QoS Performance of Optical Burst Switching in IP-Over-WDM Networks

Myungsik Yoo, *Member, IEEE*, Chunming Qiao, *Member, IEEE*, and Sudhir Dixit, *Senior Member, IEEE*

***Abstract*—We address the issue of how to provide basic quality of service (QoS) in optical burst-switched WDM networks with limited fiber delay lines (FDLs). Unlike existing buffer-based QoS schemes, the novel offset-time-based QoS scheme we study in this paper does not mandate any buffer for traffic isolation, but nevertheless can take advantage of FDLs to improve QoS. This makes the proposed QoS scheme suitable for the next generation Optical Internet. The offset times required for class isolation when making wavelength and FDL reservations are quantified, and the upper and lower bounds on the burst loss probability are analyzed. Simulations are also conducted to evaluate the QoS performance in terms of burst loss probability and queuing delay. We show that with limited FDLs, the offset-time-based QoS scheme can be very efficient in supporting basic QoS.**

***Index Terms*—Class isolation, fiber delay lines (FDLs), IP-over-WDM networks, priority, quality of service (QoS) performance.**

I. Introduction

THE BANDWIDTH demand on the Internet is surging phenomenally, due to emerging multimedia applications such as video conference, HDTV, Internet telephony, and digital audio. WDM networks [1]–[3] are attractive as an integral part of the infrastructure in the next generation Internet mainly due to their huge deliverable bandwidth. Recently, IP-over-WDM networks (or so called *Optical Internet*) [4] have received much attention as a promising approach to building the next generation Internet as it can reduce complexities and overheads associated with the ATM and SONET layers [5]. Having an all-optical WDM layer will enable a huge amount of "through" traffic to be switched in the optical domain, bypassing the potential bottleneck in electronic routers. It will also provide direct high speed/high bandwidth communication pipes, as well as transparency to bit rate and coding format.

In the next generation Optical Internet, one must address, among other issues, how to support quality of service (QoS) at the *WDM layer*. This is because current IP provides only *best effort* service, but mission-critical and real-time applications require a high QoS (e.g., low delay, jitter, and loss probability). Supporting basic QoS at the WDM layer will facilitate as well as complement a QoS-enhanced version of IP (such as IPv6). Furthermore, it is necessary not only for carrying some WDM layer traffic such as those for signaling and protection/restoration purposes, which require a higher priority than other ordinary traffic, but also for supporting certain applications directly (i.e., bypassing IP) or indirectly through other legacy or new protocols incapable of QoS support.

Much work has been devoted to supporting QoS in the Internet such as Integrated Services (IntServ) [6] and Differentiated Services (DiffServ) [7]. In addition, various service disciplines and scheduling algorithms, which govern per-hop behavior, have been proposed in the literature (for survey work, refer to [8], [9]). However, to date, there are no QoS schemes that take into account the unique properties of the WDM layer. Specifically, existing QoS schemes are based on packet switching, and mandate the use of buffer to isolate different classes of traffic. We call these the buffer-based schemes. In such schemes, the complexity of the scheduling algorithms is usually high. While these buffer-based schemes can provide a very fine granularity of QoS (e.g., per-flow or per-connection), they do not scale well when the number of connections is large. They are also difficult to apply to the WDM layer because of the following two key reasons. First, the use of electronic buffer would necessitate O/E and E/O conversions at intermediate nodes, thus losing data transparency. Second, optical RAM is not yet available even though a limited delay can be provided to optical packets via fiber delay lines (FDLs). Note that the use of FDLs does not allow the functions of the buffer-based schemes to be implemented fully.

The novel scheme to be studied in this paper is based on the concept of optical burst switching (OBS) [10]–[12]. OBS is a new but promising paradigm that combines the best of wavelength-routing and optical packet-switching in order to facilitate efficient integration of IP and WDM. The proposed QoS scheme uses an *extra offset time* instead of buffer to isolate classes of traffic. We call it the offset-time-based scheme, which is suitable for implementation in bufferless WDM networks as it does not mandate the use of any buffer (though it can take advantage of FDLs to improve QoS performance). By eliminating the complex queuing operations required in the buffer-based schemes, the complexity of the offset-time-based scheme can be much lower. In addition, the proposed scheme manages QoS on a class by class basis with a few priority levels and thus can be more scalable than the buffer-based schemes. Note that given the possibility of putting 40 or 80 separate wavelengths on a single fiber, each supporting 2.5–10 Gb/s (or OC-48–OC-192), a few priority levels (e.g., 2 to 8) should be enough provided that a proper QoS scheme (such as the offset-time-based scheme) can allocate sufficient bandwidth to high priority classes.

Manuscript received October 22, 1999; revised May 15, 2000. This work was supported in part by NSF under Contract ANIR-9801778.

M. Yoo was with the Department of Electrical Engineering, SUNY at Buffalo, NY, and Nokia Research Center, Burlington, MA USA. He is now with the School of Electronic Engineering, Soongsil University, Seoul, Korea.

C. Qiao is with the Department of Computer Science and Engineering, SUNY at Buffalo, Buffalo, NY 14260 USA.

S. Dixit is with Nokia Research Center, Burlington, MA 01803 USA.

Publisher Item Identifier S 0733-8716(00)09023-5.

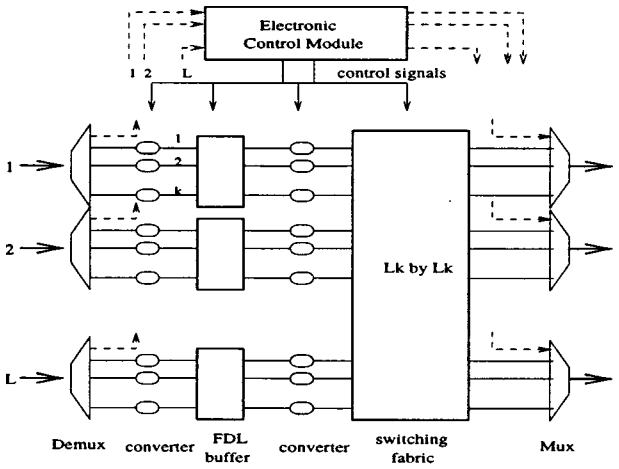

Fig. 1. Structure of optical switching node.

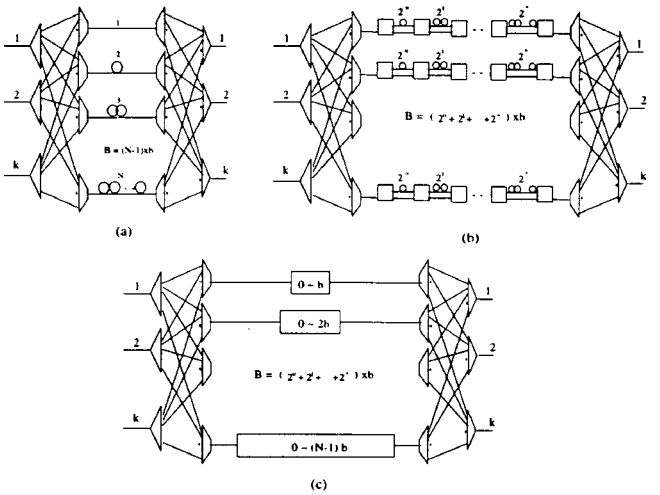

Fig. 2. Structure of FDL buffers: (a) fixed-delay, (b) variable-delay, and (c) hybrid.

The offset-time-based QoS scheme has been studied and evaluated in [13] and [14] assuming no FDLs. In this paper, we study how the scheme can be applied when FDLs are used and evaluate the resulting improvement in the QoS performance. In the following discussion, a set of FDLs used to provide a limited delay at an optical switching node will be referred to as a FDL buffer.[1] Many studies on such FDL buffers have focused on "time slotted" optical packet-switched networks [15]–[17] with synchronous switching of fixed size packets (or ATM cells). In this work, we focus on burst-switched WDM networks with asynchronous switching of variable size bursts [10]–[12].

