Routing strategies to minimize packet loss in an optical packet switched network with recirculating FDL buffers*

E. Baert, C. Develder†, D. Colle, M. Pickavet, P. Demeester

Dept. of Information Technology, Ghent University – IMEC,

Sint-Pietersnieuwstraat 41, 9000 Ghent, Belgium

e-mail: {elise.baert, chris.develder, didier.colle, mario.pickavet, piet.demeester}@intec.ugent.be.

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 Fiber Delay Line (FDL) buffers. Through simulation, we have assessed the logical performance of a single Optical Packet Router (OPR), focusing on Packet Loss Rate (PLR). By verifying that our scheduling algorithm does not alter the traffic profile characteristics from in- to output, we illustrate how the single node results can be used to assess network-wide performance.

We use the capability of assessing end-to-end PLRs to develop network-wide routing algorithms designed to minimize the maximal PLR occurring in the network. In case studies on pan-European networks, we first compare two algorithm variants and thereafter we compare the PLR-based routing algorithm with both load balancing and shortest path routing. While load balancing achieves PLRs that are multiple orders of magnitude lower than shortest path routing, the PLR-based algorithm can reach PLRs up to two orders of magnitude better. The improvement in PLR comes at the price of only a small increase in used bandwidth (a few percent).

Subsequently we show that the discussed PLR-based routing algorithm can be easily extended to multiple priorities. By introducing multiple priorities we can keep the loss rates for high priority traffic very low. However, it may lead to an increase of the obtained minimal max-PLR value for low priority traffic. But as we prove this increase to be limited, the cost of introducing multiple priorities is small.

Keywords: Optical Communication, Optical Packet Switching, Routing, Simulation, Priorities.

---

* This work has been supported by the European Commission through the IST project DAVID (IST-1999-11742) and the Flemish Government through the IWT GBOU-project "Optical Networking and Node Architectures". The work of C. Develder was supported by the Fund for Scientific Research — Flanders (FWO-Vl.). D. Colle would like to thank the IWT for its financial support for his postdoctoral grant.

† Corresponding author

1 Introducion

The rapidly increasing demand for bandwidth in telecommunication networks is met by the huge capacities provided by (D)WDM. The first step in moving from point-to-point systems to real optical networking encompasses a circuit-switched approach [1]. This approach however has difficulties dealing with highly variable traffic (both in volume and traffic pattern). Optical Packet Switching (OPS, [2-4]) offers a longer-term solution that provides better bandwidth granularity, efficiency, and flexibility.

We consider a core OPS network consisting of so-called Optical Packet Routers (OPRs) that synchronously switch fixed-length packets (as opposed to asynchronous switching of the variable-sized bursts in e.g., Optical Burst Switching, OBS [5-6]).

The major problem that needs to be addressed in any packet switched concept is contention resolution: what if multiple packets need to be switched simultaneously to the same output port of the switch? In an OPS environment, three different techniques can be identified to solve this: (i) wavelength conversion, (ii) buffering, and (iii) deflection routing. The use of wavelength conversion implies that multiple packets will be switched to the same outgoing fiber using WDM, where some of them may be forwarded on another wavelength than they entered the switch. It has been shown that this exploitation of the wavelength domain greatly reduces the need for buffering (e.g., [3], [7] and [8]). Still, using wavelength conversion alone, contention still can arise, which can most straightforwardly be solved through the use of buffering of some kind. However, since buffering in optics implies the use of Fiber Delay Lines (FDLs), also deflection routing has been proposed: some of the contending packets are sent to a “wrong” output port, forcing them to make a detour, in the hope to avoid the congested network part. Clearly, this only works when enough free capacity is available in the other parts of the network, thus for reasonably low overall network loads. The soundness of this intuitive insight has been confirmed by a comparison of the three approaches to contention resolution, showing that deflection routing is outperformed by the other two techniques [9-10].

So, to ensure the efficiency of the OPS network, and to obtain low Packet Loss Rates (PLRs) , the OPR considered in this paper will use wavelength conversion and an optical feedback buffer [11] with FDLs. The logical structure of the OPR, proposed within the framework of the European research project DAVID [12], is depicted in Fig. 1 (for the physical structure of the switching matrix proposed in DAVID, see [13]).

The OPR operates in a slotted way: at every slot time, it inspects packets arriving at its input ports, and subsequently decides which packets to forward (to the output ports or the feedback buffer) or to drop. This decision is taken by following a fixed procedure, comprising two phases: (i) for each output fiber of the OPR, elect at most w packets to be forwarded 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 based on two criteria: the priority attached to the service class the packet belongs to, and the time it already spent in the OPR. Service differentiation is based on a pure priority scheme: packets of a higher priority class are given precedence over lower priority ones. Within the set of packets with the same priority, the one which has spent the longest time in the OPR is favored. Among multiple packets sharing the same priority and with the same time spent in the OPR, one is selected randomly.

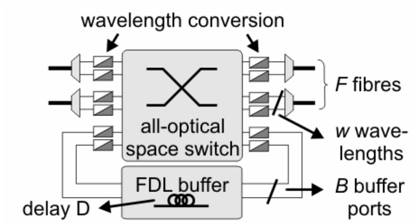

Fig. 1. Optical Packet Router architecture.

For the FDL-buffer used, two cases are considered. 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 of 1, 2, 3…,B slots are used. When a buffer with multiple FDL lengths is adopted (incr), 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. 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 fiber, that will leave the buffer at time 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 minimize the number of packets, destined for the same output fiber, leaving the optical buffer at the same time.

A detailed analysis of the performance of a single OPR, for different traffic profiles, focusing on the buffer structure and scheduling strategies can be found in [14-15]. Alternative switch structures are discussed for instance in [16].

In this paper, we discuss routing strategies to minimize the PLR in a meshed network of OPRs. In contrast with [17] where routing is only performed with the OSPF (Open Shortest Path First, [18]) protocol, we will adopt an MPLS (Multi Protocol Label Switching, [19]) approach, where routing and forwarding are separated and other than pure destination-based routing can be used, which results in a greater flexibility.

In Section 2, we explain how we assess the network-wide PLR based on simulation results for a single OPR. The routing algorithms based on PLR-estimation are discussed in Section 3. In the subsequent Sections 4 and 5, we evaluate the performance of the PLR-based routing algorithm through case studies on pan-European networks. In Section 4, we compare two PLR-based variants, and we continue in Section 5 by comparing the developed PLR-based routing algorithm with less complex load balancing and straightforward shortest path routing. In the next Section 6, we discuss the influence of introducing multiple priorities in the PLR-based routing algorithm (in all previous simulations all traffic was of one single priority). The paper is concluded in Section 7.

2 Assessing network-wide performance

To assess network-wide performance, especially in terms of PLR, a straightforward solution could be to simulate the network as a whole, using e.g., the simulation tool we have built to assess single node performance [14]. However, since this is quite time-consuming and therefore prohibitive for the iterative routing strategies we propose in the following Section 3, we assess the PLR of each individual OPR using an approximation of the PLR by an analytical formula.

In Section 2.1, we first describe an analysis of the traffic pattern at the OPR's outputs which shows that the statistical properties of the input traffic are not significantly altered, thus validating the approach of establishing end-to-end performance through analyzing each OPR in turn. In the following Subsection 2.2, we heuristically derive an analytical formula for the PLR inflicted at a single OPR, this by performing extensive simulations of a single OPR under varying traffic conditions, thereby using the single node simulator described in [14]. This formula will be used to quickly calculate the PLRs in the different iterations of the routing algorithm.

2.1 Cascadeability of the single node model.

The studies presented in [14] and [15] focused on the performance of a single node under various traffic profiles. However, in real life, such Optical Packet Routers will be interconnected in a (backbone) network. This implies that the output of a particular OPR will be the input of another one. In this context, the results of the single node studies are useful only if the input traffic profile can be assumed to be similar for all OPRs (in terms of type of packet arrival distribution, but clearly not necessarily in terms of average load). In particular the question arises whether the profile of the traffic on an output fiber of an OPR is similar to that offered at its inputs. This is the question we address in this section.

To compare the traffic profiles at the inputs and outputs of the OPR, we have traced the number of packets arriving at each of the input ports and leaving on each of the output ports, and this for each priority. Three traffic source types were considered. The first is the well-known Poisson process. The GeoOnOff source generates bursty trains of packets: 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 [20]. During the on-periods a packet is sent every timeslot.

The traffic was divided in three priority classes: 50% highest, 25% middle and 25% lowest priority traffic. For the analysis of the traces, we have focused on the number of packets n_τ that arrive in the same timeslot τ on a particular input or output fiber. To characterize the packet arrival process, the two foremost important parameters are the probability distribution of n_τ and the correlation between n_τ and $n_{\tau+lag}$ (i.e., correlation between the number of packets arriving in timeslots spaced by lag slot-times).

When switching packets from input ports to output ports, the OPR will impact the traffic profile in two ways: (i) by dropping packets, and (ii) by delaying packets. Dropping packets will mainly impact the distribution of n_τ. Delaying packets will mainly impact the correlation between the number of packets in successive timeslots: buffering will spread packets in time (especially the Balance strategy for the incr buffer structure, see introduction), thus potentially lowering the correlation of the input traffic. In Fig. 2, we have plotted, for Poisson traffic, the comparison of n_τ histograms for input and output port traffic, for the incr buffer structure with B = 64 ports, using the Balance strategy. The accompanying lag correlations between n_τ and $n_{\tau+lag}$ are presented in Fig. 3 Similar curves have been analyzed for other buffer structures and other traffic profiles leading to the same conclusions.

