The Control Handbook, 2nd Edition

by

Bill Levine, editor

74

Decentralized Control and Algebraic Approaches

Michael C. Rotkowitz
The University of Melbourne

74.1 Introduction

This chapter addresses the problem of decentralized control, where multiple controllers have access to different information, but need to achieve or optimize a global objective. Most of conventional controls analysis breaks down when information is decentralized, even in the simplest possible scenarios [Wit68].

This chapter addresses decentralized control problems in a simple unified framework. This framework is introduced in Section 74.2, where we see that a standard controls framework may be utilized, but with the addition of a particular constraint on the controller that needs to be designed. Section 74.3 then briefly reviews the parametrization of stabilizing controllers for centralized control, when one does not have this decentralization constraint.

Section 74.4 introduces quadratic invariance, an algebraic condition under which decentralized control problems may be cast as convex optimization problems. Section 74.5 looks at particular classes of problems to see when this condition holds, and to get some intuition behind when decentralized control problems may be tractable. Section 74.6 then discusses the perfectly decentralized control problem, a problem which is often of interest yet which typically does not satisfy this condition. Finally, while the rest of this chapter focuses on the case where both the system to be controlled and the possible controllers are all linear, Section 74.7 discusses some related results for nonlinear control.

74.2 Framework and Setup

We introduce a unified framework for studying optimal feedback control problems subject to decentralized information constraints.

74.2.1 Standard framework

We first review a standard framework for centralized control synthesis.

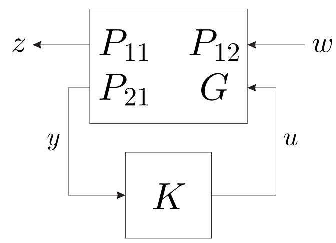

FIGURE 74.1 Standard feedback control framework

Figure 74.1 represents the standard design framework of modern control theory, and is used in many other chapters. The signal w represents the vector of exogenous inputs, those the designer has no control over, such as wind gusts if one is considering an example in aerospace, and z represents everything the designer would like to keep small, which would typically include deviations from a desired state or trajectory, or a measure of control effort, for example. The signal y represents the vector of measurements that the controller K has access to, and u is the vector of inputs from the controller that feed back into the plant. The plant is subdivided into four blocks which map w and u into z and y. The block which maps the controller input u to the measurements y is simply referred to as G, since it corresponds to the plant of classical control analysis, and so that we can later refer to its subdivisions without any ambiguity.

The design objective is to construct a controller K to keep a measure of the size of the mapping from w to z, known as the *closed-loop map*, as small as possible. There are many ways one can measure the size of a mapping, and thus this basic setup underpins much of modern controls including $\mathcal{H}_2$-control and $\mathcal{H}_\infty$-control. In this framework, a decentralized information structure may be viewed as a constraint on the structure of the controller K, as now illustrated by examples.

74.2.2 Information constraint

We now illustrate why, in this framework, decentralization may be simply encapsulated as a constraint that the controller lie in a particular subspace.

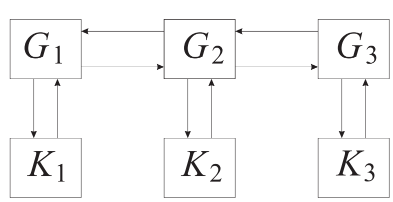

FIGURE 74.2 Perfectly decentralized control

The diagram in Figure 74.2 represents three different subsystems, each of which may effect its neighbors, and each of which has its own controller, which only has access to measurements coming from its own subsystem. In this case, if we look at the system as a whole, we need to design a controller K that can be written as

$$\begin{bmatrix} u_1 \\ u_2 \\ u_3 \end{bmatrix} = \underbrace{\begin{bmatrix} K_1 & 0 & 0 \\ 0 & K_2 & 0 \\ 0 & 0 & K_3 \end{bmatrix}}_{K} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix}$$

since each controller input may only depend upon the measurement from its corresponding subsystem. In other words, we need to design the best possible K which is block diagonal. The overall problem can be viewed as minimizing the size of the closed-loop map subject to the additional constraint that $K \in S$, where S is the set of all block diagonal controllers. This concept readily extends to any type of structural constraint we may need to impose in formulating an optimal control problem for controller synthesis. For instance, if in the above example, each controller were able to share information with its neighbors, then we would end up with a constraint set S which is tri-diagonal. In general, the ijth component of the controller is held to 0 if the ith controller is unable to see the jth measurement y_j.

