Optical Burst Switching for Service Differentiation in the Next-Generation Optical Internet

Myungsik Yoo, Soongsil University
Chunming Qiao, State University of New York at Buffalo
Sudhir Dixit, Nokia Research Center

ABSTRACT

In an effort to eliminate the electronic bottleneck, new optical switches/routers (hardware) are being built for the next-generation optical Internet where IP runs over an all-optical WDM layer. However, important issues yet to be addressed in terms of protocols (software) are how to develop a new paradigm that does not require any buffer at the WDM layer, as in circuit switching, and elimination of any layers between which exist mainly due to historical reasons. At the same time, such a paradigm should also efficiently support bursty traffic with high resource utilization as in packet switching. This article surveys design issues related to a new switching paradigm called optical burst switching, which achieves a balance between circuit and packet switching while avoiding their shortcomings. We describe how OBS can be applied to the next-generation optical Internet, and in particular how offset times and delayed reservation can help avoid the use of buffer, and support quality of service at the WDM layer.

INTRODUCTION

The explosive growth of Internet traffic provides strong incentives to exploit the huge bandwidth of fiber optic networks, which has increased dramatically due to advances in dense wavelength-division multiplexing (DWDM) technology. On the other hand, Moore's Law has dictated a relatively slower increase in the electronic processing speed, which means that data transmitted optically has to be slowed down at each node if it is to be switched electronically. Accordingly, it is only natural to find ways to build the next-generation information infrastructure, which can transport IP packets directly over the optical layer without any opto-electro-optic (O/E/O) conversion. Recently, micro-electro-mechanical systems (MEMS) and other optical switching technologies have received much attention due to their ability to eliminate O/E/O conversions, and thereby the electronic bottleneck at switching nodes. However, in order to build an all-optical network where data is kept in the optical domain at all intermediate nodes, novel protocols (software) running on top of optical switches/routers also need to be designed. One of the challenging issues is how to support fast resource provisioning, asynchronous transmission (of variable size packets, e.g., IP packets), and a high degree of statistical resource sharing for efficient handling of bursty traffic, all without requiring buffers at the WDM layer since there is no optical form of random access memory (RAM) available today. Therefore, any all-optical transport methods must avoid optical buffering as much as possible.

Another challenging issue is how to support quality of service (QoS) in the next-generation optical Internet. Given that current IP provides only *best-effort* service, but at the same time some real-time applications (e.g., Internet telephony and videoconferencing) require higher QoS than non-real-time applications (e.g., e-mail and general Web browsing), it becomes apparent that for the optical Internet to be truly ubiquitous, one must address, among other important issues, how the WDM layer can provide basic QoS support (e.g., a few priority levels). Such a WDM layer capable of basic QoS support is not only necessary for carrying vital network control and management signals related to, for example, protection/restoration at the WDM layer, but will also facilitate as well as complement a QoS-enhanced version of IP for other QoS-sensitive applications. Even though a considerable amount of effort has been and is still devoted to developing QoS schemes for asynchronous transfer mode (ATM) and IP, no one has taken into account the unique properties of the WDM layer. Specifically, existing QoS schemes (or scheduling algorithms) [1, 2], which are based on packet switching, man-

Research supported in part by National Science Foundation under ANIR-9801778.

0163-6804/01/$10.00

date buffering to achieve service differentiation. As mentioned earlier, optical RAM is not currently available, which makes it impossible to apply the existing QoS schemes (which rely on buffers) to the all-optical WDM layer.

In this article we focus on the next-generation optical Internet where IP runs over an all-optical WDM layer, which consists of WDM switches and fiber links, as shown in Fig. 1. Having an all-optical WDM layer will enable a huge amount of “through'' traffic to be switched in the optical domain, and as a result, can not only eliminate electronic bottlenecks at the IP layer, but also create high-speed/bandwidth communication pipes that are transparent to bit rate and coding format. It is expected that the continuing advances in WDM networking technology will make such a future-proof WDM layer possible.

This article surveys recent work on a new switching paradigm called *optical burst switching* (OBS). OBS is a novel approach of aggregating multiple IP packets into a burst that is then transmitted without the need for any type of buffering at intermediate nodes. It should be noted that OBS is one method of transporting IP directly over a bufferless optical WDM network, and there may exist other methods to accomplish the same objective. For the sake of simplicity we will assume that there is no adaptation layer between the IP and WDM layers for framing, QoS mapping, and so on (but guard bands may be used to delimit each assembled burst).

We first discuss the key motivations behind OBS, and describe two unique features of an OBS protocol, called Just-Enough-Time (or JET), namely the use of a “base” offset time and delayed reservation. Then we describe an offset-time-based QoS scheme for OBS, and illustrate how an “extra'' offset time can isolate service classes when making resource reservation. This is followed by a discussion of the performance issues such as service differentiated blocking probability and queuing delay. Finally, conclusions are presented.

OPTICAL BURST SWITCHING

There are three switching techniques for optical networks: wavelength routing, optical packet/cell switching, and OBS. Here, we describe their features from an optical layer perspective.

