WASPNET: A Wavelength Switched Packet Network

David K. Hunter, Mohamed H. M. Nizam, Meow C. Chia, and Ivan Andonovic, University of Strathclyde

Ken M. Guild, Anna Tzanakaki, and Mike J. O'Mahony, Essex University

John D. Bainbridge, Marc F. C. Stephens, Richard V. Penty, and Ian H. White, University of Bristol

ABSTRACT WASPNET is an EPSRC-funded collaboration between three British universities: the University of Strathclyde, Essex University, and Bristol University, supported by a number of industrial institutions. The project — which is investigating a novel packet-based optical WDM transport network — involves determining the management, systems, and devices ramifications of a new network control scheme, SCWP, which is flexible and simplifies optical hardware requirements. The principal objective of the project is to understand the advantages and potential of optical packet switching compared to the conventional electronic approach. Several schemes for packet header implementation are described, using subcarrier multiplexing, separate wavelengths, and serial transmission. A novel node design is introduced, based on wavelength router devices, which reduce loss, hence reducing booster amplifier gain and concomitant ASE noise. The fabrication of these devices, and also wavelength converters, are described. A photonic packet switching testbed is detailed which will allow the ideas developed within WASPNET to be tested in practice, permitting the practical problems of their implementation to be determined.

Wavelength-division multiplexing (WDM) is currently being deployed in telecommunications networks in order to satisfy the increased demand for capacity brought about by both narrowband services and new broadband services such as high-speed Internet. While it is thought that WDM will ultimately evolve to interconnected rings or perhaps a mesh network, the objective of the Wavelength Switched Packet Network (WASPNET) project is to gain a more long-term understanding of how optical networks will develop. WASPNET is a WDM transport network that uses optical packet switching, resulting in greater flexibility, functionality, and granularity than possible with the current generation of WDM networks. These optical packets may be used to carry asynchronous transfer mode (ATM) or IP, for example, and the network is also designed to support synchronous digital hierarchy/synchronous optical network (SDH/SONET) traffic, thus permitting a smooth upgrade path.

Optical packet switches [1–3] have attracted considerable research interest internationally due to their potential for overcoming projected difficulties with very large electronic switching cores, such as connection, pinout, and electromagnetic interference (EMI) problems. A key problem when designing packet switches of any kind is contention resolution, since multiple packets may arrive asynchronously at the same time to go to the same output. Buffering is often employed to solve this problem, but since optical random access memory (RAM) does not exist, delay lines (usually made of optical fiber) must be used to store optical packets and implement buffering. Various solutions to optical packet switching have been proposed, dictated by the buffering strategy [1].

Implement Medium to Large Buffers — The switches implemented by this technique may be cascaded to implement very large buffers, suitable for bursty traffic.

Use 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, to find their way to the destination by another route through the network. This not only provides fast and flexible routing, but also allows nodes to have no buffering. However, each packet transmitted from a node may be routed across a different path to the same destination. Some packets may wander within the network and waste bandwidth. Consequently, each packet will experience different propagation delays, and the traffic may not arrive at the destination node in sequence.

Compromise by Using a Small Amount of Buffering with Deflection Routing — There are various such 2 x 2 buffered switches consisting of a chain of 2 x 2 switch devices and delay lines.

Here a new approach, WASPNET, is proposed as a solution to the ever-increasing demand for telecommunications transport capacity. A key feature is its use of statistical multiplexing over many wavelengths to reduce the amount of packet contention, and hence the amount of optical buffering required — which is difficult to implement — for a given quality of service (QoS). It is a packet-based transport network, which is designed both to support conventional optical paths for the transport of SDH and also to switch optical packets. Unlike other optical packet network proposals, WASPNET is a reconfigurable multiwavelength transport network. Hence, it provides a smooth upgrade path from SDH over WDM while still supporting legacy SDH equipment, and possesses greater flexibility and granularity than existing WDM network proposals.

In WASPNET, not only are node design and routing considered, but also network control and operation, device fabrication, and demonstrator construction. Two possible network control methodologies were identified: the scattered wavelength path (SCWP) and shared wavelength path (SHWP) schemes. These are compared in the following section, and the problems of control, packet ordering, resilience, and sequencing are then addressed. We describe candidate packet formats, and detail the switch architecture that has been proposed. The article goes on to describe the devices that will be used in the demonstrator, followed by a description of the demonstrator, which is under construction. Finally, the last section contains the conclusions.

 0163-6804/99/$10.00

The Scattered Wavelength Path

SHWP and SCWP are multiplexing schemes for optical packets carried over WDM. In SHWP, one wavelength is allocated to each path, or optical packet path (OPP), in the WASPNET optical packet layer. Since the number of wavelengths is limited, two or more OPPs may share the same wavelength within a fiber; hence, each shared wavelength operates much like a distinct link carrying multiple packet paths. The packets' OPP 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.

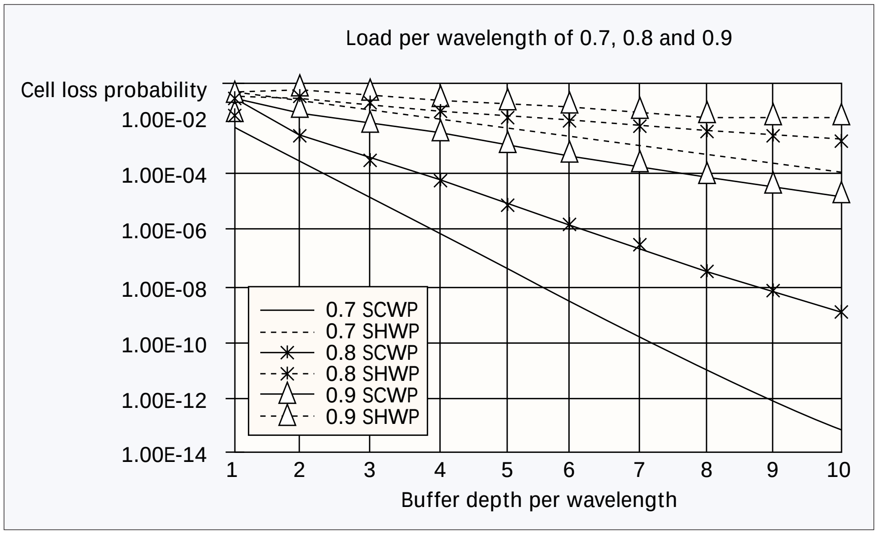

■ **Figure 1.** *Cell loss probability for a network with* N *= 16 physical input/output ports, and* n *= 4 wavelength channels per port at different loads per wavelength.*

Unlike SHWP, in SCWP each OPP is not allocated the same wavelength along its entire path; each packet 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 management factors to be described later. Suppose there are 16 wavelengths within a particular fiber link carrying a number of OPPs. Every OPP packet directed along the fiber can be transmitted on any of the 16 wavelengths; contention between any two packets being transmitted simultaneously is resolved in the first instance by transmitting each on a different wavelength. If this fails due to all wavelengths being occupied, buffering is employed to delay some of the packets until a free time slot is available. Employing wavelength to resolve contention in this way reduces the required buffer capacity.

SCWP offers higher throughput than SHWP because the dynamic wavelength allocation permits sharing of capacity between wavelengths, yielding lower packet contention probabilities and requiring smaller optical buffers at each node. Figure 1 depicts idealized analytical results that justify this, assuming uniform Bernoulli traffic and comparing the packet loss due to contention with both SCWP and SHWP at different loads. With SCWP there is one large buffer per fiber for all wavelengths, while with SHWP there is one buffer per wavelength (i.e., multiple buffers per fiber). The results indicate that a packet loss as low as 10^{-12} may be obtained at a load of 0.7/wavelength with SCWP and a buffer depth of only 9/wavelength. Since SHWP requires larger buffers for the same QoS, SCWP is identified as the optimum multiplexing format for WASPNET; for example, at a packet loss rate of 10^{-5} and a load of 0.9, SCWP offers a buffer size reduction of over three times.

