

Prioritized Routing and Burst Segmentation for QoS in Optical Burst-Switched Networks

Vinod M. Vokkarane and Jason P. Jue

Department of Computer Science, The University of Texas at Dallas, P.O. Box 830688, Richardson, TX 75083
tel: +1 (972) 883-4429, fax: +1 (972) 883-2349, e-mail: vinod@utdallas.edu, jjue@utdallas.edu

Abstract: Deflection routing and burst segmentation are two techniques for resolving contention in optical burst-switched networks. By selectively applying these techniques to bursts with different QoS requirements, we can offer a range of differentiated services at the optical layer.

© 2001 Optical Society of America

OCIS codes: (060.4250) Networks

1 Introduction

The amount of raw bandwidth available on fiber optic links has increased dramatically with advances in dense wavelength division multiplexing (DWDM). In order to efficiently utilize this bandwidth, an all-optical transport method, which avoids optical buffering while handling bursty traffic, and which supports fast resource provisioning and asynchronous transmission of variable sized packets, must be developed. Optical Burst Switching (OBS) is one such method for transporting traffic directly over a bufferless optical WDM network. In optical burst switched networks, bursts of data consisting of multiple packets are switched through the network all-optically. A control message (or header) is transmitted ahead of the burst in order to configure the switches along the burst's route. The control message may specify the duration of the burst in order to let a node know when it may reconfigure its switch for the next burst, a technique known as *Delayed Reservation* (DR) [1].

A major concern in optical switched networks is contention, which occurs when multiple bursts contend for the same link. In this paper, we will refer to the burst which arrives to the switch first as the *original burst*, and the burst which arrives to the switch later as the *contending burst*. One approach to resolving contention is to use *deflection routing*, in which the contending burst is sent to an alternate output port. *Prioritized routing* is an approach in which the choice of packet/burst to deflect is made based on priority [3].

In existing contention resolution schemes for optical burst switched networks, when contention between two bursts cannot be resolved through other means, one of the bursts will be dropped in its entirety, even though the overlap between the two bursts may be minimal. For certain applications, which have stringent delay requirements but relaxed packet loss requirements, it may be desirable to lose a few packets from a given burst rather than losing the entire burst. An approach to reducing the packet losses during contention is *burst segmentation*, in which only those packets of the original burst which overlap with a contending burst will be dropped [2]. To further reduce packet loss, the overlapping packets may be deflected rather than dropped.

Another important issue in optical burst-switched networks is how to provide differentiated service in order to support the various quality of service (QoS) requirements of different applications. In [1], an offset scheme was proposed for isolating classes of bursts, such that low-priority bursts do not cause contention losses for high-priority bursts; fixed and variable fiber delay line buffers were also utilized to further reduce blocking. In this paper, we will combine burst segmentation with prioritized deflection routing in order to offer differentiated services at the optical layer. We will assume that there are two priority classes, and that a high-priority burst is one which has low delay and loss tolerance while a low-priority burst has higher level of delay and loss tolerance.

2 Burst Segmentation And Deflection

To overcome some of the limitations of optical burst switching, burst segmentation can be used to minimize packet loss during contention. When contention occurs, a burst is divided into multiple segments, and only those packets of a given burst which overlap with segments of another burst will be dropped. If switching time is non-negligible, then additional packets may be lost when the output port is switched from one burst to another. Segmentation is used primarily to minimize loss of high priority bursts. There are two approaches for segmenting a burst when contention occurs. The first approach is to segment the tail of the original burst, and the second approach is to segment the head of the contending burst. A significant advantage of segmenting the tail of bursts rather than segmenting the head is that there is a better chance of in-sequence delivery of packets at the destination, assuming that dropped packets are retransmitted at a later time. In this paper, we will assume that the tail of the original burst will be segmented when segmentation takes place (Fig. 1). When a burst is segmented, its control message is updated accordingly. For the case in which segmentation occurs in the middle of a packet, the fractional packet is lost.

Burst segmentation can also be implemented with deflection. Rather than dropping the tail segment of the original burst, we can either deflect the entire contending burst, or we can deflect only the tail segment of the original burst. Implementing segmentation with deflection increases the probability that a burst's packets will reach the destination, and hence improves performance. At each node, one or more alternate deflection ports can be specified for each destination. The order in which the alternate deflection ports are attempted is determined by a shortest-path policy.




