209.1123*, 2012.
- [32] A. Mahajan, N. Martins, M. Rotkowitz, and S. Yuksel, “Information structures in optimal decentralized control,” in *Decision and Control (CDC), 2012 IEEE 51st Annual Conf. on*, 2012, pp. 1291–1306.
- [33] A. E. Hoerl and R. W. Kennard, “Ridge regression: Biased estimation for nonorthogonal problems,” *Technometrics*, vol. 12, pp. 55–67, 1970.

- [34] M. Yuan and Y. Lin, “Model selection and estimation in regression with grouped variables,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 68, no. 1, pp. 49–67, 2006.
- [35] H. Zou and T. Hastie, “Regularization and variable selection via the elastic net,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 67, no. 2, pp. 301–320, 2005.
- [36] G. Obozinski, L. Jacob, and J.-P. Vert, “Group lasso with overlaps: the latent group lasso approach,” *arXiv preprint arXiv:1110.0413*, 2011.
- [37] G. Raskutti, M. J. Wainwright, and B. Yu, “Restricted eigenvalue properties for correlated gaussian designs,” *The Journal of Machine Learning Research*, vol. 11, pp. 2241–2259, 2010.
- [38] E. J. Candes and T. Tao, “Decoding by linear programming,” *Info. Theory, IEEE Trans. on*, vol. 51, no. 12, pp. 4203–4215, 2005.
- [39] S. You and N. Matni, “A convex approach to sparse h-infinity analysis and synthesis,” in *Decision and Control (CDC), 2015 IEEE Annual Conf. on, Submitted.*, 2015.
- [40] Y.-S. Wang, N. Matni, and J. C. Doyle, “Localized lqg optimal control for large-scale systems,” in *Decision and Control (CDC), 2015 IEEE Annual Conf. on, Submitted.*, 2015.
- [41] —, “Localized lqr control with actuator regularization,” in *Decision and Control (CDC), 2015 IEEE Annual Conf. on, Submitted.*, 2015.

A Proof of Theorem 1

Proof: The proof of Theorem 1 centers around showing that under Assumption 1, the unique solution to the architect optimization (30) is also the unique solution of the original unconstrained optimization (29). We emphasize that at no point during the RFD process do we assume knowledge of \mathcal{A}_* or of the architect optimization problem (30).

The proof consists of two parts: we first show that if $\alpha^{\leq t, v} > 0$, the architect optimization problem (30) has a unique optimal solution \hat{U} , and control its deviation from the underlying desired controller $U_*^{\leq v}$. We then use \hat{U} and its error bound to construct a strictly dual-feasible primal/dual pair for the original RFD optimization problem (29), showing that \hat{U} is indeed its unique optimal solution as well.

Proposition 1 (Bounded Errors) *Fix a horizon $1 \leq t < \infty$, and a controller order $1 \leq v \leq t$. Assume that $\alpha^{\leq t, v}$ as defined in (33) is strictly positive, and let $\Delta := \hat{U} - U_*^{\leq v}$. Then*

$$\|\Delta\|_{act}^* \leq \frac{1}{\alpha} \left(\lambda + \left\| \left[\mathfrak{L}_{\mathcal{A}_*}^{\leq t, v} \right]^\dagger (W^{\leq t} + T^{\leq t, v}) \right\|_{act}^* + \rho \|U_*^{\leq v}\|_{act}^* \right). \quad (52)$$

