Optical Versus Electronic Switching for Broadband Networks

Joseph Y. Hui, Rutgers University
Kwok-Wai Cheung, Chinese University of Hong Kong

Abstract

The authors compare the use of electronic and optical switching for broadband networks. Their technological merits and their ability to meet traffic demand when multimedia communications become ubiquitous are considered. The authors conclude that electronic ATM switches are adequate and commercially mature for local access. An ATM-Mesh Switch with multihop routing can be used to achieve capacity of several hundreds of gigabits per second. Beyond access, the authors argue that traffic aggregation may allow switching to be performed on much longer time scales than that of ATM cells. Wavelength and mechanical fiber switching may be used economically for trunk switching, avoiding the expensive processes of opto-electronic conversions and demultiplexing.

Given the large amount of multimedia traffic that will need to be switched in the future, it is natural to ask whether photonic switching can be employed in a network to reduce the complexity and cost, and where the business case is for such a deployment [1, 2]. The main argument for the adoption of photonic switching is the elimination of the so-called electro-optic bottleneck — the interface where conversion between the optical signal and the electrical signal occurs for switching, routing, and further processing of signals.

Optical amplifiers and wavelength-division multiplexing are being adopted commercially for long distance transmission [3]. What, then, is the role of optical switching in networks with wavelength multiplexed transmission links? Proponents of all-optical switching may see electronic switching completely replaced by optical switching, wavelength routing, and wavelength conversion [4]. At the other extreme, proponents of terabit-per-second electronic switching [5] may see electronic switching as so cost effective that optics is needed only for point-to-point transport; switching, routing, and processing should be left to electronics.

Conventional wisdom is that optics is good for transport but poor for processing of signals; this is because photons can be neither slowed down, localized, nor stored conveniently. The size of photonic devices is limited by diffraction, which is on the order of one wavelength, making photonic devices much larger than electronic devices. If one accepts this statement as a fundamental limitation of optics, then the proper application of photonics in switching networks must be very carefully evaluated.

While good progress is being made on the research of optical switching, high-speed electronic ATM switches are becoming commercially mature [6, 7]. These switches are capable of supporting multimedia communication and have been deployed for enterprise networks, Internet backbones, and public switched networks. They are capable of supporting a large variety of protocols such as Internet Protocol (IP), frame relay, switched multimegabit data service (SMDS), as well as local/wide area network (LAN/WAN) emulation. Total switch capacity ranges from a typical 2.4 Gb/s for an asynchronous transfer mode (ATM) LAN to 10 Gb/s [6] for enterprise and campus networks. Switches with capacity exceeding 100 Gb/s have been announced for central office applications [7]. The later generation of ATM switches typically boasts a much larger buffer capacity per port for congestion handling.

Due to these technological and commercial factors, a consensus is emerging that switching at or above the ATM layer should be done electronically [1]. For lower layers, it is not clear whether optical switching could provide better economics and capacity. In this article, we attempt to answer the following questions. First, is there a bit rate limit for which electronic switches will become technically expensive and infeasible? If so, will this rate be needed in practice? Second, what services need to be performed at the lowest layer of the network? For example, is coarse switching of entire wavelengths [8] sufficient with SONET on top of it, or is SONET switching sufficient by itself?

Before answering these questions, the notion of layered multiplexing and switching is first explained.

0890-8044/96/$05.00 © 1996 IEEE

Authorized licensed use limited to: Chinese University of Hong Kong. Downloaded on January 22,2022 at 09:06:25 UTC from IEEE Xplore. Restrictions apply.

Layered Multiplexing and Switching

In order to facilitate end-to-end routing of information, multiplexing and switching of information for different communication sessions are necessary. These functions are similar to highway traffic routing, where vehicular traffic is multiplexed at entrance ramps and routed at traffic exchanges. However, the heterogeneity and large volume of broadband communications in comparison to vehicular travel requires a substantially more complex arrangement of multiplexing and switching methods. In particular, a layered approach is necessary, as indicated in Fig. 1.

