

On the Myths of Optical Burst Switching

Authors:

Pablo Pavon-Marino, Technical University of Cartagena (Spain). Email: pablo.pavon@upct.es

Fabio Neri, Politecnico di Torino (Italy). Email: neri@polito.it

***Abstract*—This paper discusses merits and drawbacks of the Optical Burst Switching (OBS) paradigm, which received so much attention by the academic research community. Since the topic is hot, and given the several flavors of OBS available in the literature, we first define what we mean by OBS. Then, we carefully design a reference scenario and a set of evaluation experiments that can be used to obtain quantitative answers to interesting questions in a repeatable way. Within this framework, we motivate our skepticism towards two frequent assertions about the OBS paradigm: its ability to (i) exploit the network bandwidth more efficiently than Optical Circuit Switching (OCS) alternatives, and (ii) to efficiently solve contentions without optical buffers.**

***Keywords:* Optical fiber communications, Optical Burst Switching, Optical Circuit Switching.**

I. INTRODUCTION

Since its origins more than a decade ago [1][2], the Optical Burst Switching (OBS) paradigm has attracted an enormous research interest. Most network evolution roadmaps refer to OBS as a viable technology in the long way towards the always far all-optical networks dream. Successive refinements and studies have been conducted in numerous fields that grew up related to the OBS technology. To name a few: burst assembly techniques and their effects on system performance and resulting traffic patterns, reservation protocols, or scheduling algorithms for the OBS switching node. This paper is *not* a new attempt in this line. What the authors intend is to critically revise the validity of two common claims usually accompanying OBS descriptions, and referred as unquestionable merits of this paradigm. First, given its sub-wavelength switching granularity, OBS is commonly termed to make a more efficient use of the optical transmission bandwidth than the OCS approach. Second, the ability to

solve contentions without optical buffers has been claimed to be a major advantage of the OBS paradigm, against other sub-wavelength switching paradigms (like OPS, Optical Packet Switching). Given our skepticism on these two claims, we referred to them as OBS *myths* in the title. By questioning them, we also question OBS as a candidate to substitute the current multi-layer OCS paradigm in the evolution of optical networks.

We intend to motivate our position in these questions by carefully selecting a set of benchmark experiments, and studying the cost/revenue balances of OBS and OCS networks for them. In the cost side, qualitative and quantitative considerations will be introduced in the discussion. In the revenue side, we assume that the economical profit obtained from operating a backbone network is directly dependent on the volume of the traffic demand which can be carried, maintaining a reasonable *end-to-end* traffic loss probability target.

Given the ambitious (and controversial) nature of the topic, we have put every effort to design the tests with the intention of systematically comparing the OBS and OCS alternatives. For the OBS network, we exhaustively test a large set of OBS variants, under different potential technological scenarios (represented by different optical switching fabric reconfiguration times and electronic header processing times). Of course, it was not possible to test the interplay of all the combinations of burst assemblers, schedulers and switching architectures, traffic engineering decisions, etc. The options taken in these aspects had the general guideline of finding the *upper limits* in burst loss performance that can be achieved with reasonable assumptions by the OBS networks. In other words, design configuration choices have been taken to favor the contention resolution performance (and thus the potential revenues) of the OBS optical nodes. The evaluation of OBS networks involved extensive simulations that required more than 25,000 hours of processor time. In the OCS case, the traffic matrices were used as the input to a heuristic planning algorithm to find a feasible solution to the multilayer planning problem. Section VI discusses and highlights the differences in the design of benchmark experiments with respect to previous works.

The obtained results are carefully analyzed in this paper. They suggest that the combination of (electronic) packet switching and OCS of current multilayer networks is in general able to carry the same or *more* traffic than the same networks following an OBS scheme, at a lower cost. As a

consequence, while it is naturally possible to build OBS nodes without optical buffers, our observations show that the resulting degradation of the network capacity dooms this approach both in performance and cost comparison with OCS networks.

The rest of the paper is organized as follows. Section II is devoted to clarifying what OBS paradigm is considered in this study. Section III and IV describe the OBS and OCS testing scenarios. A comparative discussion of simulation results is included in Section V. Section VI reviews some related work, motivating the main decisions taken in designing our experiments. Then, Section VII offers the authors' analysis about the merits of the OBS paradigm. Finally, Section VIII concludes the paper.

II. OUR OBS

Since OBS has been a quite fashionable research topic, several flavors and definitions of OBS can be found in the technical literature. We therefore need to first agree on what OBS is. We take a "classical" definition of OBS that is summarized in the following points.

1. *Separated transmission of header and payload.* More traditional packet switching techniques use packets that include header (addressing and control information) and payload (data carried by the packet). In OBS, the header information goes into a Burst Header Packet (BHP) transmitted in a separated control channel in the same fiber. This control channel is shared by the BHPs of bursts transmitted on the data channels belonging to the same link. The contention in the control channel is solved by buffering the BHPs in an electronic memory, if necessary.
2. *Header anticipation.* To adapt to potentially slow switching times of some optical technologies, and to the time needed to process BHPs, BHP and burst payload are separated in time by some offset.
3. *Electronic header processing.* Upon BHP arrival, the header is converted into the electronic domain for processing (all-optical BHP processing nodes were also proposed, but we refer here to an O/E/O operation of control channels). The scheduler of the node uses the addressing information, payload offset time and payload length information to allocate an output channel for the burst. The BHP is scheduled for transmission on the control channel available on the link where the burst was allocated.

4. *One-way-reservation*. BHPs reserve network (channels) and switching (input/output ports) resources in a distributed and progressive manner, according to a one-way reservation mechanism. If the reservation requested by a BHP cannot be served at a given node, *both BHP (first) and burst payload (later) are discarded*: neither BHP nor burst retransmissions are considered.
5. *Traffic assembling*. Bursts are “large” packets, typically of variable size. To compensate the overhead caused by the mandatory channel idle time during optical switching reconfiguration times (interburst gap), the burst must be somewhat larger than legacy (e.g., current Ethernet or IP) packets. Thus, bursts are assembled at the network edge, according to suitable *burstification* algorithms.
6. *Asynchronous operation*. No time-slotting is required in the network. The bursts can be of variable or of fixed size, but are not aligned at switch inputs.
7. *No optical buffering*. No buffering is available to delay burst payloads. Only fixed-length delays can be introduced at the input ports of optical nodes, for compensating the BHP processing time.

The above behavior was defined with the expectation of providing advantages with respect to the conventional O/E/O paradigm, in which all the traffic is electronically processed in a hop-by-hop and packet-per-packet basis:

1. A much lower number of O/E/O conversions are required, saving equipment cost and power consumption.
2. The cut-through switching reduces the end-to-end delay. The header anticipation helps to preserve this benefit even in presence of slow optical switching fabrics. In OBS, the offset is kept large enough so that the optical switching fabric is already reconfigured when the payload arrives to the node. Then, this reconfiguration time is added to the end-to-end propagation time only once (by an artificial delay of the payload at the ingress node), instead of one for every hop.
3. The burst aggregation yields a less demanding header processing. Moreover, the rate of number of headers per second to be processed by an OBS node is dependent on the average burst duration, but decoupled from the line rate in the data plane.

