Optical Packet Switched Networks with Recirculating FDL Buffers

Chris Develder, Elise Baert, Mario Pickavet, Piet Demeester

Abstract— **The major goal of Optical Packet Switching (OPS) is to match switching technology to the huge capacities provided by (D)WDM. We study optical packet switches with recirculating FDL buffers by comparing different FDL configurations and analysing buffer scheduling strategies in order to exploit the FDL buffer as efficiently as possible: through simulation, we assess the logical performance focusing on Packet Loss Rate (PLR). In addition, we show that through a simple priority mechanism, effective Quality of Service (QoS) differentiation can be achieved.**

Subsequently, we discuss network-wide routing algorithms designed to minimize the maximal PLR occurring in the network. In case studies on pan-European networks, we compare two algorithm variants and highlight the improvement achieved in terms of overall PLR, compared to shortest path routing and load balancing. Thus, we show that PLR-based routing achieves far better performance at the price of only a small (few percent) increase in used capacity.

Keywords— **IP-over-WDM, Optical Packet Switching, FDL Buffer, logical performance, simulation**

I. Introduction

To satisfy the ever lasting hunger for bandwidth, today's communication networks resort to the deployment of (D)WDM networking. The first step is migration from currently predominant point-to-point systems to real optical networking supporting dynamic circuit-switched paths. Optical Packet Switching (OPS, [1]) is a longer term strategy profiting from cutting edge technology: it exploits fast switching techniques to offer better granularity, efficiency and flexibility.

Towards the implementation of packet switching [2], two fundamentally different approaches exist: one can either opt for fixed, or rather variable length packets. In this paper, we focus on fixed length packets, where the OPS network is operated in a time-slotted, synchronous mode. The alternative of asynchronous switching usually is adopted for variable length packets, e.g. in the case of OBS [3].

In Section II, we detail the OPS node architecture studied, and the subsequent Section III discusses its performance. A network-wide perspective is taken in Section IV, where routing algorithms are proposed aiming at minimizing the overall Packet Loss Rate (PLR). Section V concludes the paper.

II. An Optical Packet Switch with FDL Buffer

The study presented in this paper was carried out within the frame of the European research project DAVID (Data And Voice Integration over DWDM, http://david.com.dtu.dk), aiming at proposing a viable approach towards OPS. We focus on the backbone, where so-called Optical Packet Routers (OPRs) are interconnected in a meshed DWDM network transporting fixed-length packets.

A fully non-blocking switching fabric (e.g. [4]) forms the core of the OPR. Part of this fabric's ports are connected to the OPR's neighbours via F fibres with W wavelengths each. To help solving contention, wavelength converters are foreseen at the switch's ports and B wavelength ports are reserved for connection to and from the recirculating FDL buffer.

The OPR operates in a slotted way: every timeslot, it inspects packets arriving at its input ports, and subsequently decides what packets to forward (to output fibres or FDL buffer) or to drop. This decision is taken in two phases: (i) for each output fibre of the OPR, elect at most W packets to be forward directly, (ii) from the remaining packets, elect at most B to put in the buffer; any other packet will be lost. Election of packets for forwarding and buffering is as follows: packets of a higher priority class are given precedence over lower priority ones, and within the set of packets with the same priority, the one which has already spent most time in the OPR is favoured. Among multiple packets sharing the same priority and time spent in the OPR, one is selected randomly.

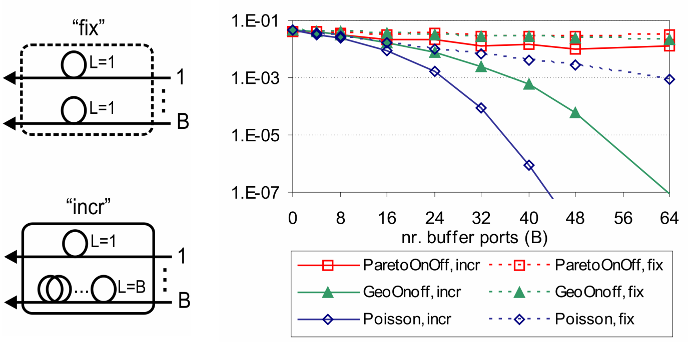

Fig. 1. Using a single FDL length (fix, dashed lines) or increasing FDL lengths (incr, full lines). The overall PLR is plotted for increasing number of buffer ports B. A uniform traffic matrix was used, for a total load of 0.95 with 50% highest, 25% middle and 25% lowest priority traffic.

