Massachusetts Institute of Technology

Department of Electrical Engineering and Computer Science

6.245: MULTIVARIABLE CONTROL SYSTEMS by A. Megretski

Lecture 1: Optimization-Based Feedback Design[1]

The goals of the MCS class

The Multivariable Control Systems class covers optimization-based approach to design, analysis, and model reduction of dynamical systems. Hardware issues or laws of nature will not be discussed, except when deriving mathematical models of systems to work with. Instead, the course is about figuring out mathematical description of good linear feedback controllers. A number of recently developed tools for automatic generation of "optimal" controllers will be discussed. Specifying an adequate criterion of optimality for such tools is a major challenge for the design engineer, since typical system specifications cannot be expressed directly in the form of a suitable optimization criteria. The main goal of the Multivariable Control Systems class is to develop good understanding and proper use of the optimization tools of modern control.

Assumptions specific for LTI feedback design

We will work with the design setup shown on Figure 1.1, common for the information processing technology. On the diagram, the plant block stands for the "outside world" – something that is observed only through the sensor channel, and can be affected only through the actuator channel. The controller block can represent any real-time information processing algorithm. The result of a dynamic interaction between the controller and the

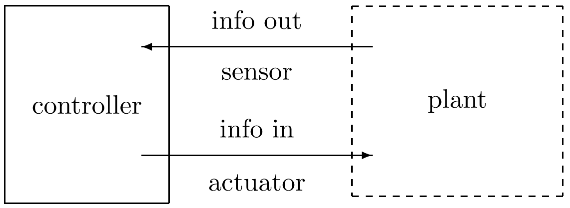

Figure 1.1: controller as an information processing unit

---

[1] Version of February 7, 2001

plant is a *feedback control system*. Controller design is considered "good" if the feedback system possesses certain "desired" features.

We will consider a special case.

- The plant is modeled as a *linear and time invariant* (LTI) system, possibly with uncertainty. Though a real-world plant is most likely to be nonlinear, most allow a fairly good linear approximation, at least as long as the control is not too aggressive. In this case the error introduced by linearization is accounted for by introducing uncertainty into the plant model.
- The controller is restricted to be LTI, too. This is compatible with assuming an LTI plant model, and means that we are considering a very restricted class of simple decision-making algorithms. In practical implementations, an LTI controller is usually a part of a more complex "hybrid" algorithm, being "supervised" by a logical controller. LTI controllers for LTI systems are relatively easy to analyze and optimize (or, one could say, they are the *only* class that allows comprehensive analysis), and do not require much of processor power to implement.
- We assume that there are no restrictions on the flow of information. In a real world situation, such restrictions may be caused by several factors:
  - limited capacity of the communication channels transmitting the information to and from the controller;
  - limited computational power of the processor implementing the controller algorithm;
  - decentralization of the decision-making due to distributed processing of information.

What makes feedback design difficult

Modern LTI feedback methodology offers a number of deep results which help making design decisions and predicting their outcome. However, one soon realizes that these results are only working within a limited environment of specific system models (LTI nominal models with several types of "uncertainty cloud" around them) and specific qualities to be optimized. This appears to be in a sharp contrast with the worlds of software design, communications, and signal processing, where complex highly nonlinear information processing algorithms are designed, tested, and implemented all the time. Below we list several reasons why the feedback design case is different.

- The controller becomes a part of a "feedback loop": its presence changes the way the process behaves. This causes potential stability problems, and makes it difficult to predict the effect a particular controller will have on the functioning of the whole system.
- The designer authority is limited to a relatively small part (the controller) of the whole system. Usually, there are severe restrictions on what can be measured and where actuators can be placed.

- The behavior of a feedback system is very sensitive to the timing of events and signal phases. In particular, only causal signal transformations can be used (hence no ideal low pass filters are allowed!).

In addition, specifications of a feedback system to be designed are usually conflicting. Here is a list of typical objectives.

- Stability of the closed loop behavior.
- Quality transition between different regimes.
- Robustness to modeling errors.
- Good rejection of external disturbances and sensor noises.
- Low complexity of the controller.
- Effective use of resources (energy, supplies, etc.)

For example, improving transient behavior of a feedback system will typically lead to worsening disturbance rejection, robustness, and overall efficiency.

LTI Control Design: Classical vs. Modern

There is some kind of "controversy" between the "classical control" and "modern control" fields. One can often hear that the "classical control" provides simpler, better, and more intuitive solutions, where the "modern control", heavy with irrelevant mathematics, produces high-order and otherwise undesirable controllers via a complicated and counterintuitive design process. The response from the other side is that the "classical control" works for very simple situations only, does not guarantee anything in most cases, and does not tell you whether a particular control design objective is achievable or not.

Most of what is mentioned in the previous paragraph appears to be true to some extent. Consider the following three approaches to LTI feedback design:

- **Optimistic Design:** first decide which closed loop behavior is wanted, then solve (approximately) backwards for the controller;
- **Gradual Refinement:** start with an initial guess, then modify it in the direction of best improvement;
- **Cautious Design:** find which feedback design qualities are easy to optimize; relate these easy-to-optimize qualities to other design qualities, which may be more difficult to optimize; design for "hard-to-optimize" qualities by using optimization algorithms for the "easy-to-optimize" qualities.

The first approach appears to be typical for the "classical" control. It produces very good results when "solving backwards" for a controller is easy. Practically, this means systems with accurate plant models, strong control authority, and good sensors. The first approach is not expected to work well on non-minimum-phase, uncertain, or multi-variable systems.

The second alternative is frequently used to improve a design which is already good. On the other hand, it is easy to miss a good design solution if you start from a bad initial guess.