With wavelength routing, lightpaths are set up between sources (ingress nodes) and destinations (egress nodes) via nodes equipped with optical WDM cross-connects (or wavelength routers). At each such wavelength router, the output wavelength (at an output port) to which an incoming signal is routed at any given time is determined solely based on the input wavelength (and input port) carrying the signal. Accordingly, wavelength routing is a form of circuit switching. In fact, under distributed signaling, two-way reservation is needed to set up lightpaths, whereby a source node sends out a control packet to make a reservation, then waits for an acknowledgment to come back before transmitting data.

One of the advantages of wavelength routing is that no optical buffer (or O/E/O conversion of data) is needed at intermediate nodes. In addition, it is a favorable choice for now since optical switches (cross-connects) based on opto-mechanical, acousto-optic, or thermo-optic technologies are currently too slow for efficient packet switching. However, one of the limitations of wavelength routing is that a lightpath takes up an entire wavelength on each link along the path, resulting in low bandwidth utilization when carrying bursty traffic streams (e.g., IP packets), since the traffic from different ingress nodes or to different egress nodes cannot statistically share the bandwidth of the lightpath. In addition, since there are not enough wavelengths in the fiber to enable full mesh connectivity, load distribution in the network may be uneven given that the traffic intensity over lightpaths varies over time. Note that the above discussion assumes that the optical nodes support only cross-connect functionality, and have no fast packet switching/routing capability. Since setting up and tearing down a lightpath would take at least several tens of milliseconds, which is comparable to (or even exceeds) the transmission time of a few megabytes of data at a high transmission rate (e.g., 2.5 Gb/s or OC-48), dynamically setting up and tearing down lightpaths on the timescale of a burst (or at the packet level) is not an efficient approach either.

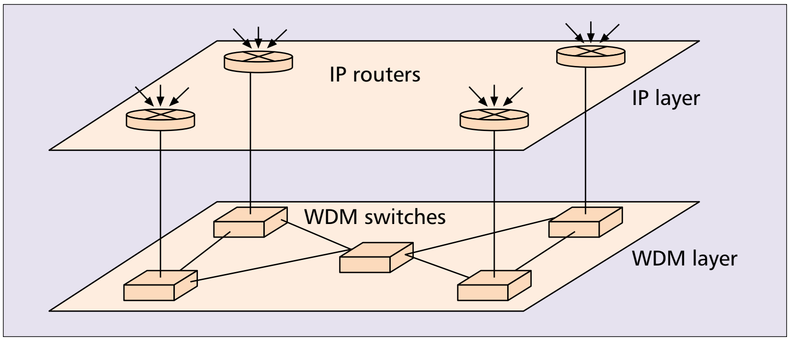

Figure 1. *IP over WDM for the next-generation optical Internet.*

When fast switches based on such technology as lithium niobate and semiconductor optical amplifiers(SOAs) become available, optical packet/cell switching will become a favorable choice because it results in better bandwidth utilization, and enables far superior traffic engineering. With optical packet/cell switching, the payload (data) is sent along with its header (control) without setting up a path. However, due to the tight coupling in time between the payload and header as well as the store-and-forward nature of packet switching, each packet needs to be buffered at every intermediate node. Because of the variations in the processing time of the packet header at the intermediate nodes, optical (or photonic) packet/cell switching also requires stringent synchronization and the complicated control that goes with it. Another problem with packet/cell switching is that the size of the payload (especially those ATM cells) is usually too small given the high channel bandwidth of optical networks, thus normally resulting in relatively high control overhead.

As discussed in [3], OBS combines the best of circuit and packet switching while avoiding their shortcomings. OBS is based on one-way reserva-

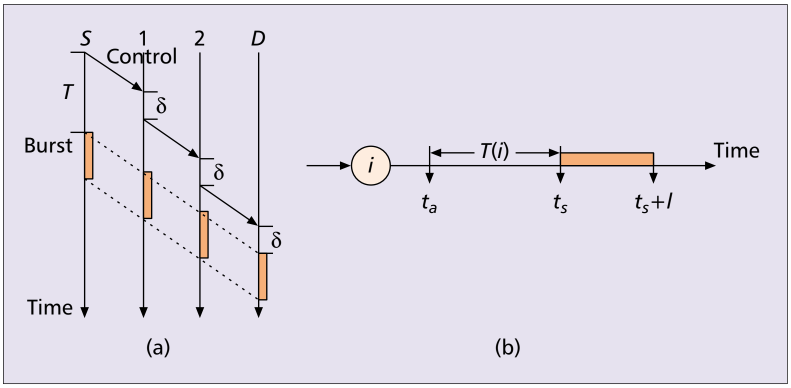

■ **Figure 2.** *The use of offset time and delayed reservation in JET-based OBS.*

tion protocols, in which a data burst follows a corresponding control (or setup) packet without waiting for an acknowledgment, resulting in much lower end-to-end latency than in circuit switching. When applied to the next-generation optical Internet, the control packet will be processed at each and every intermediate IP entity to establish an all-optical path (by configuring the WDM switches), but the corresponding burst (e.g., several IP packets) will go through only the preconfigured WDM switches at the intermediate nodes along that path.

