Service differentiation mechanisms for variable length packets in an optical switch with recirculating FDL buffer

C. Develder[1], J. Cheyns[1], M. Pickavet[1], P. Demeester[1]

1: Dept. of Information Technology (INTEC), Ghent University – IMEC, Sint-Pietersnieuwstraat 41, 9000 Gent, Belgium

Abstract: In this paper we focus on providing service differentiation for variable length packets. Apart from the well-known OBS approach using differentiated offsets to introduce multiple service classes, we consider a look-ahead approach allowing for later arriving high priority packets to pre-empt earlier arrived low priority packets, and a slotted control approach. All approaches attain service differentiation without any resource reservation, and are of limited complexity, to minimize packet processing requirements. Through simulation, we assess the quality of the approaches under varying loads, buffer dimensions and QoS algorithm parameters. The main criterion used is the packet loss rate per service class.

Keywords: WDM, Optical Packet Switching, Optical Burst Switching, service differentiation, FDL, simulation.

1. Introduction

The answer to the ever lasting hunger for bandwidth is being met by the deployment of (D)WDM networking [1]. To ensure efficient dealing with variable traffic patterns (both geographically and over time), Optical Packet Switching (OPS) has been devised. Profiting from cutting edge technology, it exploits fast optical switching techniques to offer better bandwidth granularity, efficiency and flexibility than circuit-switched approaches. To relax some of the challenges involved (e.g. synchronization, high processing overhead), the asynchronous Optical Burst Switching concept using longer, variable length data units, has been devised. In this paper, we aim at finding a suitable technique to provide service differentiation for optical switches dealing with variable length packets.

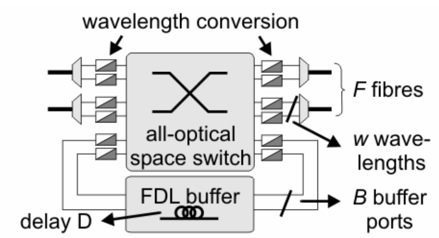

Figure 1 : Switch architecture under study.

The switch architecture we will focus on is a very generic one, and has been proposed e.g. in the European research project DAVID [2] for slotted OPS. Its functional architecture is sketched in Fig. 1: It has F input/output fibres, each carrying W wavelengths in WDM. The switch is capable of wavelength conversion, and exploits this capacity to solve contention [3]: packets arriving simultaneously and destined for the same outgoing fibre may be converted to other wavelengths to allow concurrent forwarding on the same output fibre. Where this does not suffice, an optical buffer can be used: B ports of the switching matrix are connected to Fibre Delay Lines (FDLs). In this paper, we assume that all recirculating buffer ports have the same delay D.

We will continue the paper as follows: in the next Section 2, we describe and discuss the compared QoS approaches. The methodology used is outlined in Section 3, followed by the results in Section 4. Conclusions are summarized in Section 5.

2. Approaches to QoS

The approaches we will compare are illustrated in Fig. 2 and comprise the following:

1) Header offset differentiation ("offset"): This is OBS-JET where QoS differentiation is realized through giving higher priority packets a longer header offset [4]. This way, the arrival of high priority packets is known in the switch longer beforehand and reservations can be made before low priority packets.

2) Look-ahead ("look"): The offset for different priority classes is the same but service differentiation is attained by assuming that the switch controller needs to make a decision only H after it has received the packet header. This can be achieved by having a fixed input buffer (e.g. by extending the one accounting for packet header processing). Lower priority packets can be pre-empted by higher priority packets arriving up to a time H later.

3) Slotted control ("slot"): Again without offset differentiation, the switch controller operates in a slotted mode. Each time-slot of duration T, we jointly make a decision for packets whose headers have arrived since the previous timeslot (just as a slotted OPS switch would do [5]). Thus, we can give precedence to high priority packets that arrived up to T later than low priority ones. This can be realized through a synchronous control channel (since electronic header processing is assumed, this is realistic), or at least (electronic) synchronization in the switch controller and a sufficiently large input buffer (FDL) on the data path.

