Pascal Lorenz Abbas Jamalipour Denis A. Khotimsky

IP-ORIENTED QUALITY OF SERVICE

The best-effort service for Internet delivery of data cannot be used for new multimedia applications. New technologies and new standards are necessary to offer quality of service (QoS) for these multimedia applications. Therefore, new communication architectures integrate mechanisms allowing guarantees of specific quality to services as well as high data rates for communication systems. This feature topic in *IEEE Communications Magazine* seeks to survey and present the research and engineering work currently ongoing in the field of QoS.

Quality of service must be seen as an end-to-end process. Assume, for example, that you want to establish a video conferencing over the Internet using the UMTS (the third-generation Universal Mobile Telecommunication System) as the access network for your communications. At the two very endpoints of this communications process we have the end-user terminals (e.g., a cellular phone at one end and a desktop computer connected to the wired Internet at the other). The access technology here thus comprises several systems: the local bearer service providing the service to the cellular phone user, the UMTS bearer service, and the external bearer service providing service to the desktop user.

Without the support of the required QoS indicators (e.g., delay and bandwidth in our current example) by all segments of the network from end to end, we cannot claim we have QoS support. UMTS has its own share in providing QoS, but the endpoint bearer services also need to support similar QoS indicators in order to complete the end-to-end process. Although it is possible to provide QoS with different ways of support by the individual segments in the network, it will be much more efficient and reliable to provide QoS with close interrelation between the individual segments.

The Internet Engineering Task Force (IETF) started working on providing the QoS in IP networks in the mid 1990s. Two different approaches have been introduced: integrated services (IntServ) in 1994 and differentiated services (DiffServ) in 1998. IntServ was introduced in IP networks in order to provide guaranteed and controlled services in addition to the already available best-effort service. It is an extension to the Internet architecture to support both non-real-time and real-time applications over IP. Each traffic flow in this service can be classified under one of three service classes: guaranteed, controlled load, and best effort.

Guaranteed service provides delay-bounded service agreements for voice and other real-time applications requiring severe delay constraints. Controlled load service provides a form of statistical delay service agreement (e.g., with a nominal mean delay). Finally, best-effort services are included to match current IP service, mainly for interactive burst traffic (e.g., Web), interactive bulk traffic (e.g., FTP), and background or asynchronous traffic (e.g., email).

IntServ and reservation protocols such as RSVP have failed to become an actual end-to-end QoS solution, mostly because of the scaling problems in large networks and because of the need to implement RSVP in all network elements from the source all the way to the destination.

DiffServ came to remedy the disadvantages of IntServ in providing QoS in IP networks. DiffServ aims to provide simple, scalable, and flexible service differentiation using a hierarchical model. That is, resource management is now divided into two areas: interdomain and intradomain.

In DiffServ all of the customer's local network requirements for QoS are aggregated, and then a service level agreement (SLA) is made with the network service provider. The SLA may be static (negotiated and agreed on a long-term basis, e.g., monthly) or dynamic, changing more frequently. The local network is then responsible for providing differentiated services to end users within the network. This is usually done through marking packets with specific flags used in the type of service field of IPv4 or the traffic class field of IPv6.

The cellular wireless system approaches to providing QoS are completely different from their Internet counterparts. General Packet Radio Service (GPRS) and UMTS, for example, use subscriber QoS profiles and traffic classifications, respectively, to manage QoS in their wireless systems.

At this time, we see little relation between the Internet and cellular approaches to providing QoS. They have aimed at QoS by their own approaches without much attention to each other. Providing end-to-end QoS with such a configuration would be a very difficult task, if not impossible, and therefore it will be a long time before we can see end-to-end QoS support for wireless Internet. Harmonization between the two approaches so that the service of one technology can cooperate and complement the service of the other remains the main issue toward QoS establishment for future networks.

This feature topic tries to gather related information on the topic of QoS in IP networks; the guest editors have tried to come up with a collection of tutorial papers among all the submissions to this issue so that a balanced view of the activities in different parts of the world and in academia and industry can be accommodated within our limited space.

An open call for papers for the present feature topic received an overwhelming response. A total of 58 papers from all around the world were submitted to the guest editors. These papers covered the entire spectrum of QoS concepts, algorithms, and architectures in wired, wireless, and optical environments, targeting the core as well as access network segments. The diversity of network assumptions, along with the sheer submission volume, indicates that QoS remains an increasingly important issue in modern networking.

Selecting only five papers out of the large pool of manuscripts was a formidable task. Each submission was referred to three independent peer reviewers with strong international reputations in the respective field. The comments provided by them were used as the basis for final selection. Regrettably, with an acceptance ratio of nearly 1:12, we had no choice but to leave out many high-quality submissions.

We would like to thank all the authors who responded to the call for papers, regardless of whether their paper has been included in this issue or had to be rejected due to space limitations. In the latter case we hope that their work will find its intended audience through another magazine, journal, or conference proceedings, and that the feedback provided by the peer reviews will help this happen sooner and more efficiently.

We would also like to express our sincere thanks to all the reviewers who did an excellent job, trying not only to justify their recommendations to the editors, but also to provide detailed and helpful suggestions to the authors on improving the content and presentation style. They bore on their shoulders the unexpected load of high submission volume. Their expertise and their time helped this issue happen.

The present feature issue is composed of five articles of exceptional quality that represent various viewpoints on QoS in the broad IP-oriented environment and successfully passed all stages of the selection process.

