Approaches to Optical Internet Packet Switching

David K. Hunter and Ivan Andonovic, University of Strathclyde, U.K.

ABSTRACT

Wavelength-division multiplexing is currently being deployed in telecommunications networks in order to satisfy the increased demand for capacity brought about by the explosion in Internet use. The most widely accepted network evolution prediction is via an extension of these initial predominantly point-to-point deployments, with limited system functionalities, into highly interconnected networks supporting circuit-switched paths. While current applications of WDM focus on relatively static usage of individual wavelength channels, optical switching technologies enable fast dynamic allocation of WDM channels. The challenge involves combining the advantages of these relatively coarse-grained WDM techniques with emerging optical switching capabilities to yield a high-throughput optical platform directly underpinning next-generation networks. One alternative longer-term strategy for network evolution employs optical packet switching, providing greater flexibility, functionality, and granularity. This article reviews progress on the definition of optical packet switching and routing networks capable of providing end-to-end optical paths and/or connectionless transport. To date the approaches proposed predominantly use fixed-duration optical packets with lower-bit-rate headers to facilitate processing at the network–node interfaces. Thus, the major advances toward the goal of developing an extensive optical packet-switched layer employing fixed-length packets will be summarized, but initial concepts on the support of variable-length IP-like optical packets will also be introduced. Particular strategies implementing the crucial optical buffering function at the switching nodes will be described, motivated by the network functionalities required within the optical packet layer.

INTRODUCTION

The convergence of telecommunications and data communications has caused a paradigm shift in the networking environment. The massive explosion in traffic generated by the Internet has driven the current trend, with the Internet Protocol (IP) becoming the dominant protocol for data communications as well as representing, in the longer term, a very strong candidate for the convergence of data communications with telecommunications. In tandem, the development of wavelength-division multiplexed (WDM) techniques on point-to-point links have begun to utilize the massive optical bandwidth of installed single-mode optical fibers more efficiently. WDM was initially a rapid solution to the severe route exhaustion problems brought about by exponentially increasing traffic. The future deployment of optical cross-connected WDM transport networks, initially for protection and bypass but ultimately managing optical light paths dynamically in multiple ring or mesh architectures, will potentially modify the role of the network functionalities provided by the synchronous optical network/synchronous digital hierarchy (SONET/SDH) layer. Thus, *IP over WDM* has become a very important area of study, encompassing a wide range of solutions to supporting predominantly IP traffic over WDM optical paths. Much research effort focuses on developing an elegant solution to the mismatch between the transmission capacities offered by the WDM optical layer and the processing power of routers. IP routers [1] perform four main tasks:

- *Routing*: providing network connectivity information through routing tables
- *Forwarding*: defining the output of each incoming packet (based on the routing tables)
- *Switching*: directing each packet to the proper output (defined by the forwarding process)
- *Buffering*: resolving contention by storing packets when more than one wishes to go to the same output at once, due to the unscheduled nature of their arrival

Currently, the forwarding process implies major throughput limitations, with the size of the routing tables and frequency of their updates being major issues. Such problems are currently addressed and managed, but the time needed for table lookup sets a fundamental limit on router throughput. Much work has concentrated on the development of data structures and algorithms for minimizing the lookup time given routing table and memory space constraints.

0163-6804/00/$10.00 © 2000 IEEE

There is no doubt that massive strides have been achieved in high-throughput router designs (> 1 Tb/s). Nevertheless, despite these impressive advances, there is still a fear that electronic switching systems exhibit limited upgrade flexibility. Given that WDM allows cheap and easy incremental increases of the transmission bandwidth, frequent upgrades of the transport-layer transmission capacity can be envisaged to match increasing demand, in turn placing heavy demands on the switching process.

Optical Packet Switching

A number of other approaches to obviate or ameliorate the forwarding bottleneck are being researched, utilizing optical packet or burst techniques [2], perhaps with electronic buffering, to implement IP over WDM. However, the strategy detailed here implements contention resolution (buffering) directly in the optical domain to yield WDM optical packet switching. The objective is to shift the bulk of the switching burden into the optical domain, permitting compatible scaling of the switching capability with WDM transmission capacity. Thus far, this strategy has assumed a hybrid solution, achieving decoupling between the throughput and the routing/forwarding processes. Transmission and switching are executed in the optical domain, while routing and forwarding are carried out electronically, where the relatively complex packet header processing occurs independent of the optical payload. This decoupling effectively permits the optical packet layer to support a range of networking protocols while harnessing the power of WDM transmission. However, it must be noted that this is also changing, with the recent demonstration of rudimentary header processing functions directly in the optical domain [3]. These relieve some of the burden placed on electronic processing, thereby reducing control signal setup time and managing latency more effectively.

