Noname manuscript No.
(will be inserted by the editor)

An Efficient Admission Control Mechanism for Optical Burst-Switched Networks

Igor M. Moraes · Rafael P. Laufer · Daniel de O. Cunha · Otto Carlos M. B. Duarte

Received: date / Accepted: date

Abstract This article proposes the Load-Level-Based Admission Control mechanism (LLAC) in order to provide service differentiation for optical burst-switched networks. The LLAC mechanism admits bursts of a given service class according to the network load and a class-associated parameter. Based on this parameter, called *load level*, the proposed mechanism differentiates the burst blocking probability experienced by each service class. We develop an analytical model for the proposed mechanism and evaluate its performance for different configurations through mathematical analysis. The results show that the load-level-based mechanism reduces the blocking probability of high-priority bursts by two orders of magnitude or more depending on the analyzed scenario. In addition, compared to other similar mechanisms, the load-level-based mechanism effectively differentiates the services in all analyzed configurations, requires less states in OBS nodes, and does not suffer from priority inversion.

Keywords Optical networks · Optical burst switching · Quality of service · Service differentiation

Igor M. Moraes and Otto Carlos M. B. Duarte
Grupo de Teleinformática e Automação - GTA
Universidade Federal do Rio de Janeiro
Rio de Janeiro, RJ, Brazil.
E-mail: igor,otto@gta.ufrj.br

Rafael P. Laufer
Center for Embedded Networked Sensing - CENS
University of California, Los Angeles
Los Angeles, CA, USA.
E-mail: rlaufer@cs.ucla.edu

Daniel de O. Cunha
Laboratoire d'Informatique de Paris 6 - LIP6
Université Pierre et Marie Curie - Paris VI
Paris, France.
E-mail: daniel.cunha@lip6.fr

1 Introduction

Optical Burst Switching (OBS) is an optical data transport technique proposed to ensure the efficient use of the bandwidth offered by Wavelength-Division Multiplexing (WDM) networks [1,10]. Although optical transmissions can reach the magnitude of terabits per second, currently, electronic switches do not achieve more than a few tens of gigabits per second. As the OBS nodes entirely perform the switching task in the optical domain, these nodes do not require the optical-electronic-optical (OEO) conversion as required by most WDM networks nowadays. Thus, the data transport rate of WDM networks is not limited by the OEO conversion.

There are two types of nodes in OBS networks: edge nodes and core nodes. Edge nodes perform the burst assembly process where packets with the same destination address are aggregated to be sent in a single burst. Before the burst transmission, the aggregating node sends a control packet to establish an all-optical path in an out-of-band signaling channel. When the control packet arrives at an OBS core node in the source-destination path, it is electronically converted and processed. If there are available resources, the core node reserves the required resources for the upcoming burst. Otherwise, the burst is blocked. Most signaling protocols proposed for OBS networks do not require error messages or reservation acknowledgments from OBS nodes. It is worth noting that, in OBS, the network resources are only held for the burst switching and transmission time. This is one of the main aspects that differs OBS from optical circuit switching, where an explicit control packet must be sent to tear down the circuit. Furthermore, OBS nodes do not need buffers to store and process bursts as in optical packet-switched networks.

This is an advantage because optical buffers are still experimental.

Quality-of-service (QoS) support is essential in OBS networks. Despite the bandwidth availability, the best-effort service is not able to guarantee the QoS required by new applications [4]. The main problem is that only a few tens of wavelengths are available per optical link nowadays. As a burst occupies one wavelength, future bursts may be blocked depending on the load offered to the network. In addition, the existing QoS mechanisms are proposed for packet switched networks and, at most, are based on management of electronic buffers [18]. To employ these mechanisms in OBS networks, it is necessary to convert the optical signal to the electronic domain at each edge node, which limits the data transport rate. Finally, optical memories are not yet available and bursts can only be delayed by fiber delay lines (FDLs) [16]. Such factors lead to the development of specific QoS mechanisms for OBS networks.

In order to provide service differentiation, we propose the Load-Level-Based Admission Control (LLAC) mechanism for OBS networks [6–8]. This mechanism admits bursts of a given service class according to the network load and a class-associated parameter. Through to this parameter, referred to as *load level*, the proposed mechanism differentiates the burst blocking probability experienced by each service class. The key idea is to define a maximum number of wavelengths for each service class in a link. Based on the Erlang loss model [5,14,16, 17], we develop an analytical model for LLAC. We analytically evaluate the performance of this mechanism and compare it to two previously proposed admission control mechanisms, a static and a dynamic mechanism [17]. We consider the blocking probability experienced by the service classes and the network throughput. Several configurations are tested by varying the offered load, the amount of high-priority traffic, the aggressiveness against the low-priority class, the link capacity and the number of service classes. In comparison with the other two mechanisms, LLAC requires less states in OBS nodes and provides a lower blocking probability for high-priority bursts in all analyzed configurations. Furthermore, LLAC is the only one that remains differentiating the blocking probability experienced by each class, varying the offered load, the amount of high-priority traffic, the number of classes, and the aggressiveness against the low-priority class. As seen in our analysis, it is possible for the static and dynamic mechanisms to have a high-priority class with the same or higher blocking probability than a low-priority class. This phenomenon, known as priority inversion, does not happen for the proposed mechanism.

The remainder of this paper is organized as follows. Section 2 discusses related work concerning quality of service in OBS networks. The load-level-based admission control mechanism is presented in Section 3 and its analytical model is developed in Section 4. Section 5 analytically evaluates the performance of LLAC compared to other similar mechanisms. Conclusions about this article are finally presented in Section 6.

2 Related Work

Several mechanisms have been proposed to address QoS support in optical-burst switched networks. Some approaches are based on the modification of signaling protocols [3, 13, 16] and also on burst preemption [5, 9]. These approaches, however, may increase the end-to-end delay of high-priority bursts because of the extra offset introduced [11], and are also hard to implement in practice. Another approach is to reserve a different number of wavelengths for each service class in a link. Thus, a burst belonging to a high-priority class has a lower blocking probability than a burst of a low-priority class because more resources are reserved to the high-priority class. This is the idea of the admission control mechanisms [15, 17] for OBS networks. Basically, these mechanisms differ in how the wavelengths are reserved for each service class.