OBS also facilitates IP-over-WDM integration under the multiprotocol label switching (MPLS) framework [4]. Since MPLS can decouple data forwarding from routing and support traffic engineering, it will play an important role in the next-generation optical Internet. Such IP-over-WDM integration, where MPLS is used to control an OBS WDM layer, is better than having a separate optical WDM layer (and a separate IP layer), each with its own mechanisms for addressing, routing, resource provisioning, and so on. More specifically, in the integrated network, control (e.g., routing and traffic engineering) can be performed electronically (via MPLS), but data can be switched optically. With OBS, such an extension can be made by letting each control packet carry the label information (in addition to other control information), and treating it as a "jumbo" label shim. The resulting approach, called *label OBS* (LOBS) [5], is similar to multiprotocol lambda switching (MPλS) [6], which is proposed for wavelength-routed networks. Since in the latter approach a wavelength is used as a label (or, in other words, each label switched path, LSP, is a lightpath), many MPLS defined label operations such as label stacking (or LSP merging) are difficult to perform. In LOBS, on the other hand, label stacking is easy since control packets containing label (and other) information are processed by a controller running on top of each WDM (OBS) switch (note that such controllers are not shown in Fig. 1), which works just as a labeled switching router (LSR). For a more detailed description of LOBS, readers are referred to [5]. It must be noted that MPλS and LOBS may be implemented simultaneously (but independently), although in MPλS the timescale of lightpaths will be significantly longer than in LOBS.

AN OBS PROTOCOL USING OFFSET TIME AND DELAYED RESERVATION

An OBS protocol called Just-Enough-Time (JET) works as shown in Fig. 2. A source sends out a control packet, which is followed by a burst after a "base" offset time, $T \geq \Sigma_{h=1}^{H} \delta(h)$, where $\delta(h)$ is the (expected) control delay (e.g., the processing time incurred by the control packet) at hop $1 \leq h \leq H$ (Fig. 2a, where $H = 3$ and $\delta(h) = \delta$). Because the burst is buffered at the source (in the electronic domain), no fiber delay lines (FDLs) are necessary at each intermediate node to delay the burst while the control packet is being processed. In addition, JET-based OBS can achieve efficient bandwidth utilization by using *delayed reservation* (DR) (Fig. 2b), whereby the bandwidth on the output link at node i (e.g., $i = 1, 2$) is reserved from the burst arrival time, denoted t_s, instead of from the time at which the processing of the control packet finishes, denoted by t_a (since the offset time remaining after i hops is $T(i) = T - \Sigma_{h=1}^{i} \delta(h)$, we have $t_s = t_a + T(i)$). In addition, the bandwidth will be reserved until the burst departure time $t_s + l$, where l is the burst length. If the requested bandwidth is not available, the burst is said to be blocked and will be dropped if it cannot be buffered. A dropped burst may then be retransmitted later. Note that JET-based OBS can also take advantage of any FDLs available at an intermediate node in resolving contention among multiple bursts by using the FDLs to delay a blocked burst until bandwidth becomes available, thus achieving better performance (e.g., lower burst dropping probability) than packet/cell switching where FDLs are not used for the purpose of resolving contention (at least not 100 percent of them).

For a description of other burst switching techniques and work related to OBS, interested readers are referred to [7, references therein].

QOS SUPPORT IN ALL-OPTICAL NETWORKS

As discussed earlier, an integrated IP-over-WDM solution calls for basic QoS support at the WDM layer. Although QoS concerns related to optical (or analog) transmission (e.g., dispersion, power, and signal-to-noise ratio) also need to be addressed, here we focus on how to ensure that critical data can be transported at the WDM layer more reliably than noncritical data using a QoS scheme which extends JET-based OBS. Such a QoS scheme will be referred to as the offset-time-based QoS scheme for reasons that will become apparent later.

The offset-time-based QoS scheme to be described is more scalable than per-connection (flow) QoS schemes because the former only needs to manage a few levels of priority. More specifically, burst assembly is performed at the network edge, where multiple IP packets belonging to several flows are assembled into a burst (super packet) based on their destination address and QoS parameter (e.g., class or priority). In this way, only simple and scalable QoS management is done at the WDM layer, while more complex traffic shaping and QoS scheduling, which requires a large amount of buffering (among other resources), are done at the IP layer.

Note that the offset-time-based QoS scheme also facilitates traffic engineering in an integrated IP-over-WDM network. More specifically, one of the main objectives of traffic engineering is to minimize congestion. Typically, there are two types of congestion: that caused when network resources are insufficient to accommodate the offered load due to an instantaneous high traffic load, and that caused when network resources are inefficiently utilized due to unbalanced traffic distribution. The offset-time-based QoS scheme can effectively deal with the first type of congestion as follows. When the traffic is a mixture of priority classes 0 through $(n - 1)$, for some integer $n > 1$, the offset-time-based scheme can isolate classes (discussed next), and accordingly the high-priority classes are likely to get the required bandwidth (wavelength) regardless of the offered load from the lower-priority classes. In other words, bandwidth at a given time is distributed in an hierarchical order, from the highest priority class (class $n - 1$) to the lowest priority class (class 0). Thus, even when congestion may build up temporarily, the highest priority class is not likely to be affected by congestion. The second type of congestion can usually be dealt with by constraint-based routing algorithms, which are responsible for distributing traffic evenly among the network. The offset-time-based QoS scheme is independent of these algorithms, and therefore can work well with any appropriate routing algorithms.

