

Cost-Based Burst Dropping Strategy in Optical Burst Switching Networks

Bartłomiej Klusek, John Murphy*, Liam Barry

Dublin City University, Dublin 9, Ireland

** University College Dublin, Dublin 4, Ireland*

Tel: +353-1-7005029, e-mail: klusekb@eeng.dcu.ie

ABSTRACT

Optical Burst Switching (OBS) is a new paradigm for future all-optical networks. Intentional burst dropping is one of techniques used to achieve desired quality of service. In this paper we note that some bursts are more likely to cause contention. We propose a cost function that can be used to predict the likelihood that a given burst will interfere with other traffic, then we explain how, by using this information a new burst dropping strategy can be designed. We compare our method with a random burst dropping technique and show that the cost-based approach offers a significant performance improvement.

Keywords: Optical Burst Switching (OBS), burst dropping, quality of service.

1. INTRODUCTION

It is widely assumed that all-optical networks will form the backbone of future Internet. Optical Burst Switching (OBS) [1, 2] is one of the most promising optical switching technologies, combining the best of both optical circuit and packet switching. One of challenges in designing OBS networks is providing adequate Quality of Service (QoS) for certain Internet services. For example, Internet telephony requires relatively low packet loss ratio in order to assure connection quality that will be acceptable to the end user. On the other hand, bulk data transfers, like File Transfer Protocol (FTP) can be handled by best-effort traffic.

Several methods of providing QoS in OBS networks have been proposed. For example, Yoo et al.[3] described offset-based QoS scheme. It relies on the fact that the probability of a burst being accepted depends on its offset time. By manipulating offset time depending on the traffic class of a particular burst, it is possible to achieve near total class separation and significant difference in loss probabilities for different traffic classes. Several aspects of this approach are also discussed in [4].

Another class of QoS techniques is based on intentional burst dropping. There, if a particular traffic class starts to consume more network resources than it is allowed to, then a certain proportion of its bursts is randomly dropped before it starts to interfere with other traffic. Chen et al. [5] proposed proportional QoS scheme that utilizes this method.

Random dropping of bursts, however, may not be the optimal solution. While simple and easy to implement, it will drop both bursts that almost certainly will not interfere with other bursts, and those very likely to cause a contention. It would be beneficial to differentiate incoming bursts and only drop the most problematic ones.

In [7] we presented the idea of a cost function as a way to assess the suitability of a given channel for a given burst. In this paper we extend this idea. We introduce the idea of total cost function, and use it to estimate the likelihood of a given burst interfering with other traffic. This information can then be used to only drop the worst part of the traffic in a particular class. We show that in a scenario, where an identical amount of traffic is admitted using either random or cost-based strategy, the burst loss probability for other traffic is much lower in the latter case.

The rest of the paper is organized as follows. First we present a more formalized description of the burst-dropping problem. Then we describe in detail the idea of cost function and present an example of such a function. To test the performance of our method we performed several simulations. The simulation setup and results are described in Section 4. Finally, we present our conclusions.

2. BURST DROPPING

Let us consider the following scenario. There are two traffic classes, the high priority one and the low priority or best-effort one. The maximum load offered in the high priority class is known, and there is maximum burst loss ratio that is guaranteed for this class. To achieve this objective some of the bursts belonging to the best-effort traffic may be intentionally dropped. This is done by a burst dropper. The structure of a burst dropper is shown in Fig. 1.

When a control packet arrives with a reservation request, the dropper calculates the value of a certain function

$f()$ and compares it with a pre-set threshold th . If the value of $f()$ exceeds th , the burst is dropped. Otherwise an attempt to forward the burst to the next node will be made.

Burst dropper can be:

- Random, when the value of function $f()$ is obtained from a random number generator,
- Deterministic, if for example $f(SEQ) = SEQ \bmod 10$, where SEQ - sequence number of incoming burst,

- Cost-based, if $f()$ is represented by a cost function.

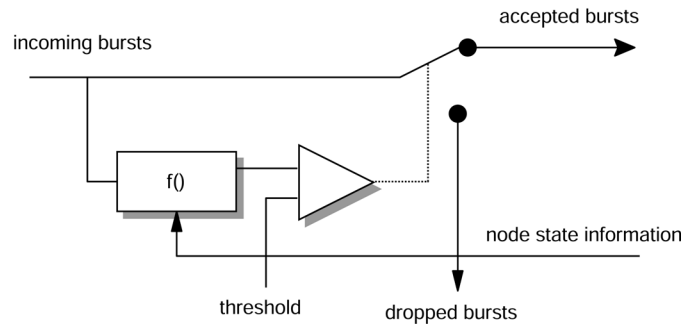


Figure 1. Structure of a burst dropper.

The value of a cost function gives the burst dropper certain knowledge. In particular the burst dropper learns if the burst would be dropped within the node anyway, due to congestion. Because of that, it is difficult to directly compare the proportion of traffic dropped by different types of burst droppers, and it is better to consider the amount of traffic that passes each dropper and is successfully forwarded to the next node.

