

Techniques for Optical Packet Switching and Optical Burst Switching

Lisong Xu, Harry G. Perros, and George Rouskas, North Carolina State University

ABSTRACT

Wavelength-division multiplexing appears to be the solution of choice for providing a faster networking infrastructure that can meet the explosive growth of the Internet. Several different technologies have been developed so far for the transfer of data over WDM. In this article we survey two new technologies which are still in the experimental stage — optical packet switching and optical burst switching — and comment on their suitability for transporting IP traffic.

INTRODUCTION

Several different technologies have been developed for the transfer of data over wavelength-division multiplexing (WDM), such as broadcast-and-select, wavelength routing, optical packet switching, and optical burst switching. Broadcast-and-select networks have been extensively studied, and several prototypes have been developed. Wavelength routing networks have already been deployed and currently represent the most promising technology for optical networks. Optical packet switching and optical burst switching are still in the research phase.

An optical packet network consists of optical packet switches interconnected with fibers running WDM. The switches may be adjacent or connected by *lightpaths*. (A lightpath is a circuit-switched connection consisting of the same wavelength allocated on each link along the path. It may consist of different wavelengths along the path if converters are present.) The user data is transmitted in optical packets, which are switched within each optical packet switch entirely in the optical domain. Thus, the user data remains as an optical signal in the entire path from source to destination. No optical-to-electrical or electrical-to-optical conversions are required.

Optical burst switching is a technique for transmitting bursts of traffic through an optical transport network by setting up a connection and reserving resources end to end only for the duration of a burst.

Both optical packet switching and optical burst switching can be used to transport IP traffic and asynchronous transfer mode (ATM) traf-

fic. The problem of transporting IP packets over an optical packet-switched network has not been addressed yet in the open literature. Furthermore, as will be seen below, there are some technologically limiting factors that need to be overcome in order for optical packet switching to become a commercially viable technology. Optical burst switching, on the other hand, seems to be a more prominent technology, at least in the near term.

In this article we survey various optical packet switch architectures and burst switching techniques. We review optical packet and optical burst switching, and present our conclusions.

WDM OPTICAL PACKET SWITCHES

A WDM optical packet switch consists of four parts: the input interface, switching fabric, output interface, and control unit. The input interface is mainly used for packet delineation and alignment, packet header information extraction and packet header removal. The switch fabric is the core of the switch and is used to switch packets optically. The output interface is used to regenerate the optical signals and insert the packet header. The control unit controls the switch using the information in the packet headers. Because of synchronization requirements, optical packet switches are typically designed for fixed-size packets.

When a packet arrives at a WDM optical packet switch, it is first processed by the input interface. The header and payload of the packet are separated, and the header is converted into the electrical domain and processed by the control unit electronically. The payload remains an optical signal throughout the switch. After the payload passes through the switching fabric, it is recombined with the header, which is converted back into the optical domain at the output interface.

In the following, we briefly describe some issues of optical packet switches. For more information about synchronization and contention resolution, the reader is referred to [1].

PACKET CODING TECHNIQUES

Several optical packet coding techniques have been studied. There are three basic categories: bit serial, bit parallel, and out-of-band signaling.


Bit serial coding can be implemented using optical code-division multiplexing (OCDM), optical pulse interval, or mixed-rate techniques. In OCDM, each bit carries its routing information, while in the latter two techniques, multiple bits are organized into a packet payload with a packet header that includes routing information. The difference between the latter two techniques is that in optical pulse interval, the packet header and payload are transmitted at the same rate, whereas in the mixed-rate technique, the packet header is transmitted at a lower rate than the payload so that the packet header can easily be processed electronically. In bit parallel coding, multiple bits are transmitted at the same time but on separate wavelengths. Out-of-band signaling coding includes subcarrier multiplexing (SCM) and dual-wavelength coding. In SCM, the packet header is placed in an electrical subcarrier above the baseband frequencies occupied by the packet payload, and both are transmitted in the same time slot. In dual-wavelength coding, the packet header and payload are transmitted in separate wavelengths but in the same time slot.

