

OPTICAL FEEDBACK BUFFERING STRATEGIES

Ronelle Geldenhuys^{1,2}, Jesús Paúl Tomillo², Ton Koonen² and Idelfonso Tafur Monroy²

¹*University of Pretoria, Pretoria, 0002, South Africa, ronelle.geldenhuys@eng.up.ac.za*

²*Eindhoven University of Technology, Den Dolech 2, 5600MB, The Netherlands, i.tafur@tue.nl*

Abstract: This paper considers the performance of fixed and incremental feedback buffers in an optical switch, and compares this with the performance of a feedback buffer configuration implementing switchable delay lines. It is shown that for a medium sized switch the switchable delay line implementation outperforms the other configurations, although for a large switch, the incremental structure has the best performance.

1. INTRODUCTION

Contention resolution in optical packet switches is implemented using fibre delay lines (FDLs) in either a travelling or a feedback (also called recirculating) configuration[1]. Travelling buffers are either input or output buffers, and have limited performance with respect to the packet loss ratio because of the huge amount of fibre required for sufficient buffering of bursty traffic. Feedback buffers have the advantage that the FDLs are reused thus decreasing the total amount of fibre required to achieve an acceptable packet loss ratio. There are however various disadvantages: the optical signal must be amplified on each recirculation resulting in an increase in amplified spontaneous emission (ASE) noise from the optical amplifiers; and the buffered signal has to traverse the switch fabric on each recirculation resulting in crosstalk and optical loss[2].

There are 2 factors supporting the implementation of feedback buffers in optical switches. The first is the improved performance of switch technology, for example decreased crosstalk and optical loss in all-optical switch technology[3]. The second is the self-similar nature of Internet traffic resulting in a situation where it is almost

impossible to buffer traffic satisfactorily, even when using feedback buffers with a lot of fibre[4].

This paper compares feedback buffering (Figure 1) with 3 configurations: fixed, incremental, and using switchable delay lines (SDLs).

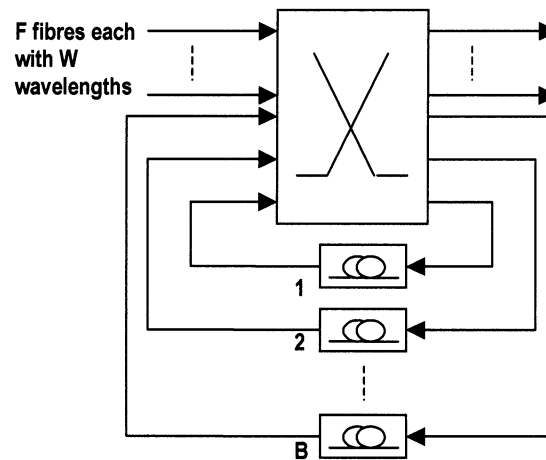


Figure 1. Schematic representation of a non-blocking switch with a feedback buffer. There are B FDLs and can be implemented in one of three configurations: 1. fixed, 2. incremental, 3. SDLs

2. TRAFFIC MODEL

Traditionally, short-range dependent models are used to model traffic, but real traffic displays burstiness on a wide range of time scales. Large-scale correlation refers to correlations that last across large time scales. Long-range dependence refers to values at any instant being positively correlated with values at all future instants. These characteristics result in Internet traffic being described as selfsimilar. Self-similarity is usually defined in statistical or qualitative terms, loosely including anything that "looks like itself" when magnified.

In this paper, three traffic models are used: Bernoulli traffic is the simplest traffic pattern; geometrically distributed ON and OFF periods are used to model bursty traffic; and a Pareto distribution is used to simulate self-similar traffic. To simulate a heavy-tailed distribution, a Pareto distribution can be used to produce "pseudo-self-similar" arrival processes[5]. This traffic has large-scale correlations but the traffic is not actually long-range dependent. When using a self-similar model to describe long-range dependence, only a single parameter is required: self-similarity is characterised by the Hurst parameter, H , which relates linearly to the shape parameter, α , of the heavy-tailed file size distribution in the application

layer. $0.5 < H < 1.0$, and as H approaches 1, both selfsimilarity and long-range dependence increases. $0 < a < 2$, and if $a < 2$ then the distribution has infinite variance, and if $a \leq 1$ then the distribution has infinite mean. According to [6] a typical value for a is 1.2.