4. No optical buffers are needed to solve contentions.

These benefits of OBS with respect to a hop-by-hop electronic packet processing (pure O/E/O network) are clearly very promising. However, OBS has not emerged as the substituting technology for legacy O/E/O networks. Transparent optical networks are currently built over the OCS paradigm. Then, OBS should show its advantages against the OCS paradigm to become the winning candidate of the next network evolution. For this reason, in this paper we focus on a comparison between OCS and OBS. While it is generally accepted that packet switching has certain advantages over OCS, the direct comparison of OBS with pure electronic packet switching is beyond the scope of this paper.

III. THE OBS TESTING SCENARIO

This section is devoted to describing and justifying the OBS testing scenario. Every effort has been put in defining precisely the relevant network elements and engineering decisions, to allow results to be repeatable.

The OBS network is composed of OBS interconnection nodes and OBS edge nodes. Interconnection nodes are connected in a mesh topology by WDM *interconnection links* that can span hundreds of km. In general, each interconnection node is connected to one edge node¹, located at the same facility, by means of short WDM links (*edge links*). Edge nodes are specific equipment, in charge of assembling the bursts from the electronic traffic demand, and injecting them into the network. All the fibers in the links have the same number of wavelengths. This corresponds to a unique wavelength grid used across the network, which is a reasonable assumption in the backbone. While interconnection links are supposed to be composed of one fiber, an edge link can have two or more fibers if the amount of ingress/egress traffic to the node requires so. This is assumed since we consider the cost of the OBS edge links negligible (as they do not require amplification and dispersion compensation equipment).

A. Traffic engineering

For traffic engineering, a GMPLS-like control plane is assumed, and to ease the paper readability, a GMPLS terminology will be used. According to this assumption, (i) higher-layers traffic connections

¹ The model can be immediately generalized to hierarchical topologies by having no traffic generated and sinked by an edge node.

take the form of LSPs (Label Switched Paths), which can be established between edge nodes, (ii) edge nodes act like OBS LERs (Label Edge Routers), and (iii) the bursts belonging to the same LSP follow the same sequence of hops from the ingress to the egress edge nodes. According to this scheme, deflection routing techniques cannot be considered as an option for solving contention, and will not be included in the study. Note that by doing this, end-to-end burst order is preserved in (bufferless) OBS networks.

In order to favor the OBS contention resolution performance, the routing and capacity of each LSP is calculated by optimally solving a multicommodity flow problem [3] which minimizes the maximum utilization of the interconnection links (see Appendix I for details). This *minimax* optimization is a common traffic engineering technique in packet and burst switching networks, since its traffic loss performance is dominated by the link utilization. The traffic between two edge nodes is allowed to be split in different routes, to permit a better OBS traffic balance. Then, one LSP per input-output node pair and per different route is established in the simulation to implement the obtained optimal routing. The *minimax* optimization program has been solved using the TOMLAB/CPLEX [4] software.

B. OBS edge node

Fig. 1 illustrates the structure of the edge nodes. In Fig. 1, we denote as L_n and L_n' the number of LSPs initiated and ending at node n respectively. The traffic demand is supposed to be composed of IP packets that are classified and sent to one of the corresponding LSPs associated with the packet destination node. Then, they are assembled in bursts. The bursts leaving the assemblers are stored in an electronic queue, separated for each LSP, waiting for their turn to be transmitted. In its turn, the bursts targeted to the node, are converted to electronic form by the O/E converters and disassembled.

For each edge node n , the number of electro-optic transmitters T_n should be large enough to carry the ingress traffic injected by the node. The dimensioning of this number of transmitters represents one of the main benefits of the OBS paradigm. Since the transmitters are *shared* by all the LSPs, a multiplexing gain is achieved which allows a higher average utilization of the transmitters pool without an excessive penalization of the average electronic injection time. In our tests, we have calculated the number of transmitters so that their average utilization is equal to 80%. A fractional

number of transmitters is rounded up. The resulting average ingress delay and memory requirements observed in the edge nodes have shown to be within acceptable values.

Unfortunately, the number R_n of receivers of the edge node n cannot benefit from the previous multiplexing gain effect, since the incoming bursts compete for the wavelengths in the last hop without (optical) buffers. To favor OBS traffic loss performance, we use an infinite number of receivers in the nodes. They are internally arranged in an ordered list, so that a receiver is used only if preceding receivers are busy. This allows us to calculate, at the end of the simulation, the minimum number of receivers for which the target loss probability was guaranteed. This is the number used in the cost analysis of the network.

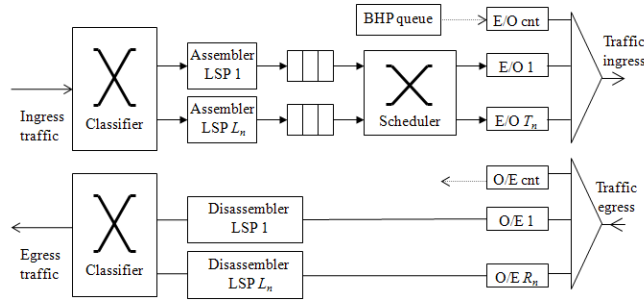


Figure 1. OBS edge node model

C. Burst assembly process

A plethora of assembling techniques has been proposed for the OBS edge nodes, based on a combination between time-based and volume-based assembling. While in general the burst size distribution depends on the traffic scenario and on the specific assembly strategy, in this paper we model the output process of the assembler as a sequence of bursts with a random independent payload duration following a truncated normal distribution, and independent exponential interarrival time. This model is in line with previous studies in the literature (see [5] for a seminal work in this topic). The minimum, maximum and average burst duration is set to be equal to 1, 100, and 50.5 times the Optical Switching Fabric Reconfiguration Time (OSFRT) of interconnection nodes. In an OBS networks, the OSFRT parameter determines the interburst gap: the mandatory idle time between the tail and the head of two consecutive payloads traversing the same wavelength in any fiber. This idle time is needed for

reconfiguring the Optical Switching Fabric (OSF) between two consecutive payloads arriving to the same input port of a node. Therefore, *short* payloads are a source of performance loss in OBS networks, since the proportion of mandatory idle time in the channels is increased. An average payload duration of 50.5 times the OSFRT means a limited $1/50.5$ (~2%) average overhead in the network, caused by OSF reconfiguration.

D. Reservation protocol

In the edge node, the transmission time of each assembled burst and the transmitter to be used is calculated by a dedicated electronic control unit (“Scheduler” in Fig. 1), when the burst leaves the assembler. It selects the transmission channel that can accommodate the payload sooner, but (i) no sooner than the *minimum waiting time in the electronic memory*, given by the initial offset of the burst, and (ii) respecting the interburst gap with the other bursts in the same channel. The initial offset of the burst depends on the OBS *reservation protocol* chosen for the network. In this paper, we include the tests for 3 different reservation protocols proposed in the literature: Just-Enough-Time (JET) [1], Only Destination Delay (ODD) [6] and Hop-by-hop Priority Increasing (HPI) [7]. JET assigns a fixed initial offset to the bursts, which shrinks hop-to-hop because of the header processing time in the nodes. The initial offset is given by the sum of processing times in the nodes along the LSP path, plus one OSFRT time. In ODD and HPI the initial offset is set to the OSFRT. Fixed length fiber delay lines are included in the input ports of the nodes (of length D in Fig. 2), to exactly compensate the header processing time (for ODD), or over-compensate it, so that the burst offset actually increases hop-by-hop (for HPI). In all the cases, the offset in every hop is kept higher or equal to the OSFRT to enable the cut-through switching at the intermediate nodes.

