1-1-2006

Supporting differentiated quality of service in optical burst switched networks

Bin Zhou

University of Central Florida

Mostafa A. Bassiouni

University of Central Florida

Find similar works at: <https://stars.library.ucf.edu/facultybib2000>

University of Central Florida Libraries <http://library.ucf.edu>

This Article is brought to you for free and open access by the Faculty Bibliography at STARS. It has been accepted for inclusion in Faculty Bibliography 2000s by an authorized administrator of STARS. For more information, please contact STARS@ucf.edu.

Recommended Citation

Zhou, Bin and Bassiouni, Mostafa A., "Supporting differentiated quality of service in optical burst switched networks" (2006). *Faculty Bibliography 2000s*. 6766.

<https://stars.library.ucf.edu/facultybib2000/6766>

Supporting differentiated quality of service in optical burst switched networks

Bin Zhou

Mostafa A. Bassiouni

University of Central Florida
School of Computer Science
Orlando, Florida 32816

E-mail: bzhou@cs.ucf.edu
bassi@cs.ucf.edu

Abstract. We propose and evaluate two new schemes for providing differentiated services in optical burst switched (OBS) networks. The two new schemes are suitable for implementation in OBS networks using just-in-time (JIT) or just-enough-time (JET) scheduling protocols. The first scheme adjusts the size of the search space for a free wavelength based on the priority level of the burst. A simple equation is used to divide the search spectrum into two parts: a base part and an adjustable part. The size of the adjustable part increases as the priority of the burst becomes higher. The scheme is very easy to implement and does not demand any major software or hardware resources in optical cross-connects. The second scheme reduces the dropping probability of bursts with higher priorities through the use of different proactive discarding rates in the network access station (NAS) of the source node. Our extensive simulation tests using JIT show that both schemes are capable of providing tangible quality of service (QoS) differentiation without negatively impacting the throughput of OBS networks. © 2006 Society of Photo-Optical Instrumentation Engineers. [DOI: 10.1117/1.2151852]

Subject terms: optical burst switching networks; wavelength division multiplexing; differentiated quality of service; dropping probability; optical network performance evaluation.

Paper 040845RR received Nov. 11, 2004; revised manuscript received May 16, 2005; accepted for publication May 31, 2005; published online Jan. 24, 2006.

1 Introduction

To provide fine bandwidth granularity and improve the utilization of wavelength division multiplexed (WDM) optical networks,¹ optical burst switching (OBS) has been proposed.^{2,3} In real-life applications with differentiated quality of service (QoS) requirements, data bursts should have different priority classes. Higher priority bursts should be given preferential treatment to reduce their drop probability and their end-to-end delay. The growing interest in introducing QoS differentiation in Internet services is motivated by the need to improve the quality of support for IP voice and video services, and in general, by the desire to provide clients with a range of service-quality levels at different prices. Since WDM optical networks are rapidly becoming the technology of choice in network infrastructure and next-generation Internet architectures, implementing QoS differentiated services and designing network protocols to support a range of service-quality levels in WDM and OBS networks have received increased recent attention. In Ref. 2, differentiated quality of service has been incorporated into just-enough-time (JET) scheduling³ by assigning different offset times to different classes. The higher priority class is given larger offset time. The drawback of this approach is that larger offset times may result in longer delays for higher priority bursts. The review in Sec. 2 presents several other proposals for improving QoS differentiation in optical burst switching (OBS) networks. We propose and evaluate two schemes to improve the QoS

differentiation in OBS networks. The two schemes are easy to implement and produce tangible QoS differentiation.

The rest of the work is organized as follows. In Sec. 2, the QoS differentiation problem in optical networks is discussed and relevant previous works are reviewed. Our first scheme, qualified just-in-time (JIT), is presented in Sec. 3. Our second scheme, prioritized random early dropping (RED), is presented in Sec. 4. In Sec. 5, the performance results of the two schemes are presented and analyzed. In Sec. 6, the conclusion of the work is given.

2 Quality of Service Differentiation Problem in Optical Burst Switching Networks

In standard OBS networks, a control packet is transmitted ahead of the data burst on an out-of-band channel to reserve a channel for the upcoming burst in each optical cross-connect (OXC) along the lightpath of the burst. There is a special offset time that is introduced at the source node between the transmission of the control packet and the data burst. During this offset time, the burst data are buffered electronically in the network access station (NAS), while the control packet propagates forward to configure each OXC along its lightpath. When the offset time expires, the burst is sent out and is switched all optically from one node to the next until it reaches the destination. It is possible, however, that the control packet fails to secure a free channel in some congested intermediate node along the lightpath. This results in dropping the data burst at the congested node. The data burst dropping probability generally increases as the load on the network increases.