CONTENTION RESOLUTION

Contention resolution is necessary in order to handle the case where more than one packet are destined to go out of the same output port at the same time. This is a problem that commonly arises in packet switches, and is known as *external blocking*. It is typically resolved by buffering all the contending packets, except one which is permitted to go out. In an optical packet switch, techniques designed to address the external blocking problem include optical buffering, exploiting the wavelength domain, and using deflection routing. Whether these will prove adequate to address the external blocking problem is still highly doubtful. Below we discuss each of these solutions.

Optical Buffering — Currently, optical buffering can only be implemented using *optical delay lines* (ODLs). An ODL can delay a packet for a specified amount of time, which is related to the length of the delay line. Thus, optical buffering is the Achilles' heel of optical packet switches. Delay lines may be acceptable in prototype switches, but are not commercially viable. The alternative, of course, is to convert the optical packet to the electrical domain and store it electronically. This is not an acceptable solution, since electronic memories cannot keep up with the speeds of optical networks.

There are many ways an ODL can be used to emulate an electronic buffer. For instance, a buffer for N packets with a FIFO discipline can be implemented using N delay lines of different lengths. Delay line i delays a packet for i time slots. A counter keeps track of the number of packets in the buffer. It is decreased by 1 when a packet leaves the buffer, and increased by 1 when a packet enters the buffer. Suppose that the value of the counter is j when a packet arrives at the buffer; then the packet will be routed to the j th delay line. Limited by the length of the delay lines, this type of buffer is usually small, and does not scale up.



■ Figure 1. An architecture with a space switch fabric.

Exploiting the Wavelength Domain — In WDM, several wavelengths run on a fiber link that connects two optical switches. This can be exploited to minimize external blocking as follows. Let us assume that two packets are destined to go out of the same output port at the same time. Then they can be still transmitted out, but on two different wavelengths. This method may have some potential in minimizing external blocking, particularly since the number of wavelengths that can be coupled together onto a single fiber continues to increase. For instance, it is expected that in a year there will be as many as 200 wavelengths/fiber.

Deflection Routing — Deflection routing is ideally suited to switches that have little buffer space. When there is a conflict between two packets, one will be routed to the correct output port, and the other to any other available output port. In this way, little or no buffer is needed. However, the deflected packet may end up following a longer path to its destination. As a result, the end-to-end delay for a packet may be unacceptably high. Also, packets will have to be reordered at the destination since they are likely to arrive out of sequence.

Below, we examine various optical packet switch architectures that have been proposed in the literature. Based on the switching fabric used, they have been classified into the following three classes: *space switch*, *broadcast-and-select*, and *wavelength routing*. For presentation purposes, we do not show the input/output interfaces and control unit used in these switches.

AN ARCHITECTURE WITH A SPACE SWITCH FABRIC

A space switch fabric architecture is shown in Fig. 1. The performance of this switch was analyzed in [2]. The switch consists of N incoming and N outgoing fiber links, with n wavelengths running on each fiber link. The switch is slotted, and the length of the slot is such that an optical

Note that the wavelength of an output port is not fixed, and varies with packets. Therefore the output interface is responsible for making it meet the requirement of the output signal.




Figure 2. The KEOPS switch with a broadcast-and-select fabric.

packet can be transmitted and propagated from an input port to an output optical buffer.

The switch fabric consists of three parts: optical packet encoder, space switch, and optical packet buffer. The optical packet encoder works as follows. For each incoming fiber link, there is an optical demultiplexer which divides the incoming optical signal to the n different wavelengths. Each wavelength is fed to a different tunable wavelength converter (TWC) which converts the wavelength of the optical packet to a wavelength that is free at the destination optical output buffer. Then, through the space switch fabric, the optical packet can be switched to any of the N output optical buffers. Specifically, the output of a TWC is fed to a splitter which distributes the same signal to N different output fibers, one per output buffer. The signal on each of these output fibers goes through another splitter which distributes it to $d + 1$ different output fibers, and each output fiber is connected through an optical gate to one of the ODLs of the destination output buffer. The optical packet is forwarded to an ODL by appropriately keeping one optical gate open and closing the rest. The information regarding to which wavelength a TWC should convert the wavelength of an incoming packet and the decision as to which ODL of the destination output buffer the packet will be switched to is provided by the control unit, which has knowledge of the state of the entire switch.

