Photonic Network Communications (2006) 11:163–172
DOI 10.1007/s11107-005-6020-1

ORIGINAL ARTICLE

On burst length distribution and preemption probabilities in OBS networks with burst segmentation

Eduardo Magaña · Javier Aracil

Received: 12 November 2004 / Revised: 3 August 2005 / Accepted 4 August 2005

Abstract Recently, preemption techniques have attracted considerable attention as a means to provide differentiated quality of service in optical burst switching (OBS) networks. This paper is focused on the analysis of preemption probabilities for bursts within the same priority class. As proposed by Vokkarane and Jue ((2003)* IEEE J Select Areas Commun 21(7): 1198–1209) an incoming burst will preempt the burst in service, within the same priority class, if the residual length of the burst in service is smaller than the incoming burst length. For a general case with wavelength conversion, the preemption probability of contending bursts with a generic service distribution, not necessarily exponential, is analyzed. First, we show that the size distribution for the preempting bursts is shifted to larger values, in comparison to the original burst size distribution. Second, we obtain an upper and lower bound for the preemption probability. Finally, the asymptotic behavior of the OBS switch is analyzed showing that preemption will always occur for a very large number of wavelengths.

Keywords Burst segmentation · Optical burst switching · Preemption

E. Magaña (✉)
Dept. de Automática y Computación,
Universidad Pública de Navarra,
31006 Pamplona,
Spain
e-mail: eduardo.magana@unavarra.es

J. Aracil
Dept. de Ingeniería Informática,
Universidad Autónoma de Madrid,
28049 Madrid,
Spain
e-mail: javier.aracil@uam.es

Introduction

The foreseeable evolution of optical networks leads to transfer modes that provide intermediate granularity between circuit and packet switching. Precisely, optical burst switching (OBS) [2] is based on encapsulation of several packets in an optical burst, which is preceded by a burst control packet (BCP). The time offset between the BCP and optical burst allows the optical switches along the path to arrange the switching matrix in order to accommodate the incoming burst. Thus, on the one hand, resources are reserved in advance, in contrast to pure packet switching, and, on the other hand, the "on-the-fly" reservation scheme is unconfirmed, in contrast to pure circuit switching. As a result, chances are that bursts can be dropped due to blocking. However, if free wavelengths along the route from source to destination are available, then the burst is swiftly transferred without leaving the optical domain.

OBS can be easily extended to provide differentiated QoS. Two different schemes for burst prioritization have been proposed: offset-based schemes and preemption-based schemes. The former are based on the principle of granting larger offset times (between burst and BCP) to high-priority bursts. By doing so, high-priority bursts are given more chances to reserve wavelengths, in comparison to low-priority bursts.

In preemption-based schemes, a burst with high priority takes the wavelength from a low-priority burst that is being transmitted. This paper is exclusively concerned with such preemption-based schemes. Usually, preemption comes together with *burst segmentation*. Instead of dropping the entire burst in service when preemption occurs, burst segmentation allows to drop *part* of a burst, so that the remaining packets may continue transmission in subsequent hops. Hence, the burst that losses the contention is segmented and only the burst tail is dropped. Namely, only the burst residual length is affected by preemption. As a result, a significant throughput improvement can be obtained [1].

The preemption-based scheme considered in this paper was proposed in [1]. More specifically, we will *only* consider the case of *bursts with the same priority*. For a bufferless switch with full wavelength conversion capabilities the following policies have been proposed in [1] for bursts *within the same priority class*:

- If the minimum of the residual lengths of the bursts in service is larger than the incoming burst length then the burst in service wins the contention. The incoming burst is dropped.
- If the minimum of the residual lengths of the bursts in service is smaller than the incoming burst length then the incoming burst wins the contention. The burst in service with minimum length is segmented and the tail is dropped.

Such policies have been proposed in [1] and correspond to the drop policy (DP) and segment and drop policy (SDP), for bursts within the same priority class. Note that the application of these policies results in a throughput increase, since the burst with a largest number of pending packets for transmission is selected. Namely, the burst with a largest unfinished work wins the contention.

In this paper, the length distribution of the preempting burst, in comparison to the original burst size distribution is analyzed, with the final goal of deriving the preemption probability, or probability that the incoming burst wins the contention. The paper considers only the same priority class, and the contention algorithm proposed in [1]. The preemption probability is relevant for OBS network engineering for several reasons. First, due to tail dropping upon preemption, gaps in the packet sequence between source and destination may occur. The preemption probability is directly related to the probability of finding a gap in a packet sequence. Second, if preemption occurs, the optical switch must drop the tail of the burst in service and then switch the contending burst to the corresponding wavelength. Furthermore, *trailers* must be generated in the segmenting switch in order to send an indication of the new burst length to the downstream nodes [3]. This implies a processing load that may be crucial for the switch performance. Note that optical switches are usually limited by the so-called "electronic bottleneck." The preemption probability now relates to the processing power that has to be incorporated to the switch.

Network scenario

Figure 1 shows an OBS network architecture. The edge nodes (burstifiers) are in charge of assembling the optical bursts, which are then relayed to the optical core, preceded by the BCP. At the core nodes, BCPs are processed electronically by the switch control unit (SCU). In case of preemption, the SCU configures the optical cross connect (OXC) matrix to switch the incoming burst to a particular wavelength. The burst in service is segmented and the tail is dropped.

Assumptions

In what follows, burst destinations will be assumed to be uniformly distributed. Thus, we focus on a single output port with N wavelengths. Furthermore, full wavelength conversion capabilities will be assumed, i.e., an incoming burst has a choice of several wavelengths (N) to be routed to. Thus, the contention takes place between the incoming burst and the N bursts in service. Note that this case also encompasses the no wavelength conversion case ($N = 1$). It will also be assumed that the switch is bufferless (no Fiber Delay Lines; FDLs) and that the offset between burst and BCP is fixed. Finally, wavelength reservation is performed in a just enough time (JET) fashion [2].

