

Constrained and Unconstrained overspill routing in optical networks: a detailed performance evaluation

Kyriakos Vlachos · Erik Van Breusegem ·
Kostas Christodoulopoulos · Didier Colle ·
Kostas Ramantas · Piet Demeester

Received: 11 July 2006 / Revised: 6 September 2006 / Accepted: 20 September 2006
© Springer Science+Business Media, LLC 2006

Abstract In this article, we present a detailed performance evaluation of a hybrid optical switching (HOS) architecture called Overspill Routing in Optical Networks (ORION). The ORION architecture combines (optical) wavelength and (electronic) packet switching, so as to obtain the individual advantages of both switching paradigms. In particular, ORION exploits the possible idle periods of established lightpaths to transmit packets destined to the next common node, or even directly to their common end-destination. Depending on whether all lightpaths are allowed to simultaneously carry and terminate overspill traffic or overspill is restricted to a sub-set of wavelengths, the architecture limits itself to *constrained* or *un-constrained* ORION. To evaluate both cases, we developed an extensive network simulator where the basic features of the ORION architecture were modeled, including suitable edge/core node switches and load-varying sources to simulate overloading traffic conditions. Further, we have assessed various aspects of the ORION architecture including two basic routing/forwarding policies and various buffering schemes. The complete network study shows that ORION can absorb temporal traffic overloads, as intended, provided sufficient buffering is present. We also demonstrate that the restriction of simultaneous packet

insertions/extractions, to reduce the necessary interfaces, do not deteriorate performance and thus the use of *traffic concentrators* assure ORION's economic viability.

Keywords Hybrid optical switching · Optical networks · Overspill routing · Photonic switching systems

Introduction

The advent of WDM technology has resulted in transmission capacities that have increased manifold in recent years. It is the router/switch throughput, however, that really transforms the raw bit rates into effective bandwidth, and commercially available switching technologies are typically capable of handling line rates of up to 40 Gb/s. Current optical networks are wavelength (circuit) switched, where optical cross-connects (OXC) are used to switch traffic [11]. Optical circuit switching (OCS) is perfectly fitted for relatively static traffic profiles. However, it is well-known that Internet traffic exhibits multifaceted burstiness and correlation structures over a wide span of time scales (short and long time variations). Therefore, the use of optical circuits to transport IP traffic results in low capacity utilization, primarily due to the low statistical multiplexing efficiency that can be achieved. Optical packet switching (OPS) and optical burst switching (OBS) [12,10] have been proposed for “*on demand*” use of capacity. Although very promising, these technologies lag behind from the lack of a true optical random access memory. This stresses the need for “*on-the-fly*” optical packet processing, which is infeasible with the current state of the art of optical

K. Vlachos (✉) · K. Christodoulopoulos · K. Ramantas
Computer Engineering and Informatics Department &
Research Academic Computer Technology Institute,
University of Patras, GR26500, Rio, Patra, Greece
e-mail: kvlachos@ceid.upatras.gr

E. Van Breusegem · D. Colle · P. Demeester
Department of Information Technology (INTEC),
IBBT-Ghent University-IMEC,
Sint-Pietersnieuwstraat 41, B-9000 Gent, Belgium.
e-mail: erik.vanbreusegem@intec.ugent.be

logic [15]. A promising solution, easier to implement than OPS and more efficient than OCS, is hybrid optical switching (HOS), which is a compromise between (electronic) packet and (optical) circuit switching. Hybrid switching combines the merits of both switching paradigms to increase link utilization efficiency, to decrease the required number of wavelengths and to constrain the processing overhead of the IP routers.

Various schemes have been proposed so far for a HOS system, including HOS [7], [18] and the hybrid optical transport network (HOTNET), [9], where optical circuit and message switching are integrated in a complementary manner to cooperatively transport a variety of traffic types efficiently. In the HOS approach, best effort traffic is transported using OBS, while high priority traffic uses OCS. At the ingress node, OBS and OCS fairly compete for all wavelengths during resource reservation process, while the core nodes support both OCS and OBS burst switching. In [8], two strategies are proposed and evaluated: (i) with no priorities between circuits and bursts and (ii) when circuits are given preemptive priority over bursts. In *HOTNET*, [8], proposed the modification of the OBS scheme to meet the operational model of a TDM wavelength routed network. Therefore, a time-slotted OBS approach is co-implemented with OCS in TDM frames and thus the two switching technologies share all transport, control and switching resource.

Other hybrid switching approaches include the *light-trail* [5], [6] and the *light-bus* concept [1], as well as the *polarization-based scheme* [2]. In the *light-trail* concept, a *light-trail* forms the basic switching entity. A *light-trail* is a lightpath that is set up between two nodes by configuring the optical shutters (ON–OFF switches) of the source nodes, the intermediate as well as the destination nodes. All nodes across a *light-trail* have access to the data and can initialize connections—transmit bursts to the other nodes. A *light-bus* is a *light-trail* with buffers and electronic control at a node. Finally, in the *polarization-based* concept, the polarization state (SOP) is used to differentiate OCS traffic from best effort IP traffic.

A significant new hybrid approach called Overspill Routing In Optical Networks (ORION) has been proposed in [13], [14] and partially evaluated in [3]. ORION is a hybrid architecture that allows full sharing of all wavelengths on a link, without using large amounts of electronic switching, or resorting to deflection routing. ORION is based on a reconfigurable wavelength switched (WS) network, which can react to long-term traffic variations by reconfiguring the wavelength paths. Additionally, in ORION packets can be transmitted during the idle periods of established lightpaths and can be forwarded either to their next node or directly to their

own end-destination if this matches the lightpath destination. In ORION, packets do not compete with OCS traffic during resource reservation and the core nodes can insert/extract packets to/from any passing through or node initiated wavelength channel.

In this article, we present the first detailed and network wide performance evaluation of the ORION architecture for the case of *constrained* and *un-constrained* overspill routing. In the latter case, all wavelengths are allowed to simultaneously carry and terminate overspill traffic by employing a separate transmitter and receiver per wavelength, while in the former case the number of packets that can be simultaneously extracted or inserted is constrained by two concentrator interfaces, one for the extraction and one for the insertion of overspill packets. A detailed ORION network simulator has been developed to evaluate both cases, where the basic features of the architecture were modeled. ORION performance is assessed with load-varying traffic sources to determine its ability to absorb temporal network overloads, which in the normal case of a pure OCS network would require additional wavelengths. We have used the ORION network simulator to assess various aspects of the architecture. In particular, we have evaluated the performance of two routing policies, called “*once ORION always ORION*” and “*lightpath re-entry*” as well as the performance of various buffering schemes in the core/edge nodes.

The rest of the article is organized as follows. Section “Overspill routing in optical networks” presents the ORION switching architecture and the two considered routing schemes. Section “ORION network simulator” describes the developed “ORION network simulator” and the implemented edge/core node modules, while Section “Network level performance evaluation—unconstrained overspill routing” presents performance evaluation results. Finally, Section “The case of constrained overspill routing” discusses and present results of the special case of ORION paradigm when the number of packets that can be simultaneously extracted or inserted is limited routing.

