Networked Sensing, Estimation and Control Systems

Vijay Gupta
University of Notre Dame

Richard M. Murray
California Institute of Technology

Ling Shi
Hong Kong University of Science and Technology

Bruno Sinopoli
Carnegie Mellon University

DRAFT v1.1b, 11 January 2010

Contents

Preface

The area of “Networked Control Systems” has emerged over the past decade as a subdiscipline in control theory in which the flow of information in a system takes place across a communication network. Unlike traditional control systems, where computation and communications are usually ignored, the approaches that have been developed for networked control systems explicitly take into account various aspects of the communication channels that interconnect different parts of the overall system and the nature of the distributed computation that follows from this structure. This leads to a new set of tools and techniques for analysis and design of networked control systems that builds on the rich frameworks of communication theory, computer science and control theory.

This book is based on a series of courses that the authors have developed over the past several years, starting with a joint course taught at Caltech in Spring 2006. These courses were typically taken by students who have a good grounding in the basic techniques of control systems but may not have a strong background in computer science or some aspects of communication theory. While the level of mathematical detail in the book should allow it to be accessible to juniors or seniors in engineering, the treatment is tuned for first and second year graduate students in engineering or computer science. Some tutorial material on estimation theory is included, as well as a brief review of key concepts in graph theory that are needed primarily in the second half of the text.

The book is intended for researchers who are interested in the analysis and design of sensing, estimation and control systems in a networked setting. We focus primary on the effects of the network on the stability and performance of the system, including the effects of packet loss, time delay and distributed computation. We have attempted to provide a broad view of the field, in the hope that the text will be useful to a wide crossection of researchers. Most of the results are presented in the discrete time setting, with references to the literature for the continuous time analogs. We have also attempted to include a review of the current literature at the end of each chapter, with an emphasis on papers that are frequently referenced by others, along with some directions for future research, when appropriate. To keep the material focused, we have chosen to only touch on material on optimization-based control (e.g., receding horizon control) or protocols for distributed systems, although these are often an integral part of complex networked control systems. References to the literature are given for readers interested in these important topics.

The book is organized into two main parts: a set of background chapters and the core material. Chapter 1 gives an introduction to the topic of networked control systems, including some driving application examples. Chapters 2–4 cover a collection of topics that are used throughout the remainder of the text. We assume familiarity with standard topics in estimation and control theory, including random processes, Kalman filtering and linear state space control theory, and provide only a quick review of this material in Chapter 2 to define the notation we will use and present some of the basic definitions and formulas. Chapters 3 and 4 complete the background chapters by giving concise overviews of the relevant results in information theory and Markovian jump linear systems, on which many of the later results of the book are built. These background chapters can be reviewed quickly for students and researchers already familiar with this material.

The core material on networked control systems is presented in Chapters 5 through 9. We begin by looking at the case of sensing, estimation and control of a single process across a communication channel, beging with the effects of rate limits in the channel in Chapter 5 and then the effects of packet loss in Chapter 6. Both of these chapters considers the cases where the communication channel affects on the measurements received from the sensor and where the channel affects both the measurements and the actuation commands. In Chapter 7 we begin to look at the problem of control over a graph, starting with an introduction to graph theory and the problem of consensus. Chapters 8 and 9 then go on to consider the distributed estimation and control problems, where one can have multiple processes, sensors, actuators, estimators and controllers distributed over a communications network. In each of these chapters on the core material we have attempted to present a unified view of many of the most recent and relevant results in network control, with the goal of establishing a foundation on which more specialized results of interest to specific groups can be covered.

The topics in the text have been taught by the authors and our colleagues in a variety of formats. In a semester-long, graduate course, it should be possible to cover most of the material in the book, assuming the students have good working knowledge of random processes, estimation theory and linear control systems. We have also used the material in the text for week-long short courses for masters and PhD students, where we cover the results in the background chapters in four 90 minute lectures, then spend 1–2 lectures on each of the remaining chapters. The material is fairly modular, so that the order of teaching the material can be varied according to the tastes of the instructor. The dependencies of the chapters are shown in Figure 1.