There are two main optical burst scheduling methods:

Fig. 1 Two network topologies.

just-in-time (JIT)⁴ and just-enough-time (JET).³ Several methods to support differentiated QoS in OBS networks have been proposed in the literature. For example, the authors in Ref. 2 propose a method to support QoS differentiation in JET by adjusting the offset time for different priority classes. The basic idea of this method is to assign a larger offset time for a higher priority class than the offset time assigned for a lower priority class. The larger offset time in JET increases the chance of securing a successful wavelength reservation for the burst. However, larger offset times also increase the delay for higher priority classes. In Ref. 5, QoS is supported by adjusting the offset time, and the lower and upper bounds on blocking probability for two burst classes are analyzed. While the method of increasing offset times works well for QoS prioritization in JET, it is not exactly a suitable method for JIT. The JIT scheduling method is less complex than that of JET and cannot fully utilize delayed reservation.³ In this work, we propose and investigate two schemes that can be used to support QoS differentiation in both JIT and JET.

Burst assembly at the ingress nodes⁶ is another method to improve QoS in JET. In this method, the window size and weight of a class determines the number of packets in the window. In Ref. 7, the QoS differentiation in JET is investigated by analyzing the loss probabilities of two classes of bursts. As was done in Ref. 2, the authors in Ref. 7 also focus on changing the offset time to support differentiated QoS. In Refs. 8 and 9, a burst segmentation method is proposed to address the differentiated QoS problem. Using the segmentation scheme, the bursts with higher priority can preempt the overlapping segments of lower priority bursts, and the preempted segments are dropped or are deflected to alternate routes. In Ref. 10, a proportional model is proposed to enhance the offset-based QoS differentiation method proposed in Ref. 2. In the proportional model, the differentiation of a particular QoS metric can be quantitatively adjusted to be proportional to the factors that a network service provider sets. The lower priority burst is intentionally dropped when the proportional differential model is violated. In Ref. 11, the authors combine a service differentiation model based on proportional resource allocation with a partially preemptive burst scheduler. The scheme improves the performance of loss and utilization while providing QoS with controllable service differentiation. In Ref. 12, the authors propose a differentiation

scheme for JET that does not assign extra time to a higher priority class, as in Ref. 2. Rather, it uses a priority queuing technique to schedule a higher priority burst earlier than a lower priority burst. Another proportional differentiation model is used in Ref. 13. In this model, the burst control packets are queued in increasing order of the burst preferred scheduling time (defined as a burst arriving time minus its differential time). Each OXC chooses its own differential time function according to its resource availability and QoS requirement. In Ref. 14, some undesired characteristics of the offset-time management mechanism for JET are identified. The authors found that the burst drop probability differentiation attained for a given offset-time value strongly depends on the distribution of the burst durations and that controlling the differentiation is difficult. In Ref. 15, assured Horizon is introduced for a coarse-grained bandwidth reservation r_i for every forwarding equivalent class (FEC) between ingress and egress. The burst assembler marks bursts as compliant and noncompliant bursts, depending on whether the burst is conforming to r_i . The noncompliant bursts are dropped when congestion occurs. In Ref. 16, a generalized latest available unused channel with void filling (LAUC-VF) algorithm is proposed. The

Fig. 2 QJIT drop probabilities for different priority levels in LongHaul network (load=12).

Fig. 3 QJIT drop probabilities for different priority levels in 5×5 mesh-torus network (load=20).

LAUC-VF algorithm aims at providing good performance by essentially choosing the wavelength with the smallest possible available window, leaving larger windows for control packets that arrive later. LAUC-VF is basically another scheduling scheme for OBS networks and is used in Ref. 16 to support differentiated services with limited buffers. In Ref. 17, an algorithm is proposed to decide about the value of the burst offset-time in JET based on the burst priority class. In Ref. 18, linear predictive filter (LPF)-based forward resource reservation is proposed for JET to reduce the burst delay at edge routers. An aggressive reservation method is also proposed to increase the successful forward reservation probability and to improve the delay reduction performance. In Ref. 19, preemptive multiclass wavelength reservation is used to provide differentiated services for the JET protocol. The preemption scheme provides QoS differentiation but complicates the control logic in OXCs. In Ref. 20, a probability preemptive scheme is proposed in which high priority bursts can preempt low priority bursts in a probabilistic manner. In Ref. 21, the authors propose a priority-based wavelength assignment (PWA) algorithm. In this algorithm, each wavelength has its own priority that can be changed according to the burst reachability. In Ref. 22, the authors propose QoS-guaranteed wavelength allocation schemes for WDM networks. In their schemes, the wavelengths are classified into different sets based on the QoS requirement, and higher priority requests can be allocated more wavelengths. In each set, there are two rules to select the idle wavelength: minimum index numbered or maximum index numbered wavelength. The authors in Ref. 22 showed that the connection loss probability for higher priority requests is improved, but the overall throughput performance of their schemes is not discussed.

In this work, we propose two new schemes to support QoS differentiation in OBS networks suitable for both JIT and JET scheduling. Compared to the previous proposals, our schemes have simple logic that can be easily implemented, even in the JIT scheduling method. Specifically, the two schemes do not increase the offset time of high priority bursts (thus do not increase their delay), do not use complex scheduling functions, do not introduce additional queuing or segmentation mechanisms, do not resort to burst

Fig. 4 Drop probability for different loads with priority=1, LongHaul network.

preemption, and do not introduce any complex modification to the lightpath setup scheme or the architecture of OXC. We have implemented a performance simulation model of the two schemes in OBS networks using JIT scheduling. Our extensive performance results show that the two methods improve QoS differentiation compared to the original JIT scheduling scheme without negatively impacting the network throughput. Since our schemes do not modify or depend on the scheduling logic of JIT, they can be also implemented in JET. In particular, they can be applied in conjunction with the improved JET scheduling algorithms recently proposed in Ref. 23.