The rest of the paper is organized as follows. Section II describes typical structures of optical switching nodes and FDL buffers under consideration. Section III presents the offset-time-based QoS scheme when FDL buffers are used. Section IV discusses the extra offset time required to achieve class isolation as well as related pretransmission delay and queuing delay. In Section V, we analyze the upper and lower bounds on the burst loss probability. The results obtained from analysis and simulations are discussed in Section VI. Finally, concluding remarks are presented in Section VII.

II. Optical Switching Node and FDL Buffer

We consider IP-over-WDM networks, where IP edge routers connect to optical switching node (OSN). Inside the WDM backbone network, each OSN is interconnected with neighboring OSNs through fiber links. Fig. 1 illustrates possible structure of an OSN. There are L incoming (and outgoing) fiber links, each of which has k wavelengths for carrying data bursts (solid lines) and one (or more) additional wavelength for carrying control packets (dotted lines). Every control packet is processed by the electronic control module inside an OSN, which generates appropriate control signals to set up the wavelength converters, FDL buffers, and switching fabric. The optical switching fabric, which functions as an Lk: Lk (i.e., Lk by Lk) nonblocking switch, switches each burst on an incoming wavelength as it arrives (i.e., without having to synchronize it with other incoming bursts).

[1]The meaning of "buffer" here is different from that in the ordinary electronic buffer.

The FDL buffers shown in Fig. 1 may be implemented as shown in Fig. 2. In Fig. 2(a), each FDL buffer has N delay lines, each of which provides a different, but fixed delay ranging from zero to the maximum buffering time (delay) $B = (N-1) \times b$, where b is the delay that a single fiber delay element (represented by a circle in the Fig. 2) can provide. This structure will be referred to as the fixed-delay FDL buffer. The FDL buffer in Fig. 2(b) also has the same number of delay lines N, but each delay line can provide a variable delay ranging from zero to $B = (N-1) \times b = (2^0 + 2^1 + \cdots + 2^n) \times b$ by using multiple wavelength-sensitive $2 : 2$ switches. This structure will be referred to as the variable-delay FDL buffer. The fixed-delay FDL buffer is simpler (i.e., has a lower cost which is proportional to the number of $2 : 2$ switches), but less effective than the variable-delay FDL buffer. The FDL buffer shown in Fig. 2(c) is a hybrid of Fig. 2(a) and (b) where each FDL can provide a variable delay, but has a different maximum delay ranging from b to $(N-1) \times b$. Note that all the FDL buffers shown in Fig. 2 have k $1 : N$ switches followed by N $k : 1$ multiplexers at the input side, and N $1 : k$ demultiplexers followed by k $N : 1$ switches connected at the output side. This allows one physical FDL to be shared by all k wavelengths, which effectively creates k virtual FDLs per physical FDL. Also, two stages of wavelength converters before and after the FDL buffer as shown in Fig. 1 are useful since they allow a burst to use different wavelengths in a FDL buffer and on an output link, thus resolving contentions among multiple incoming bursts for the same FDL and link more effectively. In the rest of the paper, we will concentrate on the hybrid FDL buffer and simply refer to it as the FDL buffer, though our discussions also apply to other FDL buffers.

III. The Offset-Time-Based QoS Scheme with FDL

A. Basics of OBS

Before we discuss the offset-time-based QoS scheme in detail, we first review how resources such as wavelengths and FDL buffers are reserved in the non-QoS capable optical burst switching (OBS) (hereafter called classless OBS), where no pri-

ority is given to any particular class of traffic. At a source edge router, a "base" offset time is calculated based on the total processing time to be encountered by a control packet (or request)[2] on its way to the destination edge router [12]. The basic idea is that the control packet will be sent first, but the corresponding burst (e.g., several IP packets) will be buffered at the source edge router for this base offset time. While the control packet is processed at each and every intermediate OSN to establish an all-optical path, the burst will cut through the preconfigured OSNs along the established path.

As to be discussed in the next subsection, to give a high priority to a burst, an "extra" offset time is also determined. The control packet carries the value of the total offset time (which is the sum of the base and extra offset time) along with other information such as destination address (or label), and burst duration. As a result of processing the control packet at each intermediate OSN, a wavelength on the desired output link will be reserved using *delayed reservation (DR)* [12], where reservation is made from the burst arrival time to the burst departure time. If no wavelength is available, the burst will be blocked and one FDL will be reserved using DR as well.[3] Note that in OBS, since every wavelength reservation is made with a known duration, the blocking time of the burst can easily be determined. Based on the blocking time, the FDL with an appropriate length (delay time) will be reserved. If the blocking time is larger than the maximum delay time B, no FDLs will be reserved and the burst will be simply dropped without wasting any buffer capacity. To facilitate our discussion, we assume that the base offset time is negligible compared to the extra offset time, and will refer to the latter as simply the offset time hereafter.

B. Isolation of Two Traffic Classes

In this subsection, we describe the QoS-capable OBS (called prioritized OBS) when applied to WDM networks with FDLs. For simplicity, we consider two classes of traffic, i.e., 0 and 1, where class 1 has a higher priority over class 0. The number of classes will be generalized to an arbitrary integer n, where $n > 2$, in Section IV. To help explain the offset-time-based QoS scheme, we distinguish between two different contentions in reserving resources (wavelengths and FDLs): the intraclass contentions caused by requests within the same class, and the interclass contentions caused by requests from different classes. Since our QoS scheme is intended to isolate class 1 from class 0, we will focus on how to resolve interclass contentions using the offset time in the following discussion.

Specifically, in order for class 1 to have a higher priority for reserving resources such as wavelengths and FDLs, an *extra* offset time, denoted by t_o^1, is given to class 1 traffic (but not to class 0 traffic, i.e., $t_o^0 = 0$).[4] To simplify our presentation, consider a single data wavelength and a single FDL, which can provide a variable delay between 0 and B. Fig. 3 illustrates why having t_o^1 can give a class 1 request a higher priority over a class 0 request when reserving resources. Specifically, let t_a^i be the arrival time of a class i request denoted by $req(i)$, and t_s^i be the service-start time (i.e., the arrival time of the corresponding data burst) whose burst length will be denoted by l_i (where $i = 0, 1$). Then we have $t_s^1 = t_a^1 + t_o^1$ and $t_s^0 = t_a^0$ (since $t_o^0 = 0$).

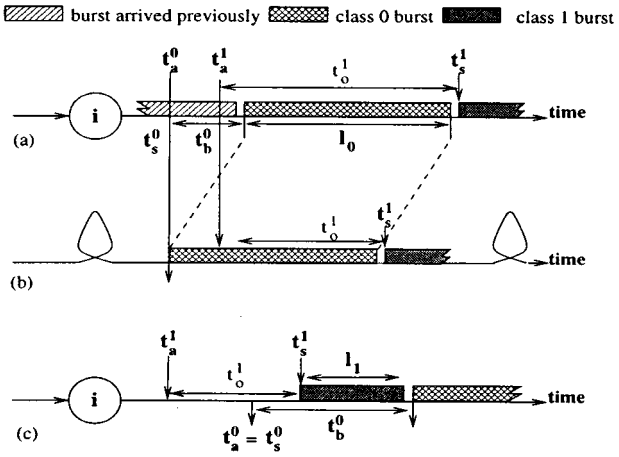

Fig. 3. Class isolation in resource reservation due to offset time.