Zhang *et al.* [17] propose two admission control mechanisms: a static and a dynamic mechanism. Both are based on the number of wavelengths occupied by each service class. Let W_i be the maximum number of wavelengths that a class i can occupy and $\omega_i(t)$ be the number of occupied wavelengths by class i at time t. A burst of class i will be admitted at time t_0 only if

$$\omega_i(t_0) < W_i. \quad (1)$$

Although these mechanisms have the same admission criterion as in Eq. 1, the wavelength reservation of each class is different for each mechanism. The static mechanism reserves a fixed set of wavelengths W_i, in a given link, for bursts of a given service class i. In other words, if the first W_i wavelengths of a link are reserved for class i, burst of class i can only occupy the wavelengths $C_j \in \{j|1 \leq j \leq W_i\}$. On the other hand, the dynamic mechanism reserves a fixed number of wavelengths W_i (instead of a fixed set) for bursts of a given service class i. Thus, a burst belonging to class i can occupy any available wavelength in a given link, as long as the admission criterion in Eq. 1 is satisfied. In addition, with the dynamic mechanism, high-priority bursts are admitted if there is at least one available wavelength. In other words, let h be the high-priority class, W_h is always equal to the link capacity. Therefore, there is

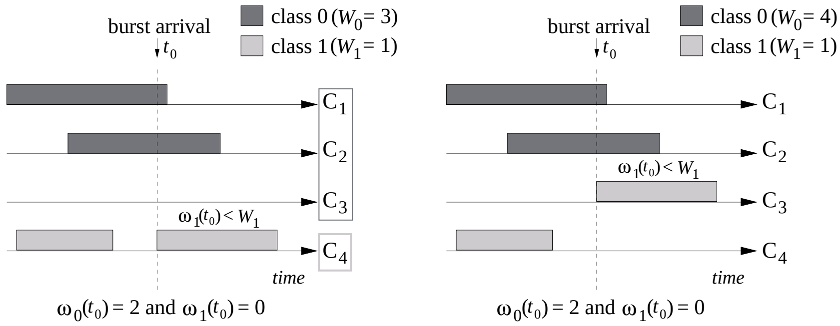

(a) Static mechanism. (b) Dynamic mechanism.

Fig. 1 Example of how the static and dynamic mechanisms work.

no guarantee that bursts belonging to a low-priority class i will be able to occupy all of its W_i authorized wavelengths because all wavelengths of a link might be already occupied by high-priority bursts. Fig. 1 shows an example of how these two mechanisms work for two service classes in a four-wavelength link. Class 0 is the high-priority class. With static mechanism, three wavelengths are reserved for bursts of class 0 ($W_0 = 3$ with C_1, C_2, and C_3 reserved). Bursts of class 1 can occupy only one wavelength ($W_1 = 1$ with C_4 reserved). With dynamic mechanism, bursts of class 0 can occupy any available wavelength ($W_0 = 4$) and bursts of class 1 can occupy, at most, one wavelength ($W_1 = 1$). For the scenario shown in Fig. 1, if a burst of class 1 arrives at time t_0, it is admitted because Eq. 1 is satisfied. With static mechanism, as shown in Fig 1a, this burst must occupy the wavelength C_4. Fig 1b shows that, with the dynamic mechanism, a burst of class 1 can occupy the available wavelengths C_3 or C_4, because there is no specific wavelengths reserved for each class.

For the static and dynamic mechanisms, a node must keep track of the number of wavelengths occupied by bursts of each service class to guarantee that the number of wavelengths occupied by bursts of a given class i does not exceed W_i. Consequently, every node must store a large number of states, which is not desirable [2]. In order to benefit high-priority bursts and also reduce the number of states stored by nodes, we propose an admission control mechanism, described in Section 3, that does not require the knowledge of what service class occupies what wavelength in a given link.

3 Load-Level-Based Admission Control Mechanism

In this section, we describe the Load-Level-based Admission Control mechanism (LLAC). We assume that the network employs a signaling protocol that does not require a positive acknowledgement for sending a burst, for example, JET (Just-Enough Time) or JIT (Just-In Time) [10]. In addition, we consider that each OBS node supports full wavelength conversion and a burst requires only one wavelength during its transmission. The use of a protocol without acknowledgements implies that all network nodes must implement an admission control mechanism. In this kind of protocol, a burst is sent after the control packet without waiting for an acknowledgment. Therefore, when a burst is sent, an edge node cannot guarantee that the number of occupied wavelengths in each link of the source-destination path satisfies the admission criterion. Just after receiving and analyzing the control packet, a node can determine if the number of occupied wavelengths is in accordance with the admission criterion at the moment of the burst arrival. Thus, to guarantee the service differentiation, LLAC must be implemented also by core nodes.

The load-level-based mechanism defines a parameter for each service class i named *load level*, l_i. The load level must be configured at each node of the network and indicates the maximum number of wavelengths that bursts of a given class i can occupy. If we define W as the number of wavelengths in a given link, the inequality $0 < l_i \leq W$ always holds for every class i. The load level is used by LLAC to differentiate the burst blocking probability experienced by each service class. Let $\omega(t)$ be the total number of occupied wavelengths at time t. A burst of a class i arriving at time t_0 is admitted if at t_0

$$\omega(t_0) < l_i. \tag{2}$$

Otherwise, the burst is blocked without sending any error message back to the edge node. Therefore, as the load level of class i increases, its burst blocking probability decreases. It is worth noting that the admission criterion of LLAC is based on the total number of occupied wavelengths, and not on the number of occu-

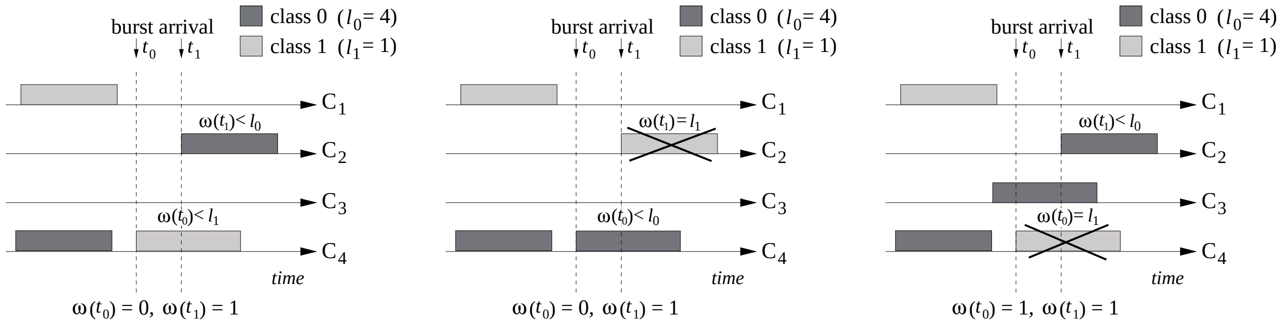

(a) Bursts of both classes are admitted. (b) Bursts of class 1 are blocked in t_1. (c) Bursts of class 1 are blocked in t_0.