Overspill routing in optical networks

ORION architecture

ORION is based on an automatically switched optical network (ASON), where lightpath connections are established with OXCs. The latter can react to long-term traffic pattern variations of hours, or even days, by reconfiguring the wavelength paths. ORION architecture allows full sharing of all wavelengths on a link and

thus significant statistical multiplexing gains can be obtained. In order to cope with short-term temporal traffic imbalances, ORION employs selective and transparent insertion/ extraction of data on wavelength paths at ORION-enabled nodes. This ability results in a network that operates preferably like a WS network, but if necessary (when congestion arises at certain flows), as a packet-switched network. The basic idea is to send IP/MPLS packets as if operating a WS network, but change the switching operation when the provided wavelength paths are temporarily overloaded. In that case, the excess data is sent as “overspill”, and the network is effectively operated in packet switching mode. By instantly switching between the two operating modes, the processing overhead on the IP/MPLS routers is maintained smaller than if functioning in a pure packet mode since most traffic is sent through a direct wavelength path, while simultaneously ORION provides maximal resource sharing, since if needed all wavelengths are accessible. ORION uses the idle periods of the established lightpaths to insert packets towards either their next hop or directly to their destinations, if this matches the lightpath destination.

Figure 1 illustrates how this principle would work in practice. Suppose that path A – C is a direct wavelength path on λ_1 , as well as A – B and B – C on λ_0 . All paths have a capacity of, say, 10 Gb/s. Under normal conditions all traffic from node A destined to node C will pass node B transparently. Now assume node A has 12 Gb/s of traffic for B. In a usual WS network this would result in loss, as there is only 10 Gb/s available, or it would require the establishment of an additional wavelength. In ORION, however, the remaining 2 Gb/s can be served by sending the data in overspill mode over wavelength λ_1 in idle periods (when there is unused capacity). These packets are marked as “overspill packets” and are treated as in a packet-switched network. Therefore, node B extracts these packets from the passing through flow and directs them to the electronic IP router of node B. Figure 2 shows an ORION-enabled node architecture, suitable to support overspill routing. Fast 1×2 optical switches are used at the output of the OXC to extract overspill packets. A single (or more) tunable laser transmitter (ORION Tx) with a simple power combiner (fiber coupler) is used to insert overspill packets to flows. A void detection module, employing a fiber delay line (FDL) and two power detection units, scouts for possible idle periods in the outgoing lightpaths and signals the IP router to insert an overspill packet.

In principle, for each wavelength channel passing through the node, a separate receiver interface and a separate tunable transmitter is needed. However, this




Fig. 1 ORION switching paradigm: In the case that λ_0 wavelength channel of A-to-B connection overloads, packet from A to B are sent in “overspill” mode over λ_1 wavelength channel, although λ_1 is a direct wavelength path from A to C

is not cost-effective, since it would require as many ORION transceivers (receivers and transmitters) as in a typical point-to-point packet-switched node, at no or limited advantage. Therefore in the node architecture shown in Fig. 2, wavelengths of the passing through lightpaths are grouped together to upload packets through a common concentrator interface. This comes with the disadvantage of possible contention, in the case that more packets than the available interfaces arrive at the same time. Deciding on how many and which wavelengths should be grouped together is a design parameter that depends on how many overspill packets are simultaneously expected. The same contention may also occur with the available tunable transmitters and the number of packets that are expected to be simultaneously inserted.

We will call *constrained overspill routing* the case when the number of lightpaths passing through a node (up to eight in Fig. 2) is higher than the number of ORION receivers and transmitters (one in Fig. 2) and *unconstrained overspill routing* when it is equal. The case of *constrained overspill routing* is separately studied in Section “The case of constrained overspill routing”.

ORION routing policies

Within the ORION architecture, several options are possible on how overspill traffic is handled. These options, termed routing policies, dictate if, how, and when packets should leave overspill mode, and go back to the wavelength switching regime. The different policies in some cases lead to different requirements in the control as well as in the architecture’s hardware. In this article we evaluate two such routing policies:

- *Once ORION always ORION*: In this routing policy a packet that has entered overspill mode remains there. Therefore, the overspill packet is treated at every node by the electronic router and thus hop-by-hop routing is performed. In *once ORION always ORION* policy, packets are stored in overspill

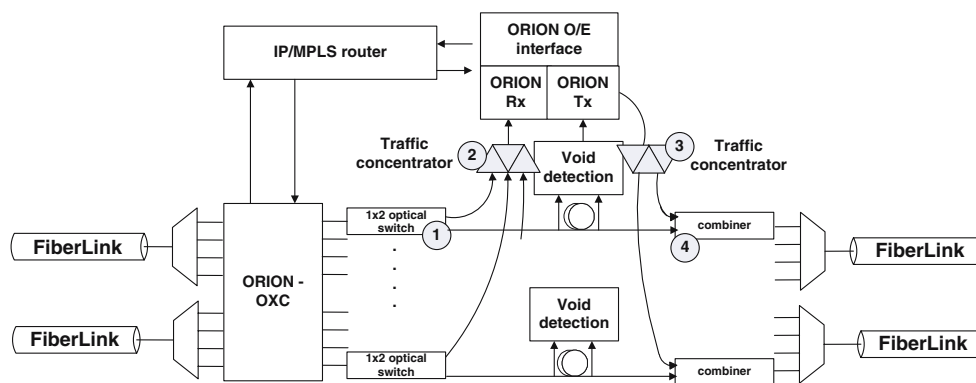


Fig. 2 ORION node architecture employing (1) a fast 1×2 optical switch for overspill packet extraction, (2) a concentrator for handling overspill packet to the electronic router, (3) a traffic

concentrator for inserting overspill packets to lightpaths and (4) a fiber combiner for overspill packet insertion

buffers and, evidently, have to look for an idle period at every hop.

Operationally, *once ORION always ORION* routing is the simplest policy, but increases the electronic processing at each node and also increases packet delay. This can eventually lead to high packet drop ratios due to buffer overflows, and thus result in a low overspill throughput.

- **Lightpath re-entry:** As an overspill packet progresses towards its end-destination in the network, it can occur that it reaches an intermediate node that has a lightpath ending at the same destination node. In that case, under the *lightpath re-entry* policy, the node inserts the overspill packet into this path. The packet is not marked as an overspill packet and is transported transparently until its destination. The advantage is that subsequent core nodes no longer see the overspill packet, which results in less processing overhead and smaller delay.

However, *lightpath re-entry* is a more complex policy as it requires state information to be maintained and additional operations to be performed for forwarding the overspill packets. For example, to facilitate matching between overspill packets and lightpath destinations, the edge/core routers have to employ multiple FIFOs (one per destination—Virtual Output Queue). Furthermore, with the *lightpath re-entry* policy packets may arrive out-of-order at the destination node, which is not the case in “*once ORION always ORION*.” A solution to this problem can be the restriction of packets belonging to the same flow to enter the same lightpath, so as to propagate through the same path. This can be accomplished by mapping entire flows (or bundles of flows) in and out of lightpath re-entry mode and not decide on a per packet basis.

Beside the aforementioned routing policies, other more complex ones also exist. For example, a combination of the above basic routing schemes that would enable the extraction of an overspill packet beyond the next hop but also before the end-destination of the lightpath, (and therefore termed as *sub-lightpath entry*), could increase network efficiency. This would enable the network to utilize all available wavelength paths at each node to further reduce overspill packet processing and thus improve overspill throughput and delay. However, in this study, we have limited ourselves to the two aforementioned policies, primarily due to the simplicity of their implementation.

ORION network simulator

In order to evaluate the ORION architecture, we have developed a discrete-event network simulator based on the ns-2 platform. Basic feature of the simulator was the design of a suitable core and edge architecture, capable of supporting overspill routing. A specific packet/header format was used for marking overspill and OCS packets, while appropriate load-varying sources were developed to simulate temporal network overloading scenarios.