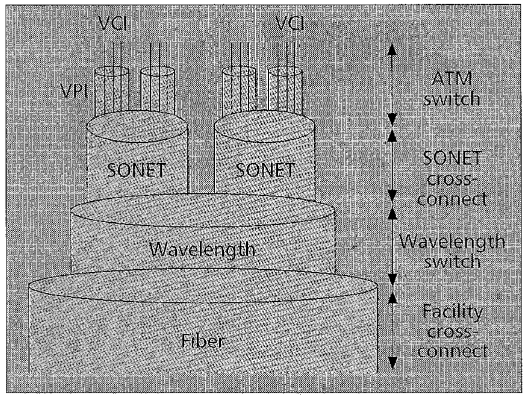

■ Figure 1. *Layered multiplexing.*

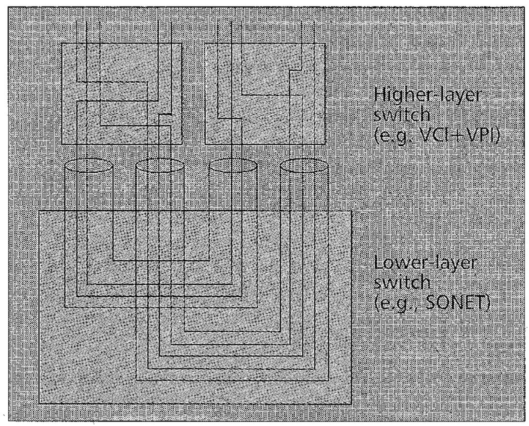

■ Figure 2. *Layered switching architecture.*

Layered multiplexing involves the concentration of traffic from users into communication pipes of increasing capacity from one layer to a lower layer. Each layer represents a different technology and protocol type. When traffic volume is small, multiplexing is usually performed electronically, while for highly multiplexed traffic the trend is to use optical multiplexing such as wavelength division multiplexing [1, 8]. This results from the well-known capability of optics: high capacity for communication but relatively slow processing.

One common multiplexing hierarchy is shown in Fig. 1. The layers are the virtual circuits, which are multiplexed into virtual paths, then into SONET pipes, then into wavelength channels, and then into fiber links. These pipes provide capacity ranging from megabits per second or less for a virtual channel, to hundreds of megabits per second for a SONET channel, to hundreds of gigabits per second and beyond for a fiber when many wavelengths are multiplexed.

Switching can be performed at each layer as an interexchange for information carried by different channels at the same layer, as shown in Fig. 2. These channels are first terminated at a switch interface, and processing related to switching is performed. There are several functions achieved by such processing. First, the channel identifier is mapped into another identifier if the channels are labeled inband, for example, the translation of the virtual circuit identifier (VCI) and virtual path identifier (VPI) for ATM connections. For physical circuits, this translation is not necessary. In either the virtual or physical case, a circuit mapping table has to be maintained which is updated when connections are initiated or terminated. Second, information may have to be buffered for the purpose of synchronizing various connected circuits across the switch and for alleviating temporary congestion at either the input or output channels of the switch. Third, and perhaps the most difficult, is the control of the switching fabric for routing information carried by many channels.

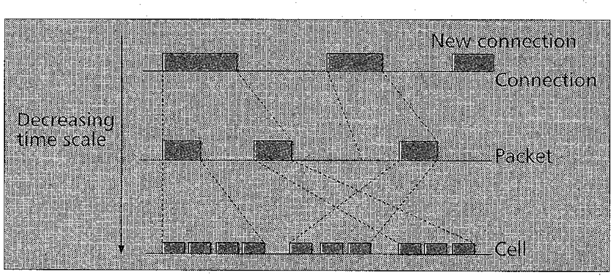

■ Figure 3. *Layered traffic and statistical gain.*

Complexity Factors for Electronics Switching

The fundamental purpose of switching is the reconfiguration of routes and their capacities as the traffic demand across a network changes over time. This fluctuation is often characterized on many time scales, as shown in Fig. 3. For TCP/IP over ATM, we may have time scales such as that for a cell (ATM layer), a packet (IP layer), or a burst (TCP layer). Over longer time scales, we have the call layer, as well as the semi-permanent circuit configured to carry multiple calls. Over time scales of months, transmission facilities may have to be reconfigured to handle differences in traffic growth for geographical areas, or disasters such as earthquakes and fiber breakage.