E. OBS interconnection node

Fig. 2 shows the scheme of an OBS interconnection node with N input and output interconnection fibers, with W wavelengths each, and one input and output link towards its associated edge node. The OSF transparently switches the arriving optical bursts from the input ports to the output ports. In this paper, we consider OSFs with full wavelength conversion. This means that a burst arriving at any input fiber and any input wavelength can be switched to any output fiber and any output wavelength.

Note that this is the most favorable case from the contention resolution point of view, and that wavelength converters are not mature devices today.

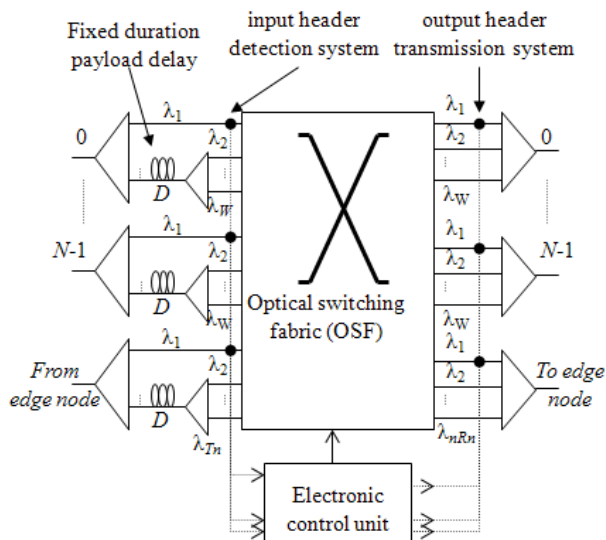


Figure 2. OBS interconnection node model

Incoming BHPs are processed by an electronic control unit, which implements the LAUC-VF (Latest Available Unused Channel with Void Filling) scheduling algorithm [8]. This choice is supported by numerous studies where this scheduler is used as an upper bound for performance in terms of the switch burst loss probability. We consider that the time required for the burst header processing T_{proc} is constant in every node and for every burst. The data channels in every input fiber go through a constant fiber delay line of length $D = \{0, T_{proc}, 1.5 T_{proc}\}$ in the JET, ODD and HPI reservation protocols respectively. The factor 1.5 in the delay for the HPI case is suggested by the authors in [7].

The contention among BHPs competing for transmission in the control channel is neglected. In real deployments, this contention could be addressed by a combination of different methods: by dropping the bursts, by increasing the number of control channels, or by buffering the headers waiting for their turn to be transmitted. Burst drop degrades system performance. The same effect occurs by increasing the number of control channels, since the number of usable data channels are reduced accordingly under our assumptions, and complexity is increased since control channels require O/E/O operation. Finally, header buffering increases the

burst processing time in a non-deterministic manner. This adds complex problems to engineer the offset between the burst header and burst payload in realistic OBS networks. In this paper, by neglecting the contention in the control channel, and by assuming a constant predictable header processing time, we provide a favorable scenario for OBS networks.

IV. THE OCS TESTING SCENARIO

In OCS networks, the traffic demand is carried on transparent *permanent* lightpaths (or with a sufficient holding time for being considered as permanent) which are optically switched in the intermediate nodes. A lightpath occupies one E/O transmitter in its initiating node, one wavelength in each traversed fiber, and one O/E receiver in the ending node. The traffic demand is routed on top of the lightpaths (which make up the so-called virtual topology), and the lightpaths are routed on top of the fiber links. For this reason, OCS networks are commonly named as optical multilayer networks. Traffic from a source edge node can in general traverse more than one lightpath, being (electronically) switched from one to another lightpath at intermediate OCS nodes. We refer to this switching and multiplexing functionality with the widely used term “traffic grooming”. The success of OCS has been its ability to decrease the electronic equipment cost with respect to pure O/E/O networks, since traffic is electronically processed only at the starting and ending nodes of a lightpath.

The OCS edge and interconnection node schemes of an arbitrary node n are shown in Fig. 3 and Fig. 4 respectively. The edge node comprises the electronic processing of the network. It is composed of an ingress classifier, an ingress scheduler with packet queues, an egress classifier, and a set of T_n E/O transmitters and R_n O/E receivers, one per each lightpath initiated and terminated at node n respectively. Input packets are first classified according to their LSP, and then switched, possibly after queuing, to their corresponding E/O transmitter. The egress electronic switch receives the traffic of the lightpaths ending in the node. Then, the egress traffic (targeted to node n) is forwarded outside the OCS network, while the groomed traffic (targeted to other nodes) is forwarded to the ingress classifier to be added to a lightpath originating from the node.

In this paper, the OCS network design is obtained by a heuristic algorithm implemented in the MatPlanWDM tool [9]. It comprises 3 consecutive steps:

Step 1. (Virtual Topology Design) The heuristic algorithm in [10] is executed. It takes the offered traffic matrix as an input parameter, and calculates (i) the number of lightpaths to be established between each pair of nodes, and (ii) how the traffic demand is routed on top of these lightpaths, (suboptimally) minimizing the number of used transceivers.

Step 2. (Lightpath Routing) The lightpaths calculated in the previous step are routed on the fibers by optimally solving an integer multicommodity flow problem [3], where the lightpaths are the commodities, and the capacity in the links is given by the number of wavelengths in the interconnection fibers. The objective function is set to minimize the average number of physical hops of the lightpaths.

Step 3. (Wavelength Assignment) The lightpaths are sequentially processed, and a wavelength is assigned to each one in a first-fit scheme. Wavelength converters are used only in the case in which a common wavelength in the route is not found.

Note that the algorithm stops after Step 2 when it does not find a feasible design for the set of lightpaths calculated in Step 1. If Step 2 succeeds, a feasible wavelength assignment can always be found in Step 3. Step 3 is then an attempt to heuristically minimize the number of wavelength converters used.

In the network design, the number of lightpaths in Step 1 is computed assuming that the maximum traffic that can be carried by a lightpath is $0.6 E$. This is equivalent to permitting a maximum average utilization of the lightpaths of 60%, and it means that the transmitters in the OCS network have an average utilization equal or lower than 60%. This is in line with existing multilayer network configurations: since packet delay performance strongly degrades at higher loads, and also accounting for the limited accuracy of traffic estimates, these networks are planned limiting the average lightpath utilization to these values. Note that the maximum average utilization planned for the OCS transmitters (60%) is lower than the one granted to the OBS transmitters (80%). The explanation of this choice comes from the different multiplexing gain effect achievable in the schedulers of the OBS and OCS edge nodes. In the OBS nodes, any transmitter is eligible to allocate any burst generated in the node, and the traffic is balanced among them. On the other hand, in OCS networks the transmitters are bound to a specific lightpath, and thus not shared by all traffic sources.