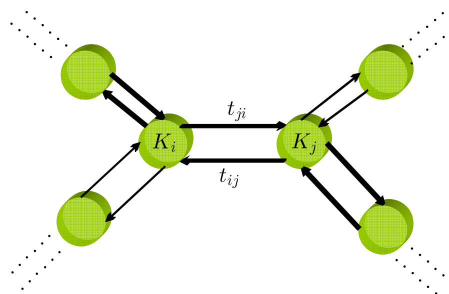

FIGURE 74.3 Network with delays

If controllers were instead allowed to communicate with each other, but with some delays, this too could be reflected in another constraint set S. This situation is represented in Figure 74.3, where the controller for a given subsystem i can see the information from another subsystem j after a transmission delay of t_{ij}. In this case, if we look at the system as a whole, we need to design a controller K that can be written as

$$\begin{bmatrix} u_1 \\ u_2 \\ u_3 \end{bmatrix} = \underbrace{\begin{bmatrix} D_{t_{11}}\tilde{K}_{11} & D_{t_{12}}\tilde{K}_{12} & D_{t_{13}}\tilde{K}_{13} \\ D_{t_{21}}\tilde{K}_{21} & D_{t_{22}}\tilde{K}_{22} & D_{t_{23}}\tilde{K}_{23} \\ D_{t_{31}}\tilde{K}_{31} & D_{t_{32}}\tilde{K}_{32} & D_{t_{33}}\tilde{K}_{33} \end{bmatrix}}_{K} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \end{bmatrix}$$

where $D_{t_{ij}}$ represents a delays of t_{ij}, and $\tilde{K}_{ij}$ represents the parts of the controller which we are free to design, since each subsystem controller must wait the proscribed amount of time before it can use information from each of the other controllers.

The set S above is called the *information constraint*, as it captures the information available to various parts of the controller. The overarching point is that the objective of decentralized control may be considered to be the minimization of a closed-loop map subject to an information constraint $K \in S$. The approach is extremely broad, as it seamlessly incorporates any type of decentralization, any control objective, and heterogeneous subsystems. It has thus come to be accepted as the canonical problem one would like to solve in decentralized control.

74.2.3 Problem formulation

The mapping from w to z that we wish to keep small in Figure 74.1, the closed-loop map, can be written as $f(P,K) = P_{11} + P_{12}K(I - GK)^{-1}P_{21}$. The problem that we would like to address may then be formulated as:

$$\begin{aligned} \text{minimize} \quad & \|f(P,K)\| \\ \text{subject to} \quad & K \text{ stabilizes } P \\ & K \in S \end{aligned} \tag{74.1}$$

The norm ($\|\cdot\|$) is any appropriate system norm, chosen based on the particular performance objectives, which could be the $\mathcal{H}_2$-norm or $\mathcal{H}_\infty$-norm as described in detail in other chapters. The information constraint S is the subspace of admissible controllers which encapsulates the decentralized nature of the system, as exemplified above. The stabilization constraint is needed in the most typical case where the signals lie in extended spaces and the plant and controller are rational proper systems whose interconnections may thus be unstable. It may not be necessary, or another technical condition may be necessary such as the invertibility of $(I - GK)$, for other spaces of interest, such as Banach spaces with bounded linear operators [RL02, RL06a].

74.3 Stabilizing Controller Parametrization

If the plant to be controlled is stable, we could use the following change of variables

$$Q = -K(I - GK)^{-1} \quad \Longleftrightarrow \quad K = -Q(I - GQ)^{-1} \tag{74.2}$$

and then allowing the new parameter Q to be stable is equivalent to the controller K stabilizing the plant P, and the set of all achievable closed-loop maps (ignoring the information constraint) is then given as

$$\{P_{11} - P_{12}QP_{21} \mid Q \text{ stable}\}. \tag{74.3}$$

This is generalized by the Youla-Kucera or YJBK parametrization [YJJ76], which gives a similar change of variables for unstable plants such that allowing the new (Youla) parameter Q to vary over all stable systems is still equivalent to considering all stabilizing controllers K, and the set of all achievable closed-loop maps is then given by

$$\{T_1 - T_2QT_3 \mid Q \text{ stable}\} \tag{74.4}$$

where T_1, T_2, T_3 are other stable systems.