Switching time will be assumed to be negligible in comparison to the burst size. Recent developments in semiconductor optical amplifiers (SOAs) support this hypothesis, with switching times in the vicinity of nanoseconds [4, 5]. Furthermore, this assumption is coherent with the technological trend towards decreasing the switching time. On the other hand, due to preemption, additional signalling is needed in the control channel to advertise the preempted burst modified length. The discussion of such signalling aspects is out of the scope of this paper.

The rest of the paper is organized as follows: second section presents the analysis and third section is devoted to results and discussion. Finally, fourth section presents the ongoing work and conclusions.

Analysis

The analysis focuses on the properties of the preempting burst size distribution. Not only these properties serve to enhance our understanding of burst segmentation but they also provide the foundation to analyze the preemption probability. More specifically, it will be shown that the length distribution of the preempting bursts is skewed to larger values in comparison to the original length distribution. Thus, the bursts that make up the output port busy periods do not come *from the same distribution*.

Burst size distribution for the preempting burst

By original size distribution we mean the size (or length) distribution of bursts *at the input* of the OBS switch. Such bursts may be arriving either from the burstifier (as happens at edge

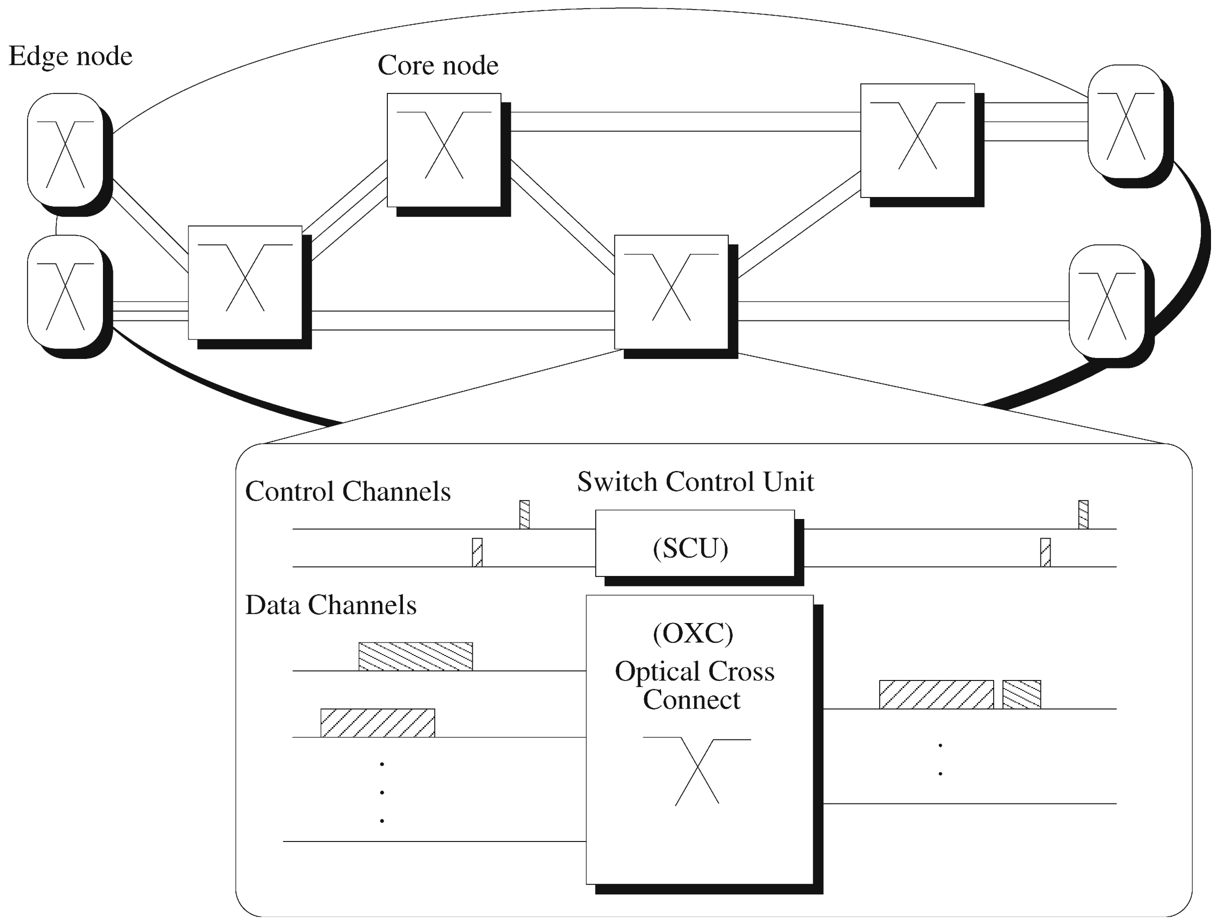

Fig. 1 Network scenario

OBS switches) or from an OBS switch upstream (as happens at core OBS switches). For the latter, truncated bursts from the upstream nodes are expected. However, note that the next theorems apply to *generic* burst size distributions with finite first moment. Thus our findings apply to a fully general case of OBS network and traffic distribution.

Let us denote by (t_0, l_0) the busy period interval for a N wavelengths output port. The arrival time of the first burst that uses the last available wavelength is t_0. Note that this implies that less than N wavelengths were occupied at time t_0^- and, right after t_0, the output port is fully occupied and contention may occur. The busy interval duration will be denoted by l_0.

Let $(t_i, l_i), i = 1, \ldots, n$ be respectively the arrival times and service times of subsequent bursts that arrive *during* the busy period, i.e., $t_0 < t_i < (t_0 + l_0)$ for all $i = 1, \ldots, n$. It must be noted that burst i wins the contention if

$$l_i > l_0 - \sum_{j=1}^{i}(t_j - t_{j-1}) = l_0 - (t_i - t_0) \qquad (1)$$

and $l_k \leq l_0 - \sum_{j=1}^{k}(t_j - t_{j-1}), k = 1, \ldots, i-1$. Fig. 2 serves to illustrate the introduced notation.

