Multiclass Wavelength Reservation in Optical Burst Switched WDM Networks[*]

Chi-Hong Loi and Wanjiun Liao[+]

Department of Electrical Engineering
National Taiwan University
Taipei, Taiwan

Abstract This paper proposes a preemptive multiclass wavelength reservation protocol to provide differentiated service for Optical Burst Switched (OBS) networks, without requiring buffers at the WDM layer. Unlike existing approaches, such as JET QoS, which suffer from implementation constraints and which may degrade to classless schemes, our mechanism is robust and supports an incremental deployment of QoS support and cooperates well with other "best-effort" reservation mechanisms like Horizon and the original JET. We maintain a usage profile for each class at the router, and implement a preemptive wavelength reservation algorithm to ensure QoS. We also conduct simulations to evaluate performance. The result shows that our approach performs best in terms of lower blocking probability and higher resource utilization, making our approach an excellent QoS mechanism for OBS networks.

Keywords: Optical Burst Switching (OBS), WDM, differentiated service, wavelength reservation

I. Introduction

Optical Burst Switching (OBS) [1,2] is a promising solution to provide terabit optical routing and to build an all-optical WDM layer for optical Internet. Based on the concept of "burst" switching, OBS groups several IP packets with the same network egress address and common attributes like quality of services (QoS) into a burst and forwards the burst through the network as a single entity. A burst consists of a burst header and a burst payload. OBS uses physically separate wavelengths/channels to transmit and switch data bursts (i.e., payloads) and their headers, with the burst header slightly ahead in time. This allows optical core routers to process the headers electronically, to set up an end-to-end optical path, and to switch data bursts optically. OBS is based on a one-way reservation protocol, in which a data burst follows its header without waiting for an acknowledgment to come back before the data transmission.

Current IP provides only best effort service to deliver variable length packets. The future Internet will demand differentiated services for multimedia applications. Thus, for the optical Internet to be truly ubiquitous, one must address, among other important issues, how the WDM layer provides differentiated service support.

There are many mechanisms in the literature [3-5] to implement QoS, mostly using buffers and scheduling. These approaches, however, incur high processing overhead at intermediate nodes for switching and mandate a certain amount of buffers. No efficient optical buffer is available today. The use of electronic buffers necessitates opto-electro-optic (O/E/O) conversions, which must be avoided in an all-optical network where data is kept in the optical domain at all intermediate nodes. This calls for new QoS mechanisms for OBS WDM networks.

[6] proposed a new approach to differentiating services without requiring buffers at the WDM layer (we call it "JET QoS" in the rest of the paper). Taking advantage of the offset time between the

---

burst header and its data burst, JET QoS assigned different "extra" offset times to different classes of traffic. The extra offset time of each higher-class must be larger than the maximum burst length of all lower-class bursts, and each higher class must take a longer reservation delay (i.e., extra offset time) than all the lower classes. As a result, a higher priority class could be isolated from lower priority classes, thus ensuring different blocking probabilities for different classes of traffic. This approach, however, has two constraints when used to provide basic QoS service. First, it requires all the routers to implement the same JET mechanism as in [7]. Second, all edge routers must use the same burst length and offset time for the same class of traffic. These two constraints limit the practical use of the JET QoS mechanism. In addition to these constraints, JET QoS service may degrade to classless (i.e., best effort) service, explained as follows. The extra offset time is assigned at the edge router, and must be strictly enforced at all intermediate core routers. However, such offset times may not maintainable at the core routers due to longer rerouting paths, congestion in the control channel, and different offsets and burst durations in a big OBS network operating on different network technologies by different organizations. The JET QoS mechanism fails under these circumstances.

This paper proposes a preemptive wavelength reservation protocol to provide differentiated service for OBS networks without requiring buffers at the WDM layer. Unlike JET QoS which has many implementation constraints and may degrade to a classless mechanism, our protocol supports an incremental deployment of QoS support and cooperates well with other "best-effort" reservation mechanisms like Horizon [1] and the original JET [7]. We maintain a usage profile for each class at the router, and implement a preemptive wavelength reservation algorithm to ensure QoS. Simulations were conducted to evaluate performance. The result shows that our approach performs best in terms of lower blocking probability and higher resource utilization, making our approach an excellent QoS mechanism for OBS networks.

The rest of the paper is organized as follows. Section 2 describes the proposed preemptive multiclass wavelength reservation protocol. Section 3 presents the simulation study to evaluate the performance of the proposed mechanism. Finally the concluding remarks are included in Section 4.

II. Multiclass Wavelength Reservation

This section describes the proposed preemptive multiclass wavelength reservation mechanism for the OBS network.

