Analytical Methods for Network Congestion Control

Synthesis Lectures on Communication Networks

Synthesis Lectures on Communication Networks is an ongoing series of 50- to 100-page publications on topics on the design, implementation, and management of communication networks. Each lecture is a self-contained presentation of one topic by a leading expert. The topics range from algorithms to hardware implementations and cover a broad spectrum of issues from security to multiple-access protocols. The series addresses technologies from sensor networks to reconfigurable optical networks.
The series is designed to:

- Provide the best available presentations of important aspects of communication networks.
- Help engineers and advanced students keep up with recent developments in a rapidly evolving technology.
- Facilitate the development of courses in this field

Analytical Methods for Network Congestion Control
Steven H. Low
2017

Advances in Multi-Channel Resource Allocation: Throughput, Delay, and Complexity
Bo Ji, Xiaojun Lin, Ness B. Shroff
2016

A Primer on Physical-Layer Network Coding
Soung Chang Liew, Lu Lu, Shengli Zhang
2015

Sharing Network Resources
Abhay Parekh and Jean Walrand
2014

Wireless Network Pricing
Jianwei Huang and Lin Gao
2013

Performance Modeling, Stochastic Networks, and Statistical Multiplexing, Second Edition
Ravi R. Mazumdar
2013

Packets with Deadlines: A Framework for Real-Time Wireless Networks
I-Hong Hou and P.R. Kumar
2013

Energy-Efficient Scheduling under Delay Constraints for Wireless Networks
Randall Berry, Eytan Modiano, and Murtaza Zafer
2012

NS Simulator for Beginners
Eitan Altman and Tania Jiménez
2012

Network Games: Theory, Models, and Dynamics
Ishai Menache and Asuman Ozdaglar
2011

An Introduction to Models of Online Peer-to-Peer Social Networking
George Kesidis
2010

Stochastic Network Optimization with Application to Communication and Queueing Systems
Michael J. Neely
2010

Scheduling and Congestion Control for Wireless and Processing Networks
Libin Jiang and Jean Walrand
2010

Performance Modeling of Communication Networks with Markov Chains
Jeonghoon Mo
2010

Communication Networks: A Concise Introduction
Jean Walrand and Shyam Parekh
2010

Path Problems in Networks
John S. Baras and George Theodorakopoulos
2010

Performance Modeling, Loss Networks, and Statistical Multiplexing
Ravi R. Mazumdar
2009

Network Simulation
Richard M. Fujimoto, Kalyan S. Perumalla, and George F. Riley
2006

Analytical Methods for Network Congestion Control

Steven H. Low

ISBN: 978-3-031-79274-8 paperback
ISBN: 978-3-031-79275-5 ebook

DOI 10.1007/978-3-031-79275-5

A Publication in the Springer series
SYNTHESIS LECTURES ON COMMUNICATION NETWORKS

Lecture #18
Series Editor: R. Srikant, *University of Illinois at Urbana-Champaign*
Founding Editor Emeritus: Jean Walrand, *University of California, Berkeley*
Series ISSN
Print 1935-4185 Electronic 1935-4193

Analytical Methods for Network Congestion Control

Steven H. Low
California Institute of Technology (Caltech)

SYNTHESIS LECTURES ON COMMUNICATION NETWORKS #18

ABSTRACT

The congestion control mechanism has been responsible for maintaining stability as the Internet scaled up by many orders of magnitude in size, speed, traffic volume, coverage, and complexity over the last three decades. In this book, we develop a coherent theory of congestion control from the ground up to help understand and design these algorithms. We model network traffic as fluids that flow from sources to destinations and model congestion control algorithms as feedback dynamical systems. We show that the model is well defined. We characterize its equilibrium points and prove their stability. We will use several real protocols for illustration but the emphasis will be on various mathematical techniques for algorithm analysis.

Specifically we are interested in four questions:

1. How are congestion control algorithms modelled?
2. Are the models well defined?
3. How are the equilibrium points of a congestion control model characterized?
4. How are the stability of these equilibrium points analyzed?

For each topic, we first present analytical tools, from convex optimization, to control and dynamical systems, Lyapunov and Nyquist stability theorems, and to projection and contraction theorems. We then apply these basic tools to congestion control algorithms and rigorously prove their equilibrium and stability properties. A notable feature of this book is the careful treatment of projected dynamics that introduces discontinuity in our differential equations.

Even though our development is carried out in the context of congestion control, the set of system theoretic tools employed and the process of understanding a physical system, building mathematical models, and analyzing these models for insights have a much wider applicability than to congestion control.