Proof: It is clear that under the assumption that $\alpha^{\leq t, v} > 0$, the architect optimization problem (30) is strongly convex, and hence has a unique optimal solution \hat{U} . Letting $\Delta := \hat{U} - U_*^{\leq v}$, and using the relation (27), the optimality conditions of the architect optimization problem (30) are then given by

$$\left(\left[\mathfrak{L}_{\mathcal{A}_*}^{\leq t, v} \right]^\dagger \mathfrak{L}_{\mathcal{A}_*}^{\leq t, v} + \rho I \right) (\Delta) - \left[\mathfrak{L}_{\mathcal{A}_*}^{\leq t, v} \right]^\dagger (W^{\leq t} + T^{\leq t, v}) + \rho U_*^{\leq v} + \lambda Z + \Lambda_{\mathcal{A}_*^\perp} \ni 0,$$

where $Z \in \partial \left\| \hat{U} \right\|_{\text{act}}$ satisfies $\|Z_{\mathcal{A}_*}\|_{\text{act}}^* = 1$, $\|Z_{\mathcal{A}_*^\perp}\|_{\text{act}}^* \leq 1$, and $\Lambda_{\mathcal{A}_*^\perp} \in \mathcal{A}_*^\perp$ is the Lagrange multiplier corresponding to the architect constraint $U \in \mathcal{A}_*$. Projecting (A) onto \mathcal{A}_* , and leveraging that $\Delta \in \mathcal{A}_*$, we then obtain

$$\left(\left[\mathfrak{L}_{\mathcal{A}_*}^{\leq t, v} \right]^\dagger \mathfrak{L}_{\mathcal{A}_*}^{\leq t, v} + \rho I \right) (\Delta) = \left(\left[\mathfrak{L}_{\mathcal{A}_*}^{\leq t, v} \right]^\dagger (W^{\leq t} + T^{\leq t, v}) - \rho U_*^{\leq v} - \lambda Z_{\mathcal{A}_*} \right). \quad (53)$$

We then have the following chain of inequalities

$$\alpha^{\leq t, v} \|\Delta\|_{\text{act}}^* \leq \left\| \left(\left[\mathfrak{L}_{\mathcal{A}_*}^{\leq t, v} \right]^\dagger \mathfrak{L}_{\mathcal{A}_*}^{\leq t, v} + \rho I \right) (\Delta) \right\|_{\text{act}}^* \leq \lambda + \left\| \left[\mathfrak{L}_{\mathcal{A}_*}^{\leq t, v} \right]^\dagger (W^{\leq t} + T^{\leq t, v}) \right\|_{\text{act}}^* + \rho \|U_*^{\leq v}\|_{\text{act}}^*$$

where the first inequality follows from (33), and the second from (53) and the triangle inequality. Rearranging terms yields the error bound (52). \blacksquare

Strict dual feasibility

In order to construct a primal/dual feasible pair for optimization (29) from \hat{U} , we first set $Z_{\mathcal{A}_*}$ to be a member of the sub differential $\partial \|\cdot\|_{\text{act}}$ evaluated at \hat{U} . We now choose $Z_{\mathcal{A}_*^\perp}$ to be

$$Z_{\mathcal{A}_*^\perp} := \frac{\left(\left[\mathfrak{L}_{\mathcal{A}_*^\perp}^{\leq t, v} \right]^\dagger (W^{\leq t} + T^{\leq t, v}) - \left[\mathfrak{L}_{\mathcal{A}_*^\perp}^{\leq t, v} \right]^\dagger \mathfrak{L}_{\mathcal{A}_*^\perp}^{\leq t, v} (\Delta) \right)}{\lambda} \quad (54)$$

In doing so, we guarantee that (\hat{U}, Z) satisfy the optimality conditions of optimization (29). What remains to be shown is the $Z_{\mathcal{A}_*^\perp}$ is an element of the sub-differential. In order to do so, we show that under the assumptions of the theorem, $\|Z_{\mathcal{A}_*^\perp}\|_{\text{act}}^* < 1$. This guarantees that Z is indeed in $\partial \left\| \hat{U} \right\|_{\text{act}}$, and that $\hat{U}_{\mathcal{A}} = 0$ for all $\mathcal{A} \notin \mathcal{M}_*$.

