ACCEPTED FROM OPEN CALL

Clos Lives On in Optical Packet Switching

Jan Cheyns, Chris Develder, Erik Van Breusegem, Didier Colle, Filip De Turck, Paul Lagasse, Mario Pickavet, and Piet Demeester, Ghent University — IMEC

ABSTRACT

While the technological evolution since Clos's seminal article on multistage switch architectures has been huge, his work and ideas still live on. In this article we discuss node architectures for optical packet switching and show how the multistage approach proposed by Clos can be adopted to solve scalability issues and construct switches with large port counts. As in the old days, the driving factors behind the introduction of multistage concepts also include economical issues: compared to a single-stage architecture, the number of components to realize the switching fabric will be reduced.

INTRODUCTION

In the era in which Clos produced his seminal paper [1], the dominant telecommunications technology was analog telephony using copper wire as the transport medium. Clos's work helped the development of large switches: the multistage architectures proved to need far fewer crosspoints. This reduction in the number of crosspoints was a very important issue, since it greatly determined the cost. Indeed, since Shockley, Bardeen, and Brattain had only developed the transistor a few years earlier, large-scale integration as we know it today was still a quite ambitious research topic. The main achievement of Clos was to circumvent technological boundaries, while also reducing the cost.

Communications technology has seen many advances since the days of Clos. Telephony is still here (albeit now mostly digital), but it is apparent that with the advent of the Internet, a large portion of traffic now consists of data rather than voice. Still, the concepts of the "old" telephony world are still in use. In essence, classical telephony is a circuit-switched concept: communication between two parties is realized by establishing a connection, which is reserved for their use only for the entire duration of their conversation. Prior to their communication, signaling takes place through the exchange of messages to set up the connection through the various switches on the path between the two parties. This same idea of connection-oriented communications prevails today, and a circuit-switched approach is also taken in so-called backbone networks to provide high-bandwidth interconnections between, for example, telephone private branch exchanges (PBXs). However, in the Internet world a packet-switched concept dominates. Instead of reserving a certain amount of bandwidth (a circuit) for a certain period of time, data is sent in packets. These packets have a header containing the information necessary for the switching nodes to be able to route them correctly, quite similar to postal services.

To provide the bandwidth necessary to fulfill the ever-increasing demand (Internet growth), the copper networks have been upgraded and nowadays to a great extent replaced with optical fiber networks. Since the advent of optical amplifiers (e.g., erbiumdoped fiber amplifiers, EDFAs) allowed the deployment of dense wavelength-division multiplexing (DWDM), the bandwidth available on a single fiber has grown significantly. Whereas at first these high-capacity links were mainly deployed as point-to-point interconnections, real optical networking using optical switches is possible today. The resulting optical communication network is still exploited in a circuit-switched manner: so-called lightpaths (making up an entire wavelength) are provisioned [2]. Optical crossconnects (OXCs) switch wavelengths from their input to output ports. To the client layer of the optical network, the connections realized by the network of OXCs are seen as a virtual topology, possibly different from the physical topology (containing WDM links), as indicated in Fig. 1. These links in the logical plane thus have wavelength capacity. To set up the connections, as in the old telephony world, a so-called control plane is necessary to allow for signaling. Enabling automatic setup of connections through such a control plane is the focus of the work in the automatically switched

 0163-6804/04/$20.00

optical network (ASON) framework. Since the lightpaths that have to be set up in such an ASON will have a relatively long lifetime (typically in the range of hours to days), the switching time requirements on OXCs are not very demanding.

It is clear that the main disadvantage of such circuit-switched networks is that they are not able to adequately cope with highly variable traffic. Since the capacity offered by a single wavelength ranges up to a few tens of gigabits per second, poor utilization of the available bandwidth is likely. A packet-switched concept, where bandwidth is only effectively consumed when data is being sent, clearly allows more efficient handling of traffic that greatly varies in both volume and communication endpoints, such as in currently dominant Internet traffic.

Therefore, during the last decade, various research groups have focused on optical packet switching (OPS), aimed at more efficiently using the huge bandwidths offered by WDM networks. The idea is to use optical fiber to transport optical packets rather than continuous streams of light, as sketched in Fig. 2. Optical packets consist of a header and a payload. In an OPS node, the transported data (payload) is kept in the optical domain, but the header information is extracted and processed using mature control electronics, as optical processing is still in its infancy. To limit the amount of header processing, client layer traffic (e.g., IP traffic) will be aggregated into fairly large packets. To unlock the possibilities of OPS, several issues arise and are being solved today. A major issue is the lack of optical random access memory (RAM), which would be very welcome to assist in contention resolution that arises when two or more packets simultaneously want to use the same outgoing switch port. Still, workarounds for the contention resolution problems have been found in optics [3]. Since the timescales at which a switch fabric needs to be (re)configured in OPS are much smaller than in, say, the ASON case, other switching technologies have been devised to unlock the possibilities of OPS.

These packet-switched networks can be operated in two different modes:

- Synchronous: Packets can only start at certain discrete moments in time; each timeslot, packets on different channels are aligned.
- Asynchronous: Packets can arrive at any moment in time, without any alignment.

In subsequent sections we will discuss the major architectures for OPS switches. To be competitive with other solutions (electronic or ASON-like), the OPS node cost needs to be limited, and the architectures should be future proof (i.e., scalable). In this context, the work of Clos has been inspiring: again, the driving factors that lead to multistage architectures were reducing switch complexity (thus cost) and circumventing technological constraints. We discuss these issues from an architectural design point, rather than elaborating on, say, packet scheduling and routing problems in multistage switches. Recent progress in solving those difficulties can be found in [4] and references therein.

