Techniques for improved scheduling in optical burst switched networks

C. Papazoglou, P. G. Sarigiannidis, Member, IEEE, G. I. Papadimitriou, Senior Member, IEEE,
A. S. Pomportsis
Department of Informatics
Aristotle University
Thessaloniki, Greece
{chrispa, sarpan, gp, apombo}@csd.auth.gr

***Abstract*—Optical burst switching (OBS) has emerged as a viable switching alternative in backbone optical networks since it can support high data rates with an intermediate granularity compared to wavelength routing and optical packet switching. At the edges of an OBS cloud, packets are assembled to form bursts which enter the network core and are switched on the fly using bandwidth previously reserved by their control packets at each node. A key problem in OBS networks is the assignment of wavelengths to incoming bursts, i.e. the scheduling of bursts. This paper proposes two new techniques which are shown to improve burst scheduling algorithms by lowering their complexity. The first proposed technique is based on a triangular estimator that defines a "drop zone"; bursts that fall into this area are considered to have a very low probability of finding a suitable wavelength and as such, no effort is made to schedule them. According to the second approach, the drop zone is defined dynamically based on the burst drop history. Simulation results show that both approaches yield burst drop rates marginally higher or identical to the LAUC-VF scheduling algorithm while reducing the number of channel or void checks and thus the algorithm complexity and execution time.**

Keywords-burst scheduling; optical burst switching; Horizon; Latest Available Unused Channel with Void Filling, bursty traffic, triangular estimator, drop history

I. INTRODUCTION

Optical burst switching (OBS) [1-12] combines the advantages of optical circuit switching (i.e. wavelength routing) and optical packet switching. OBS facilitates the transport of data units with variable sizes, called bursts over dense wavelength division multiplexed (DWDM) links without the need for path set up as in wavelength routing and without optical buffers to support the store and forward paradigm as in optical packet switching (OPS). This means that OBS networks can be implemented and deployed in a shorter term and with less sophisticated technology compared to optical packet switched networks.

The switching granularity in an OBS network lies between OPS and wavelength routing since the switching unit is larger than a packet but smaller than a wavelength. In OBS, the control information and data are transmitted separately both in time (offset) and space (separate control channel(s)). At the edges of an optical burst switching cloud (Figure 1), packets are assembled into bursts which are then assigned a wavelength and sent in the OBS network core.

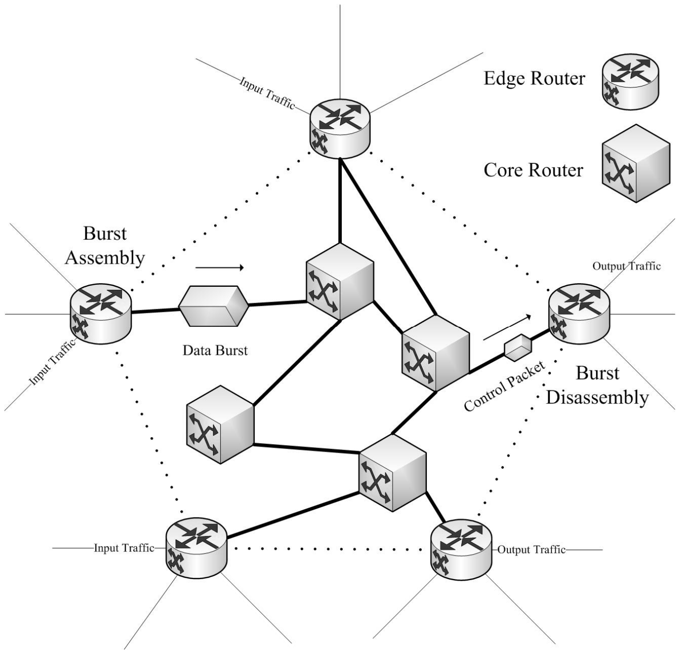

Figure 1. A general architecture of an OBS network

Each burst is preceded by a header (or control packet) which includes all information necessary for the reservation of resources (available bandwidth) at all intermediate nodes.

At each node in the burst's path, the control packet is converted into electronic form and an attempt is made to locate and reserve a wavelength that can accommodate the burst. Bursts do not wait for acknowledgements of successful bandwidth reservation but instead are transmitted shortly after their control packets. If the control packet fails to reserve the required resources the corresponding burst will be dropped. Bursts are forwarded on a hop by hop basis until they reach their destination edge router where they are disassembled into the original packets. Because the switching fabric has already been setup prior to their arrival, bursts are switched all-optically.

