

Review

A survey of quality of service differentiation mechanisms for optical burst switching networks

Nail Akar^a, Ezhan Karasan^{a,*}, Kyriakos G. Vlachos^b, Emmanouel A. Varvarigos^b,
Davide Careglio^c, Miroslaw Klinkowski^{c,d}, Josep Solé-Pareta^c

^a Electrical and Electronics Engineering Department, Bilkent University, Bilkent 06800, Ankara, Turkey

^b Computer Engineering and Informatics Department, University of Patras, 26500 Patras, Greece

^c Universitat Politècnica de Catalunya, 08034, Barcelona, Spain

^d National Institute of Telecommunications, 04894 Warsaw, Poland

ARTICLE INFO

Article history:

Received 24 March 2009

Accepted 18 September 2009

Available online 26 September 2009

Keywords:

Optical burst switching

Quality of service differentiation

ABSTRACT

This paper presents an overview of Quality of Service (QoS) differentiation mechanisms proposed for Optical Burst Switching (OBS) networks. OBS has been proposed to couple the benefits of both circuit and packet switching for the “on demand” use of capacity in the future optical Internet. In such a case, QoS support imposes some important challenges before this technology is deployed. This paper takes a broader view on QoS, including QoS differentiation not only at the burst but also at the transport levels for OBS networks. A classification of existing QoS differentiation mechanisms for OBS is given and their efficiency and complexity are comparatively discussed. We provide numerical examples on how QoS differentiation with respect to burst loss rate and transport layer throughput can be achieved in OBS networks.

© 2009 Elsevier B.V. All rights reserved.

Contents

1. Introduction.....	2
2. QoS differentiation mechanisms for OBS networks.....	3
2.1. QoS differentiation with one-way signaling.....	3
2.1.1. Edge-based QoS differentiation mechanisms.....	3
2.1.2. Core-based QoS differentiation mechanisms.....	4
2.1.3. Edge-core-based feedback mechanisms.....	5
2.2. QoS differentiation with two-way signaling.....	5
2.3. QoS differentiation with control plane methods.....	6
3. Numerical results.....	6
3.1. Performance of QoS mechanisms with one-way signaling.....	6
3.1.1. UDP offered traffic.....	6
3.1.2. TCP offered traffic.....	8
3.2. Performance of QoS mechanisms with two-way signaling.....	9
4. Conclusions.....	9
Acknowledgments.....	11
References.....	11

* Corresponding author. Tel.: +90 312 2901308; fax: +90 312 2664192.

E-mail address: ezhan@ee.bilkent.edu.tr (E. Karasan).

1. Introduction

Optical burst switching (OBS) has been widely accepted as a candidate transport architecture for the next generation optical Internet using the strengths of both optical and electronic technologies that are complementary. An optical Internet is defined as an internetwork where the link layer connections are dedicated wavelength channels which are directly interfaced to a high performance (all) optical network router. In this scenario, the high performance optical router replaces traditional ATM and SONET switching and multiplexing equipment as the essential node device that controls wavelength access, switching, routing and protection. OBS possesses significant advantages for such a router implementation. In essence, OBS technology “delayers” the complexity of many of today’s telecommunication networks and allows IP network traffic to be optimized for maximum throughput and speed. While an OBS-based Internet can be simpler to deploy and manage, it can also complement and enhance the delivery of Internet services over traditional ATM and other transport technologies.

The underlying principle of OBS is the segregation of the control and data planes by means of dedicating typically one wavelength for the control plane that uses electronic switching, whereas the remaining bulk of the wavelengths is used for the data plane and the switching in this plane takes place in the optical domain. In this architecture, client packets (e.g., IP packets) are aggregated into so-called bursts at the edge of the OBS domain. In general, bursts should be made long enough to justify the low control plane overhead per data burst, but the time required to form a burst should be kept short enough to satisfy the delay requirements of client packets. Once a data burst is formed, the OBS ingress edge node initiates a control message on behalf of the data burst. This control message is sent over the control plane to an OBS node which, after processing this message, forwards it toward the destination egress OBS node. Control messages are subject to O/E/O conversion at each OBS node on the path to destination and they serve to inform each node of the associated data burst and initiate configuration of the node so as to accommodate the data burst. The ingress OBS edge node then sends the data burst itself after a time offset over the data plane using the same path. The successful transmission of an optical burst through an OBS node depends on whether the data burst finds the OBS node already configured for switching it from the incoming fiber to the outgoing fiber. In addition to this basic OBS system description, there are variations of OBS in terms of the signaling schemes that would be used in the control plane. Signaling schemes for OBS networks are generally categorized as one-way or two-way schemes, while QoS mechanisms for OBS are different from those employed in IP networks.

