university of groningen

faculty of science and engineering

mathematics and applied mathematics

System-Level Design for Distributed Optimal Control

Bachelor’s Thesis in Mathematics

June 2024

Student: Á. Karakai

First supervisor: Prof. dr. ir. B. Besselink

Second assessor: Prof. dr. M. K. Camlibel

Abstract

A standard problem in distributed optimal control is to design a controller for each system in a network of interconnected systems to minimise the amplification of external disturbances. In addition, we often wish to constrain communication among controllers so that each control input is computed from local information. Traditional optimal control methods, such as the Youla parametrisation, are of limited utility in this setting, as they can often only encode communication sparsity as a non-convex constraint. The recently developed system-level approach aims to circumvent that issue in the case of discrete-time LTI systems. To that end, it introduces the system response, which describes the effects of disturbances on the closed-loop dynamics, and reformulates the entire design cycle in terms of that. In this thesis, we present a continuous-time adaptation of the system-level approach, from parametrisation to controller design.

Contents

Chapter 1

Introduction

Modern control problems are often concerned with networks of a very large number of interacting systems that are each equipped with their own controllers. We typically wish to design these local controllers in such a way that they stabilise the entire interconnection. On top of stability, we often have the additional objective that the global system should achieve optimal performance with respect to a chosen cost criterion.

An illustrative example is the smart electricity grid with millions of homes equipped with both solar panels and appliances with large power demand, such as charging stations for electric cars. The solar panels generate power at certain times and feed some of it back to the grid, while at other times, some houses may need more electricity and a power plant needs to be turned on. With all these fast changes, the grid must remain stable at all times. Moreover, it is desirable to optimise, for example, the cost of electricity or the environmental impact.

While our objectives are formulated in terms of the entire interconnection of systems, the controllers are attached to specific systems and their communication may be restricted by various factors such as a predetermined network topology or limited bandwidth. Therefore, we aim to design controllers such that each control input can be computed from locally available information. In the context of the smart grid example, it is easy to see that, for example, the control of a given charging station should not depend on the state of a solar panel hundreds of kilometres away. This requirement is one of the central challenges of distributed control and is one of the main topics of this thesis.

To formalise our problem setup, consider a system $G(s)$ that encodes the entire interconnection of local systems, and assume it is equipped with a feedback controller $K(s)$. That is, $K(s)$ is a virtual centralised controller, constructed as a collection of local controllers. This setup is shown in Figure 2.3. The system $G(s)$ is assumed to have two inputs, the control input u and external disturbance w, and two outputs, namely the measurement y and regulated output z. The measurement comprises the information available for the controller, while the regulated output is some quantity of interest that is to be kept minimal. Specifically, the optimal control problem is to find $K(s)$ such that it minimises the amplification from w to z while stabilising $G(s)$. To satisfy the requirement of restricted information sharing among local controllers, we may impose further constraints on the structure of $K(s)$.

Then, the distributed optimal control problem can be stated as an optimisation problem where the objective function is a chosen norm of the transfer function from w to z and the constraints ensure that $K(s)$ stabilises $G(s)$ and has the desired structure. In its direct formulation, that optimisation problem is not convex. If the structural constraints on $K(s)$ are dropped, the famous Youla parametrisation [1], [2] can be used to to obtain an equivalent convex problem. However, information sharing constraints on $K(s)$ can be expressed as convex constraints on the Youla parameter if and only if the strict condition of quadratic invariance is satisfied [3], [4]. Quadratic invariance depends on both the chosen constraints and the underlying network topology of $G(s)$, and it severely limits the range of structural constraints that we can impose on the controller.

These limitations were the main motivation behind the development of the system-level approach to distributed optimal control, presented in [5]–[7]. That theory provides a new framework for the entire design cycle in the case of discrete-time linear time-invariant (LTI) systems. The main idea is to parametrise stabilising controllers via the system response, defined as a map from disturbances to closed-loop dynamics, and to express design specifications in terms of that. In particular, the optimal control problem is reformulated such that the decision variable is the system response, and information sharing constraints are phrased as sparsity patterns of that. This leads to a convex problem for any sparsity pattern, and the enforced communication constraints carry over to the corresponding controller, which is constructed based on the computed optimal system response.

It is shown in [6] that the system-level approach is suitable for the convex parametrisation of a strictly wider class of distributed optimal control problems under information sharing constraints than the traditional Youla approach or any other known alternative.

To make these benefits available for continuous-time systems, in this thesis, we develop the continuous-time equivalent of the system-level approach. With that objective, our contributions are the following. First, through arguments analogous to those used in the discrete-time case, we introduce the continuous-time system-level parametrisation for state feedback problems through the definition and algebraic characterisation of the system response, and show that it provides a new parametrisation of all internally stabilising controllers. We extend this discussion by a new characterisation of the system response via geometric conditions on its state-space realisations.

Second, we show how information sharing constraints can be expressed as convex constraints on the system response and that they carry over to the controller, as in the discrete-time case. In addition, we present a number of minor results regarding the feasibility of such constraints, an explicit construction for a system response with communication confined to the physical network topology in a full control problem, and a direct proof that quadratically invariant subspace constraints on the controller can be expressed as linear constraints on the system response.

Third, recognising that the resulting problem is infinite-dimensional and that the methods to overcome this obstacle in discrete-time are not applicable in continuous-time, we present two approaches that solve the synthesis problem with some limitations.

Accordingly, the structure is as follows. In Chapter 2, we discuss some mathematical preliminaries and derive the Youla parametrisation. Chapter 3 introduces the system-level parametrisation, which is used in Chapter 4 to discuss communication restrictions and performance constraints. Finally, Chapter 5 covers the remaining challenges in controller synthesis and provides partial solutions for them. Moreover, it gives a brief introduction to how appropriate constraints can aid localised synthesis.

Chapter 2

Background

In this chapter, we introduce the general setup in which the system-level parametrisation is utilised. To that end, we present some mathematical preliminaries, the notion of internally stabilising controllers, and the general optimal control problem. In addition, we introduce the Youla parametrisation, a well-known and celebrated method of characterising all internally stabilising controllers of a system in such a way that the objective function in the optimal control problem becomes convex.

Finally, we discuss the idea of distributed optimal control under structural communication constraints and examine the conditions under which we can express such restrictions as convex constraints on the Youla parameter. The latter also serves as a motivation for the system-level parametrisation.

2.1 Preliminaries

In this section, we provide a brief discussion of some of the mathematical preliminaries of the theory in the rest of the thesis. All of this is standard material, based on [8] and [9].

2.1.1 Real-Rational Transfer Functions

In most of our analysis, we represent LTI systems with their transfer functions. Thus, we give a short overview of some of their properties.

First, it is well-known that LTI systems which admit standard state-space realisations have real-rational transfer functions. In the multivariable case, that means matrices of real-rational functions. We denote the space of matrices of size $m \times n$ with real-rational entries as $\mathbb{R}(s)^{m\times n}$. For a square rational matrix $T(s) \in \mathbb{R}(s)^{n\times n}$, invertibility and inverses are defined the same way as for real matrices.

Second, we call a rational function proper if the degree of the denominator polynomial is greater than or equal to that of the numerator polynomial. It is called strictly proper if that inequality is strict. A rational matrix $T(s) \in \mathbb{R}(s)^{m\times n}$ is called proper if all of its entries are proper and strictly proper if all of its entries are strictly proper. The transfer function of any LTI system that admits a state-space realisation is proper, and it is strictly proper if realisations have no direct feedthrough. The following lemma provides a useful characterisation of properness.

Lemma 1. *A rational function $T(s) \in \mathbb{R}(s)^{m\times n}$ is proper if and only if*

$$T_\infty = \lim_{s\to\infty} T(s)$$

exists, and strictly proper if and only if $T_\infty = 0$.

We call an invertible square rational matrix bi-proper if it is proper and has proper inverse.

Second, poles and stability are crucially important concepts. After cancelling common factors of the numerator and denominator polynomials, a pole of a rational transfer function is a root of its denominator, or, in the multivariable case, a root of the denominator of any of its entries. A transfer function is called stable if all of its poles are contained in $\mathbb{C}_- = \{s \in \mathbb{C} \mid \Re(s) < 0\}$. The interpretation is that the output remains bounded for any bounded input. In case $T(s) \in \mathbb{R}(s)^{n \times n}$ is invertible, it is called bi-stable if it is stable and has stable inverse.

2.1.2 The Spaces $\mathcal{RH}_\infty$ and $\mathcal{RH}_2$

With the notions of stability and properness, we can define two useful spaces of transfer functions, namely $\mathcal{RH}_\infty$ and $\mathcal{RH}_2$. To that end, we first define the space

$$\mathcal{H}_\infty^{m \times n} = \{T : \mathbb{C} \to \mathbb{C}^{m \times n} \mid T \text{ is bounded and analytic on } \mathbb{C} \backslash \mathbb{C}_-\},$$

which is a Banach space with the norm

$$\|T\|_{\mathcal{H}_\infty} = \sup_{\Re(s) \geq 0} \Big(\lambda_{\max}(T(s)^H T(s))\Big)^{1/2},$$

where $\lambda_{max}(\cdot)$ returns the maximal eigenvalue and $T(s)^H$ is the Hermitian transpose of $T(s)$. Then, we call the restriction of $\mathcal{H}_\infty^{m \times n}$ to real-rational functions $\mathcal{RH}_\infty^{m \times n}$.

Similarly, we define

$$\mathcal{H}_2^{m \times n} = \{T : \mathbb{C} \to \mathbb{C}^{m \times n} \mid T \text{ is analytic on } \mathbb{C} \backslash \mathbb{C}_- \text{ and } \|T\|_{\mathcal{H}_2} < \infty\},$$

where the $\mathcal{H}_2$-norm is given by

$$\|T\|_{\mathcal{H}_2} = \left(\sup_{\sigma > 0} \int_{\mathbb{R}} \sum_{k=1}^{m} \sum_{j=1}^{n} \left|T_{kj}(\sigma + i\omega)\right|^2 d\omega\right)^{1/2}.$$

The restriction of $\mathcal{H}_2^{m \times n}$ to real-rational functions is called $\mathcal{RH}_2^{m \times n}$. By convention, we only write the dimensions explicitly if they are not clear from the context.

It can be shown that $\mathcal{RH}_\infty$ and $\mathcal{RH}_2$ are both vector spaces over $\mathbb{R}$. Moreover, a well-known result which is used frequently in this thesis is that $\mathcal{RH}_\infty$ coincides with the set of proper real-rational and stable transfer functions, while $\mathcal{RH}_2$ is equal to the set of strictly proper real-rational and stable transfer functions. In particular, it follows that $\mathcal{RH}_2 \subset \mathcal{RH}_\infty$.

2.1.3 Well-Posed Interconnections

Recall from Section 2.1.1 that transfer functions of LTI systems that admit state-space realisations are always proper real-rational. This idea generalises to interconnections of systems. In that case, if the interconnection has a state-space realisation, then so does the system from any input to any internal signal or output in the interconnection. Accordingly, an interconnection is called well-posed if the transfer function between any two signals in the interconnection is well-defined, real-rational, and proper.

2.2 Internally Stabilising Controllers

Consider the simplest feedback control setup, where a plant $P(s)$, with input u and output y, is equipped with a controller $K(s)$, such that $u = K(s)y$.

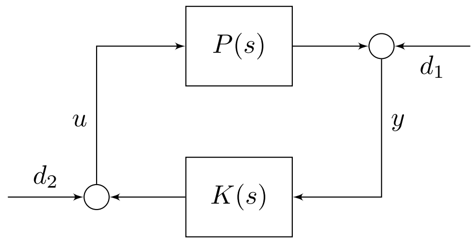

Figure 2.1: Feedback interconnection A to determine internal stability.

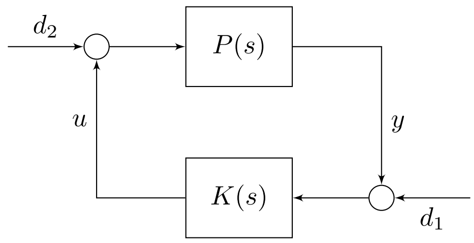

Figure 2.2: Feedback interconnection B to determine internal stability.

The first objective in practically every control problem is stability. In particular, whenever we use a feedback controller, we would like the closed-loop system to be asymptotically stable. To test that, we introduce the perturbations d_1 and d_2 to the loop, as shown in Figure 2.1, and check if the all signals in the feedback loop remain bounded for bounded d_1, d_2. The notion of an internally stabilising controller captures exactly that.

Definition 1. *The feedback controller given by a transfer function $K(s)$ internally stabilises the plant with transfer function $P(s)$ if, with the interconnection shown in Figure 2.1, the system from (d_1, d_2) to (y, u) is stable.*

Due to the utmost importance of this property, internally stabilising controllers are often referred to as admissible controllers in the optimal control setting. Naturally, if additional constraints are present, then the meaning of that term is adapted accordingly.

Although Definition 1 is used in the relevant literature [8, Chapter 5.3], [9, Chapter 6.2], its analogue with the interconnection of Figure 2.1 replaced with that of Figure 2.2 gives an equivalent definition.

