

QoS-enabled Distributed Access on Optical Burst-Switched Networks

Joan Triay*, Georgios Zervas†, Cristina Cervelló-Pastor*, Reza Nejabati†, and Dimitra Simeonidou†

*Department of Telematics Engineering, Universitat Politècnica de Catalunya (UPC)

Esteve Terradas 7, 08860 Castelldefels, Spain

†School of Computer Science and Electronic Engineering, University of Essex

Wivenhoe Park, CO43SQ Colchester, United Kingdom

{joan.triay, cristina}@entel.upc.edu, {gzerva, rnejab, dsimeo}@essex.ac.uk

Abstract—Optical Burst Switching (OBS) has been proposed as an efficient technology to cope with the ever increasing bursty traffic on the Internet. The emergence of a broad range of network-driven applications (e.g. multimedia, online gaming) introduces the need for a network environment able to provide service with different Quality of Service (QoS) guarantees. Burst loss probability, as one of the main QoS metrics on bufferless OBS networks, plays a considerable role on providing a network service with a particular level of QoS. In this paper, a novel MAC based protocol is proposed to support multiple QoS levels on OBS networks without wavelength conversion. It has been evaluated through extensive simulations and its simplistic form makes it very promising. Results indicate that the protocol can clearly provide a relative differentiation of services and guarantee null (or negligible) burst loss probability for a wide network (or offered) load range and, finally, low access delay for the highest priority traffic.

I. INTRODUCTION

Triple-play services (i.e. data, voice and video) and the new deployment of multimedia web-based services have increased the amount of bursty traffic on the Internet as well as the demand for the provisioning of Quality of Service (QoS) differentiation schemes in order to ensure their quality. The traffic diversity created by such services together with the rapid advance of optical technologies has driven the development of new optical network architectures able to provide flexible and dynamic resources allocation [1]. An interesting technology option for the future optical Internet is Optical Burst Switching [2]. OBS can satisfy the future bandwidth requirements avoiding the inefficient resource utilization of Optical Circuit Switching (OCS) and the requirements of Optical Packet Switching (OPS) in terms of optical buffers, fast processing and implementation complexity.

In optical burst-switched networks without wavelength conversion there can exist a great number of burst losses in the core of the network, specially if no other means of contention resolution are present, such as Fiber Delay Lines (FDLs). In view of this, effective contention resolution is critical on OBS networks to restrict losses to a reasonable low level. There are two main losses to address: burst collisions and overlaps. While collisions can be avoided if we restrict or assign the use of wavelengths using a virtual tree-based lambda topology, the use of FDLs or modified offset-based schemes can help

reducing the overlaps. In such a scenario, we are interested on pursuing for other methods to avoid burst switching losses.

Medium Access Control (MAC) protocols provide a means to control the transmission onto the channel depending on the state of the network, becoming a very interesting solution to provide loss-free schemes under specific network conditions. An interesting approach is the IEEE 802.6 Distributed Queue Dual Bus (DQDB) [3], a MAC standard that came up in the early 90's which provides a queue-arbitrated access to the channel guaranteeing zero losses for transmitted frames. In this sense, in [4] an adaptation of DQDB for OBS networks (DAOBS) was introduced. The original standard was enhanced by adapting it over mesh network topologies and integrating it as a wavelength-aware MAC for OBS networks. In DAOBS, a set of counters keep track of the requests for free slots and their occupation using request and burst control packets. With this information, bursts are only transmitted when there are free resources, hence overlapped bursts losses can be avoided. Results in [4] show an interesting burst loss rate performance at the expense of slightly increasing the burst access delay.

Furthermore, QoS differentiation on OBS networks is also an important issue. As the number of traffic flows in the core of the network can grow up to hundreds of thousand, a QoS approach entirely based on a per-flow differentiation would be cost-effective to manage. Thus, most of the schemes proposed for OBS are based on a per-class basis [5]. Additionally, in the per-class approach, QoS parameters can be differentiated as absolute or relative guarantees. In the first group a given class is guaranteed to experience a fixed quality in terms of a certain QoS parameter, e.g. an upper bound of the burst blocking probability. In the relative guarantee however, the quality of each class is qualitatively or proportionally guaranteed between classes. Certain types of QoS techniques applied on traditional store-and-forward electronic networks are no longer the best way to provide service differentiation in OBS unless we accept to loose the optical data transparency. Thus, other types of techniques need to be applied.