With an extensive optical packet layer, the interface to IP and other protocols is crucial. Encapsulation, the addition of delivery information to the data by the optical packet layer, will occur at interworking units (IWUs) at each interface to the electronic client layer. Encapsulation permits a range of protocols such as IP and asynchronous transfer mode (ATM) to be mapped into the optical payloads, which may be of either fixed or variable duration (see below). IP hides the complexity of the physical layer (including optical packet switching), providing a unified interface to higher layers, regardless of the underlying network type. In addition to encapsulation, the IWUs create headers for proper routing within the optical packet layer, and multiplex traffic from different input links for onward transmission in optical packets for the same destination, ensuring an entirely optical end-to-end connection path. Optical packets provide a further multiplexing tier, allowing the aggregation of traffic flows prior to transmission over the optical layer, and also potentially obviating the need for SDH as an adaptation layer for IP traffic on WDM links. Optical packet networking therefore offers a potential solution to providing both connectionless and connection-oriented networking capacities, flexible in terms of bandwidth management and future-proof with regard to bandwidth growth. No standards exist yet for mapping protocols such as IP and ATM into the optical packet layer.

As hinted above, there are two principal approaches to optical packet switching, both with applications to the Internet:

- Employing fixed-length optical packets, with many corresponding to one IP datagram, requiring IWUs to fragment and reassemble the packets either at the edges of the layer or on the inputs and outputs of the switch
- Employing a variable-length optical packet for each IP datagram

Most reported research to date uses fixed-duration optical packets. Hence, for the purposes of this article (in order to illustrate the functionality of an optical packet-switched layer), the bulk of the subsequent material will be confined to fixed-length packets where both the header and payload are encoded on the same wavelength. It is assumed that the destination switch output for each packet is derived from the header after opto-electronic conversion; the header may thus be at a lower bit rate to allow its electronic manipulation. Due to the nature of optical buffering, the payload duration is fixed, whatever its content; the network throughput is proportional to payload bit rate which may vary from 10 Gb/s and up, with easy upgrade capability. As will also become apparent, the wavelength dimension is crucial not only for transmission capacity but also in executing practical contention resolution. The article will end with some forward-looking concepts addressing the requirements to switch and buffer variable-length optical packets.

Given that WDM allows cheap and easy incremental increases of the transmission bandwidth, frequent upgrades of the transport layer transmission capacity can be envisaged to match increasing demand, in turn placing heavy demands on the switching process.

The Design of Optical Packet Switches

A generic optical packet-switched node structure (Fig. 1) consists of three subblocks:

- An input interface consisting of an (optical) synchronizer which aligns incoming packets in real time against a clock
- A switching core which routes the packets to their proper outputs and executes contention resolution
- An output interface which inserts a new header and may have to regenerate the data

Packet format is a fundamental consideration in any packet transmission system, defined by the requirements of the layer with respect to network functionalities, and is also crucial for optical domain implementations. For example, consider the packet format defined by the KEOPS project [4], upon consideration of the delay-throughput performance of the optical packet-switched layer under different traffic flows as well as consideration of the node and network routing requirements (Fig. 2). Throughout it is assumed that time is divided into equal time slots, each containing one optical packet, and the payload may contain data from 622 Mb/s to 10 Gb/s.

In summary, the header field is 14 bytes: 8 for routing information; 3 for identification of payload type; flow control information, packet numbering for sequence integrity preservation,

