WILEY

高速网络中的QoS控制

（影印版）

〔美〕约翰森·乔　郭晓雷　著

清华大学出版社

高速网络中的 QoS 控制

（影印版）

[美] 约翰森·乔 郭晓雷 著

清华大学出版社
北京

H. Jonathan Chao & Xiaolei Guo
Quality of Service Control in High-Speed Networks

EISBN: 0471003972

图书在版编目(CIP)数据

高速网络中的 QoS 控制=Quality of Service Control in High-Speed Networks/(美)乔（Chao, H.J.），郭晓雷著. —影印本. —北京：清华大学出版社，2004.4
ISBN 7-302-08246-4

I. 高… II. ①乔… ②郭… III. 计算机通信网—服务质量—质量管理—英文 IV. TP915

中国版本图书馆 CIP 数据核字（2004）第 018449 号

出 版 者：清华大学出版社　　**地　　址**：北京清华大学学研大厦
http://www.tup.com.cn　　**邮　　编**：100084
社总机：010-6277 0175　　**客户服务**：010-6277 6969
责任编辑：常晓波
封面设计：项　尧
印 刷 者：清华大学印刷厂
装 订 者：三河市李旗庄少明装订厂
发 行 者：新华书店总店北京发行所
开　　本：185×230　**印 张**：28
版　　次：2004 年 4 月第 1 版　2004 年 4 月第 1 次印刷
书　　号：ISBN 7-302-08246-4/TP・5951
印　　数：1~3000
定　　价：48.00 元

PREFACE

This book addresses the basics, theory, architectures, and technologies to implement quality-of-service (QoS) control in high-speed networks, including asynchronous transfer mode (ATM), Internet Protocol (IP), and multiprotocol label switching (MPLS) networks. For the last few years, we have witnessed an explosion of Internet demands, yielding rapid development of the corresponding technologies in the backbone network, such as terabit routers, dense wavelength division multiplexing (DWDM) transmission equipment, and optical cross-connect switches (OXCs). The success of telecommunications in supporting a wide variety of Internet services, such as multimedia conferencing and video-on-demand, depends greatly on, (in addition to high-speed transmission and switching) the reliable control in the underlying high-speed networks to provide guaranteed QoS.

QoS provision in a network basically concerns the establishment of a network resource sharing policy (including link bandwidth and buffer space) and then the enforcement of that policy. As the line speed increases to 10 or 40 Gbit/s and the number of connections in each line increases to several hundreds of thousands, implementing QoS control, under the constraints of the timing and memory requirement, becomes very challenging. Unlike other books in the area, this book not only gives a comprehensive overview of the existing QoS control techniques, but also describes practical approaches to implement the techniques in the high-speed network.

Most of the book is based on the material that Jonathan has been teaching to the industry and universities for the past decade. He taught a graduate course "Broadband Network" at Polytechnic University, NY, and used the draft of the book as the text. The book has incorporated feedback from both industry people and college students. We believe this book is timely to meet the demand of industry people who are looking for the solutions for meeting various QoS requirements in the high-speed network.

AUDIENCE

This book surveys the latest technical papers that readers can refer to for the most up-to-date development of control strategies in the high-speed network. The readers are assumed to have some knowledge of fundamental networking and telecommunications. Some of this book may require readers to have some knowledge of probability models and college-level mathematics. Since each chapter is self-contained, readers can easily choose the topic of interest for both theoretical and practical aspects. A comprehensive list of references follows each chapter. This book should be useful to software, hardware, and system engineers in networking equipment and network operation. It should be useful as a textbook for students and lecturers in electrical engineering and computer science departments who are interested in high-speed networking.

ORGANIZATION OF THE BOOK

Throughout, IP and ATM networks are used as examples. The book is organized as follows:

- Chapter 1 presents a systematic overview of QoS control methods, including admission control, traffic access control, packet scheduling, buffer management, flow and congestion control, and QoS routing. It also introduces ATM, Internet integrated services (Intserv), Internet differentiated services (Diffserv), and MPLS. It also describes the traffic characterization and QoS parameters.
- Chapter 2 explores admission control and its process of deciding whether a new connection between a source–destination pair can be accepted across a network. The decision is made subject to this new connection's QoS requirement (e.g., loss and delay) without affecting the committed resource provision for the existing connections at each node on its path.
- Chapter 3 focuses on traffic access control and its mechanism of enforcing certain criteria regarding an established connection's traffic load and burstiness, for instance, at the network ingress/entry point of this connection. If the connection crosses two subnetworks under different administration, similar control (also called traffic policing) may also be needed at the egress point of the first subnetwork.