Control, Packet Ordering, and Resilience

Despite these advantages, SCWP requires wavelength conversion and complex processing at each node. For example, packets must be sorted into order at the path's destination, since a later packet may overtake an earlier packet at an intermediate node. This is because multiple packets from the same OPP may arrive at a node simultaneously, and the node has no information available to prevent it from transmitting a later packet over a link before an earlier one.

One way of circumventing this problem might involve placing a sequence field in each packet's header to permit packet sequence within an optical path to be monitored. This is only feasible when the burst duration or connection holding time can determine the header sequence field length. In a transport network (the first section), there is no specific connection holding time, because paths are set up on a semi-permanent basis. Instead, WASPNET uses wavelength and arrival time as sequence identifiers; packets from the same OPP on the same time slot on a link are sorted with earlier packets on lower-numbered wavelengths. Hence, each node has the correct information — encoded as the packets' wavelengths — to prevent later packets in the sequence from being sent out first.

Additionally, with SCWP, restoration is computationally intensive. Suppose that each physical link carries, say, 15 active wavelengths and one spare wavelength for restoration. If a link fails between two adjacent nodes, distributed restoration can quickly reroute each affected wavelength on the link over an alternative wavelength connection, each taking a different physical route. With SCWP, the new route a packet takes will depend on its wavelength; hence, packets from the same OPP will be routed differently, propagating for different distances, and thus causing the network to lose control of OPP packet order. Even though this scheme functions well with SHWP, with SCWP it is not possible to restore each wavelength alone; instead, each OPP must be rerouted individually. Since there are generally more OPPs than wavelengths, SCWP restoration requires more iterations than SHWP.

This problem is resolved by a secondary management layer, where OPPs are bundled into sets of secondary paths known as secondary packet paths (SPPs). This technique is much like bundling ATM virtual circuits (VCs) into virtual paths (VPs). If a link fails, the restoration nodes check each packet's OPP and use this to determine the SPP from the node's database; therefore, packets need not contain an SPP field. To implement restoration, SPPs instead of OPPs are rerouted, resulting in lower iteration complexity due to the lower number of SPPs.

The complexity of control required for buffering packets and assigning them to wavelengths increases quadratically with the number of OPPs handled by the node for SCWP, and linearly for SHWP. The node control complexity of SCWP is higher than SHWP, but this is not a difficulty since the electronics available to implement a SCWP controller is sophisticated and relatively inexpensive.

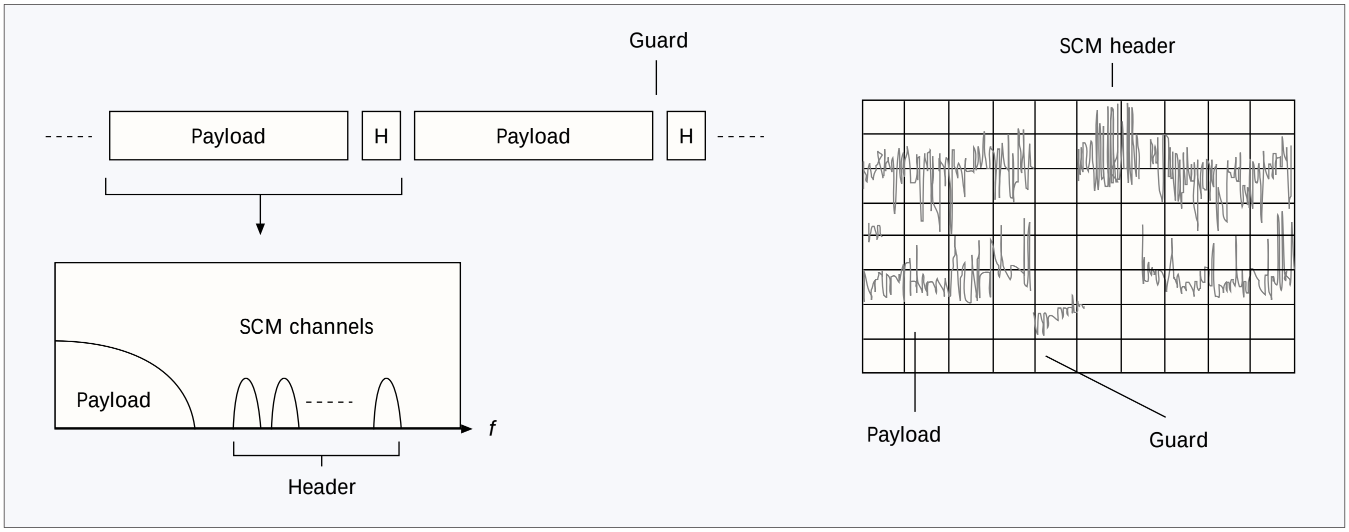

■ **Figure 2.** *Hybrid SCM, mixed-rate header format.*

Header Format

The packet header is 4 bytes long, and contains information about the payload, including an OPP identifier (24 bits), the payload type (2 bits), and priority (2 bits). The payload has been set to 256 bytes at 10 Gb/s. There are several ways of implementing the header, but a sufficiently reliable format must be chosen since a header error causes packets to be misrouted, implying error multiplication. Three options are being considered.

Subcarrier Multiplexing (SCM) — This option can transport header information with the payload using relatively inexpensive electronics. A low-bandwidth header is placed on an electrical subcarrier above the baseband frequencies occupied by the packet payload, and both are transmitted in the same time slot. Depending on the modulation depth, the header can impose a power penalty on the payload through intermodulation distortion, and a compromise must be made to ensure appropriate header detection sensitivity and minimal payload penalty. Header removal is also difficult since the SCM header passes through the photonic switching matrix with the payload. In WASPNET, the header is modified during routing, and this must be done so that the transparency of the optical baseband payload is preserved.

Transmit the Header and Payload on Separate Wavelengths — Also, demultiplex at the node input using passive optical filters. Although this offers a simple method of extracting the header, fiber dispersion will separate the two components, and each node will require delay compensation to realign the header and payload. In order to update the packet header, an additional laser is required for every payload channel, implying extra cost and complexity in header/payload power equalization. As with SCM, this method may also suffer from crosstalk since the payload and header are transmitted in parallel.

Transmit the Header Before the Payload Serially — At the expense of bandwidth utilization, this facilitates header removal/insertion and eliminates the penalty due to header and payload interaction. The header (usually at a lower bit rate) is removed by gating an optical switch at the appropriate time, and it can then be processed utilizing conventional off-the-shelf electronic circuits drawn, say, from ATM implementations.

Packet delay is an important issue for some types of payload. With an SCM header, a routing decision can only be made after an entire packet time slot has elapsed, and the delay introduced by the node is equal to the packet duration plus the header processing time. This delay can be large if long packets are being routed across many nodes. However, the mixed-rate packet format only introduces a delay equal to the header length plus the processing time. If the header length can be minimized, the processing time can fall within the remaining time period of the payload. Consequently, with minimal header time and fast processing, the network delay imposed on the payload can be kept very low.