3 Qualified Just-in-Time Scheme

The JIT scheduling protocol is used in this work to illustrate the applicability of our two proposed QoS differentiation schemes and evaluate their efficiency. As has been frequently assumed in JIT/JET scheduling,^{2-4,24} full wavelength conversion capability is assumed to be available at each node along the lightpath of the burst.

In JIT, the source node delays the transmission of a data burst by a certain amount of time after sending the control packet. The amount of this delay (called the *offset time*) is decided by the number of hops along the lightpath and the cut through time in each node. The offset time allows each hop (OXC) along the lightpath to configure its port connection for the incoming burst. The switch reconfiguration time (also called the cut through time) must be taken into consideration, because a burst is dropped if it arrives before the OXC completes its connection reconfiguration.

Normally, if the lightpath of a burst consists of m hops, the offset time t_d used in JIT can be defined as:

$$t_d \geq m^* t_p + t_\delta, \quad (1)$$

where t_p is the control packet processing time in each OXC, including O/E-E/O conversions and request/routing analysis; and t_δ is the extra delay required to assure cut through completion at the last OXC in the lightpath.⁴

Fig. 5 Drop probability for different loads with priority=3, LongHaul network.

We next explain the rational of our first scheme using a simplified high-level model for the probability of burst discarding in JIT.

Consider a burst that arrives at an OXC. Let n be the number of wavelengths operational on the destination output link of this OXC. Let β_i be the probability that the i 'th wavelength is not free at the time of burst arrival. The dropping probability of the burst in this OXC is given by:

$$P_{\text{drop}} = \beta_1 \times \beta_2 \times \beta_3 \times \dots \times \beta_n. \quad (2)$$

If $\beta_i = \beta$ for all i , then

$$P_{\text{drop}} = \beta^n. \quad (3)$$

The “go through” probability of the burst in this OXC is given by

$$P_{\text{go}} = 1 - P_{\text{drop}}. \quad (4)$$

If the lightpath of the burst has m hops, the probability that the burst will successfully reach its destination is given by

$$P_{\text{success}} = P_{\text{go}_1} \times P_{\text{go}_2} \times P_{\text{go}_3} \times \dots \times P_{\text{go}_m}. \quad (5)$$

These equations apply to an arbitrary burst from any traffic class. This means that standard JIT treats all traffic classes equally and does not provide differentiated QoS to higher priority classes. In this section, we propose a wavelength assignment scheme that skews the search of free wavelength in favor of higher priority classes. By introducing a bias in the search process, a higher priority burst will get a better chance to go through an OXC than a lower priority burst. The basic idea of the scheme is to make more free wavelengths available to higher priority bursts than to lower priority bursts. The scheme is based on a simple observation of Eqs. (2)–(4), namely, increasing the number of wavelengths n increases the chances of successful burst delivery. As the priority of the burst increases, our scheme gradually increases the number of wavelengths that can be used to switch this burst. Let W be the maximum number of wavelengths that is used for burst switching in each OXC and let P be the number of burst classes. We assume that class P has the highest priority, class $P-1$ has the second highest priority, and class 1 has the lowest priority. When

Fig. 6 Drop probability for different loads with priority=5, LongHaul network.

the control packet of a burst arrives at some OXC, it will reserve the wavelength based on its burst priority. For the burst with priority 1, the control packet is allowed to search just a fraction of the W wavelengths in this OXC. For the burst with priority 2, the control packet is allowed to search a fraction of W wavelengths that is larger than the fraction for priority 1, and so on. For a burst with priority P , all of the W wavelengths can be searched. Specifically, for the i 'th priority control packet, the number of wavelengths n_i that is searched in a hop is given by:

$$n_i = (1 - g) * W + g * i * W / P, \quad (6)$$

where g is a parameter that is assigned a value between 0 and 1. We call this method “qualified JIT” and denote it QJIT(g), where g is the controllable parameter of the scheme. The parameter g divides the search spectrum in each OXC into two parts: a base part and an adjustable part. The base part has a fixed size of $(1 - g) * W$ wavelengths, regardless of the priority level of the burst. The base part ensures that every type of burst can search some number of

Fig. 7 Drop probability for different loads with priority=1, mesh-torus network.

Fig. 8 Drop probability for different loads with priority=3, mesh-torus network.

wavelengths. The adjustable part has a size that depends on the priority level of the burst, and can reach a maximum size of g^*W wavelengths for the highest priority level. For example, if the highest priority level is $P=5$, the size of the adjustable part is 0.2^*g^*W for priority level 1, 0.4^*g^*W for priority level 2, and so on. It should be noted that the high-level model of Eqs. (2)–(5) is suitable for both JIT and JET, and the qualified JIT scheme is therefore suitable for JIT as well as JET. The qualified JIT scheme is very easy to implement and does not demand any major software or hardware resources in the OXCs. The priority level of the burst is easily passed from one hop to the other hop along a lightpath via the control packet. Implementing the adjustable search for a free wavelength implied by Eq. (6) requires minor modification to the standard JIT channel allocation scheme; the adjustable search (i.e., searching in a space of size g^*i^*W/P) actually leads to a smaller average search time. There are two important remarks about Eq. (6).

