Chapter 1

OPTICAL PACKET SWITCHING

George N. Rouskas, Lisong Xu

Department of Computer Science

North Carolina State University

Raleigh, NC 27695-7534

{rouskas,lxu2}@csc.ncsu.edu

Abstract The concept of optical packet switching (OPS) is emerging as an alternative to coarser-grained switching in the optical domain. Despite the significant technological challenges it faces, OPS holds the promise of a highly reconfigurable, bandwidth-efficient, and flexible optical layer. In this chapter we study some of the architectural and design issues for OPS networks, we examine a number of enabling technologies, and we discuss some of the ongoing research and experimental efforts.

Keywords: Optical packet switching, wavelength division multiplexing (WDM), switch architectures, contention resolution techniques

1. INTRODUCTION

Optical transmission and switching technologies based on wavelength division multiplexing (WDM) have been increasingly deployed in the Internet infrastructure over the last decade in order to meet the ever-increasing demand for bandwidth. Given that point-to-point WDM transmission technology is quite mature today, while optical switching technologies continue to evolve at a rapid pace, the result has been the creation of *opaque* optical networks in which the optical signal undergoes *optical-to-electrical-to-optical* (OEO) conversion or regeneration at each intermediate node in the network. More recently, two trends have emerged in the design and deployment of WDM networks. The first is towards increasing *transparency* in the network so as to eliminate electronic bottlenecks and enable the handling of a broad range of heterogeneous signals regardless of protocol formats, bit rates, or modulation. The second trend is towards *reconfigurability* in optical networks, such that bandwidth can be created in real time between end-users to accommodate dynamically changing traffic demands. These trends reflect the vision of a future network in which optical switching technology plays a central role and bandwidth is relatively abundant, inexpensive, and readily available to end-users.

The migration of switching functions from electronics to optics will be gradual, and will take place in several phases. Already, the first phase is underway in the form of wavelength routed networks which offer circuit switching services at the granularity of a wavelength. Due to their circuit-switched nature, wavelength routed networks can be built with commercially available optical switch technologies, such as MEMS cross-connects [Chu et al., 2002], which are still relatively slow with switch configuration times in the order of milliseconds. While wavelength routing represents a significant step in the direction of transparent and configurable optical networking, optical circuits tend to be inefficient for traffic that has not been groomed or statistically multiplexed; moreover, the circuit-switching model does not fit well within the Internet philosophy of packet switching.

The next phase in the switching evolution is likely to involve the more recent optical burst switching (OBS) paradigm [Qiao and Yoo, 1999; Baldine et al., 2002]. Because it attempts to minimize the need for header parsing and buffering at intermediate network nodes, OBS is widely viewed as a promising technology for supporting finer switching granularity in the optical domain. Since the unit of transmission and switching is a burst, which is the aggregation of a flow of data packets, OBS is more efficient than circuit switching when the sustained traffic volume does not consume a full wavelength. OBS technology is still in the stage of research and experimentation, but at least one proof-of-concept testbed has been operational for the last year [Baldine et al., 2003]; as optical switching speeds improve to microseconds or less, OBS networks are expected to become a reality within the next few years.

In the longer term, optical packet switching (OPS) [Blumenthal et al., 1999; O'Mahony et al., 2001; El-Bawab and Shin, 2002; Yao et al., 2001; Yao et al., 2000] promises almost arbitrarily fine transmission and switching granularity, evoking the vision of a bandwidth-efficient, flexible, data-centric all-optical Internet. The realization of this vision, however, faces significant challenges in that OPS requires practical, cost-effective, and scalable implementations of optical buffering and packet-level parsing. We also note that each of the three optical switching technologies (wavelength routing, OBS, and OPS) have important application domains; hence, rather than each technology replacing the previous one, it is highly likely that all three will coexist in the optical network of the future.

In this chapter, we discuss some of the critical issues in designing and implementing OPS networks. In Section 2., we describe the architecture of an OPS node, and we take a close look at the building blocks, and the corresponding optical technologies, for realizing such a node. In Section 3., we discuss switch fabric architectures and contention resolution schemes. In Section 4., we describe experimental efforts and testbeds, and we conclude the chapter in Section 5..

2. OPS NODE ARCHITECTURE

In Figure 1.1, we show the functional block diagram of a generic OPS node architecture. The architecture consists of a set of multiplexers and demultiplexers, an input interface, a space switch fabric with associated optical buffers (i.e., fiber delay lines) and wavelength converters, an output interface, and a switch control unit. Packets arriving on an input fiber are first demultiplexed into individual wavelengths and are then sent to the input interface. Each packet consists of the payload and an *optical header* which is used for routing in the optical domain; note that any network layer header (e.g., IP header) is considered part of the payload for optical routing purposes. Among other functions, the input interface is responsible for extracting the optical packet header and forwarding it to the switch control unit for processing. The switch control unit processes the header information, determines an appropriate output port and wavelength for the packet, and instructs the switch fabric to route the packet accordingly. In routing the packet, the switch may need to buffer it and/or convert it to a new wavelength. The switch controller also determines a new header for the packet, and forwards it to the output interface. When the packet arrives at the output interface, the new header is attached, and the packet is forwarded on the outgoing fiber link to the next node in its path.

OPS networks can be classified along several dimensions depending on how the above packet switching and header processing functions are implemented.