Inspection of the histograms confirms our intuitive expectations: for the priority classes suffering from higher drop probabilities (i.e., the packets with the lowest priority), the effect of these drops is a lowering of the probabilities for higher n_τ values and corresponding rise for lower n_τ, leading to a slightly lower average number of simultaneous packet arrivals for these classes (Fig. 2, (c)-(d)). For the overall number of simultaneous arrivals (Fig. 2, (a)), the losses result in a slight lowering of the probability of 32 simultaneous arrivals, and thus a relative increase for lower n_τ values. The correlation plots of Fig. 3 show that for each individual priority class, the correlation structure of the input traffic is not significantly impacted by the OPR. For the packet counts over all priority classes, we notice a reduction of the correlation for the OnOff traffic types due to the aforementioned spreading in time achieved through buffering.

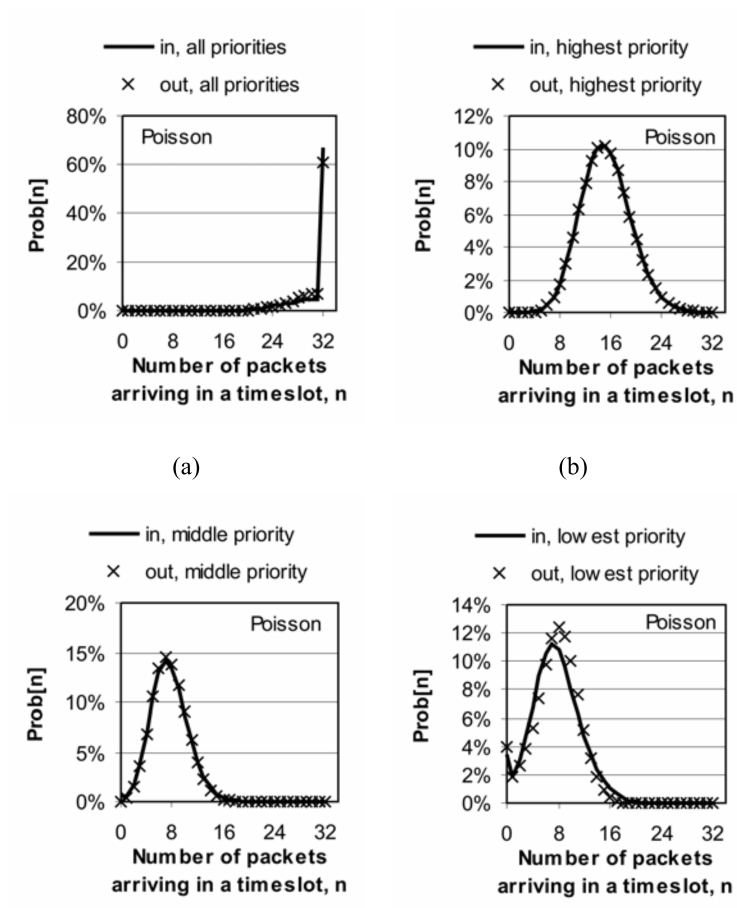

Fig. 2. Comparison of the histograms of the number of packets per timeslot (w = 32) arriving on a particular input port (solid lines) and leaving the OPR on a particular output port (marked dashed lines).

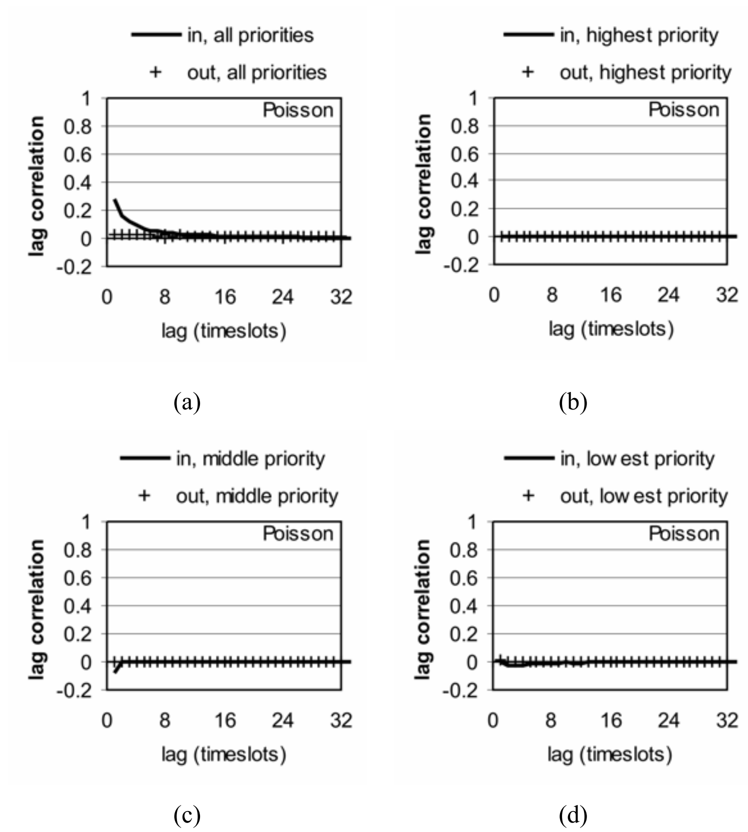

Fig. 3. Comparison of the correlation between the number of packets per timeslot (w = 32) arriving in timeslots τ and τ+lag on a particular input port (solid lines) and leaving the OPR on a particular output port (marked dashed lines).

From our probability and correlation analysis, we may conclude that it is safe to describe the input and output traffic profiles by the same model, at least to estimate packet loss rates. Indeed, the correlation structure is not significantly impacted, and nor are the probabilities of n_τ simultaneous packet arrivals. This result allows to estimate PLRs in different OPRs interconnected in a network independently. Such network-wide PLR estimation can be used to make routing decisions, as discussed in Section 3.

2.2 Capturing the single node performance in a simple formula.

The routing algorithms, proposed in the next Section 3, are of an iterative nature, and thus we need a reasonably fast method to estimate the PLRs in all nodes for each iteration step. Therefore, we will use an approximation by an analytical formula. The form of the chosen formula, and the parameter values, will be discussed in this section.

The objective is to find a formula predicting the loss on a certain output fiber, given the load offered to the OPR. Under the assumption that the traffic on such an output fiber is an aggregate of traffic coming from multiple input ports, the main factor impacting the loss for traffic passing through the OPR to this fiber will be the offered load for that output port. However, since the OPR's recirculating buffer is shared among all outputs, the loads on other output ports will also affect the PLR. To capture this correlation between packet loss rates on different output fibers, we propose a formula of the form given in Eq. (1) where L_f stands for the load on output fiber f (ranging from 0 to F–1, where F is the number of output fibers of the switch). The measure L_α is an average output load giving more weight to the higher loads, since these are the only ones that will use the recirculating buffer intensely and thus impact the losses on other fibers.

$$\text{PLR}(L_f \mid L_0 \dots L_{F-1}) \approx \text{g}(L_f, L_\alpha) \text{ , with } L_\alpha = \sqrt[\alpha]{\frac{\sum_{i=0}^{F-1} L_i^{\ \alpha}}{F}} \tag{1}$$

To get an idea on the analytical form to use for the function g, we have performed a series of simulations for each combination of output loads L_f, where each load was taken from the set {0, 0.7, 0.8, 0.9}. We analyzed the PLR curves for these given output load combinations for two sorts of traffic matrices. In a first traffic matrix, named symm, we considered the case where each input port equally contributed to the load L_f on each output fiber f (each input fiber i contributes part L_f / F). The second traffic matrix type, denoted asym, focused on asymmetrical contributions from each of the input ports to the load on output f. The formulas used to set the load of the traffic generated for a particular (input,output)-pair in case of asym is given in Eq. (2).

$$L_{i,f} = \begin{Bmatrix} L_f / 2^{1+(i-f) \bmod F} & , \ i \neq f \bmod F \\ L_f / 2^{F-1} & , \ i = f \bmod F \end{Bmatrix} \tag{2}$$

The resulting plots of the PLR in function of L_α for each of the values used for the output loads are plotted in Fig. 4 for the case of an OPR with F = 6 input/output fibers, w = 32 wavelengths per fiber, and a fix buffer structure with B = 32 recirculating buffer ports using the MinDelay strategy. By varying α we observed for reasonably large values a strong separation of the measure points for different output load values L_f and small spreading for the measure points for same (L_f, L_α)-pairs. The resulting points for the same L_f values almost fall onto the same straight line in a logarithmic plot. This observation led to the proposal of the formula given by Eq. (3), with a factor and an exponent depending only on L_f. The meaning of the factor $p(L_f)$ in that formula is the packet loss rate for $L_\alpha = L_f$, thus the case where all output ports have the same load of L_f. The packet loss rate in case of a uniform traffic matrix is plotted in Fig. 5, along with the outcome of the analytical formula we proposed for this packet loss rate as given in Eq. (4). The exponent $r(L_f)$ was chosen to be linear in L_f, as in Eq. (5). All the parameters in formulas (1, 3–5) were fit using the method of least squares, with the results listed in Table 1. The plots show that for these values, the correspondence with the simulation results is very satisfactory.

$$PLR(L_f \mid L_0 \dots L_{F-1}) = p(L_f) \cdot 10^{r(L_f) \cdot (L_\alpha - L_f)} \tag{3}$$

$$p(L_f) = PLR(L_f \mid uniform) = a \cdot L_f^{\ b} \cdot 10^{c \cdot L_f} \tag{4}$$

$$r(L_f) = u \cdot L_f + v \tag{5}$$

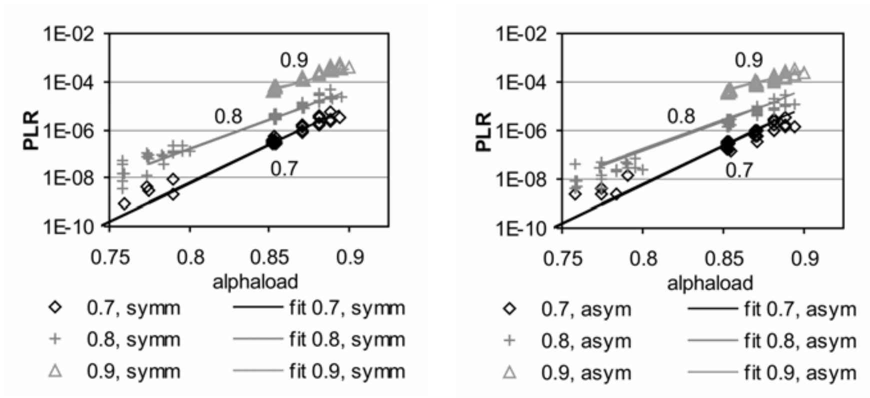

(c) (d)

Fig. 4. Packet loss rates and fits with analytical formula in (a) for the symm case where each input port contributes in the same way to the output load on a particular output port, and in (b) for the asym case with asymmetrical contributions by the input ports.