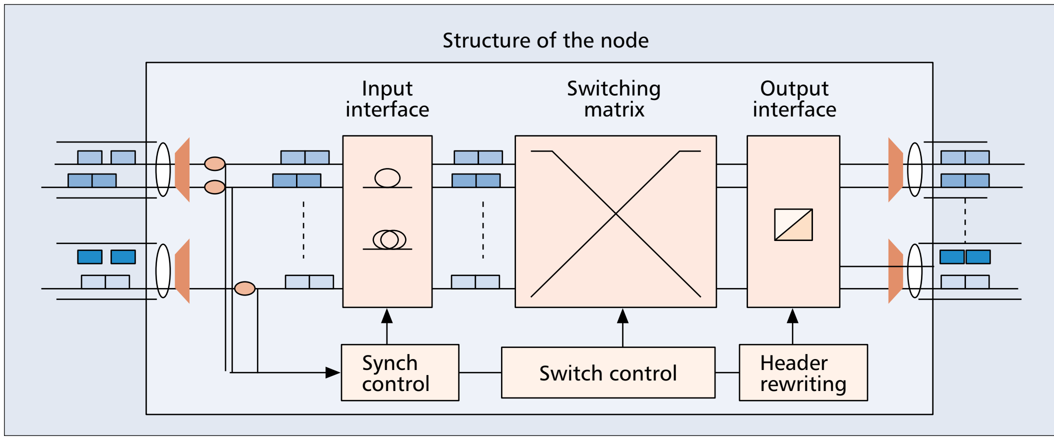

■ **Figure 1.** *A generic optical packet-switched node consisting of three subblocks: the input interface, switching core, and output interface (after KEOPS).*

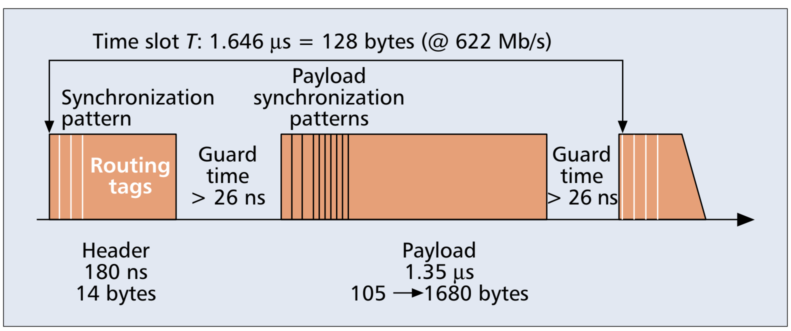

■ **Figure 2** *An example of an optical packet format (as defined by KEOPS).*

and header error checking. The format also indirectly indicates the need for optical synchronizers as the input interface to the node. It is assumed that when packets enter the switch, their boundaries are aligned so that each packet is aligned with its timeslot. Such synchronization is generally a requirement for correct switch operation, achieving low packet loss; approaches to achieving packet synchronism at the inputs to the switches constitute a separate topic and will not be discussed further. Thus, 2 bytes are required to aid in the synchronization process as "tags" uniquely identifying the start of the payload. Guard bands account for switching times of the constituent opto-electronic devices as well as payload position jitter. The payload duration results from a trade-off between transmission efficiency (the longer the payload duration, the higher the efficiency) and the practical limitations on overall length of the optical fiber delay line buffers. At this time, delay lines are the only viable approach to implementing contention resolution directly in the optical domain.

An optical packet switch has three principal functions: *switching*, *buffering*, and, optionally, *header translation* [5]. In IP-oriented systems, the latter function may be replaced by routing and forwarding, which were discussed earlier. *Switching* ensures that each packet emerges at the correct output, depending on the information contained in the packet header. Although fixed-length packets arriving on the inputs must be synchronized, there is no coordination between packet streams arriving on different inputs. Hence one or more packets may arrive during the same time slot on different inputs wishing to go to the same output. For this reason, *buffering* is required, where one or more packets are stored while others are transmitted to the desired output.

Header translation represents a central process in the ATM transmission systems strategy, while, as discussed previously, routing and forwarding are equally important for IP systems. Although these schemes offer great functionality and flexibility, they are not used in every optical packet-switching system since direct optical header translation is in the early stages of development. Thus, the header is usually at a lower bit rate than the payload to facilitate electronic decoding and interpretation of header information.

The implementation of an effective buffering strategy directly in the optical domain is fundamental to many node designs. Since optical random access memory (RAM) is not yet practical, delay lines (usually made from optical fiber) must be used to store optical packets. Various solutions to optical packet switches have been proposed, dictated by the buffering strategy [5]:

- *Implement medium to large buffers*: Switches are cascaded to implement very large buffers suitable for bursty traffic.
- *No buffers in the switch nodes, but employ deflection routing*: When multiple packets arrive destined for a given output, all but one are deflected to other outputs, finding their way to the destination node by another route through the network.
- *Use a small amount of buffering with deflection routing*: Simple 2 x 2 buffered switches consisting of a chain of 2 x 2 crosspoint devices and delay lines.
- *Use the wavelength dimension to reduce the amount of buffering*: Capacity sharing over many wavelengths reduces the amount of packet contention, which in turn reduces the amount of required physical optical buffering.

