Optimal and Robust Control of Networked Systems

James Anderson

Department of Engineering Science, University of Oxford

Abstract

This document describes the work carried out during and subsequent collaborative efforts resulting from a 5-week systems-NET *Immersive Secondment* spent at the Information Systems Laboratory in the Department of Electrical Engineering at Stanford University, California, hosted by Professor Sanjay Lall. In addition to being exposed to a research group at the cutting edge of mathematical control theory and learning how to formulate engineering problems in a rigorous manner, I was able to attend lectures and seminars in related topics, travel to nearby academic institutions to discuss my work and formulate a collaboration that will last past the end of the secondment.

Professor Sanjay Lall: http://lall.stanford.edu/
ISL, Stanford: http://isl.stanford.edu/

Problem Background

Designing and maintaining complex engineering systems is one of the great challenges for engineering in the 21st century. Systems such as the Internet, power grids, transportation networks and off-shore wind farms, for example, are all of such complexity that simply throwing computer power at them and using *ad-hoc* techniques to maintain and optimize their performance is not always possible. Furthermore, even when they are possible, such methods provide no guarantee of performance and offer little in the way of robustness. A compounding factor is that these methods are typically system-specific and do not shed light on how to redesign or *scale up the system* should the demand arise.

Systems like the examples listed above are frequently referred to as *networked* or *complex systems.* That is, they consist of multiple modules or subsystems that communicate with each other to form a larger aggregated system. This is in direct contrast to traditional *centralized systems* where a controller is designed for a single system and produces a control signal knowing (or having an estimate of) the complete state of the system.

Feedback controllers are designed with two purposes in mind. Their primary job is to safeguard against uncertainty. Uncertainty arises for numerous reasons, and the mathematical models that are necessary for design can only ever approximate the real system. Some parameters (e.g. spring constants, diffusion coefficients, etc.) may be difficult to measure accurately, and this difficulty must be taken into account when the controller is designed. Some processes defy a simple mathematical representation and thus are lumped together and treated as a single *bounded disturbance.* Feedback control specifically guards against this uncertainty. The second reason for implementing a controller is to improve the dynamic performance of a system. For instance, we may wish to ensure that the voltage across a certain component never exceeds a given value even though we do not know ahead of time how large the input may be. Another objective may be that the system settle to a steady state as quickly as possible. Different criteria can be formulated as cost functions and incorporated into the design process.

Designing controllers for networked systems is problematic in many respects. Existing computational methods scale poorly with systems size, and the cost of communication between subsystems may be expensive, meaning that a centralized solution (where each system knows the state and output of every other system) is not practically realizable. In this work we considered decentralized solutions that also take into account communication delay.

Imnersive Secondment Goals

- Establish a collaboration with Professor Lall that will offer potential for a lasting partnership beyond the duration of the Immersive Secondment.
- Formulate a research problem that will open up a new body of work on both on the mathematical and the computational front.
- Attend control and optimization classes at Stanford.
- Solicit feedback and initiate dialogue from other research groups in the area e.g. Caltech and Berkeley.
- Publish research results in a peer-reviewed journal or conference proceedings.

Research Problem: Robust and Optimal Decentralized Control

Problem Description:

In this section the standard control design problem that includes robust and optimal performance criteria will first be outlined. The networked version of the problem which was formulated on the secondment will then be described.

The general feedback control problem is concerned with the system interconnection depicted in Figure 1. We define the following signals and systems:

- **G:** System to be controlled, frequently referred to as the *plant.*
- **K:** Controller
- **w:** Exogenous input
- **u:** Control signal
- **y:** Measured outputs (available to the controller)
- **z:** Exogenous output

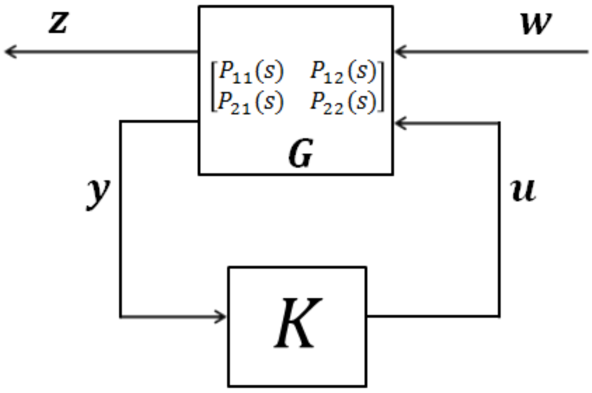

Figure 1: *The feedback interconnection of the controlled system. The plant G takes inputs from the controller and the outside world and sends measurements to the controller and the environment. The map from input w to output z is given by T_{zw}.*