A significant issue in OBS networks is the scheduling of bursts at each node they reach, i.e. the assignment of bursts to wavelengths in the desired output fiber. Previously proposed scheduling algorithms differ depending on whether or not they utilize void intervals generated by bursts scheduled in the past.

The Horizon scheduling algorithm [1, 5] does not track void intervals on wavelengths and as a result is simple to implement but makes suboptimal scheduling decisions and wastes usable bandwidth. On the contrary, the Latest Available Unused Channel with Void Filling (LAUC-VF) [6] algorithm can schedule a burst using a void interval and thus achieves the lowest possible burst drop ratio but at a cost of very high complexity. This paper proposes two new approaches at OBS scheduling which attempt to combine the advantages of both Horizon and LAUC-VF; the first approach is based on a triangular estimator which is a tool that the OBS node uses to "predict" whether an incoming burst will find an available wavelength. Estimations are based on burst features namely the offset time and length. If the outcome of the estimator is negative (which means that the probability of finding a suitable wavelength is deemed low), the control packet is not forwarded any further and the corresponding burst is dropped upon its arrival. According to the second approach, the decisions on bursts that will not be scheduled are based on two criteria: the burst length and the offset from its control packet. The proposed algorithm uses a history table to compute the burst drop probability for each class of bursts where each class is identified by its length and offset. Both approaches are shown to reduce the overall scheduling complexity in terms of the number of search operations.

This paper is organized as follows: section II presents an overview of previously proposed burst scheduling algorithms; section III presents the proposed schemes while section IV includes the numerical results which illustrate their efficiency. Finally, section V concludes the paper.

II. BACKGROUND

Channel scheduling algorithms can broadly be classified in two categories: with and without void filling. Void intervals are created in data channels as bursts are scheduled due to the fact that there is an offset time between the arrival of the control packet and the actual arrival of the burst. As their names suggest, algorithms without void filling can not schedule a burst in a void interval while void filling algorithms can. These two different approaches result in major differences in channel utilization and loss rate as well as simplicity of implementation and speed of execution [4].

The most typical example of a non void-filling scheduling algorithm is the Horizon algorithm [1, 5-6]. This algorithm maintains a single value for each data channel called the scheduling horizon. A channel's horizon is defined as the latest time at which the channel is currently scheduled to be in use [1]. Horizon can only schedule bursts if their arrival times are greater than a channel's Horizon. The operation of the algorithm is depicted in Figure 2. Although there are three channels which could potentially accommodate the burst (w1, w2 and w3) the Horizon algorithm only view two feasible wavelengths (w1 and w3) and selects w3 because it's Horizon is closer to the arrival of the burst. Simplicity in both operation and implementation is the main advantage of the Horizon algorithm [8]. However, this type of scheduling results in low bandwidth utilization and high loss rate due to the waste of channel resources.

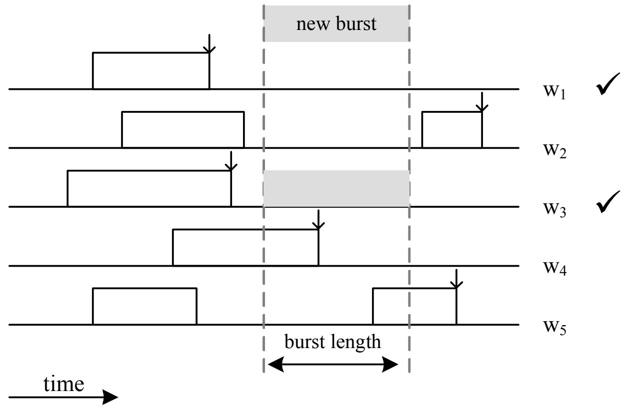

Figure 2. The Horizon burst scheduling algorithm

The most typical example of a void-filling burst scheduling algorithm is the Latest Available Unused Channel with Void Filling algorithm (LAUC-VF) [6]. The main idea is to utilize void intervals for scheduling and select the latest available unused data channel for each arriving burst. Figure 3 illustrates the operation of the scheduling algorithm. LAUC-VF identifies three suitable data channels for the new burst but schedules it on the one which minimizes the void between the end of the previously scheduled burst and the start of the new burst, i.e., the one which generates the minimum starting void. Obviously, the performance of the LAUC-VF algorithm in terms of bandwidth utilization and burst loss rate is superior to that of the Horizon algorithm. This, however, comes at a cost of a significantly higher complexity and increased memory requirements.

