The Impact of Dynamic Wavelength Assignment and Burst Aggregation in Optical Burst-Switched Networks

Ignacio de Miguel†, Michael Düser‡ and Polina Bayvel‡

† University College London (UK) and University of Valladolid (Spain), ‡ University College London (UK)

Abstract: In Wavelength-Routed Optical Burst Switched Networks (WROBS), lightpaths are dynamically established for the transmission of bursts (aggregation of packets in edge routers) through an all-optical core network. The control of a WROBS can be either centralised or distributed. In this paper, we focus on the centralised approach and analyse the timing requirements for dynamic routing and wavelength assignment (DRWA) algorithms to guarantee quality of service (QoS) for different burst aggregation techniques and network sizes. The results shown in the paper allow a network designer to determine whether a centralised control can be employed for a particular set-up.

1. Introduction.

A Wavelength-Routed Optical Burst Switched (WROBS) network architecture was recently proposed and analysed as a promising packet network architecture to provide guaranteed latencies [1-3]. It is based on the acknowledgement of the wavelength-assignment for QoS-determined bandwidth provisioning in combination with dynamic wavelength allocation. The architecture consists of optical core nodes connected to electronic edge routers, where the packets are buffered according to their destination and class of service to form a burst. Previous work has analysed the effect of traffic statistics on the edge delays, bandwidth utilisation and wavelength-reuse in these networks [2], and identified the range of parameters under which these networks have performance advantages over less complex quasi-static wavelength-routed optical networks (WRONs) [3]. In this work we analyse two functions, critical to the performance of these networks, that of burst-aggregation and dynamic wavelength-routing algorithms. Whilst most of the studies on dynamic routing and wavelength assignment (DRWA) algorithms have focused on the blocking probability or in performing computational complexity comparisons between different methods (e.g. [4-6]), the aim of this paper is to determine how fast a centralised DRWA algorithm for a WROBS should be to guarantee a pre-defined latency.

2. Burst aggregation techniques.

In the WROBS architecture, packets are buffered in edge routers, and timers are associated with each buffer. According to these timers, after a pre-defined time period (T) from the arrival of the first packet has elapsed, a request is sent in order to establish a lightpath between the source and destination edge routers for the transmission of the burst. Once a lightpath has been established, an acknowledgement is sent to the source node and packets in the buffer are transmitted. Two methods to aggregate bursts –*Limited-Size Bursts* (*LSBs*) and *Unlimited-Size Bursts* (*USBs*)– have been proposed [3]. In the first method (LSBs), the arrival of the acknowledgement to the edge router determines the end the aggregation process of the packets into a burst. Then, new packets arriving to the buffer during the burst transmission are not aggregated into this burst, but must wait for another lightpath to be established for their transmission. In the second method (USBs), new packets arriving at the buffer after the reception of the acknowledgement are considered as part of the current burst, and hence, the transmission process only finishes when the buffer is empty at which point the lightpath is deleted.

3. Lower bounds for the average edge delay.

An important parameter in WROBS is the edge delay (t_{edge}), which is defined as the waiting time for the first packet of a burst from its arrival to the buffer until its transmission. Assuming a centralised DRWA algorithm, the minimum achievable edge delay depends on twice the propagation delay from the edge router to the control node ($2t_{prop,ctrl}$) –request and acknowledgement propagation– and the processing delay in the control node, which includes the calculation time of the algorithm (t_c) and the queueing delay. For very low loads, the queueing delay is negligible, so that the average edge delay is $t_{edge} = 2t_{prop,ctrl} + t_c$. For higher traffic loads the requests will spend some time in the queue. Moreover, in order to save resources and have an easier control, it is assumed that there is at most one lightpath between source-destination edge routers per class of service. In order to fulfil this constraint, the control node only processes a request when the previous lightpath associated to the same source-destination pair has been deleted. Then, if a request to establish a connection between edge routers (s, d) is put into the control node as soon as the lightpath joining those nodes is deleted, the minimum possible average edge delay that the network can provide without prediction is achieved. Assuming an ideal DRWA algorithm (so that the blocking probability is zero), and that there is only one class of service, the system can be modelled as a closed queueing network [7]. When solving that system, the lightpath load (r), which is the average number of lightpaths established normalised by $N(N-1)$, is

