

# The Application of Optical Packet Switching in Future Communication Networks

Mike J. O' Mahony,<sup>1</sup> Dimitra Simeonidou,<sup>1</sup> David K. Hunter,<sup>2</sup> and Anna Tzanakaki  
Ilotron Engineering Centre

## ABSTRACT

Telecommunication networks are experiencing a dramatic increase in demand for capacity, much of it related to the exponential takeup of the Internet and associated services. To support this demand economically, transport networks are evolving to provide a reconfigurable optical layer which, with optical cross-connects, will realize a high-bandwidth flexible core. As well as providing large capacity, this new layer will be required to support new services such as rapid provisioning of an end-to-end connection under customer control. The first phase of network evolution, therefore, will provide a circuit-switched optical layer characterized by high capacity and fast circuit provisioning. In the longer term, it is currently envisaged that the bandwidth efficiency associated with optical packet switching (a transport technology that matches the bursty nature of multimedia traffic) will be required to ensure economic use of network resources. This article considers possible network application scenarios for optical packet switching. In particular, it focuses on the concept of an optical packet router as an edge network device, functioning as an interface between the electronic and optical domains. In this application it can provide a scalable and efficient IP traffic aggregator that may provide greater flexibility and efficiency than an electronic terabit router with reduced cost. The discussion considers the main technical issues relating to the concept and its implementation.

## INTRODUCTION

The rapidly increasing bandwidth demand, driven by the Internet, has led to a paradigm shift in the telecommunications industry from voice-optimized to IP-centric networks. In this scenario, the role of synchronous digital

hierarchy/optical network (SDH/SONET) will diminish, and the optical transport network will directly provide a global transport infrastructure for legacy and new IP services. The utilization of optical networking employing dense wavelength-division multiplexing (DWDM) in conjunction with optical cross-connects (OXC) presents many new opportunities for supporting faster and more flexible provision of legacy and IP services. A major driver for realizing this evolution is the potential ability of such networks to provide fast automatic setup and tear-down of paths across the optical network, with the capability of supporting diverse client signals on the paths. The main focus, therefore, of today's optical network planning lies in implementing a dynamically reconfigurable optical transport layer based on fast OXC coupled with a suitable control and management architecture [1]. Thus, in the near future an optical transport network (OTN) will be realized capable of supporting large numbers of high-capacity optical channels, with bit rates on the order of 10–40 Gb/s. This model of the network is illustrated in Fig. 1. The diagram presents a possible OTN structure, which comprises the interconnection of a number of OXC in a mesh topology. Since each interconnecting fiber may support many wavelengths (e.g., > 100) and there may be many fibers (e.g., 32), the OXC require the capability to support the cross-connection of many thousands of wavelength channels. This OTN, therefore, will provide wavelength paths to clients such as IP routers, SONET/SDH network elements, and ATM switches, and Fig. 1 illustrates how the network might interconnect two IP routers. In addition to the switching hardware a control layer is necessary to set up the network path, and this normally interacts with the OXC controller to initiate switching within the OXC. A signaling channel between nodes ensures that each OXC has knowledge of the network resource status, paths available, and

<sup>1</sup> University of Essex

<sup>2</sup> University of Strathclyde

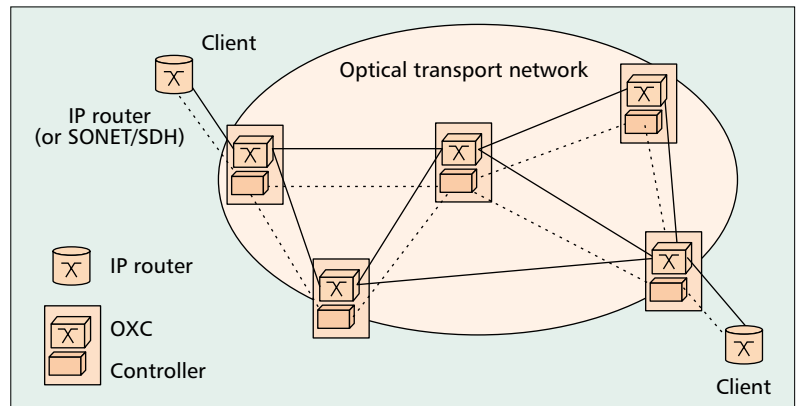
so on. Current research focuses on the use of distributed management schemes such as multi-protocol label switching (MPLS) to provide the control plane necessary to ensure fast path setup. In this type of application the label is the wavelength of the incoming signal; hence, the term multiprotocol lambda switching is more commonly used.