It is well known that client packets belong to different traffic classes with different QoS requirements in terms of performance parameters such as loss, delay, delay jitter etc. QoS provisioning refers to a collection of methods that are used in order to meet the QoS requirements of different classes. There are mainly two QoS provisioning proposals for IP layer QoS. In the integrated services (IntServ)

architecture [1], a set of extensions to the traditional best effort model of the Internet are proposed in order to provide end-to-end QoS guarantees to applications with quantitative performance requirements. In particular, the guaranteed service [2] provides absolute guarantees: an assured level of bandwidth, a firm end-to-end delay bound, and no loss due to queuing if the packets conform to some a priori negotiated contract. The integrated services model is therefore referred to as an absolute QoS model and is intended for applications with stringent real-time delivery requirements, such as audio and video applications with playback buffers. On the other hand, a relative QoS model proposed for IP networks is the differentiated services (DiffServ) model [3]. In DiffServ networks, IP packets are classified into one of a small number of aggregated classes based on the DiffServ Codepoint (DSCP) written in the Differentiated Services field of the packet’s IP header [3]. At each router in a DiffServ domain, packets from different classes receive a different Per Hop Behavior (PHB) (invoked by the DSCP) using per-class queuing and buffer management techniques. For example, high priority (HP) traffic can be isolated from low priority (LP) traffic by strict priority scheduling or deficit round robin (DRR) scheduling. However, class-based queuing and advanced scheduling techniques that are used for QoS provisioning in IP networks cannot be immediately used in OBS networks due to a lack of optical buffering with current optical technologies. It is then desirable to develop new mechanisms by which existing QoS models in IP networks can be extended to OBS domains.

In OBS networks, there also exist relative and absolute methods for ensuring QoS. In relative QoS methods, the performance of a traffic class is defined with respect to other classes. For instance, it may be guaranteed that the loss probability of bursts belonging to the HP class is lower than that of the bursts belonging to the LP class. The performance of a given class in the relative QoS model usually depends on traffic characteristics of other classes, whilst the absolute QoS model aims at irrelative QoS provisioning. The absolute QoS model requires more complex implementations in order to achieve desired levels of quality in a wide range of traffic conditions, while at the same time maintaining high output link utilization. For complexity reasons, most QoS mechanisms considered for OBS networks basically offer relative QoS guarantees. In this paper, we focus on relative QoS (i.e., QoS differentiation) methods for OBS networks, which is along the line of the DiffServ model for IP networks. A detailed review of absolute QoS differentiation methods can be found in [4].

QoS differentiation mechanisms for OBS networks differ depending on the signaling scheme used, i.e., one-way and two-way schemes. In one-way signaling (reservation) schemes (also called “tell-and-go” schemes), a setup packet is sent in advance to precede the arrival of the data burst by a time offset without having to wait for a positive acknowledgment from the nodes along the path in the OBS domain. This allows for minimization of the pre-transmission delay but can result in waste of resources since, in this scheme, a burst can travel to the final hop, but get dropped at the final hop due to lack of




Fig. 1. Illustration of one-way (in particular the JET protocol) and two-way signaling schemes for OBS networks.

resources before reaching the egress OBS node. A number of one-way signaling schemes have been proposed for OBS, including the Ready-to-Go Virtual Circuit protocol [5], Just-In-Time (JIT) [6], Just-Enough-Time (JET) [7], and Horizon [8,9]. In two-way reservation schemes (also called wait-for-reservation or “tell-and-wait”), on the other hand, end-to-end connections have to be fully established before transmission of any data can start, while resources are reserved immediately upon the arrival of the setup message. Recent research efforts like EBRP [10], and WR-OBS [11], have shown that such reservation schemes can enable the implementation of a bufferless core network with limited wavelength conversion capability by moving processing and buffering functions to the edge. An illustration of one-way (in particular the JET protocol) and two-way signaling schemes for OBS networks are given in Fig. 1.

