A Fair Packet-Level Performance Comparison of OBS and OCS

Xin Liu*, Chunming Qiao*, Xiang Yu†, Weibo Gong‡

**Department of Computer Science and Engineering, State University of New York at Buffalo, New York 14228 {xliu8, qiao }@cse.buffalo.edu*

†Department of Computer Science, Frostburg State University, Frostburg, Maryland 21532 xsyu@frostburg.edu

‡Department of Electrical and Computer Engineering, University of Massachusetts Amherst, Massachusetts 01003 gong@ecs.umass.edu

Abstract: The performance of optical burst switching (OBS) and optical circuit-switching (OCS) in terms of packet delivery and packet loss ratios is compared based on a given input traffic and comparable settings. Our results show that OBS performs better in most cases.

1. Introduction

Optical burst-switching (OBS) and optical circuit-switching (OCS) have been proposed as two promising switching paradigms for all-optical networks. They have also stirred interesting debates within the research communities. In an OCS network, data is transmitted over an existing lightpath with sufficient bandwidth. In an OBS network, packets are assembled into bursts, which are injected into the network without setting up a lightpath. Some believe that OBS is suitable for bursty data traffic as it provides statistical multiplexing gains and reduces signaling delay, thereby achieving a higher throughput than OCS, given the same network capacity; Others disagree, pointing to the potentially high burst loss probability in OBS as a main problem. Being able to determine whether OBS or OCS performs better under otherwise identical or similar settings is thus important. However, fair, quantitative performance comparisons between OCS and OBS are difficult as it may be considered to be analogous to comparing an orange and an apple. For example, a common performance metric used for OCS is the blocking probability of lightpath (establishment) requests, which however is not directly comparable to burst (or packet) loss ratio, a common performance metric used for OBS.

To the best of our knowledge, there has been no direct quantitative performance comparison between circuit and packet (or burst) switching even for electronic networks involving multiple traffic flows among multiple source and destination pairs. Although there have been a few attempts so far to compare OBS and OCS [1,2,3], the need to establish a common ground on which fair performance comparisons between OCS and OBS can be made clearly exists today more than ever as the research community continues this interesting debate over OBS and OCS.

2. Major contributions

In this paper, we compare the performance of OCS and OBS in terms of common performance metrics: packet delivery ratio and packet loss ratio. We choose not to compare their implementation costs which tend to change over time and are dependent on the demand-supply economics. We use an identical input traffic which consists of multiple flows of packets in both cases. In addition, we assume that the two types of networks have the same topology and same electronic buffering capacity at the edge, and that every switching fabric, be it a slow optical cross-connect for OCS or a fast OBS switching fabric, can perform all-optical wavelength conversion. Moreover, both networks have the same signaling functionalities and use similar distributed routing algorithms and distributed control mechanisms. More specifically, unlike most of the previous work, we assume that in OCS, lightpath requests are processed hop-by-hop along a chosen path just as burst control packets in OBS; In addition, we assume that in OBS, bursts may be stored at the edge for retransmission, just as in OCS, the data can be stored until a lightpath is found (possibly after several retransmissions of the lightpath request). We consider both the case with unlimited buffer and the case with limited buffer at the edge.

A main challenge in keeping the comparison as fair as possible is to choose the input traffic with reasonable characteristic. In this study, we assume that the packet flows arrive according to Poisson process with a certain mean arrival rate and mean duration. However, since our focus is to determine if OBS can perform better than OCS when supporting bursty traffic, we also assumed that within a given flow, packets arrive according to a self-similar process. In addition, we consider both the case where packets have a limited delay budget (i.e., deadline) and the case with unlimited delay budget (i.e., no deadline).

Another main challenge is to choose an appropriate bandwidth allocation and traffic grooming strategy for use in OCS. For simplicity (benefit of OCS), we assume that to support a bursty flow with an average bandwidth of *B* (in bps) a connection of only bandwidth *B* needs to be established (our preliminary results have shown that allocating bandwidth higher than *B* to each flow results in much worse performance for OCS). If/when *B* is lower than the per wavelength capacity (e.g., 10 Gbps), such a connection will be groomed onto an existing lightpath with sufficient residual bandwidth. In addition, to facilitate traffic grooming and increase the resource utilization, a connection from source *S* to destination *D* does not have to follow the shortest path; Instead, it may follow any of the *K*-shortest paths (where *K*=4). Furthermore, it may go through an intermediate node *I* over two separate lightpaths when a direct lightpath between *S* and *D* is not available. This helps to reduce lightpath request blocking as well, although it also tends to use a non-shortest path to route the connection, and more important, does not provide 100% end-to-end all-optical transparency as the data belonging to this flow has to go through O/E/O conversions at node *I* (even though the node has the full wavelength conversion capability). Other possible input traffic characteristics, and bandwidth allocation and traffic grooming strategies will be studied in the future.