Each output buffer is an optical buffer implemented as follows. It consists of $d + 1$ ODLs, numbered from 0 to d . ODL i delays an optical packet for a fixed delay equal to i slots. ODL 0 provides zero delay, and a packet arriving at this ODL is simply transmitted out of the output port. Each ODL can delay optical packets on each of the n wavelengths. For instance, at the beginning of a slot, ODL 1 can accept up to n optical packets, 1/wavelength, and delay them for 1 slot. ODL 2 can accept up to n optical packets at the beginning of each time slot, and delay them for 2 slots. That is, at slot t , it can accept up to n packets (1/wavelength) and delay them for 2 slots, in which case these packets will exit at the beginning of slot $t + 2$. However, at the beginning of slot $t + 1$ it can also accept another batch of n optical packets. Thus, a maximum of $2n$ packets may be in transit within ODL 2; similarly for ODL 3 through d . Let c_i denote the number of optical

packets on wavelength i , where $i = 1, 2, \dots, n$. We note that these λ_i optical packets may be on a number of different ODLs. To insert an optical packet into the buffer, we first check all the c_i counters to find the smallest one, say c_s ; then we set the TWC associated with this optical packet to convert the packet's wavelength to λ_s , increase c_s by 1, and switch the optical packet to ODL c_s . If the smallest counter c_s is larger than d , the packet will be dropped.

ARCHITECTURES WITH A BROADCAST-AND-SELECT SWITCH FABRIC

In this section we describe two different architectures with a broadcast-and-select switch fabric. In these architectures, packets from all input ports, each on a different wavelength, are combined within the switch and are broadcast to all the output ports. Wavelength selectors are then used at each output port to select a wavelength, and consequently a packet, to be sent out the switch. This type of switch fabric lends itself to multicasting.

The KEOPS Switch with a Broadcast-and-Select Fabric — This switch [3] was developed as part of the European ACTS Keys to Optical Switching (KEOPS) project. Each input and output fiber carries only one wavelength, as shown in Fig. 2. Note that the wavelength of an output port is not fixed, and varies with packets. Therefore, the output interface is responsible for making it meet the requirement of the output signal. The switching fabric consists of three blocks: encoder, buffer, and selector. The wavelength encoder block consists of N fixed wavelength converters (FWCs), one per input, and a multiplexer. The buffer block consists of a splitter, K ODLs, and a space-switching stage implemented by means of splitters, optical gates, and combiners. Finally, the wavelength selector block consists of N wavelength channel selectors implemented by means of demultiplexers, optical gates, and multiplexers. These three blocks make up the broadcast-and-select switch fabric.

The switch is slotted. At the beginning of a time slot, each wavelength converter in the wavelength encoder block converts the wavelength of the incoming packet to a fixed wavelength. The output of the N converters is combined and then distributed through a splitter

into K different ODLs. Each ODL has a different delay which is an integer number of slots. That is, ODL i has a delay of i slots. The N optical packets are stored simultaneously to the K different ODLs. At the beginning of the next slot, a maximum of $K * N$ optical packets exit from the K ODLs and up to N of them are directed to their destination output ports without any collisions. This is achieved through a combination of splitters, optical gates, demultiplexers and multiplexers. Specifically, the output signal from each ODL goes through a splitter which distributes it over N outputs. We recall that this output signal consists of N multiplexed optical packets, one for each wavelength. The signal from output j of each splitter is directed to output port j . Since there are K such splitters, there are K such output signals, of which only one is selected and directed to output port j . This selected output signal is fed into a demultiplexer, which breaks it up into the N wavelengths, of which only one is transmitted out. The operation of this broadcast-and-select switch fabric is managed by a control unit.

In this switch, an optical packet consists of a header, payload, and guard time. The header may include information about the destination, payload type, priority, and so on. The payload is the user data. A guard time is used to allow for the setup time of the optical devices in the switch. It may be inserted between the header and the payload, or between two successive packets on the same wavelength. Mixed-rate coding is used. That is, the header is encoded at a low fixed bit rate (e.g., 622 Mb/s), and the payload rate may vary from a few hundred megabits per second to 10 Gb/s. Moreover, the packet length is fixed in time, not in the number of bits. That is, the duration of the packet is fixed (e.g., 1.64 μ s), but the size of the packet is variable. This packet format has two advantages. First, the processing speed of the logic in the WDM packet switches depends on the header rate, but not on the payload rate. Second, the buffering space in the WDM packet switches, realized by means of ODLs whose length is proportional to the time length of the packet to be stored, does not depend on the payload rate.