One-way schemes are promising due to low pre-transmission delays, and they are particularly effective for long-distance wide area networks. However, due to their open-loop nature, one-way schemes may result in high burst loss rates, especially under moderate to heavy traffic loads. For one-way schemes that employ delayed reservations, sophisticated channel scheduling and void filling algorithms have been proposed to resolve contentions and efficiently utilize the available bandwidth [12]. On the other hand, two-way signaling schemes possess closed-loop connection establishment, and therefore client packets can be held at electronic edge buffers when contentions occur. Consequently, reliance on advanced contention resolution capability is relaxed for OBS nodes in two-way schemes. On the downside, two-way signaling introduces a connection establishment latency and such architectures are therefore considered more appropriate for short-distance metropolitan networks [13].

The goal of this study is to provide a survey of QoS differentiation mechanisms proposed in the literature while also describing the research that we carry out on OBS QoS under the European Commission funded project BONE (Building the Future Optical Network in Europe [14]). In Section 2, we provide a general overview of existing QoS mechanisms for OBS networks with a suitable categorization. In Section 3, we provide simulation results using some selected QoS mechanisms using both streaming and elastic traffic. We conclude in the final section.

2. QoS differentiation mechanisms for OBS networks

We first introduce QoS differentiation mechanisms for one-way and two-way signaling schemes. We note that some of the QoS mechanisms proposed for one-way signaling systems may also be used for OBS networks with two-way signaling, but we will describe these mechanisms as one-way only complying with the original intent. In contrast, few QoS provisioning schemes have been proposed for two-way reservation protocols. This is because two-way OBS has attracted relatively less attention in the literature. However, two-way OBS can be advantageous compared to one-way schemes when delay is not a primary concern for the application and/or when the network has a high traffic load resulting in a high burst loss ratio (or probability) if one-way signaling is used. Finally in Section 2.3, we present some QoS mechanisms for OBS networks that rely on signaling and routing protocols running on the control plane, which we categorize as control plane methods for QoS differentiation.

2.1. QoS differentiation with one-way signaling

The one-way reservation scheme needs additional support in QoS provisioning in order to preserve HP traffic from LP traffic during both the resource reservation process and burst transmission. One-way QoS differentiation mechanisms can be categorized as:

- edge-based: mechanisms are implemented only at the OBS ingress edge node and the core nodes are not involved,
- core-based: mechanisms are implemented only at the OBS core nodes and the edge nodes are not involved,
- edge-core-based: mechanisms require the involvement of both OBS ingress and OBS core nodes.

Most of the proposed mechanisms for QoS differentiation for OBS networks use the burst loss probability as the primary performance metric of interest. Delay is also an important metric that should also be considered since it has a substantial impact on the throughput achievable at the transport and application layers.

2.1.1. Edge-based QoS differentiation mechanisms

Basically, two mechanisms have been proposed for edge-based QoS differentiation: offset time-based and burst length-based differentiation.

Offset time-based differentiation (OTD): This QoS differentiation method is probably the most explored QoS differentiation technique in OBS networks [15]. In OTD, an extra offset time is assigned to HP bursts, resulting in an earlier reservation for HP bursts in order to favor them while the resource reservation is performed (see Fig. 2 for an illustration). OTD mechanism allows absolute isolation between HP and LP classes, i.e., no HP class burst is blocked by an LP class burst. However, to achieve almost perfect isolation, the length of the extra offset time has to be as long, at least, as a few average LP burst durations. The main advantage of OTD is its simplicity; it reduces the loss probability of HP bursts by means of their postponed transmission from the




Fig. 2. Illustration of OTD: PT is the processing time, OT is the offset time.

edge node and no differentiation mechanism is needed in the core nodes. The drawbacks of OTD are the sensitivity of the HP traffic class to burst length characteristics [16] and the need for extended pre-transmission delay that may not be tolerated by some delay-constrained applications. Moreover, the end-to-end delay for HP traffic increases as a result of increased offset times which, in turn, decreases the throughput.

Burst length-based differentiation (BLD): The underlying idea of BLD is that short bursts are more likely to fit in gaps generated by already scheduled bursts. Consequently, in BLD method, HP class is assigned shorter burst lengths than the LP class for enhancing the performance of the HP class relative to the LP class in terms of loss probabilities. In particular, HP packets are burstified using lower timer [17] and lower burst length thresholds [18] compared with the corresponding values for LP packets (see Fig. 3). Another advantage of using a shorter burstification timer for the HP class is to reduce the end-to-end-delay. BLD can also be used as a complementary method in conjunction with another differentiation mechanism, such as OTD, for improved isolation between traffic classes [19]. The downside of BLD is that the burst assembly unit is more complex and signaling overhead increases due to increased number of control packets stemming from shorter HP bursts. Moreover, in order for this method to be effective, sophisticated void filling algorithms need to be already in place at the OBS core nodes as opposed to simple-to-implement horizon-based scheduling mechanisms that do not take advantage of voids [12].