3. Simulation studies

The physical topology we simulated in this study is the NSFNET (with 16 nodes and 25 links). We assume that every node in the network consists of an electronic component providing the functions of an edge node (with an electronic buffer), and an optical switching fabric controlled by the electronic component. In an OCS network, the electronic component establishes connections for packet flows, performs traffic grooming and processes lightpath requests. In an OBS network, the electronic component assembles packets into a burst and process burst control packets. We use the published propagation delay value (on the order of ms) on each link in the NSFNET but ignore the processing delay (on the order of μs) as the former dominates.

We assume that traffic flows arrive in a Poisson process with an average arrival rate λ (with a default value of 100 flows/s), and an average flow duration of $1/\mu$ (fixed at 1 second). Define $\rho = \lambda/\mu$ to be the flow intensity. We also assume that every node in the network could be the source or the destination of a traffic flow with an equal probability. In addition, we assume that the number of wavelengths in a fiber is *w*=8, and the capacity of a wavelength is *c* =10Gbps. For simplicity, we assume a fixed packet length of 1 kb, and within a flow, packets arrive in a self-similar process with the Hurst parameter equal to 0.9 (using the algorithm in [4]), which varies from 10^6 to $4\text{x}10^6$ (packets/second). Accordingly, when simulating the OCS network, a flow requires a connection with an average bandwidth between 1 to 4 Gbps, which is between 10% and 40% of the capacity of a wavelength. When simulating the OBS network, we use a mixed time-length based burst assembly algorithm with the maximum burst assembly time and maximum burst length being 10ms and 50K packets respectively (these values are larger than usually and are chosen to reduce the simulation time) and a simple Horizon-based scheduling algorithm.

We study the following four cases based on the two constraints: packet delay (budget) and buffer size at edge.

3.1 unlimited packet delay and unlimited buffer (UDUB)

In this case, if a connection request or a burst is blocked, it will be retransmitted for an unlimited number of times until it reaches its destination (or the simulation ends). We assume that retransmission will commence only after a time out whose value is set to the maximal round-trip delay in the NSFNET (about 55ms). Since all data can be buffered, there is no packet loss. We compare OBS and OCS in terms of the ratio of the number of packets successfully delivered by the time the simulation ends over the total number of packets generated during the simulation time (alternatively, one may compare them in terms of packet throughput and delay).

3.2 limited packet delay and unlimited buffer (LDUB)

In this case, a packet will be lost if it still has not reached its destination by its deadline, which is assumed to be r times of the maximum round-trip delay (by default, r=3 or 165ms). In an OBS network, this means that a bust will be dropped if it does not arrive at the destination after (r-1) times of retransmissions (or r attempts). Similarly, in an OCS network, if a connection cannot be setup after (r-1) times of retransmissions, only the first few packets arrived before the first attempt are dropped. However, the connection request will be retransmitted until the last packet belonging to the flow is dropped. We compare OBS and OCS in terms of packet loss ratio, which is the number of packets lost over the total number of packets generated during the simulation time.

3.3 limited packet delay and limited buffer (LDLB)

In this case, packet loss can also be due to buffer overflow at the edge. When buffer size is small, packet loss is mainly due to buffer overflow. When buffer size is above a threshold, packet loss is almost entirely due to limited delay so the case resembles that of LDUB. In addition to packet loss ratio, we are also interested in knowing the threshold buffer sizes required in OBS and OCS, respectively, to minimize packet loss due to buffer overflow.

3.4 unlimited packet delay and limited buffer (UDLB)

In this case, packet loss is due only to buffer overflow at the edge. So when the buffer size is large enough, this case resembles that of UDUB. We expect that the threshold buffer sizes required in OBS and OCS, respectively, are larger than those required in the case of LDLB because some packets may be dropped due to limit delay in LDLB whereas in this case, more packets will be stored in the buffer.

4. Simulation results

In each simulation run, we have simulated 10,000 flows. We have found that simulating 100,000 flows (in the case of OCS) in each run does not differ much, and certainly does not change the relative performance of OBS and OCS. Fig. 1(a) shows that in the case of UDUB, the packet delivery ratios in the OCS and OBS networks decrease with the flow intensity. However, when the flow intensity is greater than 140, the packet delivery ratio of OBS remains close to 100% while that of OCS drops to below 99.5%. This means that more packets have to wait in the buffer in OCS for connections to be established, while OBS can send out more packets within the same amount of time due to its shorter signaling delay, and small granularity which facilitates statistical multiplexing.

Fig. 1(b) shows that in the case of LDUB, the packet loss ratios in the OCS and OBS networks increases with the flow intensity. For the same reasons mentioned above for UDUB, the packet loss ratio of an OBS network is always smaller than that of OCS network and the greater the flow intensity, the larger the difference between them.