This dynamically reconfigurable optical transport network, therefore, will enable the fast allocation of high-capacity paths to link clients. Also, the pace of development is such that the technology (of transmission and switching) will support huge numbers of optical channels (wavelengths). It might therefore seem that in this future network bandwidth is not an issue, and optical circuit switching (the technique we have been discussing) will meet all future requirements. However, this is not the case for a number of reasons. The OTN, for example, offers granularity only at the wavelength level to clients. Thus, if the traffic source is bursty, the channel capacity may be under-used, which will have an impact on the dimensioning of the network and the size of the OXCs. This argument is particularly strong as the network moves to become data- rather than voice-centric. Economics will always demand that the network resources be used efficiently. A major advantage of electronic packet switching is its bandwidth efficiency and ability to support diverse services; hence, research is now focusing on bringing the packet switching concept into the optical domain, that is, optical packet switching (OPS).

In this article the use of OPS in the future network is discussed. First a general look at its application areas is considered, however it is believed that the first application will be as an edge router interfacing the electronic IP domain and the OTN. In the succeeding sections the technical and implementation issues relating to this concept are discussed.

## OPTICAL PACKET SWITCHING

Research into OPS has been conducted over a number of years [2–4]. Pure OPS, in which packet header recognition and control are achieved

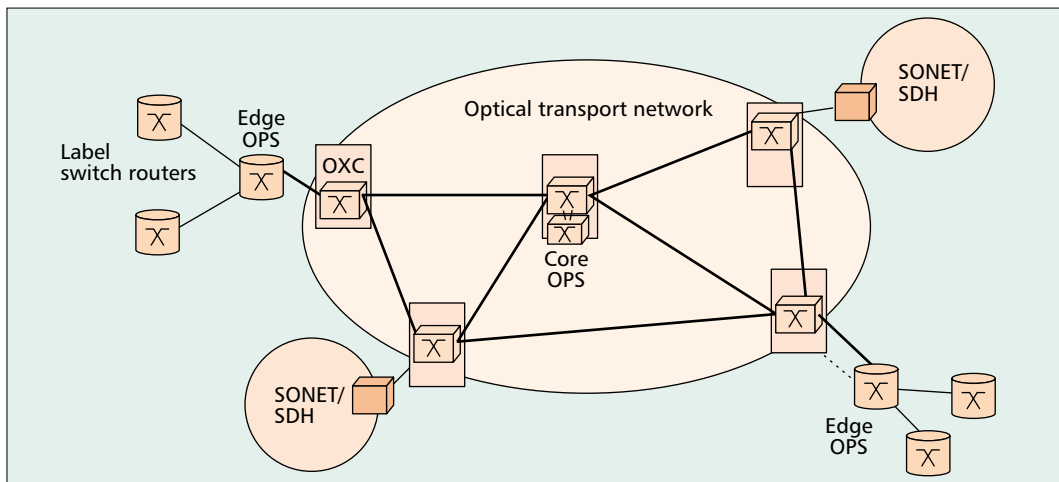


■ **Figure 1.** An optical transport network.

in an all-optical manner, is still many years away. For medium-term network scenarios, OPS using electronic control and header processing is more realistic; indeed, it is not clear what major advantages the all-optical approach has to offer over this opto-electronic approach. This article focuses on the approach in which the optical packet comprises an optical label (often realized using sub-carrier modulation techniques) attached to a payload, which may be of fixed or variable duration (other approaches, e.g., burst or flow switching, are not considered here). The client signal, such as IP packets, forms the payload, and the optical packet entity is routed through the network. Within an OPS, the packet header or label is read and compared with a lookup table. The payload is then routed to the appropriate output port with a new label attached (label swapping). An important feature is that the payload is transparently routed through the switch (i.e., stays within the optical domain), but the label processing and switch control are electronic. Some of these issues are discussed in more detail below.

## APPLICATIONS

The attractive feature of OPS is that it can appear as a natural evolution of the OTN. In particular, the OXCs developed for the OTN



■ **Figure 2.** Applications of OPS as core and edge routers.

*The first step toward optical data networking is the implementation of a network control plane, based on distributed label switched management principles such as the MPLS control model and associated with the OXC.*

can support an OPS network layer. Figure 2 illustrates a network comprising OXC and OPS elements. As shown, resources can be used in a number of ways. For example, some optical channels (wavelength paths) may interconnect high-capacity points that will fully utilize channel capacity, such as SDH rings. Other channels might be used to support optical packet transmission for efficient use of bandwidth, to either optimize resource utilization within the network or, for example, support an end-to-end point and click provisioning service where granularity may be an issue. Figure 2 therefore illustrates two key OPS application scenarios. One is the application as a core switch. Packets traveling through the network undergo switching at core nodes where ongoing route selection and label swapping take place. In this mode OPS maximizes utilization of the network resources, minimizing the total network capacity required, reducing the size of the OXCs. The second application, the main source of discussion in this article, is that of an edge router interfacing the electronic IP domain to the OTN. This is illustrated in Fig. 2, which shows the packet switch positioned as an edge router interfacing to both the OTN and IP domains. In this application the OPS provides a number of key functions required of the future OTN, as highlighted below.