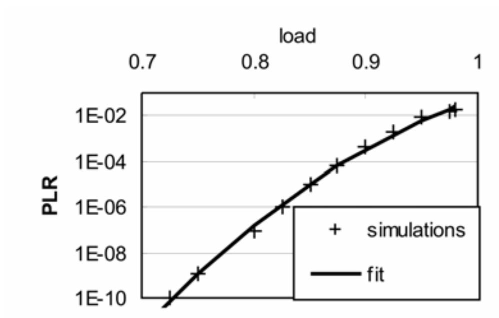

Fig. 5. Packet loss rates in case of a uniform traffic matrix for increasing loads.

Table 1. Parameter values of formulas (1, 3-5).

Parameter	***Value***	***Parameter***	***Value***
α	32.97	*a*	62.59
u	71.58	*b*	189.15
v	-82.77	*c*	-63.81

3 Routing algorithms

The PLR formula explained above (Section 2.2) enables us to assess network-wide PLR within a reasonable time. We can use this PLR-estimation technique to drive routing algorithms aiming at minimizing the maximal PLR that occurs for a given traffic demand matrix. In this section we are going to elaborate on this routing problem and the developed routing algorithms. In Section 3.1 we give a detailed problem description, followed in Section 3.2 by an in-depth discussion of the routing algorithms, which will set up paths between sources and destinations, starting from an initial shortest path routing and then trying to achieve lower PLRs by rerouting some paths.

3.1 The routing problem

The following optimization problem is considered: given the network topology (nodes and links) and capacity; the cost for nodes and links (in function of capacity), the node model (packet loss rate in terms of load), the traffic demand matrix (stating not only the required capacity between two OPRs, but also an upper bound giving the tolerable PLR for this demand), and the maximal tolerable packet loss, we want to find the (cheapest) routes fulfilling the demands and their maximal PLR requirements. The problem is illustrated in Fig. 6.

More specifically, we here want to route the demands in the network so that the maximum PLR (this is, considering the PLR of the traffic on each link of the network, the maximum of these PLRs) occurring in the network is as low as possible, and this as cheap as possible.

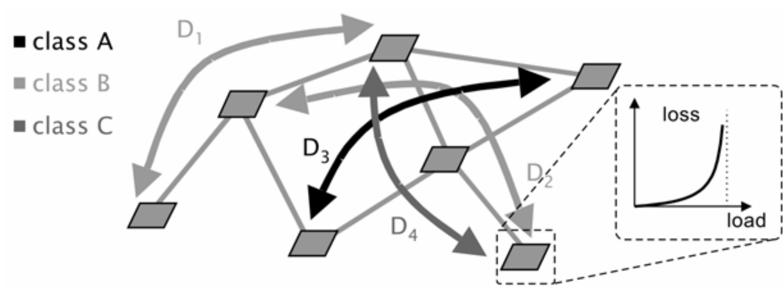

Fig. 6. Routing problem.

The general routing problem is defined as follows. The network can be represented as a directed graph G = (N,A) whose nodes and arcs represent optical packet routers and the links between them. Each arc a has a capacity c(a) which is a measure for the amount of traffic flow it can take. In addition to the dimensioned network, a given demand matrix D for each pair (s,t) of nodes gives the traffic flow from source s to destination t. Many of the entries of D may be zero, and in particular, D(s,t) should be zero if there is no path from s to t. We also have PLR requirements which state the maximum PLRs the demands may have. In our case studies we do not have specific values we want to obtain, we just iterate until we reach the lowest max-PLR possible. So, the objective is to minimize the max-PLR.

- ***Given:***
  - $G = (N,A)$ is the considered network with nodes $n \in N$ and arcs $a \in A$
  - $c(a)$ is the capacity of arc a
  - $D(s,t)$ is the demand to be set from node s to node t
  - $x_{st}(a)$ is a zero-one variable and is equal to 1 $\Leftrightarrow$ arc a belongs to the selected path for D(s,t)
  - $f_{st}(a)$ is the flow from the demand from s to t over arc a = $D(s,t)\cdot x_{st}(a)$
  - L_a is the load on arc a = $\frac{\sum_{s,t} f_{st}(a)}{c(a)}$
  - PLRa is the PLR on arc a = PLR(La|L0… L_f -1)
  - I(n) is the from-incidence of node n = set of all arcs leaving node n
  - I'(n) is the to-incidence of node n = set of all arcs arriving at node n
- ***Objective:***

- minimize MAX(PLR_a) $\forall$a

- ***Constraints:***
  - $\sum_{a \in I(n)} x_{st}(a) \le 1, \quad \forall n \in N, \forall s,t \in N$
  - $L_a \le 1, \quad \forall a \in A$
  - $\sum_{a \in I(n)} f_{st}(a) - \sum_{a \in I'(n)} f_{st}(v) = \begin{cases} D(s,t) \text{ for } n = s & (i) \\ 0 & (ii) \\ -D(s,t) \text{ for } n = t & (iii) \end{cases}$

3.2 The routing algorithm

The heuristic algorithms we developed are iterative, and consist of multiple phases, as depicted in Fig. 7. They start with calculating the shortest paths for each (source, destination)-pair. The subsequent phases will reroute some of these paths to lower the max-PLR. A Zoom-In philosophy is used [21], starting with a first phase that takes a global perspective, after which the result is refined in a second, local phase.

The global phase (Fig. 8) considers the network as a whole by giving penalties to links exhibiting high PLRs and recalculating the routes for all demands. When changing these penalties does not lower the max-PLR anymore, we go on to the second, local phase where only a single path is re-routed in each step.

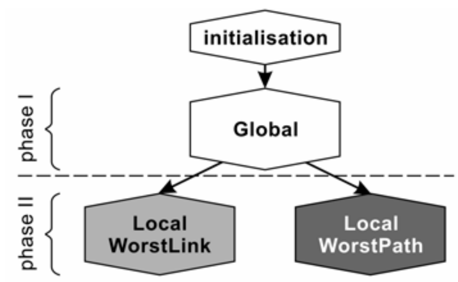

Fig. 7. Phases of the algorithm.

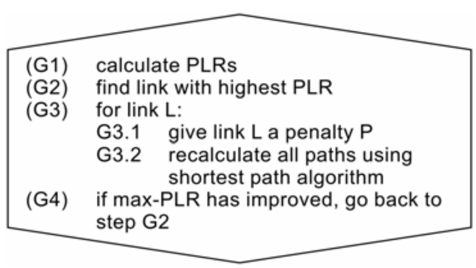

Fig. 8. Routing algorithm: algorithmic description of global phase.

In the local phase —which was added because the first phase only gives a small max-PLR reduction (see Section 4)— we compared two variants: WorstLink (Fig. 9) and WorstPath (Fig. 10). The former zooms in on the link with the max-PLR, and tries to reroute one by one (in a random order) all demands crossing this link by giving this link a high penalty. WorstPath focuses on a particular path, starting with the demand with the worst end-to-end PLR. When rerouting is successful, WorstPath continues with the path that now has the worst end-to-end PLR; but when, after a fixed number X of attempts to reroute this demand, this lowering of the max-PLR still is unsuccessful, the algorithm ignores this demand and moves on to the next demand in line (sorted on end-to-end PLR in decreasing order). Rerouting a demand in WorstPath is done by giving penalties to each of the links of the route currently followed for that particular demand (starting with the one exhibiting the highest PLR), in order to relieve the links most heavily suffering from loss.

The algorithm stops when all routes have PLRs below the requested upper bounds, or when the max-PLR cannot be improved any more (for the WorstLink variant this happens when, after changing the paths over the link with the worst PLR, this link still is the one with the max-PLR; in WorstPath this is when all paths have been tried X times without success).

Since the PLR is strongly related to the load on the links (see Section 2.2), one can expect that a load balancing algorithm —aiming at lowering the maximal load on each of the network's links— will also achieve significant PLR reduction compared to shortest path routing. Clearly, the advantage of such a load balancing approach is that the routing algorithm does not require any PLR-estimation. The results presented in the next section show that the more complex PLR-based approach can reach PLRs more than an order of magnitude lower than with load balancing.

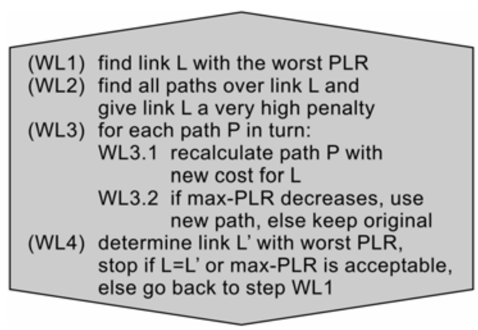

Fig. 9. Routing algorithm: algorithmic description of local phase – WorstLink variant.

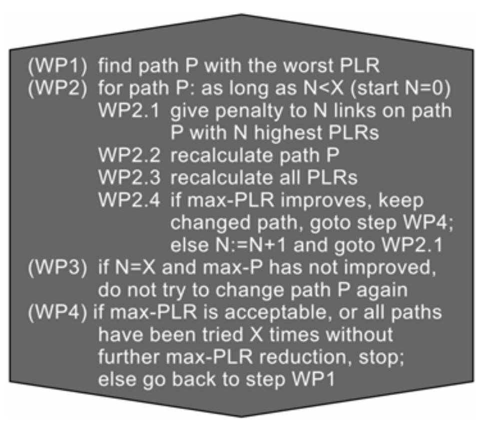

Fig. 10. Routing algorithm: algorithmic description of local phase – WorstPath variant.