III. Performance of a Single Node

A. Choosing the FDL buffer structure

For the FDL buffer structure, there are essentially two options: use a single FDL length for all B buffer ports, or adopt different FDL lengths. The latter offers greater buffer capacity for the same number of switching fabric ports.

The fixed FDL case ("fix"), and the case with FDLs of increasing lengths L=1,2,3,…B ("incr") are compared in Fig. 1 for three distinct traffic source types. The Poisson case denotes the well-known Poisson process. The GeoOnOff source generates bursty trains of packets: it is an on/off source with geometrically distributed lengths of both on- and off-periods. Self-similar traffic, labelled ParetoOnOff, was generated using on/off sources with Pareto distributed on- and off-times [5].

In accordance with intuition, we find that the buffer with increasing FDL lengths for the B buffer ports largely outperforms the buffer with a single FDL length. Yet, for the self-

The authors are with the Department of Information Technology (INTEC), Ghent University, Gent, Belgium.

This work has been supported by the European Commission through the IST project DAVID and the Flemish government through the IWT GBOU-project "Optical Node and Network Architectures". C.Develder is supported as an Aspirant of the Fund for Scientific Research – Flanders (FWO-Vl.).

similar traffic model ParetoOnOff, the differences are far less striking: adding buffer space is not that effective.

Obviously, the better logical performance of using different FDL lengths needs to be counterposed by the risk of re-ordering of packets belonging to the same flow.

B. Choosing an appropriate strategy to minimize packet loss

When the lengths of the FDLs used for each of the B buffer ports differ, not all those ports are equivalent: a decision procedure is needed to determine what FDL length will be used for a packet elected for buffering. We have compared the following strategies [6]:

- ***MinDelay ("minimize delay"):*** use the free buffer port with smallest corresponding FDL length.
- ***NoOvr ("no overload"):*** ensure that packets leaving the buffer at the same time will never overload an output port; drop the elected packet if this is unfeasible.
- ***AvoidOvr ("avoid overload"):*** same as NoOvr, but still enter packet causing overload in buffer if there is room.
- ***Balance:*** minimize the number of packets simultaneously leaving the FDL structure for the same output.

Results showed that the Balance strategy largely outperforms the others for both Poisson and the bursty GeoOnOff models (factors up to 6, resp. 3, for B=40).

IV. ROUTING STRATEGIES TO MAXIMIZE QOS PERFORMANCE

A. Assessing network-wide PLR

The previous results focused on the logical performance of a single OPR. In real life, we are interested in traffic crossing a network of OPRs. From this point of view, the single node studies' results are only meaningful if the input traffic profile can be assumed similar for all OPRs (in terms of packet arrival distribution; not necessarily avg. load), which is the case if the traffic model is not impacted by the OPR's behaviour. Simulations have shown that this is the case, thus allowing the estimation of PLRs in different OPRs independently. In addition, to be able to quickly estimate PLRs, we have derived empirical analytical formulas capturing the PLR behaviour.

B. Routing algorithms

In order to minimize the network-wide PLR, we have devised routing algorithms based on PLR estimation. The algorithm computes a route for every demand associated with a (source,destination)-pair. A heuristic approach is followed to find routes as cheap as possible, starting from an initial shortest-path routing. The first *global* phase of the algorithm considers the network as a whole and tries to minimize the max-PLR by giving penalties to links exhibiting high PLRs and rerouting all paths. The second *local* phase reroutes individual paths.

