Choosing an appropriate buffer strategy for an optical packet switch with a feed-back FDL buffer

Chris Develder, Mario Pickavet, Piet Demeester
Dept. of Information Technology (INTEC), Ghent University – IMEC
Sint-Pietersnieuwstraat 41, B-9000 Gent, Belgium
chris.develder@intec.rug.ac.be

Abstract *For an optical packet switch with a feed-back buffer consisting of Fibre Delay Lines (FDLs), we compare buffer configurations and strategies. We focus on achieved Packet Loss Ratio (PLR) for memoryless, bursty and self-similar traffic types.*

Introduction

The ever lasting increase in demand for bandwidth in today's communication networks is answered by the deployment of (D)WDM. The essentially circuit-switched approaches of wavelength-routed networks, as studied within recent research projects and standardisation bodies, are mid-term solutions that despite their relative ease of design and operation suffer from the difficulty of efficiently dealing with highly variable traffic patterns. To solve this issue, exploitation of time-division is envisaged through the introduction of Optical Packet Switching (OPS). It exploits fast optical switching techniques to provide greater bandwidth efficiency, flexibility, functionality and offer better granularity.

An important issue in packet switching is contention resolution, which requires buffering. In this paper, we present an optical packet switch architecture with a feed-back FDL buffer as proposed within the European IST project DAVID. We discuss various alternatives for the buffer structure and its operation, comparing their performance in terms of Packet Loss Rate (PLR).

Network concept and node architecture

In the DAVID project, the OPS approach is adopted both in the metro area, where a ring architecture is proposed, and the backbone area. In the backbone, Optical Packet Routers (OPRs), as depicted in Figure 1, will be interconnected in a mesh. The network transports fixed-length optical packets, and the OPRs are operated in slotted mode: packets are synchronized at the input ports.

The core of the OPR is the broadcast-and-select switching fabric based on SOA technology, see e.g. /1/. The ports of this fabric are connected to F input and output fibres, each operated in DWDM mode carrying W wavelengths. The OPR ports include wavelength converters: a packet may leave the OPR on a different lambda than it has arrived on.

A number of B ports of the switching fabric are connected to a buffer of Fibre Delay Lines (FDLs). By leading the light through FDLs of appropriate length, packets are delayed an integer number of slots. Packets leaving the FDL buffer are presented at the input ports of the switch again.

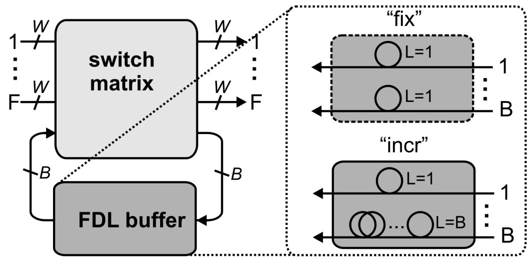

Figure 1: Node architecture of the OPR and two example buffer structures

Packet scheduling

The operation of the switch from a logical point of view consists of a procedure that is repeated every timeslot. This encompasses two phases: (i) elect W packets per output fibre to be forwarded, (ii) elect B packets that could not be forwarded and redirect them to a buffer port; any other packet is dropped.

Election of packets to be forwarded, buffered or dropped is based on the time it has already spent in the OPR's buffer. Of packets contending for the same resource (output fibre or buffer), the one which has spent most time already in the OPR is favoured, in order to avoid recirculation. Among packets having spent the same amount of time in the OPR, one is selected randomly.

Fixed versus increasing FDL lengths

An issue that needs to be addressed is what FDLs will be used to construct the buffer. A first obvious option is to use a single fibre, of a fixed length L. Alternatively, multiple fibres of different lengths could be used, thus creating a larger buffer capacity without increasing the number of switching fabric ports used for buffering purposes.

In Figure 2, we compare those approaches in terms of logical performance, i.e. PLR for an increasing number of buffer ports. In the case labelled "fix", we use the same FDL length of a single slot for each of the ports. The "incr" case uses a different FDL length for each of the wavelength ports: for the B buffer ports, lengths 1,2,3…,B are used. The graphs show loss rates for a load of 0.9 using a uniform traffic matrix, offered to an OPR with F=6 input and output fibres, and W=32 wavelengths per fibre.

As could be intuitively expected, we find that the architecture with increasing FDL lengths outperforms the single fixed-length FDL approach. The difference in PLR for B=32 wavelength ports amounts to more than two orders of magnitude (factor > 10^2) for classical traffic models such as Poisson (a Poisson process), or geometric on-off sources (geometrically distributed on- and off-times). For self-similar traffic (generated by an aggregate of on-off sources with Pareto-distributed on- and off-times, see /2/) however, adding buffer space is far less effective: the difference is limited to a factor ~4.