There has been much research on the design of optical packet switches, their common limitation being optical splitting loss. The following sections will detail progress made in the design of optical packet networks utilizing a combination of a managed amount of physical buffering in tandem with the wavelength dimension to realize moderately deep buffers. Due to space constraints, two example designs originating from the work under KEOPS within ACTS and WASPNET, a U.K.-based multi-institution collaboration, will form the focus to illustrate the strength of the approach.

Wavelength in Contention Resolution

In WDM packet networks, two possible multiplexing schemes have been identified [6]: scattered wavelength path (SCWP) and shared wavelength path (SHWP). In SHWP, each path in the optical packet layer is assigned a particular wavelength. Since the number of wavelengths is limited, two or more optical paths may share the same wavelength within a fiber; hence, each shared wavelength operates much like a distinct link carrying multiple packet paths. The packets' fields identify them within the wavelength, and if the number of wavelengths is increased and the load remains constant, the buffering requirement decreases since the load *per wavelength* is lower.

Unlike SHWP, in SCWP each optical packet is not allocated the same wavelength along its entire path, but is dynamically converted to a suitable wavelength in each link. This decision is influenced by the availability of free time slots on each wavelength when forwarding to the next node, and other network management factors. For example, if there are 128 wavelengths within a particular fiber link carrying a number of optical packets, every packet directed along the fiber can be transmitted on any of the 128 wavelengths. Contention between any two packets being transmitted simultaneously is resolved in the first instance by transmitting them each on a different wavelength. If this fails due to all wavelengths being occupied, optical buffering is employed to delay some of the packets via optical fiber memories until a free time slot is available. Employing wavelength to resolve contention in this way reduces the required buffer capacity. As an illustrative example, consider that with fixed-length packets, each wavelength carries a load of 0.8 (i.e., on each time slot each wavelength has an 80 percent probability of carrying a packet, independent of other wavelengths and time slots). 0.8 is widely regarded as a typical figure found in practice. Each packet has an equal probability of being directed to each output. The reduction in buffer size for a 10^{-9} packet loss rate is sought (i.e., on average one packet in 10^{-9} must be discarded due to buffer overflow). With SCWP, a reduction in buffer size of 70–80 percent results even with a modest number of wavelengths, while almost 90 percent reduction results with 32 wavelengths.

For a given optical buffer depth and packet loss, SCWP offers higher throughput than SHWP because the dynamic wavelength allocation permits sharing of capacity between wavelengths. For the same packet loss probability, smaller buffers can be used with SCWP, while for the same size of buffer, SCWP yields a lower packet loss probability. Figure 3 depicts analytical results that justify this, assuming the same idealized traffic as above, comparing buffer depth per wavelength with both SCWP and SHWP at different loads for a 32-input/output switch to achieve a packet loss of 10^{-9} as before. With SCWP there is one large buffer per fiber for all wavelengths, and the "buffer depth per wavelength" is the size of this buffer divided by the number of wavelengths. With SHWP, there is one buffer per wavelength (i.e., multiple buffers per fiber). Since SHWP requires larger buffers for the same packet loss, SCWP is identified as the optimum multiplexing format, offering a buffer size reduction of over three times.