AN OFFSET-TIME-BASED QOS SCHEME

We now explain how the offset-time-based QoS scheme works. In particular, we explain how class isolation (or service differentiation) can be achieved by using an "extra" offset time in both cases with and without using fiber delay lines (FDLs) at a switch.

Note that one may distinguish the following two different contentions in reserving resources (wavelengths and FDLs): intraclass contentions caused by requests belonging to the same class, and interclass contentions caused by requests belonging to different classes. In what follows, we will focus on how to resolve interclass contentions using the offset-time-based QoS scheme.

For simplicity, we assume that there are only two classes of (OBS) services, class 0 and class 1, where class 1 has priority over class 0. In the offset-time-based QoS scheme, to give class 1 a higher priority for resource reservation, an extra offset time, denoted t_o^1, is given to class 1 traffic (but not to class 0, i.e., $t_o^0 = 0$). In addition, we also 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. Finally, without loss of generality, we also assume that a link has only one wavelength for data (and an additional wavelength for control).

THE CASE WITHOUT FDLS

In the following discussion, let t_a^i and t_s^i be the arriving time and service start time for a class i request denoted $req(i)$, respectively, and let l_i be the burst length requested by $req(i)$, where $i = 0$, 1. Fig. 3 illustrates why a class 1 request that is assigned an (extra) offset time obtains a higher priority for wavelength reservation than a class 0 request. We assume that there is no burst (arriving earlier) in service when the first request arrives. Consider the following two situations where contentions among two classes of traffic are possible.

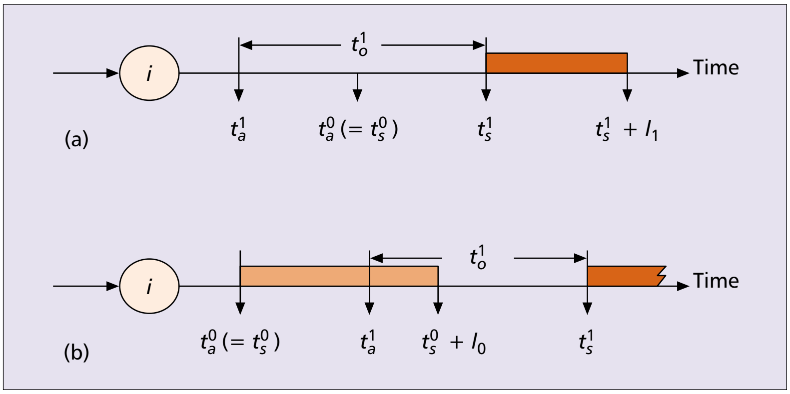

Figure 3. *Class isolation at an optical switch without FDLs.*

In the first case, as illustrated in Fig. 3a, $req(1)$ comes first and reserves wavelength using DR, and $req(0)$ comes after. Clearly, $req(1)$ will succeed, but $req(0)$ will be blocked if $t_a^0 < t_s^1$ but $t_a^0 + l_0 > t_s^1$, or $t_s^1 < t_a^0 < t_s^1 + l_1$.

In the second case, $req(0)$ arrives first, followed by $req(1)$, as shown in Fig. 3b. When $t_a^1 < t_a^0 + l_0$, $req(1)$would be blocked had no offset time been assigned to $req(1)$ (i.e., $t_o^1 = 0$). However, such a blocking can be avoided by using a large enough offset time so that $t_s^1 = t_a^1 + t_o^1 > t_a^0 + l_0$. Given that t_a^1 may only be slightly behind t_a^0, t_o^1 needs to be larger than the maximum burst length over all bursts in class 0 in order for $req(1)$ to completely avoid being blocked by $req(0)$. With that much offset time, the blocking probability of (the bursts in) class 1 becomes only a function of the offered load belonging to class 1, that is, independent of the offered load belonging to class 0. On the other hand, the blocking probability of class 0 will be determined by the offered load belonging to both classes.

THE CASE WITH FDLS

Although the offset-time-based QoS scheme does not mandate the use of FDLs, its QoS performance can be significantly improved even with limited FDLs to resolve contention for bandwidth among multiple bursts. For the case with FDLs, the variable B will denote the maximum delay an FDL (or the longest FDL) can provide. Thus, in the case of blocking, a burst can be delayed up to the maximum delay B.

