Guest Editorial
Network Support for Multipoint Communication

I. A Historical Perspective

MULTIPOINT communication is one of the oldest forms of communication among humans. It has long been recognized that communicating a *message* (in the broad sense of the word) to multiple recipients simultaneously is a very efficient method of getting the message across. Whether the message was delivered in the form of a smoke signal, political speech, religious sermon, town-hall meeting, or a classroom lecture, the *scalability* of multipoint communication was apparent. Because of the inefficiencies involved, humans tend to use one-on-one communication mostly when necessitated by privacy or other social concerns.

In the early days of electronic communication both one-on-one and multipoint communication were possible. Telephone and telegraph technology allowed *long-distance and local* point-to-point communication. Radio-frequency radio and television technology allowed *only local* multipoint communication. It was not until the deployment of communication satellites that long-distance multipoint communication became feasible but only (until very recently) as a backbone technology. This multipoint communication was mostly simplex in nature with content being disseminated from a source to multiple receivers with little or no signaling being returned by the receivers. A notable exception was the conferencing capability eventually made part of the telephone network, but that was cumbersome to set up and not scalable.

In this context it is not surprising that early research into data networks and packet switching technology focused on point-to-point communication. After all, this was a technology to provide long-distance communication, and attitudes were clearly shaped by the successes of the telephone network in carrying point-to-point conversations. Even with the development of local area network technology that could easily provide multipoint communication, the focus was on how to provide point-to-point communication with multipoint communication supported almost as an afterthought.[1]

Early on, the need to build multipoint applications was apparent. With network layer support for such applications lacking in the early packet-switched networks, many multipoint applications were developed assuming only point-to-point network support. This legacy continues today even though the inefficiencies are obvious. Notable examples of early multipoint applications include e-mail (using mailing lists), software/file distribution including file replication on mirrored sites, and Internet news distribution. Research in reliable distributed systems has also led that community to realize the need for multipoint communications. Much of the early developments (e.g., Cornell University's ISIS system) built multipoint tools or applications (or "middleware" as it would be called today) using point-to-point network support.

Our desire to carry digital audio and video over packet-switched networks is part of our desire to communicate in forms more natural and expressive than printed text. Again it is no surprise that we gravitate toward the more scalable multipoint communication form. For some time, this formed the primary impetus behind the provision of multipoint support in data networks. With the large bandwidths required by video and high fidelity audio, it was no longer feasible to accept the inefficiencies associated with providing multipoint services using point-to-point support from the network. Also, live audio and video demanded some loose notion of near-simultaneous delivery which is facilitated if the network provides built-in multipoint support. In a sense, we have come full circle with audio and video (a.k.a. radio and television) driving our communication needs, with a strong desire for scalability and thus the need for multipoint communication. The emphasis of the research community on multipoint support for the delivery of audio and video, however, delayed the development of reliable multipoint communication because certain types of errors and losses are tolerable with these media (depending on the specifics of the encoding scheme).

With at least 10–20 years of research into wide-scale network support for multipoint communication behind us, this special issue comes at a crossroads in the networking community's research agenda. We are now fully cognizant of the importance of multipoint communication and the need to efficiently support it within the network. We also are beginning to come to grips with the many applications that could potentially use this support and their sometimes distinct requirements. Examples of such services include video distribution, video conferencing, computer supported collaborative work (CSCW), wide-scale information dissemination, and support for distributed (super)computing. It is also safe to assume that among the as-of-yet unforeseen uses of networks will be several that will require multipoint communication support. The interest in multipoint communication spans Internet-style networking, B-ISDN technology including ATM, as well as application and service developers and providers.

A Note on Terminology: We use the term "multipoint" in this editorial (and in the title of the issue) to refer in general to all forms of communication with multiple participants, regardless of semantics. The term "multicast" has been used almost synonymously, even though it is often used to denote point-

Publisher Item Identifier S 0733-8716(97)02263-4.

[1] The standardization efforts for ATM seem to have followed the same path, although they occurred at a time when the need to support multipoint communication was already clear.

to-multipoint communication. We prefer the term "multipoint" because of its more general connotations.