Fig. 1. Segmentation policy for two contending bursts.




Fig. 2. NSFNET topology with 14 nodes.

The QoS scheme is implemented by selectively choosing which burst (original or contending) to segment or deflect during contention. We define the following policies for handling contention between two bursts:

- *Segment First and Deflect Policy (SFDP)*: The contending burst wins the contention. The original burst is segmented, and its segmented tail may be deflected if an alternate port is free, otherwise the tail is dropped.
- *Deflect First and Drop Policy (DFDP)*: The contending burst is deflected to an alternate port, if no alternate port is free the contending burst is dropped.
- *Deflect First, Segment and Drop Policy (DFSDP)*: The contending burst is deflected to a free port if available, otherwise the original burst is segmented and its tail is dropped, while the contending burst is transmitted.

The decision of which of these three policies to implement when a contention occurs is based on the priorities of the two bursts involved, as well as the lengths of the bursts, as shown in Table 1.

Table 1. QoS policies for various contention situations.

Condition	Original Burst Priority	Contending Burst Priority	Longer Remaining Burst	Scheme 1 Policies	Scheme 2 Policies
1	High	High	Contending	DFSDP	SFDP
2	High	Low	Contending	DFDP	DFDP
3	Low	High	Contending	SFDP	SFDP
4	Low	Low	Contending	DFSDP	SFDP
5	High	High	Original	DFDP	DFDP
6	High	Low	Original	DFDP	DFDP
7	Low	High	Original	SFDP	SFDP
8	Low	Low	Original	DFDP	DFDP

The rows of Table 1 show the various contention situations in a two priority network. *Longer Remaining Burst* indicates which of the two contending burst has a greater number of packets remaining to be transmitted from the point of contention. If all else is equal, the burst with fewer remaining packets will be given slightly lower priority in order to minimize the number of packets lost or deflected.

To better understand the table, let us consider Condition 2. The original burst has high priority while the contending burst has low priority, and the contending burst is longer than the original burst. This contention is resolved by DFDP, where we first attempt to deflect the low-priority burst, and then drop the low-priority burst if no deflection ports are available. On the other hand, in Condition 7, a high-priority contending burst contends with a low-priority original burst, and the number of packets in the high-priority burst is less than the number of packets in the tail of the low-priority burst. In this case, we segment the tail of the low-priority burst and attempt to deflect the tail if a deflection port is available; otherwise, the tail is dropped (SFDP).

We define two different QoS schemes that differ in the handling of Conditions 1 and 4, i.e., when the contention is between two bursts of the same priority, and the contending burst is longer than the original burst. Scheme 1 will first attempt to deflect the contending burst, and will segment the original burst only if no deflection port is available (DFSDP). Scheme 2 will always segment and attempt to deflect the tail of the original burst (SFDP).

3 Simulation Results

In order to evaluate the performance of the segmentation and deflection schemes, a simulation model was developed. Burst arrivals to the network are Poisson with rate λ . Burst length is exponentially distributed with average burst length of $1/\mu = 100$ ms. Transmission rate is 10 Gbps. Packet length is 1500 bytes. Switching time is 10 μ s. There is