## CONTROL PLANES

As discussed above, the first generation of OXCs will not perform packet-level processing. The entire traffic on any optical channel at an input port in an OXC is switched to an output port; thus, the optical channel supports continuous data. IP traffic, however, cannot be constructed as continuous data streams; therefore, there is a pressing requirement to develop the framework for a data/IP-aware optical transport.

At present, data/IP services are provided through networks that may include three or four different electronic multiplexing and switching layers (e.g., IP, frame relay, ATM, SONET). The multiplicity of layers produces inefficiencies, adds to the latencies of connections, and inhibits the provisioning of quality of service assurances. Worse, the layers are largely unaware of each other, causing duplication of transport protocols and management tasks.

The first step toward optical data networking is the implementation of a network control plane, based on distributed label-switched management principles such as the multiprotocol label switching (MPLS) control model and associated with the OXC. The functions of this control plane will initially be to establish and maintain optical paths within the network and, in the long term, to determine, distribute, and maintain state information associated with the OTN. This control plane will also be responsible for updating the information in the local switch controller (Fig. 1). As a result, the OXCs within the OTN will switch optical channels, in a similar way to label switched routers (LSRs) switch packets in an electronic IP network. LSRs per-

form packet-level operations using information carried on the labels attached to the data packets, while with OXCs the switching information is inferred from the wavelength (MPλS) or optical channel overhead.

In networking systems involving a number of data clients and OXCs, MPLS can provide a uniform control plane strategy in order to reduce the complexity of managing dissimilar networking systems. In these future network scenarios the question of where the boundary between the service and transport layer lies is still unanswered, but there is clearly a need to maintain topology and control isolation as well as to create an efficient interface between the optical transport and the service layers.

There are many reasons to separate network topologies and control, whether it is physical or logical for the OTN and the service layer. Some of the reasons are due to a number of important differences between electronic data routers and optical wavelength routers that necessitate special features to be implemented in the control plane. The first difference would be the bandwidth granularity, which is much coarser for an OXC than that for an IP router (wavelengths rather than packets). Because of the high-bandwidth nature of an optical connection, one would expect them to persist for a much longer duration and involve relatively infrequent connection requests when compared to per-packet routing operations. A further specific requirement for the control plane will be to maintain OTN infrastructure information in order to facilitate path selection for optical channels. Information will include fiber characteristics, amplifier positions, and signal evaluation data. This information can be collected through optical supervisory channels and optical channel overhead processing, and can be actively used for setting up optical paths and fault localization.

The most important reason, perhaps, for isolating the two layers is that they are likely to be under different administrative controls (or ownership) and policies. Under such circumstances the service provider who owns the OTN will wish to maintain full control of the network. Such an operator would not wish to give a client insight into the structure of the OTN layer since this is his/her business value.

Although the service provider does not wish to give the client knowledge of the OTN, there are client services that depend on having a view of the internal structure of the OTN. Three examples are suggested. The first involves connections diversely routed for provisioning and restoration purposes. The second involves a connection required at a future time, while the third involves being able to know which nodes can be reached via the OTN. Thus, network management features are required that allow limited internal OTN information to be accessed or manipulated by the client service layer in a manner that does not compromise the security of the operator's network.

Currently there are no router solutions that can satisfy all the above points and fit in a realistic future network solution able to carry efficiently a mixture of circuit- and packet-switched traffic into the OTN.

