PHOTONIC BUFFER ARCHITECTURE TO SUPPORT PRIORITIZED BUFFER MANAGEMENT FOR ASYNCHRONOUSLY ARRIVING VARIABLE-LENGTH PACKETS

Hiroaki Harai
Communications Research Laboratory
Koganei-shi, Tokyo 184-8795, Japan
harai@crl.go.jp

Masayuki Murata
Osaka University
Toyonaka-shi, Osaka 560-0043, Japan
murata@cmc.osaka-u.ac.jp

Abstract We study photonic packet switching architecture that enables high node throughput and provides priority services. While an ultra-high-speed address-lookup capability can be realized in the optical domain for a photonic packet switch, the optical buffer must be managed in the electrical domain. In this paper, we propose a photonic buffer architecture that can support prioritized buffer management schemes with less complexity, so that the capability for ultra-high-speed packet forwarding can be sustained in the photonic packet switch. We propose the vPBSO (Variable-length-packet-capable Partial Buffer Sharing with Overwriting) method, which is a prioritized buffer management scheme for providing DiffServ (Differentiated Services) Assured Forwarding. It can handle asynchronously arriving variable-length packets. vPBSO is based on a single queue, and its complexity is $O(p)$, where p is the number of levels of drop precedence (i.e., the number of priority classes). We show that at lower arrival rates, vPBSO provides better performance than vPBS (our extension to the existing method, called PBS, which is a prioritized buffer management method for synchronously arriving fixed-length packets) for high-priority packets, and the performance for low-priority packets is also better or at least almost the same. It is also shown that once the appropriate control parameter values for vPBSO and vPBS are determined and appropriate buffer management scheme is selected depending on the arrival rate of packets, we can achieve high-performance prioritized buffer management, which can be realized by some traffic load estimation method. vPBSO is especially effective when the arrival rate of higher priority packets is lower than that of lower priority packets, which is a likely situation in the prioritized system.

Keywords: photonic packet switch, buffer management, variable-length-packet-capable partial buffer sharing with overwriting, single queue, DiffServ

1. Introduction

To build a core network capable of handling a tremendous amount of traffic in the Internet, it is necessary to improve the node throughput (i.e., the number of packets/bits processed by a node within a unit time), as well as to increase the link capacity. Currently, we rely on electronic processing for packet forwarding at nodes such as routers. The node throughput can be improved through advancements in LSI technology known as Moore's law and via large-scale distributed/pipelined processing. While the link capacity is easily increased by bundling optical fibers, the integration and pipelined processing may limit the increase in the node throughput. As it is necessary to increase the line speed and the number of ports at nodes, the electrical limitations motivate us to introduce other technologies.

One promising way to increase the node throughput is to use MPLS (Multi-Protocol Label Switching) technology [1] over photonic technology. While MPLS requires establishment of a closed domain to utilize new, lower layer technology, it is useful to incorporate it with photonic technology in order to build a very high-speed Internet. For this purpose, GMPLS (Generalized MPLS) [2] is being developed, which can overlay a lightpath network [3]. However, there are problems of deploying GMPLS over lightpath networks. The most difficult problem is the capacity granularity: the unit of bandwidth between the edge node pairs of the GMPLS domain is the wavelength capacity. This capacity may be too large to accommodate the traffic between node pairs.

We thus introduce photonic packet switching technology under MPLS to solve the granularity problem. Photonic packet switching has an inherently finer granularity in the optical domain. The functions of photonic packet switches are roughly divided into the following five groups: address/label lookup (i.e., forwarding), switching, buffer management, buffering, and routing. To transfer very high-speed data at rates such as 160 Gbps without OEO conversion, switching and buffering must be handled in the optical domain. In transferring extremely large numbers of packets in a short period of time, however, electronic memory access is a bottleneck for packet forwarding. It is thus desired for address lookup to be operated optically. Address lookup (multi-wavelength label analysis [4, 5] and optical code label analysis [6]), switching [7, 8], and buffering [8, 9] can be handled in the optical domain. In contrast, since optical logic and optical RAM (random access memory) are still impractical, buffer management still requires electronic processing, in which the delay in the optical buffer is calculated. It is therefore important to use a less complex buffer management algorithm to avoid unexpected degradation of the photonic packet switch performance.

