

ENGINEERING INTERNET QoS

BY SANJAY JHA AND MAHBUB HASSAN,
ARTECH HOUSE, 2002, ISBN
1580533418, HARDCOVER, 350 PAGES

REVIEWER: SRILAL WEERA

Engineering Internet QoS provides a comprehensive source for QoS concepts, architectures, and algorithms. An attractive feature of this book is the broad range of QoS topics covered while not compromising on the depth of discussion of crucial topics. A wealth of useful information is packed into the 300+ pages along with copious amounts of diagrams and tables to improve readability. The breadth of coverage is evident since QoS in all major types of data networks (including mobile wireless) is addressed. Each important topic is discussed in detail and explained by way of lucid examples. For example, the RSVP section consists of 30 pages filled with an extensive discussion of the protocol (not just a simple reproduction of the RFC).

The book is designed as a text for use in a second course on networking. Some background knowledge in data communications is assumed. Since this book is partly based on industry short courses the authors have developed, it is a good self-learning tool for the industry professional as well.

Chapter 1 provides more than an "Introduction to QoS": Starting with perceptual parameters, it builds on the need for network QoS parameters and then discusses the application-specific nature of QoS. An overview of audio-video compression is presented followed by smoothing techniques used in destination buffering. The RTP and RTCP protocols are discussed complete with examples to illustrate how they support adaptive feedback and jitter calculation. Issues in providing QoS for multimedia applications on the best effort Internet are presented, setting the stage for later chapters.

The next two chapters provide a good grounding in the theory and applications of QoS. Traffic shaping, policing, and QoS management sections are well presented. A variety of schedulers including FCFS, priority round-robin, and fair queuing and its variants are explained in detail along with examples. Several advanced schedulers, including current research issues, are also discussed.

Chapter 4 contains an overview of TCP and IP protocols. The authors have included this material to minimize cross-referencing to other books, which the reader will find helpful. The latter part of this chapter deals with queue management techniques used in best effort as well as QoS-capable networks.

The RED packet drop algorithm is succinctly explained. Several queue management algorithms are discussed in detail. The IntServ model is described in Chapter 5. The functioning of an IntServ-capable router is skillfully used to illustrate the architecture. Although references to RFCs are included in the chapter, some discussion of characterization parameters would have been helpful to someone looking to decipher the associated RFCs. Nonetheless, the clear explanations should help the reader gain a thorough understanding of the subject matter. Particularly interesting is the detailed discussion of IntServ mapping of LAN QoS in the second half of the chapter. Clear diagrams are used to illustrate VLAN tagging of Ethernet frames. As with other chapters, research directions are provided at the end of the chapter for the benefit of advanced students.

Chapter 6 makes RSVP and its intricacies simple to understand. This chapter is a good starting point to study the RFCs related to Resource Reservation Protocol. (A draft version of the chapter is available online at <http://www.cse.unsw.edu.au/~sjha/qosbook/> courtesy of the authors.) The chapter covers the details of RSVP design and demonstrates its usefulness in the IntServ environment through examples and simulation. An overview of various RSVP extensions and related research is also provided.

Once the scalability issues of per-flow resource reservation are grasped, the next logical step is to investigate the pros and cons of DiffServ architecture, which is presented in Chapter 7. The PHB and DSCP concepts are deftly presented, and the data path operation performed by a DiffServ router is described in detail. Several pages are devoted to explaining experimental evaluation of premium service using a Linux testbed. Readers interested in setting up an experiment will find this section useful since test scripts and command structures are included.

The Internet is an interconnection of many networks. To provide end-to-end QoS over the interconnections of many autonomous domains, SLAs are required, which brings us to the next chapter, "Policy-Based QoS Management." The concepts of bandwidth broker and resource allocation protocol are aptly presented in the context of a policy-based network implementation. This chapter also describes Internet2 and Qbone architectures. I was pleasantly surprised to see a chapter devoted to ATM. Learning how QoS is supported in ATM networks is essen-

tial to understanding the current issues facing the Internet. Readers will also find the ATM/IP integration section useful, considering the multitude of ATM switches currently in deployment. Chapter 10 provides an easy-to-understand overview of MPLS technology and applications. Signaling protocols and traffic engineering issues are clearly explained. E-LSP and L-LSP concepts are presented, but without referencing the terms (why?). I would have preferred to see more details on network failure recovery methods (i.e., fast reroute and backup LSPs) as well as DiffServ mapping modes, load balancing, and RFC 2547-bis networks, which are of practical importance. However, MPLS is a rapidly growing area, and one has to draw the line somewhere due to space limitations. Nonetheless, this chapter provides a solid foundation. The extensive references supplied will be helpful to the readers keen to explore the subject further.

Mobile wireless is gaining attention lately, and the unique challenges in providing QoS in wireless networks are discussed in Chapter 11. A number of applications and networking technologies of the mobile wireless Internet are presented in an easy-to-understand way. Chapter 12 begins with a detailed discussion of IntServ over DiffServ followed by an overview of QoS routing, VPN and QoS, the content distribution network, as well as billing and charging for QoS. Each provides a list of key references for further reading.

From a pedagogical standpoint the book is an ample resource. Some practical insight on QoS issues in today's Internet access networks (DSL, cable) would be a good addition. More discussion on new Internet-based applications such as VoIP, VPNs, and e-commerce would certainly enrich the book. Again, space limitations may have been an issue. But the good news is that these topics are being added to the Web site.

In summary, this book is very well written and informative. The bibliography is rich and contains many valuable references. Attention to detail is evident throughout the book. For example, under the topic "Deficit Round Robin" is found a thorough discussion on pages 58-60 of the quantum-based algorithm (its pros and cons), a diagram illustrating DRR scheduling, a sample pseudo code for lab testing, as well as useful references. Pointers for further study are provided throughout the book in the sections entitled "Research Directions," a boon for graduate students and researchers alike.