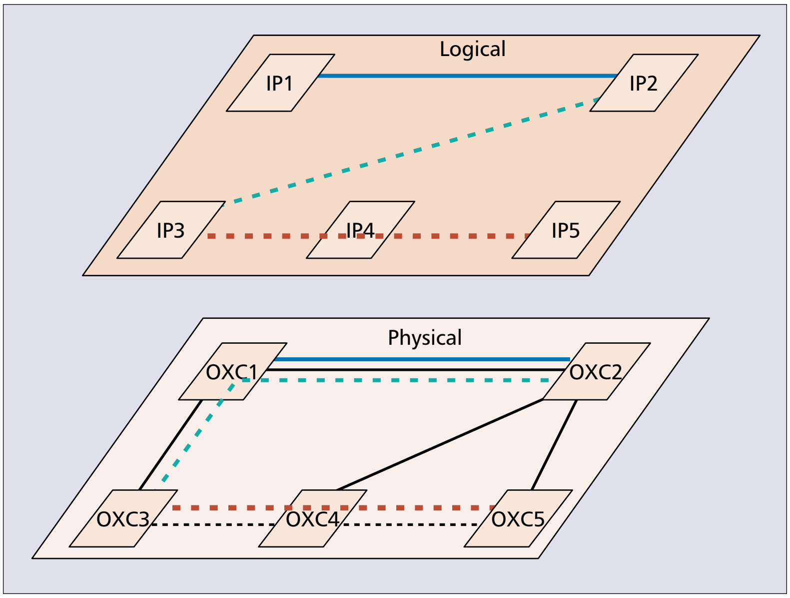

Figure 1. *Circuit switching with OXCs. Physical links (black lines) carry multiple wavelengths in (D)WDM, logical links consist of wavelength(s) on these fibers interconnected via OXCs, such as logical link IP2-IP3 using OXC1 (dotted).*

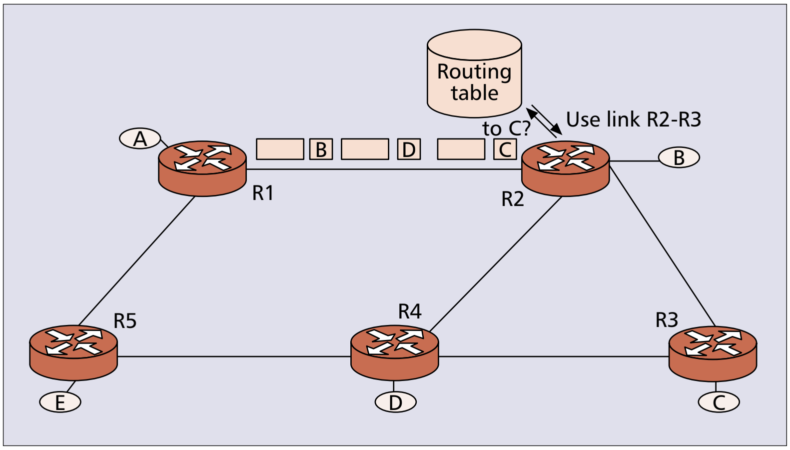

Figure 2. *Optical packet switching: a network with packets rather than the circuits in Fig. 1.*

Clos Design for Optical Networks

Obviously, similar challenges as encountered in OPS were faced for optical circuit-switched approaches. Here, we briefly outline recent work in the world of optical switching where Clos's ideas proved to be quite influential.

A thorough analysis of photonic multistage switches performing switching in space, time and wavelength dimensions was given by Thompson and Hunter [5]. Various combinations of these dimensions are studied, culminating in a full-blown OPS switch based on Clos networks of elementary switching blocks switching in one domain. Special attention is given to limitations in wavelength and time switching.

Other extensions of Clos's work have focused on expandability. Indeed, in Clos's multistage switches, there is a tight coupling between the size of the central submatrices and the number

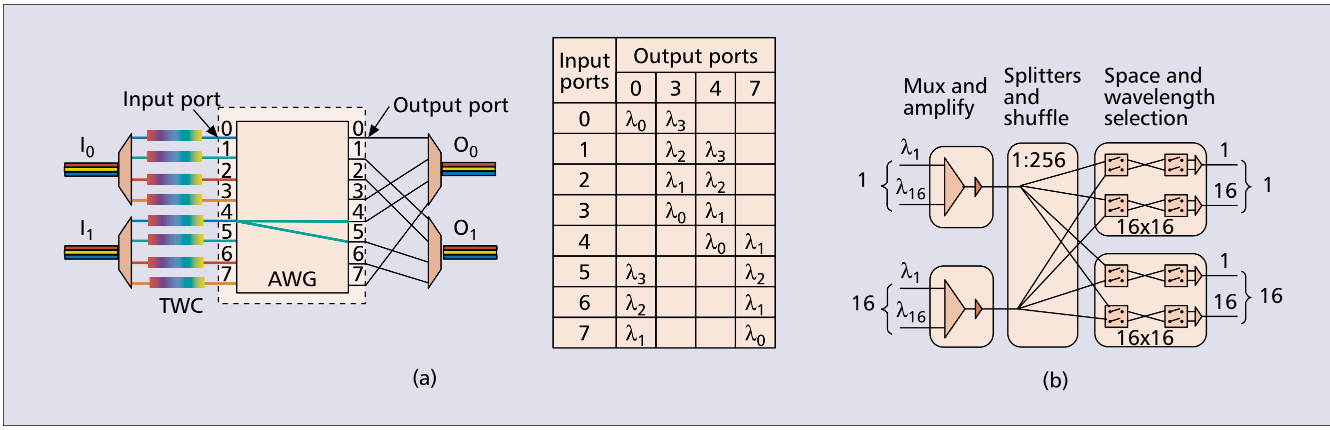

Input ports	Output ports			
	0	3	4	7
0	λ_0	λ_3		
1		λ_2	λ_3	
2		λ_1	λ_2	
3		λ_0	λ_1	
4			λ_0	λ_1
5	λ_3			λ_2
6	λ_2			λ_1
7	λ_1			λ_0