Overspill packet format

The feature that needs to be mentioned at the physical level of communication is the presence of an ORION label to distinguish overspill packets. Overspill packets were encoded using a custom packet format shown in Fig. 3. The ORION label, Field 1, specifies whether the packet is in overspill mode or not, while header field 2 (GMPLS label) defines the lightpath (LSP) that the packet is loaded on. In the actual case of an optically




Fig. 3 ORION packet form

labeled switching network, both the GMPLS and the ORION labels can be encoded in the optical domain [16]. A core node extracts and processes only these two optical labels, while payload remains in the optical domain.

ORION edge router architecture

Figure 4 shows the edge router (ER) architecture, implemented in the ORION simulator. Typical blocks for IP/GMPLS processing have been omitted for simplicity. The ER handles requests from flows and establishes lightpath connections. In the current implementation, the RSVP protocol is used for setting up the lightpaths. The ER maintains a table of the active LSPs and their associated network paths. Thus, it is aware of the links and the intermediate nodes that are being used by each flow. Further, the ER employs a “void detection” module that *listens* to a FDL in order to detect idle periods in the LSPs. This FDL should be at least the maximum packet size.

The data of each flow is forwarded to a separate per lightpath Random Early Detection (RED) queue [4]. The RED queue detects incipient congestion by computing the estimated queue size \bar{q} ; if \bar{q} exceeds a pre-defined threshold q_{min} it drops (i.e., turns to overspill mode) incoming packets with probability P_a . The average queue size is calculated for each packet arrival by $\bar{q}_{new} = (1-w_q) \cdot \bar{q}_{previous} + w_q \cdot q$, where $\bar{q}_{previous}$ and \bar{q}_{new} are the previous and newly-computed average queue size, q is the current queue size and w_q is a weighting factor. As \bar{q} varies from q_{min} to q_{max} , the individual packet dropping probability, P_b , varies linearly from 0 to p_{max} :

$P_b = p_{max} \cdot \left(\frac{\bar{q} - q_{min}}{q_{max} - q_{min}} \right)$, where p_{max} is the maximum dropping probability. The final packet dropping probability, P_a , increases slowly with the number of packet arrivals since the last drop:

$$P_a = \begin{cases} \frac{P_b}{1-C \cdot P_b} & \text{if } q_{min} \leq \bar{q} \leq q_{max}, \\ 1 & \text{if } \bar{q} > q_{max} \end{cases} \quad (1)$$

where C is the number of arrivals since the last drop.

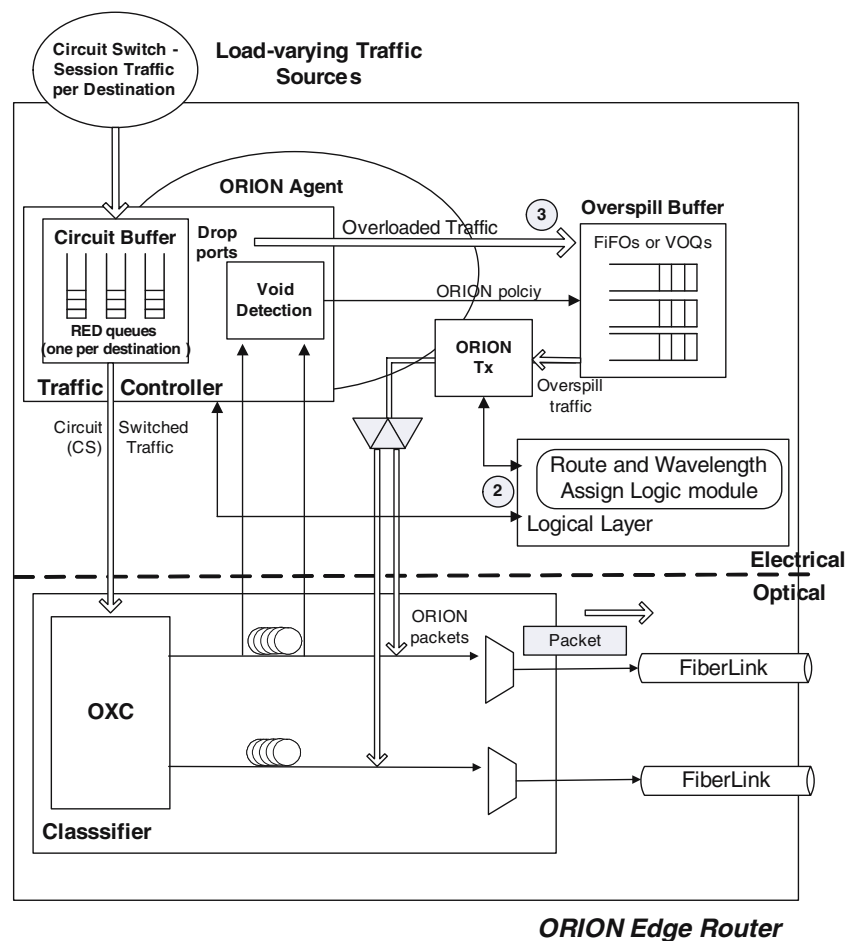
In the developed ER architecture, the drop ports of all the RED queues are connected with the overspill mechanism, which collects the “salvaged” packets and

stores them in dedicated “overspill buffers” (see Fig. 4). The numbers in Fig. 4 refer to the packet fields, shown in Fig. 3, that are being processed in each particular action. Thus, the Route and Wavelength Assign (RWA) logic module processes the GMPLS label (Field 2) to assign the packets to the proper outgoing wavelength. Similarly, the ORION agent processes the IP/MPLS label (Field 3) of the “overspill” packets in order to identify their destination IP address and store them in the proper queue in the overspill buffer. The implementation of this buffer depends on the routing policy. For the *once ORION always ORION* policy, a FIFO queue is used per outgoing link, while for *lightpath re-entry* policy, the buffer is implemented with multiple FIFOs (a separate virtual queue per destination).

In general, both the lightpath and overspill mode can have a buffer for resolving temporary overloads. This leads to the following possible buffering schemes:

- *No-Buffering (NB)*: When a packet of a connection finds its lightpath occupied (temporal overload), it is immediately put in overspill mode, provided that ORION is enabled. In the case of No-Buffering (NB), there is no overspill buffer and thus the overspill packet has only one immediate attempt to find an idle wavelength. If this is unsuccessful, or ORION is disabled, the packet is dropped.
- *Circuit-buffering (CB)*: In this case, each lightpath has an associated RED queue, termed as Circuit Buffer (CB). When a packet of a connection is dropped from the RED queue, due to temporal overloading (and ORION is enabled) it tries to go immediately into overspill mode. If this is not immediately feasible (no overspill buffer), or ORION is not supported, the packet is discarded.
- *Overspill-Buffering (OB)*: The dual of the previous regime: a packet immediately tries to enter its lightpath. If unsuccessful, the packet becomes overspill and is stored in an overspill buffer (OB), while the node starts looking into the passing through LSPs, to find an idle gap to insert it. As described earlier, for *once ORION always ORION* policy, OB is implemented as a single FIFO per outgoing link, while for *lightpath re-entry*, VOQs are used. If the OB buffers overflow the most recent packet is discarded.
- *Circuit-Buffering and Overspill-Buffering (CBOB)*: This is the combination of CB and OB, where RED queues are used for LSP data and FIFO (or VOQ queues, depending on the policy) are used for overspill packets. Packets dropped from the lightpath RED queues are forwarded to the overspill buffer.

