SCALABILITY IN IP-ORIENTED NETWORKS

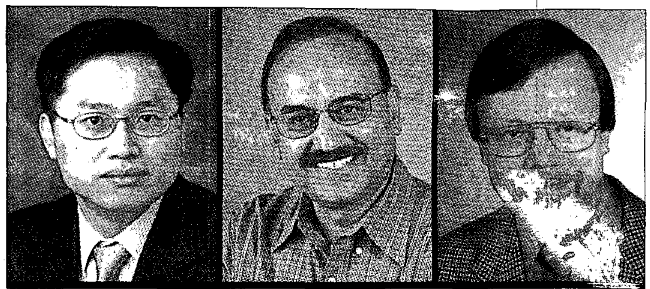

Nirwan Ansari *Amitabh Mishra* *Heinrich J. Stuettgen*

With the exponential growth of the Internet, the number of nodes in the Internet is exploding, and thus scalability is becoming an ever increasingly critical concern in this evolving environment. For such large networks with huge numbers of devices and nodes, it is virtually impossible to maintain the complete state of the network and control what is happening in the network very tightly. For a network system, scalability depends on the architecture of the network node hardware, application software, underlying operating systems, interconnection topology, and communication protocols. *Due to the heterogeneity of network equipment, network architectures, network-based services, and service providers, the scalability challenges have become even more relevant.* Since large networks such as the Internet have some fundamental scalability limitations when it comes to managing individual traffic flows, the original simplicity of the Internet architecture was based on the idea of soft states, implying that network nodes would not go through elaborate protocols to keep all states updated and current at all times, but rather refresh states periodically. Based on these basic ideas, researchers started to develop ideas for soft quality of service (QoS) in the mid-'90s. The key concepts here are traffic aggregation and relative prioritization rather than per-state fine-grained control.

The articles in this feature topic address a number of critical and relevant issues studied within industry and academia. The area where scalability concerns are most pressing is the design and implementation of QoS control in the real world. When we talk about the real world, our attention is drawn toward the scalability of switching systems (switches and routers) that are integral parts of the access and core networks of the Internet. The performance scalability of these systems is also affected by the software that is responsible for call processing/packet processing, administrative and management support, and data handling. We expect that the scalability of networked software systems to be an area of intense research during the next few years.

The introductory article, written by M. Welzl, describes a trade-off between the granularity of services and the scalability of QoS provisioning represented by IntServ and DiffServ. After a brief review of the trade-off between scalability and QoS, the article describes three approaches — combining IntServ and DiffServ, dynamic packet state, and congestion control and QoS — to alleviate this strict relationship. Good results can be achieved with each of these three approaches, thus leading to the conclusion that trading QoS for scalability is not necessarily the only option.

Following this introduction, N. Christin and J. Liebeherr present a new service architecture that can support stronger per-class service guarantees than "classic" DiffServ can. The main principle of this architecture is to use an adaptive service rate allocation to traffic classes, realized by an algorithm combining scheduling and buffer management. Since it can be implemented with low complexity and requires maintaining little state information, it possesses desirable scalability.

The article by B. Zhang and H. Mouftah first presents a survey of existing work proposed to address forwarding state scalability for providing multicast in the IP networks. Then the authors extend an existing multicast routing protocol, Multicast Extension to Open Shortest Path First (MOSPF), by introducing tunnel support to improve scalability in terms of forwarding state and computational overhead. Finally, simulation results demonstrate that the extension can significantly reduce forwarding state and overhead without adverse effects or extra control overhead.

"Proportional Service Differentiation: A Scalable QoS Approach," by Y. Chen, M. Hamdi, D. Tsang, and C. Qiao, presents an overview about recent research work on proportional service differentiation, which possesses desirable properties in terms of controllability, consistency, and scalability. This article discusses the various QoS metrics and their corresponding implementation complexity. A uniform and scalable solution with absolute QoS bounds for delay and loss rate by employing the proportional model is also proposed.

The article by S. Sengupta, D. Saha, and V. Kumar compares IP-over-WDM and IP-over-OTN architectures from an economic standpoint using real-life network data. It shows how IP-over-OTN architectures can lead to substantial reduction in capital expenditure. The authors also compare the two architectures from the perspective of scalability, flexibility, and robustness.