To develop a less complex buffer management system for photonic packet switches to be deployed in the future Internet infrastructure, it is important to support two features: variable-length packets, and the diversification of applications, which are the main topics of this paper. Since IP packets are of variable lengths, by supporting the first feature we can avoid additional fragmentation from a variable-length packet into multiple fixed-length packets, which is better than supporting only fixed-length packets. Recently, some researchers described buffer management schemes for photonic packet switches to support asynchronously arriving variable-length packets [10, 11, 12]. One of the authors of this paper also proposed a buffer management scheme with wavelength conversion of variable-length optical packets to steeply improve performance in terms of packet loss probability [13]. Callegati described a kind of PBS (Partial Buffer Sharing) [14] called WA/TD (Wavelength Allocation and Threshold Dropping) for variable-length packets [15], which is also capable of wavelength conversion. All-optical wavelength conversion at high speeds is difficult and/or costly, although the introduction of supercontinuum light source makes wavelength conversion feasible [13]. Here we only focus on buffer management of the optical fiber delay line buffer.

The diversification of applications implies the need for priority services in order to provide high-quality services rather than best-effort services. DiffServ (Differentiated Services) [16] is a typical example. Priority queueing is a practical method for DiffServ. The current LSI technology has an enough capability to support multiple queues such as those for CBQ (Class-Based Queueing) [17] and DRR (Deficit Round Robin) [18], in which only one packet at a time is dequeued from one of multiple queues, which are currently used for DiffServ. We cannot expect such sophisticated and complicated queue management schemes in the photonic packet switch because the lack of optical RAM makes it difficult to manage multiple queues. However, we can still implement the DiffServ capability based on a single queue in photonic packet switches as we present in this paper.

To provide DiffServ AF (Assured Forwarding) [19] in photonic packet switches, the authors extended the PBS method and proposed the PBSO (Partial Buffer Sharing with Overwriting) method in Refs. [20, 21]. These approaches are based on a single queue, and their complexities are $O(1)$ and $O(p)$, respectively, where p is the number of priority classes. Since p is typically 2 or 3, and thus much smaller than the number of buffer size B, these methods are more suitable for buffer management in photonic packet switches than other prioritized methods such as HOL (Head-of-the-Line priority queueing) [22] and Push Out [14], of which complexities are $O(B)$. Note that PBSO provides a smaller packet loss probability than PBS does. However, PBSO as well as PBS only supports syn-

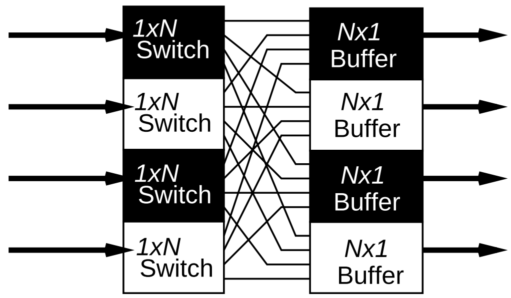

Figure 1. $N \times N$ photonic packet switch architecture ($N = 4$)

chronously arriving fixed-length packets, it is not suitable as a target for this paper.

In this paper, we propose a photonic buffer architecture that can support the vPBSO (variable-length-packet-capable Photonic PBSO) algorithm, which is an extension of PBSO. It can handle asynchronously arriving variable-length packets. vPBSO provides a different level of drop precedence for DiffServ Assured Forwarding, and its complexity is still $O(p)$. We show that at lower arrival rates, vPBSO provides better performance than vPBS (our extension of PBS for synchronously arriving fixed-length packets) for high-priority packets, and the performance for low-priority packets is also better or at least almost the same. It is also shown that once the appropriate control parameter values for vPBSO and vPBS are determined and appropriate buffer management scheme is selected depending on the arrival rate of packets, we can achieve high-performance prioritized buffer management, which can be realized by some traffic load estimation method. We also show that vPBSO is especially effective when the arrival rate of higher priority packets is lower than that of lower priority class packets, which is a likely situation in the prioritized system.

This paper is organized as follows. In Section 2, we briefly summarize the photonic packet switch architecture. In Section 3, we propose photonic buffer architecture and the vPBSO algorithm to provide different levels of drop precedence for DiffServ AF. Section 4 is devoted to our performance evaluation of vPBSO and vPBS. We then present our conclusions in Section 5.

2. Photonic Packet Switch Architecture

2.1. Overview

Figure 1 shows photonic packet switch architecture. Our buffer management is applied to the optical buffer in this architecture. The output-buffered $N \times N$ packet switch consists of the N pieces of $1 \times N$ bufferless packet switches followed by the N pieces of $N \times 1$ buffers. Every pair of a $1 \times N$ switch and an $N \times 1$ buffer is optically interconnected in a fully meshed manner. The $1 \times N$ bufferless packet switches make the address lookup function faster by providing photonic address lookup functions [4, 6] to the packet switch. They can handle packet switching of asynchronously arriving variable-length packets with precedent activity [5]. As a result, the architecture provides ultra-high node throughput to the packet switch. The $N \times 1$ buffers are used to avoid packet collision and to reduce the packet loss probability.