2.1.2. Core-based QoS differentiation mechanisms

QoS differentiation in core nodes takes place during contention resolution and is accomplished most typically via a burst dropping policy. The contention resolution usually is assisted by a mechanism such as wavelength conversion, Fiber Delay Lines (FDL) buffering or deflection routing [20]. The following core-based burst mechanisms have been proposed for QoS differentiation in OBS networks.

Preemptive dropping (PD): In PD, when an HP burst arrives at the core node and cannot find a free wavelength in the destination fiber, the resources already reserved for an LP burst are overwritten to accommodate the forthcoming HP burst by means of preempting the LP burst. Several variations of the preemption mechanism can be found in the literature, and both relative and absolute QoS models are supported by this technique. The preemption is of full type when the entire LP burst is preempted [21], whereas

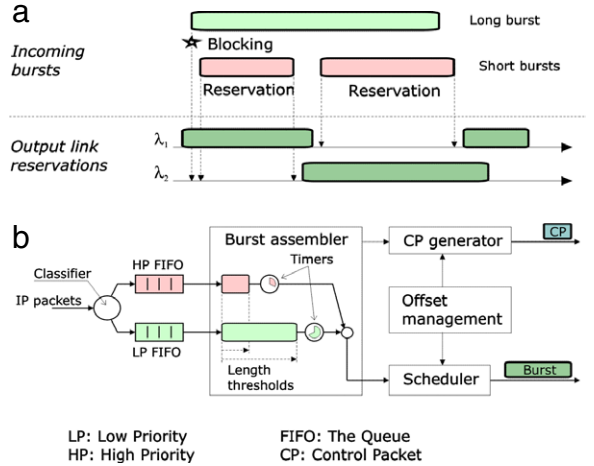


Fig. 3. (a) Illustration of BLD mechanism, (b) Block diagram of the burst assembly unit needed for BLD.




Fig. 4. Full preemption and partial preemption illustrated for two incoming lines and one outgoing line.

in the partial preemption method only the portion of the burst which conflicts with the HP burst is discarded [22], as shown in Fig. 4. Partial preemption allows for more efficient resource utilization compared with the full preemption scheme. Its drawback, however, is the additional complexity in the burst assembly process, since this technique requires additional information about data segments in the burst to be carried and processed in core nodes. One other drawback of PD is the creation of so-called phantom bursts; burst control packets associated with preempted LP bursts continue to travel towards their destination nodes, reserving resources at each downstream node of the path and thus leading to a waste of network resources.

Threshold-based dropping (TD): TD mechanism provides more resources (e.g. wavelengths, buffers) to HP bursts than LP bursts according to a certain threshold parameter. When the occupancy of the associated resource is above a threshold, LP bursts are discarded while HP bursts are accepted as long as resources are available. Fig. 5 illustrates a burst dropping scheme with a wavelength threshold, which is called Wavelength Threshold-based Dropping (WTD), for a system with four wavelengths per fiber. In this example, the wavelength threshold is two, and LP bursts finding more than two wavelengths of the output link occupied are dropped, whereas HP bursts are admitted as long as one of the wavelengths is free. The downside of WTD is that the throughput of LP bursts may be reduced substantially, especially when the traffic is dominated by LP traffic. Adaptation of the wavelength threshold to the




Fig. 5. Illustration of WTD mechanism.

traffic mix is a viable option but introduces complexity in the core OBS nodes. Moreover, this approach requires partitionability of resources and it cannot be used in a bufferless single wavelength OBS system. Threshold-based algorithms have been proposed for optical packet switching systems with wavelength and buffer thresholds, and similar solutions are applicable to OBS networks as well [20].

Intentional dropping (ID): One can extend a preemption-based QoS differentiation scheme to provide absolute QoS. ID mechanism maintains the performance objectives of HP bursts at certain levels by intentionally dropping LP bursts using an active discarding scheme. In [23], an HP burst preempts an LP burst with a probability p when the incoming HP burst overlaps an already scheduled LP burst. The parameter p is adjusted according to loss rate measurements at the core OBS nodes. By suitable adaptation of this parameter, the burst loss rate for HP traffic can be controlled to wander around a desired loss rate, providing absolute QoS to HP bursts at the expense of performance reduction for LP traffic.