## THE OPTICAL PACKET SWITCH AS EDGE ROUTER AND AGGREGATOR

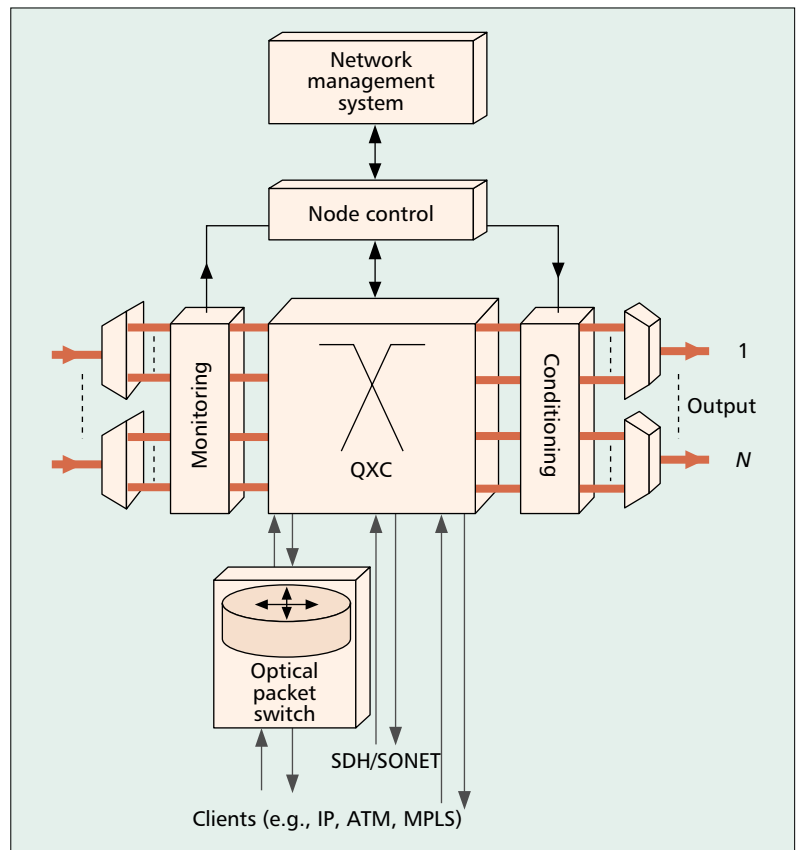
To overcome these limitations, an OPS solution can be used to facilitate efficient provisioning of packet services through the predominantly circuit-switched OTN infrastructure. The OPS will fit in a realistic network scenario where circuit- and packet-switched traffic will be transported together through the OTN. The optical packet switching functionality will then coexist with wavelength routing provided through the OXCs. In this case, fast switching will be provided for the packet traffic where granularity below the wavelength level is required, while slow wavelength switching and routing will be facilitated at the same time. Fast switching and packet traffic aggregation for efficient bandwidth utilization will mainly be performed at the edge of the network (the interface with the IP/ATM domain) where dynamic and fast wavelength allocation for packet traffic will be required. Under this scenario, the OPS router will be an edge network device, which will function as a topological and logical interface between the service and transport layers. The OPS router can directly interface with the OXC, which will make a set of static wavelength and fiber routes available to the OPS traffic. This is illustrated in Fig. 3, which shows an OXC making up a central switch fabric capable of interconnecting the demultiplexed input wavelength channels to the appropriate outgoing fibers. Interconnection is controlled through the management and control subsystems. The OPS is positioned in the add-drop ports of the OXC and accesses wavelength channels dedicated to packet switching.

The external electronic routers and OPS will handle the same granularity (per packet), which will lead to an integrated control plane between the IP and the OPS domains. At the same time, the OPS will maintain information on the configuration, the physical infrastructure, the topology and scale of the OXC transport. Therefore, the proposed OPS will be able to isolate the OTN from the service layer while interfacing fully with both layers:

- With the data/IP domain through integrated management control
- With the OTN by maintaining information on the configuration, the physical infrastructure, the topology and scale of the OXC transport

An additional benefit of the OPS will be due to the increased granularity over pure DWDM networks, which permits more efficient use to be made of the core network. One of the main disadvantages of an OTN is that there is currently no mechanism to provide direct access to the OTN with bandwidth granularity that is finer than a whole wavelength. Providing this finer granularity is central to creating a network that is efficient, from the perspective of the operator, and cost effective, for the operator's customer.

A schematic diagram of the OPS functionality as an edge aggregator/router is presented in Fig. 4. Here the OPS will provide an aggregation mechanism in the external OTN nodes that can accept packet type transport (i.e. IP and ATM)



■ Figure 3. Interfacing of the OPS with the OXC.

from a number of sources and map onto optical packets. These optical packets will be of variable length, which will be an integral multiple of a chosen time unit. The aggregating nodes will then map the optical packets onto appropriate wavelengths for transport over the OTN to de-aggregating nodes that can either be egress points from the network or intermediary nodes that further map the optical packets onto new wavelength paths. During this process, the OPS will run a protocol capable of discovering the OXC network topology, and thus will be able to combine aggregation with QoS provisioning within the OTN.

The optical router proposed here will provide a more scalable and efficient IP traffic aggregator compared with similar electronic Terabit router solutions. Furthermore, it will take full advantage of the capacity, scalability and functionality provided by the optical layer, a function that cannot be provided by an electronic router solution.