We see that these parametrizations allow the set of achievable closed-loop maps to be expressed as an affine function of a stable parameter, and thus allow our objective function in our main problem (74.1) to be cast as a convex function of that parameter. However, the information constraint $K \in S$ will typically not be simple to express in the new parameter, and this will ruin the convexity of the optimization problem.

74.4 Quadratic Invariance

We have seen that we can employ a change of variables that will make our objective convex, but that will generally cause the information constraint to no longer be affine. We thus seek to characterize problems for which the information constraint may be written as an affine constraint in the Youla parameter, such that a convex reformulation of our main problem will result.

The following property, first introduced in [RL02], provides that characterization.

Definition 1 *The set S is* ***quadratically invariant*** *with respect to G if*

$$KGK \in S \qquad \text{for all} \;\; K \in S$$

In other words, given any admissible controller K, the composition KGK has to be admissible as well. When this condition holds, it follows that a controller being admissible is further equivalent to the linear-fractional transformation we encountered earlier lying in the constraint set [RL06a, RL06b]:

$$K \in S \quad \Longleftrightarrow \quad K(I - GK)^{-1} \in S \tag{74.5}$$

We can see immediately from (74.2) that for the stable case this results in the equivalence of enforcing the information constraint on the controller or on the new parameter:

$$K \in S \quad \Longleftrightarrow \quad Q \in S \tag{74.6}$$

and it can be shown that when the plant is unstable, another change of variables can be made such that this equivalence still holds [RL06b].

Thus when the information constraint S is quadratically invariant with respect to the plant G, the optimal decentralized control problem (74.1) may be recast as the following:

$$\begin{aligned} \text{minimize} \quad & \|T_1 - T_2 Q T_3\| \\ \text{subject to} \quad & Q \text{ stable} \\ & Q \in S \end{aligned} \tag{74.7}$$

which is a convex optimization problem.

74.5 Examples

This section looks at particular classes of information constraints to see when this quadratic invariance condition holds, to identify those decentralized problems which are amenable to convex synthesis. We see that this algebraic condition often has intuitive interpretations for specific classes of problems.

74.5.1 Structural Constraints

We first look at structural constraints, or sparsity constraints, where each sub-controller can see the measurements from some subsystems but not from others. This structure can be represented with a binary matrix K^{bin}. For instance, $K^{\text{bin}}_{kl} = 1$ if the kth control input u_k is allowed to be a function of the lth measurement y_l, and $K^{\text{bin}}_{kl} = 0$ if it cannot see that measurement. The information constraint S is then the set of all controllers which have the structure proscribed by K^{bin}; that is, all of the controllers such that none of the sub-controllers use information which they cannot see.

A binary matrix G^{bin} can similarly be used to give the structure of the plant. For instance, $G^{\text{bin}}_{ij} = 1$ if G_{ij} is non-zero and the ith measurement y_i is affected by the jth control input u_j, and $G^{\text{bin}}_{ij} = 0$ if it is unaffected by that input. Given this representation of the structure of the plant and the controller constraints, we have the following result:

S is quadratically invariant with respect to G if and only if

$$K^{\text{bin}}_{ki}\, G^{\text{bin}}_{ij}\, K^{\text{bin}}_{jl}\, (1 - K^{\text{bin}}_{kl}) = 0 \qquad \text{for all } i,j,k,l. \tag{74.8}$$

Figure 74.4 illustrates this condition. The condition in (74.8) requires that, for arbitrary i, j, k, l, if the three blocks on the bottom are all non-zero (or allowed to be chosen non-zero),

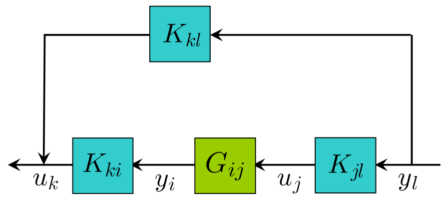

FIGURE 74.4 Structural quadratic invariance

then the top block must be allowed to be non-zero as well. In other words, if there is an indirect connection from a measurement to a control input, then there has to be a direct connection as well.

When this condition is met, the problem is quadratically invariant, and we can recast our optimal decentralized control problem as the convex optimization problem in (74.7).