As we move along the multiplexing hierarchy shown in Fig. 1, a statistical averaging effect is achieved when many connections are concentrated into a channel. In essence the burstiness of the connections is averaged out. At the semi-permanent circuit level, the randomness of the call arrival process is also largely averaged out.

ATM switching has been touted as a technological breakthrough for broadband integrated services digital network (B-ISDN) due to its ability to deal with flexible and variable bit rates for multimedia applications. However, the complexity of ATM switching grows rapidly as the total traffic through a switch increases. The purpose of this section is to look at the various cost factors of ATM switching realized by electronics switching.

The capability of a switch is often measured in terms of three parameters: first, the line speed (S) for each port of a switch; second, the number of ports (N); and third, the size of the data unit (D) for which processing is required for either the protocol functions or routing within a switch.

For given S, the hardware complexity such as crosspoint count of multiplexing and switching grows at a rate faster than linearly in N [9]. Also, switches with large N may have relatively long transmission pathways, which make high-speed transmission and switch synchronization difficult. Without multistaging, a single-stage switch using a shared memory, a shared bus, or a cross-bar has a complexity which grows as N^2. With the use of multistaging such as the buffered banyan net-

Authorized licensed use limited to: Chinese University of Hong Kong. Downloaded on January 22,2022 at 09:06:25 UTC from IEEE Xplore. Restrictions apply.

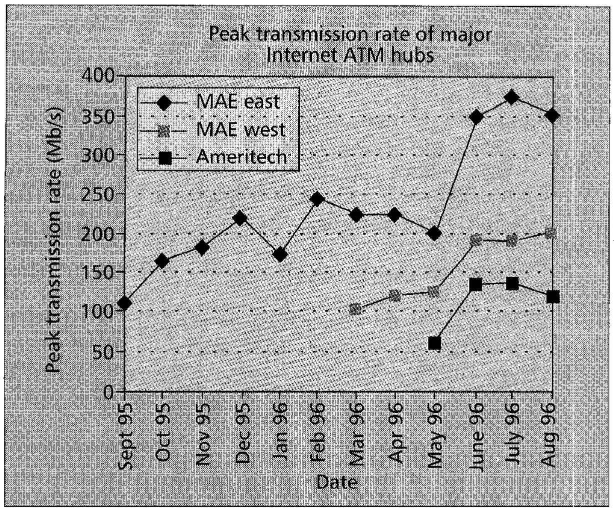

■ Figure 4. *Traffic for ATM hubs.*

work, crosspoint complexity grows as $N \log N$. However, buffered banyan networks suffer from internal blocking unless uneven traffic is somehow randomized, and extensive internal buffering is required, which incurs additional delay. Batcher-banyan switches avoid internal blocking and buffering at the expense of a hardware complexity of $N (\log N)^2$.

Electronic switching has mostly adopted complementary metal oxide semiconductor (CMOS) technology which allows high hardware complexity, but with a line speed S limited to roughly OC-3 speed (155 Mb/s) in practice. Higher line speeds are achieved by either using parallel bus [5,10] or more power-consuming electronics than CMOS. To achieve 155 Gb/s aggregate (N times S) throughput, $N = 1000$ is required for OC-3 interfaces. For the high line speed achievable using parallel transmission, it is generally difficult to achieve a large N [10]. For a local access switch placed at a central office, a large ratio of line speed to user application speed may require expensive hierarchical traffic concentration and multiplexing prior to switching. Such a high per-line speed is perhaps better suited to switching highly concentrated trunk lines.

Let us examine commercially available ATM switches [6, 7] to shed insight into how large an electronic ATM switch may be built. However, such insights should be viewed skeptically because one could argue also that commercially available ATM switches could be much improved, or are built small because of insufficient traffic demand.

Most commercial ATM switches use single-stage technologies for the purpose of simple control and traffic management. They typically have a total switch throughput from 1 Gb/s to 10 Gb/s [6]. Larger switches achieving a capacity close to 100 Gb/s use multistage switching fabric or interconnect smaller switches of several gigabits per second [7].