- Chapter 4 presents a historical overview of different packet scheduling algorithms. It also describes under each scheme how to regulate packet transmission orders among a set of connections multiplexed at a network node so that their QoS requirements can be satisfied.
- Chapter 5 is dedicated to discussing the implementation of packet fair queuing, an intensively studied packet scheduling algorithm due to its desirable capability of providing QoS guarantees (e.g., on delay) for connections with diversified QoS requirements. Practical examples covered in this chapter include designs based on sequencers, priority content-addressable memory (PCAM), and random access memory (RAM).
- Chapter 6 presents buffer management, which controls the access of incoming packets to the buffer space and decides which packet should be discarded when, for instance, the buffer is full or a threshold-crossing event happens.
- Chapter 7 explains flow control and congestion control. Flow control addresses the needs of speed matching between a source–destination pair or any two nodes on the path of a connection. Congestion control addresses the regulation of traffic loading across a network for congestion avoidance and recovery.
- Chapter 8 covers QoS routing and describes its process of deciding which path (in a connection-oriented network like ATM) or which next-hop node (in a connectionless network like IP) should be chosen for a new connection among multiple physical paths across a network. The decision may depend on the resource availability along the path and whether the QoS requirement of the new connection can be met.
- Chapter 9 describes basic architecture and conceptual model of Diffserv in details. It describes the network boundary traffic conditioning and per hop behaviors functions used to support Diffserv.
- Chapter 10 covers MPLS technology. It includes the basic concepts, such as the label stack, route selection, penultimate hop popping, and label-switched-path (LSP) tunnel, Label Distribution. It also describes the MPLS mechanism to support Diffserv and two applications of MPLS: traffic engineering and virtual private networks (VPNs).
- The Appendix briefly describes Synchronous Optical Networks (SONET) and ATM for readers who need to attain a basic understanding of the physical layer and link layer protocols of the high-speed network.

ACKNOWLEDGMENTS

We are indebted to all who contributed to this book. Especially, we thank Professor Sugih Jamin (University of Michigan), Dapeng Wu (Carnegie Mellon University), Dr. Jun-Shik Hong, Dr. Li-Sheng Chen (Alcatel), Dr. Necdet Uzun (Cisco), Dr. Yau-Ren Jenq (Fujitsu), Dr. Soung-Yue Liew, Dr. Zhigang Jing (Polytechnic University), Dr. Roberto Rojas-Cessa, and Dr. Taweesak Kijkanjanarat. We thank Dr. Jing and Dr. Liew for contributing Chapters 9 and 10, and Mike Barrett and Dr. Mukesh Taneja for reviewing the chapters.

Jonathan wants to thank his wife, Ammie, and his children, Jessica, Roger, and Joshua, for their love, support, encouragement, patience and perseverance. He also thanks his parents for their encouragement. Xiaolei would also like to thank his wife, Mei Feng, for her love and support.

We have done our best to accurately describe the QoS control methods, technologies, and implementation architectures. If any errors are found, please send email to chao@poly.edu. We will correct them in future editions.

H. JONATHAN CHAO
XIAOLEI GUO

September 2001

CONTENTS

CHAPTER 1

INTRODUCTION

The broadband integrated services digital network (B-ISDN) with standardized asynchronous transfer mode (ATM) is envisaged to support not only current services, but also new services with varying traffic characteristics and quality-of-service (QoS) requirements [1–7, 11]. In simple terms, ATM is a connection-oriented packet switching and multiplexing technique that uses short fixed-size cells to transfer information over a B-ISDN network. The short cell size of ATM at high transmission rates is expected to offer full bandwidth flexibility and provide the basic framework for guaranteeing QoS requirements of applications with a wide range of performance metrics, such as delay and loss.

Meanwhile, the advent of broadband networking technology has dramatically increased the capacity of packet-switched networks from a few megabits per second to hundreds or even thousands of megabits per second. This increased data communication capacity allows new applications such as video conferencing and Internet telephony. These applications have diverse QoS requirements. Some require stringent end-to-end delay bounds; some require a minimal transmission rate; others may simply require high throughput. As use of the Internet diversifies and expands at an exceptional rate, the issue of how to provide necessary QoS for a wide variety of different user applications is also gaining increasing importance [8, 11, 15–20]. This book attempts to clarify the QoS issue and examines the effectiveness of some proposed network solutions.

In short, QoS depends on the statistical nature of traffic. An appropriate service model should be defined and some network QoS control methods should be engineered to meet a range of QoS performance requirements (e.g., throughput, delay, and loss), which are usually represented as a set of QoS parameters associated with the service model. Section 1.1 describes the nature of traffic. Network technologies are presented in Section 1.2. Section 1.3 describes QoS parameters. QoS control methods for traffic management are discussed in Section 1.4. A summary is given in Section 1.5.