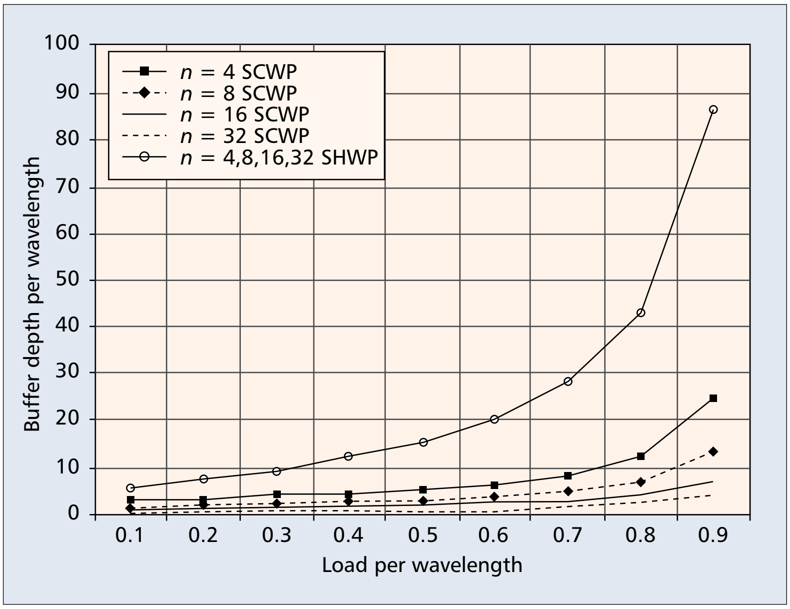

Figure 3. *An idealized comparison of buffer depth per wavelength for both SCWP and SHWP at different loads for a 32-input/output switch to achieve a packet loss of 10^{-9}, with fixed-length packets.*

Therefore, optical packet nodes that harness the power of the wavelength dimension implement switching at the packet level per wavelength per time slot. The first example is the KEOPS broadcast and select switch (Fig. 4), where wavelength is used only within the switch fabric, rather than on external links as with SCWP. While wavelength is not used to resolve contention as described above, it is crucial in facilitating the internal operation of this switch. The $N \times N$ switch consists of three sections: the *wavelength encoder*, the *buffer and broadcast section*, and the *wavelength selector block*. The wavelength-encoding block consists of N wavelength converters, one per input, encoding each packet on a fixed wavelength with a unique wavelength for each input. The cell buffer block comprises a number of fiber delay lines followed by a space switch stage. The gates select packets from the correct delay lines and send them to the correct outputs under electronic control. The last block, the wavelength selector block, consists of N demultiplexers which forward the different outputs to optical gates before the signals are recombined, thus selecting packets from the correct inputs. The broadcasting principle makes a copy of each packet at any delay, enabling flexible management of the buffered packet flows. A copy is also available at each output; hence, this architecture supports multicasting. Experimental characterization of a

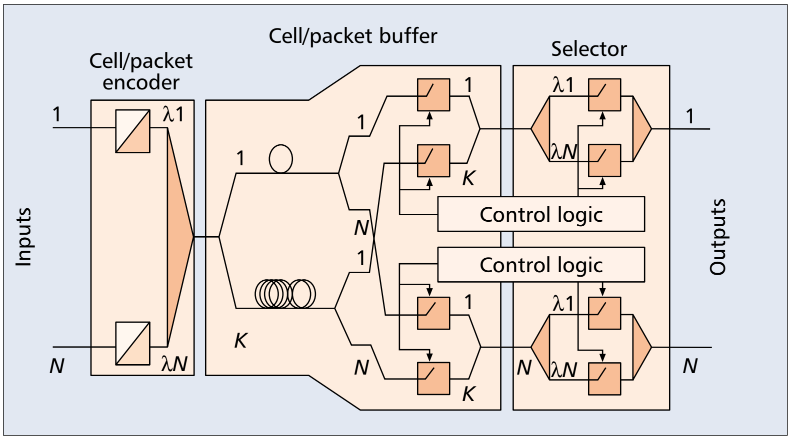

Figure 4. *The KEOPS broadcast and select switch.*

16 x 16 matrix operating at 10 Gb/s achieved a sensitivity penalty of below 1 dB with respect to the back-to-back configuration [4].

The WASPNET design [6] adopts a different architecture, overcoming the relatively large splitting losses of the above switch. The core components are tunable wavelength converters (TWCs) and arrayed waveguide gratings (AWGs). Although an equivalent feed-forward design has been developed, Fig. 5 shows the feedback configuration using shared feedback delay lines to buffer the contended packets, chosen to enable the implementation of priority routing at the expense of requiring a larger (2*N* x 2*N*) AWG. (Ignore the *N* x *N* space switch for the time being.) At each input, packets are converted to the appropriate wavelength to be switched to the correct AWG output; the contended packets are routed to the recirculating loops for buffering, while the straight-through packets are routed to switch outputs. Only one packet may exit from each multiplexer in the loop simultaneously. Control is implemented assuming a first-in first-out (FIFO) algorithm