74.5.2 Symmetry

We briefly consider the problem of symmetric synthesis. Suppose that we need to design the best symmetric controller; that is, the best controller such that $K_{kl} = K_{lk}$ for all k, l, and that the information constraint S is the set of all such symmetric controllers. If the plant is also symmetric; that is, if $G_{ij} = G_{ji}$ for all i, j, then KGK is symmetric for any symmetric K. Thus, $KGK \in S$ for all $K \in S$, the problem is quadratically invariant, and the optimal symmetric control problem may be recast as (74.7).

74.5.3 Delays

We now return to the problem of Figure 74.3, where we have multiple nodes/subsystems, each with its own controller, and each subsystem i can see the information from another subsystem j after a transmission delay of t_{ij}.

We similarly consider that the inputs to a given subsystem j may affect other subsystems after some delay, and denote the amount of time after which it may affect another subsystem i by the propagation delay p_{ij}.

The overall problem of controlling such a network with propagation delays, with controllers that may communicate with transmission delays, is depicted in Figure 74.5.

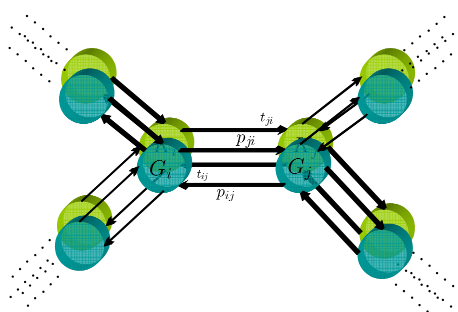

FIGURE 74.5 Network with delays

When this problem is tested for quadratic invariance, one first finds that the following condition is necessary and sufficient:

$$t_{ki} + p_{ij} + t_{jl} \geq t_{kl} \qquad \text{for all } i, j, k, l \tag{74.9}$$

This is reminiscent of condition (74.8) for structural constraints, as it similarly requires that any direct path from y_l to u_k must be at least as fast as any indirect path through the plant. This condition can be further reduced to a very simple intuitive condition [RCL05], as long as we may assume that the transmission delays themselves satisfy the triangle inequality; that is

$$t_{ki} + t_{ij} \ \geq \ t_{kj} \qquad \text{for all } k, i, j. \tag{74.10}$$

This is typically a very reasonable assumption, as it states that information is transmitted between nodes in the quickest manner available through the network. If the inequality failed for some k, j, one would want to reroute the transmissions from j to k along the faster route such that the inequality would then hold.

If the triangle inequality among transmissions does hold, then condition (74.9), and thus quadratic invariance, is reduced to simply:

$$p_{ij} \ \geq \ t_{ij} \qquad \text{for all } i, j. \tag{74.11}$$

In other words, for any pair of nodes, information needs to be transmitted faster than the dynamics propagate. When this simple condition holds, the problem is quadratically invariant, and the optimal decentralized control problem may be recast as the convex problem (74.7).

This very intuitive result has a counterintuitive complement when one considers computational delays as well. Suppose now that the ith controller cannot use a measurement from the jth subsystem until a 'pure' transmission delay of $\tilde{t}_{ij}$, representing the time it takes to send the information from one subsystem to the other, as well as a computational delay of c_i, representing the time it takes to process the information once it is received.

While intuition might suggest that these two quantities would end up being added and then replacing the right-hand side of equation (74.11), if we now assume that the pure transmission delays satisfy the triangle inequality, the condition for quadratic invariance becomes:

$$p_{ij} + c_j \ \geq \ \tilde{t}_{ij} \qquad \text{for all } i, j \tag{74.12}$$

with the computational delay on the other side of the inequality.

This shows that, regardless of computational delay, if information can be transmitted faster than dynamics propagate, then the optimal decentralized control problem can be reformulated as a convex optimization problem. If we consider a problem with multiple aerial vehicles, for example, where dynamics between any pair of subsystems will propagate at the speed of sound, this tells us that transmissions just have to be faster than that threshold for the optimal control problem to be recast as (74.7).

These results have also been extended to spatio-temporal systems [RCL10], including the special case of spatially invariant systems.

74.6 Perfectly Decentralized Control

We now revisit the problem of Figure 74.2, where each controller can only use the measurements from its own subsystem, and thus the information constraint is block diagonal. This problem is never quadratically invariant, and will never satisfy condition (74.8), except for the case where the subsystems do not affect one another; that is, except for the case where G is block diagonal as well.

