Optical Packet Switched Networks with Recirculating FDL Buffers

Chris Develder, Elise Baert, Mario Pickavet, Piet Demeester

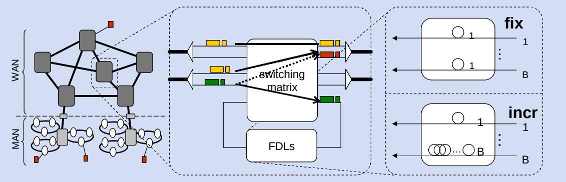

We focus on an OPR that is part of a meshed backbone network (WAN). The inputs of an OPR are fibres, each carrying multiple wavelengths in (D)WDM. The wavelengths are used to transport fixed-length packets, to be switched in slotted mode by the OPR.

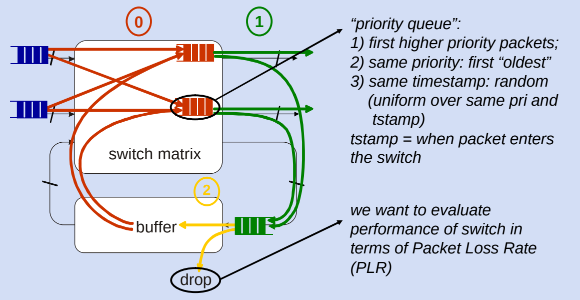

The algorithm followed by the OPR consists of two phases, which are repeated every timeslot. Based on the knowledge of how many packets arrived in the slot, the OPR (1) elects at most W packets to be directly forwarded along each of the F output fibres; (2) from the packets that could not be directly forwarded, at most B are selected to be sent to one of the B recirculating FDL ports.

Optical Packet Switches with recirculating FDLs

To match switching technology to the huge bandwidth provided by (D)WDM systems, Optical Packet Switching (OPS) has been proposed: it is a longer term strategy (as opposed to optical circuit switching) profiting from cutting edge technology to offer e.g. better bandwidth granularity and flexibility.

One approach of OPS uses fixed-length packets handled by synchronous Optical Packet Routers (OPRs), operating in a slotted mode: each timeslot, the OPR inspects packets at its input ports and decides what to do with them (forward, buffer, drop). We focus on an OPR with an optical buffer consisting of recirculating Fibre Delay Lines (FDLs), as in the backbone network proposed in the European DAVID project.

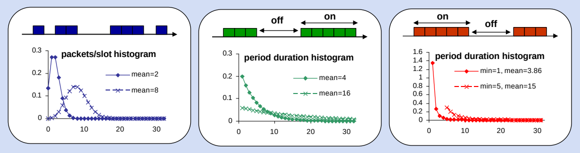

To assess the performance of an OPR, we use different traffic patterns, ranging from the classical memoryless Poisson traffic, over on/off traffic using geometrical distribution to traffic exhibiting long-term correlations (aggregate of Pareto On/Off sources).

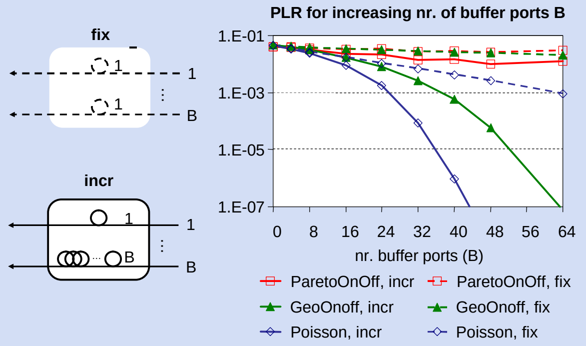

As can intuitively be expected, using the "incr" structure for the recirculating FDL buffer results in lower PLRs.

Choosing a buffer configuration

For the block of FDLs, there are essentially two options: (i) use a single FDL length for all B buffer ports, or rather (ii) deploy different lengths (resulting in a buffer capable of storing more packets for the same number of switch fabric ports devoted to buffering). Thus, by using e.g. the "incr" buffer, the Packet Loss Ratio (PLR) can be lowered with multiple orders of magnitude. Only when traffic exhibits long-range correlations (e.g. the ParetoOnOff model), the limited capacity of optical buffering is not effective (whatever the buffer structure).