2.2. Fundamental Work of Buffer Management

We use optical fiber delay lines (FDLs) to compose an $N \times 1$ optical buffer. Figure 2(a) shows 4×1 optical buffer architecture. Each optical buffer consists of B fiber delay lines ($B = 4$ in this example), an $N \times (B + 1)$ optical switch, a coupler, and a buffer manager. The lengths of the B fiber delay lines, $\{d_0, d_1, \ldots, d_{B-1}\}$, are multiples of unit length D $(0, D, \ldots, (B-1)D$, respectively). The buffer provides discrete-time delays from 0 to $(B-1)D$.

Unlike the RAM buffer for an electronic node system such as an IP router, the optical straightforwardness property of the FDL buffer causes difficulties in the traditional store and forward approach. It is necessary that an appropriate FDL be selected for each arriving packet before it arrives at the FDL buffer. In $N \times N$ output-buffered photonic packet switch, up to N packets simultaneously arrive at an $N \times 1$ optical buffer. When we employ fiber delay lines, we must implement a buffer management system to calculate all delays for the N packets within time $l_{\min}$, which corresponds to the minimum packet length. The allowable processing time for one packet is thus $l_{\min}/N$, which is identical to the processing time for one packet in round-robin scheduling. We thus develop a buffer management system by using round-robin scheduling to enable simple handling of asynchronously arriving variable-length packets.

Figure 2 shows the behavior of the buffer manager when packets arrive. The buffer manager has an internal clock (with a frequency of $1/T$) that represents a unit time for round-robin scheduling. The buffer manager monitors the arrival of every packet at every input port and determines the delay time for each arriving packet up to $N = 4$ within time T. To handle arriving packets continuously, the cycle of time T must be less than $l_{\min}$ ($T \leq l_{\min}$). This constraint on the monitoring point is necessary to prevent packets from arriving at the optical buffer before the delay for each packet is determined.

As shown in Fig. 2 (a), we now assume that four packets arrive at the buffer during the same cycle. The buffer manager determines that packets A, B, and C will be switched to delay lines d_0, d_2, and d_3, respectively and packet D will be discarded due to buffer overflow (i.e., it is switched to a sink). The buffering and discarding functions are enabled by driving the optical switch. Figure 2 (b) shows the relative positions of the packets from the output port just after packet

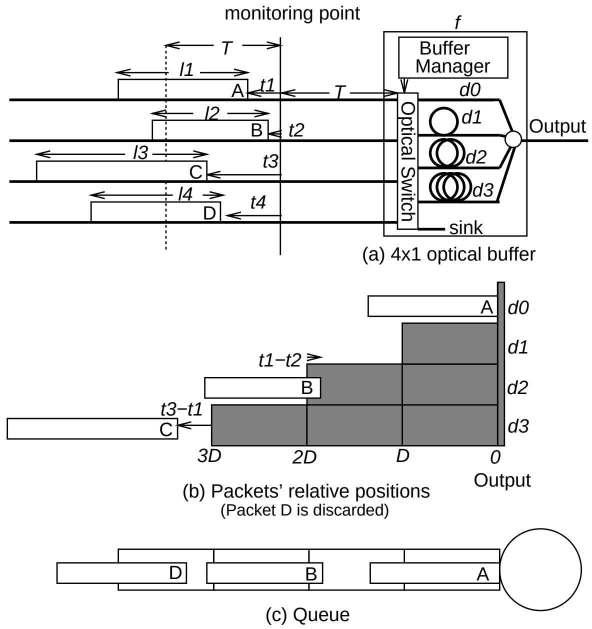

Figure 2. Fundamental behavior of packets arriving at the optical buffer

A is switched to delay line d_0. As illustrated in the figure, the three packets depart from the buffer without collision. Note that the departure order of the packets is different from the arrival order because of the round-robin scheduling. In this example, packet B arrives at the buffer before packet A, but packet A departs first. However, the arrival order of the packets at each port is identical to the departure order, except for the discarded packets. The consistency of the arrival/departure order gives better performance for higher layer applications than when there is inconsistency. This is because multiple packets in a flow/connection (e.g., TCP) usually follow the same route in the Internet.