Scheduling differentiation (SD): Another group of mechanisms supporting QoS provisioning in core nodes is based on a queuing and scheduling management of burst control packets. One can properly order the processing of burst control packets so that HP reservation requests can be processed earlier, and consequently they will be more likely served than the LP reservation requests. Some of the proposed burst control packets scheduling mechanisms are based on well-studied methods in IP networks. For instance, in [24] burst control packets are processed based on their priorities, while in [25] a fair packet queuing algorithm, which regulates access to the reservation manager for different classes of services, is employed. A drawback of SD in OBS networks is the increased delay. Moreover, an additional offset-time has to be introduced in order to give time for delaying the burst control packets for reordering purposes.

2.1.3. Edge-core-based feedback mechanisms

The involvement of both edge and core nodes in QoS service differentiation are rather rare. A feedback-based architecture for connection-oriented OBS networks for both congestion avoidance and service differentiation was proposed in [26]. The proposed architecture is based on setting up HP and LP connections between pairs of OBS edge nodes and using an explicit-rate distributed rate control mechanism. In this architecture, resource management (RM) packets, in addition to burst control packets, are sent over the control channel to collect the information about the available bit rates for high and

low priority traffic using a modification of the Available Bit Rate (ABR) mechanism in Asynchronous Transfer Mode (ATM) networks. Core OBS nodes, on the other hand, calculate an effective capacity off-line for each of their OBS interfaces based on their contention resolution capabilities. These nodes then run an online explicit rate allocation algorithm to dynamically allocate the overall effective capacity of the OBS node in a max–min fair fashion to the HP OBS connections using that particular link. While doing so, they observe the rates of HP and LP connections only by using their burst control packets. In this sense, the allocation algorithm is carried out only in the control plane. The remaining capacity from the HP connections, if any, is then allocated again using max–min fairness principles to LP OBS connections. Such a resource allocation mechanism is said to be prioritized max–min fair. Finally, the explicit rate fields of RM packets are filled by the core nodes on their way from the destination back to the source. Receiving back the RM packets with information on these two explicit rates for each of the two OBS connections, a scheduler at the ingress node is used for arbitration among HP and LP bursts destined for probably different egress edge nodes. The overall architecture is called Differentiated ABR (D-ABR) and, in this architecture, service differentiation is achieved without having to use large offset times for high priority traffic and when the resources are unpartitionable. Numerical results in [26] show promise but the proposed architecture requires the implementation of an entirely new protocol for the OBS control plane and a relatively complex scheduler at the ingress edge node for shaping and scheduling overall data burst traffic.

2.2. QoS differentiation with two-way signaling

Some of the mechanisms proposed for QoS differentiation OBS networks with one-way signaling can also be used for two-way signaled OBS networks with some modifications, e.g., burst-length based differentiation, preemptive dropping, threshold-based dropping, etc. There are also QoS differentiation methods that are specifically proposed for two-way signaling schemes, e.g., Efficient Burst Reservation Protocol (EBRP) [27].

Recall that in two-way schemes, a setup message is transmitted from the source to the destination node to reserve resources (bandwidth) for a time duration equal to the burst size. Burst transmission starts only upon the successful establishment of an end-to-end connection. The reservation can be delayed as in one-way schemes. It has been shown that the blocking rate of the setup request increases with the number of hops that the message traverses as well as the burst size [23]. In order to increase the burst acceptance probability for HP bursts, using a variable reservation duration, which depends on the priority class of the burst and the number of traversed hops, is proposed. This duration may exceed the actual burst size and it is communicated to all nodes across the route of the path during connection establishment. Hence, the Reservation Duration (RD) parameter is determined as a variable parameter for each Forwarding Equivalence Class (FEC). Various functions that can be used for selecting

the initial value of the RD field have been proposed in [27], including:

$$RD(T_{data}) = k_i \cdot T_{data}, \quad k_i \geq 1, 1 \leq i \leq C \quad (1)$$

where k_i is the reservation over-provisioning parameter for the priority class i that the burst belongs to, C is the number of priority classes and T_{data} is the burst transmission duration. RD parameter is updated at each node as the setup packet travels along the path. In order to differentiate services and to support QoS, each FEC is assigned a different priority class i and initializes its RD field according to Eq. (1). Assuming that FEC $i-1$ has higher priority than FEC i , we have $k_{i-1} > k_i$. This QoS mechanism uses RD for QoS differentiation in such a way that HP bursts request to reserve resources at time scales longer than their actual duration, and thus they experience a lower blocking probability.