Theorem 1. *The feedback controller given by the transfer function $K(s)$ internally stabilises the system $P(s)$ with the definition using the interconnection of Figure 2.1 if and only if it does so with the definition using the interconnection of Figure 2.2.*

Proof. Consider the interconnection of Figure 2.1. Assuming well-posedness of the feedback loop, direct calculations reveal the relations

$$\begin{aligned} y &= (I - P(s)K(s))^{-1}(d_1 + P(s)d_2), \\ u &= (I - K(s)P(s))^{-1}(d_2 + K(s)d_1). \end{aligned}$$

Hence, the system from (d_1, d_2) to (y, u) is described by

$$\begin{bmatrix} y \\ u \end{bmatrix} = H_A(s) \begin{bmatrix} d_1 \\ d_2 \end{bmatrix},$$

where

$$H_A(s) = \begin{bmatrix} (I - P(s)K(s))^{-1} & (I - P(s)K(s))^{-1}P(s) \\ (I - K(s)P(s))^{-1}K(s) & (I - K(s)P(s))^{-1} \end{bmatrix}. \tag{2.1}$$

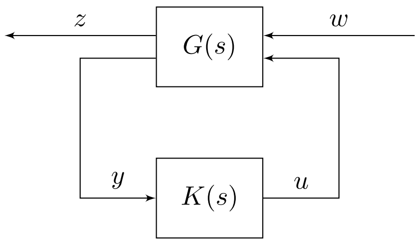

Figure 2.3: The interconnection considered in the general optimal control problem.

On the other hand, the interconnection of Figure 2.2 yields

$$y = (I - P(s)K(s))^{-1}(P(s)K(s)d_1 + P(s)d_2),$$
$$u = (I - K(s)P(s))^{-1}(K(s)P(s)d_2 + K(s)d_1).$$

Then, we have

$$\begin{bmatrix} y \\ u \end{bmatrix} = H_B(s) \begin{bmatrix} d_1 \\ d_2 \end{bmatrix},$$

with

$$H_B(s) = \begin{bmatrix} (I - P(s)K(s))^{-1}P(s)K(s) & (I - P(s)K(s))^{-1}P(s) \\ (I - K(s)P(s))^{-1}K(s) & (I - K(s)P(s))^{-1}K(s)P(s) \end{bmatrix}. \tag{2.2}$$

Comparing (2.1) and (2.2), it is clear that only the diagonal blocks differ. By symmetry, it suffices to consider the upper left block. For that, we immediately have that

$$(I - P(s)K(s))^{-1} = I + (I - P(s)K(s))^{-1}P(s)K(s),$$

and hence $(I - P(s)K(s))^{-1} \in \mathcal{RH}_\infty$ if and only if $(I - P(s)K(s))^{-1}P(s)K(s) \in \mathcal{RH}_\infty$. Then, the statement of the theorem follows directly. □

The concept of an internally stabilising controller prompts the definition of a stabilisable plant.

Definition 2. *Consider the interconnection in Figure 2.1. The plant $P(s)$ is called stabilisable if there exists a feedback controller $K(s)$ that internally stabilises $P(s)$.*

Many control problems deal with plants with additional inputs and outputs. That is, plants that have an exogenous input or disturbance w and regulated output z, on top of the control input u and measurement y. The system $G(s)$ in Figure 2.3 is an example of that. In this case, we may modify the definition of internal stability to require stability of the transfer function from (w, d_1, d_2) to (z, y, u). As above, stabilisability refers to the existence of an internally stabilising feedback controller. By [10, Theorem 2], the conditions for an internally stabilising controller in this setup can be simplified to resemble those for the simpler interconnection above, assuming the plant is stabilisable. Namely, we have the following result.

Theorem 2. *Consider the interconnection of Figure 2.4, and assume $G(s)$ is stabilisable. Then, the transfer function from (w, d_1, d_2) to (z, y, u) is stable if and only if the transfer function from (d_1, d_2) to (y, u) is.*

The proof of Theorem 2 is omitted here; it can be found in [10, Chapter 4].

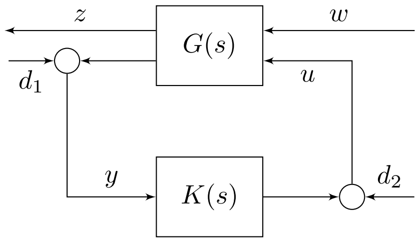

Figure 2.4: The interconnection of Figure 2.3 with external perturbations to determine stability.

2.3 Optimal Control

On top of achieving stability, optimal control also aims for efficiency. That is, we typically wish to minimise some cost criterion associated with the closed-loop dynamics of the system over time. Accordingly, in its most general formulation, the optimal control problem is to minimise a chosen norm of the transfer function from the disturbance w to the regulated output z of the interconnection in Figure 2.3, where $G(s)$ is the plant and $K(s)$ is the controller [8]. Assuming the plant is of the general form

$$G(s) = \begin{bmatrix} G_{11}(s) & G_{12}(s) \\ G_{21}(s) & G_{22}(s) \end{bmatrix},$$

so that

$$\begin{bmatrix} z \\ y \end{bmatrix} = \begin{bmatrix} G_{11}(s) & G_{12}(s) \\ G_{21}(s) & G_{22}(s) \end{bmatrix} \begin{bmatrix} w \\ u \end{bmatrix},$$

the optimal control problem is formalised as

$$\begin{aligned} &\min_{K(s)} && \left\| G_{11}(s) + G_{12}(s)K(s)(I - G_{22}(s)K(s))^{-1}G_{21}(s) \right\| \\ &\text{subject to} && K(s) \text{ internally stabilises } G(s). \end{aligned} \tag{2.3}$$

We assume $G(s)$ is stabilisable, otherwise the control problem is not feasible. Then, in accordance with Theorem 2, $K(s)$ internally stabilises $G(s)$ if and only if the transfer function from (d_1, d_2) to (y, u), as in Figure 2.4, is stable. Since that part of the closed-loop dynamics is determined by $G_{22}(s)$ and $K(s)$, we have internal stability if and only if $K(s)$ satisfies Definition 1 with the plant $P(s) = G_{22}(s)$. Thus, the internally stabilising controllers of $G(s)$ and $G_{22}(s)$ coincide.

It is immediate that (2.3) is a non-convex optimisation problem. Therefore, while state-space methods are known for certain cases, for example, the linear quadratic regulator (LQR) problem [11, Chapter 10], a convex reparametrisation is highly desirable. A well-known approach is to utilise the Youla parametrisation, originally presented in [1], [2]. In what follows, we briefly discuss the Youla parametrisation, following [9, Chapter 6.2]. For the sake of conciseness, we suppress the dependence on s here.

Consider the interconnection shown in Figure 2.1, and let $U, V, \tilde{U}, \tilde{V} \in \mathcal{RH}_\infty$ form a doubly coprime factorisation of P over $\mathcal{RH}_\infty$. That is,

$$P = VU^{-1} = \tilde{U}^{-1}\tilde{V},$$

where U and $\tilde{U}$ are bi-proper, and there exist corresponding Bézout coefficients $X, Y, \tilde{X}, \tilde{Y} \in \mathcal{RH}_\infty$ such that

$$\begin{bmatrix} X & Y \\ -\tilde{V} & \tilde{U} \end{bmatrix} \begin{bmatrix} U & -\tilde{Y} \\ V & \tilde{X} \end{bmatrix} = I. \tag{2.4}$$

Note that such doubly coprime factorisation exists for any system that is stabilisable via a feedback controller [9, Chapter 3.3]. With that, the Youla parametrisation is introduced by the following theorem.

Theorem 3. *With the feedback interconnection of Figure 2.1, the controller K internally stabilises P if and only if there exists $Q \in \mathcal{RH}_\infty$ such that*

$$K = (X + Q\tilde{V})^{-1}(-Y + Q\tilde{U})$$

or, equivalently,

$$K = (-\tilde{Y} + UQ)(\tilde{X} + VQ)^{-1}.$$

Moreover, K is well-posed if and only if $X_\infty + Q_\infty \tilde{V}_\infty$ or, equivalently, $\tilde{X}_\infty + V_\infty Q_\infty$ is nonsingular.

Proof. Let H denote the transfer function from (d_1, d_2) to (y, u). As shown in the proof of Theorem 1,

$$\begin{aligned} H &= \begin{bmatrix} (I - PK)^{-1} & (I - PK)^{-1}P \\ (I - KP)^{-1}K & (I - KP)^{-1} \end{bmatrix} \\ &= \begin{bmatrix} I & P \\ 0 & I \end{bmatrix} + \begin{bmatrix} P \\ I \end{bmatrix} K(I - KP)^{-1} \begin{bmatrix} I & P \end{bmatrix} \\ &= \begin{bmatrix} I & P \\ 0 & I \end{bmatrix} + \begin{bmatrix} V \\ U \end{bmatrix} U^{-1}K(I - PK)^{-1}\tilde{U}^{-1} \begin{bmatrix} \tilde{U} & \tilde{V} \end{bmatrix}, \end{aligned}$$

and, by Definition 1, internal stability is equivalent to $H \in \mathcal{RH}_\infty$. Note that (2.4) implies that the factors on its LHS are bi-stable. It is easily verified that if T and T' are two square transfer functions and T is bi-stable, then T' is stable if and only if $T'T$ or, equivalently, TT' is. Hence, $H \in \mathcal{RH}_\infty$ if and only if

$$\begin{aligned} \tilde{H} &= \begin{bmatrix} 0 & I \\ I & 0 \end{bmatrix} \begin{bmatrix} X & Y \\ -\tilde{V} & \tilde{U} \end{bmatrix} \begin{bmatrix} 0 & I \\ I & 0 \end{bmatrix} H \begin{bmatrix} 0 & I \\ -I & 0 \end{bmatrix} \begin{bmatrix} U & -\tilde{Y} \\ V & \tilde{X} \end{bmatrix} \begin{bmatrix} 0 & I \\ I & 0 \end{bmatrix} \\ &= \begin{bmatrix} \tilde{U} & -\tilde{V} \\ Y & X \end{bmatrix} H \begin{bmatrix} \tilde{X} & V \\ \tilde{Y} & -U \end{bmatrix} \in \mathcal{RH}_\infty. \end{aligned}$$

We compute $\tilde{H}$ as the sum of

$$\begin{bmatrix} \tilde{U} & -\tilde{V} \\ Y & X \end{bmatrix} \begin{bmatrix} I & P \\ 0 & I \end{bmatrix} \begin{bmatrix} \tilde{X} & V \\ \tilde{Y} & -U \end{bmatrix} = \begin{bmatrix} \tilde{U}\tilde{X} & \tilde{U}V \\ Y\tilde{X} + U^{-1}\tilde{Y} & YV - I \end{bmatrix}$$

and

$$\begin{aligned} &\begin{bmatrix} \tilde{U} & -\tilde{V} \\ Y & X \end{bmatrix} \begin{bmatrix} V \\ U \end{bmatrix} U^{-1}K(I - PK)^{-1}\tilde{U}^{-1} \begin{bmatrix} \tilde{U} & \tilde{V} \end{bmatrix} \begin{bmatrix} \tilde{X} & V \\ \tilde{Y} & -U \end{bmatrix} \\ &= \begin{bmatrix} 0 \\ I \end{bmatrix} U^{-1}K(I - PK)^{-1}\tilde{U}^{-1} \begin{bmatrix} I & 0 \end{bmatrix}. \end{aligned}$$

Hence,

$$\tilde{H} = \begin{bmatrix} \tilde{U}\tilde{X} & \tilde{U}V \\ Y\tilde{X} & YV - I \end{bmatrix} + \begin{bmatrix} 0 \\ I \end{bmatrix} (U^{-1}\tilde{Y} + U^{-1}K(I - PK)^{-1}\tilde{U}^{-1}) \begin{bmatrix} I & 0 \end{bmatrix},$$

where the first term is stable by the assumptions in the doubly coprime factorisation. Hence, defining

$$Q = U^{-1}\tilde{Y} + U^{-1}K(I - PK)^{-1}\tilde{U}^{-1},$$

$H \in \mathcal{RH}_\infty$ if and only if $Q \in \mathcal{RH}_\infty$.

The details of the inversion of the expression to recover K require the theory of linear fractional transformations, which is out of the scope of this thesis. Comprehensive discussions are found in, for example, [9, Chapter 5.2] and [8, Chapter 10]. That theory, [9, Proposition 5.6] in particular, implies that the map $K \mapsto Q$ is bijective if and only if $I + P_\infty Q_\infty$ is nonsingular. If $I + P_\infty Q_\infty$ is nonsingular, K is given by

$$\begin{aligned} K &= -\tilde{Y}\tilde{X}^{-1} + (U + \tilde{Y}\tilde{X}^{-1}V)Q(I + \tilde{X}^{-1}VQ)^{-1}\tilde{X}^{-1} \\ &= (-\tilde{Y}\tilde{X}^{-1}(\tilde{X} + VQ) + (U + \tilde{Y}\tilde{X}^{-1}V)Q)(\tilde{X} + VQ)^{-1} \\ &= (-\tilde{Y} + UQ)(\tilde{X} + VQ)^{-1}. \end{aligned}$$

The other claimed expression in the theorem follows by direct computation. □

Theorem 3 allows us to express all stabilising controllers of the system G in terms of the Youla parameter Q. Let $P = G_{22}$, and consider the setup of Theorem 3. With that, the objective function in (2.3) can be rewritten as

$$\begin{aligned} &\left\| G_{11} + G_{12}K(I - G_{22}K)^{-1}G_{21} \right\| \\ &= \left\| G_{11} + G_{12}(-\tilde{Y} + UQ)(\tilde{X} + VQ)^{-1}(I - G_{22}(-\tilde{Y} + UQ)(\tilde{X} + VQ)^{-1})^{-1}G_{21} \right\| \\ &= \left\| G_{11} + G_{12}(-\tilde{Y} + UQ)\left((I - G_{22}(-\tilde{Y} + UQ)(\tilde{X} + VQ)^{-1})(\tilde{X} + VQ)\right)^{-1} G_{21} \right\| \\ &= \left\| G_{11} + G_{12}(-\tilde{Y} + UQ)\left((I + (\tilde{U}^{-1} - \tilde{X} - VQ)(\tilde{X} + VQ)^{-1})(\tilde{X} + VQ)\right)^{-1} G_{21} \right\| \\ &= \left\| G_{11} - G_{12}\tilde{Y}G_{21} + G_{12}UQ\tilde{U}G_{21} \right\|. \end{aligned} \tag{2.5}$$

Hence, the optimisation problem (2.3) can be equivalently phrased as

$$\begin{aligned} \min_{Q} \quad & \left\| G_{11} - G_{12}\tilde{Y}G_{21} + G_{12}UQ\tilde{U}G_{21} \right\| \\ \text{subject to} \quad & Q \in \mathcal{RH}_\infty. \end{aligned} \tag{2.6}$$

Hence, the Youla parametrisation provides a way to formulate the general optimal control problem as a convex optimisation problem. This is a major advantage regarding the existence and computation of solutions.