Remark 1. When $g=0$, the adjustable part of Eq. (6) van-

Fig. 9 Drop probability for different loads with priority=5, mesh-torus network.

Fig. 10 QJIT throughput for different values of g , LongHaul network.

ishes and the scheme becomes equivalent to the standard JIT scheme. In other words, $QJIT(g=0)$ is identical to the standard JIT scheme. As the value of g increases, data bursts with higher priority can get better treatment, since the size of their adjustable search space increases and hence their “go through” probability gradually improves. It is obvious that larger values of g will be more effective in providing differentiated QoS services. However, higher values of g could lead to severely deteriorated performance for low priority classes and could adversely impact the overall throughput of the network. Ideally, we should choose a value of g that provides a good compromise between differentiated QoS and network throughput. The ideal value of g should provide tangible improvement in the QoS of high priority traffic without negatively impacting the overall throughput of the system.

Remark 2. The value n_i in Eq. (6) should be rounded to

Fig. 11 QJIT throughput for different values of g , mesh-torus network.

Fig. 12 QJIT overall drop probability for LongHaul network.

Fig. 13 QJIT overall drop probability for mesh-torus network.

an integer number. For effective QoS differentiation, different values of i (i.e., different priorities) should be mapped to different integer values of n_i . This will ensure that a higher priority burst will have a search space strictly larger in size than that of a lower priority burst and will therefore have a smaller blocking probability. To satisfy this constraint, there should be a limit on the value of P (number of supported priorities). Simply, the value of P should not exceed the value of the product g^*W so that different values of i in Eq. (6) map to different values of n_i . This is formally described by the following constraint on the number of priorities P that can be supported by the QJIT scheme.

$$P \leq g^*W. \quad (7)$$

The QJIT scheme and the logic used in Eq. (6) were inspired by our earlier work on reducing the dropping probability of handoff requests in the base stations of cellular wireless networks.²⁵ Abstractly speaking, the adjustable term of Eq. (6) is a generalization of the guard channels that are exclusively dedicated to handoff requests to give them priority over new call cellular requests. Our extensive tests have shown that the QJIT scheme can improve the differentiated QoS performance of optical burst switching networks while maintaining the overall throughput of the network. Our performance tests presented in Sec. 5 show the effectiveness of the QJIT scheme in providing improved QoS differentiated services in optical burst switching networks. It should be mentioned that while Eq. (6) gives the size of the search space for a given priority, it does not require a fixed set of wavelengths to be searched for that priority. In our QJIT scheme, the starting point of the search process is randomly selected. Even though the size of the search space is fixed, the set of wavelengths searched by QJIT is randomized. Unlike the schemes given in Ref. 22, by randomly selecting the sets for different priority bursts, the QJIT can maintain a healthy total network performance while improving QoS for higher priority bursts.

4 Prioritized Random Early Dropping Scheme

Our second scheme adapts the concept of random early discard (RED) to the OBS environment and prioritizes the levels of discarding based on the priority levels of bursts. We call this scheme prioritized RED or PRED.

The RED concept²⁶ has received considerable interest in electronic packet switching networks, and RED routers have been widely deployed in various commercial applications and in the Internet. There have been numerous studies that support or oppose RED,^{26–28} present schemes for tuning RED parameters,²⁹ propose modified versions of RED,³⁰ and develop analytical models for RED performance.³¹ The basic idea of RED is that routers proactively discard incoming packets with probabilities that depend on the size of the router's queue. The TCP congestion control algorithm³² reacts to lost packets by throttling the transmission rate of TCP senders. Studies have shown that well-configured RED routers have the potential to avoid severe congestion and improve the overall throughput while maintaining a small queuing delay within each router. The PRED scheme performs random proactive dropping different from the burst discarding mentioned in Ref. 33. In Ref. 34, the authors proposed a priority-based wave-

Fig. 14 Drop probability for LongHaul network with $W=12$, $P=12$, and $g=0.5$.

Fig. 15 Drop probability for mesh-torus network with $W=12$, $P=12$, and $g=0.5$.

length assignment scheme for all-optical networks. In their scheme, the higher priority requests can get the channel reserved by lower priority bursts. Unlike the scheme in Ref. 34, the dropping policy of PRED is not preemptive and is not triggered by the arrival of higher priority requests.

Our PRED differentiated QoS scheme for JIT uses proactive burst dropping with a discarding probability that decreases as the burst priority level increases. The goal of burst discarding in PRED is not to avoid congestion or throttle TCP senders as in RED (although these could be positive side effects of PRED). Rather, PRED uses proactive discarding to improve the QoS of higher priority bursts at the expense of some deterioration of the QoS of lower priority bursts. A major difference between RED and PRED is that our PRED scheme restricts burst discarding to the original source node of the burst, while RED allows any router in the path of a packet to proactively drop it. Specifically, all proactive discarding in PRED is done in the network access station (NAS) of the source node that gen-

Fig. 16 Drop probability distribution using PRED for LongHaul network. The values α_k of the proactive burst drop probabilities used in the source NAS are 0.8, 0.6, 0.4, 0.2, and 0.0 for priority levels 1, 2, 3, 4, and 5, respectively.

