Book Reviews

The Mathematics of Internet Congestion Control—R. Srikant (Boston, MA: Birkhäuser, 2004). *Reviewed by Derong Liu*

Computer networks have experienced an explosive growth over the past years and with that growth have come severe congestion problems. For example, it is now common to see Internet gateways drop 10% of the incoming packets because of local buffer overflow [5]. Congestion is a state of sustained network overload, when the demands for resources exceed the supply for an extended period of time. A congestion event will cause a significant number of packets to be lost consecutively. When a user detects a dropped packet, it typically retransmits the dropped packet. Thus, congestion slows down network traffic flow and it tends to get more severe once it starts. Congestion avoidance and control has been a hot research topic in the past two decades and researchers have investigated algorithms that maximize the throughput of the network and minimize the packet-loss rate for networks including the Internet.

To many control researchers and engineers, research in computer networks, especially the Internet, is not very close. The main reason is that research in computer networks typically involves much less mathematical rigor and derivations than what we are used to see in the control literature. This book may help to close this gap. This book is about the application of some well-known results in control theory to the study of Internet congestion control. Readers will find mathematical equations and derivations throughout the book. Results that are familiar to control researchers including convex optimization, Lyapunov stability theory, and the Nyquist criterion are the main tools used in this book.

The origin of the problem studied in this book came from Jacobson [5] where the first Internet congestion control algorithm was published for implementation as part of the transmission control protocol (TCP). Mathematical tools familiar to control engineers are used in this book to understand the dynamics of Jacobson's algorithm and to improve the algorithm.

Chapter 1 provides a brief introduction to the problem of Internet congestion control that is studied in this book. Two simple examples based on the algorithm developed in [2] are used in this chapter to illustrate the problem definition and the mathematics involved in the problem formulation. The Second Method of Lyapunov is used as a tool in arriving at the analysis results of the two examples.

Chapter 2 is concerned with the problem of resource/bandwidth allocation which is viewed as an optimization problem [7]. With the introduction of a general class of utility functions, approaches developed in the literature for convex optimization are applied to solving the present optimization problem for resource allocation. Specific cases of network resource allocation considered include minimum potential delay fairness, proportional fairness, and max–min fairness [6]. Some existing results from convex optimization are introduced in an Appendix at the end of the chapter.

Chapter 3 studies the case of the decentralized resource allocation problem. The resource allocation problem introduced in Chapter 2 which is formulated as a constrained optimization problem is solved in this chapter as an unconstrained optimization problem that incorporates the constraints using a penalty term in the utility function. This chapter provides a solution in the form of a primal algorithm and another in the form of a dual algorithm. Further extensions of these algorithms are also provided in this chapter including multipath routing and multirate multicast congestion control. The Lyapunov stability theory is summarized in an Appendix at the end of the chapter.

Chapter 4 analyzes several protocols currently used in the Internet. In particular, the adaptive window flow control algorithm [5]is analyzed in detail. The algorithm of [5] is discussed in two phases: A slow-start phase and a congestion avoidance phase. The relationship between the buffer size, the length of time that each phase lasts, and the network throughput is derived. Other congestion control algorithms including the TCP-Vegas [1], random early detection [4], explicit congestion notification [3], and high-throughput TCP [8] algorithms are analyzed briefly. All these cases are analyzed under the assumption of no feedback delays.

Chapter 5 studies the stability of congestion control algorithms under the assumption of feedback delays. The single link case is studied in this chapter for various congestion control algorithms introduced in Chapters 3 and 4 and the analyses are done using linear approximation around an equilibrium point. The Nyquist criterion is used in the analysis which is summarized in an Appendix at the end of this chapter.

Chapter 6 extends the results of Chapter 5 to the case of a general network. Again, the stability of various congestion control algorithms are analyzed and the multivariable Nyquist criterion is summarized in an Appendix at the end of the chapter.

Chapter 7 provides analysis results for the global stability of the congestion control algorithm for single link and single flow. The subject of this chapter is open to further study and the chapter only presents analysis results for the proportionally-fair controller. Other controllers in the case of single link and single flow as well as in the case of general network topology are open research problems.

In Chapter 8, stochastic models are used as opposed to the deterministic models of the previous chapters. In particular, the stochastic models of a single link accessed by congestion controlled sources is considered in this chapter. The deterministic equations are derived which are limits of the stochastic models under the condition that the number of users is large. The weak law of large numbers is summarized in an Appendix at the end of the chapter.

Chapter 9 studies the stability analysis at the connection-level with flow-level resource allocation for congestion control. The stability of weighted proportionally-fair controllers and priority resource allocation schemes are analyzed. Some currently open research problems are also mentioned in this chapter.

Chapter 10 considers the case where users cannot be characterized by a concave utility function as in previous chapters. In this case, the users will require an average bandwidth and a quality of service guarantee from the network. These users are not willing to use the network if the amount of bandwidth available is not sufficient and they do not use any excess bandwidth either. Congestion control problems in such network environments are totally different from the ones studied in previous chapters. Analysis results regarding resource sharing, probing, and distributed admission control are obtained.

Finally, Chapter 11 provides some concluding remarks as well as some references for further reading. Some key open problems are also summarized in this concluding chapter.

The book is intended for graduate students and researchers with backgrounds in systems theory and computer networks. Most of the

The reviewer is with Department of Electrical and Computer Engineering, University of Illinois at Chicago, Chicago, IL 60607 USA (e-mail: dliu@ece.uic.edu).

Digital Object Identifier 10.1109/TAC.2004.841398

needed background knowledge is provided as Appendices at the end of several chapters. The book is an important source for graduate students and researchers who conduct research in: 1) design of new protocols for congestion control, and 2) analysis of existing and new protocols for congestion control. In addition, the book is also a good source for control researchers and engineers who look for problems in the field of computer networks. However, researchers in the field of control theory may need to review some basic theory of computer networks in order to understand well the problems studied in this book. The author of this book provides a list of errata which are posted at http://comm.csl.uiuc.edu/~srikant/Papers/errata.pdf.

REFERENCES

[1] L. S. Brakmo and L. L. Peterson, "TCP Vegas: End to end congestion avoidance on a global internet," *IEEE J. Selected Areas Commun.*, vol. 13, no. 8, pp. 1465–1480, Oct. 1995.

[2] D.-M. Chiu and R. Jain, "Analysis of the increase and decrease algorithms for congestion avoidance in computer networks," *Comput. Networks ISDN Syst.*, vol. 17, no. 1, pp. 1–14, Jun. 1989.

[3] S. Floyd, "TCP and explicit congestion notification," *ACM Comput. Commun. Rev.*, vol. 24, no. 5, pp. 10–23, Oct. 1994.

[4] S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance," *IEEE/ACM Trans. Networking*, vol. 1, no. 4, pp. 397–413, Aug. 1993.

[5] V. Jacobson, "Congestion avoidance and control," *ACM Comput. Communication. Rev.*, vol. 18, no. 4, pp. 314–329, Aug. 1988.

[6] J. M. Jaffe, "Bottleneck flow control," *IEEE Trans. Commun.*, vol. COM-29, no. 7, pp. 954–962, Jul. 1981.

[7] F. P. Kelly, "Charging and rate control for elastic traffic," *Eur. Trans. Telecommun.*, vol. 8, pp. 33–37, Jan. 1997.

[8] G. Vinnicombe, "On the stability of networks operating TCP-like congestion control," in *Proc. 15th IFAC World Congr. Automatic Control*, Barcelona, Spain, July 2002.