Fig. 2 Examples of how the load-level-based mechanism works.

pied wavelengths for bursts of class i, as occurs in the static and dynamic mechanism. This is the key point of the load-level-based mechanism. Instead of considering the number of occupied wavelengths per class, we only care about the total number of occupied wavelengths in the link. Therefore, with LLAC a node stores fewer states than with other mechanisms. The load-level-based mechanism only stores the load level of each class and the total number of occupied wavelengths. Let n be the number of classes and W be the capacity of a link in wavelengths, a node stores $n+1$ states with LLAC. On the other hand, a node stores $n+W$ with static or dynamic mechanisms because they must keep track of the number of wavelengths occupied by bursts of each class.

Fig. 2 shows three examples of how LLAC works for two service classes in one link with four wavelengths ($W = 4$). The high-priority class is class 0. The load level of classes 0 and 1 are $l_0 = 4$ and $l_1 = 1$, respectively. For the situation illustrated by Fig. 2a, when bursts belonging to class 1 and class 0 arrive, respectively, at times t_0 and t_1 they are admitted because Eq. 2 is satisfied. In Figs. 2b and 2c bursts belonging to class 1 are blocked because at its arrival time there is already one wavelength occupied and the load level of class 1 is $l_1 = 1$. In this example, bursts of class 1 are only admitted when no wavelength is occupied at its arrival time. This shows that LLAC is more aggressive against low-priority classes than the static and dynamic mechanisms. These two mechanisms admit bursts according to the number of wavelengths occupied by each service class in a link. The proposed load-level-based mechanism admits bursts based on the total number of occupied wavelengths in a link, no matter what class occupies a given wavelength. Furthermore, in this example, a node stores three states with LLAC. If the static or dynamic mechanisms were employed in this example, a node would store six states.

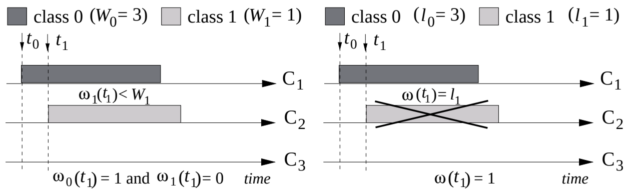

(a) Dynamic mechanism. (b) LLAC mechanism.

Fig. 3 Load-level-based mechanism vs. dynamic mechanism.

Figs. 3a and 3b show the differences of the load-level-based and the dynamic mechanisms for a scenario of one link with three wavelengths ($W = 3$). We assume $W_0 = 3$ and $W_1 = 1$ for the dynamic mechanism and $l_0 = 3$ and $l_1 = 1$ for LLAC. In Figs. 3a and 3b, a burst of class 0 arrives at instant t_0 and is allocated in wavelength C_1. After that, a burst belonging to class 1 arrives at t_1. In the dynamic mechanism, this burst can be allocated in wavelengths C_2 or C_3 because $W_1 = 1$ and there are no wavelengths occupied by bursts of class 1. On the other hand, LLAC blocks bursts of class 1 because $l_1 = 1$ and one wavelength is already occupied by a burst, no matter its class. In this example, LLAC does not block bursts of class 1 in only one situation: when the arrival time of a burst belonging to class 1 is equal to t_0. On the other hand, the dynamic mechanism blocks bursts of class 1 when no wavelength is occupied by another burst of class 1, or when all wavelengths are occupied by bursts of class 0. Moreover, in comparison with LLAC, the static mechanism is also less aggressive against class 1 traffic, because this mechanism always reserves at least one wavelength for this class. The aggressiveness against low-priority traffic is in fact the main advantage of the LLAC mechanism. The probability that a burst belonging to class 0 finds a wavelength occupied by a burst of class 1 at an instant t is small because LLAC admits a lower number of bursts belonging to class 1. As a consequence, the contention for network resources occurs, most of the time, for bursts of class 0. Thus, the burst-blocking probability of class 0 is almost exclusively a function of the amount of high-priority traffic. We show in Section 5 that this choice is adequate and that the proposed mechanism reduces the blocking probability of the high-priority traffic.

4 The Analytical Model

In this section, we present the analytical model developed for the LLAC mechanism based on the Erlang loss model [5,16,17]. We assume that the burst link arrival is a Poisson process with rate λ and the burst transmission time is exponentially distributed with mean $1/\mu$ for all service classes, where μ represents the service rate of one wavelength. In addition, a burst requires the reservation of only one wavelength for its transmission, no matter its service class.

A link is modeled as a *M/M/W/W* queue, where W is the link capacity in wavelengths. As shown in Fig. 4, each link can be represented as a continuous-time Markov chain. Each chain state ω[1] represents the number of occupied wavelengths ($\omega = 0, 1, 2, \ldots, W$).

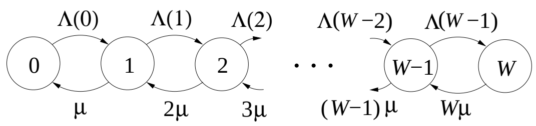

Fig. 4 The state diagram for the LLAC mechanism.

Let n be the number of service classes, λ_i be the arrival rate of bursts of class i offered to a node, and $\lambda_i(\omega)$ be the burst arrival rate of class i offered to a link, after applying the LLAC mechanism. With LLAC, a burst belonging to class i is admitted if the number of occupied wavelengths at the instant of the burst arrival is less than the load level of class i, l_i. Thus, the burst arrival rate of each class i, after applying the admission criterion, is given by

$$\lambda_i(\omega) = \begin{cases} \lambda_i, \text{ if } \omega < l_i \\ 0, \text{ otherwise} \end{cases} \tag{3}$$

The total burst arrival rate offered, $\Lambda(\omega)$, is equal to the sum of the burst arrival rates of the n classes, after verifying the LLAC admission criterion. Then, $\Lambda(\omega)$ is given by

$$\Lambda(\omega) = \sum_{i=0}^{n-1} \lambda_i(\omega), \text{ for } \omega = 0, 1, 2, \ldots, W-1. \tag{4}$$

[1] For simplicity, we consider $\omega(t) = \omega$

From the flow balance equations, derived from the state diagram presented in Fig. 4, the steady-state probabilities π_ω of each chain state ω are calculated and can be expressed by

$$\pi_\omega = \frac{\pi_0}{\omega!\mu^\omega} \prod_{k=0}^{\omega-1} \Lambda(k), \quad \omega = 1, 2, 3, \ldots, W \tag{5}$$

and

$$\pi_0 = \frac{1}{1 + \sum_{j=1}^{W} \frac{1}{j!\mu^j} \prod_{k=0}^{j-1} \Lambda(k)}. \tag{6}$$

The probability B_i of a burst belonging to class i be blocked is the probability of the chain be in a state $\omega \geq l_i$, where l_i is the load level of class i at the time the burst arrives in a given link. Therefore, from Eqs. 5 and 6, we have

$$B_i(\rho_i, l_i, W) = \sum_{\omega=l_i}^{W} \pi_\omega = \sum_{\omega=l_i}^{W} \frac{\frac{1}{\omega!\mu^\omega} \prod_{k=0}^{\omega-1} \Lambda(k)}{1 + \sum_{j=1}^{W} \frac{1}{j!\mu^j} \prod_{k=0}^{j-1} \Lambda(k)}, \tag{7}$$

where ρ_i is the load offered to the network by bursts of class i, which is given by $\rho_i = \lambda_i/(\mu * W)$.