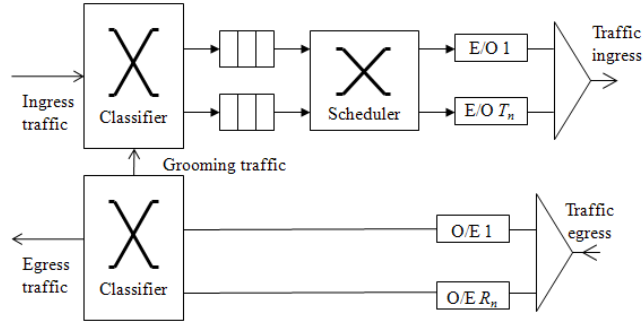


Figure 3. OCS edge node model

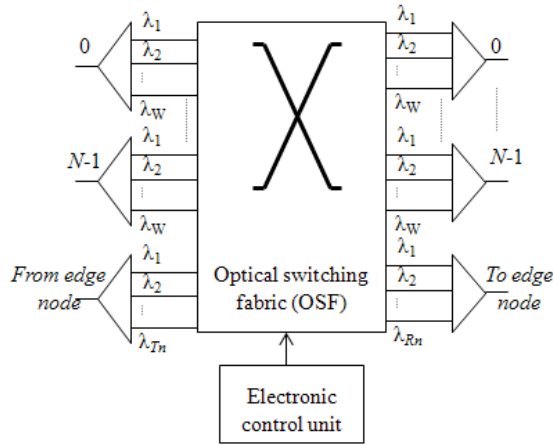


Figure 4. OCS interconnection node model

V. RESULTS COMPARISON

This section presents the results of our performance analysis. The results for OBS have been obtained by means of simulation in the OMNET++ [11] platform. Five samples were produced for each simulation point. The transient time (during which performance is not observed) is calculated as 110% of the longest end-to-end propagation time among the LSPs in the network. Each sample's running time was set so that at least one million bursts were received in the least loaded LSP. The confidence intervals, calculated from the samples with the *t*-Student method, confirmed the statistical validity of the results.

A total of 4374 simulation points were evaluated, covering every combination of the following 7 independent parameters:

1. (*PT*) Physical topology: Three reference network topologies were tested together with their traffic matrices, as described in the literature: Internet2 [12], NSFNET [13] and EON [14]. Table I summarizes some major data from these topologies (edge nodes attached to the interconnection nodes are not considered).
2. (*W*) Number of wavelengths per interconnection fiber: $W=\{20,40,80\}$. For OBS one of these wavelengths is used by the control channel, thereby reducing the number of data channels to $W-1$.
3. (*TI*) Traffic intensity: Nine traffic intensities were tested. The lowest traffic intensity corresponds to that for which the traffic (without the overhead) in the most congested interconnection link is equal to 10% of the link capacity, following the *minimax* routing engineered in OBS. This corresponds to the total demand volume in Erlangs shown in the last row of Table I, for $W=20$ wavelengths. The traffic volume for higher traffic intensities $TI=\{10\%, 20\%, \dots, 90\%\}$ and number of wavelengths are the corresponding multiples of these quantities.
4. (*RP*) Reservation protocol: As mentioned in Section III, three reservation protocols proposed in the literature are evaluated: JET, ODD and HPI.
5. (*CV*) Coefficient of variation of the payload duration distribution: $CV=\{0, 0.75, 1.5\}$. The value $CV=0$ produces fixed size bursts; it is useful for performance assessment, but not natural for the OBS paradigm.
6. (T_{proc}) Header processing time: $T_{proc} = \{1 \mu\text{s}, 10 \mu\text{s}\}$. They intend to model a slow and a fast electronic unit.
7. (*OSFRT*) Optical switching fabric reconfiguration time: $OSFRT=\{5 \text{ ns}, 1 \mu\text{s}, 1 \text{ ms}\}$. They approximately represent the switching time capabilities of the fabrics based on three potential enabling technologies: Semiconductor Optical Amplifiers (SOAs), microring resonators (MRRs) and Micro-Electro-Mechanical Systems (MEMS) respectively.

Table II summarizes the list of variables. The set of independent parameters was selected to exhaustively cover a reasonable range of possible OBS network configurations (*PT*, *W*, *RP*), traffic parameters (*TI*, *CV*), and technological scenarios (T_{proc} and *OSFRT*). The simulations were executed in

the Ben-Arabi Supercomputing facility, located in Murcia (Spain), using more than 25,000 hours of processor time.

The OCS planning was conducted as described in Section IV, for the same network configurations and traffics volumes defined by the network topology (PT), wavelengths per fiber (W) and traffic intensity (TI). Therefore, we are in a position to compare the cost-performance merits of OBS and OCS in the same network and traffic scenario. The following sections show and discuss our main findings.

TABLE I. INFORMATION ON THE TOPOLOGIES TESTED

	Internet2	NSFNET	EON
Reference	[12]	[13]	[14]
Nodes	9	14	18
Unidirectional links	26	42	66
Average in degree [link]	1.44	1.5	1.83
Average link length [km]	1009	1080	742
Max number of hops LSP [link]	4	5	5
Traffic volume $TI=0.1$, $W=20$ [E]	17.2	30.7	42.6

TABLE II. VARIABLE NAMES

Acronym	Values	Description
PT	{Internet2, NSFNET, EON}	Physical topology
W	{20, 40, 80}	Number of wavelengths in the interconnection fibers
TI	{10%, ..., 90%}	Traffic intensity
RP	{JET, ODD, HPI} (OBS)	Reservations protocol in the OBS network JET: Just-Enough-Time [1] ODD: Only destination delay [6] HPI: Hop-by-hop Priority Increasing [7]
CV	{0, 0.75, 1.5} (OBS)	Coefficient of variation of the payload duration distribution in the OBS network
T_{proc}	{1 μ s, 10 μ s} (OBS)	Header processing time in OBS nodes
$OSFRT$	{5 ns, 1 μ s, 1 ms}	Optical switching fabric reconfiguration time in OBS

	(OBS)	nodes
--	-------	-------

A. Maximum network throughput

In this section, we analyze the results for the maximum network throughput, as a measure of the maximum revenue that can be obtained from a network infrastructure. For OBS networks, we define the maximum network throughput as the maximum traffic intensity (TI) parameter, for which all the LSPs showed an *end-to-end* traffic loss probability lower than 10^{-4} (the bit loss probability is derived from the burst loss probability by weighting the losses by the burst length). In the OCS case, the network throughput is defined as the TI for which a feasible multilayer planning was found which fully carried the corresponding traffic demand. As a main observation, results obtained for the OBS setups showed that in the vast majority of the cases, the network throughput was the same for the different RP , CV , T_{proc} and $OSFRT$ values, being 10% the maximum variation of TI between the worse and the best OBS setup. In other words, the end-to-end traffic loss performance was approximately the same in all the tested technological scenarios, reservation protocols and traffic variability situations. This indicates that different header anticipation strategies and different switching speeds have little impact on the maximum network throughput. Table III helps us to observe the effects of the rest of the network parameters in the maximum network throughput. For each physical topology and wavelengths per fiber setup, it collects (i) the worse and best network throughput for OBS configurations varying the parameters RP , CV , T_{proc} and $OSFRT$, and (ii) the network throughput for OCS.