In all other cases where subsystems may have some affect on others, we thus cannot parametrize all of the admissible stabilizing controllers in a convex fashion, and cannot cast the optimal decentralized control problem as a convex problem such as in (74.7). However, a Youla parametrization can similarly be used, and while (74.6) does not hold, as the information constraint on the controller is not equivalent to enforcing it on the Youla parameter, it is instead equivalent to a quadratic equality constraint on the parameter [Man93]:

$$K \in S \quad \Longleftrightarrow \quad W_2 + QW_4 - W_1Q - QW_3Q \; = \; 0 \tag{74.13}$$

for stable operators W_1, W_2, W_3, W_4. When returning to the optimal decentralized control problem, this equality constraint replaces the final $Q \in S$ constraint of (74.7). The problem is no longer convex due to the quadratic term, but the overall difficulty is transformed to one well-understood type of constraint, for which many methods exist to approximate optimal solutions.

Other structural constraints, which are neither block diagonal nor quadratically invariant, can be similarly parametrized by first converting them to a perfectly decentralized problem [Rot10].

74.7 Nonlinear Decentralized Control

The framework discussed thus far assumes that the operators, both the plant to be controlled and the possible controllers that we may design for it, are all linear, and when applicable, time-invariant as well. A similar convex parametrization of stabilizing decentralized controllers exists even when the systems are possibly nonlinear and possibly time-varying (NLTV) [Rot06]. The condition allowing for the parametrization then becomes

$$K_1(I \pm GK_2) \in S \qquad \text{for all} \;\; K_1, K_2 \in S.$$

When the plant is stable, the stabilizing controllers may be parametrized similarly to (74.3) [DL82], and when the plant is unstable, the stabilizing controllers may typically be parametrized similarly to (74.4) [AD84]. Similar to quadratic invariance, the above condition then yields the equivalence of the controller and the feedback map satisfying the information constraint (74.5), which then gives the equivalence of the controller and the parameter satisfying the constraint as in (74.6). The convex parametrization of all causal stabilizing decentralized controllers then results, analogous to the linear case with quadratic invariance.

While this condition may appear quite different from quadratic invariance, they actually both reduce to the same conditions when one considers the classes of sparsity constraints or delay constraints, and so these results extend to all of the cases covered in Sections 74.5.1 and 74.5.3.

References

1. V. Anantharam and C. A. Desoer. On the stabilization of nonlinear systems. *IEEE Transactions on Automatic Control*, 29(6):569–572, 1984.
2. C. A. Desoer and R. W. Liu. Global parametrization of feedback systems with nonlinear plants. *Systems and Control Letters*, 1(4):249–251, 1982.
3. V. Manousiouthakis. On the parametrization of all stabilizing decentralized controllers. *Systems and Control Letters*, 21(5):397–403, 1993.
4. M. Rotkowitz, R. Cogill, and S. Lall. A simple condition for the convexity of optimal control over networks with delays. In *Proc. IEEE Conference on Decision and Control*, pages 6686–6691, 2005.
5. M. Rotkowitz, R. Cogill, and S. Lall. Convexity of optimal control over networks with delays and arbitrary topology. *International Journal of Systems, Control, and Communications*, 2:30–54, 2010.

6. M. Rotkowitz and S. Lall. Decentralized control information structures preserved under feedback. In *Proc. IEEE Conference on Decision and Control*, pages 569–575, December 2002.
7. M. Rotkowitz and S. Lall. Affine controller parameterization for decentralized control over Banach spaces. *IEEE Transactions on Automatic Control*, 51(9):1497–1500, September 2006.
8. M. Rotkowitz and S. Lall. A characterization of convex problems in decentralized control. *IEEE Transactions on Automatic Control*, 51(2):274–286, February 2006.
9. M. Rotkowitz. Information structures preserved under nonlinear time-varying feedback. In *Proc. American Control Conference*, pages 4207–4212, 2006.
10. M. Rotkowitz. Parametrization of all stabilizing controllers subject to any structural constraint. In *Proc. IEEE Conference on Decision and Control*, December 2010. submitted.
11. H. S. Witsenhausen. A counterexample in stochastic optimum control. *SIAM Journal of Control*, 6(1):131–147, 1968.
12. D.C. Youla, H.A. Jabr, and J.J. Bonjiorno Jr. Modern Wiener-Hopf design of optimal controllers: part II. *IEEE Transactions on Automatic Control*, 21(3):319–338, 1976.