In the algorithm descriptions above, we only considered traffic with one priority, but extension to traffic with multiple priorities is straightforward: the different phases can be run consecutively for all classes/priorities (starting with the class with the most strict packet loss demand, as lower priority traffic does not influence higher priority traffic). In Section 6 results of simulations with traffic of multiple priorities are presented.

Note that our heuristic is a local search technique, the descent method: a path-change is only accepted when the objective function (i.e., lower the max-PLR value) improves, which avoids cycling but could lead to a solution that is a local minimum far from the optimal solution. To prevent from ending up with such a local minimum, we perform a couple of different simulations (in average five simulations per demand matrix) for the same inputs (the same network, capacity, costs, demand matrix), but for a different (random) order of the demands. We observe that we indeed sometimes become another ultimate max-PLR when routing the demands in a different order. If this happens, we keep the lowest max-PLR value.

4 Routing algorithm performance

To evaluate the performance of our routing algorithms, we carried out case studies on two meshed pan-European networks, whose topological characteristics are summarized in Table 2. The first, illustrated in Fig. 11, connects 19 European cities. A joint effort from the IST project LION and COST action 266 resulted in topologies for a pan-European fiber-optic network [22]. We used such a network, sketched in Fig. 12, as a second test-case. It is somewhat sparser than the first: it has the same number of links, but interconnects 27 cities.

Given these topologies and the traffic demand matrices (bandwidth and PLR requirements), we want to find routes for the traffic that satisfy these demands. In the following results the networks comprise bi-directional links (one direction equals one fiber) and costs are in function of fiber length and capacity. Each of the fibers consists of 32 wavelengths and the buffer contains 32 FDLs, all of a length of one timeslot (a fix buffer, see introduction). In this section we assume Poisson traffic with a single priority (for which we want to determine the lowest PLR possible). We studied various demand patterns, both uniform and non-uniform.

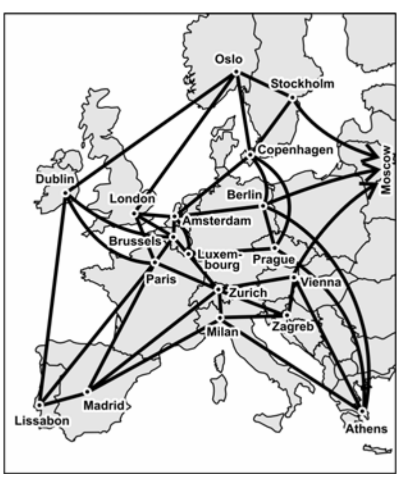

Fig. 11. The dense 19-node network.

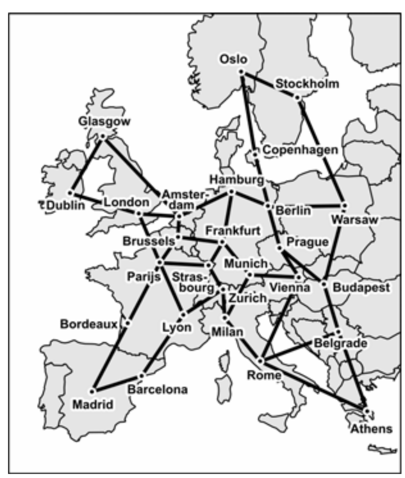

Fig. 12. The sparse 27-node network.

Table 2. Network scenarios.

Parameter	*Network 1 (Dense)*	*Network 2 (Sparser)*
Nr. of nodes	19	27
Nr. of links	40	40
Avg. node degree	4.21	2.93
Min. node degree	3	2
Max. node degree	6	5

A first set of simulations was used to compare the two alternatives (WorstLink and WorstPath) of the local part of the algorithm. These simulations are described in this section.

For the two variants of the algorithm, Fig. 13 shows the evolution of the max-PLR for a random (non-uniform) demand between all nodes from the second pan-European network, in function of the number of iterations, where a single iteration is an (attempt to) change a single path. Since we adopt a heuristic descent method, the maximal PLR does not increase for successive iterations.

When we compare the WorstLink and WorstPath version, we notice that the two versions give the same final result, for the plotted case, but WorstPath reaches this minimum max-PLR faster, i.e., after fewer iterations, than WorstLink. This observation is valid for the majority of the simulated cases. The reason why WorstPath reaches the minimum faster is that it always tries to change the worst path (i.e., the path with the worst end-to-end PLR) first. In WorstLink this is not necessarily so: the algorithm tries to change all the paths over the worst link in a random order, even if these paths are not the worst paths (anymore — if a path already has been changed and now another link suffers from the worst PLR).

When looking at the number of iterations the algorithm needs to end (see Table 3), WorstPath always needs more iterations to stop. WorstPath tries, if it is not successful, to change every path a fixed number of X times but WorstLink stops when all paths over the worst link were tried and this link has still the max-PLR over the complete network, so it does not try to change all possible paths.

In terms of the routes followed for each of the demands, we noticed that unsurprisingly this sometimes proves to be different for the two versions of the algorithm. This was to be expected, since the order in which the demands are (re)routed differs between the WorstLink and WorstPath variants. In terms of overall link capacity used, WorstPath is somewhat worse of (i.e., between 2% and 5% more used capacity) than WorstLink.

Although we see in most cases that WorstLink and WorstPath give the same final result and that WorstPath reaches this minimum max-PLR faster than WorstLink, also a few exceptions to this general rule were observed. A first exception that occurs is that WorstLink sometimes reaches the (same) minimum faster than WorstPath. A more important exception is illustrated below. Fig. 14 shows simulation results for the dense pan-European network: it plots the evolution of the max-PLR for an increasing number of iterations for another non-uniform demand.

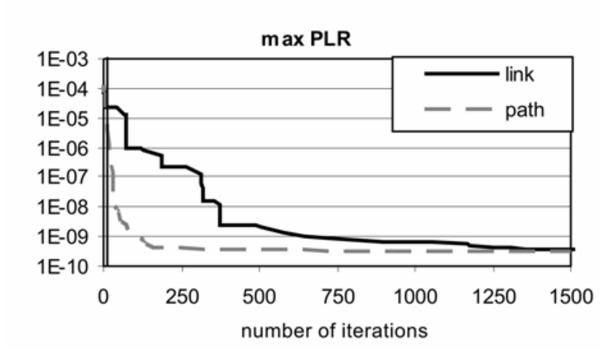

Fig. 13. Evolution of max-PLR (in function of number of iterations) for a random demand on the second pan-European network. The full line is the WorstLink version of the algorithm, the dashed line the WorstPath version. Only the first 1500 iterations are shown, since after that number the max-PLR doesn't change anymore. The vertical line on the graph marks where the global phase of the algorithm ends and the local phase begins.

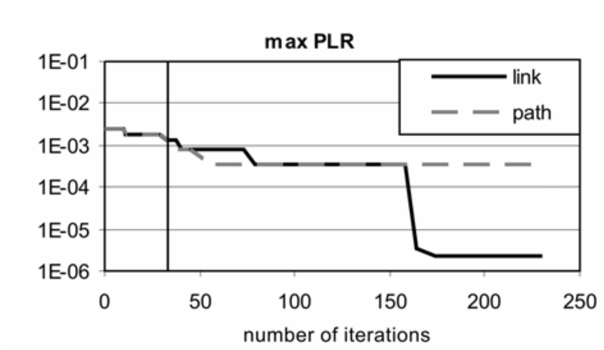

Fig. 14. Evolution of max-PLR (in function of number of iterations) for a random demand on the first pan-European network: we omitted the iterations beyond the first 250 because no further lowering of the max-PLR was attained. The vertical line marks the start of the local phase of the algorithm.

Comparing the WorstPath and WorstLink variants, we see that the WorstPath version never reaches the same minimum as WorstLink, so here the WorstLink version gives a much better solution (i.e., lower max-PLR) than WorstPath. An explanation for this worse performance of WorstPath is that after trying to change a path X times without success, this path is never tried again. However, it is possible that changing this path at a later stage, after having changed a few other paths (with a lower PLR), would lead to a lower max-PLR. The WorstLink algorithm indeed can re-consider a path tried before in a next iteration, when the worst link of that path becomes again link with overall max-PLR (this may happen after another link has exhibited the max-PLR which has successfully been lowered by rerouting some path(s)). If and when this difference in final result occurs, depends on the network, the demands, and the randomly chosen order in which the demands are routed.

In terms of capacity-use, these exceptional cases with respect to best PLR performance, do obey the formerly reached conclusion on capacity-use: we also notice a higher capacity-use with WorstPath.

We can thus summarize the conclusions of our WorstPath versus WorstLink algorithm variants as follows: WorstLink always achieves the lowest PLR and it also has the lowest penalty (in terms of extra capacity-use compared to shortest path routing). In the following, we will therefore apply the WorstLink variant.

Table 3. Number of iterations the algorithm needs to stop for WorstLink and WorstPath (for different demands).

scenario	*nr. of iterations to end* *WorstLink*	*WorstPath*
2-uni 0,38	320	2013
2-rt 0,36	1789	2564
1-uni 0.64	432	1872
1-S 0.54	186	733

5 Performance of PLR-based routing compared to shortest path routing and load balancing

To evaluate the developed PLR-based algorithms, we compared the max-PLR found with the WorstLink variant with the max-PLR for the initial shortest path routing based on link costs. Clearly, if the reduction in PLR would amount to only a small fraction of the original value, the usefulness of PLR-based routing would be marginal. To measure the cost of using a more intelligent routing strategy, we look at the price paid in terms of excess capacity-use when using the PLR-based routing.