This switch architecture can be extended to the case where it has M input and output fibers, and each input and output fiber carries n wavelengths. This is achieved by demultiplexing the signal from each incoming fiber to the n wavelengths, and then treating the switch as if it has $n * M$ input wavelengths instead of N presented above. At the output side, each group of n wavelengths can be combined together through a combiner onto the same output fiber.

A Switch with a Broadcast-and-Select Fabric and Recirculation Buffer — This switch architecture was proposed in [4], and a modified version is shown in Fig. 3. The idea of using a recirculation buffer comes from an ATM switch known as the *starlite* switch. As in the previous switch architecture, there is a single wavelength for each input and output fiber, and the wavelength of an output port varies with packets. The broadcast-and-select switch fabric is implemented through a coupler which combines up to M input wavelengths and then distributes the combined signal to N tunable optical filters (TOFs) and M




Figure 3. A switch with broadcast-and-select fabric and recirculation buffer.

fixed optical filters (FOFs). Note that M is larger than N . The input to the coupler comes from N input wavelengths and M wavelengths which are part of the feedback process, explained below.

The switch is slotted, and is controlled by a control unit. At the beginning of each time slot, the control unit knows the destination output ports of the incoming optical packets from the input ports and the 1 time slot delay line. Accordingly, it instructs the TWCs at the input ports, the tunable optical filters at the output ports, and optical gates. Up to M optical packets are fed into the coupler; according to their destinations, up to N of them are passed through the TOFs and out to the output ports, and the remaining packets are recirculated through an ODL. The recirculated optical packets are fed back to the coupler at the beginning of the following slot.

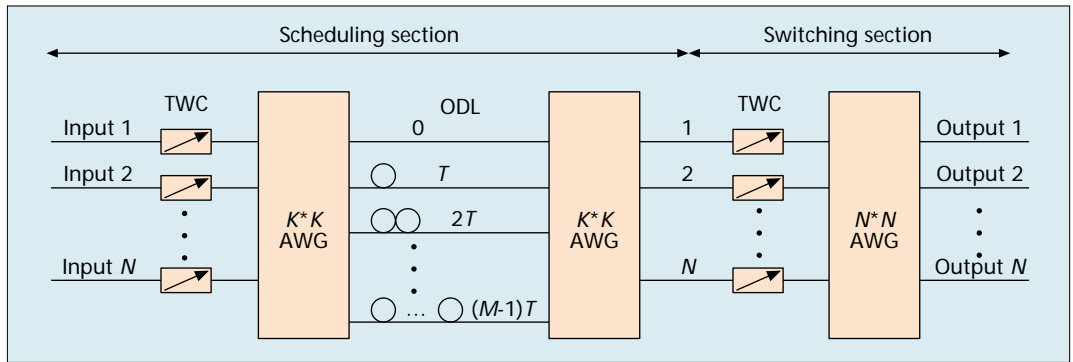
ARCHITECTURES WITH A WAVELENGTH ROUTING SWITCH FABRIC

In this section we describe three architectures based on wavelength routing switch fabrics. The switching procedure in these three architectures can be divided into two phases. In the first phase, packets are sent to ODLs for contention resolution; in the second phase, packets are routed to the correct output ports through the wavelength routing switch fabric.

An Input-Buffered Switch — This switch was proposed in [5] and is shown in Fig. 4. Each incoming and outgoing link carries a single wavelength. The wavelength of an output port varies with packets. The switch consists of the scheduling section and the switching section.

The scheduling section is used for contention resolution and is composed of N TWCs, one for each incoming wavelength, two $K \times K$ arrayed waveguide gratings (AWGs), and M ODLs, where $K = \max(N, M)$. An A WG is a wavelength routing device that can route optical signals from different input ports to different output ports based on their wavelength. The combination of all these optical devices provides for optical buffering of N

In this switch, an optical packet consists of a header, a payload, and a guard time. The header may include information about the destination, payload type, priority, and so on. The payload is the user data. A guard time is used to allow for the set-up time of the optical devices in the switch.