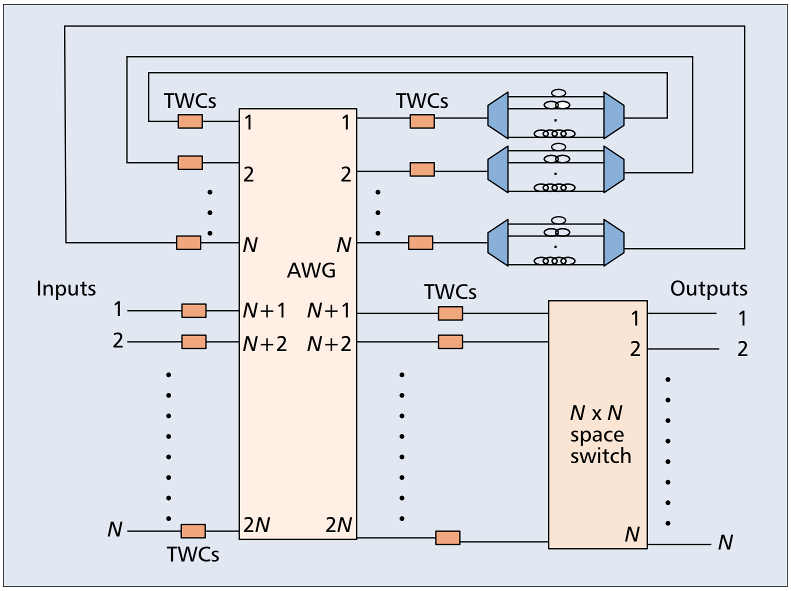

Figure 5. *The WASPNET feedback optical packet switch, incorporating tunable wavelength converters and an arrayed waveguide grating.*

and ensures that no more than one packet may exit from each combination of fiber delay lines, a demultiplexer and a multiplexer; otherwise, it will be routed to other alternative AWG outputs or, if necessary, to longer-delay lines. If a higher-priority input packet is destined for the same output as the buffered packet, the buffered packet will be switched to the feedback delay lines again, and the input packet with higher priority is switched straight through. On a network level, the FIFO strategy ensures that packet ordering is maintained throughout a connection.

Without the strictly nonblocking one-to-many *N* x *N* space switch, the above architecture has single-wavelength inputs and outputs, but when it is included along with its associated wavelength converters, the architecture can handle WDM inputs/outputs in order to implement SCWP for assistance in contention resolution. All the inputs are wavelength demultiplexed to a number of parallel planes, one for each wavelength channel, each plane containing the optical packet switch of Fig. 5. A combiner instead of a multiplexer merges the wavelength channels from all planes, allowing dynamic wavelength allocation at each plane. The AWG in each plane routes packets to the space switch, which then switches each of them to the correct output port at the correct wavelength via the final stage of wavelength conversion. A testbed has been constructed to demonstrate the full functionality of this node, and no appreciable power penalty was measured after 11 cascades.

Variable-Length Optical Packet Switching

Thus far, it has been assumed that nodes handle fixed-length packets. Another potential approach involves investigating optical packet-switched node architectures that are suitable for use as the basis of IP routers, capable of switching and buffering variable-length packets optically. This is a demanding task, and here a strategy will be outlined that potentially yields a viable architectural concept to achieve this goal [7].

In keeping with the philosophy detailed above, each node may accept or transmit multiple packets simultaneously on each input or output fiber using WDM, enhancing its throughput. This also implies that the IP and WDM layers may be combined in future networks, simplifying network management and producing a highly flexible and functional packet switching layer.

High hardware complexity is a potential difficulty, and here four tiers of the design strategy can be identified to combat this:

- Asynchronous operation is permitted on router inputs, so it is not necessary to synchronize packets to time slot or byte boundaries prior to entering the switch. Synchronizers represent considerable hardware overhead.
- Again, sharing of capacity among wavelengths (as with SCWP) is used to assist in contention resolution, reducing the physical buffering requirement.