Packet delineation is the process of identifying the packet boundaries and is another key focus of the project. High-level data link control (HDLC) in X.25 networks use a 0111110 flag to mark the start and end of a packet, whereas ATM networks use a complicated finite state machine solution, searching for the header error control (HEC) byte of the incoming data. Both techniques require either byte insertion or scrambling of the payload to ensure that it does not contain the identification byte and cause false delineation. The SCM header format facilitates packet delineation in the following way. Within the node the optical packet is converted to the electrical domain, and the header is removed by high-pass filtering the signal. The header is electrically processed, and a flag or HEC byte can be identified to signify a packet boundary. In this way, the payload does not have to be inspected or altered, thus enhancing the transparency and reducing the latency of the network.

A hybrid technique of utilizing out-of-band and mixed-rate formats has been investigated (Fig. 2). The packet format takes advantage of SCM for packet delineation and the frequency-division multiplexing (FDM) of short SCM channels to minimize the duration of the header. The payload and header are temporally separated to allow for easier header removal/reinsertion at routing nodes and to remove the latency in each routing node. Which format the header will take is at present a subject of study.

Switch Architecture

A novel aspect of the WASPNET switch [4] is its use throughout of feedback delay lines with wavelength-selective routers (e.g., arrayed waveguide gratings, AWGs). Feedback delay lines allow the implementation of multiple packet priorities, while wavelength routers exhibit low loss and low crosstalk, offering superior systems performance. While this switch uses

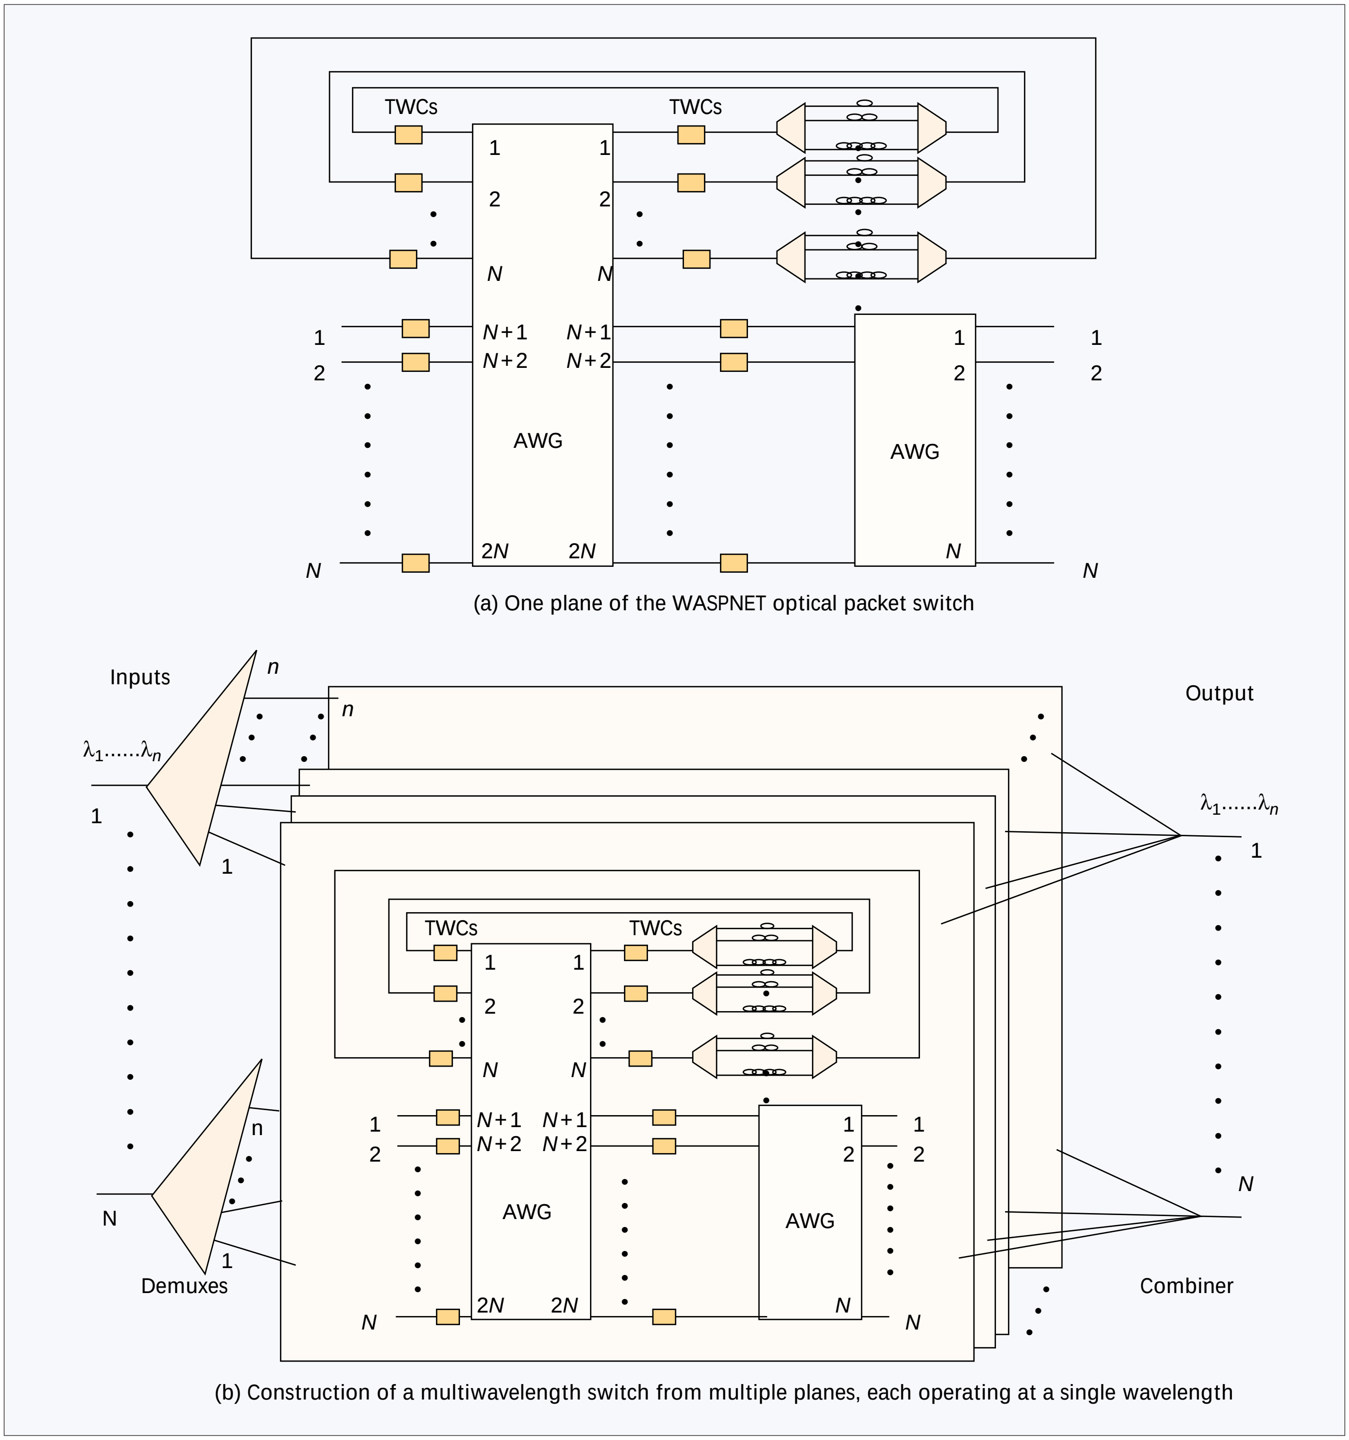

(a) One plane of the WASPNET optical packet switch

(b) Construction of a multiwavelength switch from multiple planes, each operating at a single wavelength

■ **Figure 3.** *THE WASPNET WDM optical packet switch based on wavelength routers.*

multi-input wavelength routers, other proposals using 1 x N routers have also been made in the ACTS KEOPS project [2]. Furthermore, switches have been proposed with wavelength routers and feed-forward delay lines [3], not permitting implementation of packet priorities. Further work within the project will examine the alternative technology of active space switching.