Figure 3. *Two well-known OPS architectures: a) an AWG-based switch used in STOLAS; b) the broadcast-and-select switch proposed in DAVID.*

of peripheral submatrices. In [6] MacDonald proposes to "distribute" the functionality of the central matrices into the peripheral matrices. In this way all building blocks of a node are equal (*SKOL* node), and adding one of these standard matrices can expand nodes. It alleviates the modularity problem of Clos architectures: the size of the building blocks depends on the final (maximal) size of the switch to be implemented and thus encompasses initial overbuilding. By distributing the central stages of a classical Clos architecture over SKOL in- and output modules, even though overbuilding is still required, the cost of an initial (partial) matrix configuration is significantly reduced.

For circuit-switched approaches, various authors start from Clos's ideas to exploit particular traffic characteristics to reduce the switch matrix sizes. The authors of [7] continue earlier work by others to reduce switch size for bidirectional traffic; a connection between A and B always implies a connection from B to A. Exploiting this bidirectionality allows significant cost cuts from traditional Clos networks. Similar approaches have been proposed for designs of multicast switches.

From a technological point of view, the multistage approach has been demonstrated in various domains. Micro-electromechanical systems (MEMS), using tiny mirrors (range of some tens of microns) to switch light from input to output ports, have also exploited Clos's basic ideas [8]. Such Clos-based MEMS solutions to date show rather poor reliability, especially when compared to electronic switches [9], but this is likely to improve as technology matures (meanwhile, it can be alleviated by adding some redundancy). Still, Clos design can be an important factor in lowering optical losses in MEMS optical switches [10].

To switch in the wavelength domain, fiber Bragg gratings (FBGs) are quite suitable because of their wavelength selective reflective properties [11]: wavelength switches can be realized by putting FBGs in series or parallel, and tunable approaches are also possible. Using them as building blocks in a Clos network, a large OXC can be built. Size limiting factors are physical impairments, including insertion loss and crosstalk.

Also, lithium-niobate-based switches have been proposed in a multistage architecture [12]. Since these switches are able to switch fast, they may be suitable for OPS. These switches have shown good behavior, particularly regarding number of crosspoints and insertion loss.

Despite a reasonably widespread use of the Clos approach for optical circuit-switched approaches, we found few traces in research literature of OPS architectures. In the next section we highlight the major OPS architectures, and subsequently illustrate how Clos's seminal work also proved useful in this OPS context.

Multistage Approaches to OPS

Node Architectures for OPS

One of the best known, or at least quite impressive, optical switching technologies is MEMS using tiny mirrors to deflect light from a particular input to a particular output port. Both 2D variants (where mirrors are either tilted up or lie down and let light pass) and 3D variants (with mirrors tilting along two axes) have been demonstrated. While the characteristics in terms of optical signal quality distortion are quite good, this approach is not feasible in an OPS concept where very fast switching times (range of nanoseconds) are mandatory. Two widespread approaches are one based on arrayed waveguide grating (AWG) with tunable wavelength converters (TWCs), and another based on a broadcast-and-select (B&S) concept using, for example, SOA technology.

The AWG approach is also studied in the European research project STOLAS [13]. An AWG component has the interesting feature that when light is inserted via one of its input ports, the output port it will come out of depends on the wavelength used. In Fig. 3a, a signal entering on input port 4 will end up at either output port 4 or 5, when it is on the blue or green wavelength, respectively. Thus, by providing wavelength converters at the AWG's inputs, we can exploit the structure as a space switch. What wavelength to use to reach a particular output from a given input can be found by a table lookup operation: a sample table and switch are illustrated in Fig. 3a.

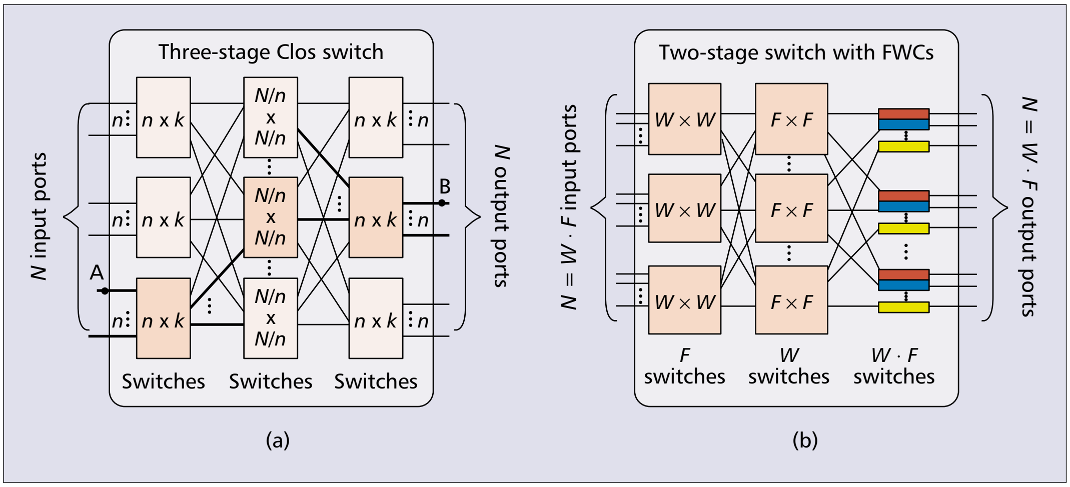

Figure 4. *Multistage switches: a) a three-stage Clos switch; b) a switch with two switching stages and fixed output wavelength converters.*