The first article, "QoS Management in Trunk-and-Branch Switched Ethernet Networks," by K. Rege, S. Dravida, S. Nanda, S. Narayan, J. Strombosky, and M. Tandon, focuses on QoS management in the next-generation broadband access network (NGBAN) based on multihop switched Ethernet network architecture. The authors claim that this network architecture and the proposed QoS management framework are generally applicable and will form the basis for the last mile independent of the physical layer. The proposed QoS management utilizes distributed policing, scheduling, and flow control, and provides a scalable and flexible architecture.

In the second article, "Quality-of-Service Mechanisms in IP-over-WDM Networks," A. Kaheel, T. Khattab, A. Mohamed, and H. Alnuweiri provide a discussion on the difference between QoS provisioning in IP networks and in optical networks. The topic of IP over WDM is discussed with consideration of the physical layer characteristics and the limitations of the optical domain. Three major optical switching methods (wavelength routing, optical packet switching, and optical burst switching) are considered for the different QoS proposals surveyed in the article.

In the third article, "QoS in Best-Effort Networks," B. Wydrowski and M. Zukerman point out that the problems faced by today's best-effort network are not fundamental ones, and discuss how the current research into congestion control has uncovered that best-effort networks can deliver the performance required for a range of applications, with a few modifications. They also discuss that best-effort networks lead to a pricing mechanism that charges users in accordance to the stress they place on the network.

The fourth article, "Overview of Constraint-Based Path Selection Algorithms for QoS Routing," by F. A. Kuipers, T. Korkmaz, M. Krunz, and P. Van Mieghem, addresses QoS routing in a QoS architecture. A high-level overview of the main solutions available in the literature for constraint-based path selection is provided in the article with an emphasis on the restricted shortest path and multiconstraint path problems.

The fifth article, "A Practical Approach for Providing QoS in the Internet Backbone," by X. Xiao, T. Telkamp, V. Fineberg, C. Chen, and L. M. Ni, proposes a practical approach to QoS provisioning that considers not only technical but also economic factors. The Internet architecture and Internet service provider billing models are outlined in the article. The proposed approach makes use of good network design, DiffServ, traffic protection, traffic engineering, and traffic management techniques.

Finally, as the guest editors of this feature topic, we hope the readers find it interesting and consider it a useful guide in research and development activities toward QoS provisioning in next-generation IP networks and the new telecommunications era.

BIOGRAPHIES

PASCAL LORENZ [SM'00] (lorenz@ieee.org) received his Ph.D. from the University of Nancy, France. Between 1990 and 1995 he was a research engineer at WorldFIP Europe and Alcatel-Alsthom. He is a professor at the University of Haute-Alsace and responsible for the Network and Telecommunication Research Group. His research interests include QoS, wireless networks, and high-speed networks. He was Program and Organizing Chair of IEEE ICATM '98, ICATM '99, ECUMN 2000, ICN '01, ECUMN '02. Since 2000 he is a Technical Editor of *IEEE Communications Magazine*. He is secretary of IEEE ComSoc Communications Systems Integration and Modeling Technical Committee. He is a member of many international committees and has served as a guest editor for a number of special issues, including *Telecommunication System*, *IEEE Communications Magazine*, and *LNCS*. He is a member of many conference technical program committees, and has organized and chaired several technical sessions. He has given tutorials in major international conferences. He is the author of two books and 90 international publications in journals and conferences.

ABBAS JAMALIPOUR [S'86, M'91, SM'00] (a.jamalipour@ieee.org) has been with the School of Electrical and Information Engineering at the University of Sydney, Australia, since 1998, where he is responsible for teaching and research in wireless data communication networks and mobility, traffic, and QoS management in wireless IP networks and satellite systems. He holds a Ph.D. degree in electrical engineering from Nagoya University, Japan. He is author of the book *The Wireless Mobile Internet — Architectures, Protocols and Services* (Wiley, 2002); author of the first technical book written about the networking aspects of LEO satellites, *Low Earth Orbital Satellites for Personal Communication Networks* (Artech House, 1998); author of a chapter in the prestigious *Wiley Encyclopedia of Telecommunications and Signal Processing*, (John Proakis, Ed., 2002); and author of a chapter in *Next Generation Wireless Networks* (Sirin Tekinay, Ed., Kluwer, 2001). He has authored over 80 papers in major journals and international conferences, and given short courses and tutorials at major international conferences. He has served on several major conference technical program committees, and organized and chaired many technical sessions and panels at international conferences including a symposium at IEEE GLOBECOM 2001. He is Vice Chair of the Satellite and Space Communications Committee and the Asia Pacific Board, Coordinating Committee Chapter, IEEE ComSoc. He has organized several special issues on the topic of 3G and beyond systems as well as broadband wireless networks in IEEE magazines and journals. He is a technical editor of *IEEE Wireless Communications*.

DENIS A. KHOTIMSKY (dkhotimsky@lucent.com) received his Dipl.Eng. degree in computer engineering from Moscow Aviation Institute in 1987, and a Ph.D. degree in computer science from the University of California, Santa Barbara, in 1996. In 1986–1992 he worked for Moscow's Institute for Automated Systems (IAS) and Sprint Networks-Russia. Since 1996 he has been a member of technical staff at Bell Laboratories, Lucent Technologies, where he has been involved in architecture development of the Lucent ATLANTA™ chipset (presently a product of Agere Systems) and is currently focusing on MPLS-centric switching systems with advanced QoS capabilities. He holds two U.S. patents with several other applications pending and is a recipient of Bell Labs President's Gold Award for the development of the ATLANTA ATM port controller. He regularly contributes as an author, reviewer, and/or TPC member to professional journals and international conferences, including IEEE INFOCOM, GLOBECOM, and ICC.