In this paper, we give an insight into an enhanced QoS MAC protocol for OBS networks that provides zero burst losses in the core of the network for any class of service, hence losses only occur at the edge of the network. Differentiation among classes is provided using a multi-queue system with priorities

in the channel access module. In addition, the distributed nature of the protocol permits that higher priority bursts from a node can also preempt lower priority ones from other nodes on the network and be placed ahead in time to be transmitted.

The remainder of the paper is as follows. Section II gives a short introduction about different existent QoS architectures for OBS networks. In section III, the proposed MAC protocol is introduced and its QoS enhancements described in detail. Results through simulations are analyzed in section IV, and finally section V presents the conclusions of the paper.

II. BACKGROUND ON QOS FOR OBS

In Optical Burst Switching, QoS differentiation can be achieved using many different strategies to guarantee several QoS parameters. Regarding burst loss rate (BLR) differentiation, offset-based schemes [6] rely on the fact that bursts with a greater offset time should see more free resources on the network and be scheduled for switching with greater probabilities than those bursts with a smaller offset time. Thus, high priority bursts are given a greater offset. Although this scheme increases the delivery rate of high priority bursts, it also increases the latency of them, and for this reason, it is not a suitable scheme when both, delay and BLR need to be guaranteed at a reasonable level.

In the preemption-based QoS scheme [5], high priority bursts are able to take over (preempt) the resources occupied by low priority bursts, while these last ones can never preempt high priority bursts. Hence, on average, high priority bursts see more available resources which results in a lower BLR. Besides, a probability parameter can be used to adjust the differentiation level between QoS classes.

Resource restriction-based schemes exclusively reserve a subset of the available resources for high priority traffic only. An example is the Wavelength Grouping [7] that can establish pre-reserved wavelengths for high priority bursts, and only be used by them even though the wavelengths are available for low priority traffic. Intuitively, the more wavelengths reserved for high priority traffic, the lower its average BLR will be.

With respect to delay-based service differentiation, many different schemes can also be applied. Nonetheless, the results and applications strongly depend on the burst assembly process [8]. In this sense, large bursts become the main delay component and can sometimes make negligible the propagation delay or offset delay. On the other hand, if bursts are small, in the order of hundreds of kbytes, offset delay and propagation are the two main end-to-end latency components.

III. QOS-ENABLED DISTRIBUTED ACCESS FOR OBS

In this section we will focus on describing the enhancements applied to DAOBS in order to provide QoS differentiation of bursts. For a more complete description of the original DAOBS and its operation the reader is also referred to [4]. One of the advantages of the proposed protocol is its low complexity. Complexity can arise an untimely issue in the design process of a new QoS service scheme and may not

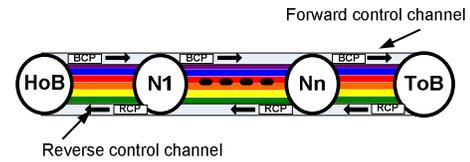


Fig. 1. Example of a DAOBS bus. If node $N1$ needs to transmit to node ToB , it requests for free slots to HoB , or simply waits for a free slot.

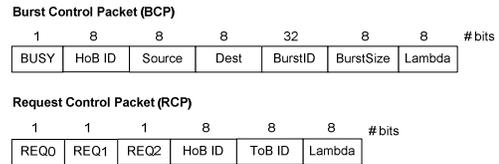


Fig. 2. BCP and RCP packet formats.

justify the increased performance for supporting it. In QoS-enabled DAOBS, the hardware complexity is low. We only need to provide for the additional electronic RAM space to store the bursts of different classes in local and distributed priority queues. Furthermore, the scheduling complexity is again virtually null as the protocol algorithm just works as a set of counters which can easily be either monitored or updated.

Next, we briefly describe the operation of DAOBS and illustrate the enhancements required to provide QoS differentiation.