Fig. 4 ORION edge router architecture. The numbers refer to the ORION packet fields that are being processed



Note that, in the ER, and in the case of *lightpath re-entry* policy, overspill packets are looking for possible idle periods in their own, initial LSP. In contrast, in core routers (CR), overspill packets are looking for an idle period in any LSP passing through the node, heading for the same destination. After this general overview, we provide a short description of the specific modules present in the ER architecture. The design is based on work presented in [17].

Route and Wavelength Assign (RWA) logic module: This module calculates the routing paths using an RWA algorithm (shortest path in our simulation) and is responsible for establishing the lightpaths. It also maintains a table with the network virtual topology.

ORION Agent: This includes the *Traffic Controller* module that handles circuit switched requests by communicating with the *RWA and Logic* module. The traffic controller establishes LSPs and forwards the traffic stored in the CB buffer.

When a time-gap in a LSP is detected (*Void Detection* module), ORION agent inserts an overspill packet from the OB. The decision of which packet to insert is taken by the *ORION policy enforcing point* module.

Thus, depending on the ORION policy, the overspill packet is retrieved either from the matching VOQ or from the single FIFO.

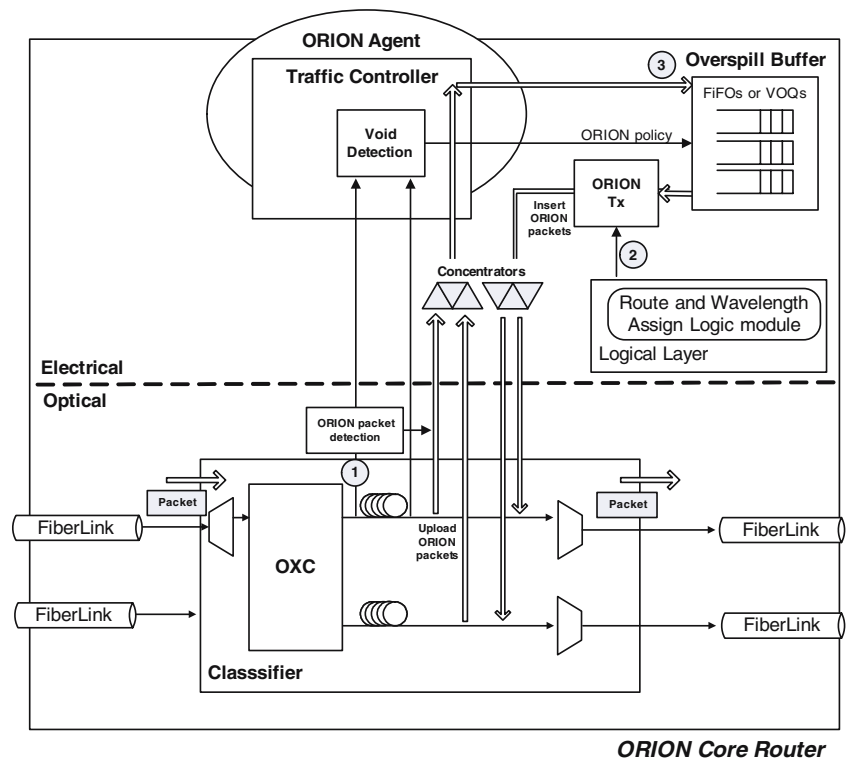
Void Detection module: The void detection module consists of two power detection points positioned at the edges of a FDL unit. The module detects the possible idle periods in the lightpaths and signals the insertion of overspill packets.

Classifier module: The *Classifier* module is responsible for receiving and forwarding packets depending on their header information. In the ORION network simulator, the *Classifier* accesses the packet header fields 1 and 2 (see Fig. 3) and performs the appropriate forwarding/routing actions. The *Classifier* module also performs all the typical functions for circuit-switched traffic that are omitted from Fig. 4 for simplicity.

ORION core router architecture

The operation of the CR is similar to that of the ER but it additionally has to accommodate the detection and extraction of overspill packets. Figure 5 shows the architecture of the CR that was modeled. The CR

Fig. 5 ORION core router architecture. The numbers refer to the ORION packet fields that are being processed



employs a concentrator module that groups together all passing lightpaths to upload overspill packets via a set of O/E interfaces to the electronic domain. The numbers in Fig. 5 refer to the packet fields that are being processed. Thus, the overspill packet detection module processes Field 1 of the packet header, to determine if the incoming packet is in overspill mode and if it is, it extracts and stores it in the OB buffer according to its IP/MPLS label (Field 3). Similar to the ER, the RWA module processes the GMPLS label (Field 2) to assign the packet to the proper outgoing wavelength.

For the execution of the routing/forwarding ORION policy, the CR maintains a table with the active light-paths and their associated link-paths, similar to that maintained at the ER. Control actions are identical to those of the ER, except for the handling of locally inserted traffic. Thus, all possible buffering architectures of NB and OB are supported.

Load-varying sources

The load-varying sources were developed for testing purposes. Each source-destination pair was modeled with a separate traffic source. Each source generates packets according to a Poisson process, with packet sizes drawn from a typical Internet mix packet size distribution (40, 520, 1500 Bytes of 50%, 37.5%, and 12.5% occurrence, respectively). The traffic load of a source, is

defined in the experiments as the ratio of $\overline{ON}/(\overline{OFF})$, which can exceed 1.0 to simulate path overloading. In that case, the excess traffic is switched to the overspill mode and thus data is dropped from the CB to the overspill buffer (OB). Every source selects its load randomly—according to a uniform distribution with a given average value (*X*-axis in the performance graphs) and a standard variation equal to 0.5. This randomly chosen load is maintained throughout a simulation cycle.

Network level performance evaluation—unconstrained overspill routing

The evaluation of the ORION architecture was performed on the NSFnet network topology, shown in Fig. 6a. All links were assumed to be bidirectional, with a 1 Gbps capacity per wavelength, to reduce simulation times. With these assumptions, we have first measured the minimum number of wavelengths required to support all source-destination pairs with a 1% packet loss ratio, in the case of a pure wavelength switched (WS) network and in the case of a point-to-point, packet-switched network (P2P). Figure 6b shows the corresponding results. It can be seen that for an average load of 1.4, the required wavelengths are 60 and 45 for the case of a WS and P2P network. In order to assess the overspill mechanism for absorbing temporal

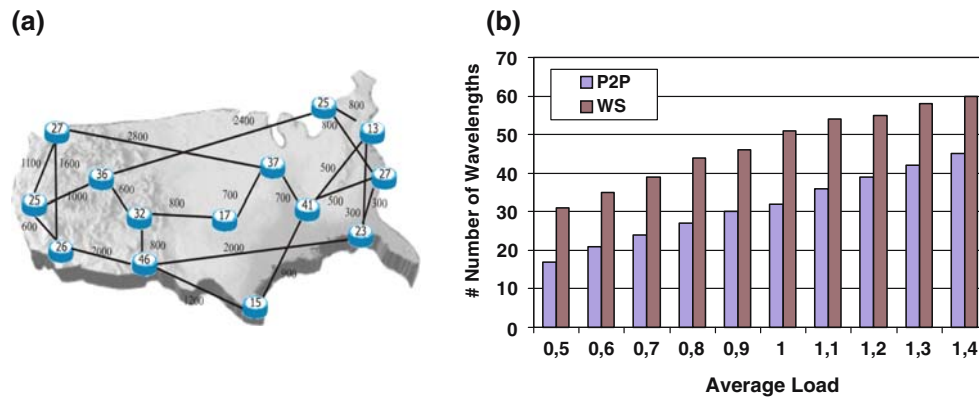
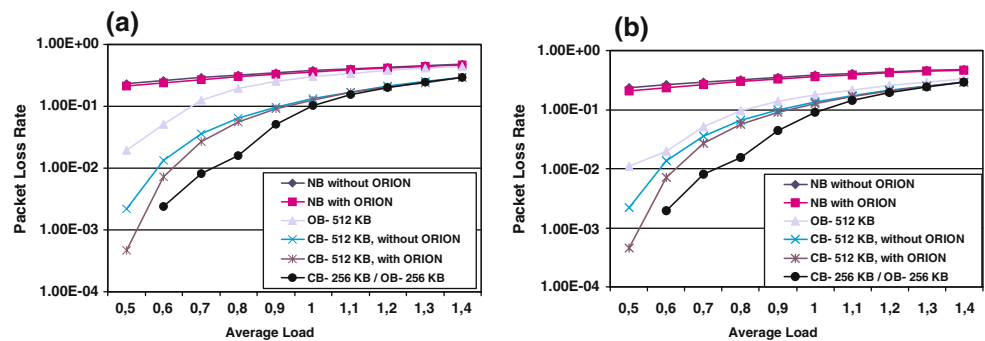