We will assume that the plant, G, can be modelled as a finite-dimensional, Linear Time Invariant (LTI) system which takes the form:

$$\dot{x}(t) = Ax(t) + B_1 w(t) + B_2 u(t)$$
$$z(t) = C_1 x(t) + D_{11} w(t) + D_{12} u(t)$$
$$y(t) = C_2 x(t) + D_{21} w(t) + D_{22} u(t)$$

With this realization, $x(t) \in \mathbb{R}^n$ is the state vector, $z(t) \in \mathbb{R}^m$ is the output and $y(t) \in \mathbb{R}^p$is the measured output. To lighten notation we will introduce the notation

$$G \triangleq \begin{bmatrix} A & B_1 & B_2 \\ C_1 & D_{11} & D_{12} \\ C_2 & D_{21} & D_{22} \end{bmatrix} = \begin{bmatrix} A & B \\ C & D \end{bmatrix},$$

where it is assumed that the matrices are all of compatible dimensions. The control problem is now to design a control signal $u(t)$ such that the mapping from w to z is made as small as possible. In this setting the input signal $w$can be thought of as being a disturbance on the system which we observe through the output signal z. This can be written as

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

$$u = K(s)y$$

where s denotes the fact that the partitioning of G above represents input output subsystems rather than static matrices. Specifically this notation indicates that we have taken Laplace transforms of the time domain system and now have a frequency domain realization of the form

$$\boldsymbol{G = C(sI - A)^{-1}B + D.}$$

The control signal u is generated via the controller, K (which we are to design) ,which is also a dynamical system of the form

$$\begin{aligned}\dot{\boldsymbol{x}}_{\boldsymbol{k}}(\boldsymbol{t}) &= \boldsymbol{A_k x_k(t) + B_k y(t)}\\ \boldsymbol{u(t)} &= \boldsymbol{C_k x_k(t) + D_k y(t)}\end{aligned}$$

Here $x_k(t) \in \mathbb{R}^k$ is the state of the controller. Thus the controller takes as its inputs the measured output from the plant and produces as an output the control signal $u(t)$ which is an input to the plant, i.e. closing the loop between the plant and controller.

The closed loop transfer function from input to output is given by the *Linear Fractional Transform* (LFT):

$$\boldsymbol{T_{zw}(G, K) \triangleq G_{11} + G_{12}K(I - G_{22}K)^{-1}G_{21}}$$

i.e. $z = T_{zw}(G, K)w$.

We can now state the optimal control problem:

$$\begin{aligned}&\boldsymbol{minimize}\ \|\boldsymbol{G_{11} + G_{12}K(I - G_{22}K)^{-1}G_{21}}\|\\ &\boldsymbol{s.t.} \qquad \boldsymbol{K\ stabilizes\ G}\end{aligned}$$

The optimization problem above states that we are searching for a controller K that stabilizes the plant and makes the mapping from input to output small in the sense of some norm $\|T_{zw}\|$. Robustness is accounted for by the fact that some of the input signals in the input vector w may represent noise, i.e. Gaussian white noise signals. The idea of minimizing the closed loop norm makes intuitive sense if you think of the output signal z as being the error between, say, the target/desired velocity of a vehicle and the actual velocity when considering a cruise-control system.

Classical $\boldsymbol{H_2}$ Solution:

Let us assume that the norm we wish to minimize is the H_2-norm. If we further assume that w is Gaussian white noise and that $D_{11} = 0$ and $D_{22} = 0$ then the optimal controller K that solves the optimization problem above is given by

$$\boldsymbol{K = \begin{bmatrix} A + B_2F + LC_2 & -L \\ F & 0 \end{bmatrix},}$$

with

$$\boldsymbol{L = -YC_2^*, \quad F = -B_2^*X,}$$

Where X and Y are positive definite solutions to the Algebraic Riccati equations:

$$\begin{aligned}\boldsymbol{A^*X + XA - XB_2B_2^*X + C_1^*C_1 = 0,}\\ \boldsymbol{AY + YA^* - YC_2^*C_2Y + B_1B_1^* = 0.}\end{aligned}$$

There are several standard technical assumptions that must be made for this solution to exist and to be well posed. However, for clarity of presentation they are omitted here; the full details can be found in any standard text such as Zhou, Doyle and Glover '96 or Dullerud and Paganini '00.