KEYWORDS

communication networks, congestion control, projected dynamics, convex optimization, network utility maximization, Lyapunov stability, passivity theorems, gradient projection algorithm, contraction mapping, Nyquist stability

Contents

Preface

The congestion control mechanism has been responsible for maintaining stability as the Internet scales up by many orders of magnitude in size, speed, traffic volume, coverage, and complexity over the last three decades. Our primary goal is to develop a coherent theory of Internet congestion control from the ground up to help understand and design the equilibrium and stability properties of large-scale networks under end-to-end control.

In addition, we have two broader purposes in mind. First we wish to introduce a set of system theoretic tools and illustrate their application to concrete problems. Second we wish to demonstrate in depth the entire process of understanding a physical system, building mathematical models of the system, analyzing the models, exploring the practical implications of the analysis, and using the insights to improve a design. Even though our development is carried out in the context of congestion control, these basic analytical tools and the research process are much more broadly applicable.

The Internet, called ARPANet at the time, was born in 1969 with four nodes. The Transmission Control Protocol (TCP) was published by Vinton Cert and Robert Kahn in 1974 [14], split into TCP/IP (Transmission Control Protocol/Internet Protocol) in 1978, and deployed as a standard on the ARPANet by 1983. An Internet congestion collapse was detected in October 1986 on a 32-kilobits-per-second (kbps) link between the University of California Berkeley campus and the Lawrence Berkeley National Laboratory that is 400 yards away, during which the throughput dropped by a factor of almost 1,000 to 40 bits-per-second (bps). Two years later Van Jacobson implemented and published the congestion control algorithm in the Tahoe version of TCP [26] based on an idea of Raj Jain, K.K. Ramakrishnan, and Dah-Ming Chiu [27]. Before Tahoe, there were mechanisms in TCP to prevent senders from overwhelming receivers, but no effective mechanism existed to prevent the senders from overwhelming the network. This was not an issue because there were few hosts, until the mid-1980s. By November 1986 the number of hosts was estimated to have grown to 5,089 [1], but most of the backbone links have remained 50–56 bps since the beginning of the ARPANet. Jacobson's scheme adapts sending rates to the congestion level in the network, thus preventing the senders from overwhelming the network.

Jacobson anticipated even in his original paper [26] the network environments in which his algorithm will perform poorly: "... TCP spans a range from 800 Mbps [megabits per second] Cray channels to 1200 bps packet radio links." The algorithm worked very well over a network with relatively low transmission capacity, small delay, and few random packet losses. This was mostly the case in the 1990s, but as the network speed underwent rapid upgrades (see Figure 1), as Internet exploded onto the global scene beyond research and education, and as wireless

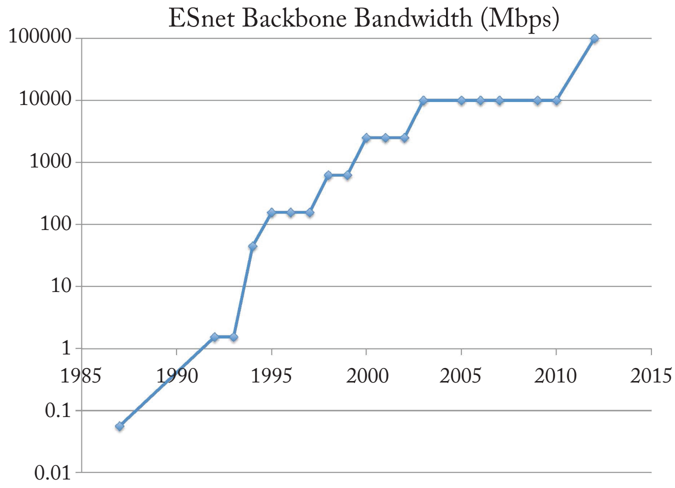

Figure 1: Highest link speed of U.S. Department of Energy's Energy Sciences Network (ESnet) from 1987 (56 kbps) to 2012 (100 Gbps) [2].

infrastructure was integrated with and mobile services proliferated on the Internet, the strain on the original design started to show. This motivated a flurry of research activities on TCP congestion control in the 1990s. A mathematical understanding of Internet congestion control started in the late 1990s with Frank Kelly's work on network utility maximization [28]. An intensive effort ensued and lasted for a decade to develop a theory to reverse engineer existing algorithms and understand structural properties of large-scale networks under end-to-end congestion control, systematically design new algorithms based on analytical insights, and deploy some of these innovations in the field.

This book is a personal account of that effort, focusing on the theory development.