Figure 4a and b illustrate class isolation at a switch equipped with FDLs where contention for both wavelength and FDL reservation may occur. In Fig. 4a, let us assume that when $req(0)$ arrives at t_a^0 $(= t_s^0)$, the wavelength is in use by a burst that arrived earlier. Thus, the burst corresponding to $req(0)$ has to be delayed (blocked) for t_b^0 units. Note that the value of t_b ranges from 0 to B, and an FDL with an appropriate length that can provide a delay of t_b will be chosen. Accord-

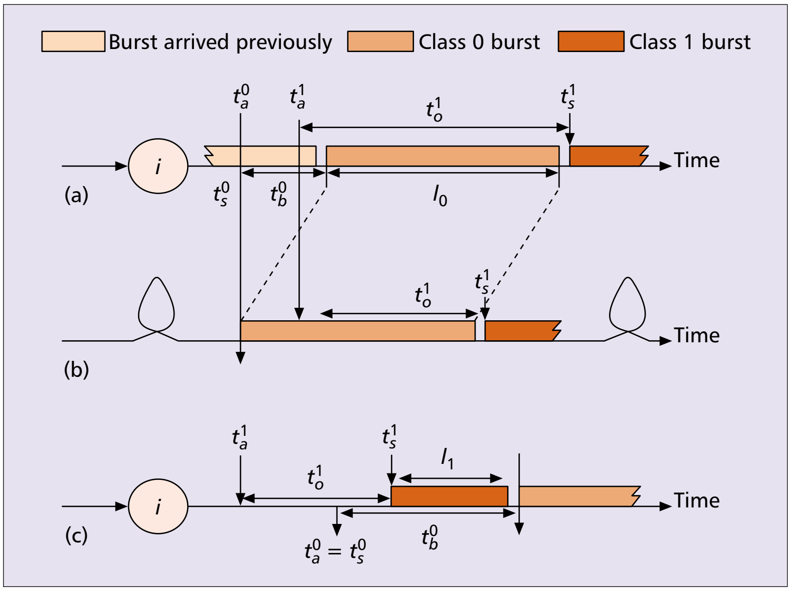

Figure 4. *Class isolation at an optical switch with FDLs.*

ingly, if $t_b^0 < B$, the FDL is reserved for a class 0 burst, as shown in Fig. 4b (the burst will be dropped if t_b^0 exceeds B), 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. 4a. Now assume that *req*(1) arrives later at t_a^1 (where $t_a^1 > t_a^0$) and tries to reserve the wavelength. *req*(1) will succeed in reserving the wavelength as long as the offset time, t_o^1, is so long 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) in Fig. 4a, it is obvious that *req*(1) would not have interclass contention caused by *req*(0). This illustrates that class 1 can be isolated from class 0 when reserving wavelength because of the offset time. Of course, without the offset time ($t_o^1 = 0$, and thus $t_a^1 = t_s^1$), *req*(1) would be blocked for $t_a^0 + t_b^0 + l_0 - t_a^1$, and it would depend entirely on the use of FDLs to resolve this interclass contention.

Similarly, Fig. 4b illustrates class isolation in FDL reservation. More specifically, let us assume that *req*(0) has reserved the FDLs as described earlier, and because t_o^1 is not long enough, *req*(1) is blocked in wavelength reservation and thus needs to reserve the FDLs. In such a case, *req*(1) will successfully reserve the FDLs if the offset time is still long enough to have $t_s^1 = t_a^1 + t_o^1 > l_0$. Otherwise (i.e., if $t_s^1 < t_a^0 + l_0$), *req*(1) would contend with *req*(0) to reserve the FDL and would be dropped.

As shown in Fig. 4c, if *req*(1) comes first and reserves the wavelength based on t_o^1 and delayed reservation (DR), and *req*(0) comes afterward, *req*(1) is not affected by *req*(0) . However, *req*(0) will be blocked when either $t_a^1 < t_a^0 < t_s^1$ but $t_a^0 + l_0 > t_s^1$, or $t_s^1 < t_a^0 < t_s^1 + l_1$. Similarly, if *req*(1) arrives first, it can reserve the FDL first regardless of whether *req*(0) succeeds in reserving the FDL. As mentioned earlier, this implies that class 1 can be isolated from class 0 in reserving both wavelength and FDL by using an appropriate offset time, which explicitly gives class 1 higher priority than class 0. As a result of having low priority for resource reservation, class 0 bursts will have relatively high blocking and loss probabilities.

t_{diff}	$0.4 \cdot L$	L	$3 \cdot L$	$5 \cdot L$
R	0.3296	0.6321	0.9502	0.9932

Table 1. *Offset time difference and class isolation degree.*

QUANTIFYING THE (EXTRA) OFFSET TIME

It is important to note that in the case of two classes, 0 and 1, the offset time assigned to class 1 needs not be longer than the maximum duration of a burst belonging to class 0 so as to achieve 100 percent class isolation. In other words, one can achieve near 100 percent class isolation, which will result in good service differentiation, with a reasonable offset time only a few times the average duration of a burst in class 0.

The above discussion is also applicable to multiple (e.g., $n > 2$) classes as well. For example, let us consider two adjacent classes, class i and class $(i - 1)$. Define two terms, offset time difference and degree of isolation, where the former, denoted t_{diff}, is the difference in the offset times assigned to class i and class $(i - 1)$ (i.e., $t_{\text{diff}} = t_o^i - t_o^{i-1}$), and the latter, denoted R, is the lower bound on the isolation degree, or, in other words, the probability that *req*(i) will not be blocked by req($i - 1$) (arrived earlier). Table 1 shows the values of t_{diff} to achieve a certain degree of isolation for wavelength reservation at a switch without FDLs. From the table, one can see that if $t_{\text{diff}} = 3L$, where L is the average duration of a burst in class $(i - 1)$ (where the burst duration is assumed to have an exponential distribution), more than 95 percent of class $(i - 1)$ traffic is isolated from class i traffic.