- *Synchronous vs. asynchronous switch operation.*
  In a synchronous OPS network [Guillemot et al., 1998], time is slotted, and the switch fabric at each individual node can only be reconfigured at the beginning of a slot. All packets in a synchronous network have the same size, and the duration of slot is equal to the sum of the packet size and the optical header length (plus appro-

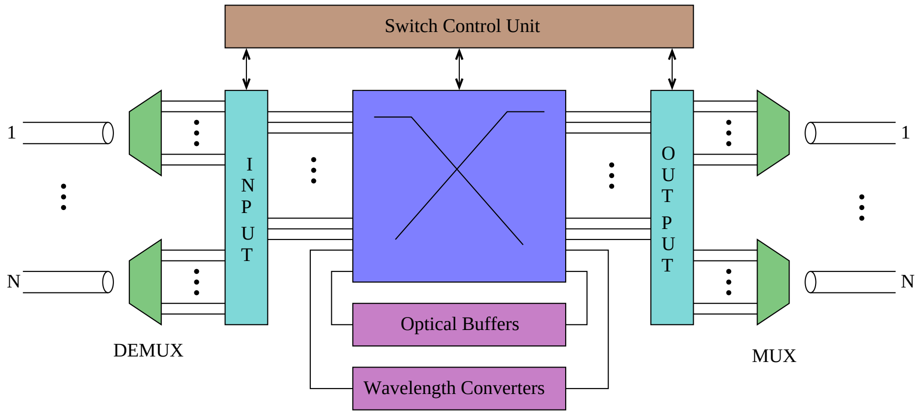

Figure 1.1 OPS node architecture

priate guard bands). Note that, due to variable link propagation delays, packets arriving at a node over different interfaces may not be aligned with the local clock. Therefore, it is the responsibility of the input interface (refer to Figure 1.1) to synchronize arriving packets and align them with switching time slots. Synchronous optical switching fabrics, much like their electronic counterparts, are easier to build and operate, hence synchronous OPS networks have received more attention from the research community.

In an unslotted network [Tancevski et al., 2000], packets are of variable size, switch operations may take place at any point in time, and there is no need to align arriving packets at the switch input. Unslotted OPS networks are more flexible and robust than slotted ones, and they do not require segmentation or reassembly at the edges of the network.

- *Electronic vs. optical control.*
  The optical packet header contains information for routing the packet in the OPS network. Today, the lack of fast, scalable, and robust optical bit-level processing technologies means that electronic processing of the header is the only practical approach. However, all-optical header processing is an active research area, and considerable progress has been made in some critical functions, such as all-optical label swapping (AOLS) [Blumenthal, 2001]. AOLS is a promising approach for routing packets in optics, in which packets are encapsulated with a small optical label as they enter the optical network. The packets are routed based on the information carried in the optical label, which is erased and rewritten

at each OPS node, while the elctronic packet header remains intact with the payload throughout the OPS network.

- *Optical header format.*
  There are two main approaches to optical header formatting: bit serial and subcarrier multiplexed; these are discussed in more detail in the next subsection.

- *Switch fabric architectures.*
  A wide variety of switch fabric architectures has been proposed for OPS networks, both for fixed-size and variable-size packets; due to their central role in the overall operation of an OPS node, we study a number of switch architecture designs in the next section.

- *Contention resolution strategies.*
  When two packets from different input port/wavelength pairs must be switched to the same output port/wavelength pair at the same time, contention arises. In this case, the switch controller and the switch fabric must employ some strategy to resolve the contention. Output port contention can be resolved in three dimensions: wavelength (using converters), time (using fiber delay lines), or space (using deflection routing); strategies that combine more than one dimensions are also possible. We discuss and compare various contention resolution strategies in Section 3.1.

  We note that, in addition to the data plane, contention is also possible in the control plane. Contention in the switch control unit may result in header loss or a significant delay such that the packet payload may precede the header; in either case, the packet has to be discarded. Therefore, proper buffer size dimensioning and efficient buffer management are of utmost importance; since, however, established techniques can be used for this purpose [Yao et al., 2003; Papadimitriou et al., 2003], we will not consider contention in the control plane any further.

2.1 ENABLING TECHNOLOGIES

Optical packet switching is still in its infancy compared to its electronic counterpart. We now discuss five functions of critical importance to the realization of practical OPS nodes.

Optical switch fabrics. The switch fabric at an OPS node must be capable of rapid reconfiguration on a packet-by-packet basis. At data rates of 40 Gbps and beyond, this requirement implies that switching times have to be on the order of a few nanoseconds. Other critical requirements include scalability of the technology to high port counts, low loss and crosstalk, and uniform operation across all signals independent of the path from input to output port; moreover, issues such as reliability, energy usage, and temperature independent operation are also important.

Today, most optical switch fabrics, including those based on opto-mechanical, thermo-optic, or acousto-optic technologies, are limited to switching speeds in the millisecond or microsecond range. Two promising technologies include semiconductor optical amplifier (SOA) switches and electro-optic lithium niobate ($LiNbO_3$) switches, both capable of switching speeds in the nanosecond rage. However, both technologies have limitations that must be overcome before it becomes possible to build high-performance, reliable, and cost-effective optical packet switches. For a recent comprehensive survey of optical switch fabric technologies, the reader is referred to [Papadimitriou et al., 2003].