There are N inputs and outputs (Fig. 3a); the present architecture includes 4N tunable wavelength converters (TWCs), a 2N x 2N wavelength router, and an N x N wavelength router. Hence the number of TWCs — and the hardware requirement generally — rises linearly with N, although the amount of hardware increases quadratically in many other architectures, such as the broadcast-and-select switch [5]. The architecture can be said to implement a shared buffer, since each delay line is not associated with a particular output, but implements buffering for all outputs.

Each input packet is wavelength-converted, so the first wavelength router switches them either to the second router, or onto the correct delay lines, depending on the delay required for buffering. Half the first router's outputs each feed a tunable wavelength converter, a demultiplexer, a number of fiber delay lines, and a multiplexer or combiner which feeds back to its input. Each input of this router only carries one wavelength, since each TWC before it only allows single-wavelength conversion. Hence the controller must ensure that several packets do not leave a multiplexer in the feedback loop simultaneously; this is achieved by switching packets to the correct first wavelength router outputs, storing packets in longer delays if necessary.

Scheduling, as in several other switches (e.g., [2, 6]) comprises two phases; in the first phase, packets are routed to the correct delay lines to avoid output contention, and in the second packets are routed to the correct switch outputs. However, here both phases are implemented by the same wavelength router so that if a packet is preempted by a higher-priority packet upon leaving a delay line, it can be redirected to a delay line again instead of proceeding directly to the correct output. This architecture has single-wavelength inputs and outputs, but can easily be modified to handle WDM inputs/outputs (Fig. 3b). All the input wavelength channels are demultiplexed to one parallel plane for each input wavelength, each containing an $N \times N$ optical packet switch (Fig. 3a) as is common practice, although here multiple output wavelengths can leave one plane simultaneously.

Each packet leaving a plane must be at the correct wavelength to prevent wavelength contention. To permit this, there is an additional wavelength router, with a tunable wavelength converter on every input. The first wavelength router sends each packet to that input of the second wavelength router which will carry it to the correct output, while being at the correct wavelength (Fig. 3a); hence, multiple packets may exit on an output of the second wavelength router, each at different wavelengths. Since each of its inputs can only carry one packet at once, there is potential for blocking here, and if this must be totally eliminated, the second wavelength router could be replaced by a many-to-one space switch.

The switch is designed so that it can also handle circuit-switched continuous SDH/SONET paths, allowing WASPNET to be introduced gracefully while still supporting legacy SDH/SONET equipment. To do this, the switch supports a semi-permanent path for the SDH/SONET signal, while still allowing packet switching to take place in the remainder of the switch.

To gain some appreciation of switching system performance, a simple power penalty calculation was performed at a bit error rate (BER) of 10^{-10} and a bit rate of 2.5 Gb/s, with AWGs as wavelength routers. The number of delay lines is assumed to be equal to the number of inputs and outputs, yielding an appreciation of performance.

An 8 x 8 broadcast-and-select switch [5] has a power penalty of 0.5 dB, less than an 8 x 8 WASPNET architecture (1 dB). However, the latter outperforms the former in 16 x 16 and 32 x 32 configurations, since it has lower splitting/combining losses for large switch sizes, requiring less amplification, decreasing the ASE noise. Indeed, for 32 inputs, the power penalty of the broadcast-and-select switch cannot be obtained due to error flooring. The proposed switch has a power penalty of 2.5 dB for 32 inputs due to the low insertion loss of the AWG devices.

Worst-case assumptions are made throughout; for example, all AWG crosstalk signals are at the same wavelength and polarization as the main signal. It is assumed that there is no penalty for wavelength conversion — this has been achieved in practice; indeed, interferometric wavelength converters can improve the signal shape, improving the above figures significantly. In any case, the results for the WASPNET switch clearly represent a significant improvement over the broadcast-and-select switch.

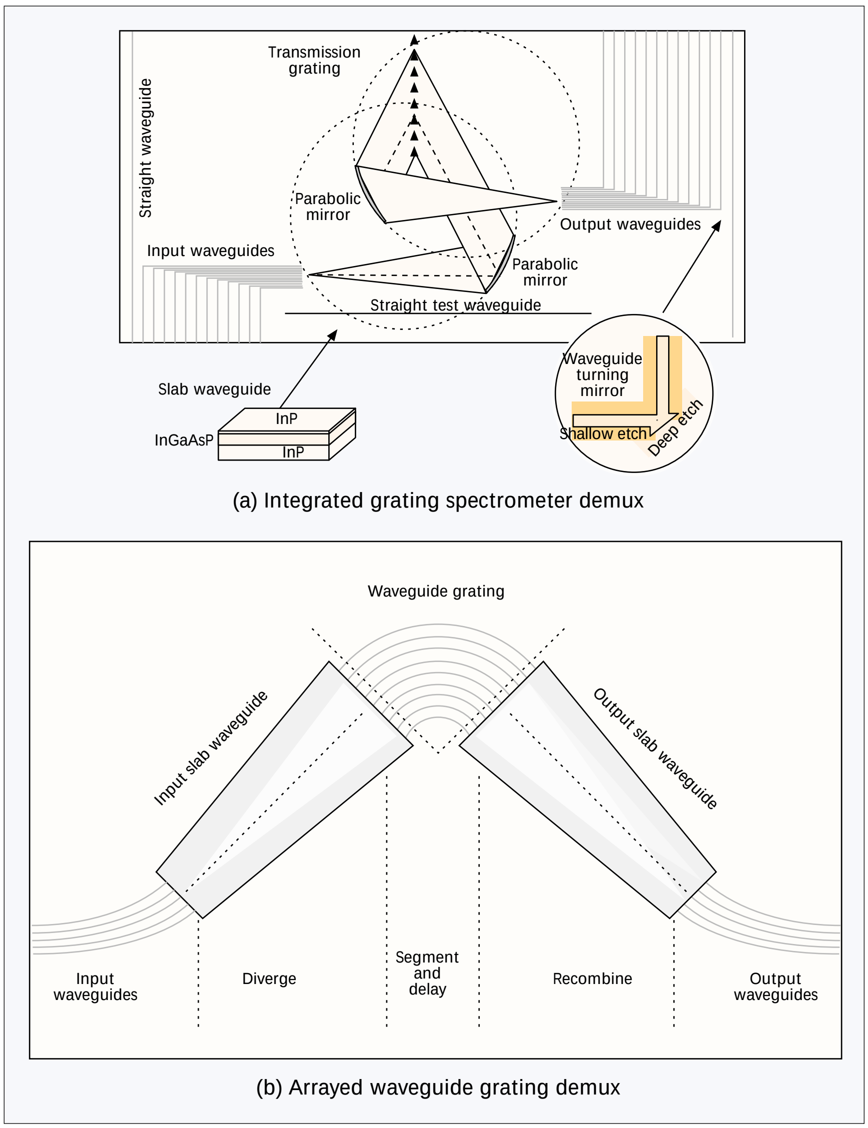

■ **Figure 4.** *Two types of wavelength router fabricated in WASPNET.*

Wavelength Router Devices

The most basic requirement in the above switch architecture — indeed in any WDM system — is to separate or combine signals on different wavelengths. There are a variety of methods for achieving this, and the choice is dictated by the needs of the system, be they technical, physical, or commercial. The most straightforward approach to the combination and separation of multiwavelength signals is to use optical combiners/splitters. In systems such as the WASPNET switch, where many wavelength channels are used, it is not appropriate to use a simple splitter or combiner approach. It is instead more convenient to have some method of separating and combining wavelengths in parallel without incurring splitting loss. Such a method has been in existence for over 200 years: the diffraction grating.