Let us assume that when $req(0)$ arrives at t_a^0 $(= t_s^0)$, the only data wavelength is in use by a burst that arrived earlier. Thus, the burst corresponding to $req(0)$ has to be delayed by a blocking time, t_b^0 units. If $t_b^0 < B$, the FDL will be reserved for the class 0 burst from t_s^0 to $t_s^0 + l_0$ as shown in Fig. 3(b),[5] and the wavelength will be reserved from $t_s^0 + t_b^0$ to $t_s^0 + t_b^0 + l_0$ as shown in Fig. 3(a). Now assume that $req(1)$ arrives later at t_a^1 (where $t_a^1 > t_a^0$) and tries to reserve the wavelength for the corresponding class 1 burst. Clearly, $req(1)$ will succeed in reserving the wavelength as long as the offset time t_o^1 is long enough so that $t_s^1 = t_a^1 + t_o^1 > t_a^0 + t_b^0 + l_0$. Note that had $req(1)$ arrived earlier than $req(0)$ as shown in Fig. 3(c), $req(1)$ would succeed in reserving the wavelength (from t_s^1 to $t_s^1 + l_1$ using delayed reservation), but $req(0)$ would not. This shows that class 1 can be isolated from class 0 when reserving "wavelength" due to the use of the offset time.

Similarly, Fig. 3(b) illustrates class isolation in "FDL" reservation due to the use of the offset time. More specifically, let us assume that $req(0)$ has reserved the FDL as described earlier, and in addition, because t_o^1 is not long enough (i.e., if $t_s^1 = t_a^1 + t_o^1 < t_a^0 + t_b^0 + l_0$), $req(1)$ would be blocked in wavelength reservation and thus needs to reserve the FDL in order to delay the corresponding class 1 burst. In such a case, $req(1)$ will successfully reserve the FDL if t_o^1 is nonetheless long enough to have $t_s^1 = t_a^1 + t_o^1 > t_a^0 + l_0$. Intuitively, this is because the head of class 1 burst will be able to enter the FDL after the tail of class 0 burst enters. Otherwise (i.e., if $t_s^1 < t_a^0 + l_0$), $req(1)$ will not succeed and the class 1 burst will be dropped. Note that if $req(1)$ arrives first, it can reserve the FDL, but $req(0)$ will be blocked (and dropped) (either when $t_a^1 < t_a^0 < t_s^1$ but $t_a^0 + l_0 > t_s^1$, or when $t_s^1 < t_a^0 < t_s^1 + l_1$).

The above discussion implies that class 1 can be isolated from class 0 in reserving both wavelengths and FDLs by using an appropriate offset time, which essentially gives class 1 a higher priority over class 0. Accordingly, it is possible to allocate the

[2]Request and control packet will be used interchangeably hereafter.

[3]Dropped if no FDLs are used or available.

[4]Intuitively, giving an offset time is similar to allowing someone to purchase a ticket to certain event *in advance* instead of *at the door* only.

[5]The burst will be dropped if t_b^0 exceeds B.

entire resources to class 1 when necessary (e.g., when congestion occurs). On the other hand, as a result of having a low priority in resource reservation, class 0 bursts will have a relatively high blocking (and loss) probability.

IV. Offset Time and End-to-End Delay

A. Offset Time Requirement

In this subsection, we first quantify the offset times required for class isolation when reserving wavelengths and FDLs. Consider a general case where the number of classes is an arbitrary integer $n > 1$. Let $t_{i,j}(> 0)$ denote the difference in the offset times assigned to class i and class j (i.e., $t_{i,j} = t_o^i - t_o^j$), where $0 \le j < i \le n-1$, and $0 \le R_{i,j} \le 1$ denote the probability that $req(i)$ will not be blocked by $req(j)$ (arrived earlier). Hereafter, $t_{i,j}$ and $R_{i,j}$ will be referred to as the *offset time difference* and the *degree of isolation*, respectively, between class i and class j.

We may generalize the discussion related to Fig. 3 as follows. In order for class i to be completely isolated from class $(i-1)$ (i.e., $R_{i,i-1} = 1.0$) in FDL reservation, $t_{i,i-1}$ needs to be larger than the maximum burst length of class $(i-1)$ since in the worst case one may have $t_a^i \approx t_a^{i-1}$. Similarly, to achieve complete class isolation in wavelength reservation, $t_{i,i-1}$ needs to be larger than the sum of B and the maximum burst length of class $(i-1)$ since in the worst case one may have $t_a^i \approx t_a^{i-1}$ and t_b^{i-1} can be as large as B. This implies that for a given offset time difference $t_{i,i-1}$, different degrees of isolation between class i and class $(i-1)$ will be achieved in FDL and wavelength reservations.

More specifically, in FDL reservation, the relationship between $R_{i,i-1}$ and $t_{i,i-1}$ only depends on the distribution of the burst length of class $(i-1)$. Let us assume that the burst length (or service time) of class $(i-1)$ is exponentially distributed with a mean L_{i-1}. Then, the probability that a class $(i-1)$ burst has a length no longer than a given value t is $(1 - e^{-\mu_{i-1} \cdot t})$, where $\mu_{i-1} = 1/L_{i-1}$. Given that a class i burst will not be blocked as long as the length of the class $(i-1)$ burst arrived earlier is no greater than $t_{i,i-1}$, we have $R_{i,i-1} = (1 - e^{-\mu_{i-1} \cdot t_{i,i-1}})$, or

$$t_{i,i-1} = t_o^i - t_o^{i-1} = -\ln(1 - R_{i,i-1}) \cdot L_{i-1}. \qquad (1)$$

In wavelength reservation, a class i burst will not be blocked as long as the length of class $(i-1)$ burst is no greater than $t_{i,i-1} - B$. Hence, we have $R_{i,i-1} = (1 - e^{-\mu_{i-1} \cdot (t_{i,i-1} - B)})$, or $t_{i,i-1} = -\ln(1 - R_{i,i-1}) \cdot L_{i-1} + B$. Table I shows some of the values of $t_{i,i-1}$ needed to achieve a given $R_{i,i-1}$ in FDL and wavelength reservations. For example, in order to achieve 99% isolation in FDL reservation, we need $t_{i,i-1} = 5 \cdot L_{i-1}$. However, in wavelength reservation, only 95% degree of isolation is achieved for the same $t_{i,i-1}$ assuming $B = 2 \cdot L_{i-1}$. Note that if class i is isolated from class $(i-1)$ in reserving FDL (or wavelengths) with $R_{i,i-1}$, then it is also isolated from class $(i-2)$ with a higher degree, that is, $R_{i,i-2} > R_{i,i-1}$, because $t_{i,i-2} > t_{i,i-1}$. Similarly, class i is also isolated from all other lower priority classes with $R_{i,j} > R_{i,i-1}$, where $j = i-3, i-4, \ldots, 0$.

TABLE I
Isolation Degree and Offset Time Difference

$R_{i,i-1}$	0.6321	0.9502	0.9932
$t_{i,i-1}$ (FDL)	L_{i-1}	$3 \cdot L_{i-1}$	$5 \cdot L_{i-1}$
$t_{i,i-1}$ (λ)	$L_{i-1} + B$	$3 \cdot L_{i-1} + B$	$5 \cdot L_{i-1} + B$

B. Pretransmission and Quereing Delays

For real-time applications, it is imperative to discuss the end-to-end delay, which consists of three components in the offset-time-based QoS scheme: propagation delay from source to destination P, pretransmission delay t_p (due to offset time), and quereing delay in the FDLs t_q (due to blocking). Among them, P is not related to the use of the offset time.

Although t_p is affected by the use of offset time, we first note that the use of a "base" offset time [12] does not increase the end-to-end delay as compared to packet switching because it merely substitutes for the total processing delay to be encountered by the corresponding control packet. On the other hand, when compared to circuit switching, the pretransmission delay (and thus end-to-end delay) is about $2P$ shorter (which translates into about tens of milliseconds if a burst travels coast to coast) as the offset-time-based QoS scheme is based on one-way reservation.