A. Description of the QoS-enabled MAC protocol

In the present protocol proposal we conceive an optical burst-switched network without wavelength conversion. Furthermore, the network data channel is time-sliced, e.g. slotted OBS, and a constant-based offset scheme is used like in [9]. Although this architecture can increase the hardware complexity of the OBS network node (e.g. FDLs at every input data port), the scheduling and burst processing modules can be much simpler.

The basic piece of network that can manage the distributed protocol is a bus or light-path between two nodes like in Fig. 1. Conceptually, this optical bus has many similarities to a light-trail [10]. However, DAOBS is also a burst MAC with fairness among nodes belonging to the optical bus and ensures the delivery of the burst without using buffers or wavelength converters in the middle. In such a case, we consider two unidirectional control channels which can be in-fiber (i.e. using a specific wavelength): the downstream or forward channel, which forwards Burst Control Packets (BCP) and goes from the Head of the Bus (HoB) node to the Tail of the Bus (ToB) node, and the upstream or reverse channel, on which Request Control Packets (RCP) are forwarded from the ToB to the HoB. The HoB node is responsible for generating and forwarding a BCP on the DAOBS bus at every time slot. The bus is normally composed by several extra nodes between the HoB and ToB. All these nodes can transmit bursts to downstream nodes of this bus according to the operation of the DAOBS protocol, that is, a node requests for free slots to the upstream and HoB nodes in order to transmit bursts to downstream nodes on the bus.

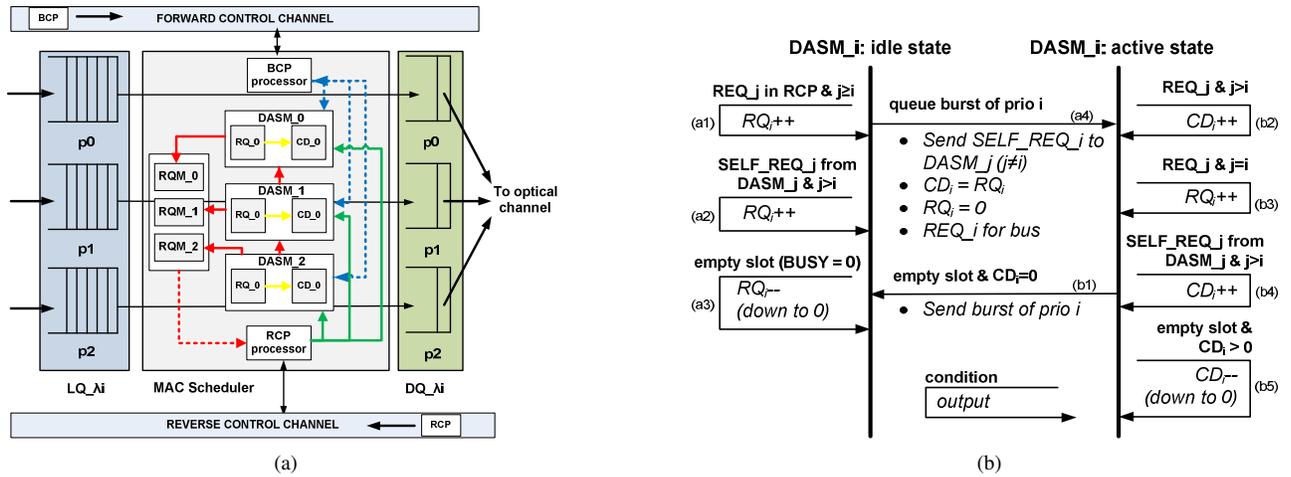


Fig. 3. (a) DAOBS entity with QoS on wavelength λ_i , and (b) DASM flow at priority i .

BCPs have a BUSY bit for announcing the occupation of the upcoming slot, whereas RCPs have a REQ bit for each burst priority class. Fig. 2 shows an example of the packet format for both, the BCP and RCP with three priority classes. In the case the BUSY bit is equal to 1, then the rest of fields in the BCP (*BurstID*, *Source*, *Destination*, etc.) are meaningful since the upcoming slot is occupied by a burst. RCPs let know the rest of upstream nodes the request for a free slot to transmit a burst of a certain priority class.