Note that (1) is fulfilled iff the minimum of the residual lengths of the bursts in service at the time of arrival of burst i is smaller than burst i length. *Assume that preemption occurs* and let (t_*, l_*) denote the arrival time and length of the burst that wins the contention. For example, assume that bursts $i = 1, \ldots, n-1$ do not fulfill (1) but burst n does. Then, bursts $i = 1, \ldots, n-1$ are dropped from the switch and burst n wins the contention, i.e., $t_* = t_n$, $l_* = l_n$ and $l_n > l_0 - \sum_{j=1}^{n}(t_j - t_{j-1}), l_k \leq t_0 - \sum_{j=1}^{k}(t_j - t_{j-1}), k = 1, \ldots, n-1$.

It will now be shown that the length distribution of burst (t_*, l_*) is shifted to larger values in comparison to the original burst size distribution. Intuitively, the preempting burst is likely to be a large burst. This is shown in the following theorem.

Theorem 1 *Let (t_0, l_0) be the starting time and the length of a newly started busy period respectively. Let F be the common burst length (or service time) distribution.* Let us assume that preemption occurs *and let (t_*, l_*) represent the arrival time and length of the burst that wins the contention, then $P(l_* > x) > 1 - F(x)$ for all $x > 0$.*

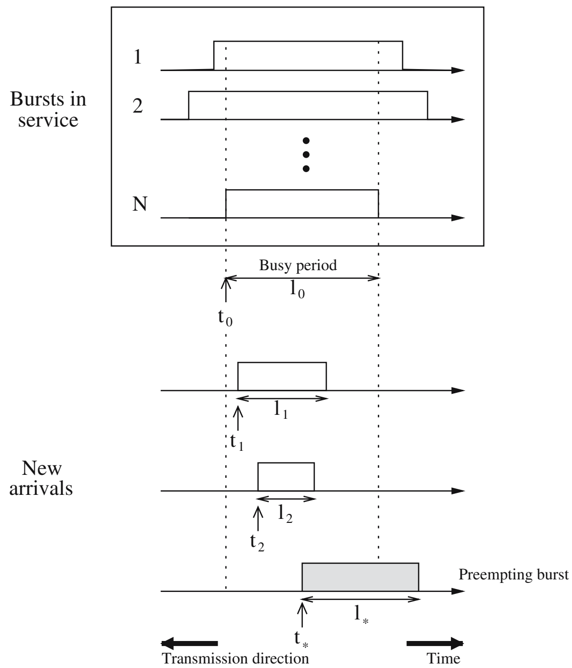

Fig. 2 Notation

Proof Let us consider the denumerable set $\Phi_n = \{(t_1, l_1), \ldots, (t_n, l_n)\}$, $n = 1, 2, \ldots$ of all possible arrivals of n bursts *during* the busy period (t_0, l_0), such that burst n wins the contention. Let L be a random variable with distribution F. First, note that

$$P(l_* > x) = \sum_{n=0}^{\infty} \sum_{\Phi_n} P(l_n > x|\Phi_n) P(\Phi_n)\pi_n, \qquad (2)$$

where the operator $\sum_{\Phi_n}$ represents the sum over all possible sets of n arrivals (Φ_n) and π_n is the probability of n arrivals in (t_0, l_0). Note that π_n is a Poisson measure and $\sum_{n=0}^{\infty} \sum_{\Phi_n} P(\Phi_n)\pi_n = 1$. This is a consequence of considering conditional probabilities to the event that preemption occurs. On the other hand, if

$$\begin{aligned} P(l_n > x|\Phi_n) &= P\left(L > x|L > l_0 - \sum_{j=1}^{n}(t_j - t_{j-1})\right) \\ &> P(L > x|L > 0) = P(L > x), \end{aligned} \qquad (3)$$

then, substitute (3) into (2) to obtain:

$$\begin{aligned} P(l_* > x) > P(L > x) \sum_{n=0}^{\infty} \sum_{\Phi_n} P(\Phi_n)\pi_n &= P(L > x) \\ &= 1 - F(x) \end{aligned} \qquad (4)$$

and the theorem is proved. □

Note that Theorem 1 applies *to any burst length distribution F with finite first moment*, regardless of possible truncation in the nodes upstream. The theorem provides a characterization of the busy period for an incoming burst that contends with the bursts in service. Assume that arriving bursts to a switch have size distribution F. For the access switch case, F depends on the burst assembly mechanism that is adopted at the network edges [6]. For the core switch case, F depends on the original burst size distribution produced by the burstifier and also on possible truncation happening in nodes upstream. Theorem 1 states that, *if preemption occurs*, the next incoming burst will not find a busy period composed by F-distributed bursts, since preempting bursts have a size distribution that is shifted to larger values. Actually, Equation (3) implies that l_* is not distributed according to F but, on the contrary, the distribution of l_* provides larger service time values.

Next, we will show that the preempting burst size distribution converges to the original distribution as the number of wavelengths N increases. Let A_N be the output port residual life of the system (N wavelengths). This is defined as the port blocking time for a randomly arriving burst. Let us assume that preemption occurs and let us consider the output port busy period (t_0, t_0). Wavelengths will be fully occupied for the duration of the busy period. Let F be the burst length distribution and let (t_*, l_*) be the tuple representing the arrival time and length of the preempting burst.

Convergence in distribution to the original burst size distribution will be proved in three steps. First, it will be shown that the preempting burst size distribution is a decreasing function with the number of wavelengths. Second, a lower bound will be found. Third, convergence in distribution will be proved using the latter results.