Fig. 17 Drop probability distribution using PRED for the LongHaul network. The values α_k of the proactive burst drop probabilities used in the source NAS are 1.0, 0.85, 0.55, 0.25, and 0.0 for priority levels 1, 2, 3, 4, and 5, respectively.

erated the burst. This restricted mode of discarding has the advantage that the discarded bursts will not waste any bandwidth resources in the core of the optical networks.

Let α_i be the probability used by PRED at the source NAS to discard a newly incoming burst whose priority level is i (larger i means higher priority level). To improve QoS differentiation in OBS networks, the values of the discarding probability should satisfy the following constraint:

$$\alpha_1 > \alpha_2 > \dots > \alpha_p, \quad (8)$$

where P is the number of priority levels in an optical burst switching network, as explained in Sec. 3. The proactive discarding of Eq. (8) is only applied to the local bursts assembled at this NAS. The NAS may also be servicing transit bursts that have come from some external OXCs and are being routed to other external OXCs. These transit

Fig. 18 Drop probability distribution using PRED for the LongHaul network. The values α_k of the proactive burst drop probabilities used in the source NAS are 0.9, 0.7, 0.5, 0.3, and 0.0 for priority levels 1, 2, 3, 4, and 5, respectively.

Fig. 19 Overall average drop probability corresponding to Fig. 16.

Fig. 20 Overall average drop probability corresponding to Fig. 17.

bursts have already escaped the PRED proactive discarding in the originating NAS where they were assembled. These transit bursts have also already consumed some network bandwidth resources during their partial trip toward their destination. By proactively discarding local bursts and not transit bursts, more bandwidth will be available to transit bursts in each OXC. This increases the likelihood that transit bursts will reach their destination without wasting the resources that they have already used prior to reaching the current OXC.

As in standard RED, the proactive discarding in PRED should not take place if the load on the OXC is not heavy. This is because at light loads, most bursts are expected to reach their destination successfully, and the burst dropping probabilities for all priority levels are nearly zero. We have adopted a simple mechanism to disable/enable proactive discarding in PRED. In OBS, each NAS uses some buffers to hold assembled bursts until they are sent to the local OXC. PRED does not discard arriving bursts when the free space in these buffers is greater than or equal to some threshold. When the free buffer space is less than the threshold, new bursts are subjected to the prioritized probabilistic discarding of Eq. (8). Notice that when the network is lightly loaded, the amount of free buffer space will be relatively large, and the PRED proactive discarding is disabled. Consequently, the actual discarding probabilities will be smaller than the probabilities α_i given in Eq. (8).

5 Performance Evaluation Results and Analysis

Figure 1 shows the two network topologies that are used in our simulation (U.S. LongHaul network with 28 nodes and 5×5 mesh-torus network).

In our simulation, a static lightpath between any two nodes is established using the shortest path first method, as was done in Refs. 3, 31, and 35. Notice that the longest shortest path has seven hops in the LongHaul topology and four hops in the 5×5 mesh-torus topology. The labels on the links of the LongHaul network represent the relative integer ratios of the lengths of the fiber cables of these links. For example, the delay of a link with label 5 is half the delay of a link with label 10. Similar to Ref. 31, the traffic used in our tests is uniformly distributed among all

nodes. This means that all nodes have equal likelihood to be the source of a data burst, and for a given source node, all other nodes in the network have equal likelihood to be the destination node. The number of priority levels in our simulation tests is $P=5$, and the traffic is equally distributed among all five classes (except the scenario in Fig. 14 and Fig. 15, which have $P=12$ priority levels). Bursts with priority P have the highest priority and those with priority 1 have the lowest priority.

In our simulation tests, assembled bursts arrive according to a Poisson distribution with controllable arrival rate λ . For each burst, the source and destination nodes are randomly selected as explained before. Our simulation tests used parameter values similar to those typically used in the literature: the cut through time in each OXC is 2.5 msec, the link delay per hop is 3 msec for the Mesh network, the burst length is 50 μ sec (equivalent to a burst of 250 Kbits at 5 Gbits/sec), and the control packet processing time t_p at each hop is 50 μ sec. For the LongHaul network shown in Fig. 1, the delay of a link in milliseconds is 0.5 multiplied by the length label of that link. Thus, the delay for a link

Fig. 21 Overall average drop probability corresponding to Fig. 18.

with length 6 is 3 msec. The number of wavelengths used in each OXC is $W=40$ (except-Fig. 14 and Fig. 15, which use 12 wavelengths). Each point in the performance graphs reported in this work was obtained by averaging the results of six simulation tests using different randomly generated seeds. Each simulation was run for a sufficiently long time to obtain stable statistics; the total number of bursts processed in each simulation test ranged from 4 million bursts at low arrival rates to 12 million bursts at high loads for the LongHaul network, and from 10 million bursts at low arrival rates to 20 million bursts at high loads for the mesh-torus network. The unit of time (denoted ut) used in the graphs presented in this work is equal to 0.05 msec. Thus, a load λ of 12 bursts/ ut is equivalent to 60 Gbits/sec.

