INTRODUCTION

In this issue we present some interesting new work in the areas of modeling and scheduling for high speed networks. We have an improved model of network traffic, an optical ring network that makes effective use of the potential optical bandwidth, a dynamically stable congestion-control protocol, and distributed algorithms for channel assignment in multi-hop radio networks.

Real broadband computer network traffic has extreme variability over time scales ranging from milliseconds to months. This variability is not captured by conventional Poisson or Markov Modulated Poisson Process (MMPP) traffic models. In an expansion of a paper presented at SIGCOMM '93, Lin, Tsai, Huang, and Gerla have developed a new traffic model, the Hierarchical Arrival Process (HAP) model. The HAP model is a generalization of current on-off traffic models, and is shown to be equivalent to a special class of MMPP's. Delay under HAP traffic can be at least an order of magnitude greater than the delay under Poisson traffic, depending on parameters and load, and congestion may persist for minutes. HAP's dramatic short-term behavior can explain the occasional congestion in real networks. Three algorithmic methods and simulations are applied to evaluate the queuing performance under HAP traffic.

Two fundamental problems need to be dealt with in an optical data transmission system operating on counter-rotating rings: packet contention and packet transmission control. In "An Optical Data Double Ring Network," Imrich Chlamtac and Andrea Fumagalli propose an optical ring network that can take full advantage of the potential optical bandwidth by eliminating performance bottlenecks associated with the electronic processing of transmitted data. Their solution does not require multiple transmitters and receivers at each node to resolve contention, and offers route-through, fault-tolerant switches. Performance evaluation shows that this ring network can asymptotically achieve the ideal capacity of an all-optical contention-free ring.

There is some consensus that congestion-control protocols for broadband packet networks should be simplified to minimize switching overhead, and that prevention is preferable to reaction, due to the high costs associated with retransmission. In "Warp Control: A Dynamically Stable Congestion Protocol and its Analysis," Kihong Park presents a distributed, end-to-end congestion control protocol for use in high-traffic, packet-switched networks. (A shorter version of this paper was presented at SIGCOMM '93.) Park's work differs from existing rate-based access control schemes in two respects. First, a time-stamp based measure of network state called *warp* is defined, and is shown to be an estimator of network utilization. Second, congestion is modeled explicitly using with a *load-throughput function*, which relates effective throughput to actual load. A fundamental conflict in achieving both stability and optimal throughput is investigated, giving new insights into the difficulties of congestion control.

There has been a tremendous growth in wireless and mobile communication systems in recent years. In the final paper of this issue, Errol L. Lloyd and S. Ramanathan present "Efficient Distributed Algorithms for Channel Assignment in Multihop Radio Networks." The authors present new distributed algorithms for the *channel assignent problem* in Time Division Multiple Access (TDMA) networks. Their algorithms compare favorably with existing centralized schemes, are parameterized to allow a tradeoff between schedule quality and efficiency, are easy to implement, and have low communication overhead.

We are pleased to offer our readers this selection of articles, and would like to express our thanks both to the authors and to the JHSN and SIGCOMM '93 reviewers.

Howard E. Motteler

Member, JHSN Editorial Board