2.4 Optimal Control for Distributed Systems

Modern control problems are often concerned with engineering systems that involve a very large number of interconnected and interacting parts. Although we may still model the interconnection of local systems as one global system, we typically aim for distributed control, where each part is equipped with its own controller instead of employing one centralised controller.

In many realistic scenarios, communication among local controllers is constrained by a predetermined network topology, limited bandwidth, and similar factors. Therefore, we wish to design distributed controllers such that each local control action only depends on information from a prescribed subset of local systems and controllers. If G is the global system and K is the corresponding collection of distributed controllers into a centralised controller, we can impose such constraints by requiring K to satisfy a desired sparsity pattern. Namely, $K_{ij} = 0$, where the indices correspond to blocks of appropriate dimensions, signifies that subsystem i is controlled independently of subsystem j, and hence no communication is needed between the corresponding controllers.

In problem (2.3), such sparsity requirements can be directly imposed as linear subspace constraints on the decision variable K. However, as established in Section 2.3, that optimisation problem is not convex and warrants the reformulation (2.6) in terms of the Youla parameter Q. At the same time, Theorem 3 suggests that constraints on the structure of K can not be as simply expressed in terms of Q.

To discuss when and how subspace constraints, such as a sparsity pattern, on K can be translated to convex constraints on Q, we introduce the concept of quadratic invariance (QI), based on the original paper on the topic [3], but adapted to the use case at hand.

Definition 3. *Let $\mathcal{S}$, a set of rational transfer functions, define a constraint on the controller K. We call $\mathcal{S}$ quadratically invariant under G_{22} if*

$$KG_{22}K \in \mathcal{S}$$

for all $K \in \mathcal{S}$.

It is shown in [3] and [4] that a linear subspace constraint $\mathcal{C}$ for the controller K can be imposed on Q in a convex manner if and only if $\mathcal{C}$ is quadratically invariant under G_{22}.

Suppose $\mathcal{C}$ is a subspace of all real-rational transfer functions of appropriate dimensions and is quadratically invariant under G_{22}, and define $h(K) = -K(I - G_{22}K)^{-1}$. Then, by the results in [3], $\mathcal{C} = h(\mathcal{C})$ and h is involutive, meaning that $h(h(K)) = K$. Therefore, $K \in \mathcal{C}$ if and only is $h(K) \in \mathcal{C}$. Define $g(Q) = \tilde{Y} - UQ\tilde{U}$, so that (2.5) implies $g(Q) = h(K)$. Noting that g is an invertible affine map since U and $\tilde{U}$ are nonsingular, the optimal control problem (2.3) with a subspace constraint $\mathcal{C}$ on K that is quadratically invariant under G_{22} can be formulated, following [6], as

$$\begin{aligned} \min_{Q} \quad & \left\| G_{11} - G_{12}\tilde{Y}G_{21} + G_{12}UQ\tilde{U}G_{21} \right\| \\ \text{subject to} \quad & Q \in \mathcal{RH}_{\infty}, \quad g(Q) \in \mathcal{C}. \end{aligned} \tag{2.7}$$

That is, if the constraints on the controller are quadratically invariant under G_{22}, we can formulate the constrained optimal control problem as a convex optimisation problem in the Youla parameter. Since the constraints are preserved under the involution h, they carry over to the controller K. That said, constraints that encode a sparsity pattern for K are never quadratically invariant under a plant G_{22} if the underlying system is strongly connected [6], i.e., information can propagate from any subsystem to any other subsystem. It is important to emphasise that the network topology of a strongly connected system need not be a complete graph. For example, a chain of systems with bidirectional communication can be strongly connected, depending on the specifics of the local systems.

Hence, the Youla parametrisation is only useful in special cases of distributed optimal control. Chapter 4 shows how these limitations are overcome using the system-level parametrisation.

Chapter 3

System-Level Parametrisation

Motivated by the limited utility of the Youla parametrisation for distributed optimal control, the series of papers [5]–[7] introduces a new approach to parametrise all internally stabilising controllers. The core idea of this new approach is to consider the feedback interconnection of Figure 2.3 and represent a controller by the resulting closed-loop response to disturbances. In the following, we present the details of this idea, adapted to continuous-time systems. We only discuss the state feedback case, but emphasise that [6] extends the approach to output feedback problems in a way that suggests the possibility of a similar extension in continuous time as well.

We consider finite-dimensional linear time-invariant systems of the form

$$\Sigma : \begin{cases} \dot{x} = Ax + B_1 w + B_2 u \\ z = C_1 x + D_{11} w + D_{12} u \\ y = x, \end{cases} \tag{3.1}$$

where x, w, u, z, y are the state, external disturbance, control input, regulated output, and measurement, respectively. In addition, we denote the external disturbance on the state as $\delta_x = B_1 w$. Write the Laplace transforms of these variables as $\hat{x}, \hat{w}, \hat{u}, \hat{z}, \hat{y}, \hat{\delta}_x$. Then, using the notation of [8], the system Σ can be written in Laplace domain as

$$G(s) = \begin{bmatrix} G_{11}(s) & G_{12}(s) \\ G_{21}(s) & G_{22}(s) \end{bmatrix} = \left[\begin{array}{c|cc} A & B_1 & B_2 \\ \hline C_1 & D_{11} & D_{12} \\ I & 0 & 0 \end{array}\right], \tag{3.2}$$

so that

$$\begin{bmatrix} \hat{z} \\ \hat{y} \end{bmatrix} = \begin{bmatrix} G_{11}(s) & G_{12}(s) \\ G_{21}(s) & G_{22}(s) \end{bmatrix} \begin{bmatrix} \hat{w} \\ \hat{u} \end{bmatrix}.$$

As in Figure 2.3, we connect a feedback controller

$$K(s) = \left[\begin{array}{c|c} A_K & B_K \\ \hline C_K & D_K \end{array}\right] \tag{3.3}$$

to the system such that $\hat{u} = K(s)\hat{y} = K(s)\hat{x}$. Then, we encode the closed-loop dynamics by the transfer functions from $\hat{\delta}_x$ to $\hat{x}$ and $\hat{u}$. This prompts the following definition.

Definition 4. *Consider the LTI system* (3.1) *equipped with a feedback controller. Its system response is the pair of transfer functions* $\big(R(s), M(s)\big)$ *such that*

$$\begin{bmatrix} \hat{x} \\ \hat{u} \end{bmatrix} = \begin{bmatrix} R(s) \\ M(s) \end{bmatrix} \hat{\delta}_x.$$

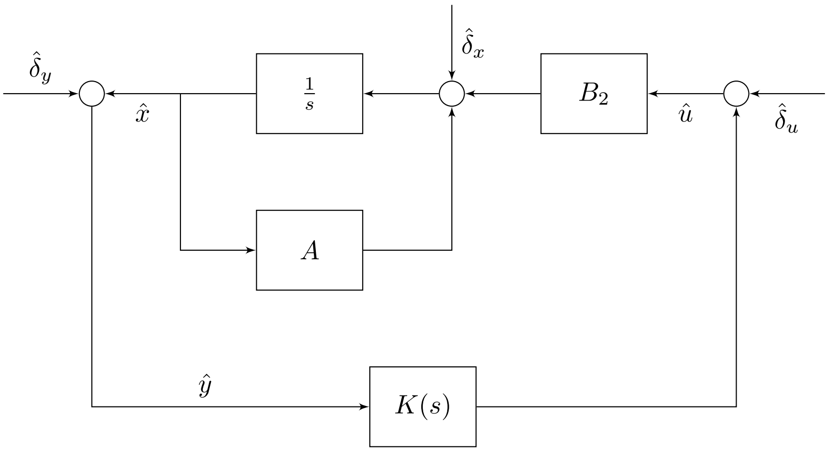

Figure 3.1: Block diagram of $G_{22}(s)$ together with the disturbance $\hat{\delta}_x$ and the feedback controller $K(s)$.

Given a controller $K(s)$, (3.1) implies

$$(sI - A)\hat{x} = \hat{\delta}_x + B_2\hat{u} = \hat{\delta}_x + B_2K(s)\hat{x}, \tag{3.4}$$

which immediately leads to

$$\begin{aligned} R(s) &= (sI - A - B_2K(s))^{-1}, \\ M(s) &= K(s)(sI - A - B_2K(s))^{-1}. \end{aligned} \tag{3.5}$$

The system responses achievable by an appropriately chosen internally stabilising controller $K(s)$ are given by the following theorem.

Theorem 4. *The conditions*

1. $\begin{bmatrix} sI - A & -B_2 \end{bmatrix} \begin{bmatrix} R(s) \\ M(s) \end{bmatrix} = I,$
2. $R(s), M(s) \in \mathcal{RH}_2,$

characterise all system responses of the system (3.1) *achievable by an internally stabilising feedback controller* $K(s)$ *of the form* (3.3).

Proof. Necessity of the first condition is immediate from (3.5). Consider the interconnection shown in Figure 3.1. By Definition 1 and Theorem 2, the controller given by $K(s)$ is internally stabilising if and only if the system from $(\hat{\delta}_y, \hat{\delta}_u)$ to $(\hat{y}, \hat{u})$ is stable. It is clear from the proof of Theorem 1 that we may replace $\hat{y}$ with $\hat{x}$ in that test. Since we would like to examine the effect of $\hat{\delta}_x$, we exploit Theorem 2 and consider the equivalent condition of stability of the transfer function $H(s)$ from $(\hat{\delta}_x, \hat{\delta}_y, \hat{\delta}_u)$ to $(\hat{x}, \hat{u})$. To determine $H(s)$, we note that

$$(sI - A)\hat{x} = B_2\hat{u} + \hat{\delta}_x = B_2(\hat{\delta}_u + K(s)(\hat{\delta}_y + \hat{x})) + \hat{\delta}_x,$$

which leads to

$$\hat{x} = R(s)(B_2\hat{\delta}_u + B_2K(s)\hat{\delta}_y + \hat{\delta}_x),$$

and

$$\hat{u} = \hat{\delta}_u + K(s)(\hat{x} + \hat{\delta}_y) = \hat{\delta}_u + K(s)(R(s)(B_2\hat{\delta}_u + B_2K(s)\hat{\delta}_y + \hat{\delta}_x) + \hat{\delta}_y).$$

Hence,

$$H(s) = \begin{bmatrix} R(s) & R(s)B_2K(s) & R(s)B_2 \\ K(s)R(s) & K(s)R(s)B_2K(s) + K(s) & K(s)R(s)B_2 + I \end{bmatrix}. \tag{3.6}$$

Since $K(s)$ is internally stabilising, all blocks of $H(s)$, in particular $R(s)$ and $M(s) = K(s)R(s)$, are contained in $\mathcal{RH}_\infty$.

To show that $R(s)$ and $M(s)$ are strictly proper, note that it follows from the state-space realisation (3.3) that $K(s)$ is proper with $K_\infty = D_k$. Then,

$$\begin{aligned} R_\infty &= \lim_{s\to\infty} (sI - A - B_2 K(s))^{-1} \\ &= \left(\lim_{s\to\infty} sI - A - B_2 K(s)\right)^{-1} \\ &= \lim_{s\to\infty} (sI - A - B_2 D_k)^{-1} \\ &= 0, \end{aligned} \tag{3.7}$$

where we use that, for sufficiently large $|s|$, $R(s)$ is invertible, and that taking the inverse is a continuous transformation. Hence, $R(s)$ is strictly proper, and so is $M(s) = K(s)R(s)$. □

Conversely, any pair of transfer matrices $\big(R(s), M(s)\big)$ that satisfies the conditions of Theorem 4 leads to an internally stabilising controller, as formalised by the following theorem.

Theorem 5. *Suppose $R(s)$ and $M(s)$ satisfy*

1. $\begin{bmatrix} sI - A & -B_2 \end{bmatrix} \begin{bmatrix} R(s) \\ M(s) \end{bmatrix} = I,$
2. $R(s), M(s) \in \mathcal{RH}_2$.