■ Figure 4. An input-buffered switch.

individual buffers, each of which has M positions. If there are available buffer spaces, a packet entering input i of the first AWG will appear at output i of the second AWG after a specified delay. The length of the delay is determined by the wavelength of the packet when it enters the first AWG. Specifically, each TWC converts the wavelength of an incoming optical packet so that the optical packet, when routed through the AWG, joins the ODL with the appropriate delay. The delay of an optical packet is selected using the following two rules: first, no two optical packets may appear at the same slot at the same switch output; second, no two optical packets may appear at the same buffer output at the same slot.

The switching section is used for switching optical packets to their destination output ports and is made up of an AWG and TWCs. The TWCs are used to assign the optical packet the right wavelength corresponding to the desired output port.

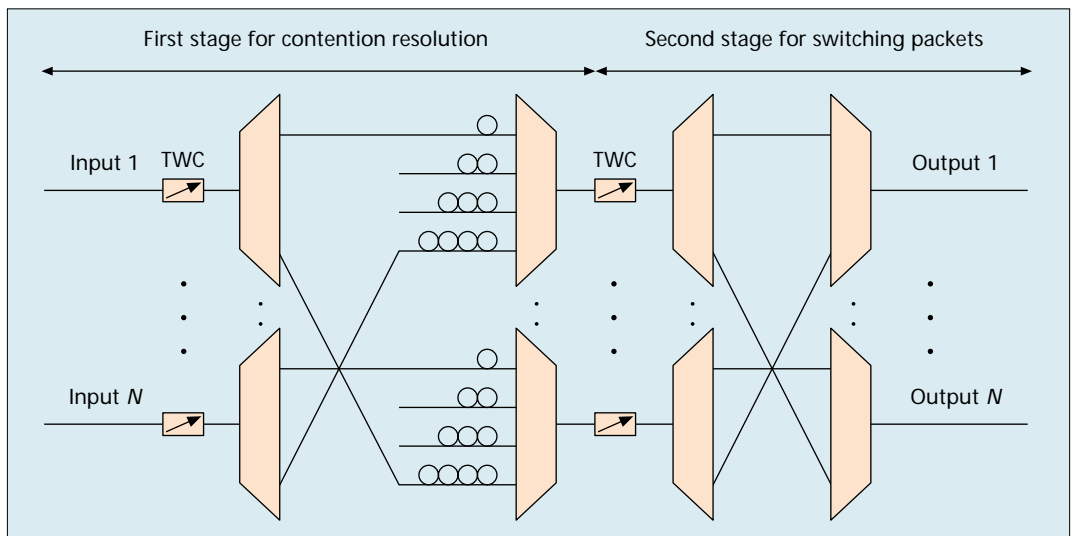
The switch suffers from head-of-line blocking, which is inherent in input buffering switches. For example, suppose that optical packet 1 in input i must be routed to output 1, while optical packet 2 behind optical packet 1 in input i must be routed to output 2. If optical packet 1 must be delayed for one time slot, optical packet 2 has to be delayed for at least one time slot due to the second rule, even though optical packet 2 goes to a different output port. However, if opti-

cal packet 1 has to be delayed by more than one time slot, optical packet 2 need only be delayed by 0 slot as long as packet 2 has no conflict at switch output 2.

In [5] the authors also proposed an output-buffering switch with a wavelength routing fabric.

An Input-Buffered Switch with a Distribution Network — This switch was developed as part of the KEOPS project [3], and is shown in Fig. 5. Each incoming and outgoing fiber carries a single wavelength. The wavelength of an output port varies with packets. The switch consists of two stages: contention resolution and switching. In the first stage, through the demultiplexer, each input port is connected with at least one ODL in each of N ODL sets. The TWC in the first stage decides to which ODL an optical packet will be sent. The second stage is used for switching optical packets to the correct output ports. Through the demultiplexer in the second stage, each ODL is connected to each output port. The TWC decides to which output port an optical packet will be sent.

Logically the first stage can be divided into two parts: distribution and input buffer. The distribution part distributes optical packets from the same input to different input buffers. Note that if we remove the distribution part (i.e., each demultiplexer can only connect to one ODL set), the switch becomes identical to the one



■ Figure 5. An input-buffered switch with a distribution network.