In addition, if class i is isolated (with degree R) from class $(i - 1)$, which in turn is isolated from class $(i - 2)$, class i is also isolated from class $(i - 2)$ (with a degree higher than R). Interested readers may refer to [8, 9] for derivation of the (extra) offset time(s) needed to achieve a given degree of class isolation when multiple classes must be supported, and/or when FDLs are used.

PERFORMANCE ISSUES

Two OBS approaches have been evaluated: *classless* and *prioritized*. In classless OBS, no (extra) offset time is assigned to any class (i.e., no priority, as in current IPs best-effort service), whereas in prioritized OBS, a longer offset time will be assigned to a higher priority class.

To show the effect of class isolation achieved by the offset-time-based QoS scheme, performance results on the blocking probabilities and queuing delays experienced by bursts in four service classes, which contend for the same output link at a switch, will be presented. These results have been obtained by assuming that there are k wavelengths for bursts (i.e., data) on the output link, and the switch is capable of wavelength conversion. In addition, each of these four classes generates an equal amount of traffic (load), and in particular, bursts in each class have the same average duration of L. In prioritized OBS, the offset time difference between any two adjacent classes is the same (and will be denoted by t_{diff} as before). Interested readers may refer to

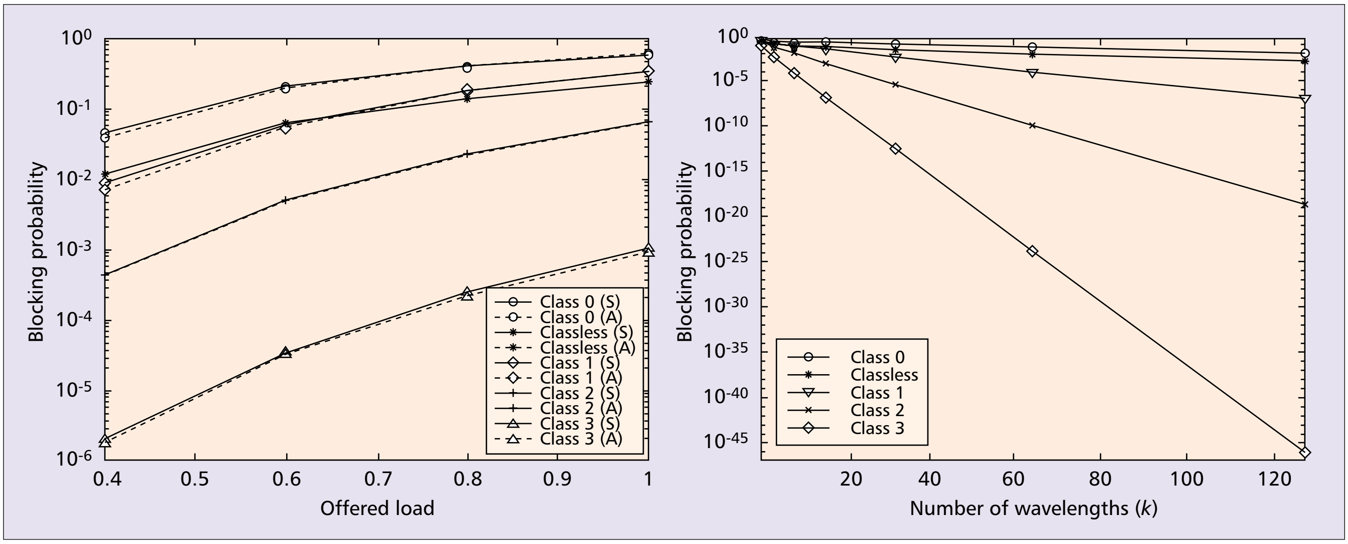

Figure 5. *a) Differentiated blocking probabilities; b) the effect of number of wavelengths on service differentiation.*

[8, 9] for descriptions of the analytic as well as simulation models used to obtain the following results, as well as additional results not presented in this article.

BLOCKING PROBABILITY

Figure 5a shows the blocking probability as a function of the offered load when the number of wavelengths (k) is 8. The dotted and solid lines show the results from analysis assuming 100 percent isolation and those from simulation assuming the offset time difference (t_{diff}) is $3L$ or equivalent, about 95 percent isolation, respectively. As can be seen, analysis and simulation results are closely matched, meaning that $t_{\text{diff}} = 3L$ is sufficient for all practical purposes. The service differentiation can be observed in that a class with higher priority experiences lower blocking probability. In addition, although not explicitly shown, these (and other) results indicate that the average blocking probability in prioritized OBS (over all four classes) is the same as the overall blocking probability in classless OBS. In other words, the offset-time-based QoS scheme effectively achieves service differentiation while maintaining the same overall blocking probability.