A. Protocol Fundamentals

Suppose that a switch has a total of m wavelengths per output link to serve data bursts. Considering the characteristics of traffic, we classify bursts into k different classes, say $c_1, c_2, \cdots c_k$. To differentiate service to different classes of bursts, each class is assigned a service priority. Without loss of generality, the priority of classes $c_1, c_2, \cdots c_k$ is assumed to be in an ascending order of $c_1 < c_2 < \cdots < c_k$. The higher the priority, the lower the blocking probability. This implies that if class c_j has priority over class c_i, class c_j bursts are allowed to use more resources than class c_i bursts. The switch assigns each class a usage limit, defined as a percentage of system utilization (in terms of the number of wavelengths) the class is allowed to use. Let p_i be the usage limit assigned to class c_i requests. Thus, $\sum_{i=1}^{k} p_i = 1$, and $p_1 < p_2 < \cdots < p_k$.

In our protocol, the switch assigns a usage limit to each service class, and maintains a usage profile per class to monitor their current usages. Based on the profile table, the switch can determine if there is an *eligible* wavelength for a new request. A wavelength is eligible for a request, say R, if it is not assigned to any other request during the burst duration of request R. The entry of the usage profile records a predefined usage limit, the current usage, and a list of granted requests with the following triple: burst duration, outgoing wavelength, and a predefined timer. Each header burst carries the offset time to the data burst, and the burst length. Let l, s, and e be the burst length, the start time and the end time of the data burst, respectively. Thus, the start time, s, is equal to the current time plus the offset time carried in the burst header, and $e = s + l$. The burst duration is maintained in the format of (the start time, the end time) of the data burst. The outgoing wavelength is an eligible or preempted wavelength scheduled to transmit the data burst. The predefined timer records the maximum tolerable time to wait for the receipt of the data burst, in an attempt to cope with network faults and the preemption allowed by our protocol.

B. In-Profile Verification

A class of traffic is said to be in profile if its current usage does not exceed a predefined limit; otherwise, the class is out of profile. We can use the following approach to determining if a class is in profile. Assume there are n class c_i requests granted. Let l_x be the burst length of a class c_i request R_x, $x = 1 \cdots n$, and μ_i, the current total usage of class c_i.

$$\mu_i = \frac{\sum_{x=1}^{n} l_x}{m \times (T - t_0)},$$ where t_0 is the current time and

T is the last finished time defined as $T = \max_{x=1 \cdots n} \{e_x\}$, e_x is the end time of request R_x. Class c_i is said to be in profile if $\mu_i \leq p_i$; otherwise, the class is out of profile.

C. Burst Preemption

Assume that a switch has newly received an in-profile class request R_x with start and end times of s_x and e_x, respectively. Let C_O be a set of out-of-profile classes, defined as

(1) $C_O = \{c_i \mid \mu_i > p_i, i = 1, 2, \cdots k\}$

(2) $\forall c_i \in C_O, \exists R_v \in c_i, s_x \geq s_v$ and $e_x \leq e_v$.

where μ_i and p_i are the current usage and the usage limit of class c_i, respectively; s_v and e_v are the start and end time of request R_v, respectively. These two conditions imply that every class in C_O must be out-of-profile (by condition (1)), and must, at least, include a request previously granted but its data burst overlapping in time with the in-profile request newly received (condition (2)). In other words, the wavelength scheduled to any request of a class in C_O, say R_p, can be used to serve R_x if R_p is preempted.

Assume that no eligible wavelength is available to serve R_x. The preemption process proceeds as follows. The switch preempts a wavelength from the class with the lowest priority in C_O, and updates the current usage accordingly. Assume that c_a is the lowest priority class in C_O, and contains a set of requests $\{R_v \mid s_x \geq s_v, e_x \leq e_v, v = 1, 2, \cdots m\}$. Request R_p in c_a is the victim to be preempted if $l_p \leq l_i$ and

$e_p \geq e_i$, $i = 1 \cdots m$.

D. Operation Overview

A switch keeps monitoring its usage profile table. Upon receiving a class c_i request, the switch first attempts to identify an eligible wavelength for the request. If the attempt succeeds, the request is granted, and the usage profile of class c_i is updated; otherwise, the following takes place. The switch first examines if the class to which the request belongs is in profile, using the in-profile verification algorithm described in Sec. 2.2. If it is in profile, the switch preempts a previous granted request from an "out-of-profile" class using the burst preemption algorithm described in Sec. 2.3; otherwise, request R is rejected and the data burst is just simply dropped.