$$r = \frac{1}{N(N-1)}\left(\frac{t_{WHT}}{t_c}\right)\left[\sum_{i=0}^{N(N-1)-1}\frac{(t_{WHT}/t_c)^i}{i!}\Big/\sum_{i=0}^{N(N-1)}\frac{(t_{WHT}/t_c)^i}{i!}\right],$$

where N is the number of edge routers in the network and t_{WHT} is the wavelength holding time, i.e. the time a lightpath is reserved for a burst transmission. Using the results in [3], for an homogeneous situation (equal traffic load and propagation delays between all source-destination pairs), the average traffic load (n) between edge routers (defined as the ratio between the buffer average input bitrate and the output bitrate, b_{core}) is given by $n = r\,[1-(2t_{prop,ctrl}/t_{WHT})]$, and $t_{edge} = t_{WHT}/r$ for LSBs and $t_{edge} = (t_{WHT}/r) - t_{WHT} + 2t_{prop,ctrl}$ for USBs. In general, it is not possible to obtain a closed form solution for t_{edge} as a function of n, but numerical results are easily calculable (Fig. 1). This model is not valid for very low traffic loads, where it overestimates the edge delay, as equations shown for t_{edge} are not valid for low traffic loads (see appendix in [3]).

In order to validate the analytical model, optical networks with N = 8, 14 and 20 edge routers were simulated. The traffic arriving at each buffer was assumed to be bursty with ON/OFF periods of duration determined by Pareto distributions ($P(t)=\alpha A^{\alpha}/t^{\alpha-1}$), setting α=1.5, the minimum length of the ON periods (A) to 4000 bytes, and adjusting the minimum length of the OFF periods to give the desired traffic load (n). The capacity of the lightpaths was assumed to be b_{core} = 10 Gbps. The network diameter of all networks was assumed to be 1000 km ($2t_{prop,ctrl}$ = 5 ms), and the delays between edge routers were randomly selected subject to the diameter constraint. When a timer set to a value of T ms at each buffer expires, a request is sent to the control node. The control node uses an Earliest Deadline First (EDF) policy to select the next request to process. The calculation time of the algorithm was distributed according to a Beta(5,5) distribution, which gives values bounded by 0 and 1 ms, and with average t_c = 0.5 ms. A uniform average traffic load between all source-destination pairs was also assumed. Figure 1 shows the average edge delay predicted by the model (solid lines) and simulation results (with 95% confidence intervals) for T = 0 ms. Results are in good agreement with the analytical model, except for very low traffic loads as expected (where the real lower bound tends towards $2t_{prop,ctrl} + t_c$). A logarithmic scale has been used in Fig. 2 to show low traffic loads. For LSBs, the maximum value of the lower bound of the average edge delay is $t_{edge} = t_c \cdot N(N-1)$ except for very high traffic loads (typically higher than 0.9 for big networks). For USBs, there is a critical traffic load where the lower bound for the average edge delay reaches a maximum, which is, approximately, $t_c \cdot N(N-1)$. Simulations were also carried out assuming different traffic loads for every source-destination pair (according to a beta distribution), and also using lower values of the minimum length of the ON periods (400 and 40 bytes). The results also matched the analytical model except for low traffic loads.

4. Impact of timers on the average edge delay.

As previously stated, timers are associated with each buffer at edge routers to determine when to send a request. To evaluate its impact, simulations were performed for networks with 8 edge routers and timers of T = 10, 20, 30 and 40 ms. As shown in Fig. 3, the edge delay obtained is $t_{edge} \approx \max\{T +$

$2t_{prop,ctrl} + t_c$, lower bound for t_{edge}}, therefore in order to achieve a required average edge delay ($t_{edge\_required}$) independently of the traffic load in a network with N edge routers, the average calculation time of the DRWA algorithm should hold $t_c < t_{edge\_required}/[N(N-1)]$, and timers should be set to $T = t_{edge\_required} - 2t_{prop,ctrl} - t_c$.