TABLE III. OBS VS OCS NETWORK THROUGHPUT COMPARISON

W	Topology	max TI OBS (min/max)	max TI OCS
20	Internet2	20% / 30%	50%
	NSFNET	20% / 30%	50%
	EON	20% / 30%	50%
40	Internet2	30% / 40%	50%
	NSFNET	30% / 40%	50%
	EON	30% / 40%	50%
80	Internet2	50% / 50%	50%
	NSFNET	50% / 60%	50%
	EON	50% / 50%	50%

Obtained results are similar for the three topologies. Unlike the OCS case, the maximum network throughput results in OBS networks are clearly dependent on the number of wavelengths in the links. This is a well-known effect, since OBS relies on the wavelength dimension to solve contentions. Conversely, the OCS paradigm is able to carry the traffic associated with $TI=50\%$, for the three numbers of wavelengths. It is interesting to see that for 20 and 40 wavelengths per fiber, OCS is capable of managing the optical bandwidth *more efficiently* than OBS (that is, it is able to deliver more traffic with the same link capacities). Only with 80 wavelengths per interconnection link, and for the NSFNET topology, the OBS paradigm improved the results of the OCS scheme, and thus was able to carry more traffic on the same infrastructure (but with a channel utilization in OCS not larger than 60%).

Finally, a more careful analysis of the 10% network throughput difference varying the $OSFRT$, T_{proc} , RP and CV parameters showed no clear trend, but a faint dependence on the reservation protocol. The set-ups based on the HPI reservation protocol exhibit in general better performance. This was somewhat expected given the original target of HPI of improving the end-to-end burst loss performance in the network by enforcing a more uniform distribution of the burst losses with respect to the number of hops.

B. Cost vs. revenue comparison

The cost side of the comparison between OBS and OCS network infrastructures combines qualitative and quantitative arguments. In order to discuss the former, we compare the complexity of the node schemes in Figs. 1, 2, 3 and 4:

- Optical switching equipment: The complexity of the OSF is clearly higher in the OBS case, since it is expected to be reconfigurable in a burst-by-burst manner. Moreover, we assume an OBS fabric with *fast* (per-burst) full wavelength converters, devices that today are at best in a prototype stage. In contrast, OCS fabrics are off-the-shelf equipment. Moreover, our results showed that no slow wavelength converters were needed in the considered OCS networks for most of the cases. Only in the case of $TI=\{40\%,50\%\}$, they were needed, but then, at most 2% of the lightpaths required

wavelength conversion. This confirms previous results found in the literature that showed that wavelength conversion was rarely required in the static planning of multilayer networks [15].

- Electronic processing: The differences in the electronic side of the node network do not allow a definitive statement. On the one hand, OBS networks add the assembling/disassembling stage to the edge node, and require fast implementations of the scheduling algorithm to reconfigure the interconnection nodes. On the other hand, the groomed traffic in OCS nodes requires extra switching capacity in ingress and egress switches. Also, the classifier and scheduler fabrics in the edge OCS nodes must process IP packets one by one, while in OBS nodes, the fabric controller should only take decisions on bursts that in general aggregate several IP packets, softening the processing requirements.

To complete the cost vs. revenue discussion we focus on two quantitative cost measures: (i) the number of E/O and O/E transceivers (at edge nodes, and on control channels –one per fiber– in OBS interconnection nodes), and (ii) the sum of grooming traffic in the OCS nodes normalized to the volume of the offered demand. The latter value represents the average extra amount of switching capacity in ingress and egress fabrics in Fig. 3 respect to Fig. 1, required to cope with traffic grooming.

Table IV collects the results obtained for the three reference topologies, $W=\{20,40,80\}$ wavelengths per fiber, and $TI=\{10\%,\dots,60\%\}$ traffic intensities. Column *Vol* shows the traffic demand volume, which is the same for both OBS and OCS networks. To allow the results to be comparable, we do not report the absolute number of transceivers (transmitters plus receivers) used, but the ratio between the number of transceivers and the demand volume (E) (*Tr/E OBS* and *Tr/E OCS* columns). The columns *ExtraTr* and *ExtraSw* contain the percentage of extra transceivers and extra electronic switching equipment (caused by traffic grooming), that OCS setups require with respect to OBS counterparts. Blank cells in the table mean that the OBS or OCS network was unable to carry the associated traffic demand.

TABLE IV. OBS VS OCS COST COMPARISON

INTERNET2	NSFNET	EON
-----------	--------	-----

W	TI	Vol (E)	Tr/E OBS	Tr/E OCS	ExtraTr	ExtraSw	Vol (E)	Tr/E OBS	Tr/E OCS	ExtraTr	ExtraSw	Vol (E)	Tr/E OBS	Tr/E OCS	ExtraTr	ExtraSw
20	0.1	17.1	9.2	5.0	-45.6%	29.4%	30.7	8.7	5.5	-36.1%	45.8%	42.6	9.0	5.6	-37.3%	38.4%
	0.2	34.3	6.9	3.9	-43.2%	9.4%	61.3	6.7	4.2	-36.3%	19.3%	85.2	6.8	4.6	-32.3%	25.6%
	0.3	51.4	6.1	3.8	-38.4%	6.4%	92.0	6.0	3.8	-36.0%	9.8%	127.8	6.0	4.1	-31.4%	19.5%
	0.4	68.6		3.6			122.6		3.7			170.4		3.9		
	0.5	85.7		3.6			153.3		3.6			213		3.7		
40	0.1	34.3	6.1	3.9	-35.6%	9.4%	61.3	5.8	4.2	-27.2%	19.3%	85.2	6.0	4.6	-22.8%	25.6%
	0.2	68.6	4.9	3.6	-26.6%	3.2%	122.6	4.8	3.7	-22.7%	6.6%	170.4	4.8	3.9	-19.1%	14.0%
	0.3	102.8	4.5	3.6	-21.1%	2.9%	183.9	4.4	3.6	-19.1%	4.2%	255.7	4.5	3.6	-19.0%	6.9%
	0.4	137.1	4.3	3.5	-20.1%	2.1%	245.2	4.3	3.5	-17.1%	3.5%	340.9	4.3	3.4	-20.8%	0.9%
	0.5	171.4		3.5			306.5		3.5			426.1		3.6		
80	0.1	68.6	4.4	3.6	-17.9%	3.2%	122.6	4.0	3.7	-6.4%	6.6%	170.4	4.4	3.9	-10.0%	14.0%
	0.2	137.1	3.8	3.5	-9.3%	2.1%	245.2	3.4	3.5	3.1%	3.5%	340.9	3.8	3.4	-10.3%	0.9%
	0.3	205.7	3.6	3.4	-5.5%	1.6%	367.8	3.3	3.5	6.2%	2.5%	511.3	3.6	3.5	-2.1%	4.6%
	0.4	274.2	3.5	3.4	-3.8%	1.4%	490.5	3.2	3.4	8.5%	1.9%	681.7	3.5	3.4	-3.7%	0.8%
	0.5	342.8	3.5	3.4	-2.5%	1.1%	613.1	3.1	3.4	9.1%	1.4%	852.2	3.5	3.5	-0.3%	3.0%
	0.6						735.7	3.1								

Quite similar results were obtained in the three reference topologies. The first relevant observation is that in the majority of the $\{W, TI\}$ scenarios in which both OBS and OCS are able to carry the traffic demand, the Tr/E efficiency of OCS is better than the Tr/E efficiency of OBS, and thus the $ExtraTr$ cost column holds a negative value. Only in the NSFNET tests with 80 wavelengths per fiber, could OBS show slightly better results (9%). Then, we can conclude that, given a traffic demand, OCS is in general able to carry it *using less transceivers* than OBS. The Tr/E efficiency of the OBS paradigm is degraded by the O/E/O conversions in control channels. This effect is more evident for low traffic volumes, and a lower number of wavelengths, since one control channel per fiber is used, whatever the traffic volume or number of wavelengths. In the OCS case, the Tr/E efficiency is almost independent of W and TI , with a slight improvement for higher values of both. This is explained by the discrete nature of the lightpath capacities. OCS networks can allocate more efficiently those volumes of traffic that are much larger than the lightpath capacity.