The basic physics of diffraction gratings allows them to act as a set of parallel filters, providing the functionality of a 1 x N demultiplexer. In addition, a grating can have both multiple inputs and outputs, which transforms it to an N x N wavelength router, bringing more exotic functionality. Each input may connect to any output purely by choice of wavelength, and if multiwavelength signals are sent to each input port, complete interconnection of ports can be achieved.

Devices that incorporate gratings for multiplexing, demultiplexing, and wavelength routing need to be compact and reliable. Integrated components are therefore preferable to the elaborate and bulky mechanical mountings used in conventional grating spectrometers. Two types of grating-based wavelength routers are being studied; both make use of semiconductor optical waveguide technology to confine light within small chips. The first takes the basic components of an optical spectrometer and places them onto a semiconductor chip in two-dimensional form (Fig. 4a).

The combined multiwavelength input signals are coupled into one of the input waveguides. This waveguide ends, allowing the light to propagate freely within the semiconductor. The diverging beam is then captured by a parabolic reflector, which both reflects and collimates the light ready for incidence on the grating. The grating, which is constructed from triangular recesses etched into the chip, separates the wavelengths in angle. Light emerging from the grating is then collected and focused by another parabolic reflector onto a set of output waveguides that lie at the focal points dictated by positions of maximum interference of each wavelength [7].

An alternative design is the AWG device, currently being used within the testbed, which has an array of optical waveguides as a grating (Fig. 4b). Each guide in the array differs from its nearest neighbors by a fixed length, constituting a fixed phase delay. The light from each waveguide is recombined at the other side of the array, where output guides are placed at the relevant position to couple an interference maximum for a particular wavelength [8].

Both these devices are compact and multiport. The integrated spectrometer is a 12 x 12 device, and the arrayed waveguide has 16 x 16 ports, although devices have been fabricated with as many as 128 ports. The functionality of these demultiplexer/router devices could form the basis for WDM systems such as WASPNET that utilize the multiwavelength aspect not only for increasing capacity, but for complex network functions as well.

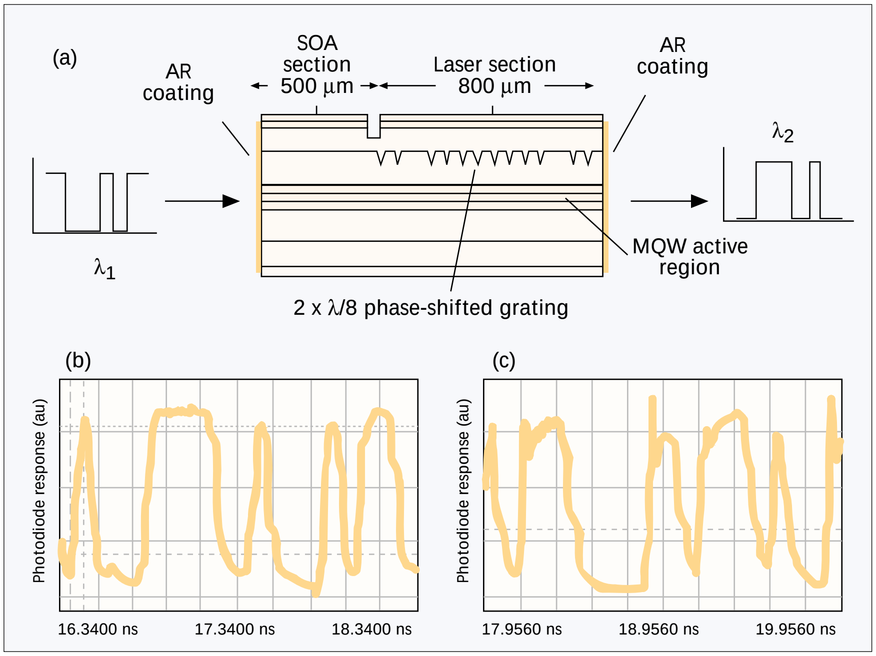

Figure 5. *(a) The SOA/DFB laser wavelength converter; (b) the input 9.953 Gb/s signal at 1559 nm; (c) the converted signal at 1553.5 nm.*

Wavelength Converters

Wavelength routing via AWGs is fundamental to the WASPNET node described above, although active space switching is also under consideration. Whichever option is eventually shown to be most effective (in terms of cost, complexity, and efficiency), a device to convert information from one wavelength to another is certain to be required. In the former case (i.e., in the switch architecture discussed above) these wavelength converters will be used after demultiplexing the input fiber wavelengths to give each packet the correct wavelength for subsequent routing by the AWG. In the latter case, wavelength conversion is necessary to help avoid wavelength contention when two wavelengths from different input fibers require switching to the same output fiber.

The wavelength converter itself must have several attributes. In order to function at the ever-increasing data rates of modern communications systems, it must operate at high speed, function with minimal sensitivity penalty to BER, and be cascadable. For the purposes of wavelength routing, the output wavelength of the device must be tunable, preferably over the whole Erbium-doped fiber amplifier window. Other requirements ideally include low input optical power operation, low component count, and low cost.

Under the WASPNET project, several wavelength conversion devices are being investigated. Here we discuss the use of an integrated semiconductor optical amplifier (SOA) with a distributed feedback (DFB) laser, a multiwavelength semiconductor laser, and also counter-propagating cross-gain modulation in SOAs for all-optical wavelength conversion.

The Integrated SOA/DFB Laser

This integrated device, fabricated by Nortel Technology (Fig. 5a), consists of an optical booster amplifier (SOA) monolithically integrated with a DFB laser. The output of the DFB section is at a discrete single wavelength λ_2. When an intensity-modulated digital signal at λ_1 requires conversion to λ_2, it is injected into the SOA section of the device, where it is boosted in power before injection into the DFB laser cavity. Here, when the logic level of the input signal at λ_1 is low (a 0), there is no effect on the DFB section, which continues to lase and thus emits a high logic level (a 1). However, when the input signal is high, the lasing of the DFB section is suppressed due to gain saturation, and thus the converter gives out a logical 0. Optical filtering is required at the output of the DFB laser section to separate the input and converted signals.

The presence of the booster amplifier allows the effective operation of the device for very low input powers. At the bit rate of 2.488 Gb/s, only 75 μW of coupled input power is required for 12 dB extinction ratio operation. Figures 5b and 5c show input and converted patterns at the higher bit rate of 9.953 Gb/s, which requires the higher power of 0.7mW, although this power is much lower than comparative schemes [9].

One disadvantage of the converter as described is the fact that the DFB laser section is fixed in wavelength. The output wavelength can be temperature tuned by ~0.1 nm/°C, although this technique is slow and limited. More promising is using a multisection DFB laser, which has been demonstrated at the University of Bristol to allow tuning of the converted signal over ~6 nm.

THE MULTIWAVELENGTH GRATING CAVITY LASER

The multiwavelength grating cavity (MGC) laser may be used as a wavelength converter, which shows many highly functional attributes for use in a WDM network node [10]. These are the ability to convert to up to four independent wavelengths (which can be designed to be on the International Telecommunication Union grid, 100 GHz apart) simultaneously, and also to combine space switching and wavelength conversion in one integrated device (Fig. 6a).