Then, there exists a state feedback controller $K(s) = M(s)R(s)^{-1}$ that achieves (3.5)*. Moreover, $K(s)$ is an internally stabilising controller.*

Proof. We first show that $R(s)$ is invertible. To that end, define $R'(s) = sR(s)$ and consider

$$\begin{aligned} \lim_{s\to\infty} R'(s) &= \lim_{s\to\infty} I + AR(s) + B_2 M(s) \\ &= I, \end{aligned}$$

which shows R'_∞ is invertible. Then, by [12, Proposition 3.3], $R'(s)$ is invertible and $R'(s)^{-1}$ is proper. Consequently, $R(s) = \frac{1}{s} R'(s)$ is invertible, meaning that $K(s) = M(s)R(s)^{-1}$ is defined. In addition, $K(s) = (sM(s))(\frac{1}{s} R(s)^{-1})$ is proper, and hence it can be realised as in (3.3).

We claim that $R(s)$ and $M(s)$ satisfy (3.5) with the above choice of $K(s)$. The feedback control law $\hat{u} = K(s)\hat{x}$ leads to

$$\begin{aligned} (sI - A)\hat{x} &= B_2 \hat{u} + \hat{\delta}_x \\ &= B_2 (M(s)R(s)^{-1}\hat{x}) + \hat{\delta}_x, \end{aligned}$$

which implies

$$(sI - A - B_2 M(s)R(s)^{-1})\hat{x} = \hat{\delta}_x.$$

Plugging in $K(s) = M(s)R(s)^{-1}$ and $\hat{x} = R(s)\hat{\delta}_x$ shows that $R(s)$ satisfies (3.5), and the corresponding result for $M(s)$ follows from the construction of $K(s)$.

To examine internal stability, we introduce the perturbations $\hat{\delta}_y$ and $\hat{\delta}_u$ as in Figure 3.1. Since (3.5) holds, the transfer function from $(\hat{\delta}_x, \hat{\delta}_y, \hat{\delta}_u)$ to $(\hat{x}, \hat{u})$ is $H(s)$ as in (3.6). For $K(s)$ to be internally stabilising, all blocks of $H(s)$ need to be contained in $\mathcal{RH}_\infty$. That is clear for the first and the third column. The remaining two blocks can be rewritten as follows.

$$\begin{aligned} R(s)B_2K(s) &= R(s)(B_2M(s)R(s)^{-1} + R(s)^{-1}) - I \\ &= R(s)(sI - A) - I, \end{aligned}$$

and

$$\begin{aligned} M(s)B_2K(s) + K(s) &= M(s)(B_2M(s)R(s)^{-1} + R(s)^{-1}) \\ &= M(s)(sI - A). \end{aligned}$$

Since $R(s), M(s) \in \mathcal{RH}_\infty$ and both are strictly proper, it follows that $R(s)(sI - A) - I \in \mathcal{RH}_\infty$ and $M(s)(sI - A) \in \mathcal{RH}_\infty$. Hence, the controller specified by $K(s)$ is internally stabilising. □

Therefore, the system response provides a way to parametrise all internally stabilising controllers of a system of the form (3.1). We call this approach system-level parametrisation.

Theorem 5 above is more general than the corresponding result in [6] in that it shows that any controller with transfer function $K(s) = M(s)R(s)^{-1}$ is internally stabilising. We present the specific controller implementation used in [6] in Chapter 4, as it behaves particularly well regarding constraints on the system response.

The system-level parametrisation is compatible with the usual notion of stabilisability in the sense that the affine subspace of possible system responses, given by the conditions of Theorems 4 and 5, is nonempty if and only if the plant is stabilisable. This is formalised and proven in the following lemma, adapted to continuous time from [6, Lemma 6].

Lemma 2. *The conditions*

1. $\begin{bmatrix} sI - A & -B_2 \end{bmatrix} \begin{bmatrix} R(s) \\ M(s) \end{bmatrix} = I,$
2. $R(s), M(s) \in \mathcal{RH}_2$

are feasible if and only if the pair (A, B_2) is stabilisable.

Proof. if. Suppose the pair (A, B_2) is stabilisable. Then, there exists a real matrix F such that $A + B_2F$ is Hurwitz. Let $K(s) = F$, which implies $R(s) = (sI - A - B_2F)^{-1}$ and $M(s) = F(sI - A - B_2F)^{-1}$. It is readily verified that 1. holds for this choice of $K(s)$.

Then, $R(s)$ and $M(s)$ are stable since $A + B_2F$ is Hurwitz. In addition, the equations (3.7) with $K(s) = D_K = F$ show that $R(s)$ and $M(s)$ are strictly proper.

only if. Assume there exist strictly proper and stable $R(s)$ and $M(s)$ such that

$$\begin{bmatrix} sI - A & -B_2 \end{bmatrix} \begin{bmatrix} R(s) \\ M(s) \end{bmatrix} = I.$$

Then, since all poles of $R(s)$ and $M(s)$ are in the open left-half plane, $[sI - A \quad - B_2]$ is right-invertible for all s with $\Re(s) \geq 0$. In other words, $[sI - A \quad - B_2]$ has full row rank for all such s. The Hautus test then implies that the pair (A, B_2) is stabilisable. □

The system-level parametrisation should not be seen as merely a different choice of parameters but rather as a new perspective on quantifying closed-loop behaviour. Accordingly, it elicits a new design approach aimed at the system response itself, instead of the transfer function $K(s)$ of the controller or $Q(s)$ of the feedback loop between outputs and inputs as in the naive or the Youla approach, respectively [6]. This new perspective on the entire design cycle is sometimes called the system-level approach.

While Theorem 4 gives a characterisation of all achievable system responses in terms of the affine space they live in, we may also formulate similar conditions in terms of the state-space realisations of $R(s)$ and $M(s)$. The following result gives an equivalent geometric characterisation in that setting.

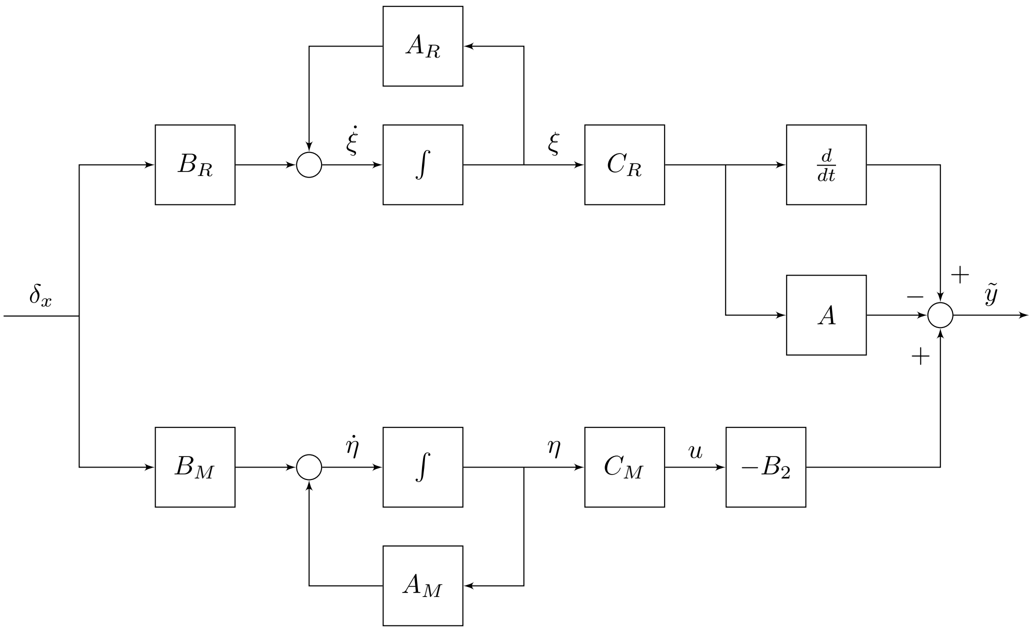

Figure 3.2: State-space realisation of $[sI - A \quad -B_2][R(s)^T \quad M(s)^T]^T$.

Theorem 6. *Suppose $R(s)$ and $M(s)$ are strictly proper and have state-space realisations (A_R, B_R, C_R) and (A_M, B_M, C_M), respectively. Define the matrices*

$$\tilde{A} = \begin{bmatrix} A_R & 0 \\ 0 & A_M \end{bmatrix}, \quad \tilde{B} = \begin{bmatrix} B_R \\ B_M \end{bmatrix}, \quad \tilde{C} = \begin{bmatrix} C_R A_R - A C_R & -B_2 C_M \end{bmatrix}.$$

Then, the following are equivalent.

1. $\begin{bmatrix} sI - A & -B_2 \end{bmatrix} \begin{bmatrix} R(s) \\ M(s) \end{bmatrix} = I,$
2. $\langle \tilde{A} \mid \operatorname{im} \tilde{B} \rangle \subset \ker \tilde{C}$ *and* $C_R B_R = I,$
3. $\operatorname{im} \tilde{B} \subset \langle \ker \tilde{C} \mid \tilde{A} \rangle$ *and* $C_R B_R = I,$
4. *there exists an $\tilde{A}$-invariant subspace $\mathcal{V}$ such that* $\operatorname{im} \tilde{B} \subset \mathcal{V} \subset \ker \tilde{C}$ *and* $C_R B_R = I.$

Proof. The equivalence of 2., 3., and 4. follows from the results that the reachable subspace $\langle \tilde{A} \mid \operatorname{im} \tilde{B} \rangle$ is the smallest $\tilde{A}$-invariant subspace containing $\operatorname{im} \tilde{B}$ [11, Chapter 3.2] and that the unobservable subspace $\langle \ker \tilde{C} \mid \tilde{A} \rangle$ is the largest $\tilde{A}$-invariant subspace contained in $\ker \tilde{C}$ [11, Chapter 3.3].

For the equivalence of 1. and 2., consider the block diagram in Figure 3.2, which shows the realisations of $R(s)$ and $M(s)$ in cascade interconnections with $(sI - A)$ and $-B_2$, respectively. The former is described by

$$\Sigma_R : \begin{cases} \dot{\xi} &= A_R \xi + B_R \delta_x \\ y_R &= (C_R A_R - A C_R) \xi + C_R B_R \delta_x, \end{cases}$$

while the latter is given as

$$\Sigma_M : \begin{cases} \dot{\eta} &= A_M \eta + B_M \delta_x \\ y_M &= -B_2 C_M \eta. \end{cases}$$

We consider the combined system shown in the diagram, from input δ_x to output $\tilde{y} = y_R + y_M$. With the matrices defined in the theorem, the combined system has state-space realisation $(\tilde{A}, \tilde{B}, \tilde{C}, C_R B_R)$ in the internal state $z = [\xi^T \quad \eta^T]^T$.

It follows that statement 1. is satisfied if and only if $\tilde{y} = \delta_x$ for $z(0) = 0$ with any input δ_x. In particular, this condition at $z = 0$ is equivalent to $C_R B_R = I$. That leads to $\tilde{C}z + \delta_x = \delta_x$ for every possible value of z with any input δ_x, which holds if and only if $\langle \tilde{A} \mid \operatorname{im} \tilde{B} \rangle \subset \ker \tilde{C}$. □

With the system-level parametrisation at hand, we may turn our attention back to the general optimal control problem (2.3). Considering the plant given by (3.2), the optimisation problem reduces to

$$\begin{aligned} \min_{R(s), M(s)} \quad & \left\| D_{11} + \begin{bmatrix} C_1 & D_{12} \end{bmatrix} \begin{bmatrix} R(s) \\ M(s) \end{bmatrix} B_1 \right\| \\ \text{subject to} \quad & \begin{bmatrix} sI - A & -B_2 \end{bmatrix} \begin{bmatrix} R(s) \\ M(s) \end{bmatrix} = I, \quad R(s), M(s) \in \mathcal{RH}_2. \end{aligned} \tag{3.8}$$

Both the objective function and the constraints are affine in this formulation, and thus we obtain a convex optimisation problem. While convexity is also achieved by the Youla parametrisation in the unconstrained case, the structure of $Q(s)$ is not related to the achieved dynamics and information exchange topology as clearly as that of the system response. In other words, with the Youla approach, design goals are not direct characteristics of the decision variable, whereas, with the system-level parametrisation, we aim to design the system response itself, which allows us to enforce a wider range of constraints. The specifics are discussed in Chapter 4.

Chapter 4

Constraints

Given that the main motivation behind the system-level parametrisation is to overcome the limitations of the Youla parametrisation when it comes to distributed optimal control, possibilities to impose constraints on the system response are a topic of utmost importance. As discussed in Section 2.4, sparsity constraints on the controller can be used to enforce distributed control with only local information exchange. With the Youla parameter, such constraints turn out to be non-convex in any system where information can eventually reach any state from any other state. In this chapter, we present how similar constraints can be imposed using the system-level parametrisation and that quadratic invariance is no longer necessary for constraints to be convex.

The reasons that allow us to surpass the limitations of the Youla approach are twofold. First, the change of perspective explained in Chapter 3, namely that we aim to design the system response itself, means that we can express structural requirements as subspace constraints directly on the decision variables of problem (3.8). Second, as we show below, once the desired system response is found, the controller $K(s)$ can be implemented in a way that it inherits the structural properties of the system response.

In addition to sparsity requirements, we can also impose performance constraints on the closed-loop system. That is, it is possible to place an upper bound on a chosen norm of the transfer function from disturbances to regulated output. The purpose may be suboptimal control or a sufficient performance requirement with respect to a norm different from the one used for optimisation. Such constraints are discussed in Section 4.4.

4.1 Locality Constraints

A common design objective in distributed control problems is that controllers should only rely on locally available measurements and information from a select subset of other controllers. This may be required by the physical characteristics of the global system and the available communication network or motivated by factors such as reduced computational and hardware costs.