The solution to this problem is one of the highlights of modern control theory. Moreover, there are multiple methods for computing the solutions to the two Riccati equations required to construct the controller. In contrast, for the decentralized case we have neither an analytic solution or scalable numerical methods for solving the problem. Next, the decentralized problem that was constructed while I was at Stanford is presented.

The Decentralized Problem:

The plant controller configuration was shown graphically in Figure 1 for the standard, centralized case. In that setup, the plant, G (i.e. the system to be controlled), consists of a single system to be controlled by a single controller. In the network case the goal is to construct multiple systems that interact with each other which are controlled via multiple controllers. The problem is described as decentralized because no individual controller in the network knows the full state information of all the subsystems. Thus, control decisions must be made using only local information but with the aim of minimizing a global cost.

For general network structures the decentralized control problem is known to be non-convex. This implies that finding optimal solutions with a reasonable amount of computing power is unrealistic. Even for a subclass of problems for which we know there exists a convex optimal solution (*c.f.* Rotkowitz and Lall '06), finding the solution is still an incredibly difficult task.

The 2-player output feedback problem is the *simplest* decentralized control problem that can be formulated, yet it was only this year that a solution to the problem was found (by Lessard and Lall, '14). The problem set-up is described in Figure 2. Consider two systems, G_1 and G_2, which interact with each other. The first system G_1 is not influenced by G_2, but it does exert influence over G_2. The controller C_1 that regulates the output of G_1 receives information only from G_1, whereas the controller C_2 that regulates G_2 receives information about both G_1 and G_2. Thus the first controller can influence both systems, whereas the second controller can only influence the second system. Mathematically this can be represented as

$$\boldsymbol{G_{22}(s) = \begin{bmatrix} P_{11}(s) & 0 \\ P_{21}(s) & P_{22}(s) \end{bmatrix}, K(s) = \begin{bmatrix} K_{11}(s) & 0 \\ K_{21}(s) & K_{22}(s) \end{bmatrix}},$$

using the notation introduced to describe the centralized problem described earlier. Thus the decentralized control problem is to construct a controller K with this triangular structure that minimizes the cost $\|T_{zw}\|$.

The problem that we will aim to formulate and solve is the 2-player problem with communication delays. Most real-world systems will have to endure delays in being able to send and receive data. The delays may occur for several reasons: large physical distance between the subsystems, low bandwidth, slow communication infrastructure between subsystems or high traffic flow across the network.

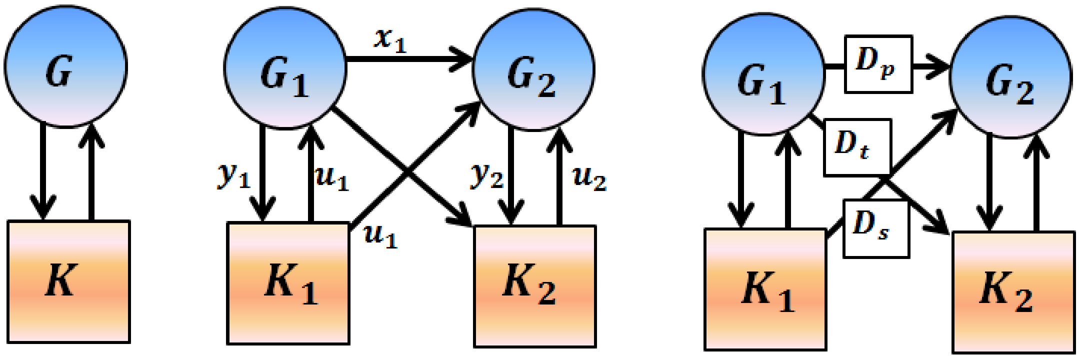

Figure 2: The standard centralized problem (left), the decentralized 2-player problem (middle) and the delayed 2-player, decentralized problem (right).

Whatever the cause of the delay, it is important that the systems being controlled are robust to its effects, i.e. the performance of the system must not be impaired in the presence of delays. From a mathematical point of view, the delayed and decentralized problem is much harder than the non-delayed counterpart. If the delays obey some simple inequalities it is known that the problem is convex (this is shown using the notion of *quadratic invariance,* Rotkowitz and Lall '06). However, obtaining the optimal solution will still be very challenging as the introduction of the delays means that our system descriptions are now *infinite-dimensional*, which is a much harder class of system to work with, both computationally and analytically.