A switch may grant a request either with an eligible wavelength, or a preempted wavelength. Once the request is granted, the switch records the information of burst duration and outgoing wavelength for the request in the corresponding usage profile. To prevent a granted request from being preempted by any switch in the data channel, thereby wasting resources in the reserved path, we associate each pair of (burst duration, outgoing wavelength) with a predefined timer. The timer is activated at the requested start time of a burst, and in the middle of any burst transmission when data packet is not received in time as expected. On expiry of the timer, the switch assumes an occurrence of a fault (either a physical fault or a preemption) if no data burst has been received. It then removes the switching information of the associated burst, and makes the wavelength available for other requests.

III.   Performance Evaluation

This section presents the simulation results to compare the performance of the proposed mechanism with classless (i.e., best effort) and JET QoS [6]. The classless mechanism may be Horizon [1] or the original JET [7]. The following results are obtained with the original JET. The JET QoS is the original JET with different extra offset times assigned to different classes of bursts.

We consider bufferless switches with m wavelengths in each output link. Each switch is assumed to be capable of full wavelength conversion. We assume there are k classes, all of which generate bursts with an exponential inter-arrival time and exponential burst duration. To simplify the computation and without loss of generality, the

847

simulation is based on the assumption that all sources have the same arrival rates (i.e., $\lambda_1 = \lambda_2 = \cdots = \lambda_k = \lambda$) and service rates (i.e., $\mu_1 = \mu_2 = \cdots = \mu_k = \mu$).

This experiment investigates the blocking probability of different mechanisms as a function of the offered load in a single bufferless WDM switch.

The offered load here is defined as $\dfrac{\sum\limits_i \lambda_i}{m\mu}$, where m is the number of wavelengths in each link, λ_i is the arrival rate of class c_i and μ is the service rate of each burst. We first consider two classes only, namely, classes 1 and 2, in an attempt to observe the service differentiation offered by each mechanism. We let class 2 have priority over class 1, and assign the usage limits of 0.0 and 1.0 to classes 1 and 2, respectively. Thus, class 2's traffic can preempt class 1's traffic when necessary.
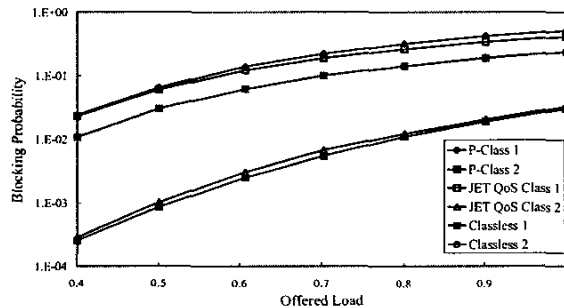

Figure 1. Blocking probabilities of the three mechanisms

Fig. 1 shows that the blocking probabilities of the three approaches when each output link has 8 wavelengths. These blocking probability curves increase as the offered load increases. Both JET QoS and the proposed preemptive approach (denoted P in the figure) provide service differentiation for multiclass traffic. While the difference is small, the blocking probability of class-2 traffic in our approach is always lower than that in JET QoS. Fig. 2 shows the blocking probabilities of the three approaches when the offset delay time becomes invalid due to network congestion. It can be observed that our approach still provides differentiated service for classes 1 and 2 traffic, while JET QoS degrades into a classless scheme. Fig.

3 shows the overall blocking probabilities of the three mechanisms. It can be seen that the three curves overlapped, obeying the conservation law of the system.
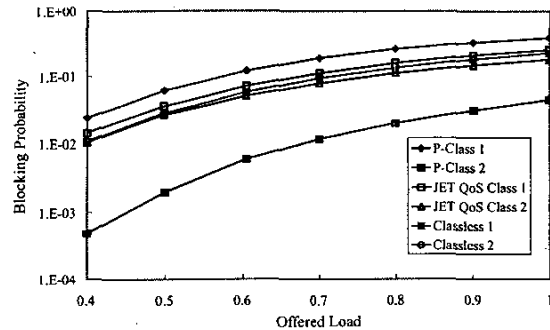

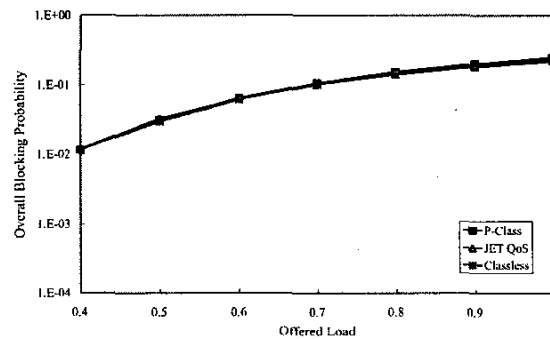Figure 2. Blocking probabilities of the three mechanisms with network congestion