In the proposed architecture, a DAOBS entity in a network node is identified by its *input port + output port + wavelength*, and a DAOBS bus by its *HoB + ToB + wavelength* on the whole network. Therefore, a node can have multiple entities depending on the number of wavelengths per output port. We should note that this number depends on the DAOBS instantiation method. On a WDM network with wavelength continuity constraint, the instantiation can be done by running a graph coloring algorithm or using a standard optical reservation protocol such as GMPLS. In the present paper, without wavelength conversion, the setup of DAOBS entities follows a tree-based wavelength topology, so that collisions from two different input ports are avoided.

A QoS-enabled DAOBS entity in a node is composed of the following components for each priority class i :

- A distributed access state machine (DASM).
- A request control machine (RQM).
- A local queue (LQ).
- A distributed queue (DQ).

Fig. 3a depicts the relationship between all these elements for a DAOBS entity on wavelength λ_i with three burst priority classes, and the interconnection and relationship between the control, state machines and packet processors.

At each priority, the LQ enqueues bursts waiting to access the optical channel, whereas the DQ is a one-position queue that stores the next burst to transmit on the bus. The collection of DQs from the nodes that belong to the DAOBS bus form the so-called virtual distributed queue, hence when a burst at a certain node gets into one of its DQs is like accessing a global network distributed FIFO queue. At each priority there

is also a request control machine that monitors the requests triggered for that priority. And finally, the distributed access state machine is responsible for monitoring and managing the counting process of the protocol. There is a DASM for each priority. In the example given in Fig. 3a these are $DASM_0$, $DASM_1$ and $DASM_2$, wherein the greater the appended number, the higher its priority. As it is shown in the figure, higher priority DASM can act and trigger changes to lower priority ones in order to ensure that high priority bursts receive a relative better service level.

Each DASM can be in two different states as shown in Fig. 3b. In the *idle* state, the OBS node, for a certain DAOBS bus and priority, has nothing to transmit, whereas in the *active* state, the node has successfully made a request for a free slot to transmit a burst and it is waiting for it. Apart from this, each DASM has two counters: a request counter (RQ) and a count-down counter (CD). On the one hand, the RQ monitors the number of requests made by downstream nodes on the optical bus and by higher priority DASMs in the own DAOBS entity. On the other hand, the CD counts the number of free slots the current node is not allowed to use before being given access to transmit a burst. Initially, all entity counters are reset to zero.

While a DASM at priority i is *idle*, it monitors the RCPs on the reverse control channel and BCPs on the forward one and increases (see Fig. 3b action (a1)) or decreases (a3) the RQ_i for every $REQ_j = 1$ in the RCP of a priority $j \geq i$, and for every $BUSY = 0$ in the BCP, respectively. Similarly, if the $DASM_i$ receives a $SELF_REQ_j$ signal from a higher priority $DASM_j$ ($j > i$) (a2), then the RQ_i is also increased.

As soon as a burst is enqueued into the LQ of a certain DAOBS entity and priority i , if the $DASM_i$ is *idle* it switches to the *active* state following the transition (a4) shown in Fig. 3b. The same happens if after returning from a successful burst transmission there are more bursts to transmit in the LQ. This state transition (a4) triggers the following events: a $SELF_REQ_i$ signal is sent to the remaining $DASM_j$ ($j \neq i$) in the entity; the value of the RQ_i is dumped to the CD_i ($RQ_i = CD_i$); then the RQ_i is reset to 0, and finally a REQ_i signal is sent to the RQM of that same priority in order to set

the REQ bit to 1 in the next upcoming free RCP received from the downstream node.

In the *active* state, the DASM at priority i continues monitoring the BCPs and RCPs. For any REQ bit of priority $j \geq i$ (b2) or *SELF_REQ_j* signal from a *DASM_j* $j \geq i$ (b4), CD_i is increased by 1. These two actions let higher priority bursts on the DAOBS bus to be placed ahead and thus transmitted before, even between bursts from two different nodes of the same bus. In (b3), for every REQ bit of priority $j = i$, the RQ_i is increased by 1. Likewise, for every empty slot on the forward control channel (b5), and while the $CD_i > 0$, the CD_i is decreased by 1 (down to 0). Finally, the transmission of the burst from the DQ at priority i occurs when CD_i reaches 0 and an empty slot is coming from the upstream link. This involves the transition (b1) of the DASM from *active* to *idle*.