Theorem 2 *The preempting burst size distribution $P(L_* > x)$ is an strictly decreasing function with the number of wavelengths N.*

Proof Let us consider the denumerable set $\Phi_n = \{(t_1, l_1), \ldots, (t_n, l_n)\}$, $n = 1, 2, \ldots$ of all possible arrivals of n bursts within a busy period (t_0, l_0), so that burst n wins the contention, i.e., $t_* = t_n$ and $l_* = l_n$. Note that

$$\begin{aligned} P(l_n > x|\Phi_n) &= P(l_n > x|l_n > A_N) \\ &= P(L > x|L > A_N). \end{aligned} \qquad (5)$$

This is a consequence of the fact that these probabilities are conditional to the event that burst n preempts the wavelength, namely the burst n length (l_n, which has the same distribution that L) must be larger than the residual life of the system A_N. Furthermore, it turns out that

$$P(L > x|L > A_N) = 1, \quad x < A_N, \qquad (6)$$

$$P(L > x|L > A_N) = \frac{P(L > x)}{P(L > A_N)}, \quad x \geq A_N. \qquad (7)$$

- Concerning (6), let us consider the probability space (Ω, α, P) and the random variable $R_i(w)$ that represents the residual life of burst i in the busy period. For all $w \in \Omega$ the system residual life can be written as

$$A_N(w) = \min(R_1(w), R_2(w), \dots, R_N(w)), \quad (8)$$

$$A_{N+1}(w) = \min(R_1(w), R_2(w), \dots, R_{N+1}(w)), \quad (9)$$

$$A_N(w) \geq A_{N+1}(w) \quad (10)$$

and, thus, $A_N \geq A_{N+1}$ almost surely. Therefore

$$P(L > x | L > A_N) \geq P(L > x | L > A_{N+1}), \quad x < A_N. \quad (11)$$

- Concerning (7), let R_i represent the residual life of a burst i. Then,

$$P(A_N > x) = \prod_{i=1}^{N} P(R_i > x) \quad (12)$$

and

$$P(L > A_N) = \int_0^\infty P(L > x) dF_{A_N}(x)$$
$$= \int_0^\infty P(L > x) d(1 - P(A_N > x)) \quad (13)$$
$$= \int_0^\infty P(L > x) \sum_{i=1}^{N} f_{R_i}(x) \times \prod_{j=1, j\neq i}^{N} P(R_j > x) dx \quad (14)$$

for $x \geq A_N$. Applying integration by parts,

$$P(L > A_N) = 1 - \int_0^\infty f_L(x) \prod_{i=1}^{N} P(R_i > x) dx \quad (15)$$

and since $\prod_{i=1}^{N} P(R_i > x) > \prod_{i=1}^{N+1} P(R_i > x)$, it turns out that, by the monotonicity of the Lebesgue integral, $P(L > A_N) < P(L > A_{N+1})$. Replacing in (7) we obtain the following:

$$P(L > x | L > A_N) > P(L > x | L > A_{N+1}), \quad x \geq A_N. \quad (16)$$

Both (11) and (16) lead to

$$P(L > x | L > A_N) > P(L > x | L > A_{N+1}), \quad \forall x. \quad (17)$$

Finally, this last equation together with (2) and (5) show that $P(l_* > x)$ is strictly decreasing with N. Note that the theorem is assuming that the residual lives (R_i) of the bursts are independent *but not identically distributed*. Therefore, the theorem applies to a generic case of an OBS switch with different burst size distributions per busy period (namely, original and truncated bursts, for instance). □

The following theorem shows that there is a lower bound for the preempting burst size distribution.

Theorem 3 *$P(l_* > x)$ is lower bounded by $P(L > x)$.*

Proof This is a direct consequence of (5) that can be written as follows

$$P(l_n > x | \phi_n) > P(L > x | L > 0) = P(L > x) \quad (18)$$

Now, substitute this expression in (2) and the theorem is proved. □

Finally, both Theorems 2 and 3 lead to the convergence in probability of the burst size distribution.

Theorem 4 *For all $x > 0$, l_* converges in distribution to F as $N \to \infty$.*

Proof From Theorem 2, $P(l_* > x)$ is an strictly decreasing function. From Theorem 3 there is a lower bound equal to $P(L > x)$. On the other hand, the fact that $P(A_N > x) = \prod_{i=1}^{N} P(R_i > x) \to 0$ ensures that $P(L > A_N) \to 1$ and, thus,

$$\lim_{N\to\infty} P(l_* > x) = \lim_{N\to\infty} P(l_n > x | l_n > A_N) \quad (19)$$
$$= P(l_n > x | l_n > 0) = 1 - F(x), \quad x > 0 \quad (20)$$

and the theorem is proved. □

The previous theorems show that the burst size distribution for the preempting bursts is shifted to larger values and that convergence to the original burst size distribution is observed as we increase the number of wavelengths. Now that we have a clear understanding on the burst size dynamics for preempting and non-preempting burst distributions let us resort to the analysis of the preemption probabilities.

Preemption probabilities for $N \to \infty$

Note that the preemption probability is only affected by the system dynamics during the busy period. The following limit theorem states that preemption will occur almost surely as N tends to infinity. Nowadays, the number of wavelengths is increasing with the development of optical technology. Thus, this limit theorem is consistent with the foreseen evolution of optical technology. We wish to evaluate the preemption probability $P(L > A_N)$ where L is the burst size and A_N is the system residual life. Thus, our starting point is the complementary distribution of the system residual life

$$P(A_N > x) = \prod_{i=1}^{N} P(R_i > x), \quad (21)$$

where R_i is the residual life for each of the bursts that made up the busy period. Note that the R_is are independent but not identically distributed. Actually, bursts from different size distributions may coexist within the same busy period. This is the case for previously truncated bursts and for preempting bursts. In fact, Theorem 1 shows that the preempting bursts size distribution is skewed to larger values in comparison to the burst size distribution at the input of the OBS switch.

Theorem 5 *For any burst size distribution with finite first moment it turns out that* $P(L > A_N) \to 1$ *as* $N \to \infty$.

Proof Use (15) to obtain:

$$\lim_{N\to\infty} P(L > A_N) = \lim_{N\to\infty} \left\{ 1 - \int_0^\infty f_L(x) \prod_{i=1}^{N} P(R_i > x) dx \right\} = 1, \tag{22}$$

since $f_L(x) \prod_{i=1}^{N} P(R_i > x) < f_L(x)$, which is integrable. By dominated convergence (22) is obtained and the theorem is proved. □

The previous theorem shows that preemption will *always occur* as the number of wavelengths increases. This is an important property for OBS switches with burst segmentation and preemption capabilities. The immediate consequence of the theorem is that bursts will be continuously removed (tail dropping) and inserted in the wavelengths. This is similar to thrashing for an operating system and it constitutes a severe limiting factor for OBS performance.

