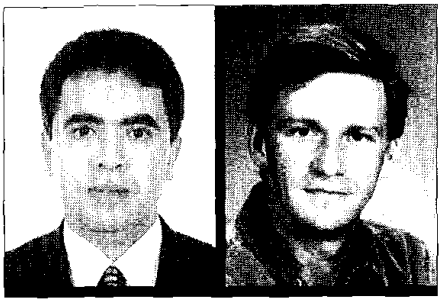

Mounir Hamdi *Nick McKeown*

SCALABLE HIGH-SPEED SWITCHES/ROUTERS WITH QOS SUPPORT

As we enter the 21st century, the Internet continues to experience extraordinary growth. By any measure, the growth is remarkable on all fronts: the number of hosts, the number of users, the amount of traffic, the number of links, the bandwidth of individual links, or the growth rates of Internet service provider (ISP) networks. This continued rapid growth, coupled with the variety of services that the Internet is expected to provide, creates two main challenges for the design and implementation of packet switches (IP Routers, ATM switches, Gigabit Ethernet switches and Frame Relay switches), queue management, and forwarding.

First, to keep pace with the growth in Internet usage, we need higher capacity packet switches with aggregate data rates of multiple terabits per second and forwarding rates of billions of packets per second. Limitations in memory bandwidth and interconnect technology, along with the challenges of system packaging and managing system power, mean that new architectures and techniques are needed to keep pace with capacity requirements. Second, so long as the network capacity is limited, there will be a desire to provision a variety of service classes for users whose traffic requires special service, such as guaranteed delay, bounded delay variation, minimal packet loss, or controlled access. While it is generally well known how to provide these services at low speed, it requires further research and investigation to determine how they can be delivered in a very high capacity network.

The purpose of this special issue is to review new techniques, architectures, and technologies that make possible the provision of differentiated qualities of service in very high capacity switches.

We have received many interesting articles for this special issue. We have selected four articles that cover a wide spectrum of issues that are essential for the provision of differentiated quality of service in high capacity switches.

The first article in this issue, by Nong and Hamdi, gives a comprehensive survey of the techniques and scheduling algorithms that have been proposed for providing quality-of-service (QoS) guarantees on scalable switches such as virtual-output-queued switches and combined input-output-queued switches. The proposed algorithms are classified under:

- Time slot assignment
- Maximal matching
- Stable matching

Their capabilities and complexities in terms of providing QoS guarantees and in terms of emulating output-queued switches is analyzed. Finally, the article presents some open problems and future directions in this area.

Editorial liaison: S. Giordano

The second article, by Minkenberg and Engbersen, describes the architecture of a switch that combines input and output queuing. The purpose of having buffers at both the input and output side is:

- To achieve good performance and QoS capability
- To reduce the complexity and interaction between the input and output arbiters

They present simulation results to demonstrate the performance of this architecture. In addition, they outline the implementation of this architecture using the PRIZMA family of switch chips.

The third article, by Chao, proposes a high-speed switch architecture with a moderate speedup. The article proposes a low-complexity dual-round-robin scheduling algorithm that achieves high throughput and appears from simulation to have a low statistical delay bound. In addition, the author uses a new token-tunneling technique to arbitrate the contenting packets at a very high speed.

Finally, the article illustrates the switch design using multiple switch backplanes and shows that it can potentially achieve a terabit/sec capacity using current technology. Finally, Shiomoto, Uga, Omotani, Shimizu, and Chimaru describe a QoS-capable IP+ATM switch/router that has been designed and constructed at NTT. The architecture consists of a core ATM switch which is extended to provide IP routing capabilities via extra packet lookup hardware added to the ATM linecards. The article also gives an overview of the various capabilities of a real high-speed, 40 Gb/s switch and how they are provided. It shows how the IP QoS can be mapped onto the inherent QoS capabilities of the core ATM switch.

BIOGRAPHY

MOUNIR HAMDI (hamdi@cs.ust.hk) received his B.S. degree in computer engineering (with distinction) from the University of Southwestern Louisiana in 1985, and the MS the Ph.D. degrees in electrical engineering from the University of Pittsburgh in 1987 and 1991, respectively. He has been a faculty member in the department of computer science at the Hong Kong University of Science and Technology since 1991, where he is now associate professor of computer science and the director of the computer engineering program. In 1999 and 2000 he held visiting professor positions at Stanford University and the Swiss Federal Institute of Technology. His general areas of research are in networking and parallel computing, in which he has published more than 120 research publications, and for which he has been awarded more than 10 research grants. Currently, he is working on high-speed networks including the design, analysis, scheduling, and management of high-speed switches/routers and on wavelength division multiplexing (WDM) networks/switches.

NICK MCKEOWN [SM] is an assistant professor of electrical engineering and computer science at Stanford University. He received his Ph.D. from the University of California at Berkeley in 1995. From 1986 to 1989 he worked for Hewlett-Packard Laboratories in their network and communications research group in Bristol, England. In spring 1995, he worked briefly for Cisco Systems. In 1997 he co-founded Abrizio Inc. (now part of PMC-Sierra), and is on several technical advisory boards. He is currently the Robert Noyce Faculty Fellow at Stanford. He researches techniques for high-speed networks with a particular interest in helping the core infrastructure run faster.