What buffer strategy?

When the lengths of the FDLs used for each of the B buffer ports differ, they are no longer equivalent. Therefore, we compared 4 strategies to choose a buffer port for a packet:

- *MinDelay:* choose the available port with shortest length
- *NoOverload:* assure that packets leaving the buffer simultaneously never overload an output port;
- *AvoidOverload:* first try NoOvr; if this fails use MinDelay
- *Balance:* spreads packets destined for same output port in time.

We showed that for non-self-similar traffic, the Balance strategy offers substantial PLR-reduction.

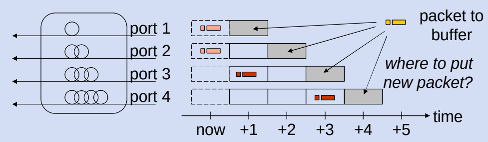

Since FDLs are shared over all output ports, we should use buffer as efficient as possible, and thus send a packet only to a single FDL. By choosing the FDL, we decide when it will re-enter the switch (and have another attempt at forwarding). If not all buffer ports have the same FDL length, we should choose the FDL wisely.

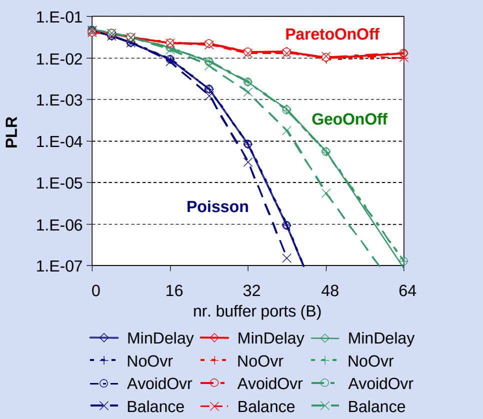

Service Differentiation

To offer service differentiation, we proposed a simple priority mechanism. Simulations showed that even under heavy load conditions, the PLR was still acceptable.

The delay was shown to be limited to a few timeslots (us range), and thus negligible compared to propagation delays (ms range).

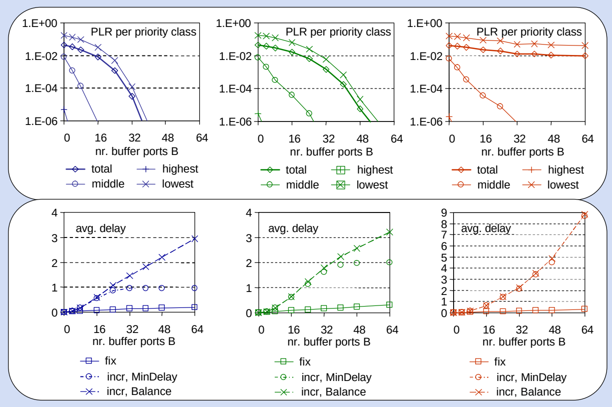

Single node results?

In real life, OPRs will be interconnected in a network. Thus, the inputs of a typical OPR will be carrying traffic coming from the outputs of another OPR. To validate the usefulness of the single node studies, we verified that an OPR does not significantly alter the traffic profile from in- to output.

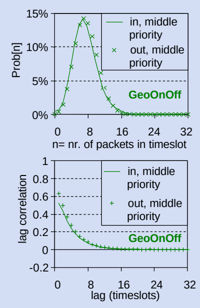

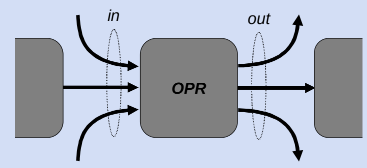

Traffic entering an OPR along its input fibres will be the output traffic of another OPR. Therefore, we compared the output traffic statistics of a single OPR to check that the given input statistics are not substantially affected. For all studied traffic patterns, this proved to be correct.