### REALIZATION ISSUES

The optical packet router will switch and buffer entities that may comprise multiple or single datagrams, or indeed only a part of one. To find the overall optimum packet transport solution for the optical edge aggregator, a number of issues need consideration, as described below.

#### THE OPTICAL PACKET

In order to reduce the number of entities that the switch must process per unit time, single or

The optical packet router will switch and buffer entities that may comprise multiple or single datagrams, or indeed only a part of one. To find the overall optimum packet transport solution for the optical edge aggregator, a number of issues need consideration.



■ Figure 4. OPS functionality as an edge aggregator/router.

multiple packets with the same destination and quality of service (QoS) class may be grouped together forming an optical packet at the edge of the network. The optical packet will be of variable length, which will be integral multiple of a unit length. While this reduces the complexity of the packet switches, it increases the complexity of the interface at the edge of the network, in fact the complexity of forming the optical packet is comparable to implementing some of the functions of an IP router. Great care must be taken when designing a packet scheduling algorithm for this type of switch, to ensure that the algorithm can be implemented in real-time by electronic control hardware. The optical packet switched network now looks very much more like a burst switched network [5], the major difference being that control information is still in-band. The implementation of the optical packet can be advantageous for the edge router application, where the optical router will perform and replace some or all of the Terabit router functionality.

With this approach, the header must also be read, and the label must be translated electronically, in the usual way. In an MPLS-based approach such as that considered here, the header translation hardware will search in a table for the label held in the packet header. The entry in the table for that particular label will contain the new label (which must then over-write the existing label in the header), and the output to which the packet must be forwarded. Label stacking is very difficult to implement in such an OPS since such an operation effectively involves changing the length of the header.

#### BUFFER MEMORY IMPLEMENTATION

To preserve an all-optical data path, it would be desirable to implement the buffer memory in the OPS optically. However, optical memory is in a relatively primitive state; there is no such thing as optical random access memory (RAM), and it is necessary to resort to fiber delay lines for memory. If these become unduly long (tens or hundreds of kilometers), they become very costly, bulky, and difficult to stabilize with respect to temperature. Here, a compromise is proposed where electronics and optics share the buffering. Optics is used for very short delays, which form the vast majority of storage, and electronics is used for longer delays. The amount of electronic memory, with its costly electrical-to-optical and optical-to-electrical interfaces, is reduced. If a packet must be delayed more than the longest optical delay, it is passed to the electronic memory.

This approach is particularly suitable for the edge aggregator because of the ability to make use of the electronic memory already in place to perform the electronic router functionality. In this case, although the switch will be able to operate without using the optical delay line memories, it can be shown that the optical memories can be used for the majority of the buffering (i.e., low delays) while only a fraction of the electronic memory is used for the higher delays, representing the minority of buffering.

Initial simulation studies have demonstrated the validity of this technique, as illustrated in Fig. 5a, which shows the probability that a randomly chosen byte stored within an output-buffered packet switch is experiencing a delay greater than a given value within the switch. In

an effort to mimic real Internet traffic, each output buffer is fed with self-similar traffic multiplexed from 404 sources with truncated Pareto on/off periods. Each burst coming from a source consists of packets generated according to a packet length distribution obtained from real traffic. The alpha parameter is 1.1, and the mean total load per output is 80 percent. For this type of traffic, these results apply for any output buffer experiencing a load of 80 percent, on any size of switch.

Suppose that optical delay line buffering handles all optical delays of 6750 bytes or less (i.e., a maximum fiber length of 1.08 km at 10 Gb/s); then, on average, only 10 percent of the bytes in memory at any time are experiencing a larger delay, and are buffered in electronic memory. If the maximum optical delay is increased to 11,500 bytes, this corresponds to a maximum fiber length of 1.84 km and only 1 percent of bytes in electronic memory. Hence there is a trade-off between electronic and optical memory use, which can lead to a significant reduction in electronic memory by employing only short lengths of shared delay line optical buffers.