Fig. 6 (a) 14-node NSFnet backbone network topology with the number of crossing lightpaths of every node (the shown distances are in km). (b) Number of wavelengths per average load needed

to support the communication of all source-destination pairs with shortest path routing in the case of wavelength switching (WS) and point-to-point, packet switching, (P2P) for 0.01 loss ratio

Fig. 7 Packet loss ratio for (a) “once ORION always ORION” and (b) “lightpath re-entry” routing policy for the various supported buffering schemes versus average source load. The number besides the name of each buffering scheme denotes the size of each in KBytes



traffic imbalances, we used in our experimental analysis the minimum case of only 31 wavelengths. It is then feasible to compare the gain in using fewer wavelengths in the network as well as the yielding loss ratios in both cases.

In the simulations carried out, a separate RED queue was maintained for each source-destination pair and their total size at an ER is depicted as “CB” buffer. The chosen RED queue parameters are: $w_q = 0.001$, $p_{max} = 0.2$ and $q_{min} = 0.7 \cdot q_{max}$ where q_{max} is the RED queue size. Further, in all experiments carried out, the opto-electronic conversion time was set equal to 0.01 ms, while the FDL length in the void detection module was set equal to 0.04 ms [13]. The addition of this FDL length inside all core nodes resulted in an average propagation delay through the network of 13.1 ms.

The key performance metrics measured include packet loss ratios for the various buffering schemes, overspill packet throughput, and end-to-end delay. All these were measured versus the given average source load for both ORION policies. In this section, it is assumed that all wavelengths can simultaneously carry and terminate (*unconstrained case*) overspill packets and all

have an individual receiver/transmitter interface so that no overspill packet contention can occur. The case of constrained overspill routing, where the number of O/E interfaces is smaller than the number of wavelengths, is studied in Section “The case of constrained overspill routing”.

Figure 7a, b shows the average packet loss ratio of the *once ORION always ORION* and the *lightpath re-entry* policies for the various buffering schemes. It can be seen that the cases of NB, either with or without ORION support, exhibit the highest loss ratios. This was expected since when no buffering is available, packets are dropped immediately after the temporary overload of the WS path. As expected, the loss curves for all cases that do not use ORION (*NB without ORION* and *CB-512 KB without ORION*) are identical in both Fig. 7a, b, since these schemes do not depend on the routing policy (actually there are no overspill packets).

With respect to the Overspill-Buffering scheme with 512 KB buffer size (OB512 KB), the *lightpath re-entry* policy outperforms *once ORION always ORION*. In this scheme, a large part of incoming traffic (20%–40%) switches to overspill mode, since there are no CB buffers.

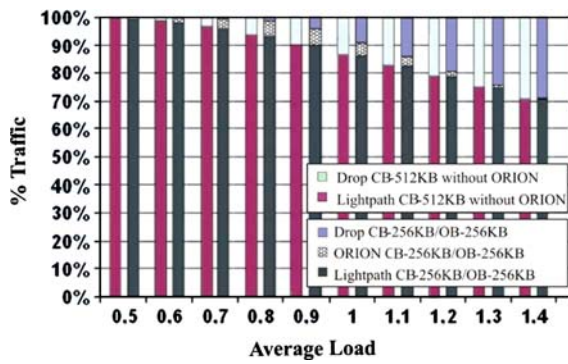


Fig. 8 Traffic statistics for lightpath re-entry policy. (a) CB = 512 KB without ORION and (b) CB = 256 KB, OB = 256 KB

As a result, OB buffers become congested and thus hop-by-hop routing (*once ORION always ORION policy*) tends to drop more packets. On the other hand, overspill packets that happen to re-enter a lightpath are treated more favorable in the sense that they are transparently forwarded to their end-destination.

Regarding the two last cases, *CB512 KB with ORION* and *CB256 KB/OB256 KB*, both routing policies exhibit similar performance characteristics. This is because overspill traffic is a small percentage of the total traffic and thus the difference in the loss probabilities between the two policies is small.

With respect to the performance of the different buffering schemes, *CB256 KB/OB256 KB (CBOB buffering)* significantly outperforms the other for the same total buffer size. The comparison of the *CBOB* scheme with the *CB512 KB without ORION* scheme reveals the positive effect of overspill routing. It is evident from Fig. 7, that using the same 512 KB buffer, but dividing it into two equal parts for CB and OB, results in a reduction of the loss ratio when the overspill mechanism is efficiently utilized (for average traffic loads smaller than 1.0). Figure 8 shows the corresponding traffic statistics for the *lightpath re-entry* policy, for the cases of *CB512 KB without ORION* (left columns) and *CB256 KB/OB256 KB* (right columns) schemes. In particular, Figure 8 illustrates the percentage of packets transported as circuit traffic or as overspill packets (either dropped or successfully transported). It can be seen that the overspill mechanism is more effective for moderate traffic loads (between 0.8 and 1.0), while, when the traffic load exceeds 1.2 both loss ratios converge (drop percentages for both CB-512 KB without ORION and CB-256 KB/OB-256 is equal to ~29%). This is reasonable, since as the average traffic load increases, all lightpaths become saturated, void filling cannot be performed and thus the beneficial effect of overspill routing eventually disappears.

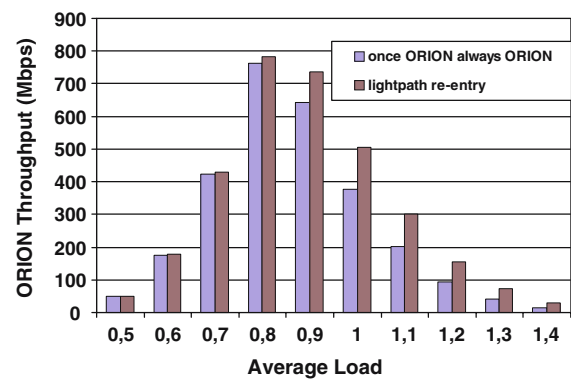


Fig. 9 Throughput of overspill traffic for the two policies, for the *CB256 KB/OB256 KB* scheme

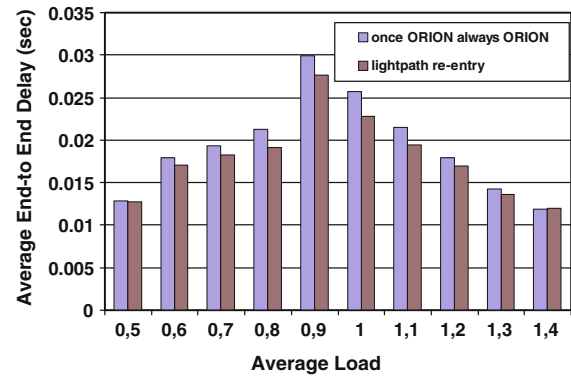


Fig. 10 Average delay of overspill packets for the two policies, for the *CB256 KB/OB256 KB* scheme