Optical buffering. The lack of an efficient way to store information in the optical domain represents a major difficulty in the design of OPS nodes. Research has focused on ways of emulating electronic RAM capabilities through the use of fiber delay lines (FDLs) to delay optical signals [Hunter et al., 1998a; Tancevski et al., 2000; Hunter et al., 1998b; Chlamtac et al., 1996]. FDLs are fibers of fixed length, and can hold a packet for an amount of time determined by the speed of light and the length of the FDL. Hence, unlike electronic RAM, FDLs cannot store a packet indefinitely, and, once a packet has entered an FDL, it cannot be retrieved until it emerges on the other side. Furthermore, FDLs can be bulky and expensive, and introduce quality degradation to optical signals. As a result, the design of optical buffers that mitigate the effects of these limitations has emerged as an important research area for OPS. Among the important issues to be considered in designing FDL architectures include packet loss, cost, control complexity, packet reordering, and signal loss along the FDL. We discuss FDL buffer architectures in the next section.

Wavelength conversion. Wavelength conversion [Ramamurthy and Mukherjee, 1998; Elmirghani and Mouftah, 2000] is the ability to convert an optical signal on a given input wavelength to some other output wavelength. One of the main applications of wavelength conversion is as a mechanism for contention resolution that can dramatically improve the utilization of resources in an optical network, especially in highly dynamic traffic environments such as OPS. Consequently, wavelength converters have become integral to the design of optical buffer and switch architectures for OPS networks. Wavelength translation can be achieved by OEO conversion; however, all-optical wavelength conversion is desirable for OPS. Important features of all-optical converters include large wavelength span, fast setup time, high signal-to-noise ratio for cascadeability, and bit-rate transparency. All-optical converter approaches include the use of cross-gain modulation (XGM) or cross-phase modulation (XPM) in SOAs, and wave mixing techniques. Unfortunately, none of the existing techniques exhibits all the desired properties listed above; for a more detailed discussion and comparison of wavelength converter technologies, see [Elmirghani and Mouftah, 2000].

Packet delineation and synchronization. Packet delineation is required for both synchronous and asynchronous networks, and its purpose is to determine the beginning and end of the arriving packet. Current approaches perform delineation electronically as follows: a splitter taps a small amount of power from incoming packets and passes it to a bit-level synchronization circuit which locks the incoming bits in phase with the local clock in order to read the header information. Since this operation must be performed for each incoming packet, the circuit must be able to synchronize the header with its clock within a few bit times.

In addition to bit-level synchronization, OPS nodes in slotted networks must also synchronize incoming fixed-size packets to the local switching slots. This slot-level synchronization is accomplished by passing each incoming packet through a cascade of fiber delay lines and optical switches, in order to delay the packet by a sufficient amount of time for it to align with the beginning of a slot. This scheme introduces losses and crosstalk, resulting in a significant power penalty over long paths. A different strategy takes advantage of the fact that the propagation delay in a highly dispersive fiber depends on the signal wavelength. Each incoming packet is therefore passed through such a fiber, after its wavelength is first converted to achieve the desired delay.

Optical header format and processing. There are two main approaches to formatting the optical header associated with a packet [Blumenthal, 2001]. In the *bit-serial* approach, the header is transmitted serially on the same wavelength; a guard band is placed between the header and payload to allow for the removal and reinsertion of the header at intermediate OPS nodes. The second method uses *subcarrier multiplexing*, in which the header is situated slightly higher in the spectrum than the payload bandwidth, and is subcarrier multiplexed with the baseband payload. Both approaches have relative advantages and disadvantages, and both are being pursued in the lab.

As we mentioned earlier, currently, the processing of the header is performed electronically [Guillemot et al., 1998; Hunter et al., 1999]. All-optical header processing [Dorren et al., 2003] is an area of research that has received considerable attention, but the technology is still in the very early stages. In order to optically process headers, two functions have to be developed in optics: optical correlators to read a header, and all-optical flip-flop memory to store the header information. Currently, these functions have been demonstrated for headers containing only a few bits worth of information, limiting the switch size to only 1×2; for a review of optical header process techniques, see [Dorren et al., 2003]. One area in which significant progress has been made is in all-optical label swapping, which refers to techniques used to extract and replace the optical header without the need for OEO conversion of the payload. The interested reader is referred to [Blumenthal, 2001] for a description of all-optical label swapping technologies for both bit-serial and subcarrier multiplexed headers.

3. OPTICAL PACKET SWITCH ARCHITECTURES

3.1 CONTENTION RESOLUTION SCHEMES

As we discussed in the previous section, contention in the data path of an OPS node can be resolved using one of three methods or combination thereof: optical buffering, wavelength conversion, or deflection routing.

Optical buffering. The most straightforward method for resolving output port contention is to exploit the time dimension. Specifically, one of the contending packets (i.e., those arriving on the same wavelength at the same time and requesting the same output port) is routed through the switch fabric, while the rest are sent to an FDL. When the stored packet(s) emerge from the FDL, the whole process is repeated.

Similar to their electronic counterpart, optical buffers may be placed at the input, output, or both, of a packet switch. However, to compensate for the lack of a true "random access" property, a number of optical buffer arrangements have been proposed, such as single- or multi-stage FDLs, feed-forward or feed-backward connections, etc. Each of these arrangements can be used to implement a variety of packet switch architectures, and some representative examples are discussed in the next subsection.

