Quality of Service Routing

Marwan Krunz   *Ibrahim Matta*

Quality of service (QoS) routing is a class of solutions that are aimed at identifying cost-efficient network paths that have sufficient resources to meet certain performance and reliability constraints. It represents a radical shift from the traditional connectivity-based approach of currently deployed intradomain (e.g., OSPF) and interdomain (e.g., BGP) routing protocols to a constraint-based approach. Service architectures developed in the last decade, including ATM, integrated services (IntServ), and more recently differentiated services (DiffServ), all require (albeit implicitly) QoS routing solutions to ensure efficient network operation. For connection-oriented services (e.g., ATM, IntServ), QoS routing can be used to find appropriate paths for connection requests. Recently, several proposals have been put forth on how to integrate QoS routing into a DiffServ framework. For example, in one proposal that advocates integrating the traffic engineering and scalable reservation aspects of MPLS into the DiffServ architecture, QoS routing is suggested as a mechanism to establish "MPLS paths" onto which aggregates of IP flows can be transported with given QoS. In this context, MPLS paths are dynamically established between ISPs as part of their service level agreements (SLAs).

This Feature Topic consists of five tutorial articles that address some of the important problems in QoS routing and their potential solutions. The selected articles by no means encompass the full spectrum of QoS routing issues; rather, they are intended to give the reader a sample of the issues encountered when moving toward a constraint-based routing paradigm. The five articles address three forms of communication: *anycast*, where one communicating endpoint (the client) is specified and the other endpoint (the server) is selected from many "equivalent" hosts; *unicast*, where both endpoints are specified; and *multicast*, where an endpoint simultaneously sends data to several endpoints. Efficiency is addressed through various forms of adaptation, including rearrangement of multicast trees, and maintenance of guaranteed-bandwidth backup paths. To achieve scalability, the featured articles present various techniques, including path caching to reduce computation cost, and dynamic multipath routing to reduce communication overhead while overcoming stale information in network state.

In "QoS Routing for Anycast Communications: Motivation and an Architecture for DiffServ Networks," Hao, Zegura, and Ammar consider QoS routing for anycast communication. They propose, for DiffServ networks, a network-aware application-layer architecture that uniquely integrates aspects of server and path selection subject to QoS constraints. The arti-

cle explores many design choices, including server selection and sorting algorithms (e.g., select a feasible server whose domain has the most available resources), signaling protocols (e.g., initiate the reservation request starting from server domains or from the client), and the granularity of resource reservation (e.g., the reservation could be split and resources reserved from several server domains to the client.)

The second article, "QoS Routing Granularity in MPLS Networks" by Lin, Hsu, and Hwang, investigates hybrid path caching schemes that are intended to improve the scalability of unicast QoS routing in MPLS networks. Essentially, the proposed schemes lie between the traditional coarse-grained (per-destination) approach and the fine-grained (per-flow) approach. One of these schemes caches, for each source-destination pair, one common path plus few per-flow paths. A second scheme caches, for each source-destination pair, a few class-based paths. These hybrid caching schemes can reduce the cache size without adversely affecting the blocking rate of bandwidth requests.

QoS routing solutions often rely on global dissemination of link metrics (e.g., available bandwidth over the link) to compute appropriate paths. Such dissemination, provided by extended versions of link state protocols, often incur a rather high communication overhead. In "Localized Adaptive Proportioning Approach to QoS Routing," Nelakuditi and Zhang present a "localized" QoS routing approach for unicast connections that uses multiple paths between a source and a destination. The approach is localized in the sense that the source node dynamically adapts the proportions of traffic over the set of candidate paths. Localized proportional routing can improve blocking performance even with infrequent network state updates.

Fault tolerance and path restorability are yet another facet of QoS routing. If a network element (node or link) fails, all QoS paths traversing that element have to be rerouted. To reroute network flows while maintaining their original QoS requirements, backup paths have to be provisioned. In "Restorable Dynamic QoS Routing," Kodialam and Lakshman describe several algorithms for restorable QoS routing. They develop a mathematical framework under which different models of restoration (local vs. end-to-end), failure modes (link or node), and link usage information (partial vs. complete) are formulated. The proposed schemes have direct applications in MPLS networks.

Finally, in "Managing Group Dynamics and Failures in QoS Multicasting," Striegel and Manimaran provide a nice overview of QoS routing for multicast connections. They first discuss

each stage in the life cycle of a multicast session, identifying the various tradeoffs and research issues. Then they study in more depth two important aspects of QoS multicast, group dynamics and failure recovery. Under group dynamics, they discuss how members with QoS requirements join or leave the multicast group, the rearrangement of the multicast tree, and tree (and core node) migration. Under failure recovery, they identify the issues and trade-offs involved in recovering a failed core node (e.g., core reselection, local vs. global recovery).

We received 21 submissions in response to our CFP, which were all sent for external review. The selection process was tough, and we were not able to accommodate several good papers. We would like to express our thanks to the authors for responding to this Feature Topic and to the reviewers for the valuable time they spent reviewing the papers.

BIOGRAPHIES

MARWAN KRUNZ [M] (krunz@ece.arizona.edu) received his Ph.D. degree in electrical engineering from Michigan State University in 1995. From 1995 to 1997 he was a postdoctoral research associate with the Department of Computer Science and the Institute for Advanced Computer Studies (UMIACS), University of Maryland, College Park. In January 1997 he joined the University of Arizona, where he is currently an associate professor of electrical and computer engineering. His research interests lie in the field of computer networks, especially in its performance and traffic control aspects. His recent work has focused on the provisioning of QoS over wireless links, QoS routing (path selection, state aggregation), traffic modeling (video, Web), bandwidth allocation, video-on-demand systems, and power-aware protocols for ad hoc networks. He has published more than 50 journal articles and refereed conference papers. He is a recipient of the National Science Foundation CAREER Award (1998–2002). He currently serves on the editorial board for *IEEE/ACM Transactions on Networking*, *Computer Communications Journal*, and *IEEE Communications Interactive*. He is a guest co-editor of a special issue on Hot Interconnects (*IEEE Micro*, January 2002). He was Technical Program Co-chair for the 9th Hot Interconnects Symposium (Stanford University, August 2001). He has served and continues to serve on the executive and technical program committees of many international conferences. He served as a reviewer and panelist for NSF proposals, and is a consultant for several corporations in the telecommunications industry.

IBRAHIM MATTA [M] (matta@cs.bu.edu) received his Ph.D. in computer science from the University of Maryland at College Park in 1995. He is currently an assistant professor at the Computer Science Department of Boston University. He leads the QoS Networking Laboratory and is a member of the Web and InterNetworking Group (WING). His research involves the design and analysis of QoS and wireless architectures and protocols. His recent projects investigate QoS routing, Internet topology and traffic analysis, and Internet traffic controllers. He received the National Science Foundation CAREER Award in 1997. He was guest co-editor of a special issue on Reliable Transport Protocols for Mobile Computing in the *Journal of Wireless Communications and Mobile Computing*, 2002. He has served on the technical program committees of many conferences including INFO-COM, ICNP, GLOBECOM, and ICDCS. He is Technical Program Co-chair of the first International Workshop on Wired/Wireless Internet Communications (WWIC 2002). He is Publications Chair of IEEE INFOCOM 2003, and was Tutorial and Panel Chair of the 9th Hot Interconnects Symposium 2001. He served as session organizer and chair, reviewer, and panelist for NSF networking grant proposals, and was the representative of the IEEE Technical Committee on Computer Communications (TCCC) for GLOBECOM 1999. He is a member of ACM.