

Multistage Architectures for Optical Packet Switching Using SOA-Based Broadcast-and-Select Switches

C. Develder, J. Cheyns, M. Pickavet, P. Demeester

Ghent University, Gent, Belgium.

Abstract:

Optical packet/burst switches can be scaled to high port counts using multi-stage architectures. To reduce the number of switching elements in SOA-based broadcast-and-select architectures, we deploy only a few stages while exploiting the WDM dimension.

1. Introduction

To satisfy the ever-lasting bandwidth hunger, (D)WDM is adopted. Long-term strategies for optical networking, that will replace the currently predominant point-to-point systems, envisage Optical Packet Switching (OPS) and/or Optical Burst Switching (OBS), fully exploiting the potential of advanced optical switching technologies. The European research project DAVID [1] aims at proposing viable approaches towards OPS. It uses a broadcast-and-select switch based on SOAs [2].

This paper illustrates how OPS/OBS switches can be scaled to high port counts using a multistage architecture. Section 2 introduces the two most wide-spread node architectures for OPS/OBS and highlights the factors limiting scalability. In Section 3 we review the three-stage Clos architecture, and show how to eliminate a switching stage by using the WDM dimension through wavelength converters. In Section 4 we present a case study for the DAVID architecture showing the advantage of the two-stage architecture in terms of number of switching elements (SOA gates). Section 5 concludes the paper.

2. Node architectures for OPS/OBS

The core functionality of an optical packet switch is to selectively transmit packets from a particular input port to a particular output port. "Port" implies a certain wavelength on a certain fibre. For "packet" switching, two fundamentally different approaches exist: one can either opt for fixed length, or rather variable length packets. The network can be operated in a time-slotted manner or an asynchronous mode. Usually a slotted approach is taken for fixed length packets, whereas asynchronous operation is adopted for variable length packets. We reserve the term OPS for fixed-length packet switching using a slotted operation, whereas OBS clearly is a case of unsynchronized switching of variable length packets.

2.1. Switch fabric architectures

An optical packet switch from a generic viewpoint comprises three parts [3]: an input interface, a switching fabric, and an output interface. Two architectures dominate the OPS/OBS approaches proposed in recent publications and research projects: (i) broadcast-and-select (B&S) architectures, (ii) Arrayed Waveguide Grating (AWG) architectures.

B&S has been proposed in e.g. the European research projects KEOPS and DAVID. The switching fabric proposed by the latter is depicted in Figure 1 [2]. The first stage multiplexes different wavelengths into a single fibre, and jointly amplifies them to compensate for the subsequent power splitting stage. For each output wavelength, two switching stages are foreseen: the first selects one of the F input fibres, and the second selects a single wavelength among the W available ones. Advantages of B&S include that it is non-blocking, and that it can perform multicasting.

The optical switching technique based on an AWG is adopted by e.g. the WASPNET project, and the more recent STOLAS [4]. The wavelength of a signal offered at one of the AWG's input ports determines via what output port it will leave: with Tuneable Wavelength Converters (TWCs) at its inputs, an AWG can be used as a switching fabric. To construct a fabric for F fibres, each carrying W wavelengths, in principle the tuneable wavelength converters have to range only over the W wavelengths in use. Unfortunately, the resulting switch then is a blocking one: there is no guarantee that all packets can be forwarded to a certain output, even if we have only to switch W (or less) packets to each of the output fibres. However, a non-blocking fabric is obtained with converters tuneable over $F \cdot W$ wavelengths. In this case, additional converters (fixed output wavelength) at the outputs are needed.

2.2. Scalability

Both the B&S and AWG architectures have limited scalability. The factor limiting the port count for the B&S architecture is the splitting ratio: each incoming signal is split in the broadcast stage over each of the $N=F \cdot W$ output ports. For the AWG-based approach, the number of output ports is limited by the tuneability range of the TWCs, which have to be tuneable over as many wavelengths as there are output ports ($N= F \cdot W$). A possible solution is the adoption of multistage architectures, as discussed in the next section.




Fig. 1. The DAVID broadcast-and-select architecture.




Fig. 2. Three-stage Clos architecture.




Fig. 3. Two-stage architecture with TWCs.