Switches with capacity of 1 or 2 Gb/s are used mostly for ATM LAN and campus networks. Up to 10 Gb/s ATM switches have been successfully deployed for applications such as frame relay services and Internet hubs [11]. At present, they are more than adequate for such applications. A plot of the traffic volume of the top three Internet hubs in the United States is shown in Fig. 4. While traffic growth almost doubles every half year, peak traffic volume as of August 1996 is well under 0.5 Gb/s. Rapid growth of multimedia applications on the Internet may require the use of ATM switches with capacity exceeding 100 Gb/s, as we shall show later.

One practical limit for ATM switches is the number of user-network interfaces (UNIs) terminating on an access switch. Such interface cost usually dominates the cost of a switch. Current ATM interface cards terminating on a switch typically allow four OC-3 ATM ports per card. This small number is probably a result of the large amount of processing for functions such as buffering, protocol adaptation, and congestion control. A 9.6 Gb/s ATM switch will require at least 16 slots, and a 155 Gb/s ATM switch will require at least 250 slots. This seems to suggest that a 100 Gb/s ATM access switch may already have physical size and management complexity similar to a large-capacity access switch for telephony.

While one may argue that technology may support a much higher line speed and larger total switch capacity [5], the desired size of a switch depends more on economic factors such as user population and traffic demand. Typically, most access switches for telephony serve a local population of 100,000 or less [12]. Is the current commercial 100 Gb/s ATM switch sufficient when a community of at least this size adopts multimedia communications extensively? Can optics help if a larger size switch is needed?

Taffic Demand and Switching for Multimedia Communications

How large should the capacity of an ATM switch be when multimedia applications replace telephony as the predominant communication mode? During peak traffic periods, the total bandwidth demand in the United States can exceed 1 Tb/s for 64 kb/s telephony, assuming 200 million phones in the country each generating 0.1 erlang of traffic [12]. Consider a switch capable of supporting 1 million people, each generating 0.1 erlang of voice traffic. To transport 100,000 erlangs of 64 kb/s telephony, the switch needs a capacity of 6.4 Gb/s. Commercial telephone access switches usually have a total switch capacity of around 1 Gb/s [12].

If 64 kb/s voice applications are replaced by video applications, a switch serving a user population of 1 million, each online 10 percent of the time, would require a capacity of 150 Gb/s for Motion Picture Experts Group type 1 (MPEG 1)-quality video of 1.5 Mb/s, or 600 Gb/s for MPEG 2 quality video of 6 Mb/s. One assumption we make in our traffic analysis is that a transmission capacity of 10 Mb/s is sufficient for most multimedia applications, because limited storage and processing at terminals and servers are perhaps more severe system bottlenecks than transmission.

For such a scenario, our commercially available 100 Gb/s switch could be inadequate. Optical switching may have to be considered seriously, if electronics indeed prevent switches from achieving a capacity beyond 100 Gb/s. To move from 64 kb/s voice to 6 Mb/s video, aggregate traffic demand will increase by a hundredfold for the entire United States to more than 100 Tb/s. High-capacity switching for trunks connecting distant ATM switches will become necessary.

It is argued here that electronics ATM access switching remains viable in such a situation. It is quite possible that ATM switches with capacity approaching 1 Tb/s could be built commercially in the near future. However, one could also use a multihop approach in building larger ATM switches out of smaller modules.

If the capacity requirement per switch exceeds 100 Gb/s, a large switch could be built by locally connecting M switch modules using trunks, as shown in Fig. 5. For a small number of switch modules, a full interconnection topology could be used. We call this network the ATM-Mesh Switch.

Consider the engineering of trunks for the ATM-Mesh Switch. We assume a uniform distribution of destination ports for traffic, which could be achieved if the inputs to the switch are randomly permuted. We assume that the load of the inputs is an average fraction F_a of their capacity. We also

Authorized licensed use limited to: Chinese University of Hong Kong. Downloaded on January 22,2022 at 09:06:25 UTC from IEEE Xplore. Restrictions apply.