The device integrates an InP passive grating section with MQW active gain regions on the input (master) and the four output (slave) waveguides. The grating section incorporates parabolic collimating mirrors and a transmission grating etched into a passive slab waveguide. This introduces four wavelength-selective lasing cavities formed between the master and slave waveguides and each cleaved facet, centered at 1557.13 nm with 3.7 nm channel spacing. On biasing, the laser can emit on up to four different wavelengths simultaneously depending on which slaves are driven in conjunction with the master.

All-optical space switching and wavelength conversion may be achieved in the MGC laser by injecting an optical signal into the master stripe when the device is biased near threshold (Fig. 6b). Suppression of the corresponding lasing wavelength by gain saturation in the master stripe imprints an inverted copy of the input signal onto the MGC output in a similar manner to the integrated SOA/DFB laser wavelength converter. By monitoring each slave output, simultaneous wavelength conversion and 1 x 4 space switching may be achieved.

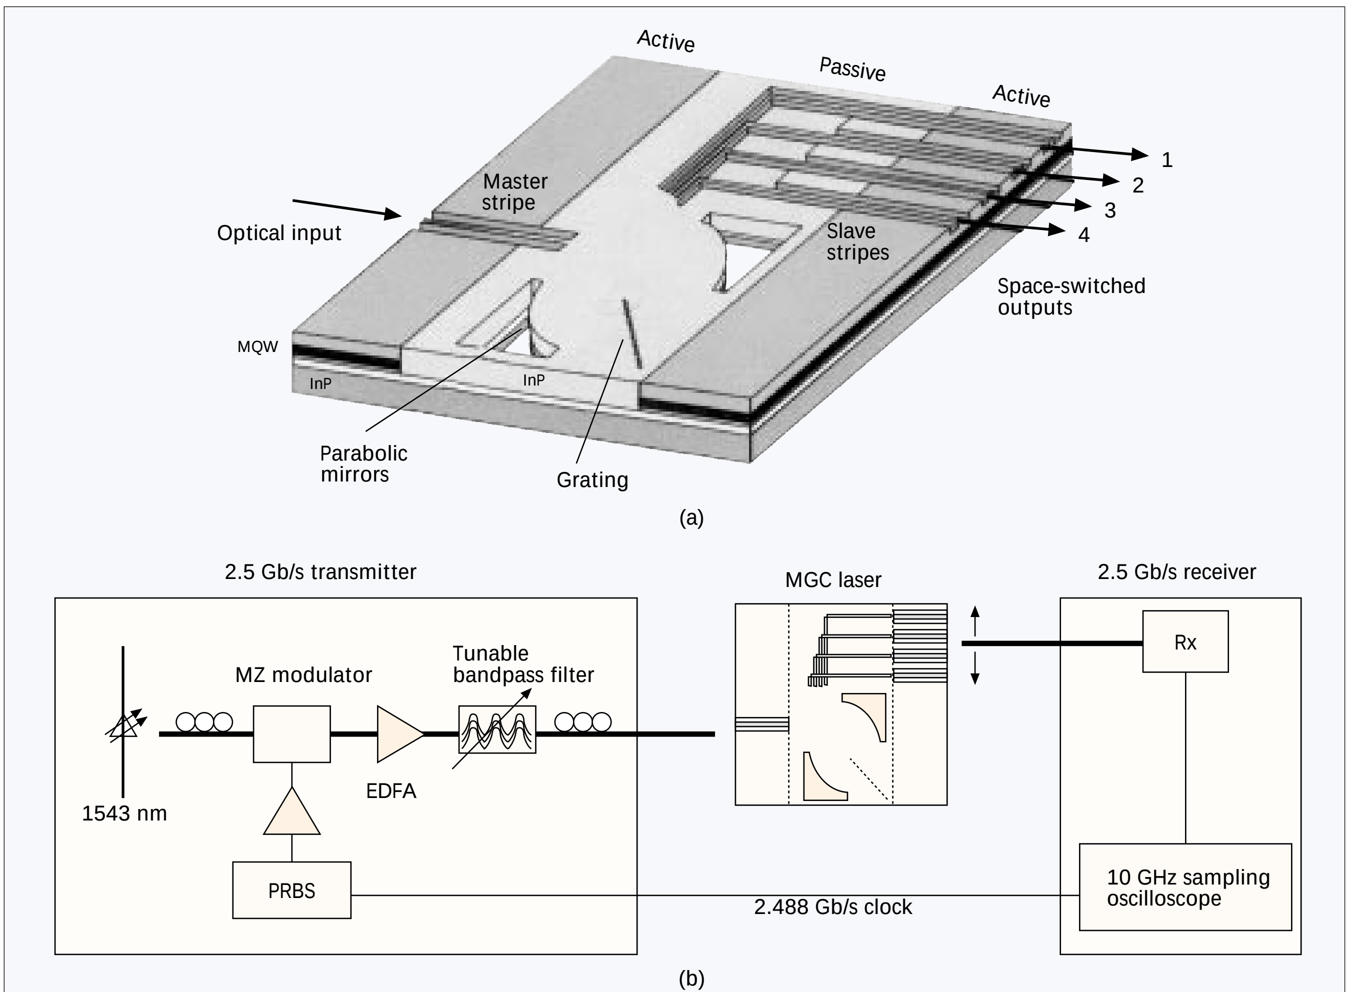

Figure 6. *(a) A diagram of the integrated MGC wavelength conversion device; (b) its use in an experimental setup demonstrating 2.488 Gb/s wavelength conversion and simultaneous 1 x 4 switching.*

CROSS-GAIN MODULATION IN SOAS

Cross-gain modulation (XGM) in SOAs is an alternative wavelength conversion technique (Fig. 7), and the one currently used in the testbed. The amplitude-modulated input signal at wavelength λ_1 induces gain variation in the SOA, causing translation of the input data to the constant wave (cw) wavelength λ_2; the converted signal is inverted. When the cw and signal wavelengths copropagate, an output filter is required to select λ_2, whereas with counter-propagation the output filter is eliminated. Clearly, this indicates the feasibility of fast tunable wavelength conversion to a single or multiple wavelength(s) (broadcasting). However, one disadvantage of counter-propagation is the reflection of a portion of the input signal at the converter output due to the residual SOA facet reflectivities [11]. The residual input signal is a source of optical crosstalk, and can be alleviated by appropriate architecture design [12].

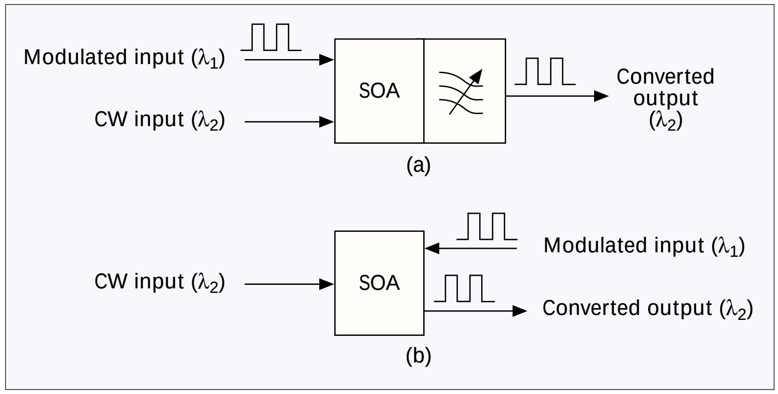

Figure 7. *Wavelength conversion using XGM in SOAs: (a) copropagation; (b) counter-propagation.*

THE OPTICAL PACKET TESTBED

Although system modeling will be used to evaluate the cascadability of a given node architecture, a photonic testbed yields insights into the practical problems and is currently being constructed at the University of Essex. The primary purpose of the testbed is to investigate the optical performance of a network comprising a number of cascaded nodes, each separated by a transmission link incorporating one or more optical amplifier repeaters. The total distance between the nodes can vary from 50–250 km, representing realistic internodal distances (Fig. 8). The principle of operation is similar to a circuit-switched recirculating loop in that the data is passed through a fixed chain of amplifiers a number of times to simulate a longer system. The optical packet testbed is a packet-switched version that allows the operator to predetermine the path of an individual packet through the switching matrix of each node and monitor it at various points throughout the network.