Classes Attended at Stanford

During the secondment I attended I the following classes:

- EE365: Stochastic Control, taught by Professor Sanjay Lall
  http://web.stanford.edu/class/ee365/
- EE364b: Convex Optimization II, taught by Professor Stephen Boyd
  http://stanford.edu/class/ee364b/

EE365 is a new course being taught for the second year. The class emphasized practical solutions to problems that involve decision processes that are subject to uncertainty. The primary computational tool is *dynamic programming* (DP). It is possible that DP will be used to obtain numerical insight into the solution of the decentralized control problem that was formulated above.

Professor Boyd's Convex Optimization class has set the standard for teaching advanced optimization concepts to graduate students in engineering, mathematics and operations research. The class is problem-orientated but with a strong mathematical core. This is the first time this class has been offered in a few years so I was very lucky to attend. The accompanying textbook, slides and full set of video lectures are commonly used by other instructors wishing to teach similar courses. I met with Professor Boyd to discuss some of my research related to the decentralized problem. His feedback was extremely positive: he described it as a 'great problem to work on'.

Professor Stephen Boyd: http://web.stanford.edu/~boyd/

California Institute of Technology & Berkeley

In addition to spending time at Stanford I also travelled to Caltech in Pasadena, where I spent the day at the Control and Dynamical Systems (CDS) option in the Department of Computing + Mathematical Sciences. I was hosted by Professor John Doyle and his group. Whilst there I presented my research at the Doyle group meeting and discussed my previous research and the proposed project with Professor Doyle and his PhD students. Two of the students in John's group are working on topics directly related to the decentralized problem that we formulated at Stanford. In particular, Nikolai Matni is looking at a discrete-time decentralized control problem formulation, while Yuh-Shyang Wang is investigating decentralized control for systems using spatial information. It was very helpful to discuss how these different ideas and approaches can link up to tackle the challenge of scalable network control.

Professor John Doyle: http://www.cds.caltech.edu/~doyle/wiki/
Nikolai Matni: http://www.cds.caltech.edu/~nmatni/home/Welcome.html
Yuh-Shyang Wang: http://www.cds.caltech.edu/~yw4ng/

Recently one of Professor Lall's former PhD students, Dr Laurent Lessard has taken up a post-doctoral position at UC Berkeley. During my stay at Stanford, Laurent visited several times and will form part of the ongoing collaboration. Laurent provided a lot of helpful guidance and insight into related work during the time we spent formulating the delay problem.

Dr Laurent Lessard: http://www.laurentlessard.com/publications/

Long Term Goals

There are several long-term goals resulting from this Immersive Secondment. On the purely theoretical side there are numerous possibilities:

1. Consider more general networks than the 2-player chain.
2. Investigate analytic results when using a different norm to optimize over, for example the H_∞-norm.
3. Explore numerical approaches to obtaining the controller, for example Linear Matrix Inequality (LMI) methods.

It is also very important that such work should have some practical use. Over the course of the next year a mobile robotics test-platform will be assembled at Oxford. It is hoped that the theory developed as a result of this secondment will be implemented on the test-platform. One of the objectives will be to perform some sort of coordinated movement task where the robots are allowed to communicate locally with each other. The delays in the problem description described earlier will then represent the communication delay between the robots.

In the long term it is hoped that in combination with the outcomes listed above it will be possible to obtain industry support and put together an early-career grant to submit to the Engineering and Physical Research Council.

In the two months since the secondment finished I have subsequently met with Professor Lall and Dr Lessard at MTNS 2014 and The 2014 American Control Conference respectively.

References

K. Zhou, J. C. Doyle and K. Glover, *Robust and Optimal Control,* Prentice Hall, Upper Saddle Hill, NJ, 1996.

G. E. Dullerud and F. Paganini, *A Course in Robust Control Theory: A Convex Approach*, Texts in Applied Mathematics, Springer 2000.

M. Rotkowtz and S. Lall, *A Characterization of Convex Problems in Decentralized Control,* IEEE Transactions on Automatic Control, 51(2), 2006.

L. Lessard and S. Lall, *Optimal Control of Two-Player Systems with Output Feedback*, To appear in IEEE Transactions on Automatic Control. Available online at http://lall.stanford.edu/publications/tpof .