The main advantage of the B&S architecture clearly is its inherent multicast capability, which the AWG approach lacks. However, the asset of the AWG-based architecture is that it relies on a passive component and does not suffer from splitting losses as the B&S does.

The B&S approach is deployed in the recent research project DAVID [14]. The switch fabric's architecture is sketched in Fig. 3b. It comprises several subblocks. In the first block a couple of input ports that use different wavelengths are multiplexed into a single optical fiber. Each of these fiber signals is broadcast through a splitting stage to each of the output ports. Using two successive SOA stages, a single wavelength signal is kept per output port. The first SOA array is used to select only one of the input fiber signals for each output port. The second selection stage uses an SOA array and a wavelength-selective component to keep only a single wavelength per output port.

The main advantage of the B&S architecture clearly is its inherent multicast capability, which the AWG approach lacks. However, the asset of the AWG-based architecture is that it relies on a passive component and does not suffer from splitting losses as the B&S does.

Clos Applied to OPS

In both the B&S and AWG approaches, scalability issues will arise, as will be discussed further in this article. A solution is to employ multistage architectures similar to the ones proposed by Clos. Let us first define the terminology on blocking we will adopt in the remainder of the article. We call a switching architecture *strictly nonblocking* when it is always possible to connect any idle input port to any idle output port irrespective of other connections that already present. A switch is *rearrangable nonblocking* if it is possible to connect any idle input port to any idle output port, but some of the existing connections have to be reconfigured in order to do so. After the reconfiguration all connections are functional again. When a switch cannot guarantee to be able to always connect an idle input to an idle output port, it is said to be *internally blocking*.

The scheme of a three-stage Clos switch is outlined in Fig. 4a for an $N \times N$ switch. The N input ports are grouped per n, and both the first and third stage have a switching fabric for each such group of n ports. The second stage contains k switches, which each in turn are connected to each of the N/n first- and third-stage switches. Thus, the first stage comprises $n \times k$ switches, the second $N/n \times N/n$, and the last one $k \times n$. Now, a crucial question is how many second-stage switches have to be foreseen. The answer to this question depends on the degree of blocking we want to realize.

In circuit switching, it is clear that the lifetimes of circuits may overlap, but the start and end times will most likely not coincide: thus, once we have chosen to route a connection from input A to output B along a certain second-stage switch, we have to stick to this choice for the entire duration of the connection. Thus, the switch needs to be strictly nonblocking. However, with synchronous OPS, we have a packet switching concept where the switch adopts a slotted mode of operation; that is, each timeslot the packets at the inputs are inspected and switched jointly to the appropriate output. The next timeslot, all these packets are finished, and the switch may be completely reconfigured. It is clear that in this case of synchronous OPS, it is sufficient to have a rearrangeable nonblocking switch: for each slot in turn we can choose the second-stage switch.

As Clos showed in his seminal paper, to achieve strict nonblocking, the number of second-stage switches needs to be $k \geq 2n - 1$. However, for a rearrangeable nonblocking switch, $k \geq n$ suffices. (For a theoretical proof of these conditions, we refer to [15].) Thus, we need only about half as many second-stage switches for synchronous OPS than for asynchronous OPS.

Now, in OPS, part of the solution to contention resolution (see above) is to employ wavelength conversion: when two or more packets need to be switched to the same outgoing fiber, one or more of them may be converted to another wavelength to allow their simultaneous transmission on the output fiber. So in packet switching we are not interested in on which exact wavelength channel the packet is put; we only want the correct output fiber. This allows a simplification of design: if we choose to have all outputs of a third-stage switch going to the same output fiber (thus, e.g., $n = W$, with W the number of wavelengths per fiber), we can replace the

