CONTROL TECHNIQUES FOR COMPLEX NETWORKS

Power grids, flexible manufacturing, cellular communications: interconnectedness has consequences. This remarkable book gives the tools and philosophy you need to build network models detailed enough to capture essential dynamics but simple enough to expose the structure of effective control solutions and to clarify analysis.

Core chapters assume only prior exposure to stochastic processes and linear algebra at the undergraduate level; later chapters are for advanced graduate students and researchers/practitioners. This gradual development bridges classical theory with the state of the art. The workload model that is the basis of traditional analysis of the single queue becomes a foundation for workload relaxations used in the treatment of complex networks. Lyapunov functions and dynamic programming equations lead to the celebrated MaxWeight policy along with many generalizations. Other topics include methods for synthesizing hedging and safety stocks, stability theory for networks, and techniques for accelerated simulation.

Examples and figures throughout make ideas concrete. Solutions to end-of-chapter exercises are available on a companion Web site.

Sean Meyn is a professor in the Department of Electrical and Computer Engineering and director of the Division and Control Laboratory of the Coordinated Science Laboratory at the University of Illinois. He has served on the editorial boards of several journals in areas of systems and control and applied probability, and he is coauthor with Richard Tweedie of *Markov Chains and Stochastic Stability*, which won the 1994 ORSA/TIMS Best Publication in Applied Probability Award.

CONTROL TECHNIQUES FOR COMPLEX NETWORKS

SEAN MEYN
University of Illinois, Urbana-Champaign

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo, Delhi

Cambridge University Press
32 Avenue of the Americas, New York, NY 10013-2473, USA

www.cambridge.org
Information on this title: www.cambridge.org/9780521884419

First published 2008

Printed in the United States of America

A catalog record for this publication is available from the British Library.

Library of Congress Cataloging in Publication Data

Meyn, S. P. (Sean P.)
Control techniques for complex networks / Sean Meyn.
p. cm.
Includes bibliographical references and index.
ISBN 978-0-521-88441-9 (hardback)
1. Computer networks. 2. Control theory. I. Title.
TK5105.5.M49 2008
004.6–dc22 2007035250

ISBN 978-0-521-88441-9 hardback

Contents

Cambridge University Press
978-0-521-88441-9 - Control Techniques for Complex Networks
Sean Meyn
Frontmatter
More information

Cambridge University Press
978-0-521-88441-9 - Control Techniques for Complex Networks
Sean Meyn
Frontmatter
More information

List of Illustrations

Preface

A representative of a major publishing house is on her way home from a conference in Singapore, excited about the possibility of a new book series. On the flight home to New York she opens her blackberry organizer, adding names of new contacts, and is disappointed to realize she may have caught the bug that was bothering her friend Alex at the café near the conference hotel. When she returns home she will send Alex an email to see how she's doing and to make sure this isn't a case of some new dangerous flu.

Of course, the publisher is aware that she is part of an interconnected network of other business men and women and their clients: Her value as an employee depends on these connections. She depends on the transportation network of taxis and airplanes to get her job done and is grateful for the most famous network today that allows her to contact her friend effortlessly even when separated by thousands of miles. Other networks of even greater importance escape her consciousness, even though consciousness itself depends on a highly interconnected fabric of neurons and vascular tissue. Communication networks are critical to support the air traffic controllers who manage the airspace around her. A supply chain of manufacturers makes her book business possible, as well as the existence of the airplane on which she is flying.

Complex networks are everywhere. Interconnectedness is as important to business men and women as it is to the viruses who travel along with them.

Much of the current interest in networks within physics and the biological sciences is phenomenological. For example, given a certain degree of connectivity between individuals, what is the likelihood that a virus will spread to the extinction of the planet? Degree and mode of connectivity in passive agents can combine to form images resembling crystals or snowflakes [463].

The main focus within our own bodies is far more utilitarian. Endocrine, immune, and vascular systems adjust chemical reactions to maintain equilibria in the face of ongoing attacks from disease and diet. In biology this is called *homeostasis*. In this book, the regulation of a network is called *control*.

It is not our goal to take on biology, computer science, communications, and operations research in a single volume. Rather, the intended purpose of this book is an introduction to a rapidly evolving engineering discipline. The examples come from applications in which complexity is real, but less daunting than that found in the human brain. We describe methods to model networks in order to capture essential structure, dynamics, and uncertainty. Based on these models we explore ways to visualize network behavior so that effective control techniques can be synthesized and evaluated.

Modeling and control. The operator of an electric power grid hopes to find a network model that will help form predictions of supply and demand to maintain stability of the power network. This requires the expertise of statisticians, economists, and power engineers. The resulting model may provide useful simulations for forecasting, but will fail entirely for our purposes. This book is about control, and for this it is necessary to restrict to models that capture essential behavior, but no more.

Modeling for the purposes of control and the development of control techniques for truly complex networks has become a major research activity over the past two decades. Breakthroughs obtained in the stochastic networks community provide important tools that have had real impact in some application areas, such as the implementation of MaxWeight scheduling for routing and scheduling in communications. Other breakthroughs have had less impact due in part to the highly technical and mathematical language in which the theory has developed. The goal of this book is to expose these ideas in the simplest possible setting.

Most of the ideas in this book revolve around a few concepts.

(i) The *fluid model* is an idealized deterministic model. In a communication network a unit of “fluid” corresponds to some quantities of packets; in a power network this might correspond to a certain number of megawatts of electricity.

A fluid model is often a starting point to understand the impact of topology, processing rates, and external arrivals on network behavior. Based on the fluid model we can expose the inherent conflict between short-sighted control objectives, longer-range issues such as recovery from a singular external disruption, and truly long-range planning such as the *design* of appropriate network topology.