However, the use of an "extra" offset time, which isolates a high priority class from lower priority classes, does increase the end-to-end delay. In fact, when using the extra offset time, the additional pretransmission delay is a function of the number of classes n and the offset time difference $t_{i,i-1}$ (or $R_{i,i-1}$). Assuming that all the offset time differences are equal (i.e., $t_{n-1,n-2} = t_{n-2,n-3} = \cdots = t_{1,0} = t_{\text{diff}}$), then the longest additional pretransmission delay (equal to the longest extra offset time t_o^{n-1}), which is incurred by the highest priority class $(n-1)$, is $t_p = t_o^{n-1} = (n-1) \cdot t_{\text{diff}}$. Given that in today's voice and video communications, the interactive applications require a few hundreds of milliseconds end-to-end delay bounds, such delay may be negligible. As an example, assume that all classes have equal mean burst length L (i.e., $L_{n-1} = L_{n-2} = \cdots = L_0 = L$), and the average burst size is 15 kbytes (containing few tens of average sized IP packets), which at 10 Gb/s, results in $L = 12\ \mu$s. When $n = 6$ and $t_{\text{diff}} = 5L$ (corresponding to a high degree of isolation), the longest pretransmission delay is $t_o^5 = 300\ \mu$s. Note that to support such delay-sensitive applications using the offset-time-based QoS scheme, the number of classes and the average burst length should be considered as key design parameters.

On the other hand, the quereing delay t_q, which is a function of blocking probability (and maximum delay time of an FDL B), reduces with the offset time difference. This is because as the offset time difference increases, the degree of isolation increases, resulting in reduced blocking probability and thus low quereing delays for high priority classes. Note that as far as end-to-end delay concerns, t_p is a fixed delay, whereas t_q is a variable delay. It implies that in the offset-time-based QoS scheme, high priority classes incur low t_q by introducing t_p, which can be negligible as described above. Thus, by virtue

of the extra offset time, the end-to-end delay of high priority classes can be low as well as predictable.

V. Analysis

In this section, we analyze the loss probabilities of two OBS protocols, namely *classless* OBS and *prioritized* OBS. As mentioned earlier, the latter distinguishes n classes where class i has higher priority (longer offset time) over class j if $i > j$, while the former does not. Assume that class i generates bursts according to an exponentially distributed arrival rate with an average λ_i and an exponentially distributed service rate with an average $\mu_i = 1/L_i$. The traffic intensity of class i is thus given by $\rho_i = \lambda_i/(\mu_i \cdot k) = r_i/k$, where $r_i = \lambda_i/\mu_i$, and the total traffic intensity is $\rho = \Sigma_{i=1}^{n} \rho_i$. For simplicity, we consider an OSN similar to the one shown in Fig. 1, and focus on a given output link having k wavelengths and a single FDL buffer containing N FDLs.

Let us first derive the lower and upper bounds on the loss probability in classless OBS. In classless OBS, the upper bound on the loss probability pb_u is reached when there is no FDL buffer, i.e., $B = 0$, and this upper bound is determined using the following Erlang's loss formula ($M/M/k/k$) [18], where $r = \rho \cdot k$

$$pb_u = \mathcal{B}(k, \rho) = \frac{1/k! \cdot r^k}{\sum_{m=0}^{k} 1/m! \cdot r^m}. \tag{2}$$

The lower bound is determined by assuming that the switch has an FDL buffer which can be roughly modeled as $M/M/k/D$ (where D is to be determined). More specifically, let N denote the number of *physical* FDLs in the FDL buffer as shown in Fig. 2(c). Since the number of *virtual* FDLs is $d = N \cdot k$, the maximum number of bursts in the switch (including those being transmitted and those being delayed) is $D = k + d = k + k \cdot N$. According to the $M/M/k/D$ model [19], the lower bound on the loss probability pb_l is

$$pb_l = \mathcal{PB}(k, \rho, D) = \frac{r^D}{k^{D-k} \cdot k!} \cdot p_0 \tag{3}$$

where

$$p_0 = \left(\sum_{n=0}^{k-1} \frac{r^n}{n!} + \sum_{n=k}^{D} \frac{r^n}{k^{n-k} \cdot k!} \right)^{-1}.$$

We now determine the lower and upper bounds on the loss probability of class i in prioritized OBS, denoted by pb_l^i and pb_u^i, respectively, where $i = n-1, n-2, \ldots, 0$. For the following derivation, it is assumed that class i is completely isolated from class $(i-1)$, i.e., $R_{i,i-1} = 1.0$. In addition, assume that the *conservation law* holds,[6] that is, the overall performance (i.e., loss probability and throughput averaged over all classes) of the network stays the same regardless of the number of classes and the degree of isolations. In particular, let $\rho_{n-1,j}$ be the sum of the traffic intensity from class $(n-1)$ through class j, i.e., $\rho_{n-1,j} = \Sigma_{i=j}^{n-1} \rho_i$. Due to class isolation, traffic from lower priority classes (class $(j-1)$ through class 0) does not affect the loss probability of higher priority classes (class $(n-1)$ through class j). In other words, the loss probability averaged over class $(n-1)$ through class j, denoted by $pb^{n-1,j}$, only depends on the sum of traffic intensity from class $(n-1)$ through class j. Thus, according to the conservation law, one may apply (2) and (3) to obtain the upper and lower bounds as $pb_u^{n-1,j} = \mathcal{B}(k, \rho_{n-1,j})$ and $pb_l^{n-1,j} = \mathcal{PB}(k, \rho_{n-1,j}, D)$, respectively.

We may also represent the average loss probabilities using a weighted sum of loss probability of each class. More specifically, let $c_i = \rho_i/\rho$ be the ratio of the traffic intensity in class i over the total traffic intensity. Then the averaged upper and lower bounds on the loss probability over class $(n-1)$ through class j can also be written as $pb_u^{n-1,j} = \Sigma_{i=j}^{n-1} c_i \cdot pb_u^i$ and $pb_l^{n-1,j} = \Sigma_{i=j}^{n-1} c_i \cdot pb_l^i$, respectively. Again, according to the conservation law, we have

$$\mathcal{B}(k, \rho_{n-1,j}) = \sum_{i=j}^{n-1} c_i \cdot pb_u^i \tag{4}$$

for upper bound, and

$$\mathcal{PB}(k, \rho_{n-1,j}, D) = \sum_{i=j}^{n-1} c_i \cdot pb_l^i \tag{5}$$

for lower bound.

We can now obtain each pb_u^i and pb_l^i by starting with the highest priority class (i.e., $i = n-1$). Since class $(n-1)$ has the highest priority, and is completely isolated from any other classes, the upper bound (when a bufferless switch is used) is $pb_u^{n-1} = \mathcal{B}(k, \rho_{n-1})$ as in (2), while the lower bound (when a buffered switch is used) is $pb_l^{n-1} = \mathcal{PB}(k, \rho_{n-1}, D)$ as in (3). To obtain the upper bound for class $(n-2)$, we have

$$pb_u^{n-1,n-2} = \mathcal{B}(k, \rho_{n-1,n-2})$$

and

$$pb_u^{n-1,n-2} = \sum_{i=n-2}^{n-1} c_i \cdot pb_u^i$$

according to (4). By equating the two, we obtain

$$pb_u^{n-2} = \frac{\mathcal{B}(k, \rho_{n-1,n-2}) - c_{n-1} \cdot pb_u^{n-1}}{c_{n-2}}.$$

Similarly, according to (5), we obtain the lower bound as

$$pb_l^{n-2} = \frac{\mathcal{PB}(k, \rho_{n-1,n-2}, D) - c_{n-1} \cdot pb_l^{n-1}}{c_{n-2}}.$$

This procedure can be applied repeatedly. In general, the upper bound on the loss probability for a given class j, where $0 \le j \le n-2$, is

$$pb_u^j = \frac{\mathcal{B}(k, \rho_{n-1,j}) - \sum_{i=j+1}^{n-1} c_i \cdot pb_u^i}{c_j} \tag{6}$$

[6] As to be discussed later, our simulation results have verified that this assumption is valid when the traffic intensity is high.