As discussed in Section 2.4, if we collect local controllers into a centralised feedback controller $K(s)$, then we may directly enforce local computation of each control input $\hat{u}_i$ by imposing an appropriate sparsity pattern on $K(s)$. Namely, $K_{ij}(s) = 0$ implies that control input i is independent from the state of system j.

As noted above, we wish to solve the optimal control problem by directly designing the system response, and, accordingly, enforcing local information transfer on $R(s)$ and $M(s)$. To that end, we formalise the notion of a locality constraint.

Definition 5. *Let $\mathcal{L} \subset \mathbb{R}(s)^{m\times n}$ be a subspace defined by a chosen sparsity pattern of rational matrices of an appropriate size $m \times n$. That is, let $\mathcal{I}$ be a collection of pairs of*

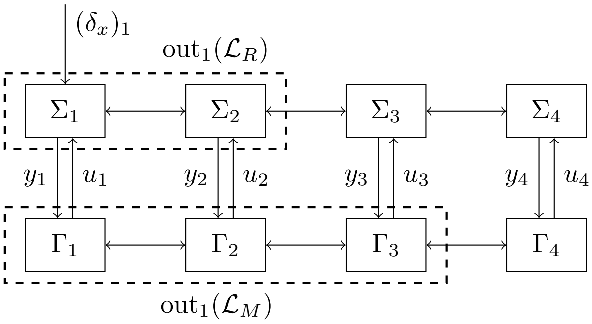

Figure 4.1: An example of a locality constraint for four interconnected systems Σ_i with controllers Γ_i. The disturbance on the state of Σ_1 is only allowed to affect the states of Σ_1 and Σ_2 and the control inputs of Σ_1, Σ_2, Σ_3.

indices and

$$\mathcal{L} = \{H(s) \in \mathbb{R}(s)^{m \times n} \mid H_{ij}(s) = 0 \quad \forall (i,j) \in \mathcal{I}\}.$$

Then, for a transfer function $H(s)$, the condition $H(s) \in \mathcal{L}$ is called a locality constraint. Moreover, the incoming and outgoing sets of system i with respect to a locality constraint $\mathcal{L}$ are defined as

$$\mathrm{in}_i(\mathcal{L}) = \{j \neq i \mid (i,j) \notin \mathcal{I}\}$$

and

$$\mathrm{out}_i(\mathcal{L}) = \{j \neq i \mid (j,i) \notin \mathcal{I}\},$$

respectively.

We further specify this concept to a system response $\big(R(s), M(s)\big)$ by constructing two sets of pairs of indices $\mathcal{I}_R$ and $\mathcal{I}_M$ and the corresponding subspaces $\mathcal{L}_R$ and $\mathcal{L}_M$, respectively. If the states or inputs are not scalars, then indices correspond to blocks of appropriate dimensions. With that, we define $\mathcal{L} = \mathcal{L}_R \times \mathcal{L}_M$ and impose a locality constraint on the system response by requiring

$$\begin{bmatrix} R(s) \\ M(s) \end{bmatrix} \in \mathcal{L}.$$

Accordingly, we construct the incoming and outgoing sets as $\mathrm{in}_i(\mathcal{L}) = \mathrm{in}_i(\mathcal{L}_R) \cup \mathrm{in}_i(\mathcal{L}_M)$ and $\mathrm{out}_i(\mathcal{L}) = \mathrm{out}_i(\mathcal{L}_R) \cup \mathrm{out}_i(\mathcal{L}_M)$.

The interpretation is that disturbances entering system j do not affect the state of system i for $(i,j) \in \mathcal{I}_R$, and control input i is independent of disturbances entering system j for each $(i,j) \in \mathcal{I}_M$. An example is shown in Figure 4.1. Since our goal is to restrict information propagation among systems and not disturbance decoupling, we assume that $(i,i) \notin \mathcal{I}_R \cup \mathcal{I}_M$ for any system index i, unless explicitly stated otherwise.

Confining the effect of each disturbance to a local region of the global system comes with an additional benefit. Namely, in some cases, it allows us to optimise the response to any given disturbance locally, with lower-dimensional transfer functions at hand. The global system response $\big(R(s), M(s)\big)$ can then be constructed by inserting the local responses into the appropriate components and exploiting the superposition principle of LTI systems [7], [13]. That is, the synthesis problem not only remains tractable to compute using convex optimisation, but appropriate locality constraints can also aid decentralised synthesis such that its complexity is independent of the scale of the global system. Section 5.3 discusses this in more detail.

Remark 1. *In general, the control input affects the dynamics through the matrix B_2. That is, a given block of the state may be influenced by several blocks of the input, depending on the structure of B_2. To take that into account, we may either impose constraints on*

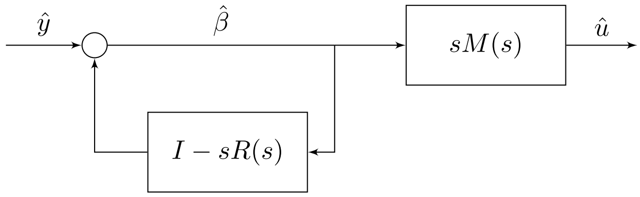

Figure 4.2: Proposed centralised implementation of the controller $K(s) = M(s)R(s)^{-1}$.

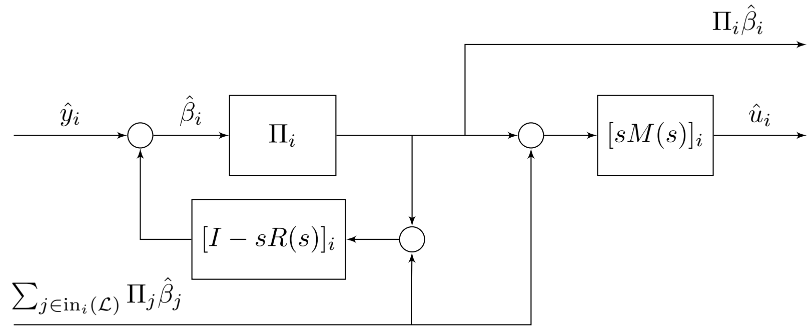

Figure 4.3: Proposed distributed controller implementation. The map Π_i embeds $\hat{\beta}_i$ in block i of a zero vector with the same dimension as $\hat{\beta}$, and $[\cdot]_i$ denotes block row i.

$B_2M(s)$ instead of $M(s)$ or assume that control inputs can be locally shared, in accordance with the structure of B_2.

In case the global system is an interconnection of systems with no feedthrough in their coupling, the matrix B_2 is block diagonal, where the diagonal is interpreted in terms of possibly non-square blocks if B_2 is non-square. Then, $M_{ij}(s) = 0$ implies $[B_2M(s)]_{ij} = 0$. Conversely, if $[B_2M(s)]_{ij} = 0$, then $M_{ij}(s)$ can be set to zero in the implementation of controller i with no effect on the dynamics. Hence, if $B_2M(s)$ satisfies some locality constraints, then $M(s)$ can be modified such that it does too, while preserving internal stability.

It is immediate that locality constraints on the system response are convex constraints on the decision variables $R(s)$ and $M(s)$ in the optimisation problem (3.8). Therefore, any locality constraint on the system response can be enforced in a convex manner, irrespective of the structure of the global system. Nevertheless, it is not guaranteed that a stabilising controller with the desired structure exists. In other words, the intersection of the chosen constraints with those of (3.8) may be an empty set. That issue is explored in Section 4.2.

Since the main motivation for imposing locality constraints on the system response is to restrict communication among local controllers, it is natural to ask whether the imposed constraints carry over from the system response to the corresponding controller. The answer is positive, at least with a specific controller implementation. This claim is formalised by the following theorem, which is a generalisation of [13, Theorem 1].

Theorem 7. *Suppose the system response $\big(R(s), M(s)\big)$ satisfies a locality constraint $\mathcal{L} = \mathcal{L}_R \times \mathcal{L}_M$. Then, the controller $K(s) = M(s)R(s)^{-1}$ can be implemented in a distributed manner such that controller i, which computes $\hat{u}_i$, only measures state i and only receives information from each controller j for $j \in \text{in}_i(\mathcal{L})$. That is, $K(s) \in \mathcal{L}_R \cup \mathcal{L}_M$.*

Proof. Consider the interconnection shown in Figure 4.2. We claim that it constitutes a centralised implementation of the controller $K(s)$. Indeed, observe that we have

$$\hat{\beta} = \hat{y} + (I - sR(s))\hat{\beta}, \tag{4.1}$$

and hence

$$\hat{\beta} = \frac{1}{s} R(s)^{-1} \hat{y},$$

where $\hat{y} = \hat{x} + \hat{\delta}_y$. Therefore, recalling from the proof of Theorem 5 that $\frac{1}{s}R(s)^{-1}$ is proper, the feedback loop within the controller is well-posed. Moreover, it is clear that the proposed implementation has transfer function $K(s) = M(s)R(s)^{-1}$.

It follows from $R(s) \in \mathcal{L}_R$ that the block $I - sR(s)$ only depends on $\hat{\beta}_i$ and $\hat{\beta}_j$ such that $j \in \text{in}_i(\mathcal{L}_R)$ to compute component i of its output. Similarly, $M(s) \in \mathcal{L}_M$ implies that the block $sM(s)$ only uses $\hat{\beta}_i$ and $\hat{\beta}_j$ such that $j \in \text{in}_i(\mathcal{L}_M)$ to compute $\hat{u}_i$.

Hence, the centralised controller shown in Figure 4.2 can be decomposed into distributed controllers as follows. Relying on the i-th block row of (4.1), each controller i computes its own state $\hat{\beta}_i$ using the output of the corresponding subsystem and $\hat{\beta}_j$ for each $j \in \text{in}_i(\mathcal{L}_R)$, which it receives from the corresponding local controllers. It computes the control input $\hat{u}_i$ based on $\hat{\beta}_i$ and $\hat{\beta}_j$ for each $j \in \text{in}_i(\mathcal{L}_M)$, also received from the other controllers in $\text{in}_i(\mathcal{L}_M)$. This construction is shown in Figure 4.3.

It follows that the distributed implementation yields the same control inputs for all subsystems as the centralised implementation and only requires each local controller to broadcast its own state $\hat{\beta}_i$ to $\text{out}_i(\mathcal{L})$ and, conversely, to read the states from $\text{in}_i(\mathcal{L})$. Thus, communication is constrained as desired. □

The internal loop in the implementation of $K(s) = M(s)R(s)^{-1}$ given by Figure 4.2 can be interpreted as an observer for disturbances integrated in time. In the distributed setting, we then have that local controllers share the estimates $\hat{\beta}_i$ of the integrated disturbances entering the corresponding systems, so that the control inputs $\hat{u}_i$ can be computed via $sM(s)$. This also provides an intuitive explanation for the locality of the proposed controller implementation: the observer part of controller i depends only on those disturbances that can affect the output of system i, and $\hat{u}_i$ is computed from estimates of those disturbances that are allowed to influence it.

Remark 2. *Replacing $I - sR(s)$ and $sM(s)$ in the proposed controller implementation with $I - R(s)$ and $M(s)$, respectively, leads to the same transfer function $K(s)$ from $\hat{y}$ to $\hat{u}$. However, in that case, the transfer function from $\hat{y}$ to $\hat{\beta}$ is $R(s)^{-1}$, which is not proper. Therefore, the internal feedback loop is ill-posed.*

A particularly intuitive class of locality constraints, presented in [13], is the following. We consider N interconnected systems that interact along the network topology given by a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where each vertex in $\mathcal{V} = \{1, 2, \ldots, N\}$ corresponds to an individual system and $(i, j) \in \mathcal{E}$ if and only if system j directly influences system i when no control is applied. In other words, $(i, j) \in \mathcal{E}$ if and only if the (i, j) block of A is nonzero. In this setting, we define the following concepts.

Definition 6. *For two subsystems i and j of a global system with network topology given by $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, the distance from j to i, denoted* $\text{dist}(i, j)$*, is the length of the shortest path in $\mathcal{G}$ from vertex j to vertex i. For subsystem i, the d-outgoing set is*

$$\text{out}_i(d) = \{j \in \mathcal{V} \mid \text{dist}(j, i) \leq d\}.$$

Similarly, the d-incoming set is

$$\text{in}_i(d) = \{j \in \mathcal{V} \mid \text{dist}(i, j) \leq d\}.$$

This allows us to define an important class of locality constraints.

Definition 7. *Given a network topology $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and a nonnegative integer d, a d-localised constraint is a locality constraint with $(j, i) \in \mathcal{I}$ whenever $j \notin \text{out}_i(d)$.*

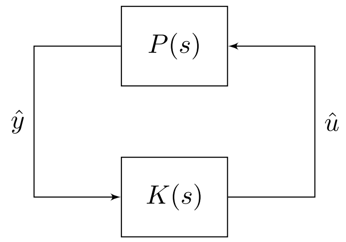

Figure 4.4: Feedback interconnection of $P(s)$ and $K(s)$.

It is important to note that d may be chosen separately for $R(s)$ and $M(s)$. If a d_R-localised constraint is imposed on $R(s)$ and a d_M-localised one on $M(s)$, then we call the system response $\big(R(s), M(s)\big)$ (d_R, d_M)-localised. In particular, $d_M > d_R$ is feasible in more scenarios than $d_M \leq d_R$, as appropriate control inputs are typically needed to decouple the $d_R + 1$st states. In that case, each system i such that $\text{dist}(i, j) = d_R + 1$ can be seen as a boundary that is controlled such that it stops a disturbance entering system j from propagating further [13].