Wavelength conversion. With this method, when two or more packets contend for the same output port and wavelength, the wavelength of all but one of the packets is converted to another wavelength, thus resolving the contention. If such a capability is available, then only when all wavelengths of an output port are busy does it become necessary to buffer contending packets. As a contention resolution method, wavelength conversion has some highly desirable properties in that it does not introduce delays in the data path and it does not cause packet resequencing.

Converters may be fixed or tunable, and they may be placed at the input and/or output of a packet switch; moreover, each port of the switch may be equipped with its own dedicated converter, or the converters may be shared by all ports. Consequently, a variety of switch architectures are possible depending on the availability and placement of converters.

Deflection routing. This method exploits the space dimension to resolve contention. Specifically, packets that lose the contention are sent to a different output port than the one requested, and hence may take a longer route to their destination. Deflection routing introduces delays in the data path and may cause packets to arrive out of order. However, it does not require additional hardware (e.g., FDLs or converters), unlike the previous two methods. On the other hand, the effectiveness of deflection routing as a contention resolution scheme depends on the traffic pattern and the density of the network topology.

The above three contention resolution schemes may be used in pure form, or they may be combined to implement more sophisticated strategies. For instance, optical buffering may be used along with either conversion or deflection routing to allow for more flexibility in resolving contention. The three pure schemes, along with the various combinations, make possible a wide spectrum of contention resolution methods that offer various tradeoffs of performance versus hardware cost and complexity. A comprehensive performance study of contention resolution methods can be found in [Yao et al., 2003]. The main finding of the study was that wavelength conversion offers the most performance benefits, and that the most efficient strategy is to combine conversion with limited buffering and selective deflection.

3.2 SWITCH FABRIC ARCHITECTURES

A wide variety of switch fabric architectures have been proposed for OPS. In general, we can classify the switch architectures as follows.

- *Single-stage vs. multi-stage switches.*
  Switches may consist of a single stage, or they may be built by appropriately cascading a set of smaller, single-stage switches [Hunter et al., 1998a; Papadimitriou et al., 2003]. Single-stage switches usually have a small number of input and output ports and small buffer capacity, and they are easy to implement and control. Due to cost (e.g., in terms of the amount of optical components required) or performance (e.g., in terms of power loss) considerations, switches with high port counts and/or large buffer capacity are usually implemented using multiple stages. Some considerations in building multi-stage switches include the number of smaller

switches required, the blocking characteristics of the architecture, and the degree of loss uniformity along the various paths from input to output.

- *Space vs. wavelength-routing vs. broadcast-and-select switches.*
  Space switch architectures are based on a non-blocking switch fabric, such as a crossbar, which is usually implemented using SOAs. A wavelength-routing switch [Chia et al., 2001] is usually based on arrayed waveguide gratings (AWGs) [McGreer, 1998], devices which implement a static permutation from input to output ports. A broadcast-and-select switch [Li et al., 2001] is usually based on a WDM passive star coupler. AWG-based switches require fewer optical components (especially SOAs) than either space or broadcast-and-select switches. On the other hand, it is straightforward to implement broadcast or multicast with a space or broadcast-and-select switch but not with a wavelength-routing switch. However, due to splitting losses, neither space nor broadcast-and-select switches may scale to large numbers of ports. We also note that, in a large, multi-stage switch, multiple technologies may be used simultaneously.

- *Feed-forward vs. feed-backward buffers (FDLs).*
  In a switch with feed-forward FDLs [Hunter et al., 1998a], a packet may be buffered only once: when such a packet emerges from the FDL after the specified delay, it is switched to an output port, and then leaves the switch. However, in a switch with feed-backward FDLs [Hunter et al., 1998a], a packet emerging from the FDL may be buffered multiple times by sending it (feeding it back) to the FDL; this situation may arise if the packet experiences contention again after emerging from the FDL. One advantage of a switch with feed-backward FDLs is that it can support priority scheduling of optical packets. That is, after leaving the FDL, an optical packet may be buffered again because of preemption by a later-arriving but higher-priority optical packet.

Due to the large number of different switch architectures that have been proposed for OPS, it is impossible to cover all of them in this chapter. In the following, we discuss some representative architectures to illustrate some of the possibilities in OPS switch design.

Single-stage space switch with feed-forward FDLs. In Figure 1.2, we present a single-stage space switch architecture with N ports, W wavelengths, and D FDLs per output port, similar to the one in [Danielsen et al., 1998a]; we can think of this as an output-queue architecture. Each incoming optical signal is first demultiplexed into the W wavelengths, and each wavelength is then converted to a wavelength that is free at the destination optical output buffer. The space switch fabric consists of splitters, optical gates, and combiners. The optical signal of each packet is split into ND identical signals, where N is the number of output ports and D is the number of FDLs per port. Once it is determined how long the packet has to be delayed in order to avoid output port collision, the packet is switched to the desired output port and corresponding FDL by closing the appropriate optical gate. It was shown in [Danielsen et al., 1998b] that with at least $W = 11$ channels per fiber, a low packet loss rate of 10^{-10} can be achieved even without the optical buffers. A more cost-efficient variant of this architecture was studied in [Eramo and Listanti, 2000]; in the new architecture, the optical buffers (FDLs) were eliminated and the tunable wavelength converters were shared among all incoming wavelengths. Obviously, the scalability of both switch variants is limited by the loss incurred by splitting each signal.