Secondly, to assess the importance of estimating the PLR accurately, we also used a load balancing algorithm to find the routes to follow. Indeed, since the main factor impacting the PLR on a link is the load generated for it (see Section 2.2), we expect that by load balancing and thus limiting the loads, we can achieve already a substantial reduction of PLR compared to shortest path routing. The load balancing algorithm used was identical to the WorstLink algorithm described above, but using the link load as cost measure to minimize: in each iteration, we strive at lowering the maximal load on each of the network's links (i.e., lower max-load instead of max-PLR).

Fig. 15 below shows, for different demands on the two studied networks (on the X-axis), the max-PLR values reached: (i) with the WorstLink variant of the PLR-based routing algorithm, (ii) with load balancing, and (iii) with shortest path routing. The type of demand is indicated on the X-axis: uni stands for uniform, rl for real non-uniform —based on the traffic forecast in [22]—, S for random demands all of the same order, and L for random patterns where a few demands are of a larger order; the 1 and 2 refer to the 19-node and the 27-node network, respectively; the last number in each label on the X-axis is the overall mean network load when using shortest path routing (range [0,1]).

Table 4 shows the numerical values of the max-PLR values plotted in Fig. 15, the ratio shortest path routing/load balancing and the ratio load balancing/PLR-based.

The results presented in Fig. 15 and Table 4 show that the max-PLR obtained with the PLR-based algorithm can lie up to multiple orders of magnitude lower than when shortest path routing is used. Clearly, the type of the demand influences the factor of improvement: for non-uniform demands the reduction factor is higher because it is then more likely to have an unbalanced load in the considered equally meshed networks. Within a certain type of demand the reduction factor depends on the mean load on the network: the higher the overall load, the less room for improvement because lack of free capacity to reroute paths. Thus, apart from cases with very high overall network loads, the reduction of the max-PLR can be huge (multiple orders of magnitude).

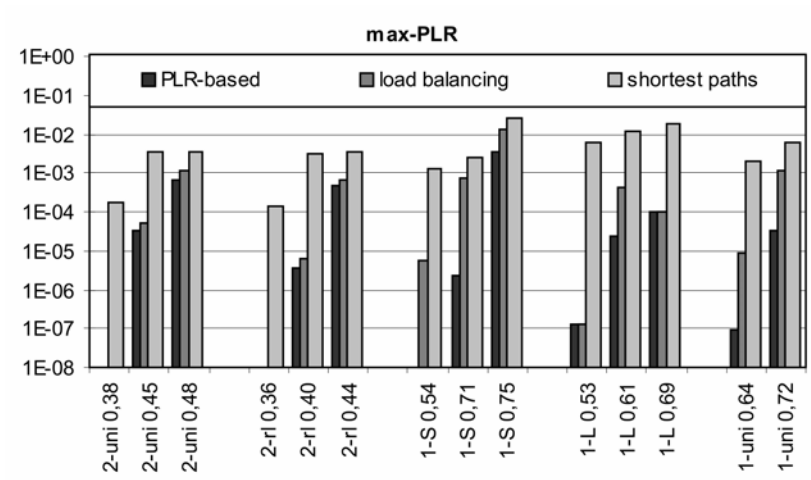

Fig. 15. Comparison of max-PLR reached with PLR-based routing, shortest path routing, and load balancing for different demands.

Table 4. Comparison of max-PLR reached with PLR-based routing, shortest path routing and load balancing for different demands.

	max-PLR			*ratio of max-PLRs*	
max-PLR	*shortest path*	*load balancing*	*PLR-based*	*shortest/ load*	*load/ PLR-based*
2-uni 0,38	1,77E-04	1,93E-10	1,21E-10	9,17E+05	1,60
2-uni 0,45	3,36E-03	5,31E-05	3,24E-05	6,33E+01	1,64
2-uni 0,48	3,48E-03	1,12E-03	6,80E-04	3,10E+00	1,65
2-rl 0,36	1,39E-04	6,89E-10	3,53E-10	2,01E+05	1,95
2-rl 0,40	3,13E-03	6,21E-06	3,19E-06	5,05E+02	1,95
2-rl 0,44	3,46E-03	6,82E-04	4,84E-04	5,08E+00	1,41
1-S 0,54	1,27E-03	5,42E-06	5,83E-09	2,34E+02	929,67
1-S 0,71	2,39E-03	7,77E-04	2,38E-06	3,07E+00	326,47
1-S 0,75	2,54E-02	1,35E-02	3,60E-03	1,87E+00	3,76
1-L 0,53	6,11E-03	1,33E-07	3,50E-08	4,60E+04	3,80
1-L 0,61	1,20E-02	4,38E-04	1,46E-06	2,74E+01	300,27
1-L 0,69	1,92E-02	1,02E-04	9,65E-05	1,88E+02	1,06
1-uni 0,64	1,94E-03	9,05E-06	4,51E-08	2,14E+02	200,67
1-uni 0,72	5,76E-03	1,19E-03	3,34E-05	4,86E+00	35,56

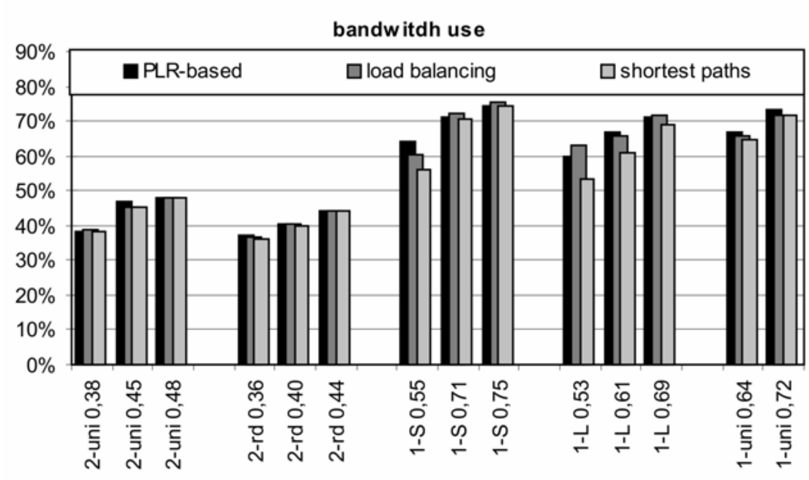

Fig. 16. Comparison of bandwidth-use (%) obtained with PLR-based routing, shortest path routing and load balancing for different demands.

The comparison of our PLR-based algorithm with load balancing shows why PLR-estimation can be useful: while load balancing offers huge improvement over shortest path routing, the PLR-based algorithm may offer additional reduction with a factor up to over two orders of magnitude (range of 1.5 to 900 for the analyzed cases). The reason is that load balancing ignores correlation between PLRs on outgoing links of the same node. This correlation stems from the sharing of the FDL buffer. For low mean network loads (0.35-0.5) the factor lies between 1.5 and 2: the load on the links is so low that the buffer is not heavily used, resulting in negligible correlation. For mean network loads above 0.5, the factor ranges from 2 to 900: the higher link loads result in heavy use of the buffer and hence correlated loss rates, since the (limited) buffer space has to be shared by all traffic crossing the node. However, there is also a bound on the improvement: with very high link loads (e.g., 0.69 of L in the first studied network) there is not much room to reroute the paths, regardless of the algorithm used.

Especially for the demand-types S and uni the PLR-based algorithm reaches max-PLR values more than one order of magnitude lower than the values obtained with the load balancing algorithm. An explanation for this can be found in the relatively high overall link load in these cases. As the PLR-based algorithm takes into account the influence of other heavily loaded output fibers, the routing is performed more accurate and lower max-PLR values are reached. For demand-type L the overall link load is also high, but the larger demands occurring here can not so easily be rerouted.

The “penalty” for the more intelligent routing (in terms of PLR) is a small increase in bandwidth-use (a few percent) compared to shortest path routing. This can be seen in Fig. 16, where the bandwidth utilization is plotted for the same demands as above (Fig. 15). Here, bandwidth utilization is defined as the sum over all links of the bandwidth used on that link, divided by the sum of the total available bandwidth on each link (a bandwidth utilization of 100% means all links in the network are fully used). The graph in Fig. 16 shows that both load balancing and PLR-based routing algorithms lead to slightly higher bandwidth utilization than shortest path routing. Clearly, this is caused by rerouting some demands from their original shortest path: more links are used to fulfill the demand. Obviously, the increase is the most pronounced in the case with some large demands (L). Still, even here the difference in bandwidth-use does not exceed 7%. Whether load balancing uses either more or less bandwidth than PLR-based routing depends on the case at hand, since the order in which demands are rerouted from their original shortest paths differs.

We can conclude that depending on how strict the max-PLR restriction for the demands are, one can either opt for the very simple shortest path routing, i.e., when no restrictions are given for the max-PLR; or, when lower max-PLR values are the goal, one can choose the more complex load balancing and finally when one wants an even better routing (in terms of PLR) one should opt for the even more complex PLR-based routing.

6 Influence of introducing multiple priorities on PLR-based routing

All simulations performed and discussed up to now considered traffic of one single priority. In this section, we are going to explore how introducing multiple priorities affects the max-PLR values reached with our developed PLR-based routing algorithm.

We start in Section 6.1 to compare the max-PLR values reached with monolithic traffic streams all of the same priority, and the values reached with the same amount of traffic but now dividing each source-destination stream into two parts of a different priority (of which 15, 30, or 45 percent highest priority traffic and correspondingly 85, 70, or 55 percent lowest priority traffic). Note that for this comparison two effects come into play: (i) the effect of splitting demands into smaller parts and (ii) the effect of using different priorities for these parts. To identify the influence of each of them, we separate them and discuss their impact individually in the following Subsections 6.2 and 6.3.

All assumptions concerning networks and node parameters in this section are the same as for the simulations with traffic of a single priority (see Section 4). In addition, we consider one extra network to perform simulations on: the well-known NSFNET (see Fig. 17). This network consists of 14 nodes and 21 links and has an average node degree of 2.93, a minimum node degree of 2 and a maximum node degree equal to 4.