II. The Challenges of Supporting Multipoint Communication

What is so hard about providing network support for multipoint communication? It is clearly a generalization of the well-understood point-to-point paradigm. As one begins to think about such a generalization several hard questions emerge.

- *In supporting multipoint communication, how much can one reuse existing point-to-point support?*
  It is clear that one can support multipoint applications, albeit extremely inefficiently, using point-to-point infrastructure only. Given that this is extremely inefficient, how much of that support can be reused? Some of the early work on multipoint communication used this reuse as one of the measures for multipoint solutions.[2] Given the importance of multipoint applications would it make more sense to "cut our losses" and think of redoing the infrastructure to support multipoint communication first and support point-to-point only as a special case? Or will we always need separate infrastructures to support the two forms efficiently?
- *What are the exact semantics of multipoint communication?*
  The semantics of point-to-point communication are very well understood. For every message there is exactly one recipient with a unique address. Reception reliability, packet ordering, and end-to-end flow control requirements are straightforward to define relative to this one receiver. The added dimension of multiple receivers complicates matters considerably.
  - *Addressing:* How should this group of receivers be addressed at each level? The group address paradigm used within IP mimics the channel concept. Are there other ways to address groups of receivers that make more sense for other applications? How should groups be identified in layers above the network layer?
  - *Reliability and Sequencing:* What does it mean to reliably transmit to a set of multiple receivers? Can some subset of the receivers fail to receive a message and delivery still be considered reliable? Of course, this depends on the semantics of the application. One important, yet to be answered, question is whether one can define a general reliability model that can fit all applications, or whether different reliability models are required to span the class of possible applications.
  - *Symmetry and Directionality:* Applications can be of the one-to-many variety or the many-to-many variety with possible distinctions being made as to whether the transmitter(s) are a subset of the receivers. Does one need to provide different notions of network support based on the type of application, or is there one form of generic support that can be used for all application forms?
- *How does one manage the heterogeneity/scalability trade-off?*
  As mentioned earlier, part of the appeal of multipoint communication is its scalability: the effort expended by a transmitter and the network grows less than linearly as the number of recipients increases. This scalability is achieved at the cost of sacrificing individual treatment of receivers. This sacrifice is most apparent in the case of receivers with heterogeneous capability and requirements or when the network paths leading to receivers have differing capabilities or traffic loads. The question here is how can we get back some of this individuality without losing the scalability benefits of multicast communication. Is network intervention required to deal with this heterogeneity problem, or are there suitable end-to-end techniques that can be used for all application forms?

[2] In fact one of the primary "selling points" for the reverse path forwarding technique, later adopted as the basis of DVMRP, was that it used the same routing tables as point-to-point routing.

III. Papers in this Issue

This issue of this Journal documents recent results of ongoing research aimed at answering the questions above in one form or another. The scope of the issues addressed by the papers illustrates the complexity of the problem. This is also made clear by the first paper by Diot *et al*. surveying the various multipoint protocols and their functions. This paper serves as an introduction and highlights recent and classic research in the area.

A. Multipoint Routing

The next set of eight papers deals with the difficult problem of routing multipoint data over packet switched networks. The first three papers in the set consider proposals for multipoint routing in the Internet. The paper by Thalerand Ravishankar considers center-based routing algorithms and develops and analyzes techniques for locating the center. The paper by Billhartz *et al.* evaluates, through extensive simulation, the CBT and PIM Internet routing protocols. The paper by Parsa and Garcia-Luna-Aceves proposes a provably loop-free routing protocol that has potential for Internet deployment.

The next three papers in the set study the issue of routing real-time multipoint traffic and the special considerations this imposes on the algorithms for constructing routing trees. The paper by Salama *et al.* compares many routing algorithms capable of guaranteeing quality of service requirements. The paper by Rouskas and Baldine considers constraining the variation of delays among receivers as well as constraining the end-to-end delay. The paper by Maxemchuk considers the issue of routing multipoint video to a heterogeneous set of receivers, each with its own bandwidth requirements.