The throughput T in a given link is the percentage of the load offered to the network that is admitted by the admission control mechanisms. Thus, T can be derived from Eq. 7 and is given by

$$T = \sum_{i=0}^{n-1} T_i = \sum_{i=0}^{n-1} \rho_i . \left[1 - B_i(\rho_i, l_i, W)\right]. \tag{8}$$

5 Results

In this section, we compare our load-level based mechanism to the static and dynamic mechanisms according to the load offered to the network, the amount of traffic of each class, the number of service classes, and the number of wavelengths per link. The efficacy of the service differentiation depends on the amount of wavelengths reserved for each service class. The more wavelengths are reserved for high-priority class, the more aggressive is the prioritization of high-priority class against low-priority class. We define level of aggressiveness as the percentage of wavelengths reserved for high-priority traffic and we verify the effectiveness of the proposed mechanism for different levels of aggressiveness against the low-priority class. We use the Tangram-II tool [12] in the analysis and, for the static and dynamic mechanisms, we consider the analytical models proposed by Zhang *et al.* [17]. The analysis considers a single node model, which admits, or not, the offered bursts to a single link. The capacity of each wavelength is 1.0 Gbps and the mean burst size is 128 kB for all service classes, resulting in a service rate $\mu = 1000$ bursts per second.

5.1 Performance for Two Service Classes

The performance of the three mechanisms is evaluated according to the load offered to the network and the amount of high-priority traffic. For this analysis, we assume two service classes: a high-priority class and a low-priority class, respectively, class 0 and class 1. The link capacity in wavelengths is $W = 16$ and, for a fair comparison, we assume that the maximum number of wavelengths that bursts of class 1 can occupy is the same for all mechanisms. As a consequence, the three mechanisms reserve the same number of wavelengths for bursts of class 0: W_0 for the static, $W_0 - W_1$ for the dynamic, and $l_0 - l_1$ for the proposed load-level based mechanism (LLAC). Furthermore, the blocking probability depends on the number of reserved wavelengths for each service class. The blocking probability for a given class decreases the larger the number of wavelengths reserved for that particular class. Hence, we analyze three different configurations varying the number of wavelengths reserved for high-priority bursts, which corresponds to the aggressiveness against bursts of class 1. For the more aggressive configuration, bursts belonging to class 1 can only occupy up to 25% of the wavelengths in a given link. In the intermediary configuration, bursts of class 1 can occupy up to 50% of wavelengths. Finally, in the less aggressive configuration, bursts of class 1 can occupy up to 75% of the wavelengths. Table 1 summarizes the parameters for the three mechanisms. The performance of the static mechanism is not evaluated for the less aggressive configuration, because the number of reserved wavelengths for class 1 would be greater than the one reserved for class 0.

5.1.1 Impact of the Offered Load

In order to evaluate the impact of the offered load on the blocking probability experienced by each service class and on the network throughput, the amount of traffic of classes 0 and 1 is fixed. For this analysis, 30% of bursts belongs to class 0 and 70% belongs to class 1. We analyze the three previously described configurations of aggressiveness against class 1.

Figs. 5 and 6 show the burst blocking probability and the throughput for the three admission control mechanisms and for the network without any QoS support, referred to as classless. For the considered configurations, all the blocking probability curves increase as the load offered to the network increase. Furthermore, the blocking probability curve of the classless network is always between the high-priority class curve, the lower one, and the low-priority class curve, the higher one. This is the conventional effect of the service differentiation.

Table 1 Parameters for the analyzed configurations.

Configuration	LLAC ($l_0 - l_1$)	Dynamic ($W_0 - W_1$)	Static ($W_0 - W_1$)
More aggressive	16 - 4	16 - 4	12 - 4
Intermediary	16 - 8	16 - 8	8 - 8
Less aggressive	16 - 12	16 - 12	-

According to Fig. 5, for all analyzed configurations, the proposed LLAC mechanism provides a lower blocking probability for bursts of class 0, as the load offered to the network increases. For example, in the more aggressive configuration, illustrated by Fig 5a, when the offered load is equal to 1.0 Erlang, the blocking probability of class 0 provided by LLAC is in order of 10^{-5}. For the other two mechanisms, the blocking probability of class 0 is in order of 10^{-3}, a two orders of magnitude reduction. For the same offered load, the blocking probability of class 1 provided by LLAC is only 37% greater than the one provided by the static or dynamic mechanisms. This better performance of LLAC is a consequence of its admission criterion, which takes into account the total number of occupied wavelengths instead of the number of wavelengths occupied by each service class. As the LLAC admits a lower number of bursts of class 1, the probability that a burst belonging to class 0 finds a wavelength occupied by a burst of class 1 at an instant t is small. As a consequence, the contention for network resources is, most of the time, between bursts of class 0.

Fig. 5 also shows that LLAC effectively differentiates the blocking probability experienced by service classes as the aggressiveness against class 1 decreases. On the other hand, the service differentiation provided by the static and dynamic is degraded. For example, in the less aggressive configuration illustrated in Fig. 5c, as the offered load increases, the blocking probabilities of classes 0 and 1 provided by dynamic mechanism tend to the blocking probability of the classless network. On the other hand, employing the proposed load-level-based mechanism, the blocking probability experienced by class 0 is about 120 times less than the one experienced by class 1 in the highest load situation. For the more aggressive configuration, Fig. 5c, the blocking probability of class 1 provided by the dynamic mechanism is greater than the one provided by the static

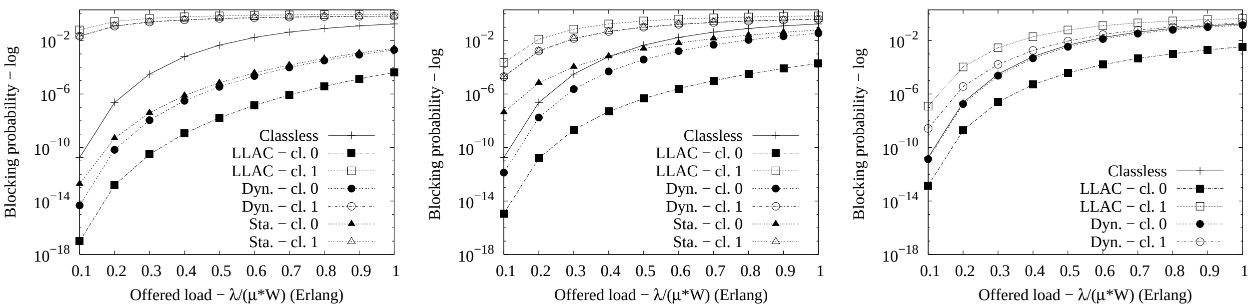

(a) More aggressive configuration. (b) Intermediary configuration. (c) Less aggressive configuration.