2.3. QoS differentiation with control plane methods

We summarize two types of mechanisms involving control plane operation which can provide service differentiation. On one hand, a hybrid signaling protocol (e.g. [28–30]) that consists of a co-operation of one-way and two-way resource reservation modes can be used to support absolute QoS. In this scenario, the establishment of end-to-end transmission paths can provide guarantees such as no losses and negligible delays inside the network, while the unreserved resources can be used to transmit the best-effort burst traffic.

A hybrid optical transport network (HOTNET) is proposed in [29], where TDM wavelength routing and slotted OBS are integrated. This hybrid optical switching architecture provides two levels of switching granularity, which results in efficient resource utilization. In HOTNET, incoming traffic flows at a source node are buffered according to their QoS requirements. Traffic is then transmitted via either pre-established optical circuits or slotted OBS, where optical circuits are tried first. If there is no sufficient bandwidth over the existing optical circuits, the remaining traffic is transmitted over slotted OBS. To guarantee the QoS requirements of the traffic flows, control plane functions such as traffic measurement, bandwidth provisioning, signaling, routing and wavelength/time-slot assignment are implemented. It is shown through simulation studies that HOTNET achieves high channel utilization while satisfying the QoS requirements of different service classes.

In [30], a hybrid OBS architecture (HOBS) is proposed where one-way burst switching and two-way circuit switching schemes are used cooperatively in order to provide QoS differentiation under a unified control plane. HOBS exploits the idle time between the bandwidth reservation and actual arrival of data, when circuit switching (wavelength routing) is used, in order to send LP optical bursts using one-way signaling. The unified control plane handles reservations for both one-way and two-way traffic. Simulation based analysis shows that HOBS achieves high throughput with a finite worst case delay.

The second type of control plane methods for QoS differentiation were originally proposed for optical packet switching networks (e.g., [31,32], the routing function

can support QoS provisioning). In particular, a properly designed routing protocol may minimize the path lengths for delay-sensitive applications, and even preserve the selection of overloaded parts of the network for loss-sensitive ones, for instance thanks to a deflection routing operation.

Table 1 summarizes the QoS differentiation mechanisms that were overviewed in this paper.

3. Numerical results

In all simulations in this paper, the NSFNET topology, which is composed of 15 nodes and 23 links, is used. Each node is assumed to be both an edge and a core node capable of generating bursts destined to any other node in the network. We assume that per-destination based burstification is used at ingress nodes. The first set of results are reported for OBS networks using one-way signaling in Section 3.1 and the results using two-way signaling are presented in Section 3.2.

3.1. Performance of QoS mechanisms with one-way signaling

In this part, we study the performances of selected QoS mechanisms when one-way signaling is used. We first consider the case of UDP offered traffic in Section 3.1.1. QoS mechanisms are then studied for TCP offered traffic in Section 3.1.2. In all simulations, network links are dimensioned with $W = 8$ channels per fiber and with a transmission rate of 10 Gbps per channel. JET one-way signaling scheme is used and the burst scheduler implements a latest available unused channel with void filling (LAUC-VF) scheduling algorithm [33].

3.1.1. UDP offered traffic

Performances of the following QoS mechanisms are evaluated when UDP traffic is offered to the OBS network:

- Offset time-based QoS differentiation (OTD): The duration of extra offset time assigned to HP bursts in the OTD mechanism is 4 times longer than an average LP burst duration.
- Wavelength threshold-based dropping (WTD): The threshold for the maximum number of wavelengths that can be occupied simultaneously by LP bursts is set to 5. On the contrary, HP bursts are allowed to access the whole pool of wavelengths.
- Preemptive Dropping (PD): A full-preemptive scheme is applied where each HP burst is allowed to preempt LP bursts if there are no free wavelengths available. We assume that the searching procedure starts from a random wavelength and is performed according to a round-robin policy.

It is assumed that the traffic is uniformly distributed between nodes. Each edge node offers the same amount of traffic to the network and the destination for each burst is uniformly chosen among all possible destinations. The offered traffic is normalized to the transmission bit rate and expressed in Erlangs. In our context, an Erlang corresponds to the amount of traffic that occupies W wavelengths. Two

Table 1
Characteristics of QoS mechanisms in OBS.