B. Routing and Wavelength/DAOBS assignment

The wavelength assignment is responsible for assigning bursts transmissions to a specific DAOBS bus. As stated before, each DAOBS bus is identified by its HoB and ToB network nodes, as well as by its wavelength. Therefore, the wavelength selection is indeed a DAOBS entity assignment.

An algorithm that takes into account the values of the counters RQ and CD and the number of bursts ahead in the LQ is run in this process. As we have previously introduced, the value of these counters and the number of bursts in the LQ determines the position of the node in the distributed FIFO queue at a certain priority class. Depending on these values, the node will take more or less time to transmit a burst on that specific bus, thus increasing or decreasing the channel access delay, and consequently, the end-to-end delay of the burst. In this process, the DAOBS bus that is expected to provide the lowest access delay is always selected. That is, for priority i

$$\min_{\lambda_j \in \mathcal{W}_n^m} (RQ_{i,\lambda_j} + CD_{i,\lambda_j} + sizeLQ_{i,\lambda_j}) \quad (1)$$

where \mathcal{W}_n^m is the wavelength/DAOBS candidates list for transmitting a burst from node n to m .

IV. RESULTS

This section analyzes the performance of the MAC protocol with QoS proposed in this paper. To this end, simulations are conducted on the well-known NSFNET network composed of 14 nodes and 21 bidirectional links. In such a scenario, there are 16 wavelengths with 10 Gbps per channel.

Regarding the setup of hardware devices, the control packet processing time and the non-blocking matrix switching time are set to 10 μs and 5 μs respectively. We assume in all the examples that the network neither has wavelength converters nor FDLs for contention resolution, hence burst transmissions are subject to the wavelength continuity constraint.

With respect to the traffic characteristics, bursts are created at each node following a volume-based algorithm with an input packet arrival Poisson process and fixed size per burst of 100,000 bytes. For simplicity, the burst destination is uniformly distributed to all the remaining nodes of the network,

so that the probability of a burst to be sent to any other node in the network is the same.

Results are gathered using the batch means method. 95% confidence intervals were also obtained, but since they are quite narrow, they have been omitted in order to improve the readability of the graphs.

The following notation is used in the graphs: burst priority classes are numbered from 0 to 2, being *class 2* the highest priority traffic. In order to account the distinct behavior of the protocol under different traffic class distributions, Table I shows the traffic configurations used throughout this evaluation. Every row counts the traffic percentage for that class over the total offered traffic load to the network. These values are representative for a great number of QoS studies.

A. Results without QoS

Initial results deal with the performance of the protocol when different local queue sizes are used and only one traffic class is transmitted on the network. Fig. 4a shows the burst blocking probability (BBP) as a function of the total offered load to the network in Erlang per wavelength (Er/wl). The LQ lengths (in number of bursts) used throughout the simulations are: 2, 5, 100 and 1000 bursts. Intuitively, the shorter the length of the LQ, the sooner the blocking probability starts rising. For sizes between 2 and 100 bursts, the results at high loads asymptotically converge to almost the same value. Only for the case in which the LQ size is equivalent to 1000 bursts we can see a clear improvement ($\sim 50\%$) of the blocking probability but at the expense of increasing nearly two orders its size. Furthermore, Fig. 4b represents the mean access delay (in ms) for the same group of LQ lengths. At the expense of decreasing the mean blocking probability, the delay experienced when the LQ size is of 1000 bursts rises up to nearly 30 ms at very high loads, so there is a tradeoff between BBP and access delay.

B. Results with QoS

Henceforth, we give an insight on the results of the protocol dealing with a number of different burst traffic classes under a worst-case scenario with a LQ size of 5 bursts. Fig. 5a shows again the BBP as a function of the offered load for the two traffic distributions from Table I. Interestingly, we can see that the differentiation between the three burst classes is very clear between the two configurations. In fact, *class 2* blocking probability is not plotted on the graph as in both cases it is zero for the whole load range. Besides, we can see in the second configuration, that when the higher priority traffic load is decreased with respect to the total, the burst blocking probability decreases for both *class 0* and *class 1*. Intuitively, the lower the *class 2* traffic load, the more resources available for the rest of classes, hence the loss probability decreases

TABLE I
TRAFFIC DISTRIBUTION CONFIGURATIONS.