- If the lowest increment of delay permitted is less than the minimum packet size, packets may be directed in a FIFO manner to the appropriate outputs. However, the hardware complexity inherent in this approach (due to the large number of delay lines) may be avoided by using void filling [8]. Here, the delay line increment may be much larger than the minimum packet length; as a result, unused gaps or "voids" are scheduled between adjacent packets on any given output. Rather than scheduling a new packet under consideration for a particular output to be transmitted immediately after the last one, it may be scheduled to fill a void on the correct output in a non-FIFO manner. It is still possible to maintain FIFO operation for each input-output pair by stipulating that a packet cannot be scheduled for transmission before the previous packet between that particular input and output.
- The use of multiple buffer stages in cascade (each with relatively few delay lines) limits the high splitting and combining losses and hardware costs inherent in implementing only a single stage with a very large number of delay lines. This approach has already been proposed for fixed-length optical packets in the *switch with large optical buffers* (SLOB) [9]. SLOB consists of a cascade of existing optical packet switches, similar to the broadcast and select switch considered above, each such stage having its own delay lines for contention resolution. The delay line lengths in each stage increase exponentially from left to right along the structure. The number of stages increases with the logarithm of the buffer depth, and SLOB emulates an output-buffered switch yielding extremely large buffer depths.

One possible optical IP router architecture based on the above strategies is composed of a series of S stages, numbered 0 to $S - 1$ (Fig. 6). Each stage has the following number of inputs and outputs, where N is the number of inputs and outputs on the architecture, and D the number of delay lines and links connecting adjacent stages:

- Stage 0: N inputs and D outputs
- Stage $S - 1$: D inputs and N outputs
- All other stages: D inputs and D outputs

Each pair of adjacent stages are interconnected by delay lines which increase in units known as *delay line granularity*. The granularity of each stage is N times that of the next stage, although other configurations are of course possible and could be tried.

Each link or delay line within the architecture can carry multiple packets at once, by means of WDM. Each packet entering a stage may leave on any other stage output, on any free wavelength. Hence, each stage functions much like a cross-connect with wavelength conversion, although operating at the packet rate. There is a specified number of wavelengths on all switching stage ports, delay lines, and links internal to the architecture, and a specified number of wavelengths on all node inputs and outputs (external to the architecture). The total number of wavelength channels between adjacent stages within the switch (i.e., number of internal wavelengths times D) must be at least equal to the number entering and leaving the switch to prevent overloading within the architecture.

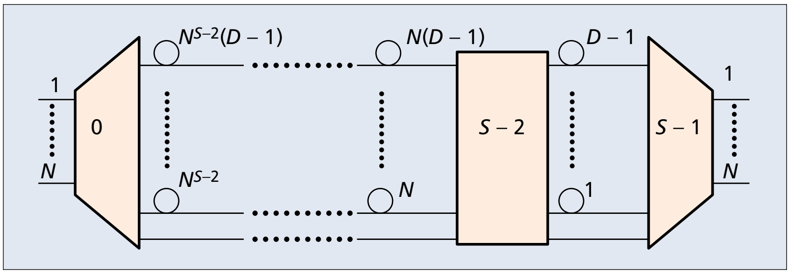

Figure 6. *A schematic of an optical packet switch for routing variable-length optical packets. It has* S *stages, and the delay line lengths are measured in units of granularity.*

Initial traffic analysis of the node under bursty and self-similar statistics has indicated the viability of the concept. For negative exponential traffic (bursty), with just two stages with 32 wavelengths throughout, 16 delay lines per stage, and 16 inputs and outputs, the packet loss was always less than 10^{-6}. Due to the deleterious effect of self-similar Pareto traffic, an appreciable packet loss was obtained for two stages with such traffic. (Simulations were carried out using self-similar traffic with $\alpha = 1.2$, mimicking real traffic, corresponding to a Hurst parameter of 0.9.) The minimum packet length was 400 bytes to give an indication of switch performance. Each simulation was carried out with the basic delay line granularity in the righthand stage equal to 400, 2000, or 5000 bytes. Results show that for three stages, 64 wavelengths, and a delay-line granularity of 5000, a satisfactory packet loss of less than 10^{-6} was achieved. Although this was a preliminary study, the results of the traffic analysis are encouraging. It must be noted that the control algorithm is brute force, a modified form of the original void-filling algorithm making a sequential search of the available routes. Thus, it is computationally intensive, and improvements must be made to improve the practicality of the control strategy.

Conclusions

Given the exponential increase in telecommunications and data traffic levels that has occurred over a relatively short period, novel approaches to creating an extensive, flexible, and future-proof optical infrastructure are required. In the longer term, one route to providing this is through an optical packet-switched WDM layer embracing many of the technological advances driven by mainstream deployment of WDM networking. Here an outline of the progress made to date, in both node design and networking concepts, toward creating a platform directly carrying optical packets on a WDM infrastructure is given. Many of the concepts assume a fixed optical packet length with a combination of physical buffering provided by optical fiber delay lines (or perhaps silica-on-silicon waveguides in the future) and wavelength to carry out the important task of contention resolution directly in the optical domain.