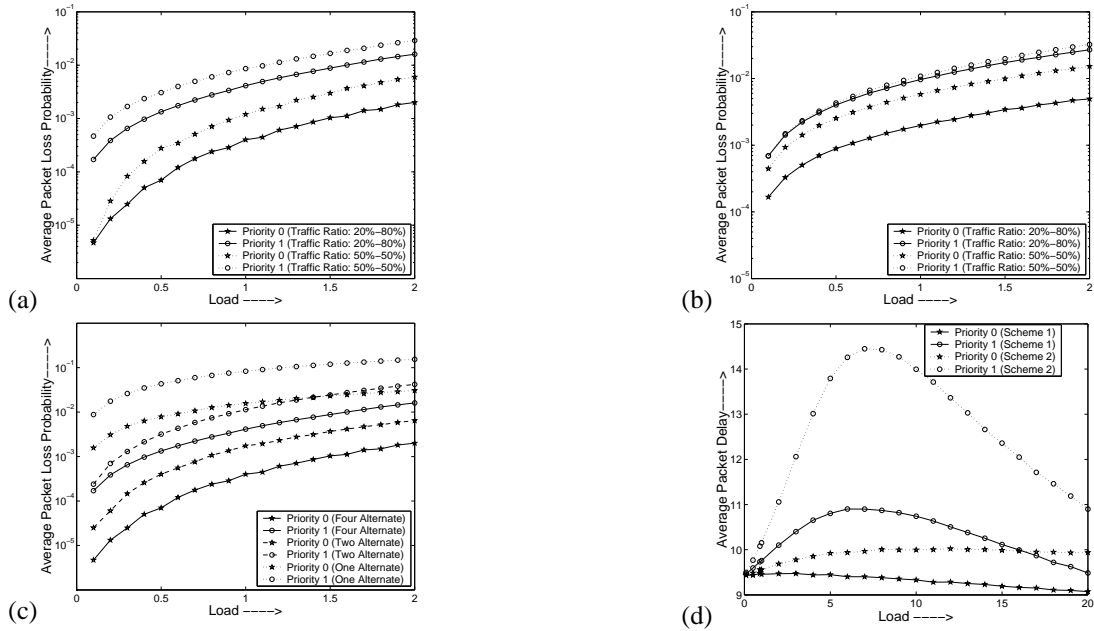


Fig. 3. Packet loss probability versus load for different traffic ratios using (a) Scheme 1, (b) Scheme 2, and (c) Scheme 1 with different number of alternate deflection ports. (d) Average packet delay versus load using Scheme 1 and Scheme 2.

no buffering or wavelength conversion at nodes. Burst arrivals are uniformly distributed over all sender-receiver pairs. Dijkstra shortest-path routing algorithm is used to find the path between all node pairs. Figure 2 shows the 14-node NSFNET on which the simulation was implemented. The distances shown are in km.

Figures 3(a) and 3(b) plot the packet loss probability versus load for high-priority (priority 0) and low-priority (priority 1) packets, using Scheme 1 and Scheme 2 respectively. Each shows packet losses for the case in which there is an equal amount of high-priority and low-priority traffic, and the case in which 20% of the traffic is high priority and 80% of the traffic is low priority. We observe that the loss of high-priority packets are lower than that for low priority packets. We also observe that Scheme 1 performs better than Scheme 2 at these loads; thus, at low loads, it is better to attempt deflection before segmentation when two bursts are of equal priority.

Figure 3(c) plots total packet loss probability versus load for different number of alternate deflection ports with 20% of high-priority and 80% of low-priority traffic. We observe that there is a significant improvement when we use two alternate deflection ports instead of one alternate ports, while there is less improvement from two to four alternate deflection ports. This result is due to the low nodal degree of NSFNET (Figure 2) and may differ for other networks.

Figure 3(d) plots total delay versus load with 20% of high-priority and 80% of low-priority traffic for the two QoS schemes. Scheme 2 has lower delays compared to Scheme 1, as Scheme 2 follows the segment-first approach rather than the deflect-first approach. The delay for high-priority bursts remains in a consistent range, while the low-priority bursts have higher delay due to multiple deflections. At very high load, bursts which are farther from their destination are less likely to reach their destination compared to those bursts which are close to their destination; thus, the average delay will eventually decrease as load increases.

4 Conclusion

In this paper we apply the concepts of burst segmentation and priority-based deflection routing algorithms to provide differentiated services in optical burst-switched networks. The high-priority bursts have significantly lower losses and delay than the low-priority bursts, and the policies which incorporate deflection tend to perform better than the policies with limited deflection or no deflection.

References

1. M. Yoo, C. Qiao, and S. Dixit, "QoS Performance of Optical Burst Switching in IP-Over-WDM Networks," *IEEE Journal on Selected Areas in Communications*, vol. 18, no.10, pp. 2062-2071, Oct. 2000.
2. V. Vokkarane, J. Jue, and S. Sitaraman, "Burst Segmentation: An Approach For Reducing Packet Loss In Optical Burst Switched Networks", UTD Technical Report UTDCS-20-01, September 2001.
3. S. Yao, B. Mukherjee, and S.J.B. Yoo, "A Comparison Study Between Slotted and Unslotted All-Optical Packet-Switched Network with Priority-Based Routing," *Proceedings, OFC 2001*, TuK2, Anaheim, CA, March 2001.