Distribution	Class 0	Class 1	Class 2
Dist. 1	50%	30%	20%
Dist. 2	70%	20%	10%

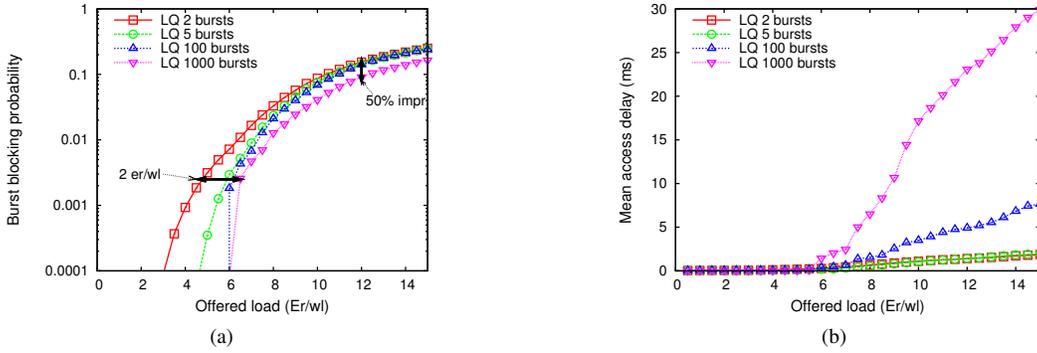


Fig. 4. Results without QoS and for different local queue (LQ) lengths, (a) Burst blocking probability, and (b) mean access delay.

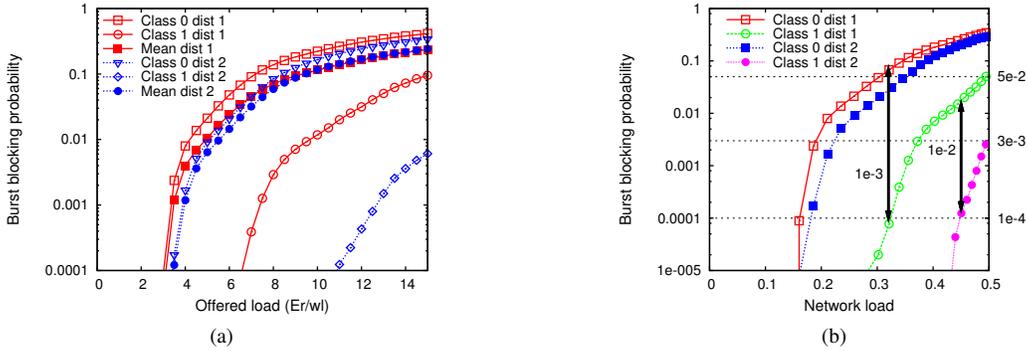


Fig. 5. Results with QoS and different traffic distribution configuration, (a) Burst blocking probability (BBP) comparison between the two traffic configurations with a LQ=5 bursts, and (b) BBP as a function of the network load.

for lower priority traffic. Nevertheless, at high load the mean BBP plots of both configurations tend to converge to the same value which ensures a predictable average BBP performance for both traffic distributions. Therefore, the traffic mix does not affect the overall loss probability.

Fig. 5b shows the blocking probability as a function of the average network load ratio calculated from Eq. 2,

$$\rho = \frac{\sum_{i=1}^n M_i}{\sum_{i=1}^n C_i} \quad (2)$$

where M_i is the traffic carried by link i , C_i is the capacity of link i , and n is the number of links (bidirectional) on the network. The network load is represented between 0 and 0.5 and computed from the link utilization results of the simulation run in Fig. 5a. As it is shown in Fig. 5b, *class 2* is still lossless for the represented network load and that is why is again omitted in the graph. The protocol provides differentiation of up to three orders of magnitude between *class 0* and *class 1* at a load of 0.3, and an improvement of two orders for *class 1* traffics between the two class distributions at load 0.45. Based on the QoS requirements shown on Table II and providing that the packet loss rate (PLR) can be approximated by the BBP for fixed size bursts, the protocol can guarantee the QoS of a diverse number of applications. High or very high loss sensitive traffic (e.g. Grid applications or Live video broadcasting) can be mapped as *class 1* for a large network load range (up to 45% for the second traffic distribution), or even be mapped as *class 2* traffic, which is loss-free.