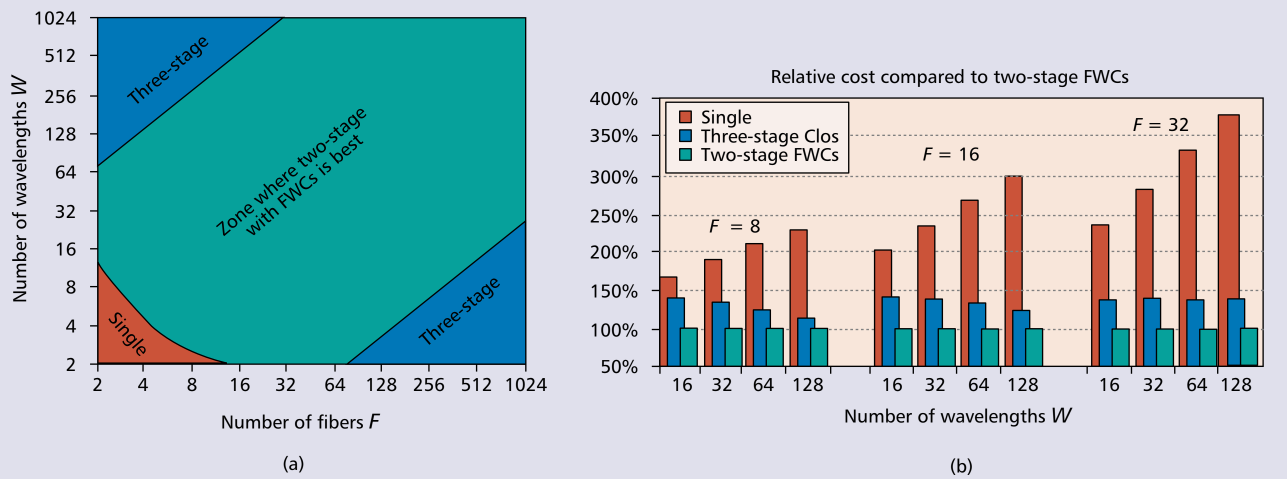

Figure 5. *The advantages of a Clos-like approach for B&S switches in OPS: a) dimensions for which a multistage architecture requires fewer SOAs; b) a comparison of the required number of SOAs per architecture.*

third-stage switch with fixed output wavelength converters (FWCs). An FWC converts any incoming wavelength into a predefined (thus fixed) wavelength. Thus, we obtain a three-stage switch architecture with only two stages comprising smaller (full) switch fabrics and one with only FWCs, as in Fig. 4b.

Reducing the Number of SOAs for a B&S Switch

The major impairment of the B&S switch architecture outlined in Fig. 3b is the splitting stage, which degrades the optical signal. It is clear that this will limit scaling this architecture to very large port counts. By combining smaller-sized switches in the Clos-like multistage approaches of Fig. 4 — obviously with some regeneration stages in between — we can overcome this problem. From a cost perspective, we may assume that the number of SOA gates used gives a good indication. Thus, in the following we compare three different architectures in terms of number of SOA gates used:

- Single-stage
- Three-stage Clos
- Two-stage with wavelength converters

The architecture of the DAVID switching fabric was discussed earlier and sketched in Fig. 3b. The number of SOA gates needed to construct a single-stage $N \times N$ switch is given in Eq. 1: for each of the N output ports, N/w gates are needed for space selection, while w gates are needed for wavelength selection. Since the switching matrix will be surrounded by wavelength converters (actually 3R regenerators), the number of wavelengths w can be optimized (and chosen different from W, the number of wavelengths on the input/output fibers) to minimize the number of SOA gates. The optimal choice is $w = N^{1/2}$, which leads to the minimal number of SOA gates for a single-stage switch.

$$s(N,w) = N \cdot (N/w + w). \qquad (1)$$

For OPS switches, we have indicated that the number of second-stage switches needed to provide a nonblocking fabric to operate in slotted mode is $k = n$. The optimization of n to reduce the number of SOA gates in the overall multistage architecture leads to the choice $n = 0.5 \cdot N^{1/2}$. In the proposed two-stage architecture the number of SOA gates can also easily be calculated.

Inspection of the number of SOA gates needed leads to the choices illustrated in Fig. 5a. On the horizontal axis, the number of I/O fibers F is indicated, whereas the vertical axis denotes the number of wavelengths W per fiber of our OPS switch. For a particular combination of F and W, the color of the zone where the point with coordinates (F,W) belongs shows the cheapest choice in terms of number of SOAs (e.g., a switch with 8 fibers and 16 wavelengths/fiber requires the fewest SOAs in a two-stage architecture with FWCs). From this graph, we see that the single-stage architecture is only advantageous for small port counts ($N = W \cdot F < 12$). The two-stage architecture with FWCs is generally the cheapest choice: the three-stage Clos architecture can only be advantageous when either the number of wavelengths per fiber or the number of input/output fibers is very large. The amount of cost reduction for a selected range of switches is shown in Fig. 5b. This illustrates that for 8–32 fibers with 16–128 wavelengths each, a three-stage Clos architecture needs 15–40 percent more SOA gates, while a single stage architecture would need 65–275 percent more SOAs than the two-stage architecture (or even be infeasible because of the high splitting factor in the broadcast block).

From this case study it is apparent that not only to solve scalability issues but also for cost reduction purposes, Clos-like approaches can be very helpful for B&S-based OPS switches.

A Strictly Nonblocking AWG-Based Switch for Asynchronous Operation

The STOLAS project uses the AWG-based approach outlined earlier in Fig. 3a. The multiple (W) wavelength channels carried in