Figure 5b shows the effect of number of wavelengths on service differentiation when the offered load is 0.8 and t_{diff} is $3L$. The blocking probability in classless OBS is 10^{-2} (10^{-3}) even when the number of wavelengths is 64 (128). However, in prioritized OBS, the blocking probabilities of classes 1–3 are much lower than those in classless OBS. As an extreme, the blocking probability of class 3 is as low as 10^{-25} (10^{-48}) when the number of wavelengths is 64 (128), which is more than 20 (40) orders of magnitude reduced from that in classless OBS. As the number of wavelengths increases, the blocking probability of higher-priority classes decreases rapidly, which shows the effectiveness of the offset-time-based QoS scheme. Note that achieving such service differentiation results in only a small increase in the blocking probability of class 0 in prioritized OBS from that in classless OBS.

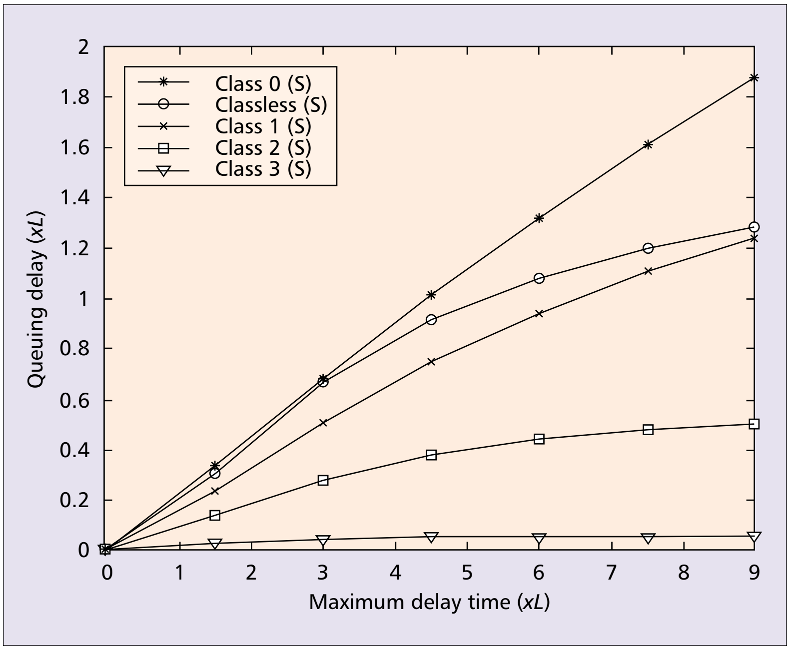

Figure 6. *Differentiated queuing delays.*

QUEUING DELAY AND END-TO-END LATENCY

The queuing delay experienced by each class (averaged over all bursts in that class) is plotted in Fig. 6 as a function of the maximum delay time (B) that can be provided via FDLs at each switch. The simulation results have been obtained by setting the offset time difference to $3L + B$, which provides about 95 percent (and higher) class isolation in bandwidth (and FDL) reservation. Since the queuing delay is proportional to the blocking probability, it is not surprising to see, as shown in the figure, that the queuing delays of different classes in prioritized OBS are differentiated, and a higher-priority class has a lower queuing delay. In addition, although the queuing delay for a given class increases with B, the dependency is weaker for higher-priority classes. In particular, class 3 encounters a low queuing delay (less than $0.02L$) even when B is as long as $9L$ since bursts in class 3 seldom need to use any FDLs due to their low blocking probability achieved via class isolation.

It is noted that using an offset time increases end-to-end latency.[1] Thus, for delay-sensitive applications, it is important to discuss the impact of the (extra) offset time, whose value depends on the number of classes (n) to be supported, and the offset time difference (t_{diff}) used between two adjacent classes. If we assume that all the offset time differences are equal, the longest additional delay introduced will be $(n - 1) \cdot t_{\text{diff}}$, which is equal to the extra offset time used by class ($n - 1$). Given that in today's voice and video communications interactive applications require that the end-to-end latency be bounded by a few hundreds of milliseconds, such an additional delay might be negligible. As an example, assume that the mean burst size is 15 kbytes (containing a few tens of average-size IP packets), which results in $L = 12$ µs at 10 Gb/s. When the number of service classes is 4, and offset time difference is $3L$ (which results in at least 95 percent class isolation), maximum additional delay is 108 µs. Even though increasing the average length of a burst, or the number of classes, and so on, may result in an additional delay longer than 0.1 ms, the key observation that can be drawn from Fig. 6 (and earlier discussion) is that such a small additional delay can significantly reduce or even eliminate queuing delays at intermediate nodes. In short, as long as the number of classes and, in particular, the average burst length are carefully engineered (e.g., during the burst assembly process), the negative impact of an extra offset time on the end-to-end latency can be negligible, and largely compensated for by its possible impact on minimizing queuing delays.

As long as the number of classes and, in particular, the average burst length are carefully engineered, the negative impact of an extra offset time on the end-to-end latency can be negligible, and largely compensated for by its possible impact on minimizing queuing delays.