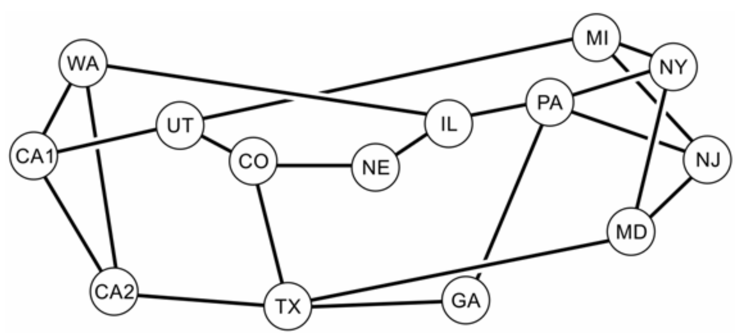

Fig. 17. The NSFNET.

Regarding the traffic assumptions there are, like with one priority, different (non-)uniform demands and the traffic is Poisson-distributed. The only difference with the single priority-case is that the traffic consists of two priorities with following 'high priority traffic/low priority traffic' ratios: 15/85 - 30/70 - 45/55. The goal was to obtain, for both priorities, max-PLRs as low as possible.

As indicated before (see Section 3.2), our PLR-based algorithm can easily be extended from one to multiple priorities. Therefore, the algorithm used in this section is the WorstLink variant of the developed PLR-based algorithm. This algorithm is consecutively run two times: first for the highest priority traffic, then for the lowest priority traffic. For the PLR calculation we always use the formula of Section 2.2, which was actually developed for one priority traffic but can also be used with multiple priorities. In the first run of the algorithm we calculate the PLR for the highest priority traffic with the formula (as lower priority traffic does not influence higher priority traffic, see introduction). In the second run, for the lowest priority, we calculate the PLR for the total amount of highest and lowest priority traffic and substract the PLR obtained for the highest priority class traffic.

6.1 Splitting demands into multiple priorities.

Fig. 18 shows, for different demands on the three studied networks (on the X-axis), a comparison of the max-PLR values reached with the WorstLink variant of the PLR-based algorithm, when (i) considering traffic of a single priority (dash-dotted line with x markers in the figure, 100/0 highest/lowest priority ratio) and (ii) when considering different ratios of traffic of two priorities (ratio high priority traffic/low priority traffic respectively dotted with diamond markers for 15/85; full with square markers for 30/70; dashed with triangular markers for 45/55 in the figure).

As before, the type of demand is indicated on the X-axis: uni stands for uniform, rl for real non-uniform, rd for random non-uniform, S for random demands all of the same order, and L for random patterns where a few demands are of a larger order. The 1 and 2 refer to the 19-node and the 27-node network, nsf refers to the NSFNET; the last number in each label on the X-axis is the overall mean network load when considering shortest path routing (range [0,1]).

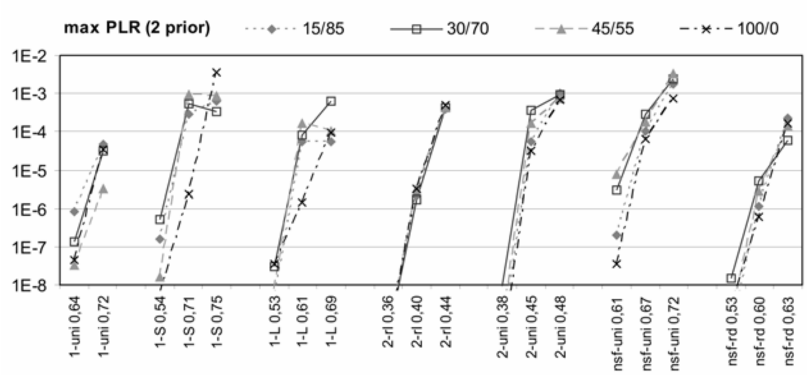

Fig. 18. Comparison of max-PLR values reached with the WorstLink variant of the PLR-based algorithm, when (i) considering traffic of a single priority (100/0 ratio highest/lowest priority ratio) and (ii) when considering different ratios of traffic of two priorities (15/85;30/70;45/55 ratio high priority traffic/low priority traffic). This for different demands on the three studied networks (on the X-axis).

Table 5 shows the minimal max-PLR values reached for the different demands and the different highest/lowest priority ratios. It also shows (in the last three columns) the ratio between the max-PLR value reached with one priority traffic and the max-PLR value reached with two priorities of traffic (respective ratios: 15/85, 30/70, 45/55).

Table 5. Comparison of max-PLR values reached with the WorstLink variant of the PLR-based algorithm, when (i) considering traffic of a single priority (100/0 ratio highest/lowest priority ratio) and (ii) when considering different ratios of traffic of two priorities (15/85;30/70;45/55 ratio high priority traffic/low priority traffic). This for different demands on the three studied networks (on the X-axis).

	max-PLR for two priorities (high/low)				*ratio of PLRs (with/no prio)*		
scenario	*15/85*	*30/70*	*45/55*	*100/0*	*15/100*	*30/100*	*45/100*
1-uni 0,64	8.28E–07	1.31E–07	3.26E–08	4.51E–08	18.36	2.90	0.72
1-uni 0,72	4.64E–05	3.04E–05	3.28E–06	3.34E–05	1.39	0.91	0.10
1-S 0,54	1.61E–07	5.09E–07	1.63E–08	5.83E–09	27.62	87.31	2.80
1-S 0,71	2.95E–04	5.54E–04	9.60E–04	2.38E–06	123.95	232.77	403.36
1-S 0,75	6.32E–04	3.31E–04	8.92E–04	3.60E–03	0.18	0.09	0.25
1-L 0,53	3.79E–09	2.96E–08	7.74E–09	3.50E–08	0.11	0.85	0.22
1-L 0,61	5.58E–05	8.09E–05	1.69E–04	1.46E–06	38.22	55.41	115.75
1-L 0,69	5.36E–05	6.45E–04	1.05E–04	9.65E–05	0.56	6.68	1.09
2-rl 0,36	3.66E–10	4.02E–10	3.37E–10	3.53E–10	1.04	1.14	0.95
2-rl 0,40	2.23E–06	1.74E–06	3.07E–06	3.19E–06	0.70	0.55	0.96
2-rl 0,44	3.96E–04	4.23E–04	4.36E–04	4.84E–04	0.82	0.87	0.90
2-uni 0,38	3.58E–10	9.55E–09	8.24E–10	1.21E–10	2.96	78.93	6.81
2-uni 0,45	5.52E–05	3.75E–04	1.62E–04	3.24E–05	1.70	11.57	5.00
2-uni 0,48	1.00E–03	9.39E–04	7.50E–04	6.80E–04	1.47	1.38	1.10
nsf-uni 0,61	1.94E–07	3.01E–06	7.72E–06	3.66E–08	5.30	82.24	210.93
nsf-uni 0,67	9.90E–05	2.97E–04	1.74E–04	6.59E–05	1.50	4.51	2.64
nsf-uni 0,72	1.80E–03	2.49E–03	3.30E–03	7.62E–04	2.36	3.27	4.33
nsf-rd 0,53	6.56E–10	1.53E–08	1.35E–09	2.69E–09	0.24	5.69	0.50
nsf-rd 0,60	1.15E–06	4.94E–06	2.81E–06	5.92E–07	1.94	8.34	4.75
nsf-rd 0,63	2.22E–04	5.89E–05	1.39E–04	1.71E–04	1.30	0.34	0.81
mean	3.84E–06	8.66E–06	4.64E–06	1.87E–06	2.05	4.64	2.48

When comparing the max-PLR values reached with traffic of one single priority (hundred percent of one priority) and the values reached with traffic of two priorities, we observe (see Table 5 and Fig. 18) that in most cases (i.e., for most of the demands) the minimal max-PLR values obtained with traffic of two priorities are higher than the values reached with traffic of one single priority. The minimal max-PLR values reached for traffic of two priorities given, are values for the lowest priority traffic, the PLR values for the highest priority traffic are much lower (<< 1E-10). So, the penalty of introducing priorities is a small increase of the minimal max-PLR value for the lowest priority traffic, while the PLR-values of the higher priority traffic can be kept very low.

This increase of the minimal max-PLR can be explained as follows. With two priorities of traffic, we first route the traffic with the highest priority as good as possible. After this, we route the lower priority traffic, but, when routing this traffic, we can not change the routes of the higher priority traffic any more, so we have less flexibility to reroute the lowest priority traffic, which results in a slightly higher minimal max-PLR value.

We however observe also some results where the opposite holds: the same amount of traffic divided into two priority classes returns better results (i.e., a lower max-PLR) than all traffic of the same priority. In most of these exceptional cases, the network load is relatively high. An explanation for this could be that we not only have to consider the effect of introducing priorities but we should also take into account the influence of the splitting of the demands into smaller parts, on the PLR-values. With high network loads we could benefit from this division of large demands into smaller parts. As all traffic of one demand is sent over the same path, dividing the traffic over two different paths could give better results because we then route smaller traffic parts. This issue is addressed in the following Subsection 6.2. Rarely (e.g., in case 1-L 0.53) the exception of reaching lower PLRs when splitting traffic into two priorities also occurs for a low network load. This can be put down to inaccurate PLR-values in this range.

Averaging the values over the different demands/networks (see last row of Table 5), we observe that the averaged max-PLR value obtained with traffic of two priorities (for the low priority traffic), is, for all high/low priority ratios, at most a factor 4.7 higher than the averaged max-PLR value reached with pure one priority traffic. Thus, for two priorities traffic there is only a small increase of the max-PLR value for the lowest priority traffic: while attaining negligible loss rates for high priority traffic, loss rate for low priority streams slightly increases but remains within the same order of magnitude.

In the following we split up the two effects that occur when introducing priorities, therefore we first (Subsection 6.2) take a look at the individual impact of splitting demands into smaller parts on the max-PLR values and next (Subsection 6.3) we study the effect of using different priorities for these smaller parts.