QoS mechanism	Implemented QoS model	Supported QoS parameter	Advantages	Disadvantages
Offset time-based QoS differentiation	relative	burst losses	<ul style="list-style-type: none"> - simple, soft operation - no need for any differentiation mechanism in core nodes 	<ul style="list-style-type: none"> - sensitivity of HP class to burst length characteristics - extended pre-transmission delay
Burst length-based QoS differentiation	relative	delay/burst losses	<ul style="list-style-type: none"> - assembly parameters can be easily setup 	<ul style="list-style-type: none"> - resulting traffic characteristics may influence network performance - requires complex void filling algorithms
Preemptive dropping	relative/absolute	burst losses	<ul style="list-style-type: none"> - fine class isolation - improved link utilization in schemes with partial preemption - absolute QoS can be achieved with a probabilistic preemptive scheme 	<ul style="list-style-type: none"> - overbooking of resources in consecutive nodes (in case of successful preemption) - additional complexity involved in burst assembly process in case of partial preemption
Threshold-based dropping	relative	bursts losses	<ul style="list-style-type: none"> - can be easily implemented 	<ul style="list-style-type: none"> - efficiency of bandwidth usage strongly depends on threshold adaptability to traffic changes
Intentional burst dropping	absolute	burst losses	<ul style="list-style-type: none"> - can provide absolute QoS 	<ul style="list-style-type: none"> - link utilization may suffer
Scheduling differentiation of burst control packets	relative	burst losses	<ul style="list-style-type: none"> - priority queuing in electrical buffers is a feasible and well- studied technique 	<ul style="list-style-type: none"> - extended delay (need for longer queuing windows and thus larger offset times to perform effectively)
Differentiated available bit rate	relative	burst losses	<ul style="list-style-type: none"> - class isolation achieved - more complex priority models than strict priority can also be enforced 	<ul style="list-style-type: none"> - requires a new rate control protocol and advanced schedulers at the edge for burst shaping
Efficient burst reservation protocol	relative	burst losses	<ul style="list-style-type: none"> - class isolation achieved 	<ul style="list-style-type: none"> - requires a complex two-way reservation protocol
Hybrid signaling	absolute	delay/ burst losses	<ul style="list-style-type: none"> - absolute end-to-end loss and delay guarantees for HP 	<ul style="list-style-type: none"> - lower statistical multiplexing gain, inefficient usage of bandwidth (fewer resources available for LP traffic)
QoS routing	absolute (delays) relative (burst losses)	delay/ burst losses	<ul style="list-style-type: none"> - introduces QoS guarantees at network level 	<ul style="list-style-type: none"> - controlling burst losses may be challenging (need knowledge about network state)

different types of burst arrival process and burst size distribution are used: (i) a Poisson burst arrival process with exponentially distributed burst lengths (called Poisson traffic model); (ii) Gaussian distributed burst inter-arrival times and burst lengths (called Gaussian traffic model) [34]. Unless otherwise stated, the mean burst duration is 32 μ s (corresponding to a burst size of 40 Kbytes at 10 Gbps). The following figures present results in terms of HP burst loss probability and LP burst loss probability; since the latter is usually at least one order of magnitude higher than the former, the overall losses are almost equal to the LP ones and are not drawn in the figures.

Fig. 6 compares OTD, WTD and PD mechanisms in terms of the burst loss probability as a function of the offered load under Poisson and Gaussian traffic models. In the simulations, HP bursts comprise 30% of the total traffic. As expected (see for example [35]), WTD mechanism presents the highest burst losses, whilst OTD and PD achieve similar performances. The reason for this behavior is that WTD mechanism has effectively fewer wavelengths available for burst transmissions in the output link than the other two mechanisms. Indeed, it provides only 5 out of 8 wavelengths for LP class bursts, while it attempts to serve

the same amount of burst input traffic. The burst loss probabilities for the HP class are also high compared with OTD and PD since HP bursts cannot preempt LP bursts in WTD. Regarding the other two mechanisms, we can see that HP traffic is served more efficiently with PD mechanism than with OTD mechanism. The explanation for this observation can be found in [36], where it is shown that the scheduling operation may be worsened by the variation of offset-times, a feature which is specific to the OTD mechanism. On the other hand, there is some deterioration in the performance of LP bursts with the PD mechanism due to the creation of phantom bursts. This effect becomes more evident at high loads when the amount of the superfluous traffic due to the preempted LP bursts intensifies the probability of burst losses. The performances under Poisson and Gaussian traffic models show similar behavior, with the latter having slightly lower burst losses. This is because the Gaussian model generates bursts with smaller length variations, which improves the efficiency of the burst scheduling [36].

