Guest Editorial Network Support for Multicast Communications

MANY networked applications rely on multicast communications for efficient operation. These applications include video conferencing, video distribution, computer supported collaborative work (CSCW), distributed games, and wide-scale information dissemination. These applications require a variety of services from the network beyond the provision of efficient routes between participants. Multicast applications require mechanisms for fair bandwidth sharing since they share the network with other multicast and unicast flows. Many applications also require the reliable transfer of data; other applications, such as video and audio applications, require low-delay or low-jitter services. Confidential data and trusted services require authentication and key management protocols, an increasingly important area of research. The papers in this special issue focus on many of those services, and their collection provides a snapshot of current research in the topics covered.

It is interesting to contrast the topics covered by this issue with those covered five years ago in the previous IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATION special issue on multicast communications (April 1997). A significant fraction of the papers published five years ago focused on intradomain multicast routing. There was also a large emphasis on protocols for providing reliable multicast based on retransmissions. Last, there was considerable emphasis on asynchronous transfer mode (ATM). Notably absent from that special issue were the topics of application-level multicast and security.

The current issue, on the other hand, places a heavy emphasis on various aspects of application-level multicast and on security. In addition, there has been a shift from retransmission-oriented reliable multicast to scalable reliable multicast and to fair congestion control schemes.

Several developments in the last few years partly explain these differences. The most important of these is the slow pace of network-level multicast deployment, which has prompted an increased interest on application-level multicast. This appears primarily due to the complexity that the original Internet protocol (IP) multicast service model imposes on intra- and interdomain routing mechanisms. This has resulted in complex and sometimes faulty routing code, ad hoc solutions for routing at the interdomain level, and a reluctance from network administrators to activate multicast. Another contributing factor to this reluctance has been the lack of access controls within the original IP multicast model. The recently proposed Source-Specific Multicast model, proposed in reaction to these problems, may prove to have a positive impact on the acceptance and deployment of native multicast.

Reliable multicast, which was a very important research topic five years ago, has diminished in importance as the problem of how to provide retransmission-based reliability has been pretty well solved. Congestion control continues to be an important research area, although during this time interesting proposals have emerged on implementing transmission control protocol (TCP)-friendly single-rate and multirate multicast congestion control.

The consequences of these developments during the past five years are numerous and significant.

First, the lack of universal, network-layer-based multicast deployment has stimulated substantial research and development of application-level infrastructures for providing multicast services. The first four papers in this issue cover various aspects of the problem. Shi and Turner present algorithms for generating multicast routes and for bandwidth dimensioning in an overlay network. Rao *et al.* present the Narada protocol that permits end systems to organize into an overlay network, where end systems continually optimize this structure by adapting to network dynamics. Liebeherr *et al.* approach the problem by setting up a multicast overlay network using Delaunay triangulation, allowing end systems to compute routes without a routing protcol. Finally, Castro *et al.* present Scribe, a scalable multicast infrastructure built on top of Pastry, a generic peer-to-peer object location and routing infrastructure.

Second, multicast security has become an active research area during the last five years, and we present two papers discussing scalable key management infrastructures for multicast applications. The paper by Chan and Chan focuses on the bandwidth requirements of the key servers. It presents a model of a distributed server architecture and determines the optimal configuration. The paper by Banerjee and Bhattacharjee presents IP-multicast oriented protocols and algorithms that produce hierarchical key management structures with constant processing and message overhead during rekey operations.

Third, as we mentioned, there has been significant progress in the research on reliable multicast and congestion control, and the next group of papers focuses on these topics. Byers *et al.* present a dissemination technique, the digital fountain, based on the heavy application of forward erasure correction codes to allow concurrent users to obtain content at any time and at different rates. Kar *et al.* develop efficient multirate congestion control algorithms that optimize a sum of receiver utilities. This optimization-based approach provides the means to ensure a wide variety of fairness objectives, including TCP fairness.

Digital Object Identifier 10.1109/JSAC.2002.803064.

Byers *et al.* present a receiver-driven congestion control algorithm that is appropriate for layered multicast. This protocol is fair to TCP connections, and is designed to mitigate the delays incurred through the use of Internet group management protocol (IGMP).

The next two papers focus on problems related to optical networks and SONET rings. Ali focuses on the problem of placing multicast nodes in wavelength-routed networks. A simple model for predicting blocking probability is derived and used to guide an iterative algorithm for the placement of multicast nodes. The second paper by Feng *et al.* focuses on the problem of implementing video conferencing services on SONET/ATM. They conclude that a new multicast virtual path (VP), multidrop VP coupled with an appropriate assignment scheme can provide very good performance.

Finally, the last paper focuses on quality-of-service (QoS). Specifically, Francini and Chiussi present an approach for providing rate guarantees to both unicast and multicast flows in multistage switches provided that they satisfy negotiated rate profiles. The approach relies on regulating access to the switch fabric and dropping copies of packets that violate the negotiated profile.