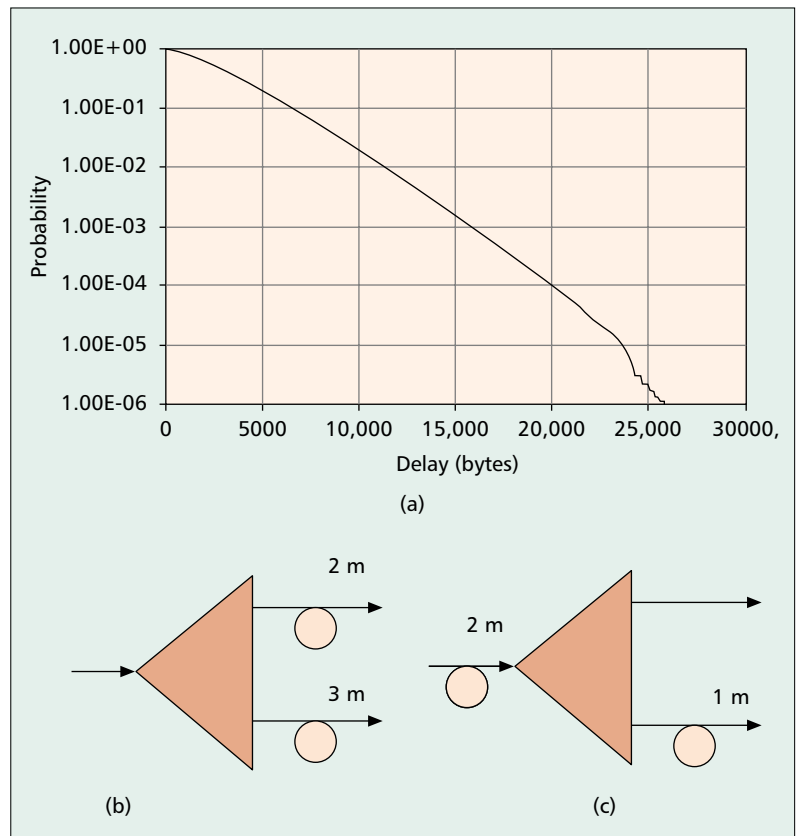
### SCHEDULING AND CONTROL

Since optical memory is implemented with delay lines and not RAM, the electronic scheduler for the architecture must direct the packets over the correct delay lines to make the architecture perform the same function as one constructed from RAM buffers. The packet scheduling algorithms for the transport solutions discussed above can be implemented using high-speed electronics, and must consider issues such as fairness, implementation of QoS classes, queue stability, and queue starvation. The trade-off between electronic and optical buffering must be determined based on cost considerations. Initial results demonstrate that using an OPS to interface with electronic routers can produce cost savings in the network.

Due to the bulkiness and expense of large amounts of delay line fiber, two techniques may be used to reduce the total length of fiber required; this impacts upon the control algorithm:

- Multiple packets on different wavelengths may pass along a specific delay path simultaneously. For example, the total length of fiber delay line required could be reduced by a factor of 16 by using one delay line with 16 wavelengths instead of 16 delay lines of equal length; in both cases the same number of wavelength converters are required (i.e., 16).
- The total length of fiber delay line memory can also be reduced by sharing fibers between different delay paths — a simple instance of this technique is shown in Fig. 5b, c [6]. This principle can be extended so that a large array of fiber delays can be replaced by multiple delay line stages, with a dramatic reduction in the amount of fiber required.

Packet scheduling algorithms should be amenable to parallel implementation in order to enable implementation on programmable gate arrays to run in real time. Also, the implementation must scale to large switches such as will be encountered in practice in future.