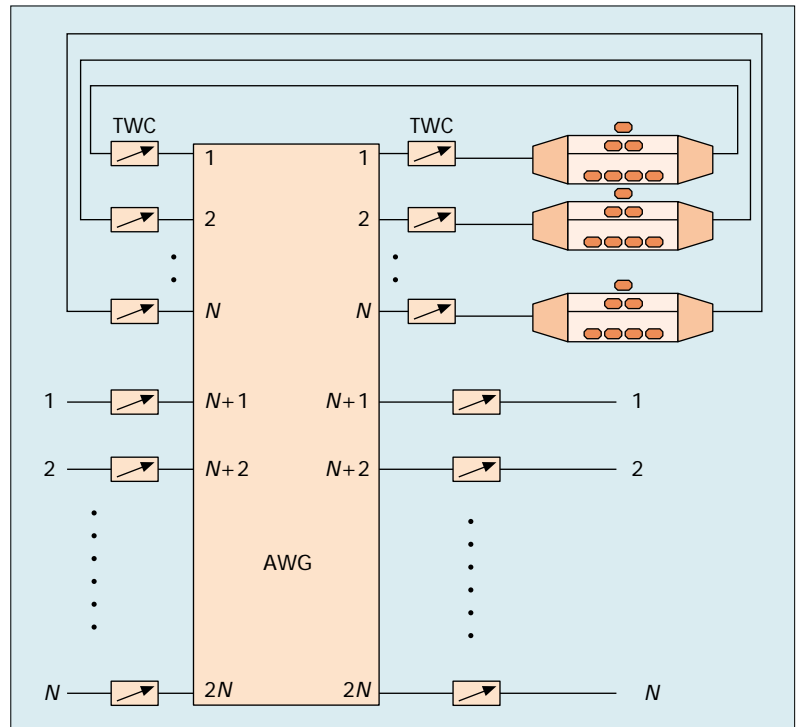
described above. The distribution part helps overcome the head-of-line blocking. Let us refer to the example of the two optical packets discussed in the previous subsection. In this case, regardless of how many time slots optical packet 1 must be delayed, it has no effect on the delay of optical packet 2, because the two optical packets could be routed to different ODL sets.

An optical packet arriving at the switch at time slot t is routed to the ODL with length d determined using the following three constraints. First, no other optical packet is scheduled to the same output port in time slot $t + d$. Second, no other optical packet is scheduled to the same TWC in the second stage in time slot $t + d$. Third, no optical packet from the same input port and to the same output port is scheduled in d with $d \geq d$. The ODL with the shortest delay satisfying these three constraints is selected.

In the WDM version of this switch, there are M incoming and outgoing fiber links, and each carries N wavelengths. N planes are located between M demultiplexers and M combiners connected to the incoming and outgoing fibers, respectively. Each plane is an $N \times N$ standard wavelength routing switch, as described above. However, in the WDM switch an outgoing fiber may carry more than one wavelength. The control part of the switch makes sure that there is no wavelength conflict in the combiner.

The WASPNET Switch — This switch was proposed as part of the Wavelength Switch Optical Packet Network (WASPNET) project [6]. The configuration of a WASPNET switch with single-wavelength inputs and outputs is shown in Fig. 6. It consists of a $2N \times 2N$ AWG, N sets of ODLs, and $4N$ TWCs. As in the previous two architectures, the switch can be divided into two phases. First, optical packets are routed to the ODLs to resolve contention, then routed to the desired output port. However, in this switch these two phases are implemented together by a $2N \times 2N$ AWG and N ODL sets. The $2N$ TWCs on the left of the AWG are used to select the AWG's output. The first N TWCs on the right of the AWG are used to select the correct ODLs for the optical packets that will be recirculated. The other N TWCs are used to convert optical packets to the wavelengths required by the switch output interface, because there are more wavelengths inside the switch than the incoming and outgoing wavelengths. One advantage of this switch is that it can support optical packet priorities. That is, after leaving the delay line, an optical packet may be delayed again because of preemption by a higher-priority optical packet.