The extra electronic switching capacity in OCS networks caused by traffic grooming is given by the average number of traversed lightpaths (or virtual hops) of a packet, minus one. This is shown in the $ExtraSw$ column of Table IV. Naturally, it is always a positive value, reflecting that in OCS networks packets are electronically switched in the ingress and egress nodes (as in OBS networks), but also at the intermediate nodes where traffic

grooming is performed from incoming lightpaths to outgoing lightpaths. Results show that the proportion of groomed traffic decreases as the volume of carried traffic increases (higher values of TI and W), and that it becomes almost negligible at moderate to high traffic intensities, for $W=80$, or $W=40$. This was again an expected effect: for higher traffic volumes, the traffic between two nodes becomes significantly higher than the capacity of a lightpath, so that it is possible to carry almost all the traffic in direct lightpaths from the ingress to the egress nodes. In its turn, for lower traffic intensities, traffic grooming is exploited by the planning algorithm to reduce the number of lightpaths in the network: the algorithm routes traffic making use of the spare capacity in already planned lightpaths, instead of allocating new ones. Note that in these cases, although the OCS networks require more electronic switching, they use less transceivers than OBS setups. In other words, there is room for reducing the electronic switching requirements by using more direct lightpaths (transceivers), if this approach was preferred by the network planner.

From Table IV, we can state that in general OCS networks can transmit the same traffic or more traffic than OBS networks, using a lower amount of transceivers. The extra electronic switching capacity caused by the groomed traffic has no significant effect, especially at medium and high loads. In its turn, OBS requires more transceivers, and fast optical switching functionalities in every interconnection node.

VI. RELATED WORK

This section is dedicated to surveying previous works found in the literature, in which the performance of OBS and OCS paradigms was comparatively assessed [16]-[24]. The scope of the works [16]-[21] and [23] is a general comparison of the OBS and OCS paradigms, as candidates for the future network evolution. In [22] and [24], the authors include comparative evaluations of OBS and OCS paradigms in the framework of a multi-granular switching network, which may combine both schemes. The investigations reported in this paper have a different focus. In particular, we do not aim at a general comparison of OBS and OCS, but motivate our position regarding two specific frequently-repeated statements, which we refer to as “myths” in the title. Since this approach seems to be new in the literature, our interest in this section is to observe the methodologies chosen in designing the experiments in previous works, and comparing them with the approach followed in this paper.

The first notable observation extracted from the surveyed studies is the diversity of the obtained network performance results. Because of this, there is a lack of agreement on the conclusions, ranging from a strong support to the merits of OBS (e.g. [16],[18],[23]), to a more neutral and somewhat pessimistic view on them (e.g.[17],[19]). We found three explanations for these discrepancies. First, the differences in the considered network settings, which can lead to significant differences in the network performance (e.g., considering or not optical buffering in the OBS network). Second, the different criteria to assess the network cost. This affects more intensely to the OBS case, which relies on technologies not (yet) commercially available. Third, the dimensioning of OBS and OCS networks has been shown to be very sensible to some network planning decisions, which can lead to significant performance differences, and trigger possibly contradictory conclusions.

On the OBS side, a major issue affecting the performance results is how burst losses are estimated in core nodes. To this end, the analyses proposed in the literature to predict the *end-to-end* losses in an OBS network require significant simplifying assumptions, raising doubts on the validity of results for real networks ([17],[21],[23]). In other works ([16],[21],[22]), the experiments produced very large burst loss probabilities (e.g. 10^{-3} and higher), or large burst ingress delays (e.g. in the order of seconds [21]). In these circumstances, the applicability of the conclusions should be questioned. Finally, the approach followed in [20] considers that edge nodes retransmit those bursts which are dropped in the interconnection nodes. However, the complexity associated with a retransmission scheme penalizes this option.

In the dimensioning of OCS networks, the aspects that most strongly affect the network design are the quality of the planning algorithm, and the use of traffic grooming as a method for efficient use of the optical equipment. Under this respect, we observed that most of the previous works use variations of a simple shortest-path planning of the network (e.g., [16],[17],[18],[19]), and/or do not consider the possibility of traffic grooming (e.g., [17],[18],[19]). By so doing, the performance on OCS networks can be degraded. A possible source of performance degradation in OCS networks that we did not consider is the one contemplated in [20], where the authors assume that the lightpath setup requests reaching an edge node are immediately followed by the data traffic. This data traffic is stored in electronic buffers in the edge node for a time equal to the two-way reservation procedure. The edge

buffer overflow caused by this effect is computed as a source of traffic losses in OCS networks. However, in our opinion, whether this degradation of the performance of OCS networks may occur is debatable. In our view, despite the higher and higher dynamicity in traffic demand of future backbone networks, a sort of connection control mechanism is expected to exist, which accepts or rejects requests. As a consequence, traffic sources would wait for the completion of the lightpath setup procedure before injecting traffic in the electronic buffers of the edge nodes.

Our approach to building the testing scenario differs from previous studies in several aspects. First, we base our evaluation of OBS on extensive simulations of OBS networks. This computational effort largely exceeds analogous studies performed in the past. By relying on system-wide simulation for OBS evaluation, we avoid the simplifications often needed to analyze the end-to-end performance of OBS networks. Second, we use a powerful multilayer planning algorithm to dimension the OCS network, which benefits from traffic grooming to minimize the number of transceivers. In order to obtain more conclusive results, we complete an exhaustive evaluation of a large set of OBS configurations, using three reference topologies, and their associated traffic matrices brought from the literature. We test different values for technology-related parameters like the header processing time or the optical switching fabric reconfiguration time in network nodes. Also, three different reservation protocols in the network are considered. This wider, but carefully selected, set of system parameters should increase the credibility of results with respect to previous studies, which commonly relied on fixed values of some or all of these parameters ([16]-[24]), were based on simple regular topologies (e.g., [19]), or non-realistic uniform traffic matrices (e.g., [16],[18],[20]). Finally, we aimed at identifying the upper limits of OBS performance, by taking design decisions that favor OBS performance when reasonable.

VII. CHALLENGING OBS MERITS

Originally, OBS attracted the attention of the community because of its potentials to reduce the equipment cost associated with the pure O/E/O paradigm, eliminating the need of a per-hop packet-by-packet processing. Since OBS could win any OBS-vs-O/E/O comparison in terms of the electronic cost and capacity scaling, it raised the interest of the community to investigate the performance side of

the question. But when compared to the OCS paradigm, the cost argument seems to be no longer favoring OBS. A common agreement in the research community exists on this, and the qualitative and quantitative observations presented in this paper support this point. Moreover, the simplicity of OCS, its low power consumption, its mature and relatively simple way of providing QoS differentiation within the carried traffic (favored by its circuit-switching nature), are all important features that also tip the scales in favor of OCS with respect to OBS. Therefore, to beat the OCS paradigm, OBS must clearly outperform OCS in the revenue part of the balance.