(ii) Refinements of the fluid model are developed to capture variability in supply, demand, or processing rates. The *controlled random walk* model favored in this book is again a highly stylized model of any real network, but contains enough structure to give a great deal of insight and is simple enough to be tractable for developing control techniques.

For example, this model provides a vehicle for constructing and evaluating *hedging* mechanisms to limit exposure to high costs, and to ensure that valuable resources can operate when needed.

(iii) The concept of *workload* is developed for the deterministic and stochastic models. Perhaps the most important concept in this book is the *workload relaxation* that provides approximations of a highly complex network by a far simpler one. The approximation may be crude in some cases, but its value in attaining intuition can be outstanding.

(iv) Methods from the stability theory of Markov models form a foundation in the treatment of stochastic network models. Lyapunov functions are a basis of dynamic programming equations for optimization, for stability and analysis, and even for developing algorithms based on simulation.

What's in here? The book is divided into three parts. The first part, entitled Modeling and Control, contains numerous examples to illustrate some of the basic concepts developed in the book, especially those topics listed in (i) and (ii) concerning the fluid and CRW models. Lyapunov functions and the dynamic programming equations are introduced; based on these concepts we arrive at the MaxWeight policy along with many generalizations.

Workload relaxations are introduced in Part II. In these three chapters we show how a cost function defined for the network can be "projected" to define the *effective cost* for the relaxation. Applications to control involve first constructing a policy for the low-dimensional relaxation, and then translating this to the original physical system of interest. This translation step involves the introduction of hedging to guard against variability.

Most of the control techniques are contained in the first two parts of the book. Part III, entitled Stability and Performance, contains an in-depth treatment of Lyapunov stability theory and optimization. It contains approximation techniques to explain the apparent solidarity between control solutions for stochastic and deterministic network models. Moreover, this part of the book develops several approaches to performance evaluation for stochastic network models.

Who's it for? The book was created for several audiences. The gradual development of network concepts in Parts I and II was written with the first-year graduate student in mind. This reader may have had little exposure to operations research concepts, but some prior exposure to stochastic processes and linear algebra at the undergraduate level.

Many of the topics in the latter chapters are at the frontier of stochastic networks, optimization, simulation, and learning. This material is intended for the more advanced graduate student, as well as researchers and practitioners in any of these areas.

Acknowledgments

This book has been in the making for 5 years and over this time has drawn inspiration and feedback from many. Some of the ideas were developed in conjunction with students, including Mike Chen, Richard Dubrawski, Charuhas Pandit, Rong-Rong Chen, and David Eng. In particular, the numerics in Section 7.2 are largely taken from Dubrawski's thesis [151], and the *diffusion heuristic* for hedging is based on a paper with Chen and Pandit [103]. Section 9.6 is based in part on research conducted with Chen [105].

My collaborators are a source of inspiration and friendship. Many of the ideas in this book revolve around stochastic Lyapunov theory for Markov processes, which is summarized in the appendix. This appendix is essentially an abridged version of my book coauthored with Richard Tweedie [368]. Vivek Borkar's research on Markov decision theory (as summarized in [70]) has had a significant influence on my own view of optimization. My interest in networks was sparked by a lecture presented by P. R. Kumar when he was visiting the Australian National University in 1988 while I resided there as a postdoctoral fellow. He became a mentor and a coauthor when I joined the University of Illinois the following year. I learned of the beauty of simulation theory from the work of Peter Glynn and his former student Shane Henderson. More recent collaborators are Profs. In-Koo Cho, David Gamarnik, Ioannis Kontoyiannis, and Eric Moulines, who have provided inspiration on a broad range of topics. I am grateful to Devavrat Shah and Damon Wischik for sharing insights on the "input-queued switch," and for allowing me to adapt a figure from their paper [435] that is used to illustrate workload relaxations in Section 6.7.1. Pierre L'Ecuyer shared his notes on simulation from his course at the University of Montréal, and Bruce Hajek at the University of Illinois shared his lecture notes on communication networks.

Profs. Cho, Kontoyiannis, Henderson, and Shah have all suggested improvements on exposition, or warned of typos. Sumit Bhardwaj, Jinjing Jiang, Shie Mannor, Eric Moulines, and Michael Veatch each spent significant hours pouring through selected chapters of draft text and provided valuable feedback. Early input by Veatch moved the book toward its present organization, with engineering techniques introduced first, and harder mathematics postponed to later chapters.

Any remaining errors or awkward prose are, of course, my own.

It would be impossible to write a book like this without financial support for graduate students and release-time for research. I am sincerely grateful to the National Science Foundation, in particular, the Division of Electrical, Communications & Cyber Systems, for ongoing support during the writing of this book. The DARPA ITMANET initiative, the Laboratory for Information and Decision Systems at MIT, and United Technologies Research Center provided support in the final stages of this project during the 2006–2007 academic year.

Equally important has been support from my family, especially during the last months, when I have been living away from home. Thank you Belinda! Thank you Sydney and Sophie! And thanks also to all the poodles at South Harding Drive.

Cambridge University Press
978-0-521-88441-9 - Control Techniques for Complex Networks
Sean Meyn
Frontmatter
More information

Dedication

It was a sad day on June 7, 2001, when Richard Tweedie died at the peak of his career. A brief survey of his contributions to applied probability and statistics can be found in [154].

In memory of his friendship and collaboration, and in honor of his many contributions to our scientific communities, this book is dedicated to Richard.

CONTROL TECHNIQUES FOR COMPLEX NETWORKS