We now describe the fundamental behavior of the buffer manager more concretely. The buffer manager maintains variable f representing the time at which all stored packets in the buffer depart and the buffer becomes idle. This time is defined relative to the starting time of the targeted cycle. In what follows, f is called the buffer occupancy. The buffer manager calculates the delays for new packets coming from all ports during each time T according to the round robin of ports $1, 2, \ldots, N$. Let l_n and t_n denote the length of a packet observed at port n and the gap (the time difference between the cycle start and the packet arrival), respectively (See Fig. 2(a)). Theoretically, $f - t_n$ is a sufficient delay given to one packet to avoid packet collision. Unfortunately, due to the discrete-time nature of the fiber delay line buffer, it is necessary for the delay

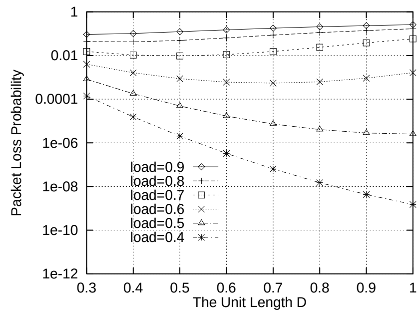

Figure 3. Performance dependence on D (unit delay normalized by mean packet length). The number of fiber delay lines $B = 50$.

given to one packet to be $\Delta_n D = \lceil \frac{f - t_n}{D} \rceil D$. The packet enters delay line d_{Δ_n} if $\Delta_n < B$, while it is discarded if $\Delta_n \geq B$. When the packet enters the delay line, the buffer occupancy f is changed to $f \leftarrow t_n + l_n + \Delta_n D$ to handle packets at the following ports appropriately. After calculating the packet delays for all ports, f is changed to $f \leftarrow \max(f - T, 0)$ to provide appropriate buffer management during the next cycle.

Due to the discrete-time nature of the fiber delay line buffer, the optical buffer has a space of $\theta = \Delta_n D - f + t_n$ between a new packet and the previous and adjacent packet. Therefore, to achieve better performance, it is important to determine the unit delay D of a fiber delay line. In Refs. [12, 13, 15], the packet loss probability is minimized by determining that D is around $0.3/\mu$ when the load is 0.8, where $1/\mu$ is the mean packet length. In addition, we observe that a larger D than $0.3/\mu$ decreases the packet loss probability when the load is smaller than 0.8 (See Fig. 3).

3. Prioritized Buffer Management for Asynchronously Arriving Variable-length Packets

3.1. Photonic Buffer Architecture for Prioritized Buffer Management

Figure 4 (a) shows the $N \times 1$ photonic buffer architecture in the case of $N = 4$. For prioritized buffer management, the $B \times 2$ optical switch is added to the fundamental $N \times 1$ photonic buffer shown in Fig. 2 (a). The first optical switch is used to discard low-priority packets as well as overflowed packets if the packets do not reach the buffer. The second switch is used to discard low-priority packets that have already entered the buffer. vPBSO can be imple-

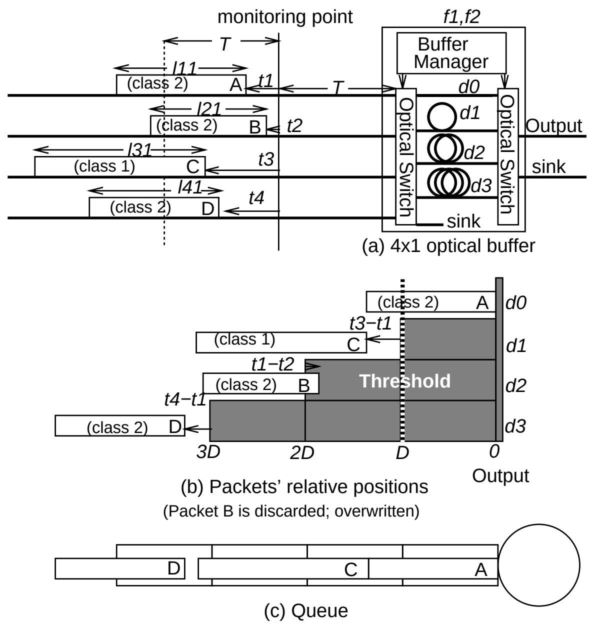

Figure 4. Behavior of packets arriving at the optical buffer (vPBSO)

mented by using this buffer architecture, which is described in more detail in the next subsection.