Two variants were studied for the local phase: WorstLink reroutes each path crossing the link with the highest PLR in turn; WorstPath reroutes the path suffering from the worst end-to-end PLR by giving a penalty to each of its links in turn.

C. Case studies on pan-European networks

To evaluate the performance of our PLR-based routing algorithms, cases studies were carried out on two Pan-European networks: a *dense* 19-node network, and a *sparse* 27-node network. For the local phase variants of our PLR-based algorithm, WorstLink outperforms WorstPath: it reaches the lowest loss rates and requires the least bandwidth. The effectiveness of the PLR-based approach is apparent from Figure 2: the PLR is lowered by a factor ranging up to 5 orders of magnitude compared to shortest path routing.

To assess the importance of accurate PLR estimation, we compared the PLR-based routing with load balancing. For the latter case, we used the same algorithms as above with link load as the measure to minimize (instead of PLR). The results showed that even though the load is a dominant factor for the PLR on a link, using a more accurate PLR measure can still lead to an additional reduction with one or two orders of magnitude.

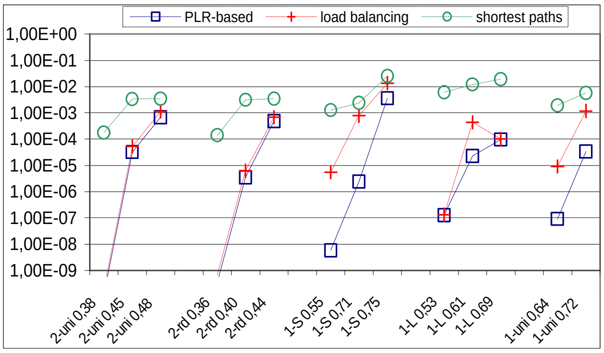

Fig. 2. Max-PLR achieved with load balancing (dash-dotted line, + markers), PLR-estimation (full line, squares) and shortest path (dashed line, circles), for different demands (network/pattern/load).

V. CONCLUSION

We have focused in this paper on an optical packet switch with recirculating FDL buffer. We assessed the logical performance through simulation and showed that multiple FDL lengths are desirable to achieve very low PLRs, and that the Balance strategy makes very efficient use of the FDL resources. A simple priority mechanism was shown to be sufficient to obtain effective QoS separation.

We have shown that the single node results are relevant for network-wide PLR assessment, since the traffic profile is impacted only very little by the OPR behaviour. We also proposed an analytical formula for quick PLR-estimation, which is useful for PLR-based routing algorithms. The usefulness of such algorithms was illustrated by case studies on pan-European networks.

REFERENCES

[1] S. Yao, B. Mukherjee, S. Dixit, "Advances in photonic packet switching: an overview", IEEE Commun. Mag, vol.38, no. 1, Jan. 2000, pp. 84–94.

[2] C. Develder, et al., "Node architectures for optical packet and burst switching", Proc. Int. Topical Meeting on Photonics in Switching (PS2002), (invited) paper PS.WeA1, Cheju Island, Korea, 21–25 Jul. 2002, pp. 104–106.

[3] C. Qiao and M. Yoo, "Optical Burst Switching (OBS) — a new paradigm for an optical internet", J. High Speed Networks, vol. 8, no. 1, 1999, pp. 69-84.

[4] D. Chiaroni et al., "First demonstration of an asynchronous optical packet switching matrix prototype for MultiTerabitclass routers/switches", Proc. of ECOC2001, Amsterdam, The Netherlands, Oct. 2001.

[5] W. Willinger, M. Taqqu, R. Sherman, and D. Wilson, "Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level", IEEE/ACM Trans. Netw., vol. 5, no. 1, Jan. 1997, pp. 71–86.

[6] C. Develder, M. Pickavet, P. Demeester, "Choosing an appropriate buffer strategy for an optical packet switch with a feed-back FDL buffer", Proc. ECOC2002, Copenhague, Denmark, 8-12 Sep. 2002.