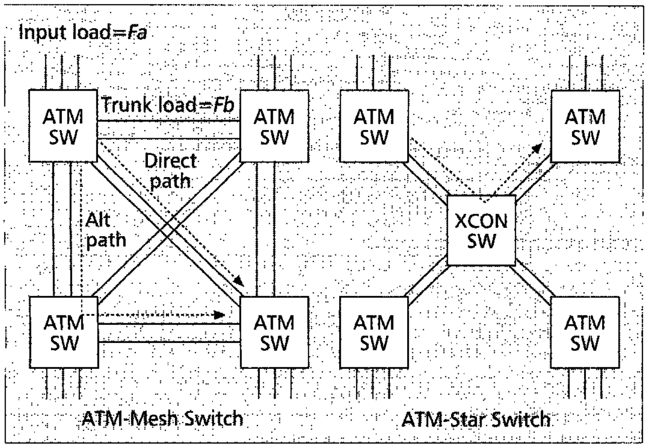

■ Figure 5. *Building large ATM switch by trunking.*

assume that each trunk could be loaded up to a fraction F_b of their capacity. These fractions are engineered close to 1 if significant statistical averaging has been achieved.

It can be readily shown that each switch must use up a fraction

$$\frac{F_a(M-1)}{F_a(M-1)+F_bM}$$

of the total capacity NS for trunking. This fraction is further divided into $M-1$ intermodule trunks, each having a capacity of

$$\frac{NSF_a}{F_a(M-1)+F_bM}.$$

trunks. For NS = 100 Gb/s, M = 5, F_a = 0.4, and F_b = 0.8, 2/7 of the capacity of a switch will be used for trunking. For a switch of total capacity 100 Gb/s, the intermodule trunks will each have a capacity of around 7 Gb/s.

If an uneven traffic pattern occurs among the M modules, multiple routes, each comprising more than one hop, may be used to connect two modules with higher than usual traffic between themselves. The multiple-hop routing could be easily achieved through the use of virtual path switching within the ATM switch itself.

For larger M, a star-type trunk network with trunks incident on a trunk switch as shown in Fig. 5 may be used. We call this the ATM-Star Switch. This configuration reduces the number of trunks from a module and allows better handling of uneven traffic flow. The added complexity is a trunk switch of very high capacity. Multiple stage circuit switch could be used employing very high speed electronic or optical switches. Almost no processing is required for information flowing through the trunk switch. Since trunking is most likely to be done optically, it could be argued that switching should be done optically to avoid the opto-electronic bottleneck.

The previous numbers suggest that trunk capacity required typically exceeds 1 Gb/s, and could be in the range of 10 Gb/s. An argument is made here that the capacity of the trunk is hundreds of times the traffic generated by a user; therefore, statistical averaging of traffic among such a large number of users will make trunk traffic almost a smooth flow. Also, very high trunking efficiency could be achieved. Therefore, slow circuit switching will be sufficient for trunk switching.

Much traffic engineering studies remain to be done for trunking and routing of multimedia traffic. For a large local switch built from interconnected modules, a spatially balanced traffic pattern is quite achievable. For such a situation, the ATM-Mesh Switch is quite sufficient without the use of a trunk switch. However, for a switched network spanning a large geographical area, spatial balance may not be achievable and traffic pattern could evolve over time, thereby requiring the use of trunk switches and more sophisticated routing algorithms.

A natural question arises: Is SONET switching sufficient for the purpose of trunking, or should an additional wavelength layer be added for which multiple SONET channels are carried by a wavelength? Two factors favor the use of wavelength switching. First, as pointed out above, it is quite likely that each trunk will exceed 1 Gb/s and run in the tens of gigabits per second range, which then will be transported over long distances by wavelength division multiplexing. Therefore, switching SONET channels in bulk using a single wavelength would be more efficient than having to demultiplex the trunk into individual SONET channels for switching. Second, the capacity of the trunk switch is expected to exceed 1 Tb/s. This could severely strain the speed and interconnection capability of electronics. Wavelength switching with high port speed but fewer ports could give a less complex switch fabric, as well as substantially reduced port cost.

Therefore, it seems that wavelength switching or mechanical switches for fibers could be an economical means to provide facility configuration at a node exchanging more than 1 Tb/s. Some of these switching systems have reached an advanced development stage, as indicated by Table 1. Doing so may eliminate the expensive process of demultiplexing bundled SONET channels before trunk switching.