3.2. vPBSO: Variable-length-packet-capable Photonic Partial Buffer Sharing with Overwriting

vPBSO is a buffer management method for providing priority control with multiple priority classes. Here, we consider two priority classes: the priority of class 1 is higher than that of class 2. Under vPBSO, the buffer manager allows every arriving packet to enter the buffer when the buffer occupancy is smaller than the threshold. When the buffer occupancy is equal to or larger than the threshold, an arriving class-2 packet is allowed to take the tail of the queue. On the other hand, an arriving class-1 packet is allowed to take the next position of class-1 packets waiting behind the threshold in the queue. If class-2 packets are waiting at the same position, some of them are overwritten. The overwritten packets are regarded as discarded. As we described, the behavior is the same as PBS if the buffer occupancy is smaller than the threshold, but different otherwise.

Now, we describe the behavior of the vPBSO buffer manager. vPBSO can be implemented by using an $N \times 1$ optical buffer ($N = 4$), as shown in Fig. 4 (a). As an example, we show the behavior in the case of four packets, A, B, C, and D, arriving with classes 2, 2, 1, and 2, respectively. The threshold is set to $Th = D$. Figure 4 (b) represents the relative positions of the packets from the output port just after packet A is switched to delay line d_0. The buffer occupancy exceeds the threshold Th after packet A enters. Accordingly, class-2 packet B is at first allowed to enter delay line d_2 as in Fig. 4 (b). However, packet B is overwritten by class-1 packet C, which is handled after packet B. In this case, the overwriting function is handled by driving the first optical switch appropriately. Namely, the first optical switch sends packet B to the sink (not d_2) and sends packet C to delay line d_1. Class-2 packet D is then allowed to enter delay line d_3 behind packet C. Figure 4 (c) shows a result of buffer management of the four packets.

As described the above, the first optical switch discards overwritten packets if the packets do not reach the buffer. On the other hand, the second switch discards overwritten packets that have entered the buffer by switching the packets to the sink. By using these two optical switches, the buffer avoids physical crash between arriving packets and corresponding overwritten packets.

The buffer manager has the following information for vPBSO:

- p thresholds, where p is the number of classes ($Th_1, Th_2, Th_3, \ldots, Th_p$). Each threshold is a multiple of the unit delay D for the fiber delay line. The condition $BD = Th_1 \geq Th_2 \geq Th_3 \geq \ldots \geq Th_p$ must be satisfied.
- Variables $f_1, f_2, \ldots, f_p$, where f_i represents the time at which a new packet with class i is granted to enter the buffer. This time is relative to the starting time of the time cycle T. The buffer occupancy b is identical to f_p.

We next describe the vPBSO algorithm as follows:

(1) For $n = 1, 2, \ldots, N$, the following procedures are operated in order if a new packet arrives at port n during a time cycle T.

(1-1) Calculate the delay $\Delta_n D = \lceil \frac{f_i - t_n}{D} \rceil D$ for an arriving packet with class i, length l_{ni}, and gap t_n. The packet enters delay line d_{Δ_n} if $\Delta_n < B = Th_1$, while it is discarded if $\Delta_n \geq B$.

(1-2) When the packet enters the delay line, the buffer occupancy f_i is changed to $f_i \leftarrow t_n + l_{ni} + \Delta_n D$ to handle packets at the following ports appropriately. Other variables $f_1, f_2, \ldots, f_p$ are changed as follows:

- If the buffer occupancy b just before the arriving packet is allowed to enter is smaller than Th_p, then set $f_k \leftarrow f_i$ for all $k = 1, 2, \ldots, p$.

- If the buffer occupancy b just before the arriving packet is allowed to enter is satisfied with $Th_j \leq b < Th_{j-1}$ $(1 < j \leq p)$ and priority class i of the arriving packet is less than or equal to j (i.e., $i \geq j$), then set $f_k \leftarrow f_i$ for $k = j, j+1, \ldots, p$.
- If the buffer occupancy b just before the arriving packet is allowed to enter is satisfied with $Th_j \leq b < Th_{j-1}(1 < j \leq p)$ and priority class i of the arriving packet is greater than j (i.e., $i < j$), then set $f_k \leftarrow f_i$ for $k = 1, 2, \ldots, j-1$. Also, set $f_k \leftarrow \max(f_{j-1}, f_k)$ for $k = j, j+1, \ldots, p$ to prevent packets with class j or lower from being placed in positions for ahead of packets with a class higher than j.

(2) After all arriving packets at all ports are handled, change $f_1, f_2, \ldots, f_p$ to provide appropriate buffer management during the next cycle.