The last two papers in the set consider other issues in multipoint routing protocol design. The paper by Shaikh and Shin presents a new algorithm based on biasing multipoint routes to pass through destination nodes, and the paper by Bauer and Varma develops and evaluates a heuristic for dynamic update of a routing tree.

B. Reliable Multipoint Transport

The issue of reliable transport of multipoint data is currently receiving significant attention in the research community. The four papers appearing in this section provide good coverage of the various issues surrounding this very important topic. The paper by Towsley *et al.* provides a framework for comparing sender-initiated (ACK-based) or receiver-initiated (NAK-based) protocols and for evaluating various NAK transmission options. The paper by Paul *et al.* describes the design and implementation of RMTP, a protocol achieving its scalability through the use of designated routers as retransmission agents. The paper by Grossglauser examines the question of how one should set timeout values in reliable multicast protocols to avoid NAK implosion. Finally in this set, the paper by Ofek and Yener describes a reliable multicast protocol designed to operate in an environment with bursty sources that combines receiver and sender-initiated techniques.

C. Multipoint Communication in ATM

Supporting multipoint communication in ATM networks provides its own set of unique challenges. These are exemplified in the next set of three papers. The paper by Armitage is an overview of the architecture and mechanisms of the MARS model, developed within the IETF to support IP multicast over ATM networks. The work described in the paper by Gauthier *et al.* develops a many-to-many shared-tree protocol for use on an ATM network. The paper by Chiussi *et al.* studies the performance of shared-memory switch architectures carrying multipoint traffic.

D. Multipoint Communication in WDM Networks

WDM (wavelength division multiplexing) networks provide a unique environment in which multipoint support needs to be provided. The paper by Tridandapani and Mukherjee considers the issue of supporting multipoint communication in a multi-hop WDM network using channel sharing. The paper by Rouskas and Ammar develops and analyzes protocols suitable for multipoint communication in single-hop WDM networks.

E. QoS and Congestion Control

The problems (beyond those of routing) of supporting real-time applications through quality of service (QoS) guarantees and congestion/admission control in a multipoint environment are explored in the next set of papers. The paper by Vickers *et al.* proposes and analyzes a new service architecture and new feedback techniques to be used in a multipoint environment that supports minimum bandwidth guarantees for transmitted video. Consideration is given to the heterogeneity in receiver capabilities. The paper by Moghé and Rubin is a study of how to provide QoS guarantees in a dynamic multipoint application. They propose the idea of application reservation, where resources required for an entire session (including future receivers) are reserved at session setup. The work by Tzeng and Siu proposes extensions to point-to-point congestion control protocols proposed by the ATM Forum for operation in a multipoint environment while preserving fairness. The paper by Shacham and Yokota considers admission control schemes for hierarchically encoded real-time streams.

F. Multipoint Collaboration Applications

Ultimately, multipoint network support is used to provide some network service. Among the most important ones are collaboration services. The paper by Gong addresses the important question of how to provide secure collaboration over the Internet. The paper describes the design and use of the Enclaves toolkit developed for such a purpose. The paper by Blum *et al.* describes a platform designed to support fast implementation of multimedia collaboration applications.

ACKNOWLEDGMENT

The Guest Editors would like to thank the authors of all the papers submitted to this issue. The wealth of research in the multipoint communication area was evident from the large number of submissions that we received. The Guest 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. A. Bush, JSAC Board Representative, for his guidance and V. Johnson (Georgia Tech), M. Foley (UCSD), and J. Doherty (University of Maryland) for their help with the issue.