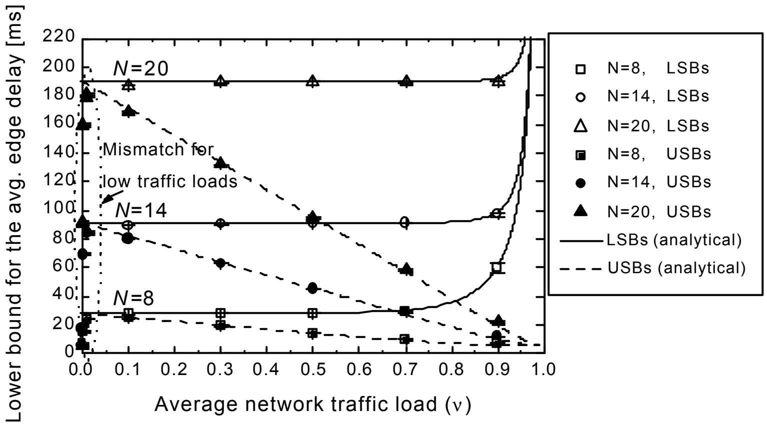

Figure 1: Lower bounds (analytical and simulation results) for the average edge delay, plotted in natural scale. The results are shown for different network sizes (8, 14 and 20 edge routers) and for both burst aggregation techniques (LSB and USB).

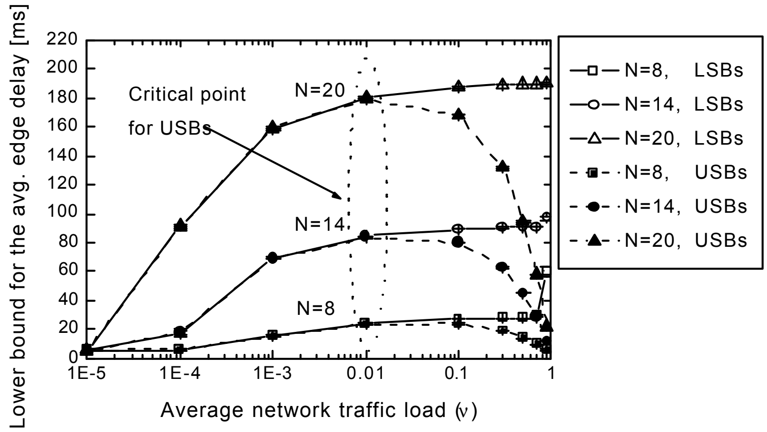

Figure 2: Lower bounds (only simulation results) for the average edge delay, plotted in logarithmic scale. The results are shown for different network sizes (8, 14 and 20 edge routers) and for both burst aggregation techniques (LSB and USB). Note that these results are exactly the same ones as in Fig. 1, the only variation is the use of a logarithmic scale.

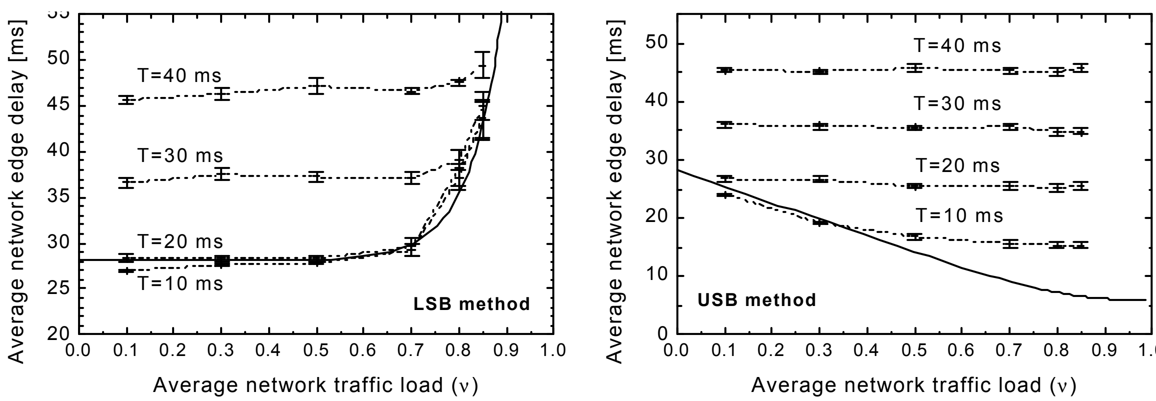

Figure 3: Impact on the average edge delay due to the utilisation of timers in a network with 8 edge routers. Solid lines show the analytical lower bounds for the average edge delay. Dashed lines show the average edge delay achieved when using different timers (from T = 10 to T = 40 ms). The left-hand side figure shows results for LSBs, the right-hand side figure shows results for USBs.