Single-stage broadcast-and-select switch with feed-forward FDLs. Figure 1.3 shows the architecture of a broadcast-and-select switch proposed as part of the European ACTS KEOPS project [Guillemot et al., 1998; Renaud et al., 1997]. The switch has N input and output ports, and it is equipped with D FDLs such that a packet can be delayed for an integer multiple of the slot time T, up to DT. The architecture in Figure 1.3 assumes that each input fiber carries only one wavelength that is different than the wavelengths carried by the other input fibers; hence the total number of wavelengths is N. The switch may be modified to handle multiple wavelengths per input fiber, by introducing an additional stage to demultiplex the input signal into individual wavelengths, and replicating the architecture shown in Figure 1.3. However, because of the power loss due to splitting, the product of the number of wavelengths times the number of ports cannot be high, limiting the scalability of the switch.

The switch operates as follows. First, the packets from all input ports are combined and distributed through a WDM passive star coupler to all D FDLs; note that in Figure 1.3, the multiplexer and splitter at the input play the role of the passive star coupler. At the output of the FDLs, optical gates are used to select the packets that have undergone an appropriate delay, of which only one packet is then selected and transmitted to the output port by yet another set of optical gates at the output ports. Note that performing broadcast or multicast is straightforward: all that is needed is for multiple output ports to select the same packet.

Single-stage wavelength routing switch with feed-backward FDLs. Figure 1.4 shows the wavelength routing switch architecture proposed

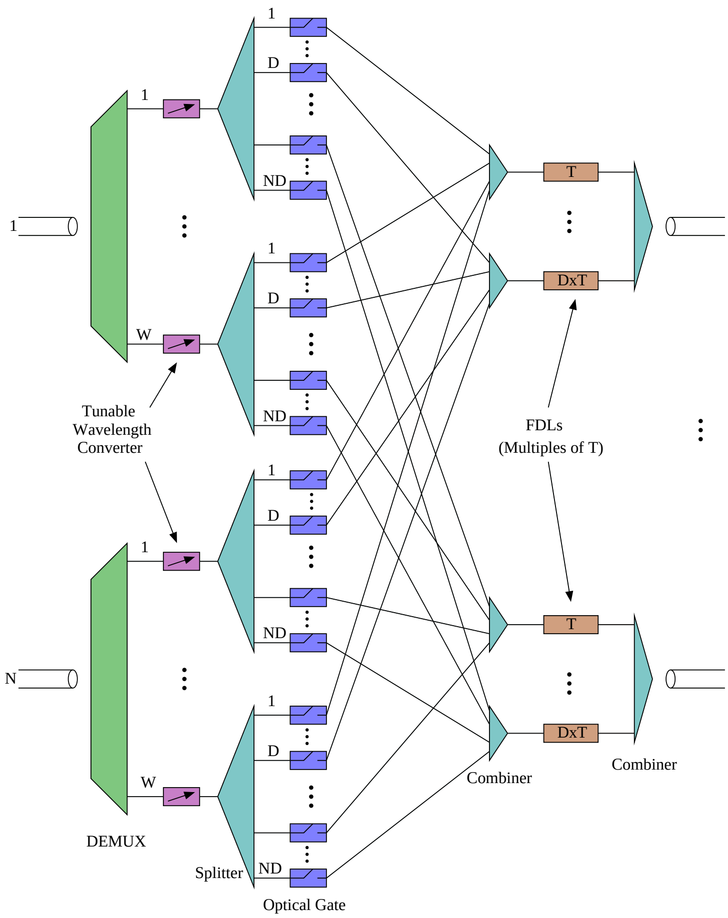

Figure 1.2 A single-stage space switch with N ports, W wavelengths, and D FDLs

as part of the WASPNET project [Hunter et al., 1999; O'Mahony et al., 2001]. The AWG is used to switch packets either to the correct output port, or to the appropriate FDLs in case of packet contention. In this switch, each input fiber carries only one wavelength. The switch can be extended to handle multiple wavelengths as follows. First, the optical signals are demultiplexed into individual wavelengths, and then they are fed to multiple planes, one for each wavelength. Each plane has the same switch architecture as in Figure 1.4. The space switch is then used to prevent wavelength contention when combining packets from different planes.

Multi-stage switch with feed-forward FDLs. If the packet size is variable, the performance of FDL buffers is poor [Tancevski et al.,