To further compare the two ORION routing policies, we have measured the throughput and average packet delay of overspill packets obtained for the best performing *CB256 KB/OB256 KB* scheme. Figures 9 and 10 show the corresponding results. It can be seen that for traffic loads higher than 0.7, the *lightpath re-entry* policy outperforms *once ORION always ORION* both in terms of throughput and delay, since a fraction of overspill traffic reaches its destination immediately and fewer packets are extracted at intermediate nodes. This insertion/extraction process of overspill packets especially in the case of *once ORION always ORION* policy increases the dropping probability and the end-to-end delay. Figure 11 shows the number of overspill packets that are extracted at intermediate nodes per second for both policies. This is the additional traffic seen by the electronic part of the CR due to the overspill mechanism. The difference in the number of extracted packets reveals the actual number of overspill traffic that was “re-entered” in the lightpaths.

Note that at low traffic loads, packets that try to switch into overspill mode constitute only a small fraction of the total traffic, explaining the low throughput and delay




Fig. 11 Number of overspill packets uploaded at intermediate nodes

levels, and the small difference between the two policies observed in Figs. 9 and 10. As traffic increases, throughput increases as well, reaching its maximum value for traffic loads close to 0.8 – 0.9. Above this average value, throughput steadily decreases for both policies (but their difference in terms of loss rate increases) and finally the performance of both policies converge, since the network saturates and thus there is no bandwidth to be re-used by the overspill mechanism.

With respect to packet delay (see Fig. 10), its maximum value corresponds to the same load range of 0.8 – 0.9. Note that the packet delay recorded in the figure refers to the delay of packets that reach their destination, while the delay of the lost packets (which is infinite) is not taken into account. Beyond this load range, overspill packets start getting dropped, and thus gradually are given limited chances to be successfully transported in overspill mode, independently of the policy enforced. However, delay for traffic loads higher than 0.9 starts decreasing, and this is due to the fact that only overspill packets that follow short paths (usually one-hop) are serviced. Almost all the other packets are dropped.

Finally, we have also investigated the effect of the total buffer size. Figure 12 shows the corresponding loss ratios of the *lightpath re-entry* policy for buffers of 512 and 1024 KB of total size. The improvement in the loss ratio differs for the various supported buffering schemes, and higher gains are observed for the combined CBOB scheme. The important finding from Fig. 12 is that the performance differences between the buffering schemes remain, regardless the size of the buffer. For example, both the CBOB schemes perform better than the *CB-512 KB with ORION* scheme, which in turns performs (slightly) better than the *CB-1024 KB without ORION* scheme. Of course, if buffer size is increased, the latter scheme will perform better, but always worse than the same scheme with the support of overspill traffic. This clearly shows that both types of buffers (CB and OB) are

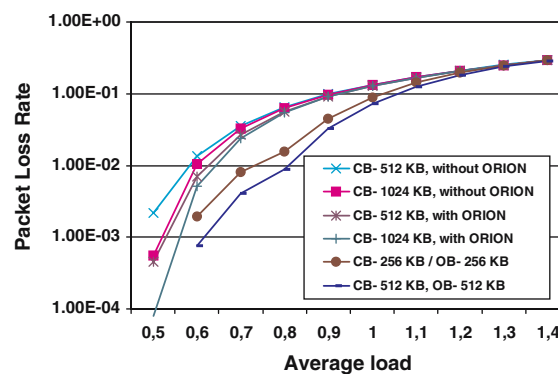


Fig. 12 Packet loss rate for lightpath re-entry policy for various buffering schemes with buffers of 512 KB and 1024 KB total size

necessary to obtain the best performance improvements from overspill routing.

Concluding, the preceding results indicate that the ORION architecture can indeed absorb temporal traffic overloads and ensure a low packet loss ratio. The proposed buffering schemes perform differently, and the best performance is exhibited when adequate buffering is provided for both overspill and circuit traffic (CBOB). Finally, with respect to the two examined overspill routing policies, lightpath re-entry policy yields the best performance with the tradeoff that it requires additional control overhead.

The case of constrained overspill routing

In this section, we consider the case of *constrained* overspill routing, where the number of receivers and tunable laser transmitters at a node limits the number of overspill packets that can be simultaneously terminated or inserted. This is done via two concentrator interfaces (as shown in Fig. 1) that groom overspill packets. Since it is not known in advance which wavelengths will temporally overload, it is reasonable to allow all wavelengths to carry overspill traffic and constrain the number of packets that can be simultaneously extracted (or inserted), rather than constrain the number of wavelengths that are allowed to carry overspill traffic. To this end, in the case of constrained overspill routing, we look into the size of the two concentrators (one for the extraction and one for the insertion of overspill packets), and analyze their effect in network performance. The size of the concentrator is defined as the number of receivers or tunable transmitters commissioned to receive or transmit overspill packets.

In order to assess the case of limiting overspill routing, we experimented with bursty traffic, since it is expected that burstiness will have a significant effect on traffic

Fig. 13 (a) Percentage of simultaneously uploaded and (b) inserted overspill packets for the lightpath re-entry routing

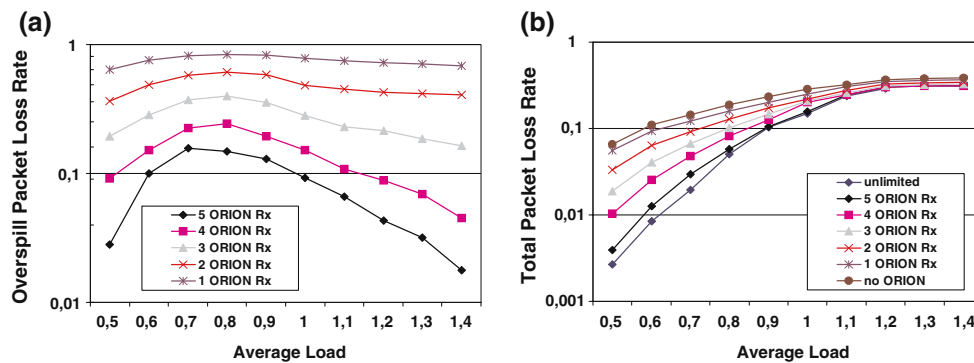
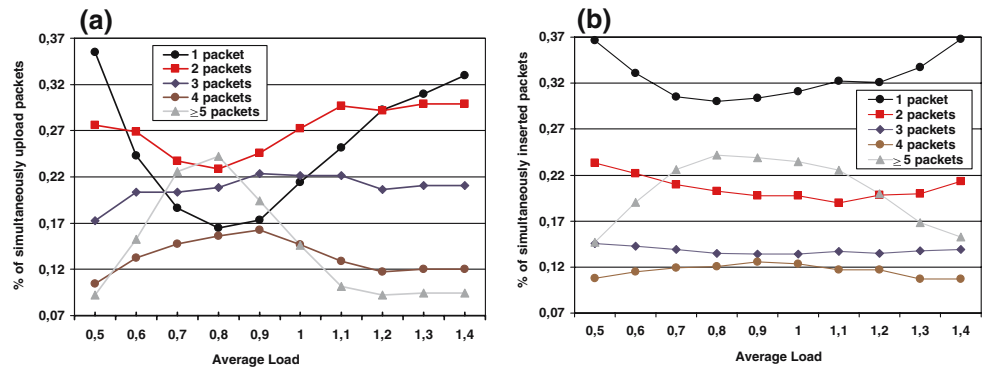


Fig. 14 (a) Overspill packet loss rate due to contention, (b) overall network loss rate for the case of *constrained* overspill traffic reception – *unconstrained* overspill traffic transmission (*constrained* Rx – *unconstrained* Tx case)

load. Thus, the traffic generating sources were modified so as to generate traffic with Pareto inter-arrival times, while packets' sizes were again drawn from the same Internet mix size distribution.