Fig. 5 Impact of the offered load: blocking probability.

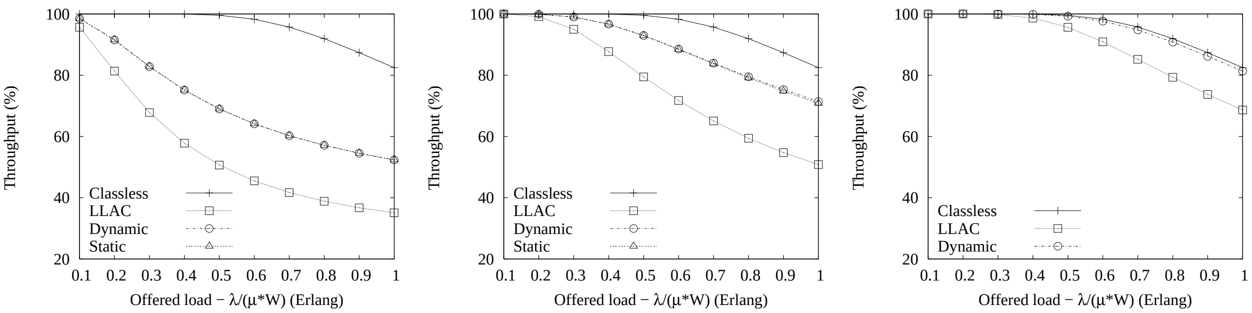

(a) More aggressive configuration. (b) Intermediary configuration. (c) Less aggressive configuration.

Fig. 6 Impact of the offered load: throughput.

mechanism, but, because of the scale, the curves are overlapped.

The better differentiation obtained by LLAC is paid by a reduction of the network throughput, as shown in Fig 6. For all configurations, the best case the classless network curve. For an offered load equal to 1 Erlang, the throughput of the classless network is equal to 82%. For the same offered load in the more aggressive configuration, the throughput is 52% for the static and dynamic mechanisms while for LLAC is 35%. On the other hand, for the intermediary configuration, the throughput for the static and dynamic mechanisms is 71% while for LLAC is 50%. These results ratifies the idea that the less the aggressiveness against class 1, the greater the throughput. Moreover, the difference between the throughput provided by LLAC and the one provided by the other two mechanisms is reduced, as the aggressiveness against class 1 decreases. For example, in the less aggressive configuration, LLAC properly differentiates the services and provides a throughput only 13% lower than the one provided by the classless network.

5.1.2 Impact of the Amount of High-Priority Traffic

The traffic amount of each service class also impacts the performance of the mechanisms. In order to analyze this impact on the blocking probability and on the throughput, the load offered to the network is fixed at 0.5 Erlang. According to the results of the previous section, the mechanisms are effective for this offered load. As in Section 5.1.1, we analyze three different configurations, considering different levels of aggressiveness against class 1.

Figs. 7 and 8 show, respectively, the blocking probability and the throughput for the three mechanisms and for the classless network. The classless curves in these figures correspond to the points $(0.5, 4.5 \times 10^{-3})$ and $(0.5, 99.5)$ of Figs. 5 and 6, respectively. For all configurations, as the high-priority traffic increases, LLAC still provides a lower blocking probability for bursts of class 0 than the other analyzed mechanisms, resulting in a better differentiation for the service classes. The higher the amount of the high-priority traffic the lower the blocking probability of class 0 with LLAC compared to the other two mechanisms because, in this situation, more bursts of class 0 contend to network resources and the proposed mechanism benefits the high-priority traffic. For example, considering the less aggressive configuration, presented in Fig. 7c, the difference between the blocking probabilities provided by LLAC for classes 0 and 1 is approximately five orders of magnitude, when 10% of bursts belong to class 0. This difference remains equal to one order of magni-

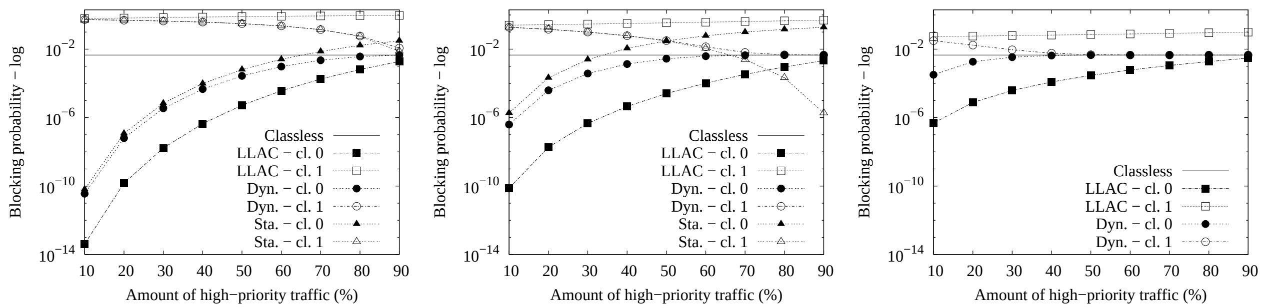

(a) More aggressive configuration. (b) Intermediary configuration. (c) Less aggressive configuration.

Fig. 7 Impact of the high-priority traffic: blocking probability.