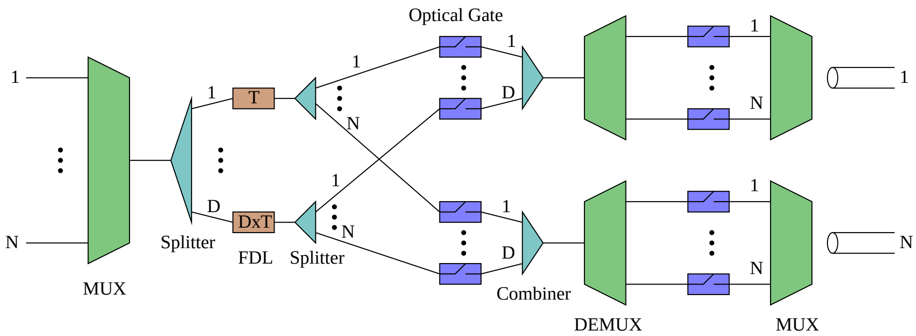

Figure 1.3 Broadcast-and-select switch with N ports, N wavelengths, and D FDLs

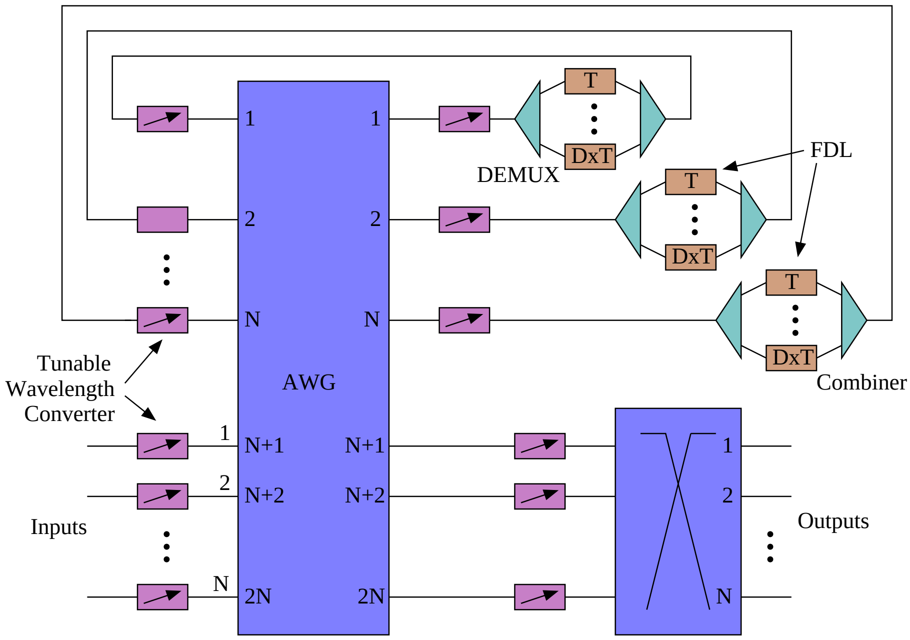

Figure 1.4 Wavelength routing switch with N ports

1999], and a low packet loss rate can be achieved only with a large buffer [Callegati, 2000]. Figure 1.5 shows one stage of the multi-stage switch proposed in [Callegati et al., 2002] to address this issue. The switch consists of three parts: an input part, a multi-stage FDL buffer, and an output part; the figure shows only stage i of the multi-stage buffer. At the input part, the set of signals from all input ports are converted to the set of W wavelengths used within the switch fabric. There is no wavelength conversion within the switch, so that a packet

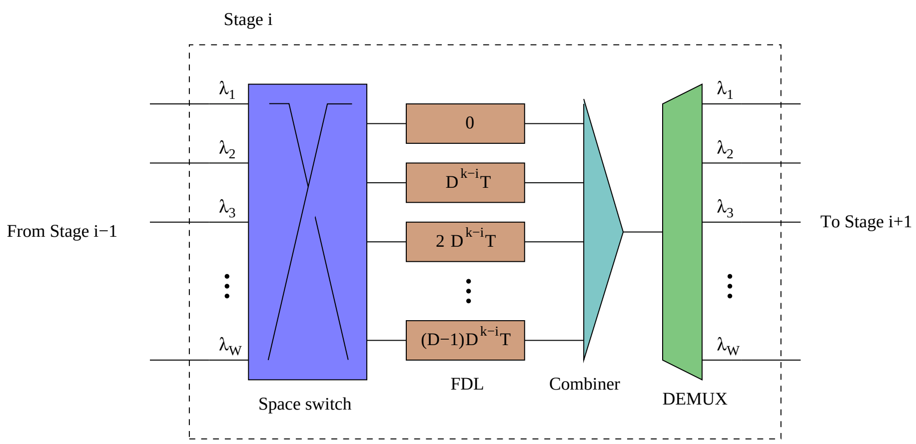

Figure 1.5 Stage i of multi-stage switch with W wavelengths and D FDLs per stage

assigned a particular wavelength at the input will emerge on the same wavelength at the output. The large buffer is implemented as multiple stages of FDLs. Let k represent the number of FDL stages, and let D denote the number of FDLs per stage. Then, for a given k and D, the FDLs are organized such that at stage $i, i = 1, \cdots, k$, the D FDLs produce delays equal to: $0, D^{k-i}T, 2D^{k-i}T, \cdots, (D-1)D^{k-i}T$, where T is the delay granularity. This arrangement makes it easy to find the indices of the FDLs to which a packet must be sent in order to realize some arbitrary delay through the switch. Finally, the output part of the switch (not shown in the figure) consists of a switch that connects any FDL and any wavelength to any of the output ports. A similar multi-stage switch with feed-forward buffers was proposed in [Hunter et al., 1998b].

Data vortex: multi-stage switch without buffers. The data vortex architecture was recently proposed in [Yang et al., 2001] to implement a large-scale switch with low latency. The objective was to minimize the number of switching operations and to eliminate the optical buffers, resulting in a cost-effective and practical implementation. The architecture of the data vortex switch is novel, consisting of a group of routing nodes which lie on a collection of concentric cylinders. The switch operates in synchronous mode, and it employs a number of sophisticated techniques such as a hierarchical routing structure and distributed traffic control. For details on the architecture and operation of the data vortex switch, the reader is referred to [Yang et al., 2001].