3. COST FUNCTIONS

The value of a cost function is expected to reflect the probability that allocating a particular burst in a particular way will interfere with future bursts. We make a distinction between a channel cost function and total cost function. The former is used to test the suitability of a given channel for the new burst. The latter returns the cost of a channel that would be actually used to handle the burst, and may include the cost of using other resources, if an advanced cost-based channel algorithm is used.

3.1 Channel cost functions

The state of a channel can be described by a set of reservations. Each reservation consists of its starting time and ending time. If we denote the starting time of the i -th reservation by S_i and its ending time by E_i then the state of the entire channel can be represented by vectors $\bar{S} = [S_1, S_2, \dots, S_n]$ and $\bar{E} = [E_1, E_2, \dots, E_n]$ where n is the number of reservations. The newly arriving burst is described by its starting and ending times, S and E , respectively. We define the cost function in relation to the channel state, new burst and possibly other variables. It indicates the probability that the new burst will interfere with the allocation of future bursts:

$$C = f(\bar{S}, \bar{E}, S, E, \dots)$$

The value of the cost function is called channel cost and in the general case the function will be applied to all available channels and the one with the lowest cost will be chosen. Some of the existing channel allocation algorithms can be described using the idea of cost functions.

For example in case of Last Available Unused Channel with Void Filling (LAUC-VF) or Min-EV[6] the value of the cost function is simply the size of the starting or ending void created by allocating the burst, respectively.

It is difficult to estimate the probability of interfering with future bursts based on the channel state alone. We will, therefore, include other information to achieve better results. The general form of our cost function is: $C = f(\bar{S}, \bar{E}, S, E, OT_{\min}, OT_{\max})$ where OT_{\min} and OT_{\max} are the minimum and maximum offset times expected at the node, respectively.

The range of possible offset times is important, because relative weights of the starting and ending voids change as a function of the burst's offset time. For example, the size of the starting void does not matter if the burst arrived with the smallest possible offset time. No other burst will be allocated in this gap anyway. Our cost function first calculates the channel price according to LAUC-VF and Min-EV, then applies appropriate weights to those results and chooses the smaller value:

$$C = L * \min\left(\frac{C_{LAUC-VF}}{OT - OT_{\min}}, \frac{C_{MinEV}}{OT_{\max} - OT}\right)$$

where L – link load, as a proportion of maximum link load.

3.2 Total cost functions

When a cost-based channel allocation algorithm [6] is used, the total cost function returns the minimum cost of handling a particular burst. In the general case this may include channel cost, FDL cost, wavelength converter cost and possibly link cost. In this paper we will consider a node with full wavelength conversion capabilities and no FDLs, so the total cost function will return the minimum of all channel cost function values.

4. SIMULATIONS

4.1 Traffic sources

We have developed an OBS extension to the ns-2 simulator, and used it to obtain the results presented in this paper. In our simulations, we used burst sources that generate bursts based on previously measured burst length and inter-arrival time distributions. To obtain the necessary data, we simulated 150 Pareto-distributed IP packet sources, with their output being fed into a burst assembler. All time parameters in this paper are defined as multiples of the maximum assembly time. Its actual value is not important, as only the proportions matter.

Note that when a used burst source generates burst length and inter-arrival time independently, the self-similarity of the original traffic is lost. In our simulations this effect proved to be beneficial, as the simulation times did not need to be excessively long.

4.2. Simulation setup

The GT-ITN tool [8] was used to randomly generate four different network topologies of the transit-stub type, later referred to as network 0, 1, 2 and 3. The generator parameters were as follows: one transit domain, three transit nodes and three stub nodes per transit node. Identical simulations were run for each network topology.

In each node, there is a burst dropper. It discards a certain proportion of low-priority traffic either in a random fashion or, when the cost-based algorithm is used, based on the cost of allocating a certain burst. This means that a burst has to traverse a certain number of burst droppers, depending on its destination.

The offset time was uniformly distributed over a range of 15 to 25 maximum burst lengths. The processing delay at each node was equal to 0.5 of maximum burst length. The total network load was equal to 377% of maximum link load, and 25% of it belonged to the high – priority class. There were 40 channels per link.

For each simulation the main parameter was the dropper threshold. In case of a random dropper it was simply the proportion of bursts to be dropped, and in case of a cost-based one it was the maximum acceptable cost of allocating a burst. Then we calculated the average burst loss probability in each class, based on the overall sent and received traffic. Those probabilities were then plotted against each other.

4.3. Simulation results and analysis

The results of our simulations are presented in Fig. 2-5. All the graphs show parametric functions, with either the percentage of dropped bursts or the maximum acceptable price as a parameter, and a pair of loss probabilities as a result. They can be interpreted as a trade off between loss probabilities in both traffic classes. It also can be noted that the cost - based algorithm is never worse than the random one and, in some cases, significantly outperforms it.

The results for both algorithms tend to converge for the lowest and the highest loss probabilities, but significant differences can be seen between those two extremes. It was to be expected that results will be identical when no bursts are dropped, regardless which algorithm makes the decision to drop no bursts at all. Then, if the loss probability in the high priority class is to be lowered, some bursts in the best-effort class have to be dropped. The random algorithm does not discriminate between bursts, and each burst is equally likely to be dropped.