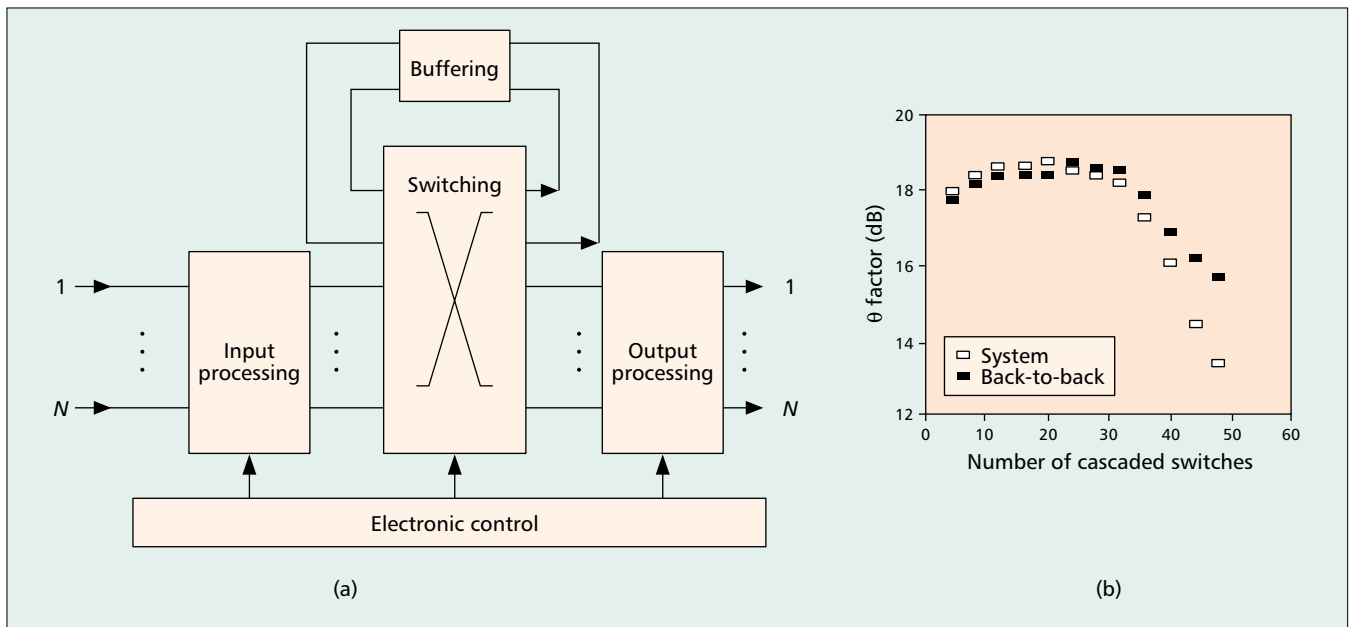


■ **Figure 5.** a) The probability that a randomly chosen byte is experiencing a delay greater than the given value in an output-buffered switch, for self-similar traffic, with a mean traffic level of 80 percent; b and c) an example of sharing delay lines. The subsystem in b) requires 5 m of fiber delay line, whereas only 3 m are required in c), because the smallest delay line has been moved in front of the splitter.

### SWITCH IMPLEMENTATION

A generic structure of the proposed packet switch consists of an input processing interface, a switching and buffering block, and an output processing module, as illustrated in Fig. 6a. The input interface performs delineation (i.e., identification of the packet start and end), packet format adaptation into the optical packet, classification into forward equivalent classes defined for the OTN, and electronic buffering. The switching and buffering blocks are responsible for the routing of the optical packets to the appropriate output ports and contention resolution, respectively, while the output interface is responsible for header reinsertion and per packet conditioning such as wavelength conversion to the appropriate OTN wavelengths, regeneration, and power equalization. The proposed architecture is based on a feedback buffering scheme to enable maximum utilization and sharing of the available buffers. The recirculating buffers used in this architecture offer the ability to support QoS classes via packet preemption. Header detection and processing are performed in the electronic domain.

Fast optical switching per packet can be performed using a switch matrix based on semiconductor optical amplifier (SOA) gates or electro-optic technology. However, both technologies are scalable only up to a limited switch



■ **Figure 6.** a) A schematic diagram of the generic structure of the optical packet switch; b) concatenation performance of the wavelength converter and AWG arrangement; back-to-back measurement shown for comparison.

dimension and require some form of synchronization at the input of the switch matrix. An alternative solution that enables fast transparent switching of individual packets enabling asynchronous operation of the switch matrix is based on tunable wavelength converters followed by a wavelength routing device such as an arrayed waveguide grating (AWG). In this case, routing of the packets to the required output ports of the switch is performed by controlling the wavelength of the incoming packets through the input wavelength conversion stage and subsequent transmission through the AWG. Optical wavelength conversion is performed through SOA-based converters using either cross-gain modulation or cross-phase modulation techniques. Using either of the two schemes, a continuous wave (CW) source is needed, and in the case of tunable wavelength conversion this source is required to be either a fast tunable laser or a switchable laser array. The tuning speed of the converter is then determined by the tuning speed of the CW signal, which can be as fast as a few nanoseconds; thus, the switching speed will also be in the nanosecond regime. The overall switch matrix scales with the dimension of the AWG router, which currently can be as high as 128 x 128. This approach was evaluated in project WASPNET [7, 8]. The concatenation performance of this configuration was evaluated through recirculating loop experiments [9], and Fig. 6b shows measured Q factor for both back-to-back and system (i.e., AWG and wavelength converter) configurations. The results demonstrate penalty-free operation for up to 25 cascaded nodes.

The buffering functionality is provided through a combination of electronic and optical buffering. The wavelength agility offered using wavelength conversion on a per packet basis enables statistical multiplexing at the fiber bandwidth capacity level. Tunable wavelength con-

verters may significantly reduce the buffering requirements by appropriately wavelength translating optical packets so that they can be stored within the same fiber delay line. This not only simplifies the buffering schemes, but also has the advantage of suppressed transfer delay and packet delay variation due to the reduction of the depth of the optical buffers [8].