6.2 The effect of splitting demands into smaller parts

To study the effect of splitting (large) demands into smaller parts, we split up the monolithic traffic streams of one priority into two smaller parts/demands of the same priority. We start with routing the demands with 15/30/45 percent of the traffic, followed by the larger demands of 85/70/55 percent of traffic of the same priority. As a result we consider now the same traffic demands as in the two priorities case, except that demands now all have the same priority. The difference with the case with two priorities is that here the smaller parts of the demands, which are routed first, can, in contrast to the highest priorities after the first run of the algorithm in the two priorities case, still be rerouted in the following iterations of the algorithm, even when bigger parts have already been rerouted.

Fig. 19 shows, for the same demands as above, the max-PLR values reached with the WorstLink version of the algorithm for demands of pure one priority traffic (dash-dotted line with x markers) and for the same demands split into two parts of the same priority: 15/85 (dotted line with diamond markers); 30/70 (full line with square markers); 45/55 (dashed line with triangular markers).

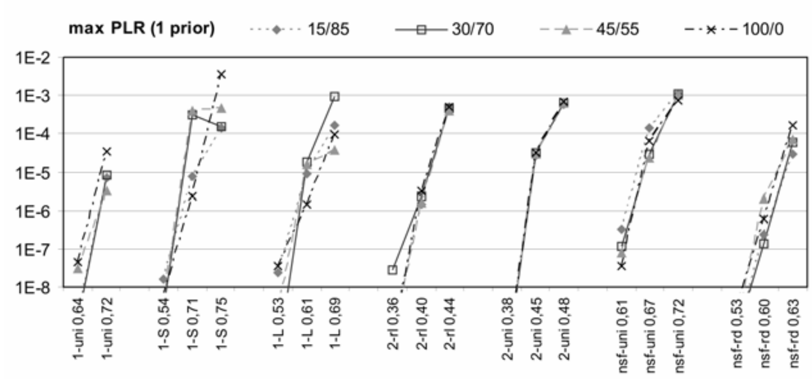

Fig. 19. Comparison of max-PLR values reached with the WorstLink variant of the PLR-based algorithm for traffic of one priority (i) for demands of 100% (100/0) and (ii) for the same demands split into two parts: 15/85; 30/70; 45/55.

Comparing now (see Fig. 19) the max-PLR values reached when splitting the demands into smaller parts (15/85, 30/70, 45/55) with the values reached with demands which are not split, we see that indeed in many cases the minimal max-PLR values reached with the demands split into two smaller parts, are lower than when demands are not split. When averaging the values over the different demands, we observe a small decrease of the mean value reached when demands are split, showing a (limited) advantage of dealing with finer granularity streams.

6.3 Routing smaller parts: the effect of introducing priorities.

Fig. 20 to 22 show, again for the same demands as above, for the different traffic demand ratios (15/85; 30/70; 45/55) the minimal max-PLR values reached when the two parts the demands are split into are (i) of the same priority, (ii) of a different (high/low) priority. Table F.6 also shows these values and a ratio of (ii) to (i).

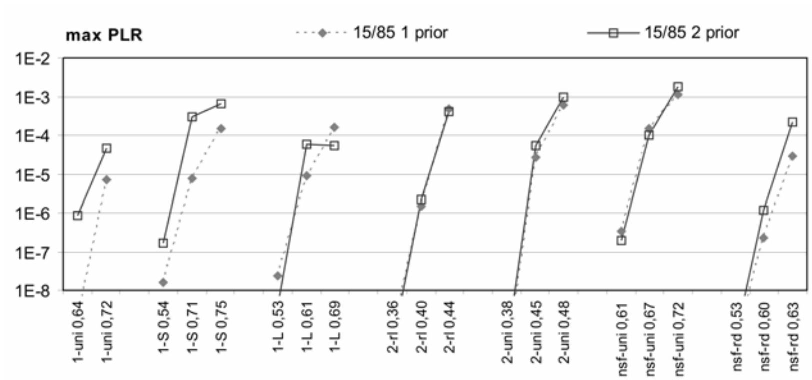

Fig. 20. Comparison of the max-PLR values reached for the traffic demand ratio 15/85 when the two parts the demands are split into are (i) of the same priority, (ii) of a different (high/low) priority.

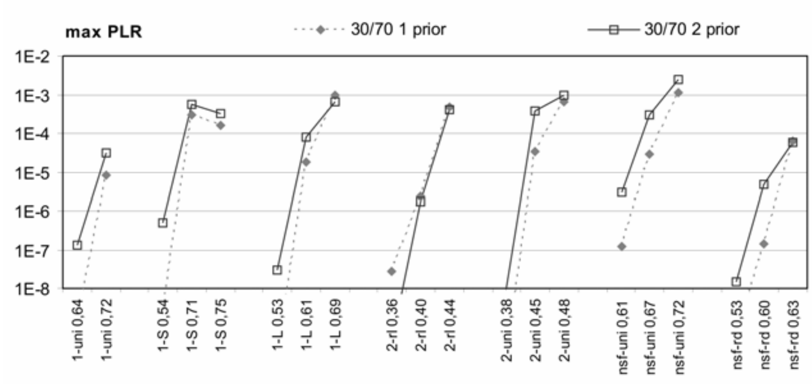

Fig. 21. Comparison of the max-PLR values reached for the traffic demand ratio 30/70 when the two parts the demands are split into are (i) of the same priority, (ii) of a different (high/low) priority.

To study the effect of introducing priorities in traffic split up in smaller parts, we compare the max-PLR values reached with the demands of one priority traffic split into two parts (15/85, 30/70, 45/55) with the max-PLR values obtained with the same demands, but where the two parts the demands are split into are of a different priority.

We see (Fig. 20 to 22 and Table F.6) that in most cases the minimal max-PLR values reached for the two priorities case (which are again the values for the lowest priority traffic as the values for the highest priority traffic are much lower) are slightly higher than the ones for split up one priority traffic. In only a few cases the minimal max-PLR value reached with traffic with two priorities is lower than with one priority traffic.

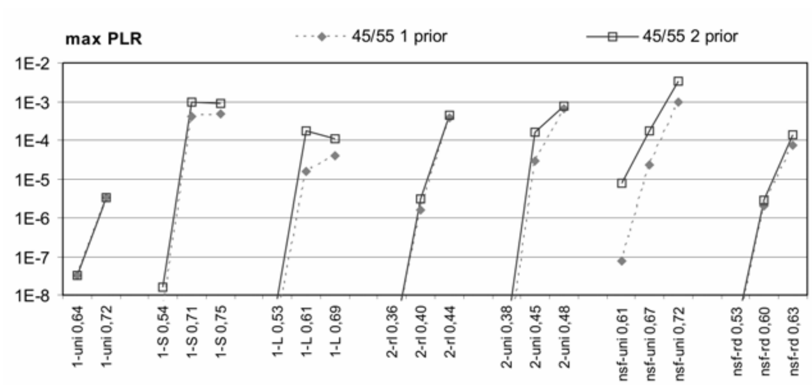

Fig. 22. Comparison of the max-PLR values reached for the traffic demand ratio 45/55 when the two parts the demands are split into are (i) of the same priority, (ii) of a different (high/low) priority.

When we average the values over the different demands and networks (for these values see the last row of Table F.6), we see that the averaged value of traffic of two priorities, is at most a factor 5.7 higher than the averaged max-PLR value reached with traffic of one priority.

An explanation for these higher values with two priorities of traffic is thus the smaller flexibility in rerouting the lowest priority traffic: once the highest priority traffic has been routed, it is left untouched when routing the lower priority.

.

Table 6. Comparison of the minimal max-PLR values reached for the different traffic demand ratios (15/85; 30/70; 45/55) when the two parts the demands are split into are (i) of the same priority, (ii) of a different (high/low) priority.

	15/85			*30/70*			*45/55*			*100/0*
scenario	*1 pri*	*2 pri*	*1 pri / 2 pri*	*1 pri*	*2 pri*	*1 pri / 2 pri*	*1 pri*	*2 pri*	*1 pri / 2 pri*	
1-uni 0.64	1.61E-09	8.28E-07	5.14E+02	1.42E-09	1.31E-07	9.23E+01	3.12E-08	3.26E-08	1.04E+00	4.51E-08
1-uni 0.72	6.82E-06	4.64E-05	6.80E+00	8.03E-06	3.04E-05	3.79E+00	3.28E-06	3.28E-06	1.00E+00	3.34E-05
1-S 0.54	1.65E-08	1.61E-07	9.76E+00	2.68E-09	5.09E-07	1.90E+02	4.03E-09	1.63E-08	4.04E+00	5.83E-09
1-S 0.71	7.78E-06	2.95E-04	3.79E+01	2.99E-04	5.54E-04	1.85E+00	4.04E-04	9.60E-04	2.38E+00	2.38E-06
1-S 0.75	1.45E-04	6.32E-04	4.36E+00	1.58E-04	3.31E-04	2.09E+00	4.77E-04	8.92E-04	1.87E+00	3.60E-03
1-L 0.53	2.31E-08	3.79E-09	1.64E-01	1.46E-10	2.96E-08	2.03E+02	2.12E-09	7.74E-09	3.65E+00	3.50E-08
1-L 0.61	8.70E-06	5.58E-05	6.41E+00	1.76E-05	8.09E-05	4.60E+00	1.50E-05	1.69E-04	1.13E+01	1.46E-06
1-L 0.69	1.60E-04	5.36E-05	3.35E-01	9.42E-04	6.45E-04	6.85E-01	3.82E-05	1.05E-04	2.75E+00	9.65E-05
2-rl 0.36	7.40E-10	3.66E-10	4.95E-01	2.85E-08	4.02E-10	1.41E-02	3.40E-10	3.37E-10	9.91E-01	3.53E-10
2-rl 0.40	1.49E-06	2.23E-06	1.50E+00	2.32E-06	1.74E-06	7.50E-01	1.59E-06	3.07E-06	1.93E+00	3.19E-06
2-rl 0.44	4.95E-04	3.96E-04	8.00E-01	4.78E-04	4.23E-04	8.85E-01	3.81E-04	4.36E-04	1.14E+00	4.84E-04
2-uni 0.38	1.21E-10	3.58E-10	2.96E+00	1.21E-10	9.55E-09	7.89E+01	1.21E-10	8.24E-10	6.81E+00	1.21E-10
2-uni 0.45	2.66E-05	5.52E-05	2.08E+00	3.24E-05	3.75E-04	1.16E+01	2.86E-05	1.62E-04	5.66E+00	3.24E-05
2-uni 0.48	5.93E-04	1.00E-03	1.69E+00	6.35E-04	9.39E-04	1.48E+00	6.34E-04	7.50E-04	1.18E+00	6.80E-04
nsf-uni 0.61	3.25E-07	1.94E-07	5.97E-01	1.19E-07	3.01E-06	2.53E+01	7.82E-08	7.72E-06	9.87E+01	3.66E-08
nsf-uni 0.67	1.43E-04	9.90E-05	6.92E-01	2.80E-05	2.97E-04	1.06E+01	2.34E-05	1.74E-04	7.44E+00	6.59E-05
nsf-uni 0.72	1.11E-03	1.80E-03	1.62E+00	1.12E-03	2.49E-03	2.22E+00	1.00E-03	3.30E-03	3.30E+00	7.62E-04
nsf-rd 0.53	4.39E-10	6.56E-10	1.49E+00	5.50E-10	1.53E-08	2.78E+01	8.16E-10	1.35E-09	1.65E+00	2.69E-09
nsf-rd 0.60	2.27E-07	1.15E-06	5.07E+00	1.40E-07	4.94E-06	3.53E+01	2.00E-06	2.81E-06	1.41E+00	5.92E-07
nsf-rd 0.63	2.80E-05	2.22E-04	7.93E+00	6.07E-05	5.89E-05	9.70E-01	7.06E-05	1.39E-04	1.97E+00	1.71E-04
mean	1.41E-06	3.84E-06	2.72E+00	1.53E-06	8.66E-06	5.67E+00	1.58E-06	4.64E-06	2.94E+00	1.57E-06