III. THE PROPOSED SCHEMES

A. The triangular estimator scheme

This scheme aims at reducing the scheduling complexity by avoiding unnecessary channel searches, i.e. searches that if performed will prove to be futile. The main motivation for this scheme was the observation that searches for suitable channels for an incoming burst are often unsuccessful, especially when the size of the burst is large and/or the offset between the burst and its control packet is small. If the burst scheduler can identify and drop such bursts upon their arrival instead of attempting to schedule them, the total number of channel searches (checks) will be reduced and significant savings in both processing time and memory operations will be achieved.

By studying the variance of the drop ratio in an OBS node in relation to the burst characteristics we concluded that there are three critical zones where the burst drop probability is unacceptably high. These are identified by the following values:

a) A burst offset smaller than 30% of the maximum offset value combined with a burst length greater than 90% of the burst length range (difference between the maximum and the minimum burst length).

b) A burst offset smaller than 20% of the maximum offset value combined with a burst length greater than 80% of the burst length range.

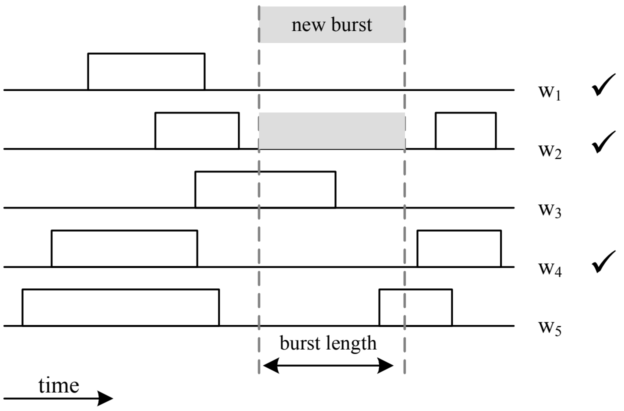

Figure 3. The LAUC-VF scheduling algorithm

c) A burst offset smaller than 10% of the maximum offset value combined with a burst length greater than 70% of the burst length range.

These three zones define the burst drop zone while the remaining area is termed as the “pass” zone. No effort will be made to schedule bursts belonging in the drop zone but bursts belonging in the pass zone will be scheduled using the algorithm that the node implements. Burst zones are depicted in Figure 4 where the darker shade represents the triangular drop zone and the lighter shade represents the pass zone. The TRiangular ESTimator (TR-EST) scheduling algorithm works as follows: upon the reception of a control packet the burst length and offset are extracted and examined and the burst’s zone is determined. If the burst falls into the pass zone it is forwarded to the node scheduler and processed as usual. Otherwise the burst is dropped. It must be noted that the TR-EST scheduling algorithm does not incur any additional processing overhead since the only extra operation is a logical test. Therefore the complexity of the scheduling algorithm is not affected.

B. *The Drop History scheme*

Although the TR-EST scheme successful identifies and discards voids that will not be able to find an available wavelength in scenarios where the network load is high, its performance and accuracy degrade when the network load is low (e.g. when the number of channels is increased) because the drop zone is fixed and is not adjusted dynamically according to traffic and/or network characteristics. For these reasons, we proposed an additional scheme in which the drop zone is not known from the beginning but instead is “constructed” dynamically. The Drop History (DH) approach works as follows:

During a learning phase the algorithm schedules bursts exactly like LAUC-VF and records the number of bursts dropped and the total number of bursts for each class (we worked with 10 different values of burst length and 10 values of burst offset = 100 classes). After enough data has been gathered in the history table, the algorithm proceeds to its main phase of operation. Before a burst is scheduled, the drop history table is checked; if the burst drop ratio for previously scheduled bursts belonging in the same class as the new one exceeds a given burst drop threshold, the corresponding burst is immediately dropped and no search for a wavelength channel is performed. Otherwise, the burst is scheduled using LAUC-VF.