Results

In this section, simulation results are provided that assess the previously derived theorems. An OBS switch with full wavelength conversion capability is assumed. We focus on a single output port with a number of wavelengths equal to N.

It will also be assumed that the product (number of input ports) (number of wavelengths per port) is large, so that the traffic directed to the output port consists of Poisson arriving bursts, despite the possible long-range dependence of incoming traffic [6]. The arrival rate λ is a parameter in order to obtain different utilization factors (low load $\rho = 0.2$ and high load $\rho = 0.8$).

The mean burst size is made equal to $12\,\mu$s, which correspond to a 15 Kbytes file size and wavelength speed 10 Gbps. It has been reported that 15 Kbytes is the average file size in the Internet [7]. Furthermore, burst sizes in the order of μs are usually considered in the literature [3,5], including prototypes [4,8]. Note that switching times in the order of some *ns* can be achieved with SOA technology [4,5], thus making the switching time negligible. On the other hand, the analysis presented in the previous section assumes an scenario with generic burst size distribution (finite first moment). However, for simulation and analysis purposes, a specific burst size distribution must be selected.

Section ‘Burst size distribution for the preempting burst’ provides simulation results that assess Theorems 1–4, with Pareto, Gaussian and Exponential burst sizes. Section ‘Preemption probabilities’ provides an assessment of Theorem 5, that comprises not only simulation but also analytical results. To this end, the Pareto distribution is selected, due to its analytical tractabilitity. The exponential distribution has little interest because the analysis is nearly trivial. On the contrary, the Gaussian distribution case is rather complex and a closed analytical expression cannot be found. Note that the theorems derived in the previous section come with the corresponding proof. In this section, we aim at providing additional results that reassure the validity of the theorems, and also provide further insight into the preemption issue.

The simulation parameters are summarized in Table 1.

A note on burst size distributions

The burst size distribution depends on the burst assembly algorithm which is used at the optical network edge nodes, where the *burst assembly* process takes place. A nearly constant burst size distribution (bytes) is provided by edge nodes that collect a fixed number of packets per burst. A possible alternative are timer-based schemes [9]. Incoming packets are stored in per-destination queues and a timer is started with the first packet in a queue. Upon timeout, packets are encapsulated in an optical burst and relayed to the first optical switch downstream. For such timer-based schemes, we have shown that the burst size turns out to be (truncated) Gaussian-distributed [6]. Other non-Gaussian burst size distributions that have also been considered in the literature are the exponential distribution [10–12], the hyperexponential distribution [13], and the Pareto distribution [13].

Burst size distribution for the preempting burst (Theorems 1–4)

Figure 3 shows $P(l_* > x)$ for different values of N number of wavelengths and utilization factors. Recall that $P(l_* > x)$

Table 1 Simulation parameters (cv is the coefficient of variation and α the decay exponent)

Burst length	Parameters	Utilization
Exponential	Mean = $12\,\mu$s	$\rho = (0.2, 0.8)$
Pareto	Mean = $12\,\mu$s, $\alpha = [1.5, 2]$	$\rho = (0.2, 0.8)$
Gaussian	Mean = $12\,\mu$s, $cv = [0.01, 0.2]$	$\rho = (0.2, 0.8)$