We conclude that introducing multiple priorities increases the minimal max-PLR value for the lowest priority traffic while for the highest priority the max-PLR value can be kept very low. As the increase of the minimal max-PLR value for the lowest priority traffic is not significant, the cost of introducing priorities is small.

7 Conclusion

We explained how results obtained for a single node could be used to assess end-to-end PLR for demands routed in a network of interconnected OPRs. We developed routing algorithms aiming at minimizing the PLR using this PLR-estimation technique. Two alternative algorithms were compared and the WorstLink variant proved to lead to the best results.

Case studies on pan-European networks illustrated that the PLR-based routing outperforms shortest path routing by lowering the max-PLR occurring in the network with multiple orders of magnitude. Compared to load balancing, which does not need PLR-estimation, the PLR-based algorithm can reach PLRs up to two orders of magnitude better. The price paid for the reduction in PLR is a higher overall bandwidth-use. Still, the amount of extra bandwidth needed compared to shortest path routing is quite limited (only a few percent).

It was also shown that the discussed PLR-based routing algorithm can be easily extended to multiple priorities. Introducing multiple priorities enables to keep loss rates negligible for high priority traffic, while it may lead to an increase of the obtained minimal max-PLR value for low priority traffic. However, the increase proved to be limited: the cost of introducing multiple priorities is small.

References

[1] B. Mukherjee, “WDM optical communication networks: Progress and challenges”, IEEE J. Select. Areas Commun., vol. 18, no. 10, Oct. 2000, pp. 1810–1824.

[2] 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.

[3] D.K. Hunter, I. Andonovic, “Approaches to optical internet packet switching”, IEEE Commun. Mag., vol. 38, no. 9, Sept. 2000, pp. 116–120.

[4] A. Hill and F. Neri, guest ed., “Optical switching networks: from circuits to packets”, IEEE Commun. Mag., vol. 39, no. 3, March 2001, pp. 107-148.

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

[6] C. Qiao, “Labeled optical burst switching for IP-over-WDM integration”, IEEE Commun. Mag., vol. 38, no. 9, Sept. 2000, pp. 104–114.

[7] D.K. Hunter et al., “WASPNET – a wavelength switched packet network”, IEEE Commun. Mag., vol. 37, no. 3, March 1999, pp. 120–29.

[8] Soren L. Danielsen, Carsten Joergensen, Benny Mikkelsen, Kristian E. Stubkjaer, “Analysis of a WDM Packet Switch with Improved Performance Under Bursty Traffic Conditions Due to Tuneable

Wavelength Converters", IEEE/OSA Journal of Lightwave Technology, vol. 16, no. 5, May 1998, pp. 729–735.

[9] S. Yao, B. Mukherjee, S. J. Ben Yoo and S. Dixit, "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, USA, Oct. 2000), pp. 235–246.

[10] H. Zang, J.P. Jue, and B. Mukherjee, "Capacity allocation and contention resolution in a photonic slot routing all-optical WDM mesh network", IEEE/OSA J. Lightwave Tech., vol. 18, no. 12, Dec. 2000, pp. 1728–41.

[11] D.K. Hunter, M.C. Chia and I. Andonovic, "Buffering in optical packet switches", IEEE/OSA J. Lightwave Tech., vol. 16, no. 12, Dec. 1998, pp. 2081–94.

[12] L. Dittman et al, "The IST project DAVID: a viable approach towards optical packet switching", to be published in IEEE J. Select. Areas Commun.
http://david.com.dtu.dk

[13] 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), vol. 6, pp. 60–61.

[14] C. Develder, M. Pickavet, P. Demeester, "Assessment of packet loss for an optical router with recirculating buffer", in Proc. ONDM 2002 6th Working Conf. on Optical Network Design and Modeling (Turin, Italy, February 4–6, 2002), pp. 247-261.

[15] C. Develder, M. Pickavet, P. Demeester, "Choosing an appropriate buffer strategy for an optical packet switch with a feed-back FDL buffer", in Proc. 28th European Conf. on Optical Communication, (Copenhague, Denmark, September 8–12, 2002), vol. 3, pp. 8.5.4 (2 pages).

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

[17] B. Fortz, M. Thorup, "Internet Traffic Engineering by Optimizing OSPF Weights", in Proc. IEEE Infocom 2000, (Tel Aviv, Israel, March 26-30, 2000), pp. 519-528.

[18] J.T. Moy, "OSPF: Anatomy of an Internet Routing Protocol", Addison-Wesley, February 1998.

[19] B. Davie, Y. Rekhter, "MPLS: Technology and applications", Morgan Kaufman Publishers, January 2000.

[20] 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.

[21] M. Pickavet, C. Develder, E. Baert, P. Demeester, "A.I. Techniques for planning telecommunication networks", Proc. International Conference on Artificial Intelligence IC-AI'02, (Las Vegas, NV, June 24-27, 2002), vol.2, pp. 608–614.

[22] S. De Maesschalck, et al., "Pan-European optical transport networks: an availability-based comparison", Photonic Network Communications, vol. 5, no. 3, May 2003, pp. 203–225.
http://ibcn.atlantis.UGent.be/projects/COST266_IST_lion/NRS/index.html

Biographies

Elise Baert received a M. Sc. degree in Electrotechnical Engineering (option Communication Techniques) in 2001 from the University of Ghent, Belgium. Since September 2001, she has been working on optical packet switched networks as a research assistant at the Department of Information Technology of the University of Ghent where she works in the Integrated Broadband Communications Network group. In this context she is involved in the European IST project “Data and Voice Integration over DWDM” (DAVID) as well as in a national research project on "Optical Networking and Node Architectures".

Chris Develder (chris.develder@intec.ugent.be) received his M.Sc. in Computer Science Engineering in 1999 from the Ghent University, Belgium. In October 1999, he joined the Department of Information Technology of the same university. There, as a researcher for the Fund for Scientific research of Flanders (FWO-Vlaanderen), he is working towards a Ph.D. in the field of network design and planning, mainly focusing on optical packet switched networks. As such, he is involved in the European IST-projects DAVID and STOLAS, as well as a national research project on "Optical Networking and Node Architectures".

Didier Colle received a M. Sc. degree in electrotechnical engineering (option: communications) from the Ghent University in 1997. Since then, he has been working at the same university as researcher in the department of Information Technology (INTEC). He is part of the research group INTEC Broadband Communication Networks (IBCN) headed by prof. Piet Demeester. His research lead to a Ph.D. degree in February 2002. From January 2003 on, 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 focussing on optical transport networks, to support the next-generation Internet. Up till now, 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.

Mario Pickavet received a M.Sc. degree in electrical engineering, specialized in telecommunications, in 1996 from Ghent University. From 1996 until 1999, he has been working as a researcher in the Broadband Communications Networks Group (Department of Information Technology, Ghent University). In 1999, he received a Ph.D. degree in electrical engineering from the same university. Since 2000, he is full-time professor at Ghent University where he is teaching telecommunication networks and algorithm design.
His current research interests are related to broadband communication networks (WDM, IP, (G-)MPLS, OPS, OBS) and include design, long-term planning and routing of core and access networks. Special attention goes to Operations Research techniques that can be applied for routing and network design. In this context, he is currently involved a.o. in the European IST projects “Data and Voice Integration over DWDM” (DAVID) and "Switching Technologies for Optically Labeled Signals" (STOLAS).

Piet Demeester received his PhD degree from the Ghent University at the Department of Information Technology (INTEC) in 1988. He became professor at the 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. He is member of the editorial board of several international journals and has been member of several technical program committees. His current interests are related to broadband communication networks (IP, G-MPLS, optical packet and burst switching, access and residential, active, mobile, CDN, grid) and include network planning, network and service management, telecom software, internetworking, network protocols for QoS support, etc. He published over 200 journal or conference papers in this field. In the past he was responsible for the research on epitaxial growth for photonic devices.