3. Multistage architectures

In the early days of network engineering, Clos proposed a multistage architecture for large switches based on smaller building blocks [5]. A sketch of an $N \times N$ switching architecture using a three-stage Clos network is outlined in Figure 2. The three stages comprise (i) N/n switches of dimension $n \times k$, (ii) k matrices $(N/n) \times (N/n)$, and (iii) N/n switches of size $k \times n$.

To be non-blocking, a lower bound on k is imposed: $k \geq 2n-1$. This minimal value can easily be determined as follows: consider a connection between input ports A and B. The worst case occurs when (i) all $n-1$ other ports of the first-stage to which A belongs are already occupied and connected to some output ports via $n-1$ second stage switches; and (ii) the $(n-1)$ remaining ports of the third-stage switch with B also are occupied, coming from $n-1$ other intermediate stage switches. Therefore the switch needs at least $(n-1)+(n-1)+1$ connection points, hence $k \geq 2n-1$. Typically, it is convenient to set $k=2n$. This boundary was determined by Clos for circuit-switched networks, implicitly assuming that connections cannot be rerouted along other (second stage) switches once they have been set-up.

Unlike Clos, we focus on OPS, operating in slotted mode. In this case, the bound for k is that of a rearrangeable non-blocking switch, which is considerably smaller: $k=n$. The proof why $k=n$ suffices is a well-known result from graph theory (the problem is equivalent to a graph coloring problem in a bipartite graph with degree n , which is n -colorable).

In case of OPS, all building block switch fabrics have symmetrical dimensions. If all wavelengths within a fibre may be considered equivalent, N input/output ports are grouped per W wavelengths into $F=N/W$ fibres, then we can eliminate the third switching stage switches if we choose $n=W$. The third “stage” then consists of W converters with a fixed outgoing wavelength, as in Figure 3. Note that the choice of $n=W$ will not necessarily be the optimal choice to construct a $(F \cdot W) \times (F \cdot W)$ switching fabric using the Clos approach; however, by eliminating the third stage, the overall cost can be reduced, as illustrated in the next section for the DAVID architecture. Note that we have not considered architectures of more than three stages: esp. for the B&S architectures, the cascade of multiple switching stages would imply a too severe signal degradation (cf. we want to avoid regeneration stages within the switching fabric).

4. Case study: limiting the number of SOAs in the DAVID architecture

In the previous sections, we discussed the DAVID switch, and illustrated how a multistage architecture can be used to construct large switch fabrics. In this section, we compare the different architectures for the case of the B&S

switch proposed within DAVID: (i) single stage, (ii) three-stage Clos, (iii) two-stage with wavelength converters. Since this architecture's cost is dominated by the SOA components, the comparison will focus on the number of SOA gates.

The DAVID switching fabric was discussed above (Figure 1). The number of SOA gates needed for 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, and w gates for wavelength selection. Since the switching matrix will be surrounded with wavelength converters (actually 3R regenerators, see [2]) the number of wavelengths w can be optimized (and chosen different from W , the number of wavelengths on the input/output fibres) to minimize the number of SOA gates. The optimal choice is $w=N^{1/2}$, leading to the minimal number of gates for a single-stage switch given in eq. (2).

For OPS switches, we have indicated that the number of second stage switches to provide a non-blocking fabric operating 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 $n=0.5 \cdot N^{1/2}$, see eqs. (3-4). For the proposed two-stage architecture, the number of SOA gates needed is given by eq. (5).

$$s_{sym}(N, w) = N \cdot ((N/w) + w) \quad (1)$$

$$s_{sym,opt}(N) = 2 \cdot N \cdot \sqrt{N} \quad (2)$$

$$c_{slotted}(N, n) = N/n \cdot s_{sym,opt}(n) + n \cdot s_{sym,opt}(N/n) + N/n \cdot s_{sym,opt}(n) = 2 \cdot N \cdot \left(2 \cdot \sqrt{n} + \sqrt{N} \cdot \frac{1}{\sqrt{n}} \right) \quad (3)$$

$$c_{slotted,opt}(N) = 4 \cdot \sqrt{2} \cdot \sqrt[4]{N} \cdot N \quad (4)$$

$$t_{slotted}(F, W) = F \cdot s_{sym,opt}(W) + W \cdot s_{sym,opt}(F) = 2 \cdot F \cdot W \cdot (\sqrt{W} + \sqrt{F}) \quad (5)$$