1.1 NATURE OF TRAFFIC

There are two main traffic types: delay-sensitive traffic and loss-sensitive traffic. Delay-sensitive traffic is characterized by rate and duration and may need real-time transmission. Examples include video conferencing, telephone, and audio/video on demand, which usually have stringent delay requirements but can accept a certain loss. Loss-sensitive traffic is characterized by the amount of information transmitted. Examples are Web pages, files, and mail. It usually has stringent data loss requirements but no deadline for completing a transmission.

There are other traffic types, such as playback traffic, multicast traffic [e.g., conferences, distributed interactive simulation (DIS), and games], and traffic aggregation [e.g., from local area network (LAN) interconnection]. Observations of LAN traffic [12] reveal its *self-similar*, or *long-range dependent*, behavior: The rate is variable at all time scales; it is not possible to define a duration over which the traffic intensity is approximately constant. These observations have been confirmed repeatedly. A plausible explanation for self-similarity is that LAN traffic results from a superposition of bursts whose duration has a heavy-tailed distribution [7, 12].

1.2 NETWORK TECHNOLOGIES

The network has been evolving to provide QoS guarantees to the users. For instance, ATM, widely adopted in the backbone network, can reserve the bandwidth and buffer for each virtual connection. Similarly, the Internet integrated service (Intserv) can also provide QoS for each flow (see definition in Section 1.2.2) in the Internet Protocol (IP) network. Internet differentiated service (Diffserv) provides different treatment for packets of different classes, instead of on a flow basis, so that it has better scalability than Intserv. Multiprotocol label switching (MPLS) allows the network providers to have better control and provision of QoS through traffic engineering policies.

1.2.1 ATM

The ATM Forum and the International Telecommunication Union (ITU) have different names for some service classes (known to the ITU as "ATM transfer capabilities"), while some classes are standardized by only one body. Hereinafter, unless otherwise stated, we use ATM Forum terminology for illustrations [5, 11, 13, 14]. Interested readers are referred to the Appendix for further details about Synchronous Optical Network (SONET) ATM protocols. The following summarizes the traffic parameters used in the ATM networks. Detailed definition can be found in Chapter 3.

The *constant bit rate* (CBR) service category applies to connections that require cell loss and delay guarantees. The bandwidth resource provided to the connection is always available during the connection lifetime, and the source can send at or below the *peak cell rate* (PCR) or not at all. A CBR connection must specify the parameters, including PCR or peak emission interval ($T = 1/\text{PCR}$), cell delay variation tolerance (CDVT), maximum cell transfer delay, and cell loss ratio. The standard defines the rate in terms of an algorithm making necessary allowances for jitter (cell delay variation) introduced between a terminal and the network interface. The chosen algorithm, called the virtual scheduling algorithm or the continuous state leaky bucket, is now standardized as the generic cell rate algorithm (GCRA).

The *variable bit rate* (VBR) service category is intended for a wide range of connections; it includes real-time constrained connections (rt-VBR) as well as connections that do not need timing constraints (nrt-VBR). (Note that CBR is normally for a real-time service.) The VBR is basically defined by its PCR, sustainable cell rate (SCR), and maximum burst size (MBS). The SCR indicates the upper bound for the mean data rate, and the MBS indicates the number of consecutive cells sent at peak rate.

The *available bit rate* (ABR) standard specifies how users should behave in sending data and resource management (RM) cells in response to network feedback in the form of explicit rate and/or congestion indications. An application using ABR specifies a PCR that it will use and a minimum cell rate (MCR) that it requires. The network allocates resources so that all ABR applications receive at least their MCR capacity. Any unused capacity is then shared in a fair and controlled fashion among all ABR sources. The ABR mechanism uses explicit feedback to sources to ensure that capacity is fairly allocated.

At any given time, a certain amount of the capacity of an ATM network is consumed in carrying CBR and the two types of VBR traffic. Additional capacity may be available for one or both of the following reasons: (1) not all of the total resources have been committed to CBR and VBR traffic, and (2) the bursty nature of VBR traffic means that sometimes less than the committed capacity is used. Any capacity not used by ABR sources remains available for *unspecified bit rate* (UBR) traffic as explained below.

The UBR service is suitable for applications that can tolerate variable delays and some cell losses, which is typically true of *transport control protocol* (TCP) traffic. With UBR, cells are forwarded on a first-in, first-out (FIFO) basis, using the capacity not consumed by other services; delays and variable losses are possible. No initial commitment is made to a UBR source, and no feedback concerning congestion is provided; this is referred to as a *best-effort* service.