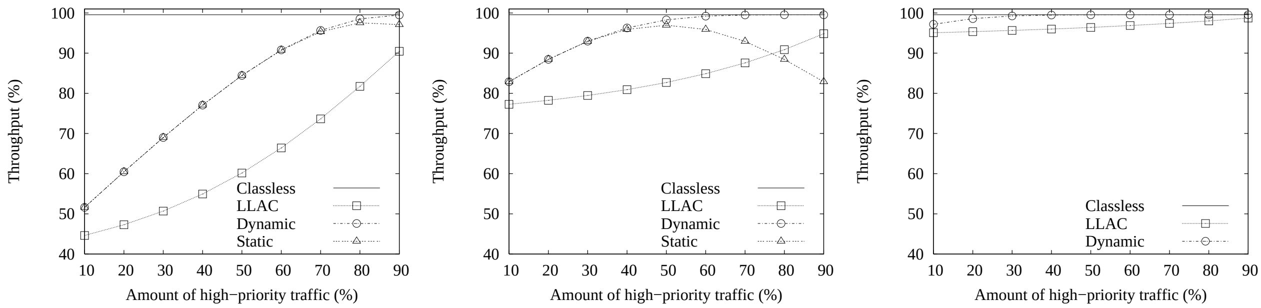

(a) More aggressive configuration. (b) Intermediary configuration. (c) Less aggressive configuration.

Fig. 8 Impact of the high-priority traffic: throughput.

tude even when 90% of bursts belong to class 0. Moreover, for LLAC, as the amount of high-priority traffic increases, the network throughput also increases because LLAC provides a lower blocking probability for bursts of class 0. For example, in the less aggressive configuration represented by Fig. 8c, the throughput provided by LLAC reaches its maximum value when 90% of bursts belonging to class 0. In this situation, with LLAC, the network throughput is equal to 98.74%.

Fig. 7 also shows that the proposed LLAC mechanism is the only one that effectively remains differentiating the services, regardless of the amount of high-priority traffic and the aggressiveness against the low-priority class. For the dynamic mechanism, the blocking probability of both classes and the throughput tend to the classless network as the amount of high-priority class increases. For the more aggressive configuration, illustrated in Figs. 7a and 8a, the blocking probabilities of both classes tends to be equal to 4.5×10^{-3}. Beyond this point the differentiation function is no more effective and the dynamic mechanism works like a classless network. According to Fig. 5a, 4.5×10^{-3} is the same value of the blocking probability for the classless network when the offered load is 0.5 Erlang. As the service differentiation provided by the dynamic mechanism is degraded, this mechanism becomes similar to the classless network and then provides a higher through. Similarly to dynamic mechanism, the performance of the static mechanism extremely depends on the amount of traffic generated by each service class. For the more aggressive configuration presented in Fig. 7a when the amount of bursts of class 0 is equal to or greater than 85% of the total traffic, the blocking probability of class 0 becomes higher than the blocking probability of class 1, a phenomenon known as priority inversion. In addition, for intermediary configuration, represented by Figs. 7b and 8b, when the traffic of class 0 reaches 50% of the total amount, the blocking probabilities of both classes are equal and the throughput is maximum because half of the link capacity is reserved for each class. After this, we can easily recognize a priority inversion.

5.2 Multiclass Performance

In this section, we evaluate the performance of the admission control mechanisms for more than two service classes. For simplicity, we assume that all service classes generate an equal amount of traffic ($\rho_0 = \rho_1 = \rho_2 = \ldots = \rho_{n-1} = \rho/n$) and the link capacity in wavelengths is $W = 16$. The performance of the mechanisms is

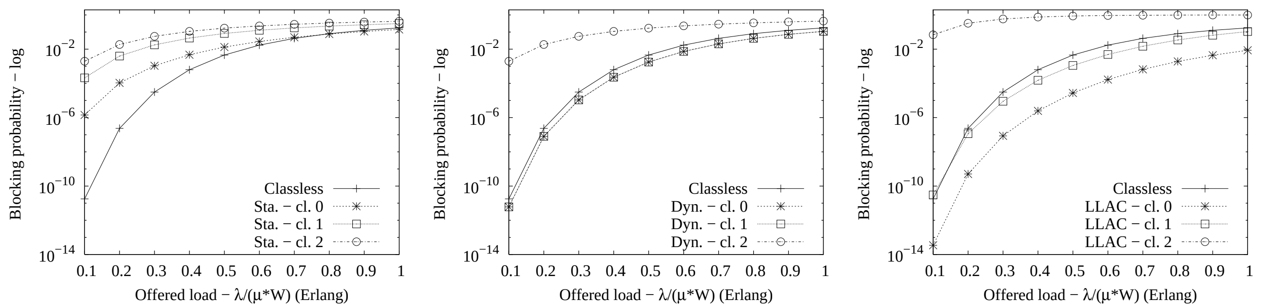

(a) Static mechanism. (b) Dynamic mechanism. (c) Load-level-based mechanism.

Fig. 9 Performance for three service classes.

evaluated according to the offered load to the network. First, we consider three service classes - class 0, class 1, and class 2. Class 0 is the high-priority class. We have $W_0 = 7$, $W_1 = 5$, and $W_2 = 4$, for the static mechanism; $W_0 = 16$, $W_1 = 14$, and $W_2 = 4$, for the dynamic mechanism; and $l_0 = 16$, $l_1 = 14$, and $l_2 = 4$, for LLAC.

As shown in Fig. 9a, the static mechanism properly differentiates the services, but the blocking probability experienced by each one of the three classes is higher than the blocking probability of the classless network. In this situation, a reduction in the offered load and in the link capacity does not imply a reduction in the blocking probability because of the exponential nature of the traffic. For example, the load offered for class 0 represents 33% of the total load and the number of wavelengths reserved for this class is 44% of the total link capacity. Even for this load reduction, the considered configuration ($W_0 = 7$) does not guarantee a lower blocking probability for class 0 than the one provided by the classless network. Therefore, the configuration of the static mechanism becomes harder as the number of classes increases. Furthermore, according to Fig. 9b, the dynamic mechanism does not properly differentiate the services because class 0 and class 1 experience almost the same blocking probability. This is a consequence of the admission criterion of the dynamic mechanism, which is based on the number of wavelengths occupied by each service class and the scenario, which considers the same amount of traffic for all classes. Thus, in this situation, a difference of 2 units between W_0 and W_1 does not differentiate the blocking probabilities of classes 0 and 1. On the other hand, according to Fig. 9c, when the offered load increases LLAC differentiates the blocking probability experienced by each service class and also provides a lower blocking probability for class 0. With LLAC, when the offered load is equal to 0.5 Erlang, bursts of class 0 experience a blocking probability two and four orders of magnitude lower than the probabilities experienced by bursts of classes 1 and 2, respectively. For the same offered load, the static mechanism provides a blocking probability for the high-priority class 6 and 13 times lower than the probabilities provided, respectively, for classes 1 and 2. The blocking probability of class 0, however, is 3 times higher than the blocking probability of the classless network. For the same 0.5 Erlang, the dynamic mechanism provides the same blocking probability for bursts of classes 0 and 1, which is two orders of magnitude lower than the one experienced by bursts of class 2.