A potentially useful strategy has been devel-

A novel approach to creating an extensive, flexible and future-proof optical infrastructure is through an optical packet-switched WDM layer embracing many of the technological advances driven by mainstream deployment of WDM networking.

oped to create new optical architectures able to route variable-length optical packets (e.g., IP packets) based on asynchronous operation, the use of wavelength to resolve contention, and void filling according to a multistage architectural concept. The architecture was simulated under bursty and self-similar traffic statistics, and preliminary results indicate that low packet loss rates can be achieved at modest hardware complexity. The control of the fabric and table lookup remain key issues. However, fragmentation and reassembly of IP packets at layer or node boundaries is not required, realizing an appreciable simplification.

All of the architectures discussed make appreciable demands on optical hardware technology, and optical performance is all-important, setting fundamental limits on the node extent and the number of nodes that can be cascaded in one optical path. Studies have shown that several stages similar to those proposed here can be cascaded, and experimental studies have demonstrated the validity of cascading up to 40 such stages with the aid of all-optical regeneration [10]. It is abundantly clear that a viable solution to all-optical regeneration is crucial if an extensive optical layer is ever to be deployed.

Acknowledgments

The authors would like to express their thanks to the members of the Broadband Networks Group at Strathclyde University for their contributions: Meow Chia, M. H. M. Nizam, Stig Quist, Chin Soon Hwa, and Andre Franzen. The authors also thank the members of both the KEOPS (ACTS funded) and WASPNET (EPSRC funded) consortia for their contributions. Grateful acknowledgment goes to Ljubisa Tancevski of Alcatel Corporate Research Center, Richardson, Texas, for encouragement and useful discussions. David Hunter is funded by the EPSRC in the form of an Advanced Fellowship.

References

[1] S. Keshav and R. Sharma, "Issues and Trends in Router Design," *IEEE Commun. Mag.*, May 1998, pp. 144–51.

[2] 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.

[3] A. Carena *et al.*, "OPERA: An Optical Packet Experimental Routing Architecture with Label Swapping Capability," *IEEE/OSA J. Lightwave Tech.*, vol. 16, no. 12, Dec. 1998, pp. 2135–45.

[4] C. Guillemot *et al.*, "Transparent Optical Packet Switching: the European ACTS KEOPS Project Approach," *IEEE/OSA J. Lightwave Tech.*, vol. 16, no. 12, Dec. 1998, pp. 2117–34.

[5] D. K. Hunter, M. C. Chia, and I. Andonovic: "Buffering in Optical Packet Switches," *IEEE/OSA J. Lightwave Tech.*, vol. 16, no. 12, Dec. 1998, pp. 2081–94.

[6] D. K. Hunter *et al.*, "WASPNET - a Wavelength Switched Packet Network," *IEEE Commun. Mag.*, Mar. 1999, pp. 120–29.

[7] D. K. Hunter, I. Andonovic, and M. C. Chia, "Optical Buffers for Multi-terabit IP Routers," *37th Ann. Allerton Conf. Commun., Control, and Comp.*, Monticello, IL, Sept. 22–24, 1999.

[8] L. Tancevski *et al.*, "A New Scheduling Algorithm for Asynchronous, Variable Length IP Traffic Incorporating Void Filling," *OFC '99*, San Diego, CA, Feb. 1999.

[9]. D. K. Hunter *et al.*, "SLOB: a Switch with Large Optical Buffers for Packet Switching," *IEEE/OSA J. Lightwave Tech.*, vol. 16, no. 10, Oct. 1998, pp. 1725–36.

[10] D. Chiaroni *et al.*, "Experimental Validation of an All-Optical Network Based on 160 Gbit/s Throughput Packet Switching Nodes," *ECOC '98*, Madrid, Spain, Sept. 20–24, 1998.

Biographies

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.

IVAN ANDONOVIC [SM] (i.andonovic@eee.strath.ac.uk) joined the University of Strathclyde in 1985 and is now a professor of broadband networks there. His interests center on guided wave architectures for optical signal processing and switching in future networks. He has been chair of IEE professional group E13, held a BT Short Term Fellowship, and been a visiting scientist at CRL in Japan.