CONCLUSION

In this article we have described the concept of optical burst switching, which has been proposed to support QoS in the next-generation optical Internet. By achieving a balance between wavelength routing and optical packet switching, OBS facilitates the integration of IP and WDM, and supports fast provisioning, asynchronous transmission (of variable-size packets), and high resource utilization without requiring buffering at the WDM layer. We have described unique features of an OBS protocol called Just-Enough-Time (i.e., the use of an offset time and delayed reservation) and, in particular, how to extend them to develop an offset-time-based QoS scheme, which can efficiently differentiate services at the optical layer without requiring any buffering (not even fiber delay lines). The class isolation using extra offset time has been illustrated with and without FDLs. It has been shown that a significant improvement in QoS performance in terms of blocking probability and queuing delay is possible by simply introducing an extra offset time whose negative impact on the overall end-to-end latency is negligible. Due to its effectiveness and scalability with a large number of wavelengths, the offset-time-based QoS scheme can not only enhance current IP best-effort service, but also supplement a QoS-enhanced version of IP in the next-generation optical Internet.

REFERENCES

[1] H. Zhang, "Service Disciplines for Guaranteed Performance Service in Packet-switching Networks," *Proc. IEEE*, vol. 83, Oct. 1995, pp. 1374–96.

[2] A. Varma and D. Stiliadis, "Hardware Implementation of Fair Queuing Algorithms for Asynchronous Transfer Mode Networks," *IEEE Commun. Mag.*, vol. 35, Dec. 1997, pp. 74–80.

[3] C. Qiao and M. Yoo, "Optical Burst Switching (OBS) — A New Paradigm for an Optical Internet," *J. High Speed Networks*, vol. 8, no. 1, 1999, pp. 69–84.

[4] D. Awduche," MPLS and Traffic Engineering in IP Networks," *IEEE Commun. Mag.*, Dec. 1999.

[5] C. Qiao, "Labeled Optical Burst Switching for IP-over-WDM Integration," *IEEE Commun. Mag.*, vol. 38, no. 9, Sept. 2000, pp. 104–14.

[6] D. Awduche *et al.*, "Multi-protocol Lambda Switching: Combining MPLS Traffic Engineering Control with Optical Crossconnects," Internet draft, Oct. 1999, work in progress.

[7] C. Qiao and M. Yoo, "Choices, Features and Issues in Optical Burst Switching," *Opt. Networks*, vol. 1, no. 2, 2000, pp. 36–44.

[8] M. Yoo and C. Qiao, "Supporting Multiple Classes of Services in IP Over WDM Networks," *Proc. IEEE GLOBECOM*, 1999, pp. 1023–27.

[9] M. Yoo, C. Qiao, and S. Dixit, "The Effect of Limited Fiber Delay Lines on QoS Performance of Optical Burst Switched WDM Networks," *ICC 2000*, pp. 974–79.

BIOGRAPHIES

MYUNGSIK YOO [M] (myoo@e.ssu.ac.kr) received his B.S. and M.S. degrees in electrical engineering from Korea University, Seoul, in 1989 and 1991, respectively, and his Ph.D. in electrical engineering from State University of New York at Buffalo, New York, in 2000. He was a senior research engineer at Nokia Research Center, Burlington, Massachusetts. He is currently an assistant professor in the school of electronic engineering, Soongsil University, Seoul, Korea. His current research interests focus on issues in the WDM optical networks and optical Internet, including optical burst switching, protection/restoration, QoS support, multiprotocol lambda switching, and control and management issues.

CHUNMING QIAO [M] is currently a tenured associate professor in the Computer and Science Engineering Department at SUNY Buffalo. He is also an adjunct professor in the Electrical Engineering Department, and the director of the Lab for Advanced Network Design, Evaluation and Research (LANDER), which conducts cutting-edge research related to optical networks, wireless/mobile networks, and the Internet. His research has been supported by a number of NSF grants including an Research Initiation Award (IRA) and Information Technology Research (ITR) award, and by Telcordia tech (formerly Bellcore), Alcatel Corporate Research Center, and Nokia Research Center. He has over 10 years of research experience in optical networks, covering the areas of photonic switching devices, WDM communications and optical packet-switching. He pioneered research on the next-generation optical Internet, in particular optical burst switching. He is the founder and chair of the recently established Technical Group on Optical Networks (TGON) sponsored by SPIE, the IEEE Communication Society's Editor-at-Large for optical networking and computing, an editor of the *Journal on High-Speed Networks* and the new *Optical Networks Magazine*, and a member of IEEE Computer Society. He received the Carnegie-Mellon Distinguished doctoral fellowship award from University of Pittsburgh, and subsequently earned his Ph.D. in computer science in 1993.

SUDHIR DIXIT [SM] (sudhir.dixit@nokia.com) received his B.E. degree from Maulana Azad College of Technology, Bhopal, India, his M.E. degree from Birla Institute of Technology and Science, Pilani, India, and his Ph.D. degree from the University of Strathclyde, Glasgow, Scotland, all in electrical engineering. He also received his M.B.A. degree from Florida Institute of Technology, Melbourne. He currently heads research in broadband networks at Nokia Research Center, Burlington, Massachusetts, specializing in ATM, Internet, optical networks, and third-generation mobile networks. From 1991 to 1996 he was a broadband network architect at NYNEX Science and Technology (now Verizon). Prior to that he held various engineering and management positions at other major companies, such as GTE, Motorola, Wang, Harris, and STL (now Nortel Europe).

---

[1] *Since the base offset time may be as small as the total processing delays encountered by a control packet, only the (extra) offset time will contribute to any increased end-to-end latency.*