Unlike the other two mechanisms, LLAC still differentiates the services for a larger number of classes ($n = 2, \ldots, 7$). Fig. 10 shows the blocking probability of seven different service classes as a function of the number of classes. The curves of each class start in different points of the x-axis as they are introduced in the analysis. In other words, the curve of the i-th class introduced in the analysis starts in i. We define that the load level of the high-priority class is $l_0 = 16$ and the load level of the low-priority class is $l_6 = 4$. The load level of all classes can be viewed directly in Fig. 10. We assume that all service classes generate an equal amount of traffic, the offered load is 0.5 Erlang, and the link capacity is equal to 16 wavelengths.

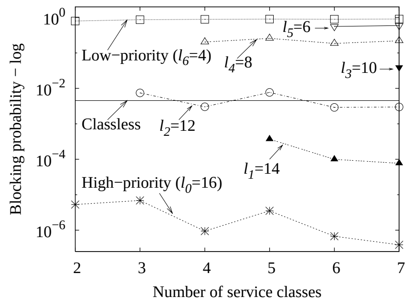

Fig. 10 Impact of the number of service classes.

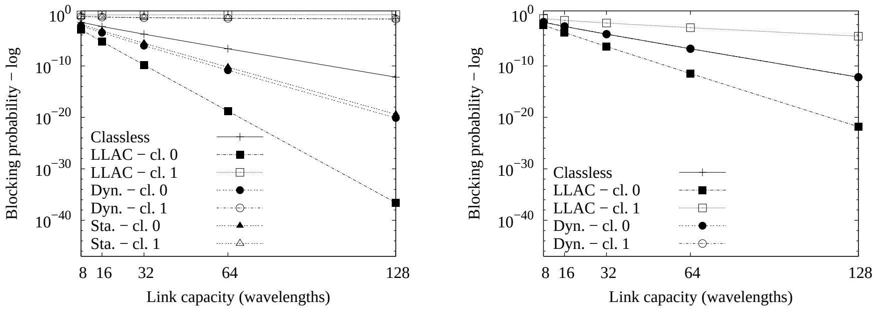

(a) More aggressive configuration.

(b) Less aggressive configuration.

Fig. 11 Impact of the link capacity.

The results presented in Fig. 10 show that the high-priority class experiences a lower blocking probability with 7 classes than with 2 classes. This happens because the higher the number of classes, the lower the amount of traffic generated individually by each service class because all classes generate the same amount of bursts. Thus, the blocking probability of class 0 decreases because less bursts of class 0 contend to wavelengths and LLAC benefits high-priority bursts. In Fig. 10, when class 2 ($l_2 = 12$) and class 1 ($l_1 = 14$) are introduced the blocking probability of class 0 increases compared to its previous value because the load levels of these two classes are closer to the load level of class 0.

The better performance for the high-priority class is paid by a higher blocking probability of the low-priority class. Nevertheless, the aggressiveness against the low-priority class is determined by its load level, which is a choice of the network operator. Thus, the level of priority for the high-priority class and the level of starvation for the low-priority class can be tuned.

5.3 Impact of the Link Capacity

Another factor that impacts the blocking probability experienced by service classes is the number of wavelengths in a link. We evaluate the performance of the three differentiation mechanisms according to the number of wavelengths in the link. We vary the number of wavelengths from 8 to 128. The two classes generate the same amount of traffic and the load offered to the network is 0.5 Erlang. We analyze two different configurations, varying the aggressiveness against class 1.

Fig. 11a shows the results for the more aggressive configuration, where bursts of class 1 can occupy up to 25% of the wavelengths in a link, regardless of the considered link capacity. For the three mechanisms, as the number of wavelengths increases, the blocking probability experienced by class 0 decreases and the blocking probability of the class 1 is almost constant. The LLAC mechanism, however, better differentiates the services, providing a lower blocking probability for bursts of class 0. When the link capacity is equal to $W = 8$, LLAC provides for class 0 a blocking probability two orders of magnitude lower than the blocking probability of class 1. When the link capacity is equal to $W = 32$, this difference increases to ten orders of magnitude. For the static and dynamic mechanisms this difference increases from 1 to 5 orders of magnitude. LLAC provides a lower blocking probability for class 0, but it can starve the traffic of class 1, depending on the number of wavelengths. When the link capacity is equal to or greater than 256 wavelengths, more than 90% of bursts of class 1 are blocked. In Fig. 11a, curves of class 1 for the static and dynamic mechanisms are overlapped.

The results for the less aggressive configuration are presented in Fig. 11b. In this configuration, bursts of class 1 can occupy up to 75% of the wavelengths. The performance of the static mechanism is not evaluated for this configuration because the number of reserved wavelengths for class 1 would be greater than the one reserved for class 0. The results ratify the better performance of the proposed load-level-based mechanism. As the link capacity increases, bursts of class 0 experience a lower blocking probability and class 1 is no more starved. For example, with LLAC, when the link capacity is equal to 32, bursts of class 0 and 1 experience blocking probabilities in order of 10^{-7} and 10^{-2}, respectively. On the other hand, the increase of the link capacity is not enough to avoid the degradation of the service differentiation provided by the dynamic mechanism. When the link capacity is equal to 32, the blocking probability of classes 0 and 1 differ in the sixth decimal place. For this configuration, curves of both classes of the dynamic mechanism and the curve of the classless network are overlapped, because bursts of both classes experience almost the same blocking probability and this configures a classless network.

5.4 The Effectiveness of the LLAC Mechanism

The number of wavelengths reserved for each service class also impacts the effectiveness of the differentiation provided by the mechanisms. The larger the number of wavelengths reserved for a service class, the lower the blocking probability experienced by this class. Thus, when the difference between the numbers of reserved wavelengths for each class is small, it is more difficult for a mechanism to differentiate the traffic. In this section, we evaluate the performance of the proposed Load-Level-Based Admission Control (LLAC) and the dynamic mechanisms when the aggressiveness against class 1 decreases. As stated before, for both mechanisms, bursts belonging to class 0 can occupy any available wavelength. Hence, considering a link capacity of 16 wavelengths, $W = 16$, we have $l_0 = 16$ and $W_0 = 16$. In addition, the load offered to the network is 0.5 Erlang, 30% of bursts belong to class 0, and the number of wavelengths that bursts of class 1 can occupy in a link, l_1 or W_1, ranges from 1 to 15.