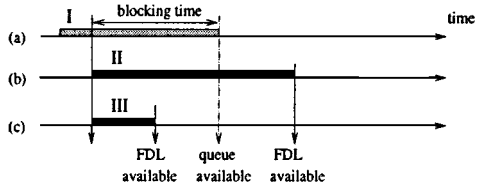

Fig. 4. Difference between an FDL and a queue.

and the corresponding lower bound is

$$pb_l^j = \frac{\mathcal{PB}(k, \rho_{n-1,j}, D) - \sum_{i=j+1}^{n-1} c_i \cdot pb_l^i}{c_j}. \quad (7)$$

Before we present simulation results in the next section, it is worth noting that a FDL cannot be accurately modeled as a queue. This is because a queue in the $M/M/k/D$ model can store any burst regardless of its blocking time, whereas an FDL has to drop a burst if its blocking time is longer than B. Of course, as B increases, it is expected that such a difference between the two will diminish. However, regardless of the value of B, the time an FDL becomes available may differ from the time a queue in $M/M/k/D$ becomes available as illustrated in Fig. 4.

More specifically, assuming that a burst I is using the only available wavelength as shown in Fig. 4(a) when another burst II (or III) arrives and subsequently reserves an FDL as shown in Fig. 4(b) [or (c)]. In the $M/M/k/D$ model, a queue becomes available (to other bursts) as soon as burst I releases the wavelength to burst II (or III) (i.e., exactly at the end of the blocking time indicated by the dashed vertical line). However, the FDL occupied by burst II (or III) becomes available (to other bursts) as soon as the tail of burst II (or III) enters the FDL. In other words, the time FDL becomes available depends on the length of burst II (or III) and could be later (or earlier) than the time a queue becomes available as shown in Fig. 4(b) [or (c)]. Thus, if the average burst length is longer than the average blocking time, the $M/M/k/D$ model tends to predict a lower blocking probability than that obtainable from simulations of the FDL. Such a case is likely to occur when the maximum delay time B is relatively short as compared to the average burst length, and/or when the number of wavelengths k is small, which results in a relatively long average blocking time. On the other hand, if the average burst length is shorter than the average blocking time, the simulated blocking probability can be lower than that predicted by the $M/M/k/D$ model, and such a case is likely to occur when B is relatively long as compared to the average burst length, and/or when k is large, which results in a relatively short blocking time.

VI. Results and Discussion

The results from analysis and simulation are presented in this section. We consider two traffic models, Poisson and self-similar traffic. We simulate Poisson traffic, which facilitates the analysis in the previous section and yields a lower bound on the loss probability and quereing delay under self-similar traffic. To simulate self-similar traffic, Pareto distribution [20] and exponential distribution are used to generate the burst inter-arrival time and the burst length, respectively.[7]

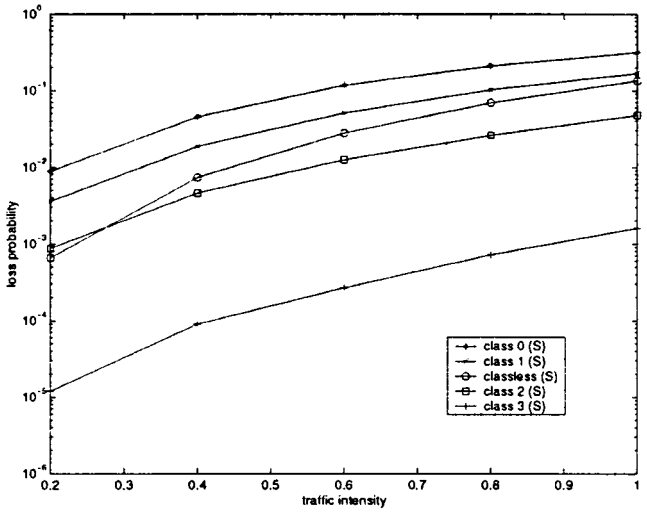

Fig. 5. Class loss probability.

The simulation (and analysis) results are obtained for a single OSN with an outgoing link having k wavelengths and an FDL buffer containing N FDLs. Without loss of generality, we assume that all sources (classes) generate an equal amount of traffic (i.e., $\rho_0 = \rho_1 = \cdots = \rho_{n-1} = \rho/n$). In addition, it is assumed that all classes have equal mean burst length (i.e., $L_{n-1} = L_{n-2} = \cdots = L_0 = L$), and all offset time differences are equal (i.e., $t_{n-1,n-2} = t_{n-2,n-3} = \cdots = t_{1,0} = t_{\text{diff}}$). The performance metrics are the burst loss probability and quereing delay as a function of traffic intensity ρ, the maximum delay time B, the offset time difference t_{diff}, the number of FDLs N, the number of classes n, and the number of wavelengths k.

A. Service Differentiation and Conservation Law

Fig. 5 plots the simulated loss probability as a function of the traffic intensity when $n = 4$, $k = 1$, $N = 3$ ($d = 3$), $B = 3L$, and $t_{\text{diff}} = 3L + B$. As can be observed by comparing the loss probabilities of the classes 0 to 3 in prioritized OBS with the average loss probability (over all classes) in classless OBS, service differentiation is achieved by taking advantage of the offset time. More specifically, class 2 and class 3 have a lower loss probability than the classless case (or average), while class 0 and class 1 have a higher loss probability because of their low priority. As to be shown later, the reduction in the loss probabilities of classes 2, 3 over the average increases with the maximum delay time B (see Section VI-B), the offset time difference t_{diff} (see Section VI-C), and other parameters such as the number of FDLs, the number of classes, and the number of wavelengths (see Section VI-D). In particular, when the number of wavelengths is large, the loss probability of class 1 can also be lower than that of the classless case as will be shown in Section VI-D. However, as the offset time difference decreases, the loss

[7]Simulation results under self-similar traffic are not presented due to space constraints. It is shown that the offset-time-based QoS scheme still achieves a good service differentiation under self-similar traffic (with Hurst parameter set to 0.8), although burstiness of traffic causes lower performance than the case under Poisson traffic.

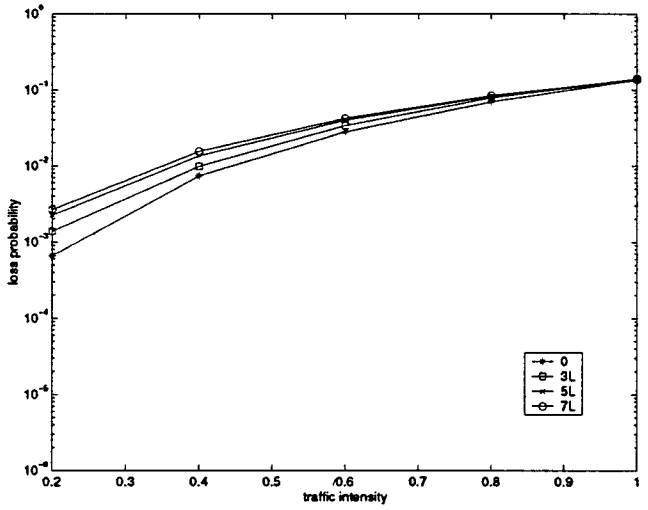

Fig. 6. Overall loss probability.