The WDM version of this switch is made up of demultiplexers, combiners, and multiple planes of wavelength routing switch fabrics, each with single-wavelength inputs and outputs. It has N input and N output fibers, each with n wavelengths. There are n planes, each corresponding to one of the n wavelengths. For example, wavelength i in each input fiber is always demultiplexed to plane i . In view of this, the inputs of each plane have the same wavelength. However, different wavelengths may appear at the output of each plane. In one time slot, the switch allows multiple optical packets to leave not only from



■ Figure 6. The WASPNET switch.

the same output of the WDM switch, but also from the same output of a single plane. An $N \times N$ AWG is inserted between the N TWCs ($N + 1, \dots, 2N$) on the right of the $2N \times 2N$ AWG and the output of a plane for this function. Now the N TWCs carry out more functions. For instance, they make a final routing decision in addition to assigning optical packets to the wavelengths required by the switch output interface. The control part of the switch makes sure that there is no wavelength conflict in the combiner.

OPTICAL BURST SWITCHING

Optical burst switching is an adaptation of an International Telecommunication Union — Telecommunication Standardization Sector (ITU-T) standard for burst switching in ATM networks, known as *ATM block transfer* (ABT). There are two versions of ABT: *ABT with delayed transmission* and *ABT with immediate transmission*. In the first case, when a source wants to transmit a burst, it sends a packet to the ATM switches on the path of the connection to inform them that it wants to transmit a burst. If all the switches on the path can accommodate the burst, the request is accepted and the source is allowed to go ahead with its transmission. Otherwise, the request is refused, and the source has to send another request later. In ABT with immediate transfer, the source sends the request packet, and then immediately after it, without receiving a confirmation, transmits its burst. If a switch along the path cannot carry the burst due to congestion, the burst is dropped. These two techniques have been adopted to optical networks. The *tell-and-go* (TAG) scheme [7, 8] is similar to the ABT with immediate transmission, and the *tell-and-wait* (TAW) scheme [8] is simi-

In the tell-and-go scheme, the source transmits the control packet and immediately after it transmits the optical burst. In this scheme, it may be necessary to buffer the burst in the optical burst switch, until its control packet has been processed.

lar to ABT with delayed transmission. An intermediate scheme known as *just-enough-time* (JET) was proposed in [9].

An optical burst switching network consists of optical burst switches interconnected with WDM links. An optical burst switch transfers a burst coming in from an input port to its destination output port. Depending on the switch architecture, it may or may not be equipped with optical buffering. The fiber links carry multiple wavelengths, and each wavelength can be seen as a channel. A burst is dynamically assigned to a channel. The control packet associated with a burst may also be transmitted over a channel, or over a nonoptical network. The burst may be fixed to carry one or more IP packets. Currently, OBS networks do not exist. An optical burst switch architecture is described in [10]. In this article various algorithms for scheduling bursts within an optical burst switch are also discussed.

In the TAG scheme, the source transmits the control packet and immediately after transmits the optical burst. In this scheme, it may be necessary to buffer the burst in the optical burst switch until its control packet has been processed. In the JET scheme there is a delay between transmission of the control packet and transmission of the optical burst. This delay can be set to be larger than the total processing time of the control packet along the path. This way, when the burst arrives at each intermediate node, the control packet has been processed and a channel on the output port has been allocated. Therefore, there is no need to buffer the burst at the node. This is a very important feature of the JET scheme, since optical buffers are difficult to implement. A further improvement of the JET scheme can be obtained by reserving resources at the optical burst switch from the time the burst arrives at the switch, rather than from the time its control packet is processed at the switch.

In [9] a variation of JET was proposed which supports quality of service. Specifically, two traffic classes were defined: real-time and non-real-time. A burst belonging to the real-time class is allocated higher priority than a burst belonging to the non-real-time class by simply using an additional delay between transmission of the control packet and transmission of the burst. The effect of this additional delay is to reduce the blocking probability of the real-time burst at the optical burst switch.

CONCLUSIONS

Optical packet and optical burst switching techniques are currently at an experimental stage. The lack of commercially viable optical buffering technology imposes a constraint on the commercial development of optical packet switches. Also, it is not clear how such switches can be deployed to carry IP traffic. For instance, what would be an ideal packet size, and how much memory is required in an optical switch are questions that still need to be addressed.

Burst switching seems to be a more commercially viable technology than optical packet switching in the near future, if it can be designed to not require optical buffering. A possible first application of burst switching could be in a metropolitan

ring which can be used to connect devices such as IP routers and head-ends of cable access networks. The performance of the various burst switching schemes, as well as the size of a burst need to be studied in the light of IP traffic.