	Switching speed	Max. size	Technology maturity	Commercial availability
Spatial switch (mechanical) (electro-optic) [13]	 ≥ ms ≥ ns	 ≥ 128 x 128 128 x 128	 Good Fair	 Yes Yes
Wavelength mux/demux (Static) [14] (Static) [14] (Dynamic) [2]	 ≥ ms	 16 x 16 32 x 1 2 x 1	 Good Good Prototype	 Yes Yes No
Wavelength-selective switch (LCD) [2] (Acousto-optic) [2] (Acousto-optic) [2]	 ≥ ms ≥ μs ≥ μs	 2 x 2 x 8λ 4 x 4 x 4λ 2 x 2 x 8λ	 Prototype Prototype Prototype	 No No No
Wavelength converter [15]	≥ ns	1 x 1 x 8λ	Experimental	No
Fast optical packet switch [15]	≥ ns	4 x 4	Experimental	No

■ Table 1. *Comparison of various optical switching technologies.*

Optical Switching for Trunk Switching

Photonic switches can be classified according to the underlying multiplexing technologies, such as space-division multiplexing, wavelength-division multiplexing, time-division multiplexing, subcarrier

Authorized licensed use limited to: Chinese University of Hong Kong. Downloaded on January 22,2022 at 09:06:25 UTC from IEEE Xplore. Restrictions apply.

multiplexing, code-division multiplexing, and so on. It is certainly possible to combine different multiplexing technologies in various ways to create novel photonic switches. For example, if space-division switches are combined with wavelength-division multiplexing, a new class of switches known as wavelength-selective switches is obtained. One can also combine time-division multiplexing with space switches to create optical packet switches. Even if pure wavelength switching is employed, one can also create novel switches called wavelength converters.

Clearly, the feasibility of these novel switches depends very much on technological maturity. For a successful implementation, both multiplexing technology and switching technology must be reliable. So far, only a very few kinds of photonic switches and multiplexers show signs of maturity. The list includes spatial switches (mechanical and electro-optic-based), wavelength multiplexers, and wavelength-selective switches. Even though a lot of the emphasis has been put on wavelength conversion and fast optical packet switching recently, success has yet to be claimed.

Table 1 compares various optical switching technologies with respect to switching speed, switch size, technological maturity, and commercial availability. Only a few kinds of optical switches that are of current interest have been listed: spatial switches, wavelength-selective switches, wavelength-converting switches, and fast optical packet switches.

From the list, it can be concluded that while all-optical networking is a very remote goal which requires substantial development efforts in optical technology, there exist mature technologies which could be used intelligently to reduce the complexity of the entire network by an order of magnitude. For example, large mechanical spatial switches (with millisecond switching speeds) are available. Combined with wavelength multiplexers and demultiplexers, it is possible to create wavelength-selective switches that can serve as a high-capacity backbone fault-tolerant routing network for SONET channels. Whether such a high-capacity optical routing layer is necessary is entirely a matter of economy.

Conclusion

Advances in transmission have far outstripped advances in processing. Subsequently, much of the information processing should be relegated to the edge of the network and at higher layers of a communication protocol.

Electronic ATM switching has been demonstrated technically and commercially to be the most economical and scalable switching methodology for multimedia communications. The simple ATM protocol devised enables electronics to handle essential protocol functions. SONET, an electro-optic interface methodology, further facilitates scalable transmission.

Beyond scalable transmission, wavelength-space-mechanical switching may allow total traffic carried by a wide area broadband network to scale up to hundreds of terabits per second, without the use of SONET switches with thousands of ports. Although wavelength division multiplexing has proven itself economical for long distance transmission, there is still a strong need for transparent switching via wavelength selection and conversion. Further traffic study is necessary to understand how wavelength routing and bandwidth provisioning could be efficiently achieved.

References

[1] C. Brackett, "Foreword, Is There an Emerging Consensus on WDM Networking?" *IEEE J. Lightwave Tech.*, vol. 14, June 1996, pp. 936–41.

[2] B. Hui, "Progress of Optical Networking and Technology in the USA.," *OFC '96 Tech. Digest*, San Jose, CA, 1996, pp. 239–40.

[3] N. Bergano and C. Davidson, "Wavelength Division Multiplexing in Long Haul Transmission Systems," *IEEE J. Lightwave Tech.*, vol. 14, June 1996, pp. 1299–1308.

[4] I. P. Kaminow *et al.*, "A Wideband All-Optical WDM Network," *IEEE JSAC*, vol. 14, 1996, pp. 780–99.