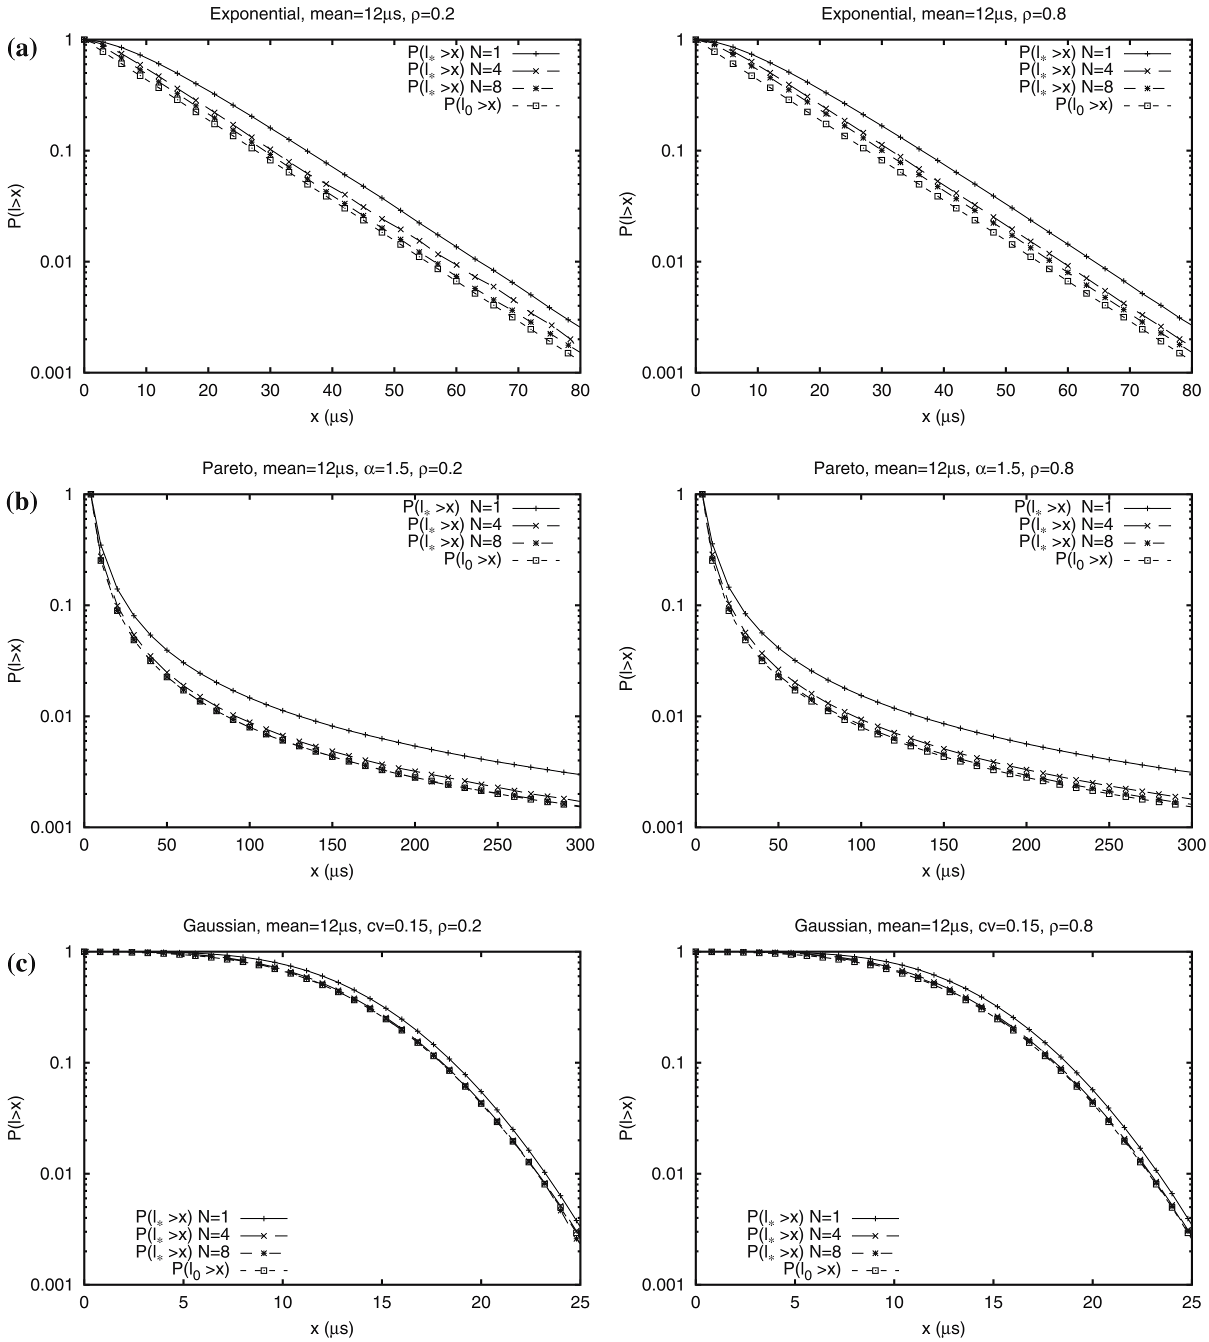

Fig. 3 Comparative between $P(l_* > x)$ and $P(l_0 > x)$ for N wavelengths and burst size distribution **(a)** Exponential, **(b)** Pareto and **(c)** Gaussian $\rho = 0.2$ (left) and $\rho = 0.8$ (right)

is the burst size distribution for the preempting bursts. In Fig. 3 this is compared to the size distribution of the incoming bursts $P(l_0 > x)$. First, Theorem 1 states that the size distribution of the preempting burst is shifted to larger values. Note that for all distributions the curves corresponding to $P(l_* > x)$ are *above* the curves corresponding to $P(l_0 > x)$, as predicted by the theorem. Second, Theorem 2 states that $P(l_* > x)$ is an strictly decreasing function. This is also verified in fig. 3(a–c) Third, Theorem 4 states that l_* converges in distribution to the original burst size distribution. From Fig. 3 we note that $P(l_* > x)$ approaches $P(l_0 > x)$ as N increases.

Preemption probabilities (Theorem 5)

Concerning the preemption probabilities, Theorem 5 states that preemption will occur almost surely as the number of wavelengths $N \to \infty$. In this section, the preemption probability for Pareto-distributed burst sizes is evaluated. Such burst size distribution is chosen since it is analytically amenable, yet realistic for a case of "single file per burst" [14]. On the other hand, while the *incoming* bursts size distribution is assumed to be Pareto, note Theorems 1–4 show that a busy period may comprise different burst size distributions, since preempting and non-preempting bursts may be present. As a result, analytical expressions for preemption probability are very hard to find, if possible. Alternatively, an upper and lower bound for the preemption probability is obtained, for any number of wavelengths N. Both bounds tend to unity as the number of wavelengths increases, thus providing an assessment of Theorem 5.

Upper bound for preemption probability

In order to derive an upper bound for preemption probability the *best case* scenario for preemption should be considered. This best case scenario corresponds to the case of *non-preempting bursts within a busy period*, since they are distributionally shorter than preempting bursts (Theorem 1). Thus, it is easier for the contending burst to win the contention.

Note that the best case implies that $P(A_N > x) = (P(R > x))^N$ (bursts that make up the busy period are independent and *identically* distributed). This is in contrast to the general case $P(A_N > x) = \prod_{i=1}^{N} P(R_i > x)$ assumed in Theorems 2 and 5.

Consider the Pareto distribution

$$P(L > x) = 1, \quad x \leq K,$$
$$P(L > x) = K^{\alpha} x^{-\alpha}, \quad x > K. \tag{23}$$

For Pareto-distributed burst sizes, the residual life for a system of N independent and identically distributed burst is given by [15]

$$P(A_N > y) = \left(\frac{(\alpha-1)(K-y)+K}{\alpha K}\right)^N, \quad 0 \geq y \geq K,$$
$$P(A_N > y) = \left(\frac{K^{(\alpha-1)}}{\alpha} y^{(1-\alpha)}\right)^N, \quad y > K. \tag{24}$$

Thus, if L is the burst length,

$$P(L > A_N) = \int_0^{\infty} P(L > X) dF_{A_N}(x) = \int_0^{K} P(L > X) dF_{A_N}(x) + \int_K^{\infty} P(L > X) dF_{A_N}(x). \tag{25}$$

For the first term in the sum, note that $P(L > X) = 1$ and, thus,

$$\int_0^{K} P(L > X) dF_{A_N}(x) = F_{A_N}(K) = 1 - P(A_N > K) = \frac{\alpha^N - 1}{\alpha^N}. \tag{26}$$

Concerning the second term in the sum,

$$\int_K^{\infty} P(L > X) dF_A(x) = \int_K^{\infty} \frac{K^{\alpha} K^{N(\alpha-1)} N(1-\alpha)}{\alpha^N} x^{-\alpha+N(1-\alpha)-1} = \frac{N(\alpha-1)}{(\alpha+N(\alpha-1))\alpha^N} \tag{27}$$

and, thus,

$$P(L > A_N) = \frac{\alpha^N(\alpha+N(\alpha-1)) - \alpha}{\alpha^N(\alpha+N(\alpha-1))}. \tag{28}$$

Note that the upper bound tends to unity as $N \to \infty$. This is consistent with Theorem 5.