At the start of the experiment, the computer updates the routing table within the node. The header address of each transmitted packet has an entry in the routing table which contains the information necessary to determine the route of the packet through the switching matrix and its new address. A graphical user interface (GUI) allows the operator to specify how many nodes are to be traversed and the respective switch paths to be taken at each circulation. A particular packet is selected for monitoring (identified by a specific address) and is transmitted together with arbitrarily

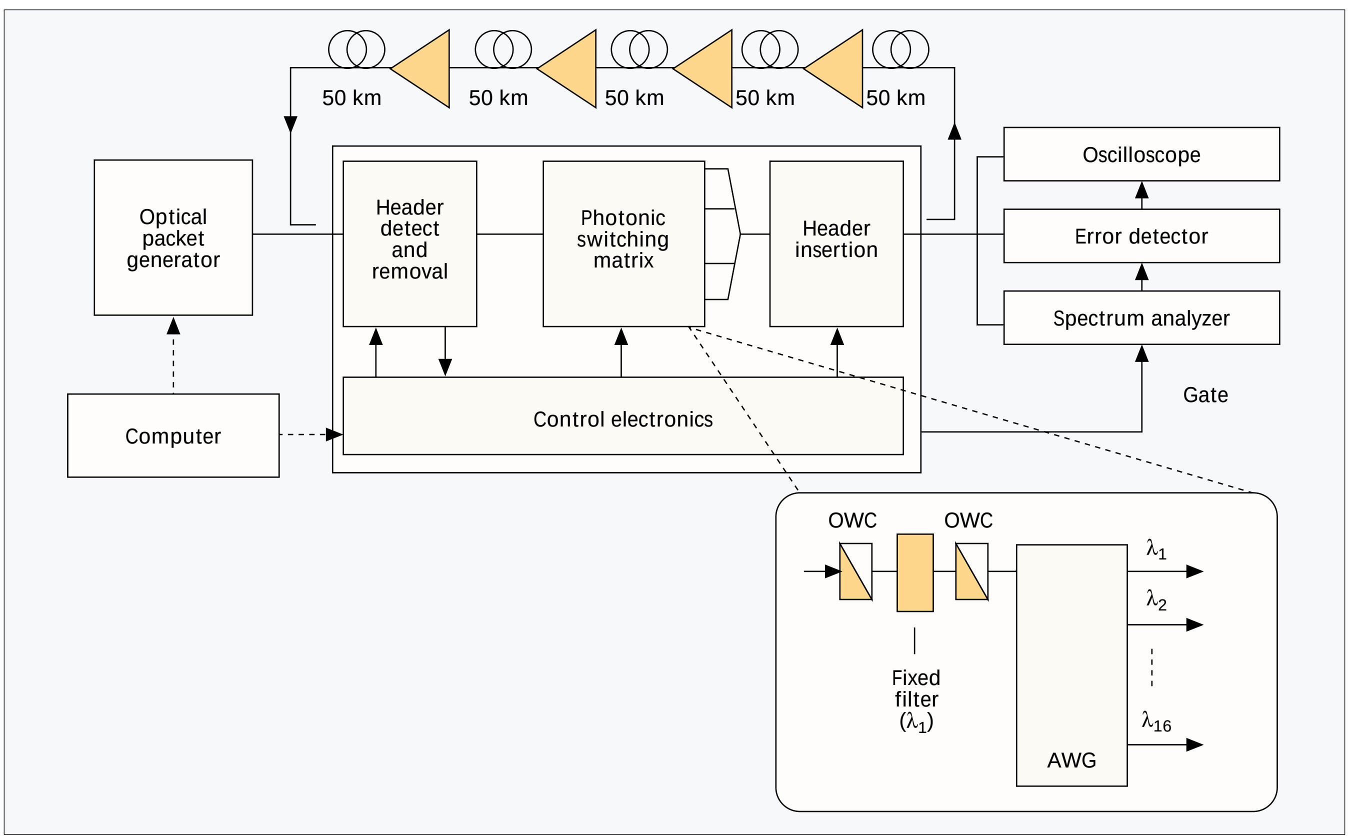

Figure 8. *A block diagram of the photonic ATM testbed.*

addressed dummy packets to fill the recirculating loop; once the loop is filled, the optical packet generator ceases transmission. At the node, the header of each incoming packet is decoded and removed by the header detection and removal module. The packets are routed through the photonic switching matrix according to the routing table within the control electronics.

The photonic switching matrix utilizes two counter-propagating optical wavelength converters (OWCs) using XGM in SOAs and a 16 x 16 AWG, as illustrated in the inset in Fig. 8. The dual-converter arrangement has the advantage of reducing the crosstalk associated with the SOAs and also ensures that the node imposes no data inversion [12]. The individual outputs of the switching matrix, which would normally be separate output ports in a real node, are combined and coupled to the header insertion module. A new header is inserted according to the routing table, and the packet leaves the node. The packet stream is fed back into the loop and also to the monitoring instruments. By appropriate gating of the instruments, the eye diagram, error performance, or optical spectrum can be obtained at the output of any specified node in the chain. The process is repeated a number of times and averaged over the duration of the experiment to achieve realistic results.

The optical packet testbed allows measurement of physical limitations associated with the node such as extinction ratio degradation and optical crosstalk, the two most debilitating effects in this type of router. Different header formats will be evaluated in the testbed for their robustness and reliability. The testbed also enables contention conditions to be established so that the ability of the node to resolve contention can be established and the physical implications associated with overcoming contention evaluated.

Conclusions

This article introduces WASPNET, a novel WDM optical packet transport network. A new network control scheme has been developed, which is flexible and simplifies optical hardware requirements. Like all packet-switching networks, WASPNET encounters packet contentions, but these are resolved not only by suitable node design but also by suitable network control, reducing node design complexity. SCWP realizes this requirement, reducing the node buffer size dramatically. Solutions to network control and management issues are presented, resolving the packet-ordering problem and overcoming the problem of restoration complexity.

Several schemes for packet header transmission are described, using subcarrier multiplexing, separate wavelengths, and serial transmission. It was shown that the SCM format facilitates packet delineation. Header implementation via other modulation formats is currently being investigated, and the most cost-effective and robust solution will be implemented in the final demonstrator.

A novel node design is introduced, based on wavelength router devices, which reduce loss, and hence booster amplifier gain and concomitant ASE noise. This is one of several variant nodes being investigated, some performing priority routing. While the power penalty of single nodes has been determined, future work will concern the feasibility of cascading many nodes, with a view to determining the network topology and hence the node size. The implementation of wavelength routers by means of integrated grating spectrometers and AWGs is described, and wavelength converters — which are essential companions to a wavelength router — are described, based on an integrated SOA/DFB laser, a multiwavelength grating cavity laser, and SOA cross-gain modulation.

The photonic packet-switching testbed will allow these concepts to be tested in practice, permitting the practical problems of their implementation to be determined. To conclude, in the WASPNET program a novel multiplexing technique for optical packets over WDM has been proposed and analyzed, and the ramifications of using this technique in practice are being investigated. Not only is node architecture being investigated, but also control, management, device fabrication, optical systems performance, and demonstrator construction. When the work under WASPNET is complete, it will provide a comprehensive blueprint for the construction and operation of future optical packet transport networks.