To that end, notice that $\|Z_{\mathcal{A}_*^\perp}\|_{\text{act}}^*$ can be upper bounded by

$$\begin{aligned} & \left(\left\| \left[\mathfrak{L}_{\mathcal{A}_*^\perp}^{\leq t, v} \right]^\dagger \mathfrak{L}_{\mathcal{A}_*^\perp}^{\leq t, v} (\Delta) \right\|_{\text{act}}^* + \left\| \left[\mathfrak{L}_{\mathcal{A}_*^\perp}^{\leq t, v} \right]^\dagger (W^{\leq t} + T^{\leq t, v}) \right\|_{\text{act}}^* \right) \\ & \leq \frac{1}{\lambda} \left(\beta^{\leq t, v} \|\Delta\|_{\text{act}}^* + \left\| \left[\mathfrak{L}_{\mathcal{A}_*^\perp}^{\leq t, v} \right]^\dagger (W^{\leq t} + T^{\leq t, v}) \right\|_{\text{act}}^* \right) \\ & \leq \frac{1}{\lambda} \delta \left(\rho \|U_*^{\leq v}\|_{\text{act}}^* + \left\| \left[\mathfrak{L}_{\mathcal{A}_*}^{\leq t, v} \right]^\dagger (W^{\leq t} + T^{\leq t, v}) \right\|_{\text{act}}^* \right) \\ & + \delta + \frac{1}{\lambda} \left\| \left[\mathfrak{L}_{\mathcal{A}_*^\perp}^{\leq t, v} \right]^\dagger (W^{\leq t} + T^{\leq t, v}) \right\|_{\text{act}}^* < 1, \end{aligned}$$

where the first inequality follows from applying the triangle inequality to (54), the second from applying definition (34), the third from applying the error bound (52), and the fourth from (40). Thus we have shown that under the assumptions of the Theorem, \hat{U} is also the optimal solution of the original problem (29) – its uniqueness follows from the *local* strong convexity of the cost function around \hat{U} . Finally, if for $\mathcal{A} \in \mathcal{M}_*$, we have that (41) holds, then $\hat{U}_{\mathcal{A}} \neq 0$. ■

Proof: [Proof of Lemma 2] The gain $\alpha^{\leq t, v}$ is bounded below by

$$\begin{aligned} & \min_{\substack{\|\Delta\|_{\text{act}}^* = 1 \\ \Delta \in \mathcal{A}_*}} \max_{\mathcal{A} \in \mathcal{M}_*} \left\| \left(\left[\mathfrak{L}_{\mathcal{A}}^{\leq t, v} \right]^\dagger \mathfrak{L}_{\mathcal{A}}^{\leq t, v} + \rho I \right) \Delta_{\mathcal{A}} \right\|_{\mathcal{H}_2} - \sum_{\mathcal{B} \neq \mathcal{A} \in \mathcal{M}_*} \left\| \left[\mathfrak{L}_{\mathcal{A}}^{\leq t, v} \right]^\dagger \mathfrak{L}_{\mathcal{B}}^{\leq t, v} \Delta_{\mathcal{B}} \right\|_{\mathcal{H}_2} \\ & \geq \rho + \min_{\mathcal{B} \subseteq \mathcal{M}_*, |\mathcal{B}| \geq 1} \max_{\mathcal{A} \in \mathcal{B}} \sigma_{\min} \left(\left[\mathfrak{L}_{\mathcal{A}}^{\leq t, v} \right]^\dagger \mathfrak{L}_{\mathcal{A}}^{\leq t, v} \right) - \sum_{\mathcal{B} \neq \mathcal{A} \in \mathcal{B}} \sigma_{\max} \left(\left[\mathfrak{L}_{\mathcal{A}}^{\leq t, v} \right]^\dagger \mathfrak{L}_{\mathcal{B}}^{\leq t, v} \right), \end{aligned}$$

where the inequalities follow from the fact that $\|\Delta\|_{\text{act}}^* = 1$ implies that there exists $\mathcal{A} \in \mathcal{M}_*$ such that $\|\Delta_{\mathcal{A}}\|_{\mathcal{H}_2} = 1$, and the definition of the respective norms. The derivation of the bound on $\beta^{\leq t, v}$ is similar, and hence omitted. ■