- If the buffer occupancy b is less than or equal to threshold Th_p, then change $f_k \leftarrow \max(f_k - T, 0)$ for all $k = 1, 2, \ldots, p$.
- If the buffer occupancy b has a value between two adjacent thresholds Th_i and Th_{i-1} ($Th_i \leq b < Th_{i-1}(1 < i \leq p)$), then set $f_l \leftarrow f_l - T$ for $l = i, i+1, \ldots, p$.
  Also, set $f_k \leftarrow \min(f_p, \max(f_k - T, Th_i))$ for $k = 1, 2, \ldots, i-1$.

The buffer manager for vPBSO is simple, because for prioritized buffer management, it does not require the number of packets stored in the buffer to be dependent on the priority class or on where the packets are located in the buffer. It only requires p variables: that is, $f_1, f_2, \ldots, f_p$. The complexity for handling a packet is $O(p)$, which depends on the number of renewals of the variables $\{f_1, f_2, \ldots, f_p\}$ in step (1-2). By applying parallel processing, the actual processing time can be minimized. When we focus on a priority class, the arrival order of the packets at each port is identical to the departure order, except for discarded packets. The consistency of the arrival/departure order provides better performance for higher layer applications than when there is inconsistency.

4. Performance of Prioritized Buffer Management

We now investigate the performance of prioritized buffer management through simulation experiments. We first show that vPBSO can provide different levels of performance in terms of packet loss probability for different priority classes. We then compare the performance of vPBSO with that of vPBS and non-priority queueing (i.e., best effort) and discuss a suitable control method of the optical buffer.

Through our evaluations, we focus on an output buffer of a photonic packet switch. We generate 10^9 packets. Packets arrive according to a Poisson process

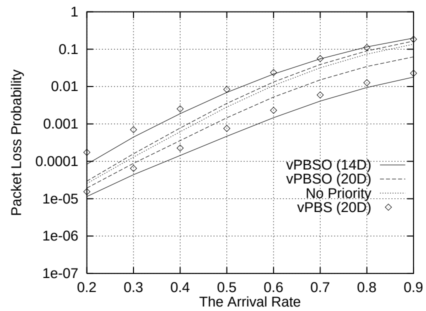

(a) $B = 25$, $D = 0.50$

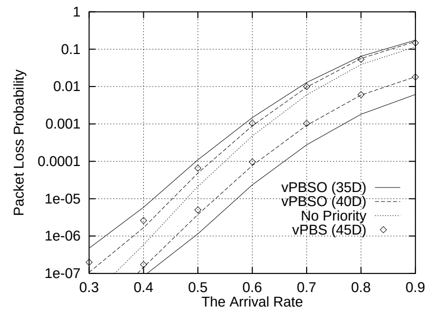

(b) $B = 50$, $D = 0.50$

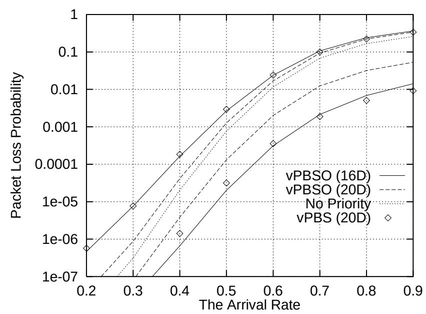

(c) $B = 25$, $D = 1.00$

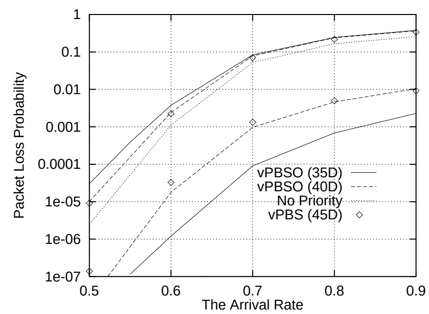

(d) $B = 50$, $D = 1.00$

Figure 5. Comparison of vPBSO and vPBS ($\rho_1 : \rho_2 = 1 : 3$)

with rates a. The class-1 packets to class-2 packets ratio is $\rho_1 : \rho_2$. The lengths of the arriving packets are distributed exponentially with mean $1/\mu = 1$. We set the numbers of fiber delay lines to $B = 25$ and 50. The unit lengths of the fiber delay lines are set at $D = 0.5$ and 1.0, because these values provide better performance than $D = 0.3$ at a lower arrival rate than 0.8 (See Fig. 3).