[5] J. Turner, "An Optimal Nonblocking Multicast Virtual Circuit Switch," *Proc. IEEE Infocom '94*, Toronto, Canada, June 1994, pp. 298–305. Information concerning the switch prototype can be found at www.arl.wustl.edu/~jst/gigaswitch/gigatech.html

[6] Information on commercial ATM switches of capacity up to 10 Gb/s can be found at the following URLs: http://www.stratacom.com, http://www.cisco.com, http://www.fore.com, http://www.synoptics.com, http://www.networks.digital.com, http://www.nortel.com, http://www.casc.com, http://www.newbridge.com.

[7] ATM switches of capacity beyond 100 Gb/s include AT&T GlobeView 200, NEC ATM Model 20, and Fujitsu FETEX 150.

[8] S. Okamoto, A. Watanabe, and K. Sato, "Optical Path Cross-Connect Node Artchitectures for Photonic Transport Network," *IEEE J. Lightwave Tech.*, vol. 14, June 1996, pp. 1410–22.

[9] J. Hui, *Switching and Traffic Theory for Integrated Broadband Networks*, Boston, MA: Kluwer, 1990.

[10] S. Butner and R. Chivukula, "On the Limits of Electronic ATM Switching," *IEEE Commun. Mag.*, this issue.

[11] Traffic statistics on major Internet hubs using ATM technologies are collected from the following http homepages:
MAE East: http://ext2.mfsdatanet.com/MAE/east.stats.html
MAE West: http://ext2.mfsdatanet.com/MAE/west.stats.html
Ameritech: http://nap.aads.net/~nap-stat/AADS.NAP/atm/.
The points on the plot represent the maximum 5-minute average on the first Monday of every month.

[12] R.F. Rey, Ed., *Engineering and Operations in the Bell System*, 2nd Ed., Bell Telephone Labs., 1983.

[13] T. Sawano *et al.*, "A High-Capacity Photonic Space-Division Switching System for Broadband Networks," *IEEE J. Lightwave Tech.*, vol. 13, 1996, pp. 335–40.

[14] Commercial products available from Mitsubishi and Lucent Technologies.

[15] F. Masetti *et al.*, "High Speed, High Capacity ATM Optical Switches for Future Telecommunication Transport Networks," *IEEE JSAC*, vol.14, 1996, pp. 979–98.

Biographies

JOSEPH Y. HUI [F] received the S.B., S.M., and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology in 1981, 1981, and 1983, respectively. After graduation he joined Bell Communications Research, Morristown, New Jersey after a brief period with Bell Laboratories in Murray Hill, New Jersey. He joined Rutgers University in 1989 as an associate professor in the Department of Electrical and Computer Engineering, where he is now a professor. He is also a professor in the Department of Information Engineering of the Chinese University of Hong Kong. His current research interests include multimedia networks, wireless information networks, and switching, protocol, traffic, and security issues for integrated broadband networks. Dr. Hui received the Presidential Young Investigator Award from the National Science Foundation, and the 1985 William Bennett Prize Paper Award from the IEEE Communications Society. His work on multiple access communications and switching networks resulted in several patents. He is the author of a book entitled *Switching and Traffic Theory for Integrated Broadband Networks*, Kluwer, 1990.

KWOK-WAI CHEUNG received the B.S. degree (1st Hon.) from the University of Hong Kong, in 1978, the M.S. degree from Yale University, New Haven, Connecticut, in 1981, and the Ph.D. degree from Caltech, Pasadena, CA, in 1987. He joined the Chinese University of Hong Kong in December 1992, and is currently a full professor in the Department of Information Engineering. In January 1996, he was appointed the founding director of the Information Networking Laboratories, established by the Faculty of Engineering, the Chinese University of Hong Kong. From 1987 to 1992, he was member of technical staff at Bell Communications Research. He has performed research in many areas of lightwave communications, including system and device technology, network architectures and protocols, local and metropolitan area networks, and loop distribution networks. He has authored or coauthored about 90 journal and conference publications, and holds five U.S. patents. He was awarded an Electronic Letters Premium in 1992.

Authorized licensed use limited to: Chinese University of Hong Kong. Downloaded on January 22,2022 at 09:06:25 UTC from IEEE Xplore. Restrictions apply.