The cost-based algorithm, on the other hand, analyzes existing reservations and uses the cost function to identify the bursts that are most likely to cause contention. Those bursts are then dropped. This results in a lower burst loss probability in the high priority class for a given proportion of dropped best-effort traffic. But, as more and more bursts are intentionally dropped, the difference between algorithms decreases again. Apparently, there is a relatively small proportion of the "worst" bursts, and after those have been dropped, it becomes difficult to achieve further gains.




Figure 2. Loss comparison of Best Effort vs. High Priority, network 0.




Figure 3. Loss comparison of Best Effort vs. High Priority, network 1.




Fig. 4. Loss comparison of Best Effort vs. High Priority, network 2.

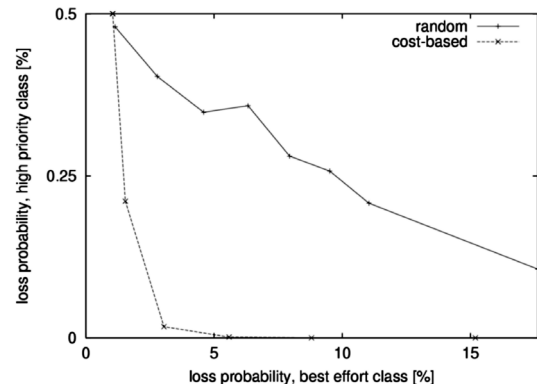


Fig. 5. Loss comparison of Best Effort vs. High Priority, network 3.

The above observations are true for all four network topologies. There are the inevitable differences in loss ratios, but the performance of the cost-based burst dropping algorithm remains consistently better.

Our results suggest that the cost-based burst dropping can be used in two cases. Firstly to ensure a given level of loss probability for the high priority class. This is likely to be done as a part of a QoS scheme. Then, using our algorithms allows us to drop less of the best-effort traffic than it would be necessary if the random algorithm was used. It has to be noted, though, that to achieve a significant change might require dropping relatively high amounts of best effort traffic. Still, that amount will be lower than it would have to be with the random algorithm.

The other use would be enforcing Service Level Agreements. If any traffic class exceeds its allocated bandwidth, some of its bursts will have to be dropped. By using a cost-based algorithm it is possible to identify and drop the bursts that are most likely to cause contention. This allows us to achieve lower loss probability for all other traffic classes, especially when relatively large amounts of non-compliant traffic have to be dropped.

5. CONCLUSIONS

In this paper we have presented a novel intentional burst dropping technique that can be used to implement QoS in OBS networks. We noted that, instead of dropping bursts in a random fashion, it is possible to find and drop only those bursts that are most likely to interfere with other traffic. A cost function was used to estimate which bursts should be dropped.

In our simulations no attempt was made to drop an exact percentage of incoming bursts. Instead, the dropper threshold was set to arbitrary values, and burst loss probabilities were measured during the simulation. This would not be acceptable in a real network. In practice it would be necessary to either find an analytical formula to calculate an appropriate threshold, or implement an adaptive algorithm that would automatically adjust the threshold value.

ACKNOWLEDGEMENTS

The support of both Science Foundation Ireland and the Informatics Research Initiative of Enterprise Ireland is gratefully acknowledged.

REFERENCES

- [1] M. Yoo, M. Jeong and C. Qiao, "A high speed protocol for bursty traffic in optical networks," in *All-Optical Communication Systems: Architecture, Control and Protocol Issues, Proc. SPIE 3230*, 79-90 (1997).
- [2] C. Qiao and M. Yoo, "Optical Burst Switching (OBS) - a new paradigm for an optical internet," *Journal of High Speed Networks*, vol. 8, pp. 69-84, Jan. 1999.
- [3] M. Yoo and C. Qiao, "A new optical burst switching protocol for supporting quality of service," in *Proc. SPIE '98 Conf. All Optical Comm. Syst.: Architecture, Control, Network Issues*, vol. 3531, Boston, 1998.
- [4] W.H. So and Y.C. Kim, "Offset time decision for supporting service differentiation in optical burst switched networks," in *Proc. COIN-PS 2002, Korea, 2002*.
- [5] Y. Chen, M. Hamdi and D.H.K. Tsang "Proportional QoS over OBS networks," in *Proc. GLOBECOM 2001*, vol. 3, pp. 1510-1514.
- [6] Jinhui Xu, C. Qiao, J. Li, and G. Xu, "Efficient channel scheduling algorithms in optical burst switched networks," in *Proceedings of IEEE INFOCOM (Mar. 2003)*, 3, 2268-2278.
- [7] B. Klusek, J. Murphy and P. Perry, "Cost-based wavelength allocation algorithms in Optical Burst Switching networks," in *Proc. APOC 2004, Beijing, China 2004*.
- [8] E. W. Zegura, K. Calvert and S. Bhattacharjee, "How to Model an Internetwork", *Proceedings of IEEE Infocom '96, San Francisco, CA*.