It is worth mentioning that, in the considered long-haul scenario with bufferless nodes and assuming the use of the same burst assembler at the edge, OTD, PD and WTD

Fig. 6. Burst loss probability as a function of the offered load comparing OTD, WTD, and PD mechanisms when 30% of bursts are HP. (a) Poisson traffic model, (b) Gaussian traffic model.

mechanisms do not introduce any significant delay to the burst. Eventually, the issue of additional delay resulting from the extra offset time in the OTD mechanism needs some comment. The extra offset time is set to 4 times larger than the mean LP burst duration and it equals 128 μ s. Such delay is quite low, when compared with the propagation delays in a fiber link (about 1 ms per 200 km). On the other hand, if long bursts of the duration of some milliseconds were transmitted in the network, the additional delay in OTD might be significant.

Fig. 7 compares OTD, WTD and PD in terms of the burst loss probability as a function of the percentage of HP load and the number of wavelengths per fiber under the Poisson traffic model. In Fig. 7(a), the ratio of HP traffic is varied with an offered load of 0.3 and 8 channels per fiber, whereas in Fig. 7(b), the number of channels per fiber is varied while the offered load is set to 0.5 and the ratio of HP traffic is set to 30%. Similar to the results in Fig. 6, PD scheme is shown to offer the lowest HP burst losses and WTD presents the worst performance for both classes under all scenarios we studied.

3.1.2. TCP offered traffic

In this section, we evaluate the performance of two edge-based QoS mechanisms for OBS networks with one-way signaling when offered traffic is TCP. To this end, we consider OTD and BLD QoS differentiation mechanisms. The simulation results are obtained using nOBS, which is an ns-2 based simulation tool for OBS networks [37]. In the simulations, it is assumed that there are 10 TCP flows between each node pair in the NSFNET topology for a total of 2100 flows. Half of the 10 TCP flows are of the HP class, and the remaining five of the LP class. We use per-destination based burstification with timer-based assembly. For the OTD scheme, we use a fixed burstification delay of 10 ms for all burstifiers. In order to avoid synchronization between TCP flows occurring when fixed timers are used, a Gaussian random variable with a standard deviation of 5 μ s is added to the fixed burstification delay of 10 ms.

In Fig. 8(a), the average normalized TCP goodputs (where the averages are taken over all HP and LP TCP flows, respectively, and then normalized with respect to the average overall goodput achieved when there is no QoS offset) for HP and LP TCP flows are plotted as the extra QoS offset changes when OTD mechanism is used. In Fig. 8(b), burst loss rates for HP and LP bursts are depicted again as the extra QoS offset increases. The average TCP goodput for HP flows first increases and then decreases as the QoS offset increases. Although the burst loss rate monotonically decreases for HP bursts as the QoS offset increases, the decrease in TCP goodput occurs due to the delay penalty caused by the excessive QoS offset. The average TCP goodput for LP flows first decreases with the increasing QoS offset, but then increases as the QoS offset becomes excessive and the total traffic generated by HP flows becomes low due to the reduced goodputs of HP flows. Meanwhile, burst loss rates for LP bursts first increase and then decrease as the QoS offset increases since less HP bursts are generated when the QoS offset is excessively large.

In Fig. 9, the performance of BLD mechanism is evaluated as the burstification timeout for HP bursts changes. In these simulations, LP bursts have a fixed burstification delay of 50 ms, whereas HP bursts have shorter burstification delays so that HP TCP flows enjoy shorter round-trip delays. Furthermore, shorter HP bursts take advantage of burst length dependent losses against longer LP bursts, arising as a result of using LAUC-VF void-filling scheduling algorithm. In Fig. 9(a), the normalized goodputs (normalized with respect to the average goodput achieved by HP flows when the burstification timeout is 1 ms) are plotted as a function of the burstification timeout for HP flows. As the burstification timeout for HP flows increases, the round-trip delays and the average burst lengths for HP TCP flows increase. Consequently, the average goodput for HP flows decreases since TCP flows have longer round-trip delays and longer HP bursts cannot fit into voids formed by earlier reservations. Furthermore, the average burst loss rate seen by HP bursts decreases since the amount of competing HP bursts