Acknowledgments

The authors wish to acknowledge funding from the U.K. Engineering and Physical Sciences Research Council (EPSRC), and also wish to thank the industrial collaborators, Fujitsu, BT, Nortel, and GPT, for their input. J. D. Bainbridge and M. C. Chia thank Nortel and Fujitsu, respectively, for personal support, while D. K. Hunter thanks the EPSRC for support in the form of an Advanced Fellowship.

References

[1] D. K. Hunter, M. C. Chia, and I. Andonovic, “Buffering in Optical Packet Switches,” *IEEE/OSA J. Lightwave Tech.*, Dec. 1998.

[2] C. Guillemot *et al.*, “A Two-Stage Transparent Packet Switch Architecture Based on Wavelength Conversion,” *Proc. ECOC ’95*, Brussels, Belgium, pp. 765–68.

[3] W. De Zhong and R. S. Tucker, “Wavelength Routing-Based Photonic Packet Buffers and Their Applications in Photonic Packet Switching Systems,” *IEEE/OSA J. Lightwave Tech.*, vol. 16, no. 10, Oct. 1998, pp. 1737–45.

[4] M. C. Chia, D. K. Hunter, and I. Andonovic, “Performance of Feedback and Feedforward AWG-Based Optical Packet Switches with WDM Inputs and Outputs,” All-Optical Networking, Architecture, Control and Management Issues, SPIE Int’l. Symp. Voice, Video and Data Commun.,, paper 3531–33 Nov. 1–5, 1998, Boston, MA.

[5] 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.

[6] Z. Haas, “The Staggering Switch — An Electronically Controlled Optical Packet Switch,” *IEEE/OSA J. Lightwave Tech.*, vol. 11, no. 5/6, May/June 1993, pp. 925–36.

[7] M. Gibbon *et al.*, “Optical Performance of Integrated 1.5 μm Grating Wavelength Demultiplexer on InP-Based Waveguide,” *Elect. Lett.*, vol. 25, no. 21, 1989, pp. 1441–42.

[8] C. Dragone, “An $N \times N$ Optical Multiplexer Using a Planar Arrangement of 2 Star Couplers,” *IEEE Photonics Tech. Lett.*, 1991, vol. 3, no. 9, pp. 812–15.

[9] M. F. C. Stephens *et al.*, “Low Input Power Wavelength Conversion at 10 Gb/s Using an Integrated Amplifier/DFB Laser and Subsequent Transmission over 375 km of Fibre,” *IEEE Photonics Tech. Lett.*, 1998, vol. 10, no. 6, pp. 878–80.

[10] M. Owen *et al.*, “All-Optical 1 x 4 Network Switching and Simultaneous Wavelength Conversion Using an Integrated Multi-Wavelength Laser,” *Proc. ECOC ’98*, Madrid, Spain, 1998.

[11] A. Tzanakaki *et al.*, “Penalty-Free Wavelength Conversion Using Cross-Gain Modulation in Semiconductor Laser Amplifiers with no Output Filter,” *Elect. Lett.*, vol. 33, no. 18, 1997, pp. 1554–56.

[12] K. M. Guild and M. J. O’Mahony, “Dual Wavelength Converter Architecture in an All-Optical Packet-Switched Routing Node,” *Elect. Lett.*, Oct. 1998.

Biographies

DAVID K. HUNTER [M] (d.hunter@eee.strath.ac.uk) is a senior research fellow at the University of Strathclyde, Scotland, and holds an EPSRC Advanced Fellowship. He obtained a B.Sc. in electronics and microprocessor engineering with first class honors in 1987, and a Ph.D. in 1991 for research into optical TDM switching architectures, both from the University of Strathclyde. His research interests include optical switch architectures, telecommunications switch modeling, network routing, and network survivability.

MOHAMED H. M. NIZAM (nizam@eee.strath.ac.uk) obtained a B.Eng. with first class honors in 1993, and a Ph.D. in 1998, both from the University of Strathclyde, becoming a member of staff at the same university in 1997. He is employed full-time on WASPNET, specializing in optical network design, control, and management. His research interests are network architecture, control and management design, performance analysis, and optical network design.

KEN M. GUILD (kmguil@essex.ac.uk) received an M.Eng. degree in electronic and electrical engineering from the University of Surrey, Guildford, in 1993. He worked as a design engineer at Alcatel Submarine Networks, Greenwich, for four years. He is currently pursuing a Ph.D. degree and working as a research officer at the University of Essex, Colchester, England. His research interests are photonic packet switching and WDM architectures in broadband networks.

JOHN D. BAINBRIDGE (J.D.Bainbridge@bristol.ac.uk) graduated from the University of Bath, United Kingdom, in 1992 with a degree in physics. He worked for three years at Nortel Optoelectronics, Paignton, as a development engineer on low-cost laser transmitters. He is just finishing a Ph.D. at Bristol University, sponsored by Nortel Technology, Harlow, on integrated optical demultiplexers and their applications to WDM systems.

MEOW C. CHIA (meow.chia@strath.ac.uk) received a Bachelor's degree in electrical and electronic engineering with first class honors from the University of Strathclyde, Scotland, in 1995. She is currently studying for a Ph.D. degree in the Electronic and Electrical Engineering Department, University of Strathclyde. Her research interests include optical packet switching node design, packet loss performance, and traffic analysis.

ANNA TZANAKAKI (annat@essex.ac.uk) received a B.Sc. degree in physics from the University of Crete, Greece, in 1992, and an M.Sc. degree in the physics of laser communications from the University of Essex, Colchester, England, in 1994. She is currently working toward a Ph.D. degree and is a research officer at the Network Research Center of the University of Essex. Her research interests are WDM networks, optical packet switching, and wavelength conversion.

MARC F. C. STEPHENS (M.F.C.Stephens@bristol.ac.uk) received a B.Sc. degree in physics (1st class honors) from Imperial College, London, in 1994. His postgraduate research into all-optical wavelength conversion techniques at the University of Bristol between 1995 and 1997 led to the award of a Ph.D. in 1998. He is currently working at Bristol on both the U.K. EPSRC WASPNET project and the European ACTS HIGHWAY project. He is a graduate member of the Institute of Physics.

RICHARD V. PENTY (Richard.Penty@bristol.ac.uk) received a Ph.D. in optical signal processing from Cambridge University in 1990, when he was appointed lecturer in the School of Physics, University of Bath. Since 1996 he has been with the Department of Electrical and Electronic Engineering, University of Bristol, becoming a reader in 1998. His current research interests include optical time-division and wavelength-division multiplexing, optical switching, wavelength conversion, semiconductor lasers, and optical routing.

MIKE J. O'MAHONY [SM] (mikej@essex.ac.uk) is head of the Electronic Systems Engineering Department at the University of Essex. He is also head of the Photonic Networks Laboratory, which is focused on the study and design of photonic networks. He is the author of over 150 papers relating to optical communications. He is project manager for the WASPNET project.

IVAN ANDONOVIC [SM] (i.andonovic@eee.strath.ac.uk), professor of broadband networks, joined the Electronic and Electrical Engineering Department at Strathclyde University in 1985, following a three-year period as a research scientist responsible for the design, manufacture, and test of optical guided wave devices for a variety of applications. His main interest centers on the determination of the role photonics will play in the evolution of future broadband networks.

IAN H. WHITE (Ian.White@bristol.ac.uk) joined the University of Bristol in 1996, previously being with the University of Bath as professor of physics. He is head of the Department of Electrical and Electronic Engineering and deputy director of the Center for Communications Research at the University of Bristol. His current research interests include local area networks using optical links, vertical cavity surface emitting lasers, high-power short-pulse laser diodes, and multiwavelength communication sources and systems.