5.1 Performance of QJIT

Figure 2 shows the burst drop probability for different priority levels at load 12 bursts/ ut (60 Gbits/sec) in the LongHaul network. Figure 3 shows the burst drop probability for different priority levels at load 20 bursts/ ut (100 Gbits/sec) in the 5×5 mesh-torus network. Notice that the mesh-torus network has more links than the LongHaul network, and it often has multiple shortest-path routes connecting the same source-destination pair. The mesh-torus network therefore requires higher total load than the LongHaul network to induce a certain level of congestion on the individual links. The horizontal axis in Figs. 2 and 3 gives the value of the parameter g . Notice that the case $g=0$ corresponds to the standard JIT scheme. The figures show that the QoS for the highest priority levels (levels 5 and 4) gradually improves as the value of g increases. The value $g=1$ gives the largest difference of drop probability between levels 5 and 1.

Figures 4–6 show the QJIT drop probability in the LongHaul network for priority values 1, 3, and 5, respectively. For each priority value, the drop probability is plotted for different values of g and load λ . Figures 7–9 show the corresponding graphs for the Mesh network. For the lowest priority (priority 1), Figs. 4 and 7 show that the drop probability increases as the value of g increases. For the medium priority (priority 3), Figs. 5 and 8 show that the burst drop probability does not change much as the value of g increases from 0 to around 0.8. When g reaches the value 0.8, there is a slight decrease in the drop probability for priority 3, then there is a significant increase in the drop probability as g increases to 1. Figures 6 and 9 show that as g increases, the drop probability for the highest priority (priority 5) decreases.

Figures 2–9 examined the QJIT drop probabilities for different priority levels and clearly showed that higher values of g are more effective in providing differentiated QoS. We now examine the overall throughput and the overall average drop probability (i.e., averaged over all priority levels). Figures 10 and 11 show the total throughput (in gigabits per unit time) at different values of g and different loads in the LongHaul and 5×5 mesh-torus networks, respectively. As evident from these graphs, the value $g=1$ degrades the throughput significantly. In general, values around $g=0.5$ to 0.6 give the best throughput performance while still providing noticeable differentiated QoS. Notice

Fig. 22 Drop probability distribution using PRED for the mesh-torus network. The values α_k of the proactive burst drop probabilities used in the source NAS are 0.6, 0.45, 0.3, 0.15, and 0.0 for priority levels 1, 2, 3, 4, and 5, respectively.

Fig. 23 Overall average drop probability corresponding to Fig. 22.

Fig. 24 Drop probability distribution using PRED for the mesh-torus network. The values α_k of the proactive burst drop probabilities used in the source NAS are 0.8, 0.6, 0.4, 0.2, and 0.0 for priority levels 1, 2, 3, 4, and 5, respectively.

Fig. 25 Overall average drop probability corresponding to Fig. 24.

Fig. 26 Drop probability distribution using PRED for mesh-torus network. The values α_k of the proactive burst drop probabilities used in the source NAS are 1, 0.8, 0.6, 0.4, and 0.0 for priority levels 1, 2, 3, 4, and 5, respectively.

Fig. 27 Overall average drop probability corresponding to Fig. 26.

that the throughput around $g=0.5$ is almost the same as the throughput of standard JIT (i.e., when $g=0$).

Figures 12 and 13 give bar charts for the average burst drop probability (i.e., averaged over all priority levels) for the LongHaul and 5×5 mesh-torus networks, respectively. Again, the value $g=1$ has the worst performance, and values around $g=0.5$ give the best overall average drop probability.

The prior results suggest that when using values of g close to 1, the QJIT(g) scheme provides strong QoS differentiation in OBS networks at the expense of degraded total network throughput performance. However, using values of g in the neighborhood of 0.5 offers definite practical advantage compared to standard JIT, namely, the throughput of the system is kept at normal levels and good QoS differentiation among priority levels is achieved.

As explained earlier in Eq. (7), the number of supported priorities P should not exceed g^*W . Since the previous results show that $g=0.5$ provides a good compromise between QoS differentiation and network throughput, the number of supported priorities P should be less than 0.5^*W . The smaller the value of P we use, the better QoS differentiation we get from the QJIT scheme. In all previous graphs, we used $P=5$ and $W=40$, and therefore the constraint in Eq. (7) was well satisfied. If the value of P exceeds g^*W , the QJIT scheme will not be able to provide strict differentiation between every pair of priorities. This is illustrated in Figs. 14 and 15, which use $P=12$, $W=12$, and $g=0.5$. The value of P is double the value of g^*W , and Eq. (6) produces the same value of n_i for two values of the priority i . As shown in Figs. 14 and 15, priority 2 and priority 3 have the same drop probability and therefore the same QoS. Similarly, priority 4 and priority 5 have the same QoS, priority 7 and priority 8 have the same QoS, and so on. Assuming the value of g is 0.5, the number of supported priorities P should preferably be less than 0.5^*W . For example, if $W=8$, the maximum number of supported priorities P is recommended to be 3 or less, and for $W=16$, it is recommended to be 5 or less.