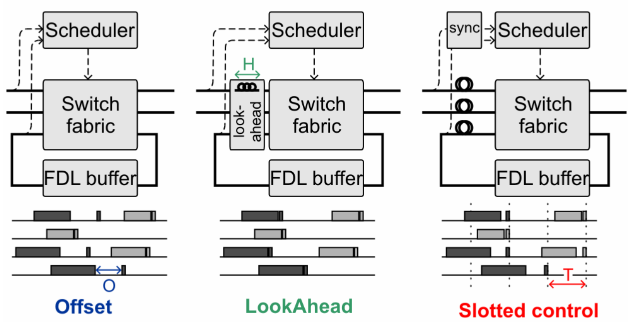

Figure 2 : The three QoS approaches we focus on.

Clearly, these approaches are not suitable for a large number of service classes: offsets, look-ahead times or slot sizes would need to be quite large and lead to unacceptable latency for high priority packets. However, since it is widely acknowledged that few (two or three) classes will be required in the core networks where switches as in Fig. 1 will be deployed, we believe the suggested approaches are valid candidates towards service differentiation.

Note that these are not the only possible approaches to providing QoS in an IP-over-WDM scenario [6]. The QoS methods analysed in the following are all based on explicit indication of the class of service (priority), which applies to the whole packet. It all are cases without segmentation [7], without a priori resource reservation (to maximize resource utilization), without intentional dropping (opposed to e.g. [8]) and without the need to revoke anything sent out on output fibres. These properties ensure that the packet scheduling algorithm's complexity is quite limited and thus restricts time- and resource consuming packet processing.

The scheduling algorithms used all follow the PostRes [9] approach, meaning that no reservations are made for buffered packets until they leave the FDL and re-enter the switch (cf. otherwise, precautions have to be taken to avoid that buffering interferes with the differentiation mechanism [9]). The scheduler thus performs the same task for newly arriving packets as for recirculated ones: (i) use LAUC-VF [10] to find an available wavelength channel on the output fiber the packet is destined for, (ii) if none is free, use LAUC to find the most suitable free FDL port, (iii) otherwise drop the packet.

3. Methodology

The parameters used for the node architecture of Fig. 1 are: F=6 i/o fibers, W=8 wavelengths per fiber, B=0..64 buffer ports. For the packet arrival process, we used Poisson arrivals and packet lengths based on a negative exponential distribution: packets have a minimal length of L/2 and mean length L (the length minus L/2 follows a negative exponential distribution). Traffic was uniformly spread over all output fibers. We considered two priority classes, where 60% of the packets were of the low priority and 40% high priority. In our simulations, we ignored header processing times and thus set the basic header offset to zero. The offset-times clearly (see Fig. 2) depend on the QoS approach taken (zero or O for offset differentiation; zero for look-ahead; in the range [0,T) for slotted control).

To assess the major differences in performance between the three analyzed QoS approaches, we focus mainly on the loss rate achieved, i.e. the fraction of packets that is lost (which should be considerably lower for high priority traffic). Since this traffic is composed of variable length packets, we are also interested in the "fairness" within a single priority class: do all packets belonging to the same priority class experience the same QoS? It is indeed a well-known fact that short packets usually have lower chances of being dropped [11].

4. Results

4.1 Influence of the number of buffer ports

In a first experiment we focused on the efficiency in exploiting the available buffer resources: we kept all parameters constant except the number of buffer ports B. Fig. 3 shows the evolution of the loss rate for an increasing number of recirculating buffer ports in case of a load of 0.8. The QoS parameters were the following: *(i) differentiated offsets*: high-priority offset O=2L; *(ii) look-ahead*: look-ahead time H=2L; *(iii) slotted control:* slot resolution T=2L. The buffer length was set to D=2L.

The *slot* approach achieves loss rates that are higher than the other two approaches, esp. when the number of recirculating buffer ports increases. Whereas the difference in overall loss rates is limited, the loss rate for high priority packets is multiple orders of magnitude bigger than for *offset* or *look-ahead*. Still, even such a simple mechanism is able to provide clear service differentiation. When comparing *offset* with *look-ahead*, the differences are small, with a slightly better performance of look-ahead.