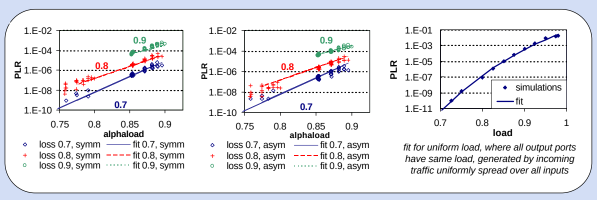

We derived a heuristic formula giving the PLR for an output link in function of the load on that link and an indication ("alphaload") of the overall load on the OPR it is an output fibre of. We fitted our formula to the simulation results for a wide set of uniform and non-uniform traffic patterns on the 6x6 node with B=32 buffer ports where output fibre loads are taken from {0,0.7,0.8,0.9}.

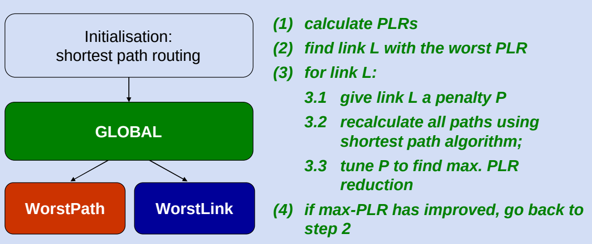

Our routing algorithm tries to route paths such that the PLR is minimal. For the local phase, we compared (i) a WorstPath variant focusing on rerouting a single path exhibiting the max-PLR, with (ii) a WorstLink variant that reroutes multiple paths crossing the same link with a high PLR. We concluded that while WorstPath usually leads to better results in a few iterations, the final optimum reached with WorstLink always results in the lowest max-PLR.

Assessing end-to-end performance

To asses end-to-end performance, we captured the PLR behaviour observed in a wide range of simulation set-ups into an analytical formula. We use this formula to estimate the PLR in every node and combine the results to find the end-to-end PLR for each traffic flow.

Routing algorithms to minimize PLR

To minimize network-wide PLR, we have devised routing algorithms based on PLR-estimation. The algorithm finds routes for every demand associated with a (source,destination)-pair. A heuristic approach is followed, starting with a *global* phase penalizing links exhibiting high PLRs and rerouting all demands accordingly. A second, *local*, phase reroutes individual paths.

Compared to shortest-path routing, we showed that for case studies on pan-European networks the maximal PLR occurring in the network can be lowered by multiple orders of magnitude. To assess the importance of accurate PLR-estimation, we compared the PLR-based algorithm with load balancing (using the same algorithm to minimize the max. load). Although link load is a dominant factor for the PLR, more accurate PLR-assessment can lead to an additional PLR reduction with one or two orders of magnitude.

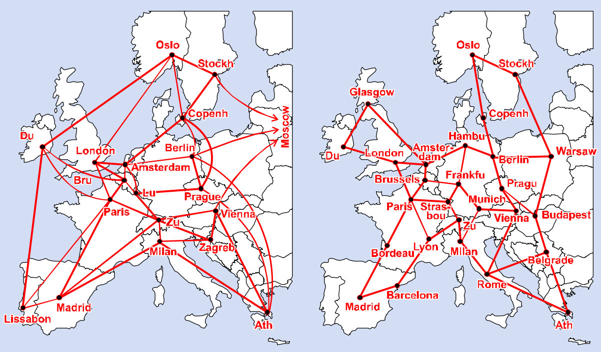

The routing algorithm was evaluated in case studies on the pan-European networks sketched above: the first (1, left) connects 19 cities with 40 links, the second (2, right) uses the same number of links to connect 27 cities.

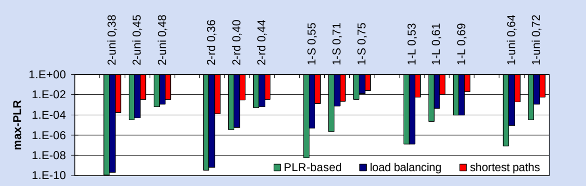

The case studies covered various demand patterns: "uni" for uniform, "rd" for random non-uniform, "S" for random with all of the same order, "L" for random patterns with a few demands of larger order; 1 refers to the 19-node network, 2 to the 27-node network. The third number is the overall mean network load (range [0,1])

UNIVERSITEIT GENT

Department of Information Technology (INTEC) – Broadband Communication Networks (IBCN)