Fig. 1(c) and Fig. 1(d) show that in the case of LDLB, the packet loss ratios in OCS and OBS decrease sharply when the buffer size of each node increases but then stabilize when the buffer size is large enough. These results show that each node needs a buffer size of at least 6Gb for an OCS network but only 1.4Gb for an OBS network to virtually eliminate packet loss due to buffer overflow. Beyond these threshold buffer sizes, the case resembles that of LDUB in that the packet loss ratios are largely due to limited delay, which are around 0.2% in OCS and 0.04% in OBS (the same as in Fig. 1(b)). Although not shown for lack of space, the results for the case of UDLB are similar except that each node needs a buffer size of at least 7Gb for an OCS network and at least 2Gb for an OBS network to minimize packet loss, and even at these buffer sizes, OCS still delivers fewer packets than OBS within the same amount of time, similar to the case of UDUB (shown in Fig. 1(a)).

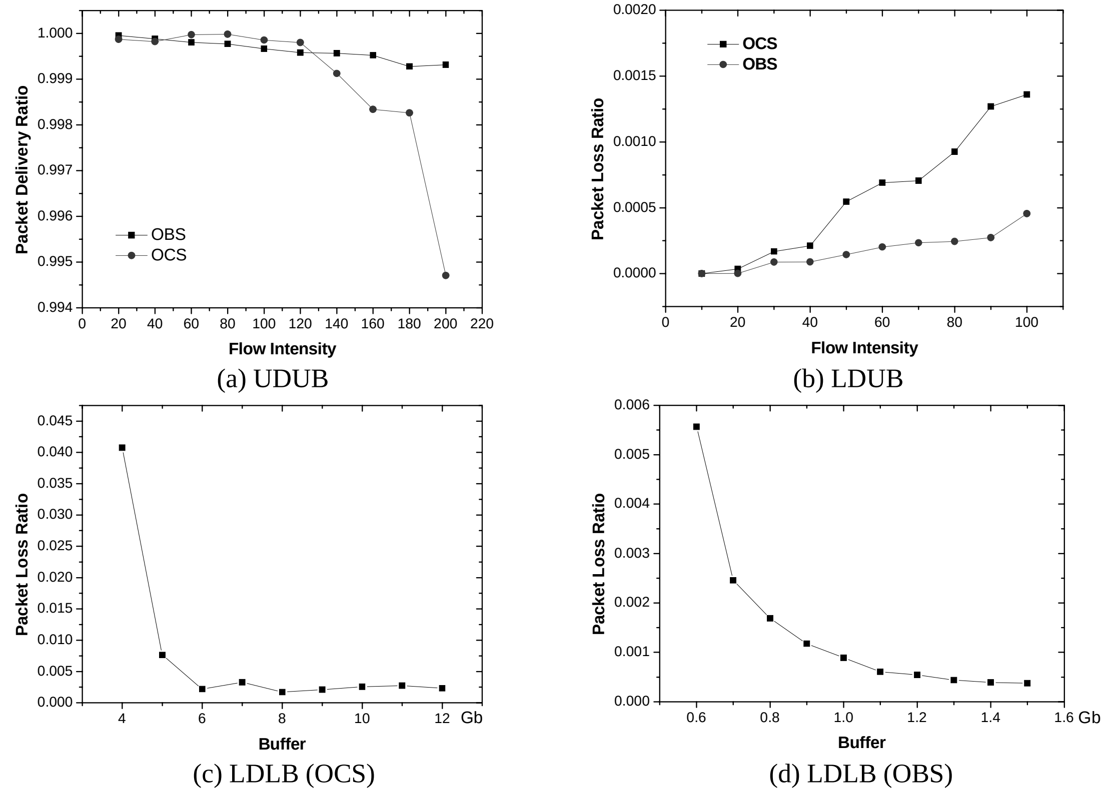

(a) UDUB (b) LDUB

(c) LDLB (OCS) (d) LDLB (OBS)

Fig. 1. The performance of OBS and OCS networks in different cases.

5. Conclusion

We have addressed the important question of whether OBS can be better than OCS in certain cases. The framework we have developed is useful in that it allows for fair performance comparisons. Our simulation codes are available for download at http://www.cse.buffalo.edu/~qiao/wobs/OBS_OCS.htm and we expect more comprehensive studies to follow.

Acknowledgement: This research is supported in part by NSF grant 0312563

6. References

1. C. Xin and C. Qiao, "A comparative study of OBS and OFS," in Proceeding of OFC, **4**, pp. ThG71-ThG73 (2001).

2. Fei Xue et al., "Performance comparison of optical burst and circuit switched networks," in OFC, **3**, pp.47-49 (2005)

3. R. Parthiban et al., "Does optical burst switching have a role in the core network?" in OFC, **3**, pp.50-52 (2005).

4. V. Paxson, "Fast Approximation of Self-Similar Network Traffic," Lawrence Berkeley National Lab, Tech. Rep. LRL-36 750, (1995)