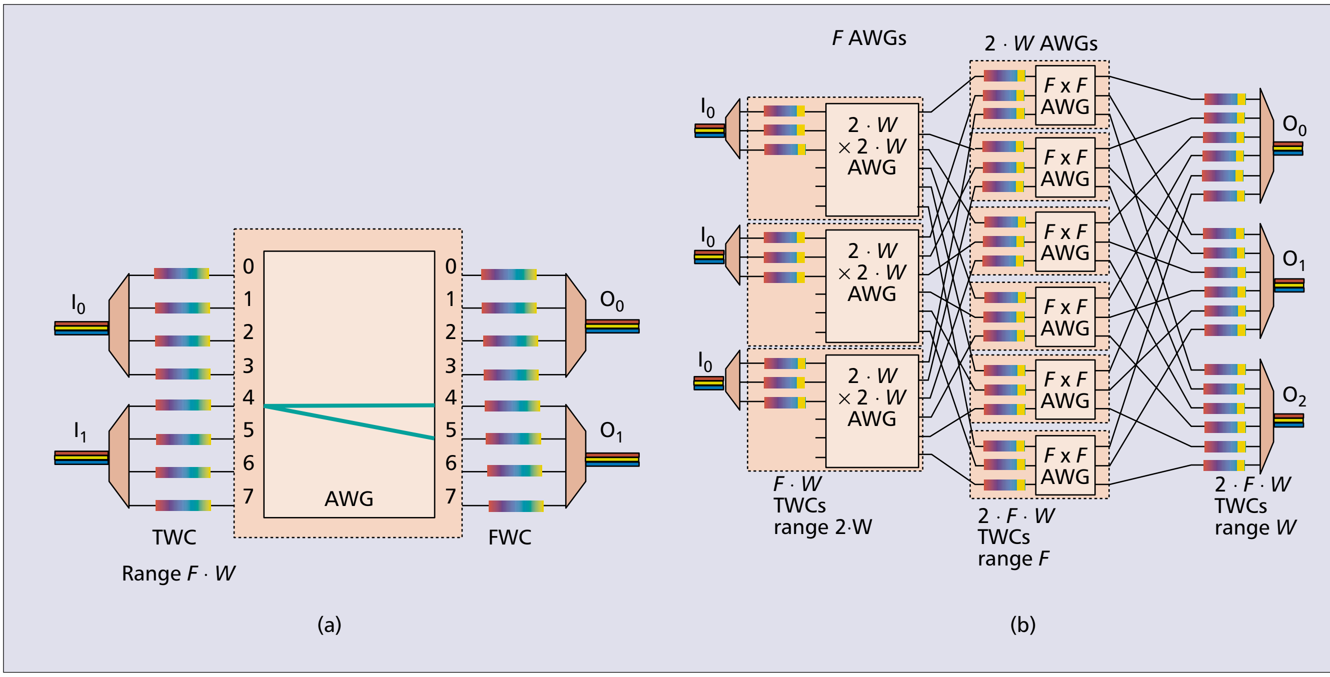

■ **Figure 6.** *A strictly nonblocking switch architecture with AWGs and wavelength converters: a) single-stage; b) multistage.*

(D)WDM on incoming fibers are demultiplexed, and each of them is led through a tunable wavelength converter to control the output port of the AWG to which it needs to be switched. The outputs of the AWG are then coupled into output fibers. Since the set of wavelengths used on in- and output fibers should be the same, the range of the TWCs should not exceed those W wavelengths. However, the design in Fig. 3a leads to an internally blocking switch. Still, when the switch is used for slotted OPS, we have shown that the internal blocking can be overcome and the performance is very close to that of a rearrangeable nonblocking switch [16]. However, for asynchronous switching, the blocking problem cannot easily be alleviated [17].

To construct a strictly nonblocking switch with an AWG for asynchronous operation, we need to increase the range of the input TWCs to $F \cdot W$, that is, as many wavelengths need to be used as there are switch ports. To limit the wavelength range on the output fibers to W, we then need to provide output wavelength converters. Those output converters can be fixed output wavelength converters (FWCs). The resulting configuration is shown in Fig. 6a.

The non-blocking switch's requirement of TWCs with range $F \cdot W$ raises a scalability issue. It is quite intuitive that the technological evolution of the range of wavelengths for tunable transmitters (the core part of a TWC) will closely follow the increase in number of wavelengths used on the fibers. Thus, for the blocking node where only a range of W is required for the TWCs, we foresee no serious scalability problem. However, when the range needs to be extended to $F \cdot W$, this may be an issue, certainly when a large number of fibers F is involved.

To overcome this scalability limit, a Clos-like multistage design can be helpful. The eventual switch design is depicted in Fig. 6b, which is similar to the generic structure of Fig. 4b presented earlier: a first switching stage comprises $W \times 2 \cdot W$ switches, a second consists of $F \times F$ switches, and the last stage only contains TWCs. As we are designing a strictly nonblocking node, the converters at the output can no longer be fixed output wavelength converters. The range of the TWCs for each of the three stages is $2 \cdot W$, F, and W. When we assume that the range of the TWCs is limited to W wavelengths, we end up with the condition that $F \leq W/2$. This is a quite realistic assumption, given the fairly broad range of wavelengths available already today. The multistage design of Fig. 6b also uses smaller-sized AWGs than the single-stage approach, which allows overcoming potential technological limitations.

To conclude this discussion of the multistage AWG architecture, note that the Clos-like approach did not lead to a reduction of the number of components or crosspoints, as in most other Clos-based approaches. Indeed, we now need $5 \cdot F \cdot W$ converters, whereas the original design (recall Fig. 6a) needed only $2 \cdot F \cdot W$, of which $F \cdot W$ were fixed output (which are cheaper). However, the demands on those $F \cdot W$ converters were quite unrealistic in terms of tunability range. Thus, the Clos approach's advantage in this case is that it enables technological limitations to be circumvented. Even though TWCs are, at this point in time, rather complex and thus expensive devices, we do expect that their cost will drop sharply. Indeed, research on these devices continues and integration of the converters with tunable lasers has already been proposed in [19, 20], allowing production at a substantially lower price. Thus, a TWC seems a viable candidate component for usage in OPS, being a technology for the mid- to long-term

Even though TWCs are, at this point in time, rather complex and thus expensive devices, we do expect that their cost will drop sharply. Indeed, research on these devices continues and integration of the converters with tunable lasers has already been proposed.

future. An additional quality of wavelength conversion particularly useful in the multistage solutions at hand is its side-effect of amplification.