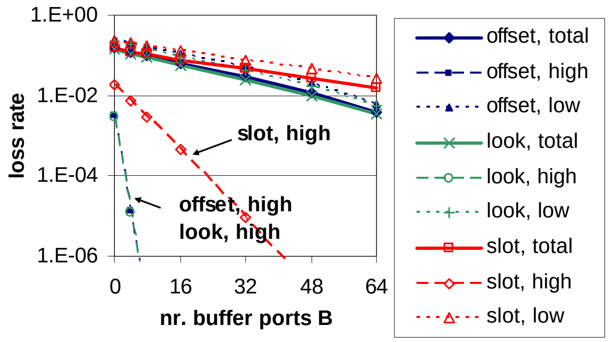

Figure 3 : Loss rates for increasing number of recirculating buffer ports B (load=0.8, F=6, W=8, 40% high priority traffic).

From a fairness perspective, results (not plotted because of space limitation) confirmed our earlier statement that short packets are subject to lower loss rates. For the look-ahead strategy, this consistently favoring of short packets is more pronounced, due to the preemptive nature of the look-ahead strategy. Packets are scheduled upon arrival, and successively arriving packets of the same (or lower) priority are scheduled taking into account this schedule. For look-ahead, this schedule may be changed when later on (less than H) a higher priority packet arrives destined for the same output fiber. Thus, the allocation of packets to fibers is not optimal, and the longer packets are the first to suffer from this effect.

4.2 Influence of the class offset

Class separation depends on the parameter setting of the various QoS approaches: the actual parameter differs for each of the proposed approaches, but we will refer to it by "class offset". For the differentiated offset approach, it is the difference O in header offset between two successive priority classes. For look-ahead, it is the look-ahead delay H. For the slotted approach, it is the slot resolution T.

To assess the influence of the "class offset", we carried out simulations for a load of 0.8, and a buffer with B=8 ports and FDL length D=4L, leading to loss rates plotted in Fig. 4(a). The "class offset" was varied from L/2 to 4L.

By increasing the class-offset, the overall loss rate rises: high priority packets are considered more important, and their loss rate drops accordingly, but at the price of disregarding more low priority packets. Clearly, there is a limit to the improvement: as soon as they don't "see" any low priority packets anymore, i.e. sufficient class isolation is achieved, the loss rate stabilizes. The point at which this isolation is achieved depends on the packet size distribution. For the negative exponential packet length distribution at hand, nearly complete isolation is reached for a "class offset" around 2L (less than 5% of the packets are longer than 2L for the distribution) for offset differentiation and somewhat earlier for look-ahead. This threshold is flattest for look-ahead. For the slotted approach, there seems to be much more room for improvement by increasing the "class offset" (thus slot size) further.

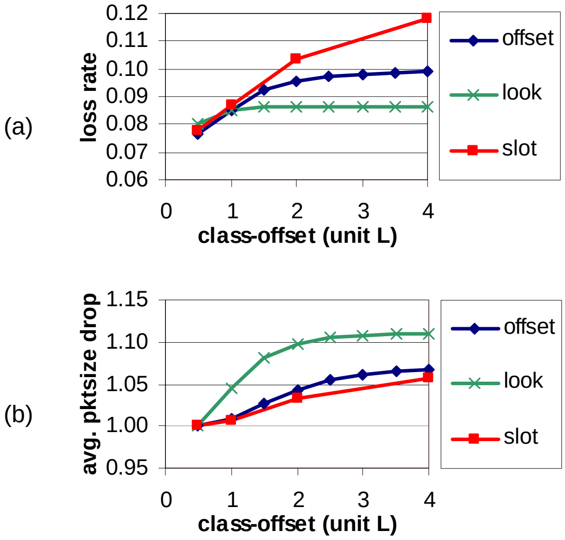

Figure 4 : (a) Overall loss rate and (b) avg. size of dropped packets for increasing class-offset O=H=T. (load=0.8, F=6, W=8, 40% high priority traffic; L=avg. packet length.).

Note that to limit latency, we should aim at restraining the "class offset": since this is related to the average packet length L, this implies that packet lengths should be limited (i.e. at least those of high priority packets).

To answer the question of intra-class fairness (how bad is the discrimination of long packets?) we plot in Fig. 4(b) the evolution of the average size of the packets dropped. With increasing "class offset", and thus class isolation, the unfairness rises (only packets longer than average are dropped), but it is far more pronounced for look-ahead.

4.3 Influence of buffer delay

By varying the delay realised on the recirculating path through the FDLs, we unsurprisingly [9] find lower loss rates for increasing delay D. However, as for the "class offset", there is a floor: when the buffer is large enough to contain almost all packets (around 2L, see above), we see no further improvement of the loss rates.