Figures 13a, b shows the percentage of 1,2,3,4 and ≥ 5 simultaneously extracted and inserted overspill packets for the lightpath re-entry policy and for Pareto traffic with a shape parameter $a=1.2$. From Fig. 13a, it can be seen that for average traffic loads in the range of 0.7–0.8, where the overspill mechanism is mostly utilized (more packets are uploaded in the intermediate nodes), the percentage of 4 and ≥ 5 simultaneously uploaded packets increases, while the percentage of 1 and 2 simultaneously uploaded packets decreases. The curves for 1 and ≥ 5 -packet uploads change rapidly (decrease/increase, respectively), while the curve for 3-packet uploads remains almost constant. In general, only one-third of the overspill traffic corresponds to a single packet upload, a percentage that drops to one-fifth, for traffic loads in the range of 0.7–0.8. It is therefore clear that the use of a concentrator at the receiver side will have an imminent effect to loss ratio.

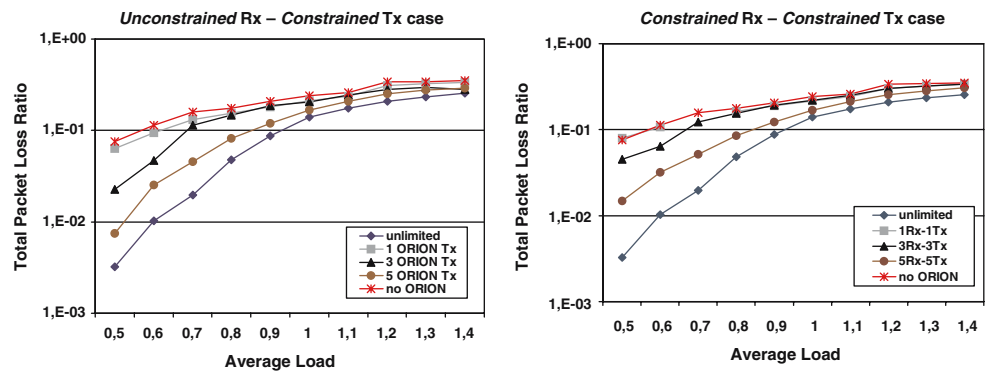
The results differ significantly for the simultaneous packet insertions. From Fig. 13b it can be seen that the percentage of single packet insertions, does not de-

crease as much as in the single packet extraction, while the percentage of ≥ 5 simultaneous packet insertions increases from 14% to $>23\%$ for a larger workload range of 0.7–1.1. The rest of the curves (2-, 3-, and 4- simultaneous packet insertions) remain almost constant. The difference between the extraction and insertion processes (Fig. 13a, b) is mainly due to the lightpath re-entry policy. In particular, a percentage of overspill packets is inserted as overspill packets (through the ORION transmitters) but extracted as circuit switching traffic (OCS). Thus, ORION transmitters are more utilized in the lightpath re-entry policy than the corresponding receivers, which means their impact on loss is higher.

Based on the above analysis, we have first measured how packet loss ratio varies when only one concentrator is employed either at the receiver or the transmitter side (*constrained* Rx – *unconstrained* Tx and *unconstrained* Rx – *constrained* Tx cases) and then how loss varies again when both the ORION receivers and transmitters are constrained (*constrained* Rx – *constrained* Tx case).

Figure 14 shows the first case, where only the ORION receivers are constrained to 1-to-5 interfaces. In particular, Fig. 14a shows the packet loss ratio due to contention in the concentrator that groups overspill traffic from all

Fig. 15 Overall network loss ratio for the case of (a) *unconstrained* overspill traffic reception – *constrained* overspill traffic transmission (*unconstrained Rx – constrained Tx*) and (b) *constrained* overspill traffic reception – *constrained* overspill traffic transmission (*constrained Rx – constrained Tx*)



passing through lightpath and terminates it to either 1, 2, 3, 4, or 5 ORION Rx interfaces. The loss ratio depicted in Fig. 14a corresponds only to the contending traffic, and which decreases at high loads due to network saturation.

Figure 14b shows the overall packet loss ratio in the network, including the packets, dropped due to buffer overflows. For comparison, in Fig. 14b we have also included the corresponding losses for the unconstrained case as well as for the case that no interfaces (no ORION) are employed at all. The important finding from Fig. 14b is that by employing one ORION receiver, performance approximates the case of a pure circuit switch network (no ORION), while five interfaces highly approximate the *unconstrained* ORION case. In the simulated network topology, using five receivers per node is a reduction of up to 82% of the receivers needed in the unconstrained case.

In the sequence, we have measured how loss varies, when the number of packet insertions is constrained with a concentrator at the transmitter side. Figure 15a shows the overall packet loss ratio, in the case that a concentrator limits the number of packets that can be simultaneously inserted to 1, 3, or 5 packets, while there is no limitation at the receiver side (*unconstrained Rx – constrained Tx* case). In this case, no packets are lost due to contention but only due to buffer overflow. From Fig. 15a, it can be seen that one transmitter resembles the case of a pure circuit switching network (no ORION), while the case of five Tx approximates again the *unconstrained* case, but not as much as in the case of five receivers in the *constrained Rx – unconstrained Tx* case (see Fig. 14b). It is therefore clear that the number of ORION transmitters is more important than the number of receivers.

Finally, Fig. 15b, shows the loss performance of the *constrained Rx – unconstrained Tx* case. It can be seen that performance has deteriorated for all curves. Nevertheless, the gain is still significant and particularly five (5) ORION transceivers (Tx and Rx) yield a loss ratio of

1.4% and 12% at a workload of 0.5 and 0.9, respectively, whereas the corresponding losses in the pure circuit-switching case (no ORION) are 7.6% and 20%.

It must be noted here again that five transceivers per node results in 70 transceivers in the NSF network, as opposed to 390 (see Fig. 6a) needed in the *unconstrained* ORION case. Thus, the gain in cost and size of an ORION-enabled node is significant. This conclusion is valid for the specific network topology under study and the wavelength-routed paths designed. For other topologies, results may differ but in any case they prove the economic viability of the architecture. In particular, it is clear that with only a small subset of ORION transceivers, a large part of the overloaded traffic can be absorbed without the deployment of new wavelengths.

Conclusions

In this article, a detailed performance evaluation of the ORION hybrid switching architecture was presented using a network simulator platform developed for this purpose. We have evaluated the case of *unconstrained* and *constrained* overspill routing, depending on whether all lightpaths are allowed to simultaneously carry and terminate overspill traffic or overspill is restricted to a sub-set of wavelengths via traffic concentrators. Further, we have assessed various aspects of the ORION architecture including two basic routing/forwarding policies and various buffering schemes. It was shown that the lightpath re-entry policy, combined with the Circuit-Buffering Overspill-Buffering (CBOB) scheme is superior to all other schemes examined in terms of loss ratio, throughput and packet delay. The complete network study revealed that ORION can absorb temporal traffic overloads, as intended, provided sufficient buffering is present. Moreover, we have shown that we can limit the number of transceivers in the architecture to a small

subset, while still maintaining a significant gain in performance. Limiting the number of ORION receivers and transmitters results in a reduction of the size and the cost of the node and thus improves ORION economic viability.