CONCLUSION

Despite the fact that both information characteristics and communication technology have greatly evolved since the time of Clos's seminal paper, his ideas on multistage switches still prove to be very useful. In this article we have focused on their application in optical networking. We have outlined a range of examples in the field of circuit switching, and then focus on Clos-like design in optical packet switching.

We presented the two most widespread architectures for OPS: broadcast-and-select switches using SOAs, and AWG-based switches. The former profits from a Clos-like multistage architecture to reduce the number of SOA gates needed and enlarge the switch size to high port counts. The AWG-based design was shown to be prone to internal blocking when the tunability range of wavelength converters is limited. To overcome this blocking problem, we showed that a multistage design inspired by Clos networks offers a viable solution.

As in the "old days," multistage approaches thus are still very useful to either reduce costs (i.e., the number of components used) or circumvent technological limitations.

ACKNOWLEDGMENTS

This work was partly financed by the European IST projects STOLAS (IST-2000-28557) and DAVID (IST-1999-11387), and the national GBOU-project IWT 010058. Jan Cheyns and Chris Develder are research assistants of the Fund for Scientific Research — Flanders, Belgium (FWO-Vlaanderen); Filip De Turck acknowledges the FWO-Vlaanderen for their support through a postdoctoral fellowship. Erik Van Breusegem and Didier Colle would like to thank the IWT for its financial support through their Ph.D. and postdoctoral grants.

REFERENCES

[1] C. Clos, "A Study of Nonblocking Switching Networks," *Bell Sys. Tech. J.*, vol. 32, Mar. 1953, pp. 406–24.
[2] R. Ramaswami and K. N. Sivarajan, "Design of Logical Topologies for Wavelength-Routed Optical Networks," *IEEE JSAC*, vol. 14, no. 5, June 1996, pp. 840–51.
[3] S. Yao *et al.*, "All-Optical Packet-Switched Networks: A Study of Contention-Resolution Schemes in an Irregular Mesh Network with Variable-Sized Packets," *Proc. SPIE*, vol. 4233, OPTICOMM 2000, Plano, TX., Oct. 2000.
[4] H. J. Chao, K.-L. Deng, and Z. Jing, "Petastar: A Petabit Photonic Packet Switch," *IEEE JSAC*, vol. 21, no. 7, Sept. 2003, pp.1096–1112.
[5] R. A. Thompson and D. K. Hunter, "Elementary Photonic Switching Modules in Three Divisions," *IEEE JSAC*, vol. 14, no. 2, Feb. 1996, pp. 362–73.
[6] R. I. MacDonald, "Large Modular Expandable Optical Switching Matrices," *IEEE Photon. Tech.. Lett.*, vol. 11, no. 6, Jun. 1999, pp. 668–70.
[7] O. Gerstel and P. Lin, "Efficient Architecture for a Duplex Multi-Stage Optical Non-Blocking Switch," *Proc. OFC 2000*, vol. 2, Baltimore, MD, Mar. 2000, pp. 350–52.
[8] I. Lelic, "Large Multistage OXC," Proc. ECOC 2001, Amsterdam, The Netherlands, Sept. 2001, pp. 540–41.
[9] L. Wosinska, L. Thylen, and R. Holmstrom, "Large-Capacity Strictly Nonblocking Optical Cross-Connects Based on Microelectrooptomechanical Systems (MEOMS) Switch Matrices: Reliability Performance Analysis," *J. Lightwave Tech.*, vol. 19, no. 8, Aug. 2001, pp. 1065–75.
[10] G. Shen *et al.*, "Architectural Design for Multistage 2-D MEMS Optical Switches," *J. Lightwave Tech.*, vol. 20, no. 2, Feb. 2002, pp. 178–87.
[11] Y.-K. Chen and C.-C. Lee, "Fiber-Bragg Grating-Based Large Nonblocking Multiwavelength Cross-Connects," *J. Lightwave Tech.*, vol.16, no. 10, Oct. 1998, pp. 1746–56.
[12] F.M. Suliman, A. B. Mohammed, and K. Seman, "A New Nonblocking Photonic Switching Network," *Proc. GLOBECOM 2001*, vol. 4, San Antonio, TX, Nov. 2001, pp. 2071–76.
[13] K. Vlachos, *et al.*, "An Optical IM/FSK Coding Technique for the Implementation of a Label-Controlled Arrayed Waveguide Packet Router," *J. Lightwave Tech.*, vol. 21, no. 11, Nov. 2003, pp. 2617–28.
[14] L. Dittmannet *al.*, Eds., "The European IST Project DAVID: A Viable Approach towards Optical Packet Switching," *IEEE JSAC*, Special Issue on High-Performance Optical/Electronic Switches/Routers for High-Speed Internet, vol. 21, no. 7, Sep. 2003, pp. 1026–40.
[15] J. Y. Hui, *Switching and Traffic Theory for Integrated Broadband Networks*, Chapter 3, "Point-to-Point Multi-Stage Circuit Switching," Kluwer, 1990, pp. 53–83.
[16] J. Cheyns *et al.*, "Routing in an AWG Based Optical Packet Switch," *Phot. Net. Commun.*, vol. 5, no. 1, Jan. 2003, pp. 69–80.
[17] J. Cheyns *et al.*, "Performance improvement of an internally-blocking optical packet/burst switch," *Proc. ICC 2003*, Anchorage, AK, vol. 2, May 11–15, 2003, pp. 1304–8.
[18] R. Hanfoug *et al.*, "Wavelength Conversion with Polarisation Labelling for Rejection and Isolation of Signals (POLARIS)," *Proc. 11th Euro. Conf. Integrated Optics*, Prague, Czech Republic, April 2–4, 2003, pp. 105–8.
[19] R.G. Broeke *et al.*, "All-Optical Wavelength Converter with a Monolithically Integrated Digitally Tunable Laser," *Proc. ECOC '02*, Copenhagen, Denmark, 2002, pp. 1–2.
[20] M. L. Masanovic *et al.*, "Monolithically Integrated Mach-Zehnder Interferometer Wavelength Converter and Widely Tunable Laser in InP," *IEEE Photon. Tech. Lett.*, vol. 15, no. 8, Aug. 2003, pp.1117–19.