Fig. 6a compares the mean access delay (in ms) for different burst traffic classes under the two traffic configurations as a function of the offered load. In both configurations, *class 2* bursts not only have the lowest access delay, but also perceive a delay almost flat for the whole load range. *Class 0* bursts have a similar delay trend for both traffic configurations and finally, *class 1* bursts (i.e. the intermediate class) get an access delay totally different depending on the traffic distribution. When *class 2* proportion is 10% and *class 1* represents the 20% of the total traffic on the network, the access delay for *class 1* resembles more the delay of *class 2*. Therefore, the delay experienced by a burst traffic class depends on the aggregate traffic value between itself and all its higher priority traffics.

Fig. 6b shows a completely different performance value. In the graph, the probability that a certain class of traffic is being transmitted from a Head of the Bus (HoB) node is counted.

TABLE II
APPLICATIONS' QoS REQUIREMENTS [11].

Apps.	Delay	Jitter	Loss sens. (PLR)
Interactive audio/video	<150 ms	<75 ms	High (<1e-3)
Inter. transaction data	<50 ms	<10 ms	High (game <5e-2) to very high (grid <1e-4)
Video/audio streaming	<2 s	<40 ms	High (<3e-3) to very high (live video <1e-4)
Legacy applications	Not spec.	Not spec.	Low

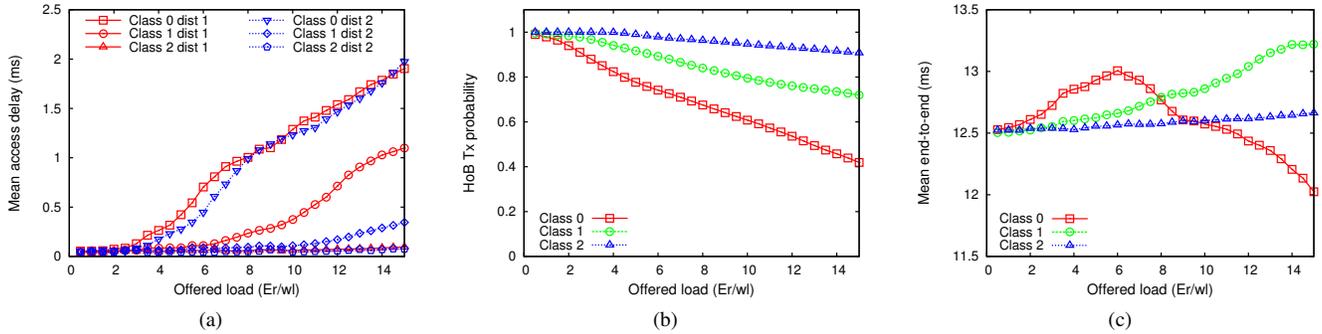


Fig. 6. (a) Access delay, and (b) transmission rate as a HoB for different burst classes, and (c) mean end-to-end delay with traffic configuration 1.

Clearly, we can observe that *class 2* bursts are mainly transmitted from HoBs, whereas the transmission rate for *class 0* and *class 1* decreases almost linearly as a function of the load. Higher priority burst traffic is more often transmitted from a HoB as this node tends to see the capacity channel of the network more available while the requests from downstream nodes are not yet received due to propagation delays.

With respect to the burst end-to-end delay, Fig. 6c shows a comparison between the three classes. In this specific scenario, an average end-to-end delay within a 1.5 ms range is maintained between the three traffic classes still providing a clear differentiation of burst loss probability. Specifically, *class 0* bursts see increased their end-to-end delay as the offered load rises, according to the increased access delay from Fig. 6a. However, we can see a change of tendency from an offered load of 6 Er/wl and on where the delay starts falling. This behavior can be explained as follows: at high loads, *class 0* bursts get a high blocking probability and those that are able to get to the destination tend to follow a shorter path, i.e. less number of hops, hence decreasing the mean delay. This tendency is supported by Fig. 6b, from which we can see that at high loads, most of *class 0* bursts transmissions are made within the DAOBS bus, and therefore closer to the destination node, which at the most can be the ToB.