4.4 Influence of the load

To verify that the approaches achieve sufficient isolation under all loads, we have analyzed the evolution of the loss rates for loads ranging from 0.1 to 0.9. When comparing the loss rates, the difference between both classes decreases for increasing load for both the differentiated offset and the look-ahead approach. (For the bufferless case, it decreases from a factor ~700 to ~60). For the slotted control approach, the difference in loss rates is smaller, but the relative difference does not diminish that much with increasing loads. It is worth noting that even for the very simple slotted control mechanism, even the lowest priority packet loss rate stays below 10^{-3} for loads as high as 0.5.

5. Conclusion

We introduced and compared three scheduling approaches that attain service differentiation for variable length packets in an optical packet switch with a recirculating FDL buffer. We compared the well-known differentiated offsets approach with a look-ahead approach that proved to achieve comparable loss rates. A slotted control approach which could simplify the burst scheduler implementation achieves almost equal overall loss rates and delays, but does not achieve the same class separation. Still, for low to medium loads, with a moderate buffer, the performance attained by slotted control may be acceptable. From a fairness point of view, the look-ahead approach most severely discriminates against longer bursts.

The robustness of each of the service differentiation mechanisms was assessed by varying the key parameters influencing their class isolation: (i) number of buffer ports, (ii) class offset, (iii) buffer delay, and (iv) load. It was concluded that (i) all approaches greatly benefit from adding buffer space, but the slotted approach does not succeed in exploiting it as efficiently as the others; (ii) the class offset needs to be set according to the burst length distribution (somewhat smaller for look-ahead); (iii) performance is improved if recirculation delay is sufficiently large; (iv) class isolation tends to slightly decrease for increasing loads.

Acknowledgment

This work has been supported by the European Commission through the IST-project DAVID (IST-1999-11387), and by the Flemish Government through the IWT GBOU-project "Optical Networking and Node Architectures". C. Develder and J. Cheyns are supported as a Research Assistant of the Fund for Scientific Research – Flanders (F.W.O.–Vl.), Belgium.

References

[1] B. Mukherjee, "WDM optical communication networks: Progress and challenges", JSAC, vol. 18, no. 10, Oct. 2000.

[2] L. Dittmann, et al., "The European IST Project DAVID: a Viable Approach towards Optical Packet Switching", to appear in JSAC.

[3] S. Yao, et al., "All-optical packet-switched networks: A Study of Contention-Resolution Schemes in an Irregular Mesh Network with Variable-Sized Packets", Proc. OptiComm 2000.

[4] M. Yoo, C. Qiao, "Supporting multiple classes of services in IP over WDM networks", Proc. Globecom 1999.

[5] C. Develder, et al., "Strategies for an FDL based feed-back buffer for an optical packet switch with QoS differentiation", Proc. COIN 2002.

[6] A. Kaheel, et al., "Quality-of-service mechanisms in IP-over-WDM networks", IEEE Comm. Mag., vol. 40, no. 12, Dec. 2002.

[7] A. Detti, et al., "Performance evaluation of a new technique for IP support in a WDM optical network: opti-cal composite burst switching (OCBS)", J. Lightwave Tech., vol. 20, no. 2, Feb. 2002.

[8] Y. Chen, M. Hamdi, D.H.K. Tsang, "Proportional QoS over OBS networks", Proc. Globecom 2001.

[9] C. Gauger, "Dimensioning of FDL Buffers for Optical Burst Switching Nodes", Proc. ONDM 2002.

[10] Y. Xiong, et al., "Control architec-ture in optical burst-switched WDM networks", JSAC, vol. 18, no. 10, Oct. 2000

[11] K. Dolzer, et al., "Evaluation of reservation mechanisms for optical burst switching", AEÜ Int. J. Electron. Commun., vol. 55, no. 1, Jan. 2001.

Glossary

(D)WDM	(Dense) Wavelength Division Multiplexing
FDL	Fibre Delay Line
LAUC	Latest Available Unused Channel
LAUC-VF	Latest Available Unused Channel with Void Filling
OBS	Optical Burst Switching
OBS-JET	Optical Burst Switching with Just Enough Time
OPS	Optical Packet Switching
QoS	Quality of Service