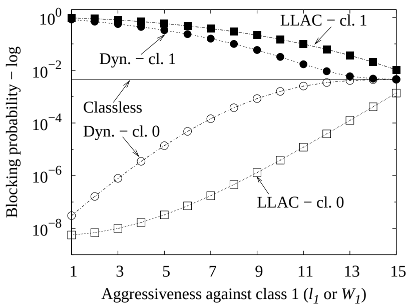

Fig. 12 Impact of the aggressiveness against the low-priority class.

As shown in Fig. 12, LLAC remains differentiating the blocking probability of each service class as the load level of the class 1 increases. Even when $l_1 = 15$, just one unit less than l_0, the two service classes experience different blocking probabilities. The blocking probability for class 0 is in order of 10^{-3} and for class 1 is in order of 10^{-2}. These results reinforce that the LLAC mechanism always differentiate the blocking probability of each class, without starving the low-priority traffic. Furthermore, the level of differentiation imposed on the services is a choice of the network operator. On the other hand, the service differentiation provided by the dynamic mechanism degrades for $W_1 \geq 13$. Therefore, for this configuration, the dynamic mechanism only differentiates the services if the difference between the parameters W_0 and W_1 is larger than 3 units.

6 Conclusion

In this paper, we presented the Load-Level-Based Admission Control (LLAC) mechanism for providing service differentiation in optical burst-switched networks. The proposed mechanism admits bursts of a given service class according to the network load and a class-associated parameter, called load level. An analytical model was derived for LLAC using a continuous-time Markov chain. The proposed mechanism also was compared to other existing mechanisms. The blocking probability experienced by each service class and the offered load were considered as performance metrics.

The LLAC mechanism, compared to static and dynamic mechanisms, provides a lower blocking probability for high-priority bursts in all analyzed configurations. In addition, LLAC effectively differentiates the blocking probability experienced by service classes even under adverse conditions, such as high offered load, high number of classes, and low level of aggressiveness against the low-priority class. As seen in the analysis, the static and dynamic mechanisms can provide for a high-priority class the same or higher blocking probability than a low-priority class. This phenomenon, known as priority inversion, does not happen for the proposed mechanism.

The better differentiation provided by LLAC is paid by a reduction in the network throughput mainly due to the reduction of the low priority-class throughput. Adjusting the LLAC parameters a network operator can control the level of aggressiveness against the low-priority class and, as a consequence, avoid the starvation of this traffic if desired. For an analyzed configuration, the proposed mechanism reduces the blocking probability of high-priority bursts by two orders of magnitude while the blocking probability of low-priority bursts only increases 37%. Furthermore, the higher the amount of high-priority traffic, the better the differentiation provided by LLAC and the lower the reduction of the throughput. This happens because more high-priority bursts contend to network resources and the proposed mechanism benefits the high-priority traffic. A scenario with a high load of bursts belonging to high-priority class is realistic because it is expected that the optical network traffic be mainly derived from voice and video applications.

Acknowledgements This work has been supported by CNPq, CAPES, FAPERJ, UOL, FINEP, RNP and FUNTTEL.

References

1. Battestilli, T., Perros, H.: An introduction to optical burst switching. IEEE Communications Magazine **41**(8), S10–S15 (2003)
2. Costa, L.H.M.K., Fdida, S., Duarte, O.C.M.B.: Incremental service deployment using the Hop By Hop multicast routing protocol. IEEE/ACM Transactions on Networking **14**(3), 543–556 (2006)
3. Farahmand, F., Zhang, Q., Jue, J.P.: A feedback-based congestion control mechanism for labeled optical burst switched networks. Photonic Network Communications **14**(3), 307–316 (2007)
4. Fawaz, M., Daheb, B., Audouin, O., Du-Pond, M., Pujolle, G.: Service level agreement and provisioning in optical networks. IEEE Communications Magazine **42**(1), 36–43 (2004)
5. Liao, W., Loi, C.H.: Providing service differentiation for optical-burst-switched networks. IEEE/OSA Journal of Lightwave Technology **22**(7), 1651–1660 (2004)
6. Moraes, I.M., Duarte, O.C.M.B.: An efficient multiclass mechanism for optical burst-switching networks. In: 5th IEEE/CreateNet International Workshop on Optical Burst/Packet Switching - WOBS, pp. 160–169 (2005)
7. Moraes, I.M., Duarte, O.C.M.B.: A multiclass admission control mechanism for optical burst-switched networks. In: 19th International Teletraffic Congress - ITC19, pp. 1937–1946 (2005)
8. Moraes, I.M., de O. Cunha, D., Bicudo, M.D.D., Laufer, R.P., Duarte, O.C.M.B.: An admission control mechanism for providing service differentiation in optical burst-switching networks. In: XII International Conference on Telecommunications - ICT'05 (2005)
9. Phuritatkul, J., Ji, Y., Zhang, Y.: Blocking probability of a preemption-based bandwidth-allocation scheme for service differentiation in OBS networks. IEEE/OSA Journal of Lightwave Technology **24**(8), 2986–2993 (2006)
10. Qiao, C., Yoo, M.: Optical Burst Switching - a new paradigm for an optical Internet. Journal of High Speed Networks **8**(1), 69–84 (1999)
11. Reviriego, P., Hernández, J.A., Aracil, J.: Analysis of average burst-assembly delay and applications in proportional service differentiation. Photonic Network Communications **14**(2), 183–197 (2007)
12. de Souza e Silva, E., Leão, R.M.M.: The TANGRAM-II environment. In: XI International Conference on Modelling Tools and Techniques for Computer and Communication System Performance Evaluation - TOOLs, pp. 366–369 (2000)
13. Tachibana, T., Kasahara, S.: Burst-cluster transmission: service differentiation mechanism for immediate reservation in optical burst switching networks. IEEE Communications Magazine **44**(5), 46–55 (2006)
14. Tintor, V., Matavulj, P., Radunovic, J.: Analysis of blocking probability in optical burst switched networks. Photonic Network Communications **15**(3), 227–236 (2008)
15. Wan, J., Zhou, Y., Sun, X., Zhang, M.: Guaranteeing quality of service in optical burst switching networks based on dynamic wavelength routing. Optics Communications **220**, 85–95 (2003)
16. Yoo, M., Qiao, C., Dixit, S.: QoS performance of optical burst switching in IP-over-WDM networks. IEEE Journal on Selected Areas in Communications **18**(10), 2062–2071 (2000)
17. Zhang, Q., Vokkarane, V.M., Jue, J.P., Chen, B.: Absolute QoS differentiation in optical burst-switched networks. IEEE Journal on Selected Areas in Communications **22**(9), 2062–2071 (2004)
18. Ziviani, A., de Rezende, J.F., Duarte, O.C.M.B.: Evaluating the expedited forwarding of voice traffic in a differentiated services network. International Journal of Communication Systems **15**(9), 799–813 (2002)