5.2 Performance of PRED

As explained in Sec. 4, PRED uses proactive discarding at the source NAS, with probability α_k to discard bursts that have priority level k . Unlike the QJIT(g) scheme, which has a single parameter g , the PRED scheme has $P+1$ parameters (P discarding probabilities and the threshold on the size of free buffer space used to enable/disable proactive discarding). However, the constraint represented by Eq. (8) greatly simplifies the parameter tuning process.

Figures 16–18 show the drop probability distribution for four scenarios with different proactive dropping parameters in the LongHaul network. The empty buffer threshold used in these tests is 10% of the total buffer space. Notice that the probabilistic discarding is disabled when the total number of free buffers is greater than the threshold. Therefore, the actual rate of proactive discarding is lower than the values of the discarding probabilities α_k . The overall average drop probability (i.e., averaged over all priority levels) corresponding to the results of Figs. 16–18 are shown in Figs. 19–21, respectively. From Figs. 16–18 we can easily see that the PRED scheme improves the level of QoS dif-

ferentiation compared to the standard JIT scheme. As the intervals among the discarding probabilities α_k increase, the drop probability difference for different priority levels also increases. Figures 19–21 show that as the QoS performance is improved by the PRED scheme, there is no negative impact on the overall drop probability. Similar results for the 5×5 mesh-torus topology are shown in Figs. 22–27.

6 Conclusions

We investigate two methods for supporting QoS differentiation in optical burst switching networks using JIT scheduling. The first scheme, qualified JIT, adjusts the size of the search space for a free wavelength based on the priority level of the burst. The second scheme uses different proactive discarding rates in the network access station (NAS) of the source node. The first scheme (QJIT) has less number of parameters and is easier to implement/tune than the second scheme. We present performance simulation tests that show that both schemes are capable of providing tangible QoS differentiation without negatively impacting the throughput of OBS networks. One future extension of the work presented is to apply the two schemes to the JET scheduling protocol. Another extension is to investigate the case in which wavelength converters are not available in all nodes of the network.^{36,37}

Acknowledgments

This work has been supported by the National Science Foundation under grant number 9980316 and by ARO under grant number DAAD19-01-1-0502. The authors are grateful to the two anonymous reviewers for their comments and suggestions that helped improve the quality of this paper.