A common argument found in the literature supporting OBS from the performance side, states that OBS is able to provide a more efficient management of the bandwidth because of its sub-wavelength switching granularity. However, our results contradict this assertion in two aspects. First, in the majority of the tests when both OBS and OCS paradigms were able to satisfy a given traffic demand, OCS required less transceivers. Second, results show that the OCS paradigm can in general carry more traffic in the same network infrastructure. In particular, across our simulation campaign, only in the NSFNET, and for the DWDM case $W=80$, OBS could reach a higher network utilization than OCS (with a loss target equal to 10^{-4}). A significant fact can illustrate the OCS efficiency: in all cases it could deliver the same amount of traffic as a pure O/E/O network with a maximum utilization per link of 50%, but enormously cutting down the electronic processing costs (as the traffic grooming figures show). Since this is achieved utilizing the lightpaths at a 60% (a realistic value in current networks), the extra inefficiency caused by the circuit-switched bandwidth delivery and the discrete nature of the lightpaths seems to be more than acceptable. In contrast to this, OBS degrades its network utilization by two unavoidable overheads (interburst gap and control wavelength), and one contention-related limitation, given by its need to solve the bursts contentions optically. This last aspect creates a further drawback on the OBS side. Results show that the OCS networks manage to provide a fairly good utilization of the network, even for a low number of wavelengths per fiber. On the contrary, OBS networks require a large amount of wavelengths to be a competing candidate, since the OBS paradigm relies on the wavelength dimension for contention resolution. Therefore, a pay-as-you-grow migration strategy which grants a fraction of the wavelength grid to an OBS network, sharing the links with OCS

equipment, is non-profitable, since the OBS part of the network would have to be grossly underutilized to be within the traffic loss target.

The results obtained contradict a frequently repeated campaign for a bufferless operation of the OBS nodes. In the authors' opinion, the extensive evaluation reported in this paper, seeking the upper performance limits of the bufferless OBS paradigm, in fact endorses the claim that OBS paradigm *cannot* be bufferless, in order to be a competing option. In plain words, our results show that a bufferless OBS alternative cannot compete with current commercial OCS networks on either side of the cost-revenue comparison. This is witnessed by current network realizations, based upon IP electronic switching at the edges of optically switched circuits.

What, in authors' opinion, should be changed in the OBS model, as described in Section II, to be considered a potential substituting candidate to OCS? Outperforming OCS in the cost side of the comparison is not foreseeable for OBS, or any sub-wavelength switching technology (like OPS). Then, the arguments must be brought from the performance side. To increase its network throughput, the OBS paradigm must improve its contention resolution performance.

Adding optical buffers to OBS nodes (hence contention resolution in the time domain), with the aim of improving its contention resolution capabilities, has already been considered in the OBS literature in numerous previous works. However, it brings another frequently mentioned OBS merit into discussion. OBS was originally presented as a valid approach to implement an optical sub-wavelength granularity switching making use of slow optical switching nodes. Then, OBS networks could be based on MEMS-based nodes, a mature and more cost-effective solution if compared to SOA-based fast optical switching devices. Long bursts (e.g., of an average length in the order of 10 ms) are needed in this case, in order to limit the overhead caused by the interburst gap. The increase in the assembling and electronic injection delays associated with longer bursts is then compensated by two techniques in OBS: the cut-through operation in the interconnection nodes and the header advancement feature. However, if we eliminate a bufferless operation from the equation, new issues come into play. Assuming that the optical buffering is based on Fiber Delay Lines (FDL), their length should be close to a multiple of the average burst duration to become effective [25]. A burst duration in the ms range, means hundreds of km fiber for building one single FDL. As an example, an FDL of

10 ms needs 2000 km of fiber, and about 25 optical amplification spans. It is hard to conceive a technological situation which can make this approach profitable.

In this paper, we refer to bufferless OBS as a *myth*, since a bufferless OBS operation has not yet been recognized in the literature as a losing option in the cost and performance comparison with OCS. On the contrary, the possibility of a bufferless OBS is commonly found as an argument favoring OBS in contrast to other paradigms like OPS (and one may wonder why is it not possible to build bufferless OPS nodes).

A second alternative that is frequently considered for improving contention performance, which can be combined with the use of optical buffers, is enforcing a slotted operation of the network. That means delivering fixed-size bursts that are optically aligned at switch inputs by means of optical synchronization stages. The authors would like to note that applying this alternative to OBS would mean abandoning its asynchronous operation, one of its hallmarks, and adopt what has been classically seen as an OPS-defining variant.

The assessment of the performance merits of optical buffering combined or not with burst alignment is outside the scope of this paper. Naturally, it is possible to increase the contention resolution performance of OBS nodes by adding a sufficient amount of optical buffering (and the associated extra optical switching requirements in the nodes). Then, these OBS network settings would be able to carry more traffic than their OCS counterparts. Therefore, for these situations a cost-performance trade-off appears, where the OBS option is a winner performance-wise. We do not intend to explore these trade-offs, which are more related to a general comparison between OCS and OBS paradigms and also bring in features often related to OPS. We prefer to limit our study to motivate our skepticism towards the two assertions referred as “myths” in this paper. Moreover, since the enabling technologies for implementing the optical burst buffering and/or alignment are not commercially available, the uncertainty in the cost side of the trade-off would not allow the drawing of clear conclusions.

VIII. CONCLUSIONS

This paper critically revised some of the merits of OBS as a competing candidate to replace OCS backbone networks. In particular, we motivate our disbelief towards two statements commonly associated with the OBS paradigm: its ability to deliver the optical bandwidth more efficiently than the OCS approach, thanks to its sub-wavelength switching granularity, and its ability to solve contention without optical buffers. To this end, we carefully designed a set of testing scenarios. By comparing cost/performance tradeoffs in them, we observed that OBS does not exhibit these frequently claimed advantages. We further argued that, to be superior to the currently widely used OCS, OBS must be empowered with features that make it very similar to OPS. We think that these investigations seriously question the suitability of OBS as a feasible intermediate technology between current OCS and the long-term goal of OPS. In other words, we foresee that, if highly dynamic (i.e., packet-by-packet) optical switching finds technical maturity and commercial deployment, it will take a form much closer to what is today called OPS than to OBS.

IX. APPENDIX I

This Appendix describes the minimax optimization problem solved to plan the OBS network. Let N be the set of nodes in the network, E the set of interconnection fibers and W the number of wavelengths in the fibers. Let $\delta^+(n)$ and $\delta^-(n)$ denote the set of links outgoing from and incoming to node $n \in N$, respectively. Let T^{sd} denote the amount of traffic in Erlangs to be carried from node s to node d . Let x_{sde} denote the traffic originated in node s and targeted to node d that traverses link e . Note that variables x_{sde} define the traffic routing. The minimax optimization to solve is given by:

Find $M \geq 0, x_{sde} \geq 0, s, d \in N, e \in E$, that minimizes M/W , subject to:

$$\sum_{e \in \delta^+(n)} x_{sde} - \sum_{e \in \delta^-(n)} x_{sde} = \begin{cases} T_{sd} & \text{if } n = s \\ -T_{sd} & \text{if } n = d, \forall n, s, d \in N \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

$$\sum_{s, d \in N} x_{sde} \leq M, \forall e \in E \quad (2)$$

Constraints (1) are the flow conservation constraints. Constraints (2) together with the objective function make variable M represent the traffic traversing the highest loaded link. Then, the objective function to minimize is M/W , equal to the maximum utilization among the links in the network.