As d-localised constraints are a special case of locality constraints, Theorem 7 can be specified for them.

Corollary 1. *Suppose the system response $\big(R(s), M(s)\big)$ is (d_R, d_M)-localised with $d_R \leq d_M$ and the controller $K(s) = M(s)R(s)^{-1}$ is implemented as in Figure 4.2. Then, the controller can be broken up into a set of distributed controllers such that controller i, which computes $\hat{u}_i$, only measures state i and only receives information from $\text{in}_i(d_M)$.*

Proof. The proof follows immediately from that of Theorem 7. □

4.2 Feasibility

In Section 4.1, we have shown that locality constraints imposed on the system response are trivially convex, regardless of the properties of the global system. Nevertheless, as with any parametrisation, feasibility of the chosen constraints is not guaranteed. In particular, we know from Theorems 4 and 5 that

$$\begin{gathered}\begin{bmatrix} sI - A & -B_2 \end{bmatrix} \begin{bmatrix} R(s) \\ M(s) \end{bmatrix} = I, \\ R(s), M(s) \in \mathcal{RH}_2 \end{gathered} \tag{4.2}$$

are necessary and sufficient conditions on the system response with an internally stabilising controller $K(s)$, which, whenever it exists, is given by $K(s) = M(s)R(s)^{-1}$. Thus, a locality constraint is feasible if and only if it can be simultaneously satisfied with (4.2). While easily testable necessary and sufficient conditions for feasibility are not known for arbitrary continuous-time systems with general constraints, we provide a number of results here that are informative in specific cases.

We first examine necessary conditions. To that end, we introduce the notion of fixed modes with respect to a given information sharing constraint. This concept is originally from [14] and is generalised in [15].

Consider the interconnection of a strictly proper plant $P(s)$ with input $\hat{u}$ and output $\hat{y}$ and a feedback controller $K(s)$, as in Figure 4.4. Assume $P(s)$ is realised as

$$P(s) = \left[\begin{array}{c|c} A_P & B_P \\ \hline C_P & 0 \end{array}\right]$$

and $K(s)$ as in (3.3). Then, the closed-loop dynamics is given by

$$\dot{x}_{cl} = A_{cl}(P, K)x_{cl},$$

where $x_{cl} = [x_P^T \quad x_K^T]^T$ is the collected state and the closed-loop system matrix $A_{cl}(P, K)$ is given by

$$A_{cl}(P, K) = \begin{bmatrix} A_P + B_P D_K C_P & B_P C_K \\ B_K C_P & A_K \end{bmatrix}.$$

Definition 8. *Consider the interconnection in Figure 4.4 with plant $P(s)$ and controller $K(s)$. Let $A_{cl}(P, K)$ denote the closed-loop system matrix, and let $\mathcal{L}$ be a locality constraint. Then, the set of fixed modes of $P(s)$ with respect to $\mathcal{L}$ is defined as*

$$\Lambda(P, \mathcal{L}) = \{\lambda \in \mathbb{C} \mid \lambda \in \sigma(A_{cl}(P, K)) \quad \forall K(s) \in \mathcal{L}\}.$$

Moreover, we define $\Lambda^s(P, \mathcal{L})$ with the additional restriction that $K(s) = D_K$ is a static feedback controller.

Fixed modes are analogous to uncontrollable closed-loop eigenvalues, taking into account that controllers need to satisfy the constraint $\mathcal{L}$. The purpose of introducing the concept is the following result from [15].

Theorem 8. *For any plant $P(s)$ and any locality constraint $\mathcal{L}$, we have*

$$\Lambda(P, \mathcal{L}) = \Lambda^s(P, \mathcal{L}).$$

Proof. See [15, Theorem 10]. □

Assuming $G(s)$ is stabilisable and taking $P(s) = G_{22}(s)$, Theorem 2 implies $K(s)$ internally stabilises $P(s)$ if and only if it internally stabilises $G(s)$. Hence, Theorem 8 leads to a necessary condition for the feasibility of a chosen locality constraint.

Corollary 2. *Assume $\mathcal{L} = \mathcal{L}_R \times \mathcal{L}_M$ defines a locality constraint on the system response $\big(R(s), M(s)\big)$ and its intersection with the affine subspace given by* (4.2) *is non-empty. Then,*

$$\Lambda^s(G_{22}, \mathcal{L}_R \cup \mathcal{L}_M) \subset \mathbb{C}_-.$$

Proof. Let the system response $\big(R(s), M(s)\big)$ satisfy the locality constraint $\mathcal{L}$ and (4.2). Theorem 7 then implies that there exists an internally stabilising controller with transfer function $K(s) = M(s)R(s)^{-1}$ such that $K(s) \in \mathcal{L}_R \cup \mathcal{L}_M$. It follows that $\Lambda(G_{22}, \mathcal{L}_R \cup \mathcal{L}_M) \subset \mathbb{C}_-$, and hence, by Theorem 8, $\Lambda^s(G_{22}, \mathcal{L}_R \cup \mathcal{L}_M) \subset \mathbb{C}_-$. □

Another approach to derive necessary conditions for feasibility is through the sparsity operator.

Definition 9. *Let $H(s) \in \mathbb{R}(s)^{Mm \times Nn}$ be a partitioned rational matrix with blocks of dimension $m \times n$. Then, the sparsity operator is defined as*

$$\mathrm{sp} : \mathbb{R}(s)^{Mm \times Nn} \to \{0, 1\}^{M \times N}$$
$$\mathrm{sp}(H(s))_{ij} = \begin{cases} 1 & \text{if} \quad H_{ij}(s) \neq 0 \\ 0 & \text{if} \quad H_{ij}(s) = 0, \end{cases}$$

where $H_{ij}(s)$ stands for a block of $H(s)$. Let $H'(s)$ be another rational matrix of the same size. Then, we say $\mathrm{sp}(H'(s)) \subset \mathrm{sp}(H(s))$ if $H'_{ij}(s) \neq 0$ implies $H_{ij}(s) \neq 0$. In addition, we define $\mathrm{sp}(H'(s)) \cup \mathrm{sp}(H(s))$ as the element-wise OR operation.

Finally, assuming compatible dimensions, the product $\mathrm{sp}(H(s))\,\mathrm{sp}(H'(s))$ is defined via the rule

$$\mathrm{sp}(H(s))\,\mathrm{sp}(H'(s))_{ij} = \begin{cases} 1 & \text{if} \quad \exists k \text{ s.t. } \mathrm{sp}(H(s))_{ik}\,\mathrm{sp}(H'(s))_{kj} = 1 \\ 0 & \text{otherwise.} \end{cases}$$

Rewriting the equation in (4.2) as

$$(sI - A)R(s) = I + B_2 M(s)$$

allows us to apply the sparsity operator on both sides. Then, the following lemma provides a way to relate the sparsity of $M(s)$ to that of $R(s)$.

Lemma 3. *Let $H(s)$ and $H'(s)$ be rational matrices of compatible dimensions. Then, we have that*

$$\operatorname{sp}(H(s)H'(s)) \subset \operatorname{sp}(H(s))\operatorname{sp}(H'(s)).$$

Proof. Assume $\operatorname{sp}(H(s)H'(s))_{ij} = 1$. Then, by definition,

$$0 \neq (H(s)H'(s))_{ij} = \sum_k H(s)_{ik} H'(s)_{kj}.$$

Hence, there exists k such that $H_{ik}(s)H'_{kj}(s) \neq 0$, which implies $\operatorname{sp}(H(s))_{ik} = 1$ and $\operatorname{sp}(H'(s))_{kj} = 1$. □

Hence,

$$\operatorname{sp}(I + B_2 M(s)) \subset \operatorname{sp}(sI - A)\operatorname{sp}(R(s)) \tag{4.3}$$

is a necessary condition for (4.2) to hold. Let $[B_2 M(s)]_{ij}$ denote an off-diagonal block of $B_2 M(s)$ and assume it is nonzero. Then, (4.3) requires that $(sI - A)_{ik} \neq 0$ and $R_{kj}(s) \neq 0$ for some k. Noting that $(sI - A)_{ii} \neq 0$ and $(sI - A)_{ik} = A_{ik}$ for $i \neq k$, that translates to either $R_{ij}(s) \neq 0$ or there exists an index k such that $R_{kj}(s) \neq 0$ and $A_{ik} \neq 0$. That is, if disturbance j affects the control action on system i, then it must also affect its state or the state of one of its in-neighbours k in the physical network topology. A special case is formalised below.

Lemma 4. *Suppose $\big(R(s), M(s)\big)$ is a (d_R, d_M)-localised system response with $d_M > d_R + 1$ that satisfies (4.2), and B_2 is a block diagonal matrix of full column rank. Then, $\big(R(s), M(s)\big)$ is $(d_R, d_R + 1)$-localised.*

Proof. Note that the full column rank of B_2 implies that the blocks on its diagonal have full column rank as well. That means $\operatorname{sp}(B_2 M(s)) = \operatorname{sp}(M(s))$. The reasoning above then completes the proof. □

As a step towards sufficient conditions for feasibility, we consider a special case. In particular, assume B_2 is block diagonal, so that $\operatorname{sp}(B_2 M(s)) \subset \operatorname{sp}(M(s))$, and has full row rank. An example of that is a full control problem, where $B_2 = I$. Then, the system response

$$\begin{aligned} R(s) &= \frac{1}{s+c} I, \\ M(s) &= -\frac{1}{s+c} B_2^{\dagger}(A + cI), \end{aligned} \tag{4.4}$$

where $B_2^{\dagger}$ is a right inverse of B_2 and $c > 0$, satisfies (4.2). Indeed, $R(s) \in \mathcal{RH}_2$ and $M(s) \in \mathcal{RH}_2$ are immediate, and we have

$$(sI - A)R(s) - B_2 M(s) = \frac{s}{s+c} I - \frac{1}{s+c} A + \frac{1}{s+c} A + \frac{c}{s+c} I = I.$$

Moreover, $B_2 M(s)$ satisfies a 1-localised constraint with respect to the physical network topology given by A. In the spirit of Remark 1, the same constraint can be enforced on the obtained $M(s)$ without changing the dynamics, simply by setting the appropriate blocks to zero. Hence, since $R(s)$ is diagonal, the system response is $(0, 1)$-localised. This immediately implies the following.

Lemma 5. *Assume B_2 is block diagonal and has full row rank, and let $\mathcal{L}$ be a locality constraint on $\big(R(s), M(s)\big)$ such that it contains every $(0, 1)$-localised system response. Then, $\mathcal{L}$ is a feasible constraint.*

Therefore, assuming B_2 satisfies the premises of Lemma 5, communication among controllers can always be restricted to the physical network topology, and such restrictions can be encoded in convex constraints on the system response. This is in stark contrast with the necessary condition of QI for convex locality constraints in the Youla approach.

While sufficient conditions for feasibility of locality constraints on the system response in the case of more general systems are unknown to us, the following result from [16] is worth stating.

Theorem 9. *Consider a plant $P(s)$ with input $\hat{u}$ and output $\hat{y}$, equipped with a to-be-designed feedback controller $K(s)$. Given a locality constraint $\mathcal{L}$ to be imposed on $K(s)$, an internally stabilising controller $K(s) \in \mathcal{L}$ exists if and only if $\Lambda^s(P(s), \mathcal{L}) \subset \mathbb{C}_-$.*

Proof. See [16, Theorem 12]. □

That is, if block $G_{22}(s)$ of the stabilisable global system $G(s)$ has no unstable fixed modes with respect to the chosen locality constraints, then an internally stabilising controller that satisfies the constraints always exists. However, the sparsity of the controller does not necessarily manifest itself in a related sparsity pattern of the system response. In other words, the converse of Theorem 7 does not hold. A trivial example is a strongly connected and open-loop stable system, where $K(s) = 0$ is a stabilising controller but the corresponding $R(s)$ is not sparse.

This shows that locality of the system response is a stricter condition than locality of the controller. Consequently, although the system-level parametrisation is less restrictive than other approaches, it still does not allow us to impose all feasible sparsity patterns of the controller of a general system as locality constraints on the system response.

4.3 QI Subspace Constraints

To illustrate the power of imposing constraints on the system response, we consider subspace constraints that are quadratically invariant with respect to the plant at hand. By the results in Section 2.4, such constraints can be enforced on the Youla parameter in a convex manner. The following theorem, adapted to state feedback from [6, Corollary 1], shows that they can also be expressed as convex constraints on the system response.

Theorem 10. *Assume the plant $G(s)$ admits a feedback controller $K(s)$ of dimension $m \times n$. Consider a linear subspace $\mathcal{C} \subset \mathbb{R}(s)^{m\times n}$ such that it is quadratically invariant with respect to $G_{22}(s) = (sI - A)^{-1}B_2$. Then, $K(s) \in \mathcal{C}$ if and only if $M(s)(sI - A) \in \mathcal{C}$, where $M(s)$ is as in* (3.5).

Proof. From Section 2.4, we know that the function $h(K(s)) = -K(s)(I - G_{22}(s)K(s))^{-1}$ preserves $\mathcal{C}$ and is involutive. Hence, $K(s) \in \mathcal{C}$ is equivalent to $h(K(s)) \in \mathcal{C}$.