Acknowledgements This work has been supported by European Commission through projects IST-LASAGNE, IST-NOBEL, NoE e-Photon/ONe+, the Greek General Secretariat for Research and Technology–GSRT via a PhD scholarship (PENED), and the Flemish government through the IWT-GBOU project “Optical Networking and Node Architectures”, and a PhD grant to E. Van Breusegem.

References

- [1] Balasubramanian, S., et al. (2004). Medium access control protocols for light trail and light bus networks. *Proceedings of 8th IFIP ONDM*, Ghent, Belgium, Feb. 2004.
- [2] Bjornstad, S., et al. (2003). A highly efficient optical packet switching node design supporting guaranteed service. *Proceedings of ECOC*, Rimini, Italy, Sept. 2003, pp. 110–111.
- [3] Christodouloupoulos, K., et al. (2006). Performance Evaluation of Overspill Routing in Optical Networks. *Proceedings of ICC 2006 Istanbul*, Turkey, June 2006.
- [4] Floyd, S., et al. (1993) Random early detection gateways for congestion avoidance. *IEEE/ACM Transaction on. Networking*, 1(4), pp. 397–413.
- [5] Gumaste, A., et al. (2004). Light-trails: an optical solution for IP transport. *OSA Journal of Optical Networks*, 3(5), pp. 261–281.
- [6] Gumaste, A., et al. (2005). Next-generation optical storage area networks: The light-trails approach. *IEEE Communications Magazine*, 43(3) pp. 72–79.
- [7] Le Vu, H., et al. (2005). Scalable performance evaluation of a hybrid optical switch. *IEEE/OSA Journal of Lightwave Technology* 23(10), 2961–2973.
- [8] Lee, G. M., et al. (2003). Performance evaluation of an optical hybrid switching system. *Proceedings of IEEE GLOBECOM '03* San Francisco, USA, Dec. 2003, Vol. 5, pp. 2508–2512.
- [9] Lee, Hyeong-II, et al. (2005). Hybrid optical transport Network (HOTNET): An optical network with hybrid switching technologies for integrated services. *IEICE Transactions on Communications*, E88-B, (10), pp. 3874–3885.
- [10] Qiao, C., et al. (1999). Optical burst switching (OBS)— a new paradigm for an optical Internet. *Journal of High Speed Networks (JHSN)* 8(1), 69–84.
- [11] Ramaswami, R., & Sivarajan, K (1998). *Optical networks: a practical perspective*. San Francisco: Morgan Kaufmann.
- [12] Renaud, M., et al. (1997). Network and system concepts for optical packet switching. *IEEE Communication Magazine* 35(4), 96–102.
- [13] Van Breusegem, E., et al. (2005). A broad view on overspill routing in optical networks: a real synthesis of packet and circuit switching? *Elsevier Journal of Optical Switching and Networking*, 1(1) 51–64.
- [14] Van Breusegem, E., et al. (2006). Overspill routing in optical networks: a true hybrid optical network design. *IEEE Journal of Selected Areas in Communications*, 24(4), pp. 13–25.
- [15] Vlachos K. et al. (2003a). Ultrafast time-domain technology and its application in all-optical signal processing. *IEEE/OSA Journal of Lightwave Technology* 21(9), 1857–1868.
- [16] Vlachos, K., et al. (2003b). An Optical FSK/IM coding technique for the implementation of a label controlled, AWG-based Optical Packet Router. *IEEE/OSA Journal of Lightwave Technology*, 21(11), pp. 2617–2628.
- [17] Wen, B., et al. (2001). Optical Wavelength Division Multiplexing (WDM) Network Simulator (OWNs): *Architecture and performance Studies*, *SPIE Optical Networks Magazine*, 2(5), pp. 16–26.
- [18] Xin, C., et al. (2003) A hybrid optical switching approach. *Proceedings of IEEE GLOBECOM '03* San Francisco, USA, Dec. 2003, Vol. 7, pp. 3808–3812.



Prof. Kyriakos Vlachos is a faculty member at Computer Engineering and Informatics Dept. of University of Patras, Greece. He received his Dipl.-Ing. and PhD in electrical and computer engineering from the National University of Athens (NTUA), Greece, in 1998 and 2001 respectively. From 1997 to 2001 he was a

senior research associate in the Photonics Communications Research Laboratory and in April 2001 he joined Bell Laboratories, Lucent Technologies, working on behalf of the Applied Photonics Group. During 2003, he joined the National Regulation Authority of Telecommunication and Postal Service of Greece (EETT), where he served as a scientific advisor. His research interests include optical networks, optical burst switching and WDM transmission systems. Prof. Vlachos is the (co) author of more than 60 journal and conference publications and holds five patents.



Erik Van Breusegem was a research scientist in the field of Broadband Communication Networks, specialized in optical networks, at Ghent University (Belgium). Erik's academic work includes over 50 scientific publications, including leading journals in his field. In the context of his Ph.D thesis (defended in April 2006), he was awarded a full scholarship by the IWT-Vlaanderen from 2001 to 2005. He also participated in several

European research projects, including IST-STOLAS and IST-NOBEL, focusing on the development of next-generation telecommunication networks. He currently holds a position as a senior associate in the Brussels office of The Boston Consulting Group.



Kostas Christodoulopoulos received the Diploma of electrical and computer engineering from National Technical University of Athens (NTUA), Athens, Greece, with specialization in telecommunications, in 2002 and the M.Sc. degree in advanced computing from Imperial College London, London, U.K., 2004. He is currently working toward the Ph.D. degree at the Department of Computer Engineering and Informatics of the University of Patras, Patras, Greece. His research interests are in the areas of traffic engineering, protocols, and algorithms for optical networks

He was granted a postdoctoral scholarship from the Institute for the Promotion of Innovation through Science and Technology in Flanders (IWT-Vlaanderen) in the period 2003-2004. His research deals with design and planning of communication networks. This work is focusing on optical transport networks, to support the next-generation Internet. Up till now, he has actively been involved in several IST projects (LION, OPTIMIST, DAVID, STOLAS, NOBEL and LASAGNE), in the COST-action 266 and 291 and in the ITEA/IWT TBONES project. His work has been published in more than 100 scientific publications in international conferences and journals.



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.

He is currently working towards his Ph.D. degree. His research interests are on optical burst-switching (OBS) architectures, algorithms and protocols.



Kostas Ramantas has received the diploma of computer engineering from Computer Engineering and Informatics Department (CEID) of the University of Patras, Greece, in 2006. He is currently working towards his Ph.D. degree. His research interests are on optical burst-switching (OBS) architectures, algorithms and protocols.



Prof. Piet Demeester (IEEE-SM) is full-time professor at the Ghent University-IBBT, Broadband Communication Networks group (www.ibcn.intec.ugent.be), where he is teaching courses in Communication Networks. His research interests include: multilayer networks, Quality of Service in IP-based multimedia networks, mobile networks, access networks, grid computing, distributed software, network and service management and applications. The research in this field resulted in over 500 international publications. In this research domain he was and is a member of several programme committees of international conferences, such as: OFC, ECOC, ICC, Globecom, Infocom, DRCN.

He was and is a member of several programme committees of international conferences, such as: OFC, ECOC, ICC, Globecom, Infocom, DRCN.