References

1. U. Black, Optical Networks, Prentice Hall, New York (2002).
2. M. Yoo, C. Qiao, and S. Dixit, "QoS performance of optical burst switching in IP-over-WDM networks," *IEEE J. Sel. Areas Commun.* **18**(10), 2062–2071 (2000).
3. C. Qiao and M. Yoo, "Optical burst switching (OBS)—a new paradigm for an optical internet," *J. High Speed Networks* **8**, 69–84 (1999).
4. J. Y. Wei and R. I. McFarland, "Just-in-time signaling for WDM optical burst switching networks," *J. Lightwave Technol.* **18**(12), 2019–2037 (2000).
5. M. Yoo and C. Qiao, "A new optical burst switching protocol for supporting quality of service," *Proc. SPIE* **3531**, 396–405 (1998).
6. M. C. Yuang, J. Shih, and P. L. Tien, "QoS burstification for optical burst switched WDM networks," *Proc. OFC*, pp. 781–783 (2002).
7. C. Gauger, K. Dolzer, J. Späth, and S. Bodamer, "Service differentiation in optical burst switching networks," *Beiträge zur 2. ITG Fachtagung Photonische Netze*, pp. 124–132 (2001).
8. V. Vokkarane and J. Jue, "Prioritized routing and burst segmentation for QoS in optical burst-switched networks," *Proc. Opt. Fiber Commun. Conf. (OFC)* (2002).
9. V. Vokkarane and J. Jue, "Prioritized burst segmentation and composite burst assembly techniques for QoS support in optical burst-switched networks," *IEEE J. Sel. Areas Commun.* **21**(7), 1198–1209 (2003).
10. Y. Chen, M. Hamdi, and D. H. K. Tsang, "Proportional QoS over OBS networks," *Proc., IEEE Globecom* **3**, 1510–1514 (2001).
11. H. C. Cankaya, S. Chacranon, and T. S. El-Bawab, "A preemptive scheduling technique for OBS networks with service differentiation," *IEEE Globecom*, pp. 2704–2708 (2003).
12. D. Q. Liu and M. T. Liu, "Priority-based burst scheduling scheme and modeling in optical burst-switched WDM networks," *Proc. Intl. Conf. Telecomm. (ICT 2002)* (June 2002).
13. D. Q. Liu and M. T. Liu, "Differentiated services and scheduling scheme in optical burst-switched WDM networks," *Proc. IEEE Intl. Conf. Networks (ICON)* (2002).
14. F. Poppe, K. Laevens, H. Michiel, and S. Molenaar, "Quality-of-service differentiation and fairness in optical burst-switched networks," *Proc. Opt. Networking Commun. Conf. (OptiComm)* (2002).
15. K. Dolzer, "Assured Horizon—an efficient framework for service differentiation in optical burst switched networks," *Proc. Opt. Networking Commun. Conf. (OptiComm 2002)* (2002).
16. M. Yang, S. Q. Zheng, and D. Verchere, "A QoS supporting scheduling algorithm for optical burst switching DWDM networks," *GLOBECOM '01 IEEE* **1**, 86–91 (2001).
17. W. H. So and Y. C. Kim, "Offset time decision for supporting service differentiation in optical burst switching networks," *Proc. COINPS* (2002).
18. J. Liu and N. Ansari, "Forward resource reservation for QoS provisioning in OBS systems," *Proc. IEEE Globecom 2002* (2002).
19. C. H. Loi, W. Liao, and D. N. Yang, "Service differentiation in optical burst switched networks," *Proc. IEEE Globecom 2002* (2002).
20. L. Yang, Y. Jiang, and S. Jiang, "A probability preemptive scheme for providing service differentiation in OBS networks," *IEEE Globecom 2003*, pp. 2689–2693 (2003).
21. X. Wang, H. Morikawa, and T. Aoyama, "Priority-based wavelength assignment algorithm for burst switched photonic networks," *OFC 2002*, pp. 765–767 (2002).
22. T. Tachibana and S. Kasahara, "QoS-guaranteed wavelength allocation for WDM networks with limited-range wavelength conversion," *IEICE Trans. Commun.* **E87-B**(6), 1439–1450 (2004).
23. J. Xu, C. Qiao, J. Li, and G. Xu, "Efficient channel scheduling algorithms in optical burst switched networks," *IEEE INFOCOM 2003*, pp. 2268–2278 (2003).
24. B. Zhou, M. Bassiouni, and G. Li, "Improving fairness in optical burst switching networks," *J. Opt. Network.* **3**(4), 214–228 (2004).
25. M. Chiu and M. Bassiouni, "Predictive schemes for handoff prioritization in cellular networks based on mobile positioning," *IEEE J. Sel. Areas Commun.* **18**(3), 510–522 (2000).
26. S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance," *IEEE/ACM Trans. Netw.* **1**(4), 397–413 (1993).
27. W. C. Park, S. J. Park, and B. H. Rhe, "RED-priority cell discarding method for improving TCP performance in ATM networks," *Joint 4th IEEE Intl. Conf. ATM (ICATM 2001) High Speed Intelligent Internet Symp.*, pp. 61–65 (2001).
28. M. May, J. Bolot, C. Diot, and B. Lyles, "Reasons not to deploy RED," INRIA, Technical Report, pp. 260–262 (June 1999).
29. M. Christiansen, K. Jeffay, D. Ott, and F. D. Smith, "Tuning RED for WEB traffic," *ACM SIGCOMM*, pp. 139–150 (2000).
30. W. Feng, D. Kandlur, D. Saha, and K. Shin, "A self-configuring RED gateway," *Proc. INFOCOM*, pp. 1320–1328 (1999).
31. C. Hollot, V. Misra, D. Towsley, and W. Gong, "A control theoretic analysis of RED," *Proc. INFOCOM*, pp. 1510–1519 (2001).
32. H. El-Aarag and M. Bassiouni, "Performance evaluation of TCP connections in ideal and non-ideal network environments," *J. Computer Commun.* **24**(18), 1769–1779, Elsevier Publishing (2001).
33. Y. Chen, C. Qiao, and X. Yu, "Optical burst switching: A new area in optical networking research," *IEEE Network*, pp. 16–23 (2004).
34. W. Peng and C. Wei, "Distributed wavelength assignment protocols with priority for WDM all-optical networks," *Proc. Computer Commun. Networks*, pp. 625–630 (2000).
35. H. Zhang, J. P. Jue, and B. Mukherjee, "A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks," *SPIE Opt. Networks Mag.* **1**(1), pp. 47–60 (2000).
36. M. El Houmaidi, M. Bassiouni, and G. Li, "Dominating set algorithms for sparse placement of full and limited wavelength converters in WDM optical networks," *J. Opt. Networking* **2**(6), 162–177 (June 2003).
37. M. El Houmaidi, M. Bassiouni, and G. Li, "Architecture and sparse placement of limited wavelength converters for optical networks," *Opt. Eng.* **43**(1), 137–147 (2004).

Bin Zhou received BS and MS degrees in computer science from the Harbin Institute of Technology, China, in 1990 and 1993, respectively. From 1993 to 1999, he worked in the China Center for Resource Satellite Data and Applications. He is currently a PhD candidate in computer science at the University of Central Florida, Orlando. His research interests include optical and wireless networks.

Mostafa A. Bassiouni received his BSc and MSc degrees in computer science from Alexandria University, and received the PhD degree in computer science from the Pennsylvania State University in 1982. He is currently a professor of computer science at the University of Central Florida, Orlando. His research interests include distributed systems, computer networks, real-time protocols, and concurrency control. He has authored more than 140 papers published in various computer journals, book chapters, and conference proceedings. His research has been supported by grants from ARO, ARPA, NSF, STRICOM, PM-TRADE, CBIS, Harris, and the State of Florida. He has served as a member of the program committee of several conferences, as the program committee chair of CSMA'98 and CSMA'2000, and as the guest co-editor of a special issue of the *Journal of Simulation Practice and Theory*.