MOSTAFA H. AMMAR, *Guest Editor*
Georgia Institute of Technology
Atlanta, GA 30332

GEORGE C. POLYZOS, *Guest Editor*
University of California at San Diego
La Jolla, CA 92093

SATISH K. TRIPATHI, *Guest Editor*
University of Maryland
College Park, MD 20742

A. BUSH, *JSAC Board Representative*

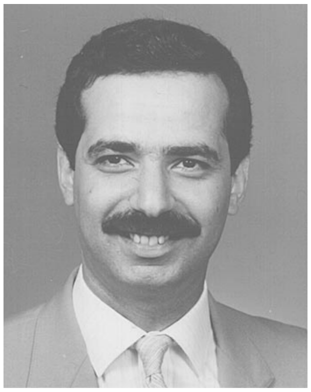

Mostafa H. Ammar (S'83–M'85–SM'95) received the S.M. and S.B. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, in 1980 and 1978 and the Ph.D. degree in electrical engineering from the University of Waterloo, Ont., Canada, in 1985.

Since 1985, he has been an Associate Professor in the College of Computing at Georgia Institute of Technology, Atlanta. From 1980 to 1982, he worked at Bell-Northern Research (BNR) as Manager of Data Network Planning. His research interests are in the areas of computer network architectures and protocols, multipoint communication, distributed computing systems, and performance evaluation. He is the co-author of the textbook *Fundamentals of Telecommunication Networks* (New York: Wiley).

Dr. Ammar is the holder of a 1990–1991 Lilly Teaching Fellowship and received the 1993 Outstanding Faculty Research Award from the College of Computing. He is a member of the editorial boards of IEEE/ACM TRANSACTIONS ON NETWORKING and *Computer Networks and ISDN Systems Journal*. He is the Technical Program Co-Chair for the 1997 IEEE International Conference on Network Protocols. He is a member of the ACM and a member of the Association of Professional Engineers of the Province of Ontario, Canada.

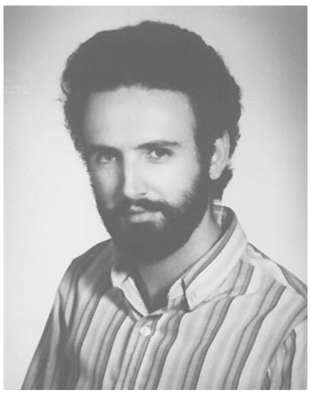

George C. Polyzos (S'80–M'88) received the Diploma in electrical engineering from the National Technical University, Athens, Greece, and the M.A.Sc. degree in electrical engineering and the Ph.D. degree in computer science from the University of Toronto, Ont. Canada.

He is an Associate Professor of Computer Science and Engineering at the University of California, San Diego (UCSD). He is also Codirector of the Computer Systems Laboratory and member of the Center for Wireless Communications at UCSD and a Senior Fellow of the San Diego Supercomputer Center. His research interests include communication network and distributed multimedia system design and computer and communications systems performance evaluation, with his most current research focusing on wireless multimedia networks and efficient multicast. He has published papers in the areas of random access protocols, local and metropolitan area networks, high performance computing, Internet traffic and performance characterization, network and operating system support for multipoint multimedia communications, and efficient multicast routing.

Dr. Polyzos is a member of the ACM and the Technical Chamber of Greece.

Satish K. Tripathi (M'86–SM'86) attended the Banaras Hindu University, the Indian Statistical Institute, the University of Alberta, and the University of Toronto. He received the Ph.D. degree in computer science from the University of Toronto, Ont. Canada.

He is a Professor in the Department of Computer Science at the University of Maryland, College Park. He has been on the faculty at the University of Maryland since 1978. From 1988 to 1995, he served as the Department Chair. In March 1997, he will join the University of California at Riverside as the Dean of Engineering and the Johnson Professor of Engineering. For the last 20 years, he has been actively involved in research related to performance evaluation, networks, real-time systems, and fault tolerance. In the networking area his current projects are on mobile computing, ATM networks, and operating systems support for multimedia information. He has written more than 100 papers for international journals and refereed conferences. He has edited books in the areas of performance evaluation and parallel computing.

Dr. Tripathi has served as the member of the Program Committee and Program Chairman for various international conferences. He has guest edited special issues of many journals and serves on the editorial boards of *Theoretical Computer Science*, *ACM/Springer Multimedia Systems*, IEEE/ACM TRANSACTIONS ON NETWORKING, and IEEE TRANSACTIONS ON COMPUTERS.