The third approach is the path taken by what is called the "modern control". Over the course of last 40 years, a number of "easy-to-optimize" objectives has been identified. As a rule, such objectives include stabilization of the feedback system plus minimization of some sensitivity factor, such as

- averaged (over all frequencies) sensitivity to some external input (LQ, LQG, H2 optimization);
- maximal (over all frequencies) sensitivity to some external input (H-infinity optimization);

Other tools available include assessment of robustness with respect to an uncertainty structure within the process model (aided by LMI optimization), optimal model/controller order reduction, and use of linear programming to optimize almost any convex objective on the set of closed loop transfer functions.

Main Elements of Modern LTI Control Theory

There are several well-connected major components of the course.

- **Stabilization via pole placement** is the simplest design technique. It allows one to produce a stabilizing controller which places the poles of the closed loop system at desired locations. Unfortunately, there is little connection between the pole location and quality of feedback. In fact, forcing the poles to become very stable (far in the right half plane) typically produces a system which is very sensitive to external signals and model uncertainty.
- **Uncertain Models of Systems and Signals** introduces formal quantities describing intensity of noises/disturbances and size of uncertainty.
- **Optimal Feedback Control design setup** formulates the controller design objectives as stability of the closed loop system plus minimization of one of the quantities introduced in the previous item.
- **Limitations of Feedback** introduces a fundamental result of modern control: an affine parameterization of all closed loop transfer matrices in a stabilized system. This result is based on the previous developments in stabilization via pole placement, and allows one to describe general limitations on the qualities of feedback control design imposed by the unstable poles and zeros of the open loop process.
- **Well-posedness of feedback design problems.** Not any optimal feedback control design setup is meaningful. Moreover, a typical beginner's setup will turn out to be ill-posed! In this section, we learn about ill-posedness, its meaning, and how to avoid it.
- **H2 optimization (LQG control, Kalman filters, etc.)** is a major optimization tool of modern control.
- **H-infinity optimization** was shown to be more convenient then the H2 optimization when dealing with optimization of robustness. It is the second major technique of LTI feedback optimization.

- **Hankel norm optimization and model order reduction** are used to reduce the complexity of designs obtained via the modern control techniques.
- **LMI/LP optimization (mu-analysis and synthesis, etc.)** is an emerging tool for solving more challenging problems of feedback design.

Complexity of optimization algorithms

The issue of "complexity" (i.e. difficulty of finding a solution) of optimization algorithms plays an important role in understanding modern control. Indeed, the state-of-the-art is such that some of the most attractive controller optimization problems will take years to solve using the tools currently available. It is very common in modern control to trade accuracy of a solution for an efficiency of the algorithm. Even among the optimization problems which can be solved efficiently, the difference in complexity of the solution can be critical in deciding suitability for real-time adaptation and re-design.

The following is a rough classification of complexity in controller optimization.

- **Analytical Solution.** The optimization algorithms in this class (among them pole placement, Hankel norm model reduction, H-Infinity and LQG optimization) require an a-priori fixed number of standard linear algebra operations (such as Shur or Singular value decompositions, matrix multiplications, etc.) to complete. These are the most efficient algorithms. Unfortunately, only a handful of controller design objectives can be optimized analytically.
- **Polynomial Time Algorithms.** The number of standard Linear Algebra operations required by the optimization algorithms in this class to get an ϵ-approximation of the optimum is polynomial with respect to $\log(\epsilon)$ (and the problem dimensions). Typically, polynomial time algorithms are available for *convex problems*, including LMI optimization and linear programming.
- **Smart Random Search.** These algorithms do not have proven polynomial time bounds for the number of iterations needed to reach a good approximation of the optimum, as in "Mu-synthesis" or stabilization with a fixed order controller.

Trademarks of modern control design

Modern control design gives a number of useful general hints for feedback control designer.

- DESIGN FOR AN LTI ENVIRONMENT (may be LTI+uncertainty), as shown on Figure 1.2. Use linearization, gain scheduling to treat nonlinear effects that cannot be ignored.
- DESIGN VIA OPTIMIZATION: choose a "cost functional" which describes the "quality" of a controller, then optimize the "cost".
- USE SUBOPTIMAL DESIGNS: it is often undesirable to push a controller optimization process to its end. Suboptimal controllers have better behavior and are easier to find.

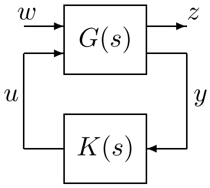

Figure 1.2: Standard design diagram

- USE SPECIAL COST FUNCTIONALS: replace the "true" design criteria with the "artificial", but "convenient" ones, to get one of the "low complexity" optimization problems: H_2, H_∞, linear programming, positive semidefinite programming (LMI's).
- USE PLANT AUGMENTATION: add artificial components to the original system model to better express the design objective.
- DO NOT FIX THE CONTROLLER STRUCTURE. For example, replace the traditional "prefilter" design setup of Figure 1.3 with the setup of Figure 1.4

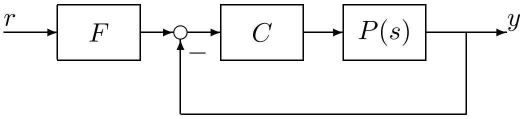

Figure 1.3: Traditional design with a pre-filter

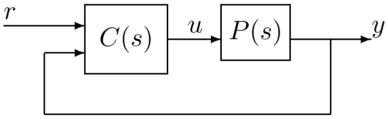

Figure 1.4: Canonical setup for design with a "pre-filter"

- SEARCH FOR A "COST" which leads to a better design in terms of the original objectives.

What to learn in 6.245:

- Which optimization problems have "easy" solutions
- How to avoid ill-posed problems
- What limits performance of any feedback and makes a system "hard to control"
- What to expect as an outcome of an optimization problem