We start in Chapter 1 with a summary of classical Internet congestion control protocols. We explain how to model them as dynamical systems using ordinary differential equations:

$$\begin{aligned} \dot{x} &= f(x(t), q(t)), & \qquad q(t) &= R^T p(t) \\ \dot{p} &= g(y(t), p(t)), & \qquad y(t) &= R\, x(t) \end{aligned}$$

and its variants, where $x(t), q(t) \in \mathbb{R}^N$, $p(t), y(t) \in \mathbb{R}^L$, and $R \in \{0,1\}^{L\times N}$ for a network with N nodes and L links. The graph structure is described by the routing matrix R. The decentralized nature of the system manifests itself in the structure of f and g:

$$\begin{aligned} \dot{x}_i &= f_i(x_i(t), q_i(t)), & \qquad q_i(t) &= \sum_l R_{li}\, p_l(t) \\ \dot{p}_l &= g_l(y_l(t), p_l(t)), & \qquad y_l(t) &= \sum_i R_{li}\, x_i(t) \end{aligned}$$

i.e., each node i (link l) updates its state $x_i(t)$ ($p_l(t)$) based only on local variables $(x_i(t), q_i(t))$ $(y_l(t), p_l(t))$. We prove the existence and uniqueness of solution trajectories to these equations. This ensures that the models are well defined. This class of network models is more general than congestion control and therefore the techniques developed here may be of wider applicability.

We prove in Chapter 2 that the equilibrium point of an arbitrary network under congestion control is the unique optimal solution of a simple convex optimization problem, called network utility maximization. Hence we can interpret congestion control as a distributed algorithm carried out by traffic sources and network resources to maximize utility over the Internet in real time. We explain several implications of this insight.

We present in Chapters 3–5 three different methods to study the global asymptotic stability of the equilibrium point, assuming there is no feedback delay. These methods are based on Lyapunov stability theorems, passivity theorems, gradient descent and contraction mapping theorems. The Lyapunov method is the basic tool for proving stability of general nonlinear systems. The passivity method allows one to analyze the stability of an interconnection of multiple dynamical systems in terms of the passivity of the component systems in open loop. The last method treats congestion control as a gradient algorithm for solving the dual of the network utility maximization.

Finally we describe in Chapter 6 the Nyquist stability method for analyzing local stability around the equilibrium point in the presence of feedback delay.

There is a large amount of literature on congestion control and we have not attempted to provide a survey. Pointers are provided at the end of each chapter only to some papers that are directly related to or extend materials covered in that chapter. We present proofs for some, but not all, of the classical results to illustrate techniques or concepts that we find particularly useful.

Many applications, including congestion control, can be modeled by a system of nonlinear differential equations of the form:

$$\dot{x} \quad = \quad (f(x(t))_{x(t)}^{+}$$

where the projection operation $(\cdot)_{(\cdot)}^{+}$ on the right-hand side ensures that the state variable $x(t)$ remains nonnegative. For example, $x(t)$ may represent the sending rates of traffic sources or the prices of an economy. The projection introduces discontinuity to the vector field, even when f itself is continuous, and complicates analysis. Analytical models often ignore projection even though nonnegative dynamics is prevalent in reality. A notable feature of this book is the careful treatment of the projected dynamics. In particular we include detailed proofs that extend standard results on the existence, uniqueness, equilibrium, and stability properties of smooth unprojected systems to discontinuous projected systems. Some of the stability proofs for congestion control algorithms modeled by projected dynamics in Chapters 3 and 4 are new. As we will see, projection mostly preserves these properties.

Steven H. Low
Pasadena, CA, June 2017

Acknowledgments

The idea of this book started when Jean Walrand of the University of California, Berkeley asked me in early 2004 to write a little book on TCP congestion control for Morgan & Claypool's Synthesis Lectures on Communication Networks, of which he was the inaugural Editor. I agreed but did not start writing until the summer of 2010, when I visited Karl Åström and Anders Rantzer at Lund University in Sweden with my extended family. That was a memorable summer! It was also when my research switched from Internet to power systems, so writing again went onto the backburner after the first draft at Lund. Major revisions were done during the summer of 2015, when I visited Janusz Bialek at Skoltech in Russia to give a short course on analytical methods for Internet and power systems, and during spring 2016, when I visited Jiming Chen, Youxian Sun, and Zaiyue Yang at Zhejiang University in China. I thank the tremendous encouragement and patience of Jean Walrand and the gentle prodding of the publisher Michael Morgan over more than a decade. It's a relief to have paid my debt. I also thank the warm hospitality of my hosts at Lund University, Skoltech, and Zhejiang University.

This book is a product of our FAST project at Caltech from 2000–2007, and I have learned a lot from my collaborators, especially John Doyle, Harvey Newman, and Fernando Paganini, and from the first generation of Netlab members, including Lachlan Andrew, Lijun Chen, Cheng Jin, George Lee, Lun Li, Mortada Mehyar, Christine Ortega, A. Kevin Tang, Jiantao Wang, David Wei, and Bartek Wydrowski. I thank the U.S. National Science Foundation (especially Darleen Fisher), Army Research Office, Air Force Office of Scientific Research, Cisco, and Caltech's Lee Center for Advanced Networking for their generous financial support. Some of us took the effort to deploy our research in the real world through a startup FastSoft. Since 2014, FastTCP has been accelerating more than 1 TB of Internet traffic every second. I experienced first hand the thrill and the challenge in crossing the gap from theory to practice and I thank my colleagues and supporters at FastSoft.