The final article, "Scalable Design of a Policy-Based Management System and Its Performance," written by K. Law and A. Saxena, studies the drawbacks of the standard Policy-Based Management (PBM) system and proposes the Unified Policy-Based Management (UPM) system. The authors compare the performance of PBM and UPM through extensive experiments. With the prototype implementation, it is shown that the proposed UPM is a high-performance scalable design with some desirable properties.

We hope this feature topic in *IEEE Communications Magazine* will help readers gain better understanding of some key scalability issues in IP communication networks. Due to the space limitation, only a fraction of the submitted papers are included in this issue. We hope other topics on scalability issues will be covered in the near future in *IEEE Communications Magazine*. We are grateful to the authors who invested a lot of work in presenting important insights into the scalability challenge. In particular, we thank the reviewers for their helpful comments, which contributed greatly to the overall quality of this feature topic. We would also like to acknowledge the assistance provided by the Editor-in-Chief and publications staff.

BIOGRAPHIES

NIRWAN ANSARI (Nirwan.Ansari@njit.edu) received B.S.E.E. (summa cum laude), M.S.E.E., and Ph.D. from the New Jersey Institute of Technology (NJIT), University of Michigan, and Purdue University in 1982, 1983, and 1988, respectively. He joined the Department of Electrical and Computer Engineering, NJIT, in 1988, and has been a full professor since 1997. He is a technical editor of *IEEE Communications Magazine* as well as the *Journal of Computing and Information Technology*, and is General Chair of ITRE 2003. He was instrumental, while serving as its Chapter Chair, in rejuvenating the North Jersey Chapter of the IEEE Communications Society, which received the 1996 Chapter of the Year Award. He served as Chair of the IEEE North Jersey Section and on the IEEE Region 1 Board of Governors during 2001–2002, and currently serves on various IEEE committees. He was the recipient of the 1998 NJIT Excellence Teaching Award in Graduate Instruction and a 1999 IEEE Region 1 Award. His current research focuses on various aspects of high-speed networks and multimedia communications. He co-authored with E.S.H. Hou *Computational Intelligence for Optimization* (1997, translated into Chinese in 2000), and edited with B. Yuhas *Neural Networks in Telecommunications* (1994), both published by Kluwer Academic.

AMITABH MISHRA [SM] is an associate professor in the Bradley Department of Electrical and Computer Engineering at Virginia Tech, which he joined in August 2000. At Virginia Tech his main thrust of research is in the area of architecture and performance of computer-communication networks. Currently he is working on the architectures of 4G/3G wireless networks (UMTS and CDMA2000), scalability issues in communication networks, routing in mobile ad hoc networks, protocols for wireless sensor networks, QoS issues in packet data networks, network-processor-based next-generation switches and routers, and fault tolerance. From 1987 to 2000 he was a member of technical staff with Lucent Technologies, Bell Laboratories, Naperville, Illinois, where his focus was on application architecture and performance. POTS, IN/AIN, ISDN, ATM, GPRS, CDMA2000, and UMTS were the major areas he worked on while with Bell Laboratories. He received his B. Eng. and M. Tech. degrees in electrical engineering from Government Engineering College, Jabalpur, India, and Indian Institute of Technology, Kharagpur in 1973 and 1975, respectively. He obtained a M. Eng. in 1982 and a Ph.D. in 1985, also in electrical engineering from McGill Universitym Montreal, and an M.S. in computer science in 1996 from the University of Illinois at Urbana-Champaign. He is a chair of communications software, chief information officer of system integration and modeling, and secretary of Internet technical committees of IEEE Communications Society.

HEINRICH J. STUETTGEN [SM] was a Fulbright student at the State University of New York at Buffalo where he obtained an M.Sc. in 1979. In 1984 he graduated as Doctor of Science from Dortmund University. He then joined the IBM R&D Laboratory at Boeblingen, Germany. In 1987 he moved to IBM's European Networking Center at Heidelberg and became manager of the Broadband Multimedia Communication Department. Since 1997 he leads NEC's Network Laboratories at Heidelberg, where he is currently general manager of the laboratories. Since 2002 he is Vice Chair of ComSoc's Technical Committee on Switching and Routing. In 2000 he was program and organizing chair of the IEEE High Performance Switching and Routing Conference at Heidelberg. Currently his main technical interest is in communication protocols for mobile and multimedia Internet communication.