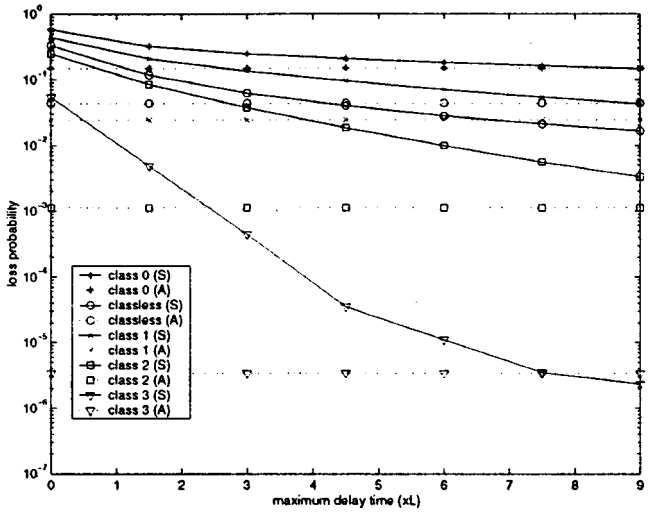

Fig. 7. Maximum delay time versus loss probability.

probability of each class in prioritized OBS reduces to that of the classless case.

With all other parameters being the same, the average loss probabilities for various offset time differences t_{diff} are plotted in Fig. 6. When the traffic intensity is low ($\rho < 0.8$), prioritized OBS tends to have a higher average loss probability than classless OBS. However, when $\rho \geq 0.8$, the average loss probabilities for various offset time differences tend to match closely with each other, implying that the assumed conservation law holds well. Since all the subsequent simulation (and analysis) results are obtained with a high traffic intensity (e.g., $\rho = 0.8$), it is justified to apply the conservation law. Accordingly, we may regard the results for classless OBS to be the same as the average over all classes in prioritized OBS.

B. The Role of the Maximum Delay Time of FDL

In addition to achieving service differentiation, it is important to keep the quereing delay (and end-to-end delay) low as we introduce FDL buffers. Thus, although a long maximum delay time (or FDL length) can reduce the loss probability, one needs to keep it as short as possible. In this subsection, we show the effect of maximum delay time on QoS performance such as loss probability and quereing delay, and, in particular, how service differentiation can reduce the quereing delay of high priority classes even with long FDLs.

Fig. 7 shows the loss probability as a function of the maximum delay time B. The parameters used are $n = 4$, $k = 2$, $N = 3$ ($d = 6$), $\rho = 0.8$, and $t_{\text{diff}} = 3L + B$. The dotted and solid lines show the theoretical *lower bound* obtained by analysis and the actual loss probability obtained from simulations under Poisson traffic, respectively. As can be seen, the loss probability of each class is the highest when B is zero (i.e., with a bufferless switch). As B increases, the loss probability approaches to, and even goes slightly below, the lower bound (for class 3 and the average). The latter phenomenon is due to the difference between an FDL and a queue as explained in Fig. 4. We also observe that the higher the priority of a class is, the more the benefit is due to the use of FDLs.

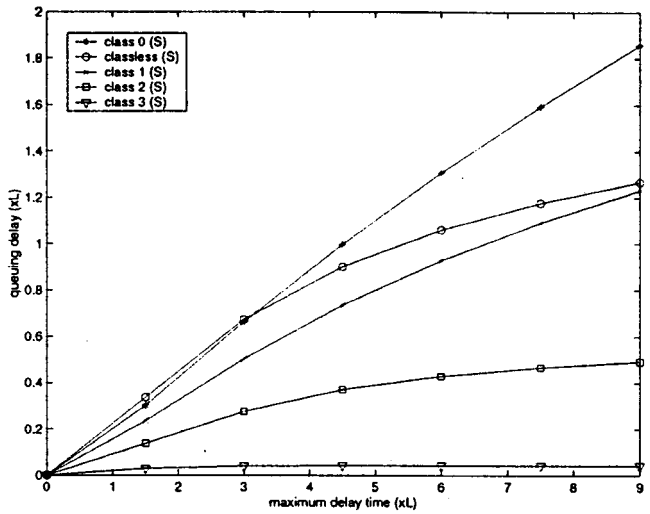

Fig. 8. Maximum delay time versus quereing delay.

The quereing delay experienced by each class (averaged over all bursts in that class) is plotted in Fig. 8 with the same set of parameters as those used in Fig. 7. Since the quereing delay is proportional to the blocking probability, it is not surprising that the quereing delays of different classes in prioritized OBS are differentiated, and a higher priority class has a lower quereing delay. In addition, although the quereing delay for a given class increases with B, the high priority classes (especially class 3) encounter a low quereing delay (less than $0.02L$) even when B is as long as $9L$ due to their low blocking probabilities achieved by class isolation.

Although not shown (due to space constraints), other simulation results show that the maximum delay time B required to achieve a given loss probability can be significantly reduced as the number of wavelengths increases.

C. The Role of the Offset-Time Difference

As mentioned earlier in Section IV, pretransmission delay is proportional to the offset time difference used for class isolation, and thus it is important to see the effect of the offset time difference on quereing delay as well as loss probability.

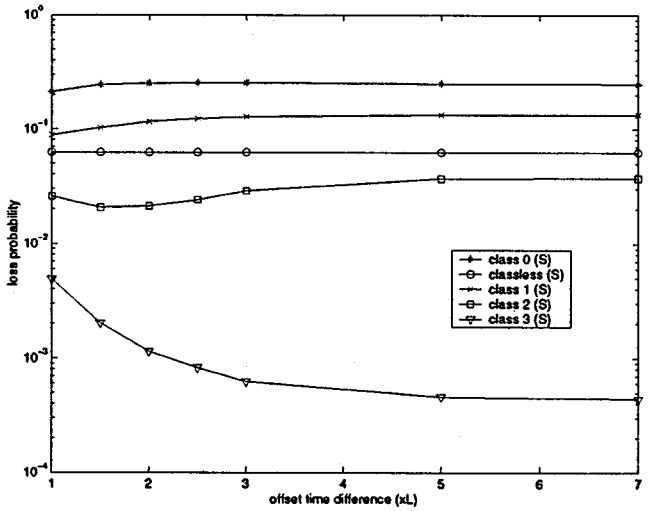

Fig. 9. Offset time difference versus loss probability.

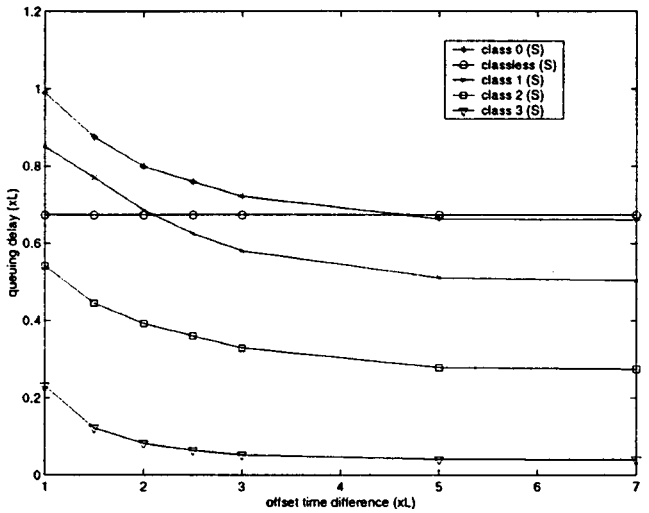

Fig. 10. Offset time difference versus quereing delay.