To prove the theorem, we show that $M(s)(sI - A) = -h(K(s))$ and note that $\mathcal{C}$ is a linear subspace. Indeed,

$$M(s)(sI - A) = K(s)(sI - A - B_2K(s))^{-1}(sI - A),$$

and

$$\begin{aligned}
&\Big(I - (sI - A)^{-1}B_2K(s)\Big)\left((sI - A - B_2K(s)\right)^{-1}(sI - A) \\
&= (sI - A)^{-1}\Big((sI - A)(sI - A - B_2K(s))^{-1} - B_2K(s)(sI - A - B_2K(s))^{-1}\Big)(sI - A) \\
&= I,
\end{aligned}$$

which implies

$$M(s)(sI - A) = K(s)\left(I - (sI - A)^{-1}B_2K(s)\right)^{-1} = -h(K(s)). \qquad \square$$

Theorem 10 thus shows that subspace constraints that are quadratically invariant under the plant can be expressed as linear constraints on the system response. As discussed in Section 2.4, QI is a necessary condition for a subspace constraint to induce a convex constraint on the Youla parameter. Hence, any convex constraint in the Youla approach that corresponds to a subspace constraint, such as a required sparsity pattern, on the controller can be equivalently imposed as a convex constraint using the system-level parametrisation.

4.4 Performance Constraints

In addition to restricting communication among controllers, we may wish to prescribe requirements on quantities associated to the overall dynamics of the closed-loop system. The motivations for that fall into two categories. First, instead of optimal control, where the desired controller minimises the norm of the transfer function from disturbances to the regulated output, we may aim for suboptimal control, where we only require the controller to keep said norm below an appropriate performance level. Second, in optimal control, it is sometimes desirable to optimise one performance metric while keeping another, perhaps a different norm of the same quantity, below a critical value.

Irrespective of the motivation, performance constraints can be imposed in a rather straight-forward fashion in the system-level approach, as shown in [6]. Namely, let $g : \mathcal{RH}_\infty \to \mathbb{R}$ be any convex function with domain of the same dimensions as $\big(R(s), M(s)\big)$. Then, since sublevel sets of convex functions are convex,

$$g(R(s), M(s)) \leq \gamma, \tag{4.5}$$

where $\gamma \in \mathbb{R}$, is a convex constraint on the system response. Hence, it can be used as an additional constraint in (3.8). Based on (3.8),

$$g(R(s), M(s)) = \left\| D_{11} + \begin{bmatrix} C_1 & D_{12} \end{bmatrix} \begin{bmatrix} R(s) \\ M(s) \end{bmatrix} B_1 \right\| \tag{4.6}$$

with a suitable norm is a notable choice of g. For suboptimal control, we choose g like (4.6) and accept any system response $\big(R(s), M(s)\big)$ that satisfies (4.2) and (4.5).

Chapter 5

Synthesis

Building on the previous chapters, we turn our attention to controller synthesis in the system-level design approach. Consider a stabilisable system $G(s)$ as in (3.2), and let $\mathcal{C}$ denote a convex constraint on the system response. Of course, $\mathcal{C}$ may as well be the intersection of several convex constraints. Then, we aim to find the solution to

$$\begin{aligned} \min_{R(s),M(s)} \quad & \left\| D_{11} + \begin{bmatrix} C_1 & D_{12} \end{bmatrix} \begin{bmatrix} R(s) \\ M(s) \end{bmatrix} B_1 \right\| \\ \text{subject to} \quad & \begin{bmatrix} sI - A & -B_2 \end{bmatrix} \begin{bmatrix} R(s) \\ M(s) \end{bmatrix} = I, \quad R(s), M(s) \in \mathcal{RH}_2 \cap \mathcal{C}. \end{aligned} \tag{5.1}$$

Once the optimal system response $\big(R(s), M(s)\big)$ is found, Theorem 7 provides a recipe for the implementation of the corresponding controller.

While (5.1) is a convex optimisation problem, it is, in general, infinite-dimensional. Thus, standard optimisation methods cannot solve it directly.

For discrete-time systems, [5] and [6] propose to utilise deadbeat control in order to turn (5.1) into a finite-dimensional problem. Specifically, by imposing finite impulse response (FIR) constraints on the system response, we obtain the finite-dimensional representations $R(z) = \sum_{t=0}^{T} z^{-t} R_t$ and $M(z) = \sum_{t=0}^{T} z^{-t} M_t$, where T is a chosen time horizon and R_t, M_t are real matrices for each t. In terms of the resulting dynamics, such constraints mean that the effects of any disturbance are extinguished in at most T time steps. It is shown in [6] that FIR constraints are always feasible for sufficiently large T if the plant is controllable and observable, and relaxations exist for the remaining cases [17].

With appropriate FIR constraints, the discrete-time optimal control problem can be readily solved. Nevertheless, this approach is not without shortcomings, as the additional constraint may alter the attainable optimal value. That difference can be made arbitrarily small by increasing the time horizon T, however, the minimal state-space realisation of the resulting controller, developed in [18], has a state-space dimension that grows linearly with T. Hence, there is an important trade-off between performance and complexity in both synthesis and implementation.

Due to the fundamental differences between discrete- and continuous-time dynamics, the above method does not work with continuous-time systems. In particular, deadbeat control is not possible for continuous-time LTI systems. In this chapter, we examine two approaches towards solving the optimal control problem (5.1) with locality constraints, but note that a universal method for general systems and constraints remains a topic of future research. Lastly, in Section 5.3, we briefly discuss how locality constraints enable localised synthesis.

5.1 Approximate Solutions

A straight-forward approach to turning (5.1) into a finite-dimensional problem is to replace $R(s)$ and $M(s)$ with finite-order approximations. This method is suggested in, for example, [19] for an input-output parametrisation similar to the system-level parametrisation, and a more general treatment of the topic is found in [20, Chapter 15.1].

First, we construct an approximation for each component of a given transfer function. To that end, we state the following result as a special case of the discussion in [21].

Lemma 6. *Let $a > 0$ and define*

$$V_n = \left\{ \sum_{i=0}^{n-1} \frac{c_i s^i}{(s+a)^n} \;\middle|\; c_i \in \mathbb{R} \right\}.$$

Then,

$$\bigcup_{n=0}^{\infty} V_n$$

is dense in $\mathcal{RH}_2$ under both the $\mathcal{H}_2$- and the $\mathcal{H}_\infty$-norm.

Proof. See [21, Theorem 1] and the accompanying discussion. □

Motivated by Lemma 6, we can define the additional constraint

$$\mathcal{F}_N = \left\{ H(s) \in \mathcal{RH}_2^{k\times l} \;\middle|\; H_{ij}(s) \in \cup_{n=0}^{N} V_n \quad \forall (i,j) \in \{1,2,\ldots,k\} \times \{1,2,\ldots,l\} \right\},$$

where N is the chosen order of the approximation, on transfer functions of dimension $k \times l$. We do not indicate k and l explicitly for these constraints, as they are always clear from the context. It is easily verified that $\mathcal{F}_N$ is convex.

Then, problem (5.1) with the additional constraint of

$$\big(R(s), M(s)\big) \in \mathcal{F}_N$$

is a finite-dimensional convex optimisation problem. Hence, a range of standard techniques can be utilised for solving it. In addition, by Lemma 6, taking the limit as $N \to \infty$ in this setting is equivalent to solving the original, infinite-dimensional problem. It is therefore advantageous to repeatedly solve the approximate problem with increasing N until the update of the optimal value becomes negligible.

An obvious downside of this method is that there is no guarantee about the order N needed to obtain a feasible optimisation problem and a solution sufficiently close to the optimum of (5.1). These limitations are similar to those of the FIR method for the discrete-time case.

5.2 Synthesis through Vectorisation

In this section, we turn our attention to the $\mathcal{H}_2$ optimal control problem, that is, problem (5.1) where the norm is the $\mathcal{H}_2$-norm. We present a method that is intended as a stepping stone towards a general algorithm for solving the $\mathcal{H}_2$ problem with the system-level parametrisation in the presence of locality constraints. It is inspired by the idea in [3], where a constrained problem in the Youla parameter is transformed into a higher-dimensional unconstrained optimal control problem. An unconstrained $\mathcal{H}_2$ problem can then be solved via well-established state-space methods; see [8, Chapter 14] for example.

Consider the optimal control problem (5.1) with the $\mathcal{H}_2$-norm and a locality constraint $\mathcal{L}$. Let $\mathcal{S}$ denote the intersection of the relevant constraints, so that

$$\mathcal{S} = \left\{ \begin{bmatrix} R(s) \\ M(s) \end{bmatrix} \in \mathcal{RH}_2 \cap \mathcal{L} \;\middle|\; \begin{bmatrix} sI - A & -B_2 \end{bmatrix} \begin{bmatrix} R(s) \\ M(s) \end{bmatrix} = I \right\}, \tag{5.2}$$

and assume $\big(R_{\text{nom}}(s), M_{\text{nom}}(s)\big) \in \mathcal{S}$ is a system response that satisfies the constraints. See (4.4) for an example of a nominal system response. Define

$$\mathcal{K} = \ker[sI - A \quad -B_2] \cap \mathcal{RH}_2 \cap \mathcal{L}.$$

Then, any admissible system response $\big(R(s), M(s)\big) \in \mathcal{S}$ can be written as

$$\begin{bmatrix} R(s) \\ M(s) \end{bmatrix} = \begin{bmatrix} R_{\text{nom}}(s) \\ M_{\text{nom}}(s) \end{bmatrix} + \begin{bmatrix} \bar{R}(s) \\ \bar{M}(s) \end{bmatrix},$$

for some

$$\begin{bmatrix} \bar{R}(s) \\ \bar{M}(s) \end{bmatrix} \in \mathcal{K}.$$

Pick $\Xi_1(s), \Xi_2(s), \dots, \Xi_p(s) \in \mathcal{K}$ for some $p \in \mathbb{Z}_{>0}$ and construct $\tilde{\mathcal{K}}$ as

$$\tilde{\mathcal{K}} = \left\{ \sum_{i=0}^{p} \eta_i(s)\Xi_i(s) \;\middle|\; \eta_i(s) \in \mathcal{RH}_\infty^{1\times 1} \right\}. \tag{5.3}$$

That is, $\tilde{\mathcal{K}}$ is the span of $\Xi_1(s), \Xi_2(s), \dots, \Xi_p(s)$ over $\mathcal{RH}_\infty$, which can be considered as a ring but not as a field, meaning that $\tilde{\mathcal{K}}$ is a module over $\mathcal{RH}_\infty$. It is immediate that $\tilde{\mathcal{K}} \subset \mathcal{K}$. Moreover, $\tilde{\mathcal{K}}$ can be seen as a real vector space and is therefore convex. In what follows, we solve the optimal control problem restricted to the affine subspace

$$\tilde{\mathcal{S}} = \begin{bmatrix} R_{\text{nom}}(s) \\ M_{\text{nom}}(s) \end{bmatrix} + \tilde{\mathcal{K}} \subset \mathcal{RH}_\infty. \tag{5.4}$$

To that end, we define the vectorisation operator.

Definition 10. *Let $\mathcal{R}$ be any ring. The vectorisation operator is defined as the function*

$$\text{vec} : \mathcal{R}^{m\times n} \to \mathcal{R}^{mn}$$
$$\text{vec}(X)_{(i-1)\cdot m+j} = X_{ji}$$

for $i = 1, \dots, n$ and $j = 1, 2, \dots, m$.

That is, $\text{vec}(\cdot)$ returns a vector of the columns of a matrix stacked on top of each other. Clearly, it is a bijection provided that the dimensions are fixed. With that, we construct

$$T(s) = \begin{bmatrix} \text{vec}(\Xi_1(s)) & \text{vec}(\Xi_2(s)) & \cdots & \text{vec}(\Xi_p(s)) \end{bmatrix},$$

and introduce the notation

$$W = B_1^T \otimes \begin{bmatrix} C_1 & D_{12} \end{bmatrix}.$$

Then, the following result holds.