Linqi Guo has worked through the entire draft carefully and corrected numerous errors. I thank him for his meticulous reading and helpful suggestions. Teaching assistants of my networking course (cs/ee 143) at Caltech have contributed some of the exercises, especially Lingwen Gan, Ben Yuan, and Changhong Zhao. Finally, I thank my family, Jenny, Zhi, Zhiyou, my parents, and my sister's family for their unwavering support and trust.

Steven H. Low
Pasadena, CA, June 2017

Notations

We collect some of the notational conventions in this book.

Let $\mathbb{R}^n$, $n \geq 1$, be the set of n-dimensional real vectors, $\mathbb{R}^n_+$ the set of n-dimensional nonnegative real vectors, and $\mathbb{R}^{n \times m}$ the set of $n \times m$ real matrices. If x is a vector or matrix then x^T denotes its transpose. By default a vector x is taken to be a *column* vector and can be specified as either

$$x = \begin{bmatrix} x_1 \\ \vdots \\ x_n \end{bmatrix} \quad \text{or} \quad x = (x_1, \ldots, x_n) \quad \text{or} \quad x = (x_i, i = 1, \ldots, n).$$

Inequalities are taken componentwise, i.e., $x \geq 0$ ($x > 0$) means $x_i \geq 0$ ($x_i > 0$) for $i = 1, \ldots, n$. If $x_i \in \mathbb{R}^{n_i}$, $i = 1, \ldots, k$, are defined then, unless otherwise specified, x denotes the vector $x := (x_i, i = 1, \ldots, k)$ with dimension $n := \sum_i n_i$. Conversely if a vector x is defined then x_i denotes its ith component in $\mathbb{R}^{n_i}$. Similarly for functions $f_i : \mathbb{R}^{k_i} \to \mathbb{R}^{m_i}$, $i = 1, \ldots, n$, and $f := (f_i, i = 1, \ldots, n) : \mathbb{R}^K \to \mathbb{R}^M$ where $K := \sum_i k_i$ and $M := \sum_i m_i$.

For a scalar function $f : \mathbb{R}^n \to \mathbb{R}$, $\frac{\partial f}{\partial x}$ is the row vector and $\nabla f(x)$ is the column vector, both with components $\frac{\partial f}{\partial x_i}$. For a vector function $f : \mathbb{R}^n \to \mathbb{R}^n$, $\frac{\partial f}{\partial x}$ is the $n \times n$ Jacobian matrix defined by

$$\left[\frac{\partial f}{\partial x}\right]_{ij} := \frac{\partial f_i}{\partial x_j}.$$

Given a set of utility functions $U_i(x_i) : \mathbb{R} \to \mathbb{R}$, $i = 1, \ldots, N$, $U_i'(x_i)$ denote their derivatives. We sometimes use $U : \mathbb{R}^N \to \mathbb{R}$ to denote the sum $U(x) := \sum_i U_i(x_i)$. Since U is separable in x_i we use $U'(x)$ to denote the *vector* $U'(x) := (U_i'(x_i),\ i = 1, \ldots, N)$.

For a scalar $a \in \mathbb{R}$, $(a)^+ := \max\{a, 0\}$; for a vector a, $(a)^+$ is defined componentwise, i.e.,

$$(a)^+ := \left([a_i]^+,\ \forall i\right).$$

For scalars $a, b \in \mathbb{R}$

$$(a)_b^+ := \begin{cases} a & \text{if } \ a > 0 \ \text{ or } \ b > 0 \\ 0 & \text{otherwise.} \end{cases}$$

If $a, b \in \mathbb{R}^n$ are vectors of the same dimension then $(a)_b^+$ is defined componentwise, i.e.,

$$\left[(a)_b^+\right]_i := (a_i)_{b_i}^+ \qquad \forall i.$$

We use $\|\cdot\|$ to denote an arbitrary norm and $\|x\|_2 := \sqrt{\sum_i x_i^2}$ the Euclidean norm. $B_\delta(x^*) := \{x \mid \|x - x^*\| \leq \delta\}$ is a *closed* ball around x^* in $\mathbb{R}^n$ unless otherwise specified. $A \subseteq B$ means A is a subset of B and $A \subset B$ means A is a strict subset of B. Given $a_i, i = 1, \ldots, n$, $\mathrm{diag}(a_i, i = 1, \ldots, n)$ denotes the diagonal matrix with a_i as its ith diagonal entry.