The loss probability and quereing delay as a function of the offset time difference t_{diff} are plotted in Figs. 9 and 10, respectively, for $n = 4$, $k = 2$, $N = 3$ ($d = 6$), $B = 3L$, and $\rho = 0.8$. The classless case (average) has a constant loss probability (and quereing delay) because it is independent of the offset time difference. For prioritized classes, the offset time difference plays an important role in service differentiation, as it will result in a low loss probability (and quereing delay) in high priority classes. As can be seen in Fig. 9, $t_{\text{diff}} < 3L$ is sufficient for high priority classes to achieve a high enough class isolation. Fig. 10 shows that the quereing delays of prioritized classes decrease with the offset time difference. Specifically, classes 2 and 3 have a lower quereing delay than the average due to their lower blocking probability. Note that the quereing delays of classes 0 and 1 tend to saturate with a large offset time difference. This is because while the low priority classes have a longer blocking time caused by higher priority classes as the offset time difference increases, the bursts whose blocking time is longer than B are dropped (and not considered when determining the average quereing delay).

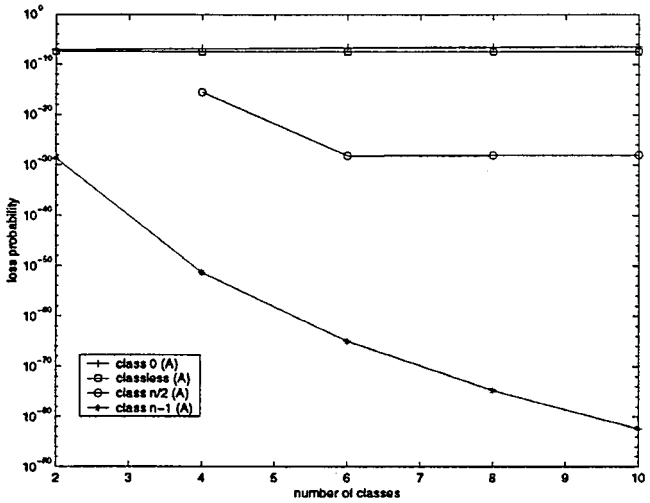

Fig. 11. Loss probability versus number of classes.

Although not shown due to space limit, other simulation results show that as the number of wavelengths increases, the same isolation degree can be maintained with significantly reduced value of offset time difference.

D. The Effectiveness of the Offset-Time-Based Scheme

In the previous subsections, we have presented the QoS performance results when the number of wavelengths (and classes) is small because otherwise the simulations take too much time to converge, especially for cases with low loss probabilities (e.g., $\leq 10^{-6}$). In this subsection, we present analytical results to evaluate the effectiveness of the offset-time-based QoS scheme with a larger number of wavelengths (and classes).

Fig. 11 shows the lower bound on the loss probability as a function of the number of classes n, when $k = 16$, $N = 4$ ($d = 64$), and $\rho = 0.8$. Since the average loss probability is independent of the number of classes, it is shown as a straight line. For prioritized classes, while the loss probability of class 0 stays a little above than the average, the loss probability of class $n/2$ and class $(n-1)$ drops as n increases. This means that higher priority classes receive better service when more classes exist. This is because even though the overall traffic intensity ρ remains the same, the traffic intensity of a high priority class, ρ_{n-1} for example, decreases proportionally when n increases. Note that the loss probability of class $n/2$ almost stops decreasing at $n = 6$. This is because when n increases, the relative priority of class $n/2$ drops as more higher priority classes undermine classes below $n/2$.

The lower bound on the loss probability as a function of the number of FDLs N is shown in Fig. 12, when $n = 4$, $k = 16$, and $\rho = 0.8$. As can be seen, all classes except the lowest priority class (class 0) have a lower loss probability than the average. Even when $N = 5$, the reduction in the loss probability as compared to the average is more than 10, 20, and 40 orders of magnitude for classes 1, 2, and 3, respectively. This means that with a limited number of physical FDLs, the QoS performance can be significantly improved when the number of wavelengths is fairly large.

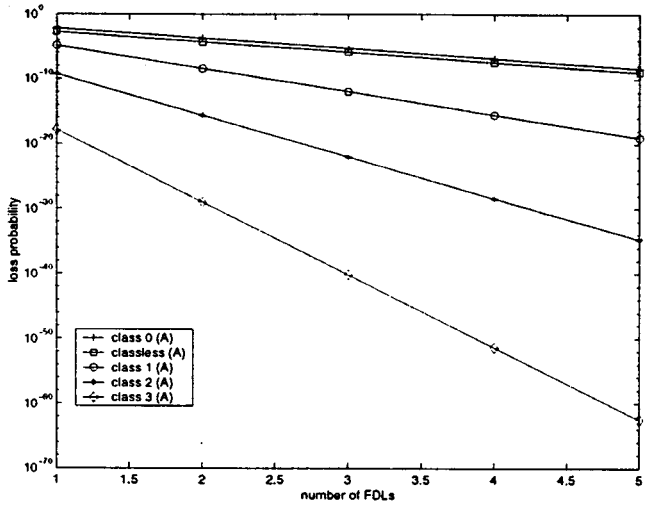

Fig. 12. Loss probability versus number of FDLs.

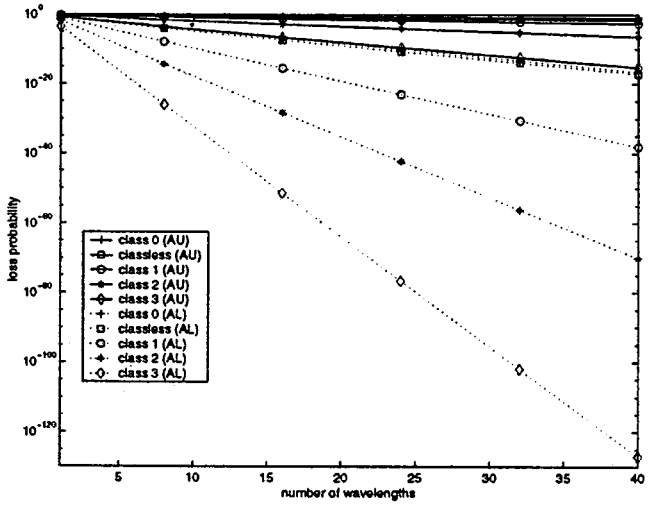

Fig. 13. Loss probability versus number of wavelengths.

Fig. 13 shows the loss probability as a function of the number of wavelengths k when $n = 4$, $N = 4$ (or $d = 4 \cdot k$), and $\rho = 0.8$. The solid lines show the upper bounds obtained from (2) and (6), while the dotted lines show the lower bounds obtained from (3) and (7). As k increases, the improvement ratio of the loss probability of class i, where $i > 0$, over the average becomes significant. For example, with $k = 40$, classes 1, 2, and 3 can achieve more than 10, 20, and 100 orders of magnitude reduction in the loss probability, respectively, while resulting in a negligible increase in the loss probability of class 0. This can be inferred from (7) since when pb_l^i (where $i > 1$) (and the average loss probability PB) is small, $pb_l^0 \approx PB$. For a given k, the difference between the lower and upper bounds on the loss probability of class i, which can be as large as 100 orders of magnitude for $i = 3$, indicates the potential improvement that is solely due to the use of four physical FDLs. In addition, as can be seen by comparing Fig. 13 to Fig. 12, increasing the number of wavelengths is a more effective way than increasing the number of FDLs to improve QoS performance because of the wavelength shared nature of the FDL buffer as shown in Fig. 2.