Lower bound for preemption probability

For the lower bound, a *worst case* scenario for preemption should be considered. The worst case scenario is presented in Fig. 4, that shows a simultaneous arrival of bursts to the switch output port. This is the case where all bursts are "new", namely *the residual length is equal to the entire burst length*. Thus, preemption will occur if the contending burst length is larger than the minimum length of the N bursts that make up the busy period.

For the worst case scenario, *should the burst size distribution be the same*, the preemption probability would simply be equal to $N/N + 1$, since the $N + 1$ bursts have the same distribution. However, note from Theorems 1–4 that the burst size distributions *may be different*. Thus, we expect the lower bound to apply only if N is large enough so that the burst size distribution is approximately the same (Theorem 4).

Simulation results

The results for the Pareto distribution are shown in Fig. 5 (upper bound) and 6 (lower bound). Note that the simulation

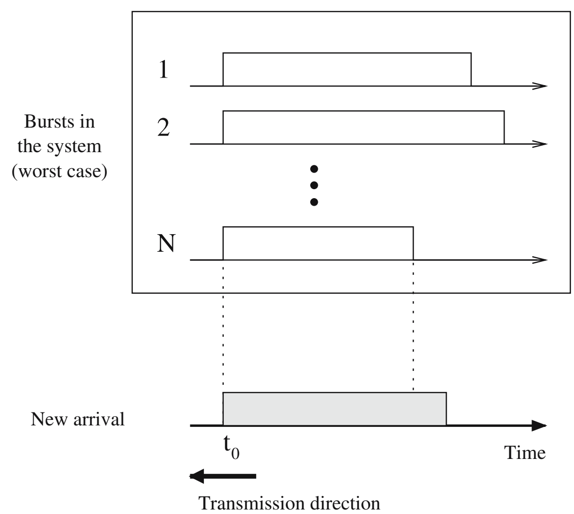

Fig. 4 *Worst case* scenario for the preemption probability

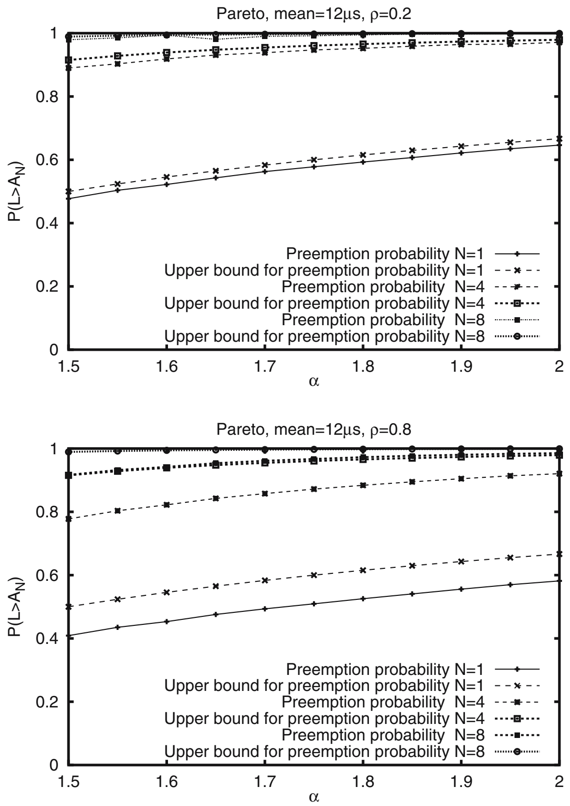

Fig. 5 Upper bound for preemption probability (N wavelengths and Pareto-distributed bursts) ($\rho = 0.2$ (top) and $\rho = 0.8$ (bottom))

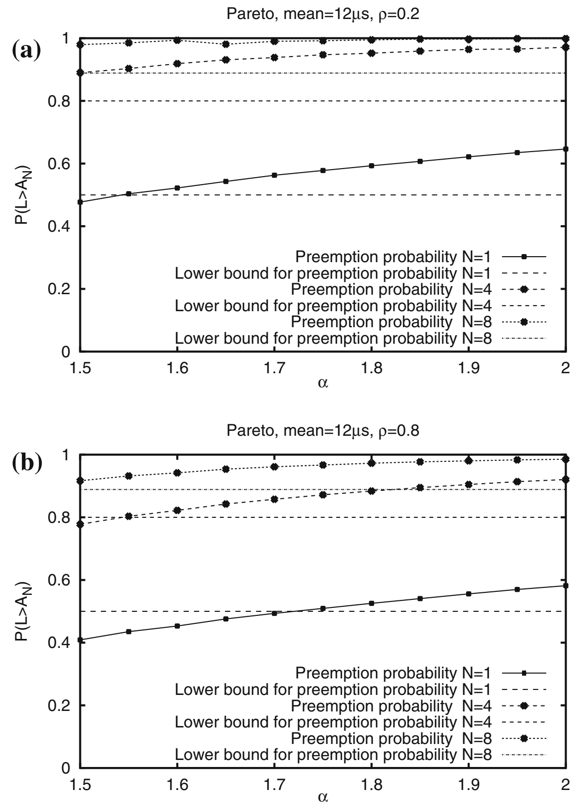

Fig. 6 Lower bound for preemption probability (N wavelengths and Pareto-distributed bursts) ($\rho = 0.2$ (top) and $\rho = 0.8$ (bottom))

and analytical results confirm Theorem 5, since the preemption probabilities tend to unity, as N increases. Furthermore, note that the lower bound applies only with values of N larger than eight. This is a consequence of Theorem 1, that states that the burst size distributions that make up the bundle differ from the original burst size distribution, with a tendency to larger values. On the other hand, as N increases, the burst size distribution converges to the original size distribution (Theorem 4). The worst case scenario depicted in Fig. 4 and the lower bound derivation assume that the size distribution of the contending bursts and the bursts that make up the busy period are the same. This only happens with large number of wavelengths (Theorem 4). Actually, the simulation results provided in Fig. 6 show that the lower bound applies for $N \geq 8$. Current state-of-art prototypes provide 8 [8] or 32 [16] wavelengths per fiber. Thus, our findings apply to current and future OBS switch architectures.