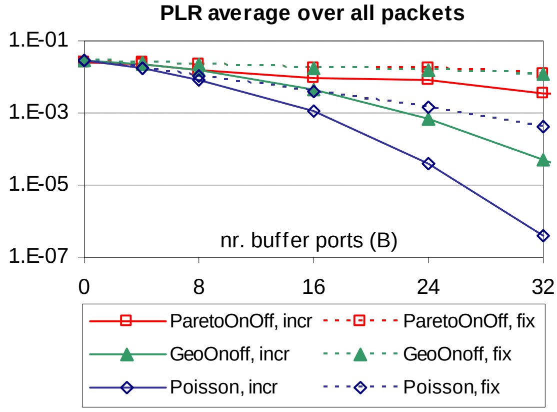

Figure 2: Comparison of using a single FDL length (fix, dashed lines) and increasing FDL lengths (incr, full lines)

The advantage in terms of logical performance for the buffer structure with different FDL lengths needs to be counterposed by (i) its more complex buffer scheduling, (ii) the fact that may introduce reordering and, to a lesser extent (iii) the fact that it necessitates a larger number of FDLs. The increase in complexity of the buffer scheduling is discussed next.

Buffer strategies

When a buffer with multiple FDL lengths is adopted, the B buffer ports are no longer equivalent. Thus, the election procedure of packets to direct to the buffer needs to determine what FDL length to use.

An obvious strategy could be to simply use the smallest FDL length for which no other packet has been elected yet; this is the one used for Figure 2. We label this approach as MinDelay. This strategy does not take into account packets put into the FDL buffer at earlier times. A more intelligent approach, denoted as Balance, inspects the buffer contents to choose an appropriate FDL length. For each available FDL length L, we count the total number of packets N_L already present in the complete buffer, destined for the same output fibre, that will leave the buffer at *now*+L slots. We choose the free buffer port with FDL length L having the smallest count N_L. Thus, the Balance strategy tries to minimise the number of packets, destined for the same output fibre, leaving the optical buffer at the same time.

In Figure 3, we compare those two strategies for increasing number of buffer ports B=0…32. The plot shows the ratio of the PLRs of the respective strategies. For the non-self-similar traffic types Poisson and GeoOnOff, we note that the Balance strategy improves the PLR by more than halving it for B=32 buffer ports. The improvement for self-similar Pareto-OnOff traffic however, is far more limited.

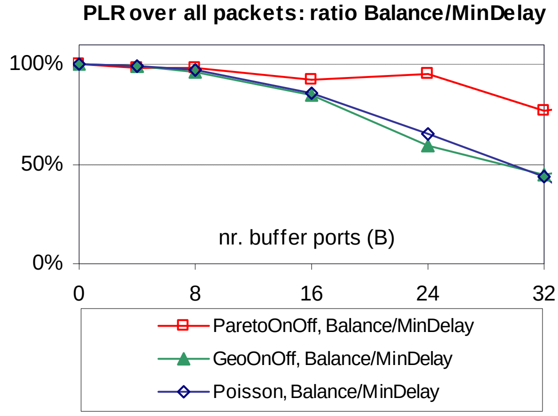

Figure 3: Comparison of using the Balance strategy versus the MinDelay strategy.

Conclusions

We have compared two buffer structures: through the use of multiple FDL lengths in a feed-back shared buffer, the PLR can be effectively cut down compared to the use of a single FDL length. By adopting a sufficiently intelligent buffer scheduling algorithm, the PLR can be further brought down. However, the effectiveness of buffering (and reduction through the aforementioned techniques) proves to be limited when considering self-similar traffic.

Acknowledgements

This work has been supported by the European Commission through the IST-project DAVID (IST–1999–11387), and by the Flemish Government through the IWT GBOU-project "Optical Networking and Node Architectures". The first author is a Research Assistant of the Fund for Scientific Research – Flanders (F.W.O.-Vl., Belgium).

References

/1/ D. Chiaroni, "Status and applications of optical packet switching", Tutorial at 27th European Conference on Optical Communication (ECOC 2001), Amsterdam, The Netherlands, Oct. 2001.

/2/ W. Willinger, et al., "Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level", IEEE/ACM Trans. on Networking, Vol. 5, No. 1, Jan. 1997, pp. 71-86.