#### SUMMARY AND CONCLUSIONS

This article presents a novel and efficient solution for a fully data/IP (i.e., IP and ATM) aware optical transport network. The current proposal is to use an optical packet switching technology in order to:

- Reduce the number of network layers, to simplify network management software and remove associated transport overheads
- Offer efficient traffic aggregation and finer service granularity (compared to current wavelength switching technology), thereby improving OTN utilization.
- Facilitate dynamic QoS-based provisioning through the OTN
- Provide domain isolation between the service and OTN networks

The proposed OPS router will be a predominantly edge network element, and will function as a topological and logical interface between the service and transport layers. The OPS will have the capability to aggregate the traffic from a large number of IP routers and ATM switches at the edge of the network, and groom them based on QoS to a number of dedicated wavelengths in the OTN. This edge-OPS device has the potential to replace terabit routers as a more scalable, efficient, and future-proof solution in the marketplace because it has the ability to provide a more scalable and efficient IP traffic aggregator than similar electronic router solutions. Furthermore, it will take full advantage of the capacity and functionality provided by the

optical layer, a function that cannot be provided by an electronic router solution. Although this article focuses on the application of OPS as an edge router, it is envisaged that core OPS nodes would also be used to ensure appropriate data transport consistent with minimizing the size of the OPS nodes. The core packet switch would perform label swapping while transparently transferring the data payload.

#### ACKNOWLEDGMENT

The authors wish to thank Dr. Robert D. Gardner of the University of Strathclyde for programming the simulator which produced Fig. 5a.

#### REFERENCES

- [1] N Ghani, S Dixit, and T. S. Wang "On IP over WDM Integration," *IEEE Commun. Mag.*, Mar. 2000, pp. 72–83.
- [2] D Chiaroni *et al.*, "A Novel Photonic Architecture for High Capacity ATM Switch Applications," *Photonics in Switching '95*, paper PTHC3, Salt Lake City, UT, Apr. 1995
- [3] D. K. Hunter and I. Andonovic: "Approaches to Optical Internet Packet Switching," *IEEE Commun. Mag.*, Sept. 2000, pp. 116–22.
- [4] A. Carena *et al.*, "OPERA: An Optical Packet Experimental Routing Architecture with Label Swapping Capability," *J. Lightwave Tech.*, vol. 16, no. 12, Dec. 1998, p. 2135.
- [5] 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.
- [6] R. A. Spanke, "Optical Time Division Switching Using Multiple Stages of Fiber Optic Delay Loops," Ph.D. thesis, Northwestern Univ., June 1994.
- [7] K. M. Guild *et al.*, "Cascading and Routing of 14 Optical Packet Switches: WASPNET," *ECOC 2000*, Munich, Germany, P4.8, Sept. 2000.
- [8] D. Hunter *et al.*, "WASPNET: A Wavelength Switched Packet Network," *IEEE, Commun. Mag.*, Mar. 1999, pp. 120–29.

- [9] A. Tzanakaki *et al.*, "Penalty-Free Concatenation of 25 Optical Cross-Connects Performing Reconfigurable Wavelength Routing," *ECOC '99*, Nice, France, Tu C3.5, Sept. 1999, pp. I-254–55.

#### BIOGRAPHIES

MIKE J. O'MAHONY [SM] (mikej@essex.ac.uk) has been a professor of telecommunications at the University of Essex since 1991. He previously held the position of head of Inland Systems Section at BT Laboratories where he was responsible for research into long-haul optical systems and networks. His current interests include the application of packet switching to future optical networks and he has authored over 200 papers in optical communications.

DIMITRA SIMEONIDOU (dsimeo@essex.ac.uk) is senior lecturer at the University of Essex and has over 10 years experience in the field of optical transmission and optical network design. Previously she held the position of principal engineer in Alcatel Submarine Networks where she was involved in the introduction of WDM technologies to submarine networks. She is an author of over 60 papers relating to optical transmission, and of 11 international patents on submerged network technologies. Her current areas of interest are transmission considerations for terabit networks, and architectural and physical layer issues for photonic packet-based networks.

DAVID K. HUNTER [SM] (d.hunter@eee.strath.ac.uk) obtained a B.Sc. in electronic and microprocessor engineering with first class honors in 1987, and a Ph.D. in optical telecommunications in 1991, both from the University of Strathclyde. He is now a senior research fellow there, researching optical networking, especially optical packet switching. He has authored or co-authored over 80 publications, is an associate editor of *IEEE Transactions on Communications*, and has been a guest editor for *IEEE/OSA Journal of Lightwave Technology*.

ANNA TZANAKAKI (annat@ilotron.com) is currently working with ilotron Ltd. on the design of optical cross-connects and associated systems. Interests include the design and optimization of tunable wavelength converters for optical packet networks and optical packet switch architectures; she is the author of over 20 publications on optical communications.

*The proposed OPS router will be a predominantly edge network element, and will function as a topological and logical interface between the service and transport layer.*