GUEST EDITOR'S INTRODUCTION: QUALITY OF SERVICE IN INTEGRATED SERVICES NETWORKS

Judging from the popular press, it would seem that the information superhighway is just around the corner. Exciting new multimedia applications are being designed for networks, and a growing number of us routinely surf the Internet to sample its multitude of images and sounds. Networking hardware technology continues to improve at a steady rate as well, making it only a matter of time before most networks are upgraded to run at megabit if not gigabit speeds.

There are many exciting challenges to be faced by the architects and builders of this evolving information infrastructure. Expectations from "the network" have *more* than kept up with the improvements in link and processor speeds. The network must accommodate a staggeringly wide range of traffic types (both in terms of message size and data rate), provide real-time performance when required, and do all of this at low cost. The trade-offs among performance, flexibility, and cost in the design of these *integrated services networks* remain poorly understood, despite the large amount of work going on in the area.

The papers in this issue deal with the important issue of how *quality of service* can be assured in integrated services networks. Given the extremely large bandwidths of optical fibers, the performance bottlenecks in integrated services networks reside in the switches and in the end-devices. While optimizing hardware and software performance at the host interface is an interesting and challenging area of research, this special issue focuses on end-to-end performance measured from the time that packets *are submitted* to the network to the time just *before they are delivered* to the host.

One can view quality of service as an "appropriate" level of assurance or guarantee from the network provider to a user on the delivery of packets. As things stand, there is no clear agreement on what the nature of these guarantees should be, nor of the network mechanisms that could pave the way for them. Many of the unresolved issues center around the willingness to pay for added flexibility and performance. The stringency with which quality of service should be guaranteed is an interesting and hotly debated issue that has not been resolved. It has, however, been recognized by most that something more than the traditional "best-effort" service is required. As the process of finding the "appropriate architecture" continues, it is important to provide simple yet useful mechanisms in the interim that mitigate the effects of rapidly increasing traffic volumes on the internet. The paper by Bohn, Braun, Claffy, and Wolff provides such a scheme for IP networks that is based on multiple levels of precedence.

Even if problems of heterogeneity and backward compatibility are ignored, there is still the fundamental question of how quality of service should be measured. In an effort to provide better guarantees than simple best-effort delivery, work has gone into schemes that provide guarantees on the "average" behavior of *aggregate* delay and loss characteristics when the network is in steady state. These guarantees may be very hard for any user to actually verify since he will not be able to measure aggregate traffic nor will he be able to determine that the network is in steady state. Also, the intervals over which the "average" is guaranteed may be too large relative to the life-time of a particular session. Thus, from the standpoint of a single user, the aggregate average guarantees are not too different from best-effort guarantees.

An approach to provide the user with more meaningful guarantees is presented in the paper by Nagarajan, Kurose, and Towsley. They propose and examine new quality of service criteria for *packet loss* that are based on averages over *finite time horizons*. The metrics are computed over the duration of the connection or over a fixed interval of time known to the user. The authors also use simple examples to demonstrate the inadequacies of guarantees on packet loss that are based on steady-state analysis.

An appropriate source model is critical to any frame-work that guarantees performance. In the paper by Yaron and Sidi, sources are characterized as Exponentially Bounded Processes. Interestingly, the model allows for bounds on the *distribution* of backlogs in *arbitrary topology networks*, and takes into account the dynamics associated with various scheduling disciplines. In this paper, the authors employ a scheduling discipline called Generalized Processor Sharing at the nodes of the network.

Scheduling at the nodes of the network is a key component of many schemes that guarantee quality of service. Schedulers generally attempt to control the queueing delay/buffer loss by class/session. Since the goal is to manage these quantities on an end-to-end basis, it is important to ensure that the output traffic processes from each node in the network be well behaved (since they form the inputs to subsequent nodes and affect performance at these nodes). Consequently, traffic shaping or rate control is often incorporated into the scheduling policies either explicitly or implicitly. The paper by Zhang and Ferrari introduces an interesting class of service policies called Rate-Controlled Service disciplines that *explicitly separate* the rate-control and scheduling functions.

The paper by Cruz and Liu attempts to understand the impact of scheduling in a more abstract and general setting. Specifically, they focus on a single connection in a discrete time system and model the scheduling policy at each node in the route of the connection as an *availability process.* This process defines the number of packets from the connection that will be scheduled for service (given that many packets are in queue) at a node in each slot.

The paper by Yates, Kurose, Towsley, and Hluchyj is a detailed study comparing various approaches to *scheduling and call admission.* Given the real-time requirements of many multimedia applications, it is not a surprise that much effort has gone into understanding how to obtain worst-case end-to-end bounds on delay, and more generally on the tails of end-to-end delay distribution. The paper focuses on how several different schemes for guaranteeing bounds on end-to-end delay perform in the context of a large tree network, and presents empirical results.

The impact of scheduling policies on buffer loss is an interesting area that is explored in the paper by Huang and Leon-Garcia. A frame-oriented policy based on Golden-Ratio hashed slot assignment is analyzed. A key feature of this analysis is that it shows optimality *independent* of the traffic statistics.

In addition to scheduling, switches have to implement *buffer management* policies to decide which packets to discard when buffers are full. In a shared memory switch this problem is particularly interesting since the memory is shared by packets on all of the outgoing links. The paper by Choudhury and Hahne is a detailed simulation study of the performance of various policies in such a shared memory switch. They find that a Selective Push-out scheme has attractive properties over simpler pushout and threshold schemes.

One of the key features of high speed switches is their ability to support multicast. The importance of this feature comes from the fact that many applications (such as distance learning) are inherently multipoint and may require real-time performance from the network. The paper of Verma and Gopal tackles the problem of *computing trees* for reserved bandwidth multipoint connections.

Ensuring quality of service in integrated services networks is very much a growing and exciting area of engineering research. It has much to offer in terms of intellectual excitement and opportunity for significant practical impact. I am sure that you will find these papers to be insightful and of the highest quality, and I thank the authors and reviewers for their efforts. I am also grateful to the editor-in-chief for giving me the opportunity to put together this special issue, and for his support for the duration of the project.

Abhay K. Parekh
IBM Corporation
MS H3-D42
PO Box 704
Yorktown Heights, NY 10598
USA