Figure 5 shows the performance of vPBSO, vPBS, and non-priority queueing for different numbers of fiber delay lines (B) and different unit lengths of the fiber delay lines (D). The ratio of class-1 packets to class-2 packets is fixed at $1 : 3$. The values in parentheses for each algorithm are the thresholds Th_2. The results above the case of non-priority queueing represent the performance of class-2 packets, while those below represent the performance of class-1 packets. The horizontal axis is the arrival rate of packets at the optical buffer. The vertical axis is the packet loss probability. We can see that vPBSO provides different levels of performance depending on the priority class, and that it improves the class-1 packet loss probability more than non-priority queueing does. For example, when we use $Th_2 = 14D$ for vPBSO in Fig. 5(a), the class-1 packet loss probability is improved by about one order of magnitude compared with that for non-priority queueing. The amount of improvement depends on the parameters B, Th_2, and D. For example, the class-1 packet loss probability with $B = 50$, $Th_2 = 35D$, $D = 1.0$ is improved by more than two orders of magnitude, as shown in Fig. 5(d).

Furthermore, from these results, we find that vPBSO with proper threshold selection provides better performance for class-1 packets than vPBS, and the performance for class-2 packets is better or at least almost the same. For example, such a condition occurs when we select $Th_2 = 14D$ for vPBSO in the case of $B = 25$ and $D = 0.5$ and compare the performance to vPBS with $Th_2 = 20D$, as shown in Fig. 5(a). On the other hand, when we select $Th_2 = 20D$ for vPBSO, the performance for class-2 packets is better than that of vPBS, while the performance for class-1 packets is worse.

We obtained similar conditions for other sets of B and D. For example, we find improvement by using vPBSO as shown in Fig. 5(c): vPBSO with $Th_2 = 16D$ outperforms vPBS with $Th_2 = 20D$. However, it is limited by a lower arrival rate (i.e., $a < 0.7$), and the performance of vPBSO deteriorates when the arrival rate is larger. This can be understood as follows. In principle, vPBSO provides higher buffer utilization than vPBS, as the authors described previously in Ref. [20]. Since we keep the effectiveness with a lower arrival rate, the performance for class-1 packets is improved while maintaining the performance for class-2 packets. However, we may have a non-overwritten block of overwritten packets in the buffer, since two packets usually have different lengths and arrive asynchronously. When the arrival rate is higher, we find that a number of non-overwritten blocks of overwritten packets unexpectedly occupy the buffer. Accordingly, the performance of vPBSO does not improve. We thus conclude that once appropriate buffer management scheme and its threshold values for vPBSO and vPBS are determined depending on the arrival rate of packets, we can achieve high-performance prioritized buffer management. This can be realized by some traffic load estimation method.

vPBSO provides better performance than vPBS when the ratio of class-1 packets to class-2 packets is small. Figure 6 shows the performance of vPBSO and vPBS when the ratio is $\rho_1 : \rho_2 = 1 : 1$. We find that for both vPBSO and vPBS, the differences between the packet loss probabilities of class-1 and class-2 packets are smaller than those in the cases shown in Fig. 5(a). Accordingly, the difference between vPBSO and vPBS also becomes smaller.

5. Concluding Remarks

We have investigated a photonic buffer architecture supporting prioritized buffer management methods in photonic packet switches to provide DiffServ Assured Forwarding. Since a photonic packet switch requires high-speed buffer management, as well as a high-speed address-lookup function, a less complex

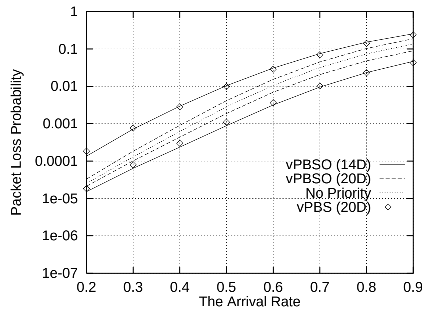

Figure 6. Comparison of vPBSO and vPBS ($B = 25, D = 0.50, \rho_1 : \rho_2 = 1 : 1$)

strategy should be adopted. We have proposed a buffer architecture and algorithm for vPBSO. It supports asynchronously arriving variable-length packets. vPBSO can be used for optical buffers since its complexity, $O(p)$ (p is the number of priority classes) is small. We have shown that at lower arrival rates, vPBSO provides better performance than vPBS for high-priority packets, and the performance for low-priority packets is also better or at least almost the same. We furthermore have shown that we achieve high-performance prioritized buffer management by appropriately selecting either vPBSO and vPBS depending on the arrival rate of the packets. vPBSO is especially effective when the arrival rate of higher priority packets is lower than that of lower priority packets, which is a likely situation in the prioritized system.

References

[1] E. Rosen, A. Viswanathan, and R. Callon, “Multiprotocol label switching architecture,” *RFC 3031*, January 2001.