Conclusions and ongoing work

This paper provides a through evaluation of fundamental properties of preemption mechanisms for OBS switches, within the same priority class. It has been shown that the preempting burst size distribution is skewed to larger values in comparison to the original size distribution. However, as the number of wavelengths increases, convergence in distribution of preempting and original burst size is observed. On the other hand, as the number of wavelengths increases, preemption is more likely to occur, eventually happening almost surely in the limiting case of $N \to \infty$. Upper and lower bounds for the preemption probabilities are derived for a case with Pareto-distributed burst sizes. Extensive analytical and simulation campaigns assess the results and provide significant insight on the dynamics of size-based preemption mechanisms, as proposed by [1].

References

1. Vokkarane, V., Jue, J.: Prioritized burst segmentation and composite burst assembly techniques for QoS support in Optical Burst-Switched networks. IEEE J. Select. Areas Commun. **21**(7), 1198–1209 (2003)
2. Qiao, C., Yoo, M.: Optical burst switching (OBS) – A new paradigm for an optical Internet. J. High-Speed Networks **8**(1), 69–84 (1999)
3. Vokkarane, V., Jue, J.: Burst segmentation: an approach for reducing packet loss in optical burst-switched networks. Optical Networks Mag. **4**(6), 81–89 (2003)

4. Ma, X., Kuo, G.: Optical switching technology comparison: optical MEMS vs. other technologies. IEEE Communs. Mag. **41**(11), 16–23 (2003)
5. Buchta, H., Gauger, C., Patzak, E., Saniter, J.: Limits of effective throughput of optical burst switches based on semiconductor optical amplifiers. In: Proceeding of Optical Fiber Communication conference OFC2003, Atlanta, GA, USA, vol. 1, pp. 215–217 (2003)
6. Izal, M., Aracil, J.: On the influence of self-similarity in optical burst switching traffic, In: Proceeding of GLOBECOM'02, Taipei, Taiwan, vol. 3, pp. 2308–2312 (2002)
7. Downey, A.: Evidence for long-tailed distributions in the Internet. In: Proceeding of ACM SIGCOMM Internet Measurement Workshop San Francisco, CA, USA vol. 1, pp. 229–241 (2001)
8. Masetti, F., Zriny, D., Verchere, D., Blanton, J., Kim, T., Talley, J.: Design and implementation of a multi-terabit optical burst/packet router prototype, In: Proceeding of Optical Fiber Communication conference OFC'02, Anaheim, CA, USA vol. 1, pp. 1–3 (2002)
9. Ge, A., Callegati, F., Tamil, L.: On optical burst switching and self-similar traffic. IEEE Commun. Lett. **4**(3), 98–100 (2000)
10. Verma, S., Chaskar, H., Ravikhant, R.: Optical burst switching: a viable solution of terabit IP backbone. IEEE Network **14**(6), 48–53 (2000)
11. Yoo, M., Qiao, C.: Supporting multiple classes of services in IP over WDM networks. In: Proceedings of GLOBECOM'99, Rio de Janeiro, Brazil, vol. 1B, pp. 1023–1027 (1999)
12. Yoo, M., Qiao, C., Dixit, S.: QoS performance of optical burst switching in IP over WDM networks. IEEE J. Select. Areas Commun. **18**(10), 2062–2071 (2000)
13. Dolzer, K., Gauger, C., Spath, J., Bodamer, S.: Evaluation of reservation mechanisms for optical burst switching. Int. J. Electron. Commun. **55**(1), 18–26 (2001)
14. Izal, M., Aracil, J.: IP over WDM dynamic link layer: challenges, open issues and comparison of files-over-lighpaths versus photonic packet switching. In: Proceedings of SPIE OptiComm 2001: Optical Networking and Communications Denver, CO, USA vol. 1, pp. 89–103 (2001)
15. Morato, D., Izal, M., Aracil, J., Magaña, E., Miqueleiz, J.: Blocking time analysis of OBS routers with arbitrary burst size distribution. In: Proceedings of IEEE GLOBECOM'03 San Francisco, CA, USA vol. 5, pp. 2488–2492 (2003)
16. Gripp, J., Duelk, M., Simsarian, J., Bhardwaj, A., Bernasconi, P., Laznicka, O., Zirngibl, M.: Optical switch fabrics for ultra-high-capacity IP routers. IEEE/OSA J. Lightwave Technol. **21**(11), 2839–2850 (2003)

Eduardo Magaña received his M.Sc. and Ph.D. degrees in Telecommunications Engineering from Public University of Navarra, Pamplona, Spain, in 1998 and 2001, respectively. He is an assistant lecturer at Public University of Navarra. During 2002 he was a postdoctoral visiting research fellow at the Department of Electrical Engineering and Computer Science, University of California, Berkeley. In 2005 he has received a national habilitation in the area of Telematics Engineering. His main research interests are network monitoring, optical burst switching and wireless networks.

Javier Aracil received the M.Sc. and Ph.D. degrees (Honors) from Technical University of Madrid in 1993 and 1995, both in Telecommunications Engineering. In 1995 he was awarded with a Fulbright scholarship and was appointed as a Postdoctoral Researcher of the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley. In 1998 he was a research scholar at the Center for Advanced Telecommunications, Systems and Services of The University of Texas at Dallas. He has been an associate professor for University of Cantabria and Public University of Navarra and he is currently a full professor at Universidad Autonoma de Madrid, Madrid, Spain. His research interest are in optical networks and performance evaluation of communication networks. He has authored more than 50 papers in international conferences and journals.