4. TESTBEDS AND EXPERIMENTAL EFFORTS

A number of experimental projects have been carried out in the area of optical packet switching over the last decade. We now briefly describe a representative set of these efforts; for details and a comprehensive description of the results of these projects, the reader is referred to the relevant publications.

The KEOPS (keys to optical packet switching) project [Guillemot et al., 1998; Renaud et al., 1997], was funded by the Advanced Communications Technology and Service (ACTS) program, a research program of the European Union, from 1994 to 1998. Partners of KEOPS included companies and universities from Europe. Building upon the results of the previous ATM optical switching (ATMOS) project, the objective of KEOPS was to analyze, develop, and demonstrate bit-rate-transparent all-optical packet switching for future all-optical networks. Two of the major results of the project were the demonstration of a 4×4 wavelength-routing switch operating at 2.5Gbps and a 16×16 broadcast-and-select switch operating at 10Gbps.

The WASPNET (wavelength switched packet network) project [Hunter et al., 1999; O'Mahony et al., 2001], funded by the U.K. Engineering and Physical Sciences Research Council (EPSRC) starting in July 1997, is a collaboration between three British Universities, the Universities of Essex, Bristol and Strathclyde, and three industrial partners, BT Laboratories, Marconi Communications, and Fujitsu Telecommunications. The objective of the project is to understand the advantages and potential of optical packet switching compared to the conventional electronic packet switching. A 8×8 wavelength-routing switch [Hunter et al., 1999], and a cascade of 14 switches operating at 2.5Gbps [Hunter et al., 1999] were demonstrated.

The DAVID (data and voice integration over WDM) project, was funded by the Information Society Technology (IST) Program, another research program of the European Union, from 1998 to 2002. DAVID was pursued by a fellowship of major operators, manufacturers, as well as leading universities, and research organizations from all over Europe. The main objective was to propose a packet-over-WDM network solution, including traffic engineering capabilities and network management, covering the entire area from MAN to WAN. For additional information and a list of relevant publications, see [DAVID Project, 2000].

All-optical label swapping (AOLS) [Blumenthal, 2001], is a new approach to routing packets in optics that combines optical network traffic engineering techniques with photonic packet switching technologies.

AOLS is being developed by the University of California at Santa Barbara in collaboration with several companies, and is funded by a DARPA NGI grant and by the DARPA sponsored Center for Multidisciplinary Optical Switching Technology (MOST). Using technology based on XPM wavelength converters, the project has demonstrated AOLS with variable-length packets at 80Gbps and optical labels at 10 Gbps [Blumenthal et al., 2003].

5. CONCLUDING REMARKS

The concept of optical packet switching, which seeks to replace the electronic switching functions by optical ones, represents a paradigm shift with the potential to revolutionize networking. Before a practical OPS network becomes a reality however, a number of technological issues must be addressed, as we discussed in this chapter. Nevertheless, given the ongoing research activities in this area, it is reasonable to expect that the key technological challenges will be eventually overcome, and some form of OPS will become possible within the next decade.

As optical technology advances and the OPS vision comes closer to reality, a number of other challenges will emerge. The constraints and new realities imposed by the optical layer and WDM transmission and switching technology will certainly affect our long-held assumptions regarding fundamental networking issues such as routing, control, packet transport, etc., which have been developed for mostly opaque electronic networks. As we carefully rethink these issues in the context of transparent OPS networks, protocol and algorithm design will certainly evolve to better accommodate the OPS technology, creating the need for proof-of-concept systems and testbeds of realistic size in which to experimentally study these new solutions. Migration strategies will also need to be developed for the network infrastructure to make the transition from electronic to optical packet switching, as well as from other forms of optical switching (e.g., wavelength routing or OBS) to packet switching. In the years ahead, as research in OPS shifts from components to integrated systems (i.e., OPS nodes and networks), we can expect a wide range of exciting research opportunities requiring interdisciplinary approaches that combine expertise in networking and optical engineering.

References

Baldine, I., Cassada, M., Bragg, A., Karmous-Edwards, G., and Stevenson, D. (2003). Just-in-time optical burst switching implementation in the ATDnet all-optical networking testbed. In *Proceedings of Globecom 2003*.

Baldine, I., Rouskas, G. N., and Stevenson, D. (2002). `JumpStart`: A just-in-time signaling architecture for WDM burst-switched networks. *IEEE Communications*, 40(2):82–89.

Blumenthal, D., Ikegami, T., Prucnal, P. R., and (editors), L. T. (1999). Special issue on photonic packet switching technologies, techniques, and systems. *Journal of Lightwave Technology*, 17(12).

Blumenthal, D. J. (2001). Photonic packet switching and optical label swapping. *Optical Networks Magazine*, 2(6):54–65.

Blumenthal, D. J., Bowers, J. E., *et al.* (2003). Optical signal processing for optical packet switching networks. *IEEE Communications Magazine*, 41(2):23–29.

Callegati, F. (2000). Optical buffers for variable length packets. *IEEE Communications Letters*, 4(9):292–294.

Callegati, F., Corazza, G., and Raffaelli, C. (2002). Exploitation of DWDM for optical packet switching with quality of service guarantees. *IEEE Journal on Selected Areas in Communications*, 20(1):190–200.