Altogether, an advantage of the present protocol with respect to other QoS differentiation schemes such as the offset-based, is that high priority traffic, not only gets the best BBP performance, but also obtains the lowest access delay, which can directly improve the end-to-end burst delay. Moreover, its simplicity is also another important feature in comparison with the burst segmentation and scheduling-based schemes.

V. CONCLUSION

This paper has presented a QoS aware Medium Access Control (MAC) protocol for optical burst-switched networks without wavelength conversion. The protocol is based on a counting and monitoring process of burst and request control packets traveling in both control channel directions and a distributed preemption-based scheme in a multi-queue access priority system. Apart from its simplicity and low implementation complexity, the protocol also offers a great performance regarding the relative differentiation of QoS among different

burst traffic classes for both QoS parameters, burst blocking probability (BBP) and end-to-end delay. Indeed, for two different traffic distribution configurations involving, respectively, 20% and 10% of the total load for high priority traffic, highest priority bursts are guaranteed zero losses and very low access latencies. Even for the intermediate traffic class, the protocol may guarantee an acceptable BBP for a diverse number of applications. This differentiation together with the simplicity of the protocol can potentially serve a broad range of services with diverse QoS requirements.

ACKNOWLEDGMENT

This work has been supported by the Euro-NF NoE, the Government of Spain and FEDER through project TEC2009-13901-C02-01, i2CAT Foundation's project TARIFA, and by the Government of Catalonia and the European Social Fund through a predoctoral scholarship (FI) and a BE fellowship.

REFERENCES

- [1] G. Zervas, et al., "SIP-enabled Optical Burst Switching Architectures and Protocols for Application-aware Optical Networks", *Computer Networks*, vol. 52, no. 10, pp. 2065-2076, July 2008.
- [2] Y. Chen, C. Qiao, and X. Yu, "Optical Burst Switching: A New Area in Optical Networking Research", *IEEE Network*, vol. 18, no. 3, pp. 16-23, May 2004.
- [3] IEEE Standards Board, *Distributed Queue Dual Bus (DQDB) Subnetwork of a Metropolitan Network (MAN)- 802.6*, Dec. 1990.
- [4] J. Triay, et al., "On Avoiding-Minimizing Burst Collisions in Optical Burst-Switched Networks without Wavelength Conversion", *Proc. of 11th ICTON*, Azores, Portugal, June 2009.
- [5] H. Overby, and N. Stol, "Quality of service in asynchronous bufferless optical packet switched networks", *Kluwer Telecommunication Systems*, vol. 27, no.2-4, pp. 151-179, 2004.
- [6] M. Yoo, C. Qiao, and S. Dixit, "QoS performance of optical burst switching in IP over WDM networks", *IEEE J. Sel. Areas Commun.*, vol. 18, no. 10, pp. 2062-2071, Oct. 2000.
- [7] Q. Zhang, et al., "Absolute QoS Differentiation in Optical Burst-Switched Networks", *IEEE JSAC*, vol. 22, no. 9, pp. 1781-1795, 2004.
- [8] J.A. Hernández, et al., "On the analysis of burst-assembly delay in OBS networks and applications in delay-based service differentiation", *Photonic Network Communications*, vol. 14, no. 1, pp. 49-62, Aug. 2007.
- [9] M. Klinkowski, D. Careglio, and J. Solé-Pareta, "Comparison of Conventional and Offset Time-Emulated Optical Burst Switching", *Proc. 8th ICTON*, pp. 47-50, Nottingham, UK, Jun. 2006.
- [10] A. Gumaste, et al. "Light-mesh - A pragmatic optical access network architecture for IP-centric service oriented communication", *Optical Switching and Networking*, vol. 5, no. 2-3, pp. 63-74, June 2008.
- [11] IST IP Nobel Phase 2: "Next generation Optical networks for Broadband European Leadership Phase 2". *Project deliverable D1.1: Architectural vision of network evolution*. Aug. 2006.