[2] E. Mannie *et al.*, “Generalized multi-protocol label switching (GMPLS) architecture (draft-ietf-ccamp-gmpls-architecture-03.txt),” *IETF Internet Draft (Work in Progress)*, August 2002.

[3] I. Chlamtac, A. Ganz, and G. Karmi, “Lightpath communications: An approach to high bandwidth optical WAN’s,” *IEEE Transactions on Communications*, vol. 40, pp. 1171–1182, July 1992.

[4] N. Wada, H. Harai, W. Chujo, and F. Kubota, “Photonic packet routing based on multi-wavelength label switching using fiber Bragg gratings,” *ECOC 2000 Technical Digest (26th European Conference on Optical Communication)*, vol. 4 (No. 10.4.6), pp. 71–72, September 2000.

[5] N. Wada, H. Harai, W. Chujo, and F. Kubota, “Photonic variable length packet routing based on multi-wavelength label switch using multi-section fiber Bragg gratings and supercontinuum light source,” *in Proceedings of Optical Amplifiers and Their Applications/ Bragg Graring, Photosensitivity, and Poling in Grass Waveguides 2001 (OAA/BGPP 2001)*, vol. JW4, July 2001.

[6] K. Kitayama and N. Wada, “Photonic IP routing,” *IEEE Photonic Technology Letters*, vol. 11, pp. 1689–1691, December 1999.

[7] D. K. Hunter and I. Andonovic, “Approaches to optical Internet packet switching,” *IEEE Communications Magazine*, vol. 38, pp. 116–122, September 2000.

[8] K. Habara, H. Sanjo, H. Nishizawa, Y. Yamada, S. Hino, I. Ogawa, and Y. Suzaki, “Large-capacity photonic packet switch prototype using wavelength routing techniques,” *IEICE Transactions on Communications*, vol. E83-B, pp. 2304–2311, October 2000.

[9] D. K. Hunter, M. C. Chia, and I. Andonovic, “Buffering in optical packet switches,” *IEEE/OSA Journal of Lightwave Technology*, vol. 16, pp. 2081–2094, December 1998.

[10] L. Tancevski, S. Yegnanarayanan, G. Castanon, L. Tamil, F. Masetti, and T. McDermott, “Optical routing of asynchronous, variable length packets,” *IEEE Journal on Selected Areas in Communications*, vol. 18, pp. 2084–2093, October 2000.

[11] A. Ge, L. Tancevski, G. Castanon, and L. S. Tamil, “WDM fiber delay line buffer control for optical packet switching,” *Proc. SPIE Vol. 4233 (OptiComm 2000)*, pp. 247–256, October 2000.

[12] F. Callegati, “Optical buffers for variable length packets,” *IEEE Communications Letters*, vol. 4, pp. 292–294, September 2000.

[13] M. Murata and K. Kitayama, “Ultrafast photonic label switch for asynchronous packets of variable length,” *Proceedings of IEEE INFOCOM 2002*, June 2002.

[14] Y. Lin and J. A. Silverster, “Priority queueing strategies and buffer allocation protocols for traffic control at an ATM integrated broadband switching system,” *IEEE Journal on Selected Areas in Communications*, vol. 9, pp. 1524–1536, December 1991.

[15] F. Callegati, G. Corazza, and C. Raffaelli, “Exploitation of DWDM for optical packet switching with quality of service guarantees,” *IEEE Journal on Selected Areas in Communications*, vol. 20, pp. 190–201, January 2002.

[16] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, “An architecture for Differentiated Services,” *RFC 2475*, December 1998.

[17] S. Floyd and V. Jacobson, “Link-sharing and resource management models for packet networks,” *IEEE/ACM Transactions on Networking*, vol. 3, August 1995.

[18] M. Shreedhar and G. Varghese, “Efficient fair queueing using deficit round robin,” *ACM Computer Communication Review*, vol. 25, pp. 231–242, October 1995.

[19] J. Heinanen, F. Baker, W. Weiss, and J. Wroclawski, “Assured Forwarding PHB groups,” *RFC 2597*, June 1999.

[20] H. Harai and M. Murata, “Prioritized buffer management in photonic packet switches for DiffServ Assured Forwarding,” *The 6th IFIP Working Conference on Optical Network Design and Modelling (ONDM 2002)*, February 2002.

[21] H. Harai and M. Murata, “Performance analysis of prioritized buffer management in photonic packet switches for DiffServ Assured Forwarding,” *Proceedings of SPIE vol. 4874 (OptiComm 2002)*, pp. 298–309, July 2002.

[22] L. Kleinrock, *Queueing systems Volume II: Computer applications*. A Wiley-Interscience Publication, 1976.