5. Summary.

Lower bounds for the average edge delay have been determined as a function of the number of edge routers, the burst aggregation method and the average calculation time (t_c) of centralised DRWA algorithms. To ensure a desired value of t_{edge} for all traffic loads, a DRWA algorithm with $t_c < t_{edge}/[N(N-1)]$ should be used. With that condition, a network using USBs can support any traffic load. When using LSBs high traffic loads (typically higher than 0.9) cannot be supported by the network, as it becomes unstable. The results in this paper allow a network designer to determine how fast a centralised DRWA algorithm should be to fulfil the requirements on delay of the class of service. For instance, in order to allow videoconference services, the end-to-end delay should be below 100 ms [8]. In that case, in a network with 20 edge routers and 1000 km of diameter, the edge delay should be lower than 80 ms (assuming 20 ms are reserved for processing at end nodes and propagation delay). Therefore, the average calculation time of the centralised DRWA algorithm should be below $t_{edge\_required}/[N(N-1)] = 80/380 \approx 0.2$ ms to provide that service independently of the traffic load the network is supporting. If the network is known to operate with high loads, a slower algorithm could be used if USBs are employed, as the lower bound for the edge delay decreases when increasing the traffic load (Fig 1).

Acknowledgements.

The authors thank D. Wischik (Univ. of Cambridge), E. Kozlovski, A. Myers and A. Zapata (UCL), and F. González (Univ. of Valladolid) for valuable comments. Financial support from the Royal Society and Marconi Communications is acknowledged. I. de Miguel is also grateful to the Optical Communications Group at Univ. of Valladolid for their help during his leave.

References.

[1] M. Düser, E. Kozlovsky, R.I. Killey and P. Bayvel. "Design Trade-Offs in Optical Burst Switched Networks with Dynamic Wavelength Allocation", *Proc. European Conference on Optical Communications 2000 (ECOC 2000)*, Vol. 2, paper Tu 4.1.4, pp. 23-24

[2] M. Düser and P. Bayvel. "Bandwidth Utilisation and Wavelength Re-Use in WDM Optical Burst-Switched Packet Networks", *Proc. IFIP/TC6 5th Working Conference on Optical Network Design and Modelling 2001 (ONDM 2001)*, Vol. 1

[3] I. de Miguel, M. Düser and P. Bayvel. "Traffic Load Bounds for Optical Burst-Switched Networks with Dynamic Wavelength Allocation", *Proc. IFIP/TC6 5th Working Conference on Optical Network Design and Modelling 2001 (ONDM 2001)*, Vol. 1

[4] H. Zang, J.P. Jue, B. Mukherjee. "A Review of Routing and Wavelength Assignment Approaches for Wavelength-Routed Optical WDM Networks", *Optical Networks Magazine*, Vol.1, No. 1, Jan. 2000, pp. 47-60

[5] S. Xu, L. Li, S. Wang. "Dynamic Routing and Assignment of Wavelength Algorithms in Multifiber Wavelength Division Multiplexing Networks", *IEEE Journal on Selected Areas in Communications*, Vol. 18, No. 10, Oct. 2000, pp. 2130-2137

[6] A. Mokhtar, M. Azizoglu, "Adaptive Wavelength Routing in All-Optical Networks", *IEEE/ACM Transactions on Networking*, Vol. 6, No. 2, Apr. 1998, pp. 197-206

[7] G. Bolch, S. Greiner, H. de Meer and K.S. Trivedi. *Queueing Networks and Markov Chains: Modeling and Performance Evaluation with Computer Science Applications.* John Wiley & Sons, 1998

[8] M. Baldi and Y. Ofek. "End-to-End Delay Analysis of Videoconferencing over Packet-Switched Networks", *IEEE/ACM Transactions on Networking*, Vol. 8, No. 4, Aug. 2000, pp. 479-492