VII. Concluding Remarks

In this paper, we have described the offset-time-based QoS scheme when applied to an optical burst-switched WDM network with limited fiber delay lines (FDLs). Unlike the buffer-based schemes, the offset-time-based scheme does not mandate any buffer to achieve class isolation, but can nevertheless take advantage of FDLs to further improve the QoS performance. The offset times required for class isolation in wavelength reservation and FDL reservation have been quantified. For both classless OBS and prioritized OBS, theoretical upper and lower bounds on the loss probability have been obtained based on $M/M/k/k$ and $M/M/k/D$, respectively. In addition, under both Poisson and self-similar traffic, QoS performance has been measured in terms of the burst loss probability and quereing delay as a function of the maximum delay time of a FDL, the offset time difference, the number of FDLs, the number of classes, and the number of wavelengths. We have shown that the offset-time-based QoS scheme can be effective in achieving a good service differentiation and in particular, for high priority classes, both the quereing delay due to blocking and the pretransmission delay due to the use of the offset time can be kept low while maintaining a low loss probability. In summary, the offset-time-based QoS scheme can efficiently support basic QoS in the next generation Optical Internet due to its simplicity, scalability, and effectiveness with the number of wavelengths.

References

[1] A. Gnauck *et al.*, "One terabit/s transmission experiment," in *Proc. IEE OFC'96*, vol. 2, pp. 407–410.
[2] R. Wagner *et al.*, "MONET: Multiwavelength optical networking," *J. Lightwave Technol.*, vol. 14, pp. 1349–1355, 1996.
[3] E. Karasan and E. Ayanoglu, "Performance of WDM transport networks," *IEEE J. Selected Areas in Commun.*, vol. 16, pp. 1081–1096, Sept. 1998.
[4] B. Arnaud, "Architectural and engineering issues for building an optical internet," in *SPIE Proc., All Opt. Commun. Syst.: Architecture, Control Network Issues*, vol. 3531, Nov. 1998, pp. 358–377.
[5] J. Manchester, J. Anderson, B. Doshi, and S. Dravida, "IP over SONET," *IEEE Commun. Mag.*, vol. 36, pp. 136–142, May 1998.
[6] R. Braden *et al.*, "Integrated services in the internet architecture: An overview," *RFC 1633*, June 1994.
[7] S. Blake *et al.*, "An architecture for differentated services," *RFC 2475*, Dec. 1998.
[8] H. Zhang, "Service disciplines for guaranteed performance service in packet-switching networks," *Proc. IEEE*, vol. 83, pp. 1374–1396, Oct. 1995.
[9] A. Varma and D. Stiliadis, "Hardware implementation of fair queueing algorithms for asynchronous transfer mode networks," *IEEE Commun. Mag.*, vol. 35, pp. 74–80, Dec. 1997.
[10] M. Yoo and C. Qiao, "Just-enough-time(JET): A high speed protocol for bursty traffic in optical networks," *IEEE/LEOS Technol. Global Information Infrastructure*, pp. 26–27, Aug. 1997.
[11] M. Yoo, M. Jeong, and C. Qiao, "A high speed protocol for bursty traffic in optical networks," in *SPIE Proc. All Optical Commun. Syst.: Architecture, Control Network Issues*, vol. 3230, Nov. 1997, pp. 79–90.
[12] C. Qiao and M. Yoo, "Optical burst switching (OBS)—A new paradigm for an optical internet," *J. High Speed Network*, vol. 8, pp. 69–84, 1999.
[13] M. Yoo and C. Qiao, "A new optical burst switching (OBS) protocol for supporting quality of service," in *SPIE Proc., All Optical Commun. Syst.: Architecture, Control Network Issues*, vol. 3531, Nov. 1998, pp. 396–405.
[14] ——, "Supporting multiple classes of services in IP over WDM networks," in *Proc. IEEE Globecom*, 1999, pp. 1023–1027.
[15] F. Masatti *et al.*, "Fiber-delay lines optical buffer for ATM photonic switching applications," in *Proc. IEEE Infocom*, 1993, pp. 935–942.

[16] I. Chlamtac *et al.*, "CORD: Contention resolution by delay lines," *IEEE J. Select Areas Commun.*, vol. 14, pp. 1014–1029, June 1996.
[17] D. K. Hunter, W. D. Cornwell, T. H. Gilfedder, A. Franzen, and I. Andonovic, "SLOB: A switch with large optical buffers for packet switching," *J. Lightwave Technol.*, vol. 16, pp. 1725–1736, Oct. 1998.
[18] L. Kleinrock, *Queueing Systems, Volume 1: Theory*. New York: Wiley-Interscience, 1975.
[19] D. Gross and C. Harris, *Fundamentals of Queueing Theory*. New York: Wiley-Interscience, 1997.
[20] V. Paxon and S. Floyd, "Wide area traffic: The failure of poisson modeling," *IEEE/ACM Trans. Networking*, vol. 3, pp. 226–244, 1995.

Myungsik Yoo (M'00) received the B.S. and M.S degrees in electrical engineering from Korea University, Seoul, Korea, in 1989 and 1991, respectively, and the Ph.D. degree in electrical engineering from State University of New York at Buffalo (SUNY at Buffalo), Buffalo, NY, in 2000.

He was a Senior Research Engineer for Nokia Research Center, Boston, MA. He is currently an Assistant Professor in the School of Electronic Engineering, Soongsil University, Seoul, Korea. His current research interests include optical burst switching for WDM networks, supporting quality of service in the next generation Optical Internet, MPLS (lambda switching), and control and management issues in the optical networks.

Chunming Qiao (S'89–M'93) received the B.S. degree from University of Science and Technology of China (USTC) in 1985. He received the Andrew-Mellon Distinguished doctoral fellowship award from University of Pittsburgh, PA, and received the Ph.D. degree in 1993.

He then joined the University at Buffalo (SUNY) in 1993. His research interests cover the two converging areas of computers and communications, including parallel and distributed computing and systems, interconnection networks and photonic switching. He has published many papers in these areas in leading journals including IEEE/ACM TRANSACTIONS ON NETWORKING, IEEE TRANSACTIONS ON COMMUNICATIONS, JOURNAL OF LIGHTWAVE TECHNOLOGY, IEEE TRANSACTIONS ON COMPUTERS, and IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS. He has pioneered research on the next generation Optical Internet and in particular, optical burst switching (OBS). His research has been sponsored by U.S. National Science Foundation, Alcatel Corporate Research Center, Nokia Research Center, and Telcordia (formerly Bellcore). His research has been featured in *OE Reports*, several invited talks, and a keynote speech. He is also an Editor of the *Journal on High-Speed Networks (JHSN)* and the *Optical Networks Magazine*.

Dr. Qiao is a member of IEEE Computer Society and IEEE Communications Society. He has served as a Co-Chair for the annual All-optical Networking Conference sponsored by SPIE (International Society for Optical Engineering) since 1997, a Program Vice Co-Chair for the 1998 International Conference on Computer Communications and Networks (IC3N), technical program committee members, and as a session organizer/chair in several other conferences and workshops.

Sudhir Dixit (SM'75–M'80–SM'94) received the B.E. degree from Maulana Azad College of Technology (MACT), Bhopal, India, the M.E. degree from Birla Institute of Technology and Science (BITS), Pilani, India, and the Ph.D. degree from the University of Strathclyde, Glasgow, Scotland, all in electrical engineering. He also received the M.B.A. degree from Florida Institute of Technology, Melbourne.

He currently heads the research in broadband networks at Nokia Research Center Burlington, MA, specializing in ATM, Internet, optical networks, and 3rd generation mobile networks. From 1991 to 1996, he was a broadband network architect at NYNEX Science and Technology (now Bell Atlantic). Prior to that he held various engineering and management positions at other major companies, e.g., GTE, Motorola, Wang, Harris, and STL (now Nortel Europe). He has published extensively, given invited talks/panels, and has patents either granted or pending. He has served in various capacities in several conferences, and has also been an ATM Forum Ambassador.

Dr. Dixit has served as a Guest Editor in IEEE NETWORK and IEEE COMMUNICATIONS MAGAZINE. Currently, he is an editor and Lightwave Series editor of the IEEE COMMUNICATIONS MAGAZINE. He is listed in several *Who's Who* publications.