Chia, M. C., Hunter, D. K., and *et. al.* (2001). Packet loss and delay performance of feedback and feed-forward arrayed-waveguide gratings-based optical packet switches with WDM inputs-outputs. *Journal of Lightwave Technology*, 19(9):1241–1253.

Chlamtac, I. *et al.*, (1996). CORD: Contention resolution by delay lines. *IEEE Journal on Selected Areas in Communications*, 14(5):1014–1029.

Chu, P. B., Lee, S.-S., and Park, S. (2002). MEMS: The path to large optical crossconnects. *IEEE Communications*, 40(3):80–87.

Danielsen, S. L. *et al.* (1998a). Analysis of a WDM packet switch with improved performance under bursty traffic conditions due to tunable wavelength converters. *IEEE/OSA Journal of Lightwave Technology*, 16(5):729–735.

Danielsen, S. L. *et al.* (1998b). Optical packet switched network layer without optical buffers. *IEEE Photonics Technology Letters*, 10(6):896–898.

DAVID Project (2000). http://david.com.dtu.dk.

Dorren, H. J. S., Hill, M. T., and *et. al.* (2003). Optical packet switching and buffering by using all-optical signal processing methods. *Journal of Lightwave Technology*, 21(1):2–12.

El-Bawab, T. and Shin, J.-D. (2002). Optical packet switching in core networks: Between vision and reality. *IEEE Communications Magazine*, 40(9):60–65.

Elmirghani, J. M. H. and Mouftah, H. T. (2000). All-optical wavelength conversion: Technologies and applications on DWDM networks. *IEEE Communications Magazine*, 38(3):86–92.

Eramo, V. and Listanti, M. (2000). Packet loss in a bufferless optical WDM switch employing shared tunable wavelength converters. *Journal of Lightwave Technology*, 18(12):1818–1833.

Guillemot, C. *et al.* (1998). Transparent optical packet switching: The European ACTS KEOPS project approach. *Journal of Lightwave Technology*, 16(12):2117–2134.

Hunter, D. K., Chia, M. C., and Andonovic, I. (1998a). Buffering in optical packet switches. *Journal of Lightwave Technology*, 16(12):2081–2094.

Hunter, D. K., Cornwell, W. D., Gilfedder, T. H., Frenzen, A., and Andonovic, I. (1998b). SLOB: A switch with large optical buffers for packet switching. *Journal of Lightwave Technology*, 16(10):1725–1736.

Hunter, D. K. *et al.* (1999). WASPNET: A wavelength switched packet network. *IEEE Communications Magazine*, 37(3):120–129.

Li, B., Qin, Y., Cao, X., and Sivalingam, K. M. (2001). Photonic packet switching: Architecture and performance. *Optical Networks Magazine*, 2(1):27–39.

McGreer, K. A. (1998). Arrayed waveguide gratings for wavelength routing. *IEEE Communications Magazine*, 36(12):62–68.

O'Mahony, M. J., Simeonidou, D., Hunter, D. K., and Tzanakaki, A. (2001). The application of optical packet switching in future communication networks. *IEEE Communications Magazine*, 39(3):128–135.

O'Mahony, M. J. *et al.* (2001). An optical packet switched network (WASPNET)- concept and realisation. *Optical Networks Magazine*, 2(6):46–53.

Papadimitriou, G. I., Papazoglou, C., and Pomportsis, A. S. (2003). Optical switching: Switch fabrics, techniques, and architectures. *Journal of Lightwave Technology*, 21(2):384–405.

Qiao, C. and Yoo, M. (1999). Optical burst switching (OBS)-A new paradigm for an optical Internet. *Journal of High Speed Networks*, 8(1):69–84.

Ramamurthy, B. and Mukherjee, B. (1998). Wavelength conversion in WDM networking. *IEEE Journal Selected Areas in Communications*, 16(7):1061–1073.

Renaud, M., Masetti, F., Guillemot, C., and Bostica, B. (1997). Network and system concepts for optical packet switching. *IEEE Communications Magazine*, 35(4):96–102.

Tancevski, L., Ge, A., Castanon, G., and Tamil, L. S. (1999). A new scheduling algorithm for asynchronous, variable length IP traffic incorporating void filling. In *Proceedings of OFC*.

Tancevski, L., Yegananarayanan, S., Castagnon, G., Tamil, L., Masetti, F., and McDermott, T. (2000). Optical routing of asynchronous variable length packets. *IEEE Journal on Selected Areas in Communications*, 18(10):2084–2093.

Yang, Q., Bergman, K., Hughes, G. D., and Johnson, F. G. (2001). WDM packet routing for high-capacity data networks. *Journal of Lightwave Technology*, 19(10):1420–1426.

Yao, S., Dixit, S., and Mukherjee, B. (2000). Advances in photonic packet switching: An overview. *IEEE Communications Magazine*, 38(2):84–94.

Yao, S., Mukherjee, B., Yoo, S. J. B., and Dixit, S. (2003). A unified study of contention-resolution schemes in optical packet-switched networks. *Journal of Lightwave Technology*, 21(3):672–683.

Yao, S., Yoo, S. J. B., Mukherjee, B., and Dixit, S. (2001). All-optical packet switching for metropolitan area networks: Opportunities and challenges. *IEEE Communications Magazine*, 39(3):142–148.