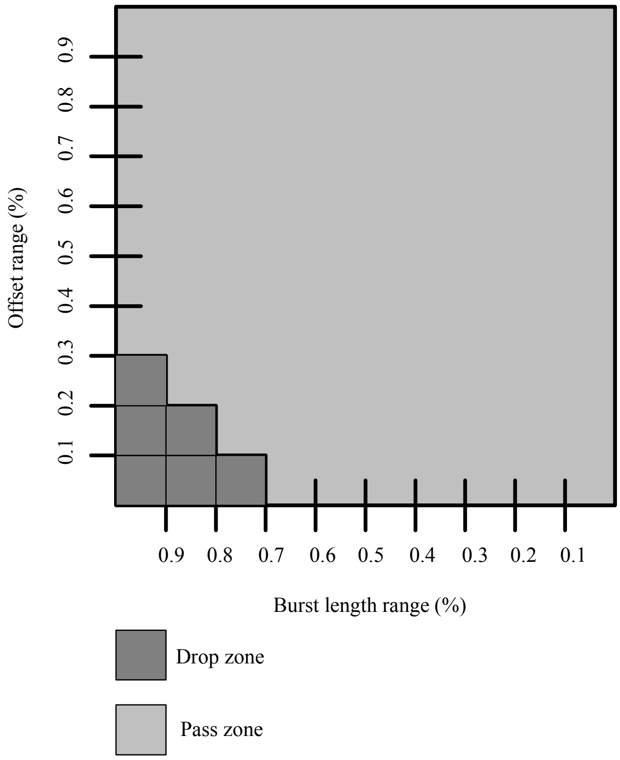

Figure 4. Pass and drop zones for the triangular estimator

IV. Simulation Results

In order to study the performance of the proposed schemes, we simulated all scheduling algorithms (Horizon or LAUC and LAUC-VF) and compared the results to TR-EST and DH in a variety of simulation scenarios. Our performance metrics were the percentage of dropped bursts (burst drop ratio) and the number of channel or void checks that had to be performed (for TR-EST and DH respectively).

In our simulations, we considered a single output fiber in an OBS node (as in [11, 12]). The line rate was assumed to be 1Gb/s. The load was fixed and set to 50%. The input traffic was modeled using a Pareto distribution for both the control packet interarrival times and the burst lengths (shape during the ON states = 1.6 and during the OFF states = 1.3). The distribution of the offset times was assumed to be uniform.

Figure 5 plots the burst drop ratio vs. the maximum offset time for Horizon (also referred to as LAUC), LAUC-VF and TR-EST. In this set of experiments we considered different values for the maximum offset time (between 100 and 500ms). The number of wavelength channels was fixed and equal to 10 for all simulations and the burst length range was 1024 – 4096 Bytes. It is evident that the drop ratio of the TR-EST algorithm closely matches that of the LAUC-VF algorithm which means that TR-EST achieves a drop ratio marginally higher than the

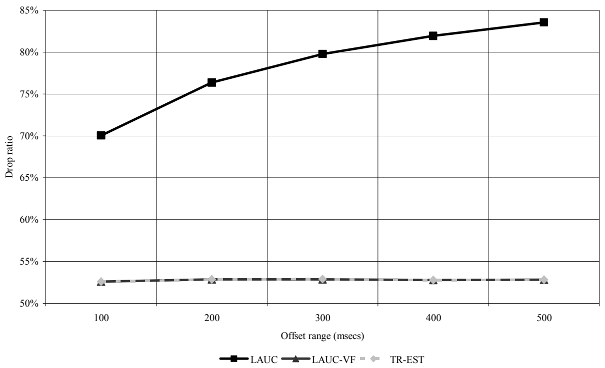

Figure 5. Burst drop ratio vs. maximum offset time

minimum that can be achieved. Figure 6 plots the number of channel checks vs. the offset time for the same three scheduling schemes. From the figure we can see that the number of checks for the TR-EST algorithm is larger than that of the Horizon algorithm but significantly smaller than that of the LAUC-VF algorithm. This gain in channel checks and therefore scheduling complexity stems from the fact that because several bursts are dropped upon their arrival the total number of search operations is reduced. Figures 6 and 7 plot the drop ratio and the number of void checks versus the burst length range for the Horizon, LAUC-VF and DH algorithms. DH is plotted for two values of the burst drop threshold (90% and 95%). This means that for these schemes bursts belonging to classes whose drop probability exceeds 90% or 95% will be immediately dropped. Both figures confirm that the DH scheme decreases the number of void checks (and hence the scheduling complexity) while maintaining an almost identical burst drop ratio.