Inspection of the number of SOA gates needed leads to the choices illustrated in Figure 4. From this comparison, clearly the single stage architecture is only advantageous for small port counts ($N=W \cdot F < 12$). The two-stage architecture with TWCs is generally the cheapest choice: the three-stage Clos architecture only can be advantageous when either the number of wavelengths per fibre or the number of input/output fibres is very large. Figure 5 illustrates that for 8–32 fibres with 16–128 wavelengths each, a three-stage Clos architecture needs 15–40% more SOA gates, while a single stage architecture would need 65–275% more SOAs than the two-stage architecture (or rather be unfeasible because of the high splitting factor in the broadcast stage, cf. Section 2.2).

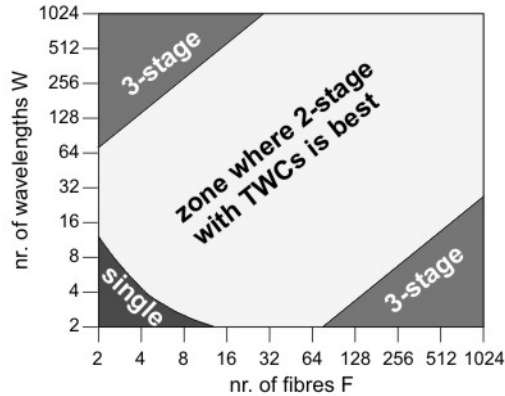


Fig. 4. Boundaries of the regions where each of the three node architectures is the most advantageous one.

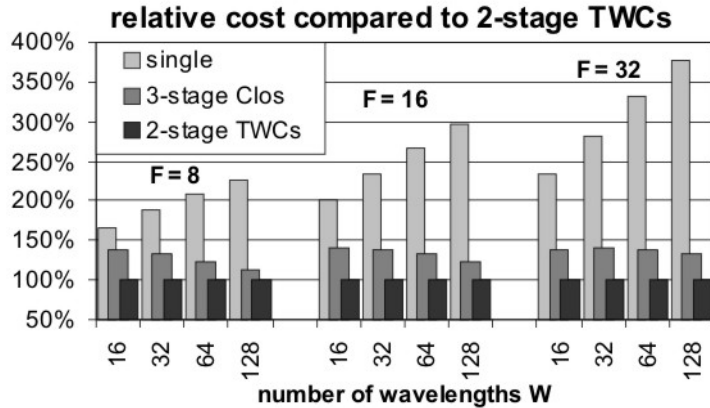


Fig. 5. Comparison of the three architectures for $F=8,16,32$ input/output fibres and $W=16,32,64,128$ wavelengths per fibre.

5. Conclusion

We have discussed two switching fabric architectures (broadcast-and-select, and AWG-based) and outlined how their respective scalability limits can be overcome by adopting a multistage approach. We have shown that for OPS we can exploit WDM to eliminate the last switching stage and replace it with wavelength converters. For the B&S architecture proposed in DAVID, we have shown that this two-stage architecture with TWCs usually is the cheapest in terms of used SOA gates; in most practical cases the three-stage architecture needs 15–40% more SOA gates, while a single stage would need 65–275% more SOAs (or even be unfeasible).

References

1. DAVID: Data And Voice Integration over DWDM, project IST-1999-11387, <http://david.com.dtu.dk>
2. D. Chiaroni, et al., “First demonstration of an asynchronous optical packet switching matrix prototype for MultiTerabitclass routers/ switches,” in Proc. 27th European Conf. on Optical Communication, (Amsterdam, The Netherlands, September 30 – October 4, 2001).
3. C. Develder, J. Cheyns, E. Van Breusegem, E. Baert, A. Ackaert, M. Pickavet, and P. Demeester, “Node architectures for optical packet and burst switching”, in Proc. Int. Topical Meeting on Photonics in Switching (PS2002), (invited) paper PS.WeA1, (Cheju Island, Korea, July 21–25, 2002), pp. 104–106.
4. STOLAS: Switching Technologies for Optically Labeled Signals, project IST-2000-28557, <http://www.ist-stolas.org>
5. C. Clos, “A study of non-blocking switching networks”, Bell System Technical Journal, Vol. 32, 1953, pp. 406-424.