IN MEMORIAM

This paper is the result of more than two years of discussions and investigations, triggered by the ideas Fabio Neri and me shared in a coffee-break of the ONDM 2008 conference. Sadly, Fabio passed away unexpectedly in April 2011, a huge loss to his family, friends and the research community. I would like to dedicate this paper to his memory.

ACKNOWLEDGMENT

The work described in this paper was carried out with the support of the BONE project (“Building the Future Optical Network in Europe”); a Network of Excellence funded by the European Commission through the 7th ICT-Framework Program. This research has been partially supported by the MEC project TEC2010-21405-C02-02/CALM, it is also developed in the framework of the projects from Fundación Seneca 00002/CS/08 (FORMA) and "Programa de Ayudas a Grupos de Excelencia de la Región. de Murcia", F. Séneca (Plan Regional de Ciencia y Tecnología 2007/2010)."

REFERENCES

- [1] C. Qiao and M.Yoo, “Optical burst switching (OBS)—A new paradigm for an optical internet,” *J. High Speed Netw.*, vol. 8, pp. 69–84, 1999.
- [2] J. Turner, “Terabit burst switching”, *Journal of High Speed Networks*, vol.8, pp. 3-16, 1999.
- [3] D. P. Bertsekas, “Network optimization: continuous and discrete models”, Athena Scientific, 1998.
- [4] <http://tomopt.com/tomlab/products/cplex/>. Last accessed 2nd Feb. 2010.
- [5] X. Yu, Y. Chen and C. Qiao, “Study of Traffic Statistics of Assembled Burst Traffic in Optical Burst Switched Networks”, in *Proceedings of Optical Networking and Communications Conference 2002 (Opticomm 2002)*, Boston (US), July 2002, pp. 149-159.

- [6] L. Xu, H. G. Perros, and G. N. Rouskas, "A simulation study of optical burst switching and access protocols for WDM ring networks," *Computer Networks*, 41, pp.143-160, Feb.2003.
- [7] B. C. Kim, J. H. Lee, Y. Z. Cho and D. Montgomery, "An Efficient Optical Burst Switching Technique for Multi-Hop Networks", *IEICE Trans. Commun.*, vol.E87-B,NO.6 June 2004.
- [8] Y. Xiong, M. Vandenhoute, H. C. Cankaya, "Control Architecture in Optical Burst-Switched WDM Networks", *IEEE Journal on Selected Areas in Communications*, vol.18, no. 10, Oct. 2000.
- [9] P. Pavon-Mariño, R. Aparicio-Pardo, G. Moreno-Muñoz, J. Garcia-Haro, J. Veiga-Gontan, "MatPlanWDM: An educational tool for network planning in wavelength-routing networks", *Lecture Notes in Computer Science*, vol. 4534, pp. 58-67.
- [10] V.R. Konda, T.Y. Chow, "Algorithm for traffic grooming in optical networks to minimize the number of transceivers", in: *IEEE Workshop on High Performance Switching and Routing 2001*, Dallas (US), May 2001, pp. 218–221.
- [11] <http://www.omnetpp.org>. Last accessed 2nd Feb. 2010.
- [12] Internet 2 Global Research Network Operations Center, Web Site, [Online], Available: <http://www.abilene.iu.edu/>
- [13] R. Ramaswami, K.N. Sivarajan, "Design of logical topologies for wavelength-routed optical networks", *IEEE J. on Selected Areas in Communications*, vol.14, no.5, pp.840–851, Jun.1996.
- [14] L. Wuttisittikulij, M.J. O'Mahony, "Design of a WDM network using a multiple ring approach", in *Proc. IEEE Global Telecommunications Conference*, Phoenix, AZ, 1997, vol. 1, pp. 551-555.
- [15] P. Pavon-Marino, R. Aparicio-Pardo, B. Garcia-Manrubia, J. Fernandez-Palacios, O. Gonzalez, F. Martin, J. Garcia-Haro, "Balancing Multifibre and Wavelength Converter Cost in Wavelength Routing Networks", *European Conference on Optical Communications (ECOC 2008)*, Brussels (Belgium), September 2008, pp. 149-156.
- [16] Fei Xue, S. J. Ben Yoo, H. Yokoyama, Y. Horiuchi, "Performance Comparison of Optical Burst and Circuit Switched Networks", in *OFC/NFOEC Conf. Tech. Dig.*, Mar. 6–11, 2005, vol. 3, 3 p.

- [17] A. Agrawal¹, T. S. El-Bawab¹, L. B. Sofman, “Comparative Account of Bandwidth Efficiency in Optical Burst Switching and Optical Circuit Switching Networks”, *Photonic Network Communications*, vol. 9, no. 3, pp. 297-309, 2005.
- [18] T. Coutelen, H. Elbiaze, B. Jaumard , “Performance Comparison of OCS and OBS Switching Paradigms”, in Proc. 7th International Conference on Transparent Optical Networks (ICTON 2005), Barcelona (Spain), July 2005.
- [19] E . Zouganeli, R. Andreassen, B. Feng, A. Solem, N. Stol, H. Kjongsberg, A. Sudbo, B. El. Helvik, R. B. Haugen, “Why bother with optical packets? An evaluation of the viability of optical packet/burst switching”, *Teletronikk* 2.2005.
- [20] X. Liu, C. Qiao, X. Yu, and W. Gong, " A Fair Packet-Level Performance Comparison of OBS and OCS," in Optical Fiber Communication Conference and Exposition and The National Fiber Optic Engineers Conference, Technical Digest (CD) (OSA, 2006), paper JThB48.
- [21] A Zalesky, “To burst or circuit switch?”, *IEEE/ACM Transactions on Networking*, vol. 17, no. 1, pp. 305-318, Feb. 2009.
- [22] M. D. Leenher, C. Develder, J. Buysse, B. Dhoedt, P. Demeester, “Performance analysis and dimensioning of multi-granular optical networks”, *Optical Switching and Networking*, vol. 6, no. 2, pp. 88-98, April 2009.
- [23] R. Parthiban, C. Leckie, A. Zalesky, M. Zukerman, and R. S. Tucker, “Cost Comparison of Optical Circuit-Switched and Burst-Switched Networks”, *IEEE/OSA Journal of Lightwave Technology*, vol. 27, no. 13, pp. 2315-2329, July 2009.
- [24] C. M. Gauger¹, E. Van Breusegem, P. J. Kühn¹, M. Pickavet, P.Demeester, “Hybrid Optical Network Architectures: Bringing Packets and Circuits Together”, *IEEE Communications Magazine*, vol. 44, no. 8, pp. 36–42, Aug. 2006.
- [25] D. Careglio, J. Sole-Pareta, S. Spadaro, “Service category-to-wavelength selection technique for QoS support in connection-oriented optical packet switching”, *Computer Networks*, vol. 51, no. 1, pp. 14-30, Jan. 2007.