The full text for this manuscript, along with additional supplemental material, is available on a companion web site:

`http://www.cds.caltech.edu/~murray/amwiki/NCS`

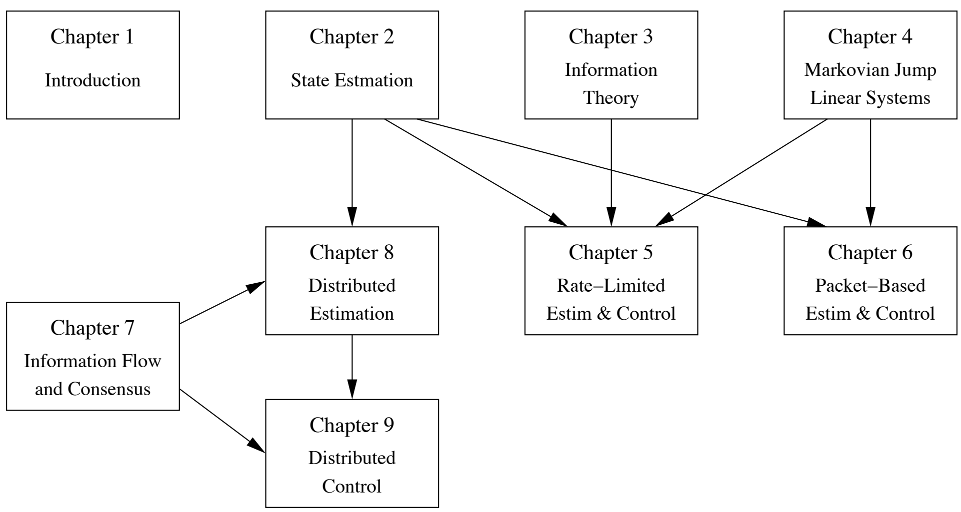

Figure 1: Dependencies of the chapters in the text.

Notation

This is an internal chapter that is intended for use by the authors in fixing the notation that is used throughout the text. In the first pass of the book we are anticipating several conflicts in notation and the notes here may be useful to early users of the text.

General mathematics

- Use $*$ for expressions that are not given explicity
- Matrix transpose: A^T

System dynamics

We focus on linear discrete time systems

$$\begin{aligned} x_{k+1} &= Ax_k + Bu_k + w_k \\ y_k &= Cx_k + v_k. \end{aligned}$$

The system is described by the state $x \in \mathbb{R}^n$, inputs $u \in \mathbb{R}^p$ and outputs $y \in \mathbb{R}^m$. Disturbances are represented by the random process w_k, which we typically take to be zero mean, white Gaussian noise with covariance matrix $\Sigma_W \geq 0$. Measurement noise is represented by the Gaussian random process v_k with covariance matrix $\Sigma_V > 0$. For systems with multiple sensors, we use the notation y^j to represent the jth output and use corresponding superscripts for the other relevant quantities.

In the few instances that we use continuous time dyanmics, these are written as

$$\begin{aligned} \frac{dx}{dt} &= Ax + Bu + w \\ y &= Cx + v \end{aligned}$$

Note that for both the continuous and discrete dynamics we leave out the direct term (Du). We should point out the first time these equations come up in a chapter whether it is easy to include the direct term or whether not everything extends directly.

Random variables and processes

- Expectation: $\mathbb{E}[X]$, $\mathbb{E}_Y[X]$ for the expectation of X over Y.
- Mean: μ, μ_X for a random variable/vector X
- Variance: σ^2, σ_X^2 for a (scalar) random variable X; Σ, Σ_X for a random vector X; Σ_{XY} for the cross-covariance

Additional mathematical notation

- Lists and sets: A index set is can be written inline as $\{X_i : i = \min, \dots, \max\}$ or as a displayed equation:
$$\{X_i\}_{i=\min}^{\max}.$$

Observer dynamics

We write $P > 0$ for the covariance of the estimation error. The observer for a discrete time linear system is written as a prediction step,

$$\begin{aligned}\hat{x}_{k|k-1} &= A\hat{x}_{k-1|k-1} + Bu_{k-1},\\ P_{k|k-1} &= AP_{k-1|k-1}A^T + \Sigma_{W\,k-1},\end{aligned}$$

followed by a correction step,

$$\begin{aligned}\hat{x}_{k|k} &= \hat{x}_{k|k-1} + K_k(y_k - C\hat{x}_{k|k-1}),\\ P_{k|k} &= P_{k|k-1} - P_{k|k-1}C^T(CP_{k|k-1}C^T + R_k)^{-1}CP_{k|k-1}.\end{aligned}$$

The gain matrix for the estimator is given by K (for Kalman). The gain matrix for a state space controller can either by L or possibly F (?).