ACKNOWLEDGMENT

The editors would like to thank the authors of all the papers submitted to this issue. The wealth of research in the area of multicast communications was evident from the large number of submissions that we received. The editors would also like to acknowledge the contribution of the many experts who participated in the reviewing process of submitted manuscripts. The quality of this issue is due in large part to the constructive suggestions that came out of the review process.

They would also like to thank Dr. W. Bux, JSAC Board Representative, for his guidance.

DON TOWSLEY, *Guest Editor*
University of Massachusetts
Department of Computer Science
Amherst, MA 01003–4610 USA

CHRISTOPHE DIOT, *Guest Editor*
Sprint ATL
Burlingame, CA 94010 USA

BRIAN NEIL LEVINE, *Guest Editor*
University of Massachusetts
Department of Computer Science
Amherst, MA 01003–4610 USA

LUIGI RIZZO, *Guest Editor*
Università di Pisa
Dip. Ing. dell'Informazione
Pisa, 56126 Italy

W. BUX, *J-SAC Board Representative*

Don Towsley (M'78–SM'93–F'95) received the B.A. degree in physics and the Ph.D. degree in computer science from the University of Texas, Austin, in 1971 and 1975, respectively.

From 1976 to 1985, he was a Member of the faculty of the Department of Electrical and Computer Engineering at the University of Massachusetts, Amherst. He is currently a Distinguished Professor at the University of Massachusetts in the Department of Computer Science. He has held visiting positions at IBM T.J. Watson Research Center, Yorktown Heights, NY (1982–1983); Laboratoire MASI, Paris, France (1989–1990); INRIA, Sophia-Antipolis, France (1996), and AT&T Labs–Research, Florham Park, NJ (1997). His research interests include networks, multimedia systems, and performance evaluation. He was Program Co-Chair of the joint ACM SIGMETRICS and PERFORMANCE '92 conference. He has organized several workshops on networking topics and has served on numerous program committees.

Dr. Towsley received the 1998 IEEE Communications Society William Bennett Paper Award and three best conference paper awards from ACM SIGMETRICS. He has been elected Fellow of the ACM. He is an Associate Editor of the *Journal of the ACM* and has previously served on several editorial boards including those of the IEEE TRANSACTIONS ON COMMUNICATIONS and IEEE/ACM TRANSACTIONS ON NETWORKING. He is a Member of ACM and ORSA, and Chair of IFIP Working Group 7.3 on Performance Modeling and Analysis.

Christophe Diot received the Ph.D. degree in computer science in 1991 from INP Grenoble, Grenoble, France.

From 1993 to 1998, he was a Research Scientist at INRIA, Sophia-Antipolis, France, working on new Internet architecture and protocols. He moved to Sprint Advanced Technology Laboratory, Burlingame, CA, in October 1998 to take the lead of the IP research group. His current interest is in the passive monitoring of the Sprint IP backbone in order to study IP traffic characteristics and to design new analytical models and traffic engineering solutions for pure packet networks.

Dr. Diot is a Member of ACM, and serves as an editor for ACM/IEEE TRANSACTIONS ON NETWORKING.

Brian Neil Levine received the Ph.D. degree in computer engineering in 1999 from the University of California, Santa Cruz. His dissertation focused on reliable multicast and multicast routing protocols.

He is currently an Assistant Professor at the University of Massachusetts, Amherst, in the Department of Computer Science. His research interests are in the areas of computer networks, network security, and group communication. His most recent interests are in IP anonymity, peer-to-peer networking, and wireless networking.

Dr. Levine is the recipient of an NSF CAREER award and the Outstanding Paper award from the 2002 ISOC Symposium on Network and Distributed Systems Security. In 2002, he is Co-Chair of the International Workshop on Networked Group Communication and the Guest Co-Editor of the special issue of *Computer Communications* on Network Security. In addition, he has served on numerous program committees for conferences and workshops including NGC, Infocom, and ICNP.

Luigi Rizzo received the Ph.D. degree in electronic engineering in 1993 from the SSSUP S.Anna, Pisa, Italy.

He is currently an Associate Professor at the Dipartimento di Ingegneria dell'Informazione, Università di Pisa, Italy, where he has been since 1991. His research is mainly in the area of computer networks, for the past 4–5 years in multicast and multicast congestion control. Some of this work has been done with colleagues in the RMRG, and funded by Cisco and Microsoft Research. In the fall of 2000 and 2001, he was a Visiting Researcher at ICSI, Berkeley. As a side effect of his research, he ended up writing a few pieces of software, related (and sometimes not) to research and teaching, some of which (e.g., dummynet) have become widely used, and are now part of FreeBSD and have been included in other popular operating systems.

Dr. Rizzo is a Member of COST 264, for which he organized the workshop NGC'99, and has served on the Program Committee for a number of conferences and workshops, including SIGCOMM, NGC, ICNP, PfHSN, Networking 2002.