BIOGRAPHIES

JAN CHEYNS (jan.cheyns@intec.Ugent.be) received an M. Sc. degree in electrical engineering (1998) and an M. Sc. degree in physics engineering (1999) from Ghent University. In 2000 he was granted a Fund for Scientific Research of Flanders (FWO-Vlaanderen) scholarship to prepare a Ph.D., toward which he is working within the research group INTEC Broadband Communication Networks (IBCN) of the Department of Information Technology (INTEC) of Ghent University. His main interests are packet/burst switched and hybrid optical networks. He is currently involved in the European IST project STOLAS and the national research project "Optical Networking and Node Architectures."

CHRIS DEVELDER (chris.develder@@intec.Ugent.be) received his M.Sc. in computer science engineering and a Ph.D. in electrotechnical engineering from Ghent University, Belgium in 1999 and 2003, respectively. Since October 1999 he has been working in the Department of Information Technology of the same university, as a researcher for FWO-Vlaanderenin the field of network design and planning, mainly focusing on optical packet switched networks. As such, he has been involved in the European IST projects DAVID and STOLAS, as well as the national research project "Optical Networking and Node Architectures."

ERIK VAN BREUSEGEM (erik.vanbreusegem@intec.Ugent.be) received his M. Sc. in computer science in 2001 from Ghent University, Belgium. In August 2001 he joined the Department of Information Technology of the same university. In January 2002 he was granted an IWT scholarship to work on a Ph.D., mainly focused on optical packet/burst switched and hybrid optical networks. He is currently involved in the European IST projects DAVID and STOLA, as well as a national research project on Optical Networking and Node Architectures. Since 2004 he is with OPNET Technologies.

DIDIER COLLE (didier.colle@intec.Ugent.be) received an M.Sc. degree in electrotechnical engineering (option: communications) from Ghent University in 1997. Since then, he has been working at the same university as a researcher in the Department of Information Technology (INTEC). He is part of the research group INTEC Broadband Communication Networks (IBCN). His research led to a Ph.D degree in

February 2002. In January 2003 he was granted a postdoctoral scholarship from the Instituut voor de aanmoediging van Innovatie door Wetenschap en Technologie in Vlaanderen (IWT-Vlaanderen). His research deals with design and planning of communication networks. This work is focused on optical transport networks to support the next-generation Internet. He has actively been involved in four IST projects (LION, OPTIMIST, DAVID, and STOLAS) and in the COST266 action. His work has been published in more than 70 scientific publications in international conferences and journals.

FILIP DE TURCK (filip.deturck@intec.Ugent.be) received his M.Sc. degree in electronic engineering from Ghent University in June 1997. In May 2002 he obtained his Ph.D. degree in electronic engineering from the same university. From October 1997 to September 2001 he was a research assistant with the FWO-Vlaanderen. He is currently a postdoctoral fellow of the FWO-Vlaanderen and affiliated with the Department of Information Technology of Ghent University. His main research interests include scalable software architectures for telecommunication network and service management, performance evaluation and optimization of routing, admission control, and traffic management in telecommunication systems.

PAUL LAGASSE (lagasse@intec.Ugent.be) received his M.Sc. degree in electrical engineering in 1969 and his Ph.D. degree in 1972, both from Ghent University. In 1981 he became a professor of electrical engineering at Ghent University where he is now head of INTEC. After originally doing research in the areas of surface acoustic waves, high-frequency technology, and opto-electronics, he is now managing a research group of more than 100 academics working in the fields of telecommunications and information technology. Currently he is also involved in assessing R&D policies and programs on a national level as chairman of the Board of the Flemish Institute for Science and Technology and on a European level as a member of the Advisory Group of the IST program.

MARIO PICKAVET (mario.pickavet@intec.Ugent.be) is a full-time professor at Ghent University, where he teaches telecommunications network courses. His current research interests are related to broadband communication networks (WDM, IP, (G-)MPLS, OPS, OBS). In this context, he is currently involved in the European IST projects DAVID, LASAGNE, and ePHOTON/ONe. He has published several international publications, both in journals and magazines, such as *IEEE JSAC*, *IEEE Communications Magazine*, *European Transactions on Telecommunications*, and *Photonic Network Communications*, and conference proceedings.

PIET DEMEESTER (piet.demeester@intec.Ugent.be) received his Ph.D. degree from Ghent University at INTEC in 1988. He became a professor at Ghent University in 1993 where he is responsible for research on communication networks. He was involved in several European COST, ESPRIT, RACE, ACTS, and IST projects. His current interests are related to broadband communication networks (IP, GMPLS, optical packet and burst switching, access and residential, active, mobile, CDN, grid) and include network planning, network and service management, telecom software, internetworking, and network protocols for QoS support. He has published over 200 journal or conference papers in this field.

As in the "old days," multi-stage approaches thus are still very useful to either reduce costs (i.e., the number of components used) or to circumvent technological limitations.