V. CONCLUSIONS

This paper presented an accurate and efficient triangle estimator for OBS networks and an improved scheduler based on a drop history table. The estimator uses a triangular formula based on the burst offset and length in order to assess the probability of successful resource reservation for an incoming burst and discards the burst if the probability is too low. The drop history based algorithm classifies bursts according to their lengths and offsets, dynamically computes

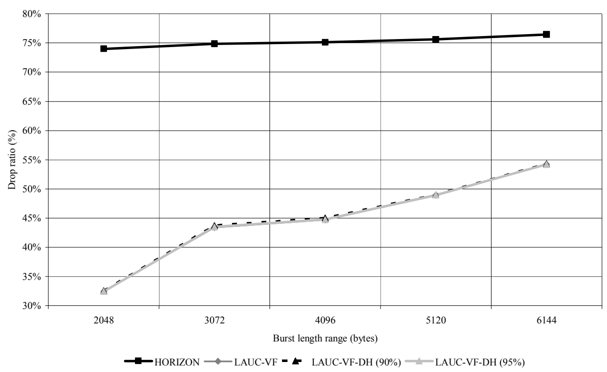

Figure 6. Burst drop ratio vs. burst length range

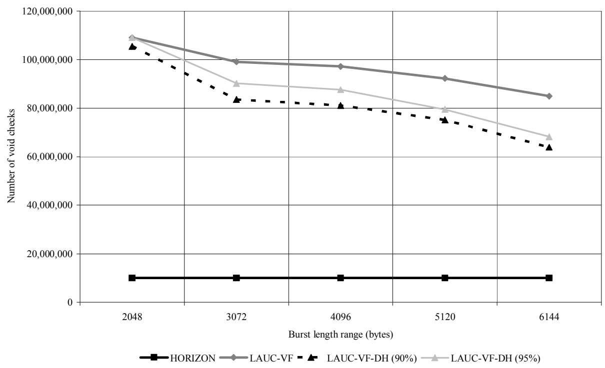

Figure 7. Number of void checks vs. burst length range

drop probabilities for each class and drops burst whose present drop probability exceeds a given threshold. Simulation results for an OBS node with self-similar traffic comparing the proposed schemes with the Horizon and the LAUC-VF scheduling algorithms clearly indicate that they have a similar performance while they reduce the number of channel or void checks and thus the overall scheduling complexity.

REFERENCES

[1] J. S. Turner, “Terabit burst switching”, Journal of High Speed Networks, 8, pp. 3-16, 1999.

[2] C. Qiao and M. Yoo, “Optical burst switching (OBS) - A new paradigm for an optical internet,” Journal of High Speed Networks, 8, pp. 69-84, 1999.

[3] C. Qiao and M. Yoo, “Choices, Features and Issues in Optical Burst Switching”, Optical Networks Magazine, 1, pp. 36-44, 2000.

[4] G.I. Papadimitriou, C. Papazoglou, and A.S. Pomportsis, “Optical switching”, Wiley Series in Microwave and Optical Engineering, Wiley – Interscience, 2006.

[5] Y. Chen and J.S. Turner, “WDM Burst Switching for Petabit Capacity Routers”, IEEE Military Communications Conference Proceedings, 1999 (MILCOM 1999), vol. 2, pp. 968-973, 1999.

[6] Y. Xiong, M. Vandenhoute and H.C. Cankaya, “Control Architecture in Optical Burst-Switched WDM Networks”, IEEE Journal on Selected Areas in Communications, 18, pp. 1838-1851, 2000.

[7] T. Battestilli and H. Perros, “An introduction to optical burst switching”, IEEE Communications Magazine, 41, pp. S10-S15, 2003.

[8] Y. Chen, C. Qiao, and X. Yu, “Optical Burst Switching: A New Area in Optical Networking Research”, IEEE Network, 18, pp. 16-23, 2004.

[9] H.M.H. Shalaby, A Simplified Performance Analysis of Optical Burst-Switched Networks, Contour-Based Priority Scheduling in Optical Burst Switched Networks, IEEE/OSA Journal of Lightwave Technology, 25, pp. 986-995, 2007.

[10] J.P. Jue and V.M. Vokkarane, “Optical Burst Switched Networks”, Optical Networks Series, Springer, 2005. ISBN: 0-387-23756-9

[11] M. Yoo, M. Jeong, and C. Qiao, “A High Speed Protocol for Bursty Traffic in Optical Networks”, Proceedings of the 3rd SPIE Conference on All-Optical Communication Systems, pp. 79-90, 1997.

[12] S. Verma, H. Chaskar, and R. Ravikanth, “Optical Burst Switching: A Viable Solution for Terabit IP Backbone”, IEEE Network, 14, pp. 48-53, 2000.