As optical devices mature and become cost-competitive, it is conceivable that by the end of the decade optics will play a much more significant role in packet switching technology than today, with many switching and routing functions incorporated into the optical domain.

REFERENCES

- [1] S. Yao, B. Mukherjee, and S. Dixit, "Advances in Photonic Packet Switching: An Overview," *IEEE Commun. Mag.*, Feb. 2000, pp. 84-94.
- [2] S. L. Danielsen *et al.*, "Analysis of a WDM Packet Switch with Improved Performance Under Bursty Traffic Conditions Due to Tunable Wavelength Converters," *J. Lightwave Tech.*, vol. 16, no. 5, May 1998, pp. 729-35.
- [3] C. Guillemot *et al.*, "Transparent Optical Packet Switching: The European ACTS KEOPS Project Approach," *J. Lightwave Tech.*, vol. 16, no. 12, Dec. 1998, pp. 2117-33.
- [4] G. Bendeli *et al.*, "Performance Assessment of a Photonic ATM Switch Based on a Wavelength Controlled Fiber Loop Buffer," *OFC '96 Technical Digest*, OFC, 1996, pp. 106-7.
- [5] W. D. Zhong and R. S. Tucker, "Wavelength Routing-Based Photonic Packet Buffers and Their Applications in Photonic Packet Switching Systems," *J. Lightwave Tech.*, vol. 16, no. 10, Oct. 1998, pp. 1737-45.
- [6] D. K. Hunter *et al.*, "WASPNET: A Wavelength Switched Packet Network," *IEEE Commun. Mag.*, Mar. 1999, pp. 120-29.
- [7] E. Varvarigos and V. Sharma, "The Ready-to-Go Virtual-Circuit Protocol: A Loss-Free Protocol for Multigigabit Networks Using FIFO Buffers," *IEEE/ACM Trans. Net.*, vol. 5, Oct. 1997, pp. 705-18.
- [8] I. Widjaja, "Performance Analysis of Burst Admission-Control Protocols," *IEEE Proc. Commun.*, vol. 142, Feb. 1995, pp. 7-14.
- [9] C. Qiao and M. Yoo, "Optical Burst Switching (OBS) - A New Paradigm for an Optical Internet," *J. High Speed Networks*, vol. 8, 1999, pp. 69-84.
- [10] J. Turner, "Terabit burst switching," *J. High Speed Networks*, vol. 8, 1999, pp. 3-16.

BIOGRAPHIES

LISONG XU (lxu2@eos.ncsu.edu) obtained a B.Sc. degree in computer science with first class honors in 1994, and an M.Sc. degree in computer science in 1997, both from University of Science and Technology Beijing, China. He is currently pursuing his Ph.D. in the Computer Science Department of North Carolina State University. His research interests include protocols for optical networks and quality of service issues in high-speed networks.

HARRY G. PERROS [M'87, SM'97] (hp@eos.ncsu.edu) received a B.Sc. degree in mathematics in 1970 from Athens University, Greece, an M.Sc. degree in operational research with computing from Leeds University, England, in 1971, and a Ph.D. degree in operations research from Trinity College Dublin, Ireland, in 1975. From 1976 to 1982 he was an assistant professor in the Department of Quantitative Methods, University of Illinois at Chicago. In 1979 he spent a sabbatical term at INRIA, Rocquencourt, France. In 1982 he joined the Department of Computer Science, North Carolina State University, as an associate professor, and since 1988 he has been a professor. During academic year 1988-1989 he was on sabbatical, first at BNR, Research Triangle Park, North Carolina, and subsequently at the University of Paris 6, France. Also, during academic year 1995-1996 he was on sabbatical at Nortel, Research Triangle Park, North Carolina.

GEORGE N. ROUSKAS [S'92, M'95] (rousakas@eos.ncsu.edu) received a Diploma in electrical engineering from the National Technical University of Athens (NTUA), Greece, in 1989, and M.S. and Ph.D. degrees in computer science from the College of Computing, Georgia Institute of Technology, Atlanta, in 1991 and 1994, respectively. He joined the Department of Computer Science, North Carolina State University in August 1994, and has been an associate professor since July 1999.