Theorem 11. *Suppose $\eta(s) \in \mathcal{RH}_\infty^p$ is the solution of*

$$\begin{aligned} \min_{\eta} \quad & \left\| \text{vec}(D_{11}) + W \,\text{vec}\left(\begin{bmatrix} R_{\text{nom}}(s) \\ M_{\text{nom}}(s) \end{bmatrix} \right) + WT(s)\eta(s) \right\|_{\mathcal{H}_2} \\ \textit{subject to} \quad & \eta(s) \in \mathcal{RH}_\infty. \end{aligned} \tag{5.5}$$

Then,

$$\begin{bmatrix} R(s) \\ M(s) \end{bmatrix} = \begin{bmatrix} R_{\text{nom}}(s) \\ M_{\text{nom}}(s) \end{bmatrix} + \operatorname{vec}^{-1}(T(s)\eta(s))$$

solves

$$\begin{aligned} \min_{R(s),M(s)} \quad & \left\| D_{11} + \begin{bmatrix} C_1 & D_{12} \end{bmatrix} \begin{bmatrix} R(s) \\ M(s) \end{bmatrix} B_1 \right\|_{\mathcal{H}_2} \\ \textit{subject to} \quad & \begin{bmatrix} R(s) \\ M(s) \end{bmatrix} \in \tilde{\mathcal{S}}. \end{aligned} \tag{5.6}$$

Proof. It can be verified by direct computation that for any compatible matrices P, Q, R,

$$\operatorname{vec}(PQR) = (R^T \otimes P) \operatorname{vec}(Q).$$

Thus,

$$\begin{aligned} &(B_1^T \otimes \begin{bmatrix} C_1 & D_{12} \end{bmatrix}) \operatorname{vec}\left(\begin{bmatrix} R_{\text{nom}}(s) \\ M_{\text{nom}}(s) \end{bmatrix} \right) + (B_1^T \otimes \begin{bmatrix} C_1 & D_{12} \end{bmatrix}) T(s)\eta(s) \\ &= \operatorname{vec}\left(\begin{bmatrix} C_1 & D_{12} \end{bmatrix} \left(\begin{bmatrix} R_{\text{nom}}(s) \\ M_{\text{nom}}(s) \end{bmatrix} + \operatorname{vec}^{-1}(T(s)\eta(s)) \right) B_1 \right). \end{aligned}$$

In addition, it follows from (5.3) and (5.4) that

$$\left\{ \begin{bmatrix} R_{\text{nom}}(s) \\ M_{\text{nom}}(s) \end{bmatrix} + \operatorname{vec}^{-1}(T(s)\eta(s)) \;\middle|\; \eta(s) \in \mathcal{RH}_\infty \right\} = \tilde{\mathcal{S}}.$$

Hence, (5.5) is precisely the vectorised formulation of (5.6). From Section 2.1.2, we observe that the $\mathcal{H}_2$-norm is invariant under vectorisation. Then, the statement of the theorem follows. □

Notice that (5.5) is an unconstrained optimal control problem of the form (2.6). That is, $\eta(s)$ can be seen as the Youla parameter for the system

$$\tilde{G}(s) = \begin{bmatrix} \operatorname{vec}(D_{11}) + W \operatorname{vec}\left(\begin{bmatrix} R_{\text{nom}}(s) \\ M_{\text{nom}}(s) \end{bmatrix} \right) & WT(s) \\ I & 0 \end{bmatrix},$$

using that a doubly coprime factorisation of $\tilde{G}_{22}(s) = 0$ is given by $0 = I0 = 0I$ with Bézout coefficients $X = \tilde{X} = I$ and $Y = \tilde{Y} = 0$. Hence, (5.5) can be solved by readily available state-space methods for unconstrained $\mathcal{H}_2$ optimal controller synthesis, such as those in [8, Chapter 14].

However, this method still has a clear limitation in that it introduces additional constraints to the optimal control problem. More specifically, at the present time, we do not know if it is possible to construct $\tilde{\mathcal{K}}$ such that it is equal to $\mathcal{K}$, which would imply $\tilde{\mathcal{S}} = \mathcal{S}$. A further examination of that is an interesting direction for future research. That said, if the optimal control problem with the original constraint set $\mathcal{S}$ is feasible, it is always possible to construct $\tilde{\mathcal{K}} \neq \emptyset$ so that (5.6) and, equivalently, (5.5) are feasible problems.

5.3 Separable Synthesis

As explained in Section 4.1, locality constraints on the system response confine the effects of disturbances to local regions in the global system. In other words, the system response to a disturbance entering a given system is allowed to be nonzero in specific components only. Formally, if $\big(R(s), M(s)\big) \in \mathcal{L} = \mathcal{L}_R \times \mathcal{L}_M$, then block column j of $R(s)$ can only be nonzero at each block row i such that $i \in \text{out}_j(\mathcal{L}_R)$, and the analogous statement holds for $M(s)$ with $\mathcal{L}_M$. This suggests that block columns of the system response can be optimised locally, only considering the block rows that are permitted to be nonzero.

Optimising the block columns of the system response independently is motivated by the superposition principle of LTI systems. That is, the system response to disturbances entering all systems is the sum of the responses to disturbances entering each one of the systems in the interconnection. Nevertheless, whether or not we can optimise the block columns independently depends on the chosen performance criterion as well. A comprehensive treatment of this topic is found in [7], which is the basis of our brief discussion below.

First, we define column-wise separable objective functions and constraints in the context of the optimisation problem

$$\begin{aligned} \min_{R(s), M(s)} \quad & g\big(R(s), M(s)\big) \\ \text{subject to} \quad & \big(R(s), M(s)\big) \in \mathcal{S}, \end{aligned} \tag{5.7}$$

where $g : \mathcal{RH}_2 \to \mathbb{R}$ is some convex function and $\mathcal{S} \subset \mathcal{RH}_2$ is some convex constraint set.

Definition 11. *Consider the optimisation problem* (5.7)*, and assume the system response is partitioned into N block columns. Then, the objective function g is called column-wise separable if there exist functions $g_1, g_2, \ldots, g_N$ such that*

$$g\big(R(s), M(s)\big) = \sum_{i=1}^{N} g_i\big(R_i(s), M_i(s)\big),$$

where $R_i(s)$ and $M_i(s)$ denote block column i of $R(s)$ and $M(s)$, respectively.

Similarly, the constraint set $\mathcal{S}$ is called column-wise separable if there exist sets $\mathcal{S}_1, \mathcal{S}_2, \ldots, \mathcal{S}_N$ such that

$$\big(R(s), M(s)\big) \in \mathcal{S}$$

if and only if

$$\big(R_i(s), M_i(s)\big) \in \mathcal{S}_i$$

for $i = 1, 2, \ldots, N$.

If the objective function and the constraints are both column-wise separable, then the optimal control problem can be solved column-wise. In that case, by our discussion above, appropriate locality constraints lead to lower-dimensional synthesis problems for each block column. Hence, the optimal system response to the disturbances that enter system j in the interconnection can be computed locally, and the complexity of such computation depends only on the number of systems in $\text{out}_j(\mathcal{L}_\mathcal{R} \cup \mathcal{L}_M)$, as opposed to the scale of the entire interconnection. In large-scale networks of systems, this property is extremely advantageous.

Lastly, we present an example of a column-wise separable constraint set and objective function. Specifically, assume $\mathcal{S}$ is as in (5.2). Then, $\mathcal{S}$ is column-wise separable. Indeed, a locality constraint can be trivially separated into constraints for each block column, $\mathcal{RH}_2$ is interpreted component-wise by definition, and the constraint

$$\begin{bmatrix} sI - A & -B_2 \end{bmatrix} \begin{bmatrix} R(s) \\ M(s) \end{bmatrix} = I$$

is column-wise separable by the rules of matrix multiplication.

For the objective function, consider

$$g\big(R(s), M(s)\big) = \left\| D_{11} + \begin{bmatrix} C_1 & D_{12} \end{bmatrix} \begin{bmatrix} R(s) \\ M(s) \end{bmatrix} B_1 \right\|_{\mathcal{H}_2}^2,$$

where B_1 is block diagonal so that the disturbances on different systems are decoupled. Note that optimising the square of the norm is equivalent to optimising the norm itself, as it is a nonnegative quantity. Then, g is a column-wise separable objective function. This can be seen from the following. By the rules of matrix multiplication, the block columns of the transfer function inside the norm can be computed independently, using that B_1 is block diagonal. Our claim then follows from the fact that the square of the $\mathcal{H}_2$-norm of a transfer function matrix is the sum of the squared $\mathcal{H}_2$-norms of its components.

Hence, assuming disturbances are decoupled, synthesis for the $\mathcal{H}_2$ problem under locality constraints can be performed locally, independent of the scale of the global system. For a variety of other separable constraints and objective functions, see [7].

Chapter 6

Conclusion

In this thesis, we have introduced the distributed optimal control problem subject to constraints on communication among controllers. Aiming to solve that problem, we have presented the Youla parametrisation and its limitations, namely the necessity of quadratic invariance, in encoding information sharing constraints as convex constraints. Motivated by these limitations, we have developed a system-level design approach for continuous-time LTI systems, analogous to the approach presented in [5]–[7] for the discrete-time case.

In particular, we have considered the state feedback case and defined the system response as the pair of transfer functions that map disturbances affecting the state to the state and the control input. We have shown that the system response can be used to parametrise all internally stabilising controllers of a given system and that such parametrisation leads to a convex formulation of the distributed optimal control problem. In that setup, communication sparsity can be enforced via convex constraints on the system response such that they carry over to the controller. In particular, any subspace constraint on the controller that can be expressed as a convex constraint on the Youla parameter can also be imposed on the system response in a convex fashion, and for a full control problem, it is always possible to restrict communication to the physical network topology.

Finally, we have explained the difficulties in solving the resulting synthesis problem, which are due to the optimisation problem being infinite-dimensional. As partial solutions, we have presented a finite-dimensional approximation method, and proposed a construction that translates the constrained $\mathcal{H}_2$ problem into an unconstrained $\mathcal{H}_2$ problem of a higher-dimensional system that can be solved with available state-space methods. While it is unclear at the moment whether the latter can be carried out without introducing additional constraints, feasibility is guaranteed provided that the original problem is feasible.

Based on the work presented in this thesis, several directions of future research can be identified. Namely, a full solution of the synthesis problem, possibilities to incorporate modular design techniques, and constraints that encode communication delays would all be interesting and useful areas of further work within the theme of system-level design.

Acknowledgements

I would like to express my sincere gratitude towards my supervisor, Bart Besselink, for his availability, for the immensely useful support and feedback I have received over the course of this project, and for our enjoyable weekly discussions.

I would also like to thank Kanat Camlibel for taking the role of my second assessor, and, last but not least, Radu Cimpean, Marieke Heidema, Armin Pirastehzad, and Melvyn Wildeboer, for being great teaching assistant colleagues in the course Linear Systems and for taking over some of my grading duties in the final days of writing this thesis.

Bibliography

[1] D. C. Youla, J. J. Bongiorno Jr, and H. A. Jabr, "Modern Wiener-Hopf design of optimal controllers - Part I: The single-input-output case," *IEEE Transactions on Automatic Control*, vol. 21, no. 1, pp. 3–13, Feb. 1976.

[2] D. C. Youla, H. A. Jabr, and J. J. Bongiorno Jr, "Modern Wiener-Hopf design of optimal controllers - part II: The multivariable case," *IEEE Transactions on Automatic Control*, vol. 21, no. 3, pp. 319–338, Jun. 1976.

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

[4] L. Lessard and S. Lall, "Convexity of decentralized controller synthesis," *IEEE Transactions on Automatic Control*, vol. 61, no. 10, pp. 3122–3127, Oct. 2016.

[5] Y.-S. Wang, N. Matni, and J. C. Doyle, "Localized LQR optimal control," in *Proceedings of the 53rd IEEE Conference on Decision and Control*, Los Angeles, CA, USA: IEEE, Dec. 2014, pp. 1661–1668.

[6] Y.-S. Wang, N. Matni, and J. C. Doyle, "A system-level approach to controller synthesis," *IEEE Transactions on Automatic Control*, vol. 64, no. 10, pp. 4079–4093, Oct. 2019.

[7] Y.-S. Wang, N. Matni, and J. C. Doyle, "Separable and localized system-level synthesis for large-scale systems," *IEEE Transactions on Automatic Control*, vol. 63, no. 12, pp. 4234–4249, Dec. 2018.

[8] K. Zhou, J. C. Doyle, and K. Glover, *Robust and Optimal Control*. Prentice Hall, 1996.

[9] L. Mirkin, *Linear Control Systems*. Lecture Notes, Technion IIT, 2022.

[10] B. A. Francis, *A course in $\mathcal{H}_\infty$ control theory* (Lecture notes in control and information sciences). Berlin: Springer, 1987.

[11] H. L. Trentelman, A. A. Stoorvogel, and M. Hautus, *Control theory for linear systems*. Springer, 2001.

[12] B. M. Shali, "Contract theory for linear control systems," Ph.D. dissertation, University of Groningen, 2023.

[13] Y.-S. Wang, N. Matni, and J. C. Doyle, "Localized LQR control with actuator regularization," in *Proceedings of the 2016 American Control Conference (ACC)*, Boston, MA, USA: IEEE, Jul. 2016, pp. 5205–5212.

[14] S.-H. Wang and E. Davison, "On the stabilization of decentralized control systems," *IEEE Transactions on Automatic Control*, vol. 18, no. 5, pp. 473–478, Oct. 1973.

[15] A. Alavian and M. Rotkowitz, "Fixed modes of decentralized systems with arbitrary information structure," in *Proceedings of the 21st International Symposium on Mathematical Theory of Networks and Systems*, Groningen, The Netherlands, Jul. 2014.

[16] A. Alavian and M. Rotkowitz, "Stabilizing decentralized systems with arbitrary information structure," in *Proceedings of the 53rd IEEE Conference on Decision and Control*, Los Angeles, CA, USA: IEEE, Dec. 2014, pp. 4032–4038.

[17] N. Matni, Y.-S. Wang, and J. Anderson, "Scalable system level synthesis for virtually localizable systems," in *Proceedings of the 2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, Melbourne, Australia: IEEE, Dec. 2017, pp. 3473–3480.

[18] J. Anderson and N. Matni, "Structured state space realizations for SLS distributed controllers," in *Proceedings of the 2017 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, Monticello, IL, USA: IEEE, Oct. 2017, pp. 982–987.

[19] L. Furieri, Y. Zheng, A. Papachristodoulou, and M. Kamgarpour, "An input–output parametrization of stabilizing controllers: Amidst Youla and system level synthesis," *IEEE Control Systems Letters*, vol. 3, no. 4, pp. 1014–1019, Oct. 2019.

[20] S. Boyd and C. Barratt, *Linear Controller Design: Limits of Performance*. Prentice Hall, 1991.

[21] A. Linnemann, "Convergent Ritz approximations of the set of stabilizing controllers," *Systems & Control Letters*, vol. 36, no. 2, pp. 151–156, Feb. 1999.