Figure 3. Overall blocking probabilities of the three mechanisms

We then extend k from two to four classes with the following priority: class 4 > class 3 > class 2 > class 1. Fig. 4 shows the class blocking probabilities of classes 1 to 4, with the usage limits of 0.75, 0.2, 0.05, and 0.0 for the four classes, respectively. The higher the offered load is, the higher the blocking probability. Note that the blocking probability is closely related to the usage limit. Either can be derived with a priority queuing system with preemption when the other is given. Thus, the service provider can assign a usage limit to a class once the guaranteed blocking probability for the class is determined.
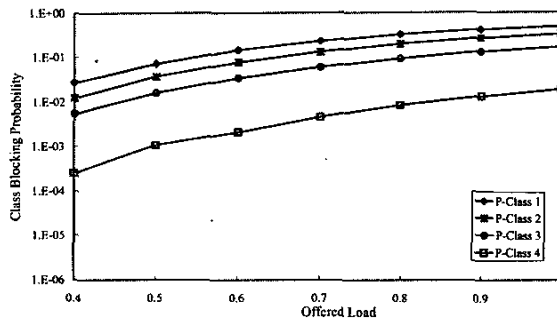
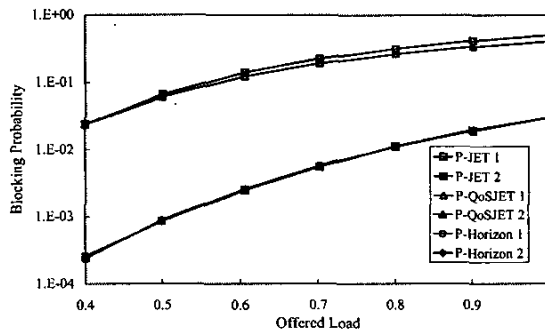Figure 4. Class blocking probabilities of four classes

Figure 5. Blocking probabilities of three reservation mechanisms in cooperation with our protocol

Fig. 5 shows the blocking probability of the proposed mechanism in cooperation with two "best effort" mechanisms, Horizon and the original JET, and with JET QoS. It shows that our approach can work well with existing reservation mechanisms and provide/improve their provision of service differentiation for different classes of traffic.

IV. Concluding remarks

In this paper, we have described a new bufferless mechanism using a preemptive wavelength reservation mechanism to differentiate services in optical burst switched WDM networks. Unlike JET QoS has many implementation constraints and which may degrade to a classless scheme, our mechanism is robust and supports an incremental deployment of QoS support and cooperates well with other "best-effort" reservation mechanisms like Horizon and the original JET, and even with JET QoS. We maintain a usage profile for each class at the router, and implement a preemptive wavelength reservation algorithm to ensure QoS. We have also conducted simulations to evaluate performance. The result shows that our approach performs best in terms of lower blocking probability and higher resource utilization, making our approach an excellent QoS mechanism for OBS

networks.

References

[1] Y. Xiong, M. Vandenhoute, and H. C. Cankaya, "Control Architecture in Optical Burst-Switched WDM Networks," *IEEE JSAC*, vol. 18, no. 10, pp. 1838-1851, Oct. 2000.

[2] C. Qiao and M. Yoo, " Optical Burst Switching (OBS) – A New Paradigm for an Optical Internet," *Journal of High Speed Network*, vol. 8, no. 1, pp. 69-84, 1999.

[3] H. Zhang, "Service Disciplines for Guaranteed Performance service in Packet Switching Networks," *Proc. of the IEEE*, vol. 83, pp. 1374-1396, Oct. 1995.

[4] A. Varma and D. Stilladis, "Hardware Implementation of Fair Queuing Algorithms for Asynchronous Transfer Mode Networks," *IEEE Communications Magazine*, vol. 35, pp. 74-80, Dec. 1997.

[5] U. Briem et al., "The Traffic Management for an ATM Switch with per-VC Queuing: Concept and Implementation," *IEEE Communications Magazine*, vol. 36, pp. 88-93, Jan. 1998.

[6] M. Yoo and C. Qiao, "Supporting Multiple Classes of Services in IP over WDM Networks," *Proc. IEEE Globecom '99*, pp. 1023-1027.

[7] M. Yoo and C. Qiao, "Just-Enough-Time(JET): A High Speed Protocol for Bursty Traffic in Optical Networks," *Dig. of IEEE/LEOS Summer Topical Mtgs. Technologies for a Global Info. Infrastructure*, pp. 26-27. Aug. 1997.

849