3. BUFFER STRATEGIES

3.1 Fixed Length Feedback Lines

F is the number of input and output fibres and B is the number of feedback fibres. The simplest feedback approach is to consider B equal length FDLs connecting output and input ports of the switch. We assume delays of one time slot for this configuration.

The effect of the buffering depends strongly on the kind of traffic used. Simulations show that for Bernoulli traffic the improvement achieved by adding an extra feedback fibre is very important, but the improvement is not so considerable for the other traffic models.

3.2 Incremental Length Delay Lines

In this configuration the delays are distributed from 1 to B time slots. Here, a scheduling strategy is required to decide to which feedback fibre contending packets should be delivered. The three strategies simulated, based on [7], are as follows:

1. *minDelay* – the delay of the packets is minimised by sending bursts through the minimum delay fibre available.
2. *noOvr* – reservation of resources is done in advance so that when a new burst arrives at the switch its destination, route and the resources used are calculated in advance.
3. *avoidOvr* – this is similar to *noOvr* except that if no buffer is found the packet/burst is not dropped, it is sent to the minimum delay line available.

3.3 Buffering with Switchable Delay Lines

The switchable delay line structure is shown in Figure 2. This structure can switch the length of a fibre and get delays ranging from 0 to $7T$. T depends on the length of the delay loops.

The following three scheduling strategies were used:

1. *useBusy* – using the minimum possible delay, first try to use a busy (already configured) SDL, and then an idle SDL. If there is no available minimum delay, the packet/burst is dropped.

2. *minBusy* – using the minimum possible delay, first try to use an idle SDL configuring it with that delay and if not possible, use a busy SDL. If this is not possible, the packet/burst is sent through a minimum delay available SDL.
3. *minIdle* – similar to *useBusy*, but if the required minimum delay is not found, then the packet/burst is sent through a minimum delay available SDL.




Figure 2. Schematic representation of a switchable delay line.

4. RESULTS

Event-driven simulation based on slotted operation was used. Both Bernoulli and self-similar traffic models are used. A Bernoulli traffic model can be used to simplify configurations where relative results are required.

A burst is a sequence of one or more packets sent in consecutive time slots from the source and all with the same destination. Two different approaches were used for modelling bursts:

1. *Packet trains*: The sequence of packets cannot be segmented and all the packets that compose the burst follow the same route as the first one[8].
2. *Packet wagons*: Packets that compose the burst are considered independently[9].

The influence of the following factors are shown in Figures 3 – 8:

- Traffic type and load
- Number of feedback fibres
- Fixed versus incremental FDLs
- Buffering strategies
- Switch size – Medium: $F = 4$, $W = 8$; Large: $F = 8$, $W = 32$

Note that for the incremental delay structures the strategy *avoidOvr* is an upgraded version of *noOvr* and always behaves better, so to improve readability of

figures *noOvr* is not plotted.

Figure 3 shows the performance with Bernoulli traffic. B=0 means no feedback is used and it is plotted as a reference curve. Differences with 1 feedback fibre are insignificant, but with 2 fibres the incremental strategy *avoidOvr* has the best results.




Figure 3. Performance of a medium sized switch simulated with Bernoulli traffic. F=4, W=8.

Figure 4 shows the same configurations with self-similar traffic and using packet trains. For B=1 results for fixed feedback and SDL are almost the same. For B=2 the best algorithm is *minBusy* using SDL and the difference is bigger as the traffic load is smaller.




Figure 4. Performance of a medium sized switch simulated with self-similar traffic. F=4, W=8.

Figure 5 compares self-similar performance assuming wagons. For $B=1$ results are almost identical and the best configuration for $B=2$ is *minBusy* using SDL.




Figure 5. Performance of a medium sized switch simulated with self-similar wagons. $F=4$, $W=8$.

Figure 6 shows the performance of a large switch with $F=8$ and $W=32$ using Bernoulli traffic. For $B=1$ the fixed feedback structure behaves slightly better than the one based on SDL. For $B=2$ the difference is more notable and the best algorithm is *avoidOvr*, except for loads very near to 1, where the fixed delay structure is preferred. Figure 7 shows similar results for self-similar traffic.




Figure 6. Performance of a large switch simulated with Bernoulli traffic. $F=8$, $W=32$.




Figure 7. Performance of large switch simulated with self-similar wagons. $F=8$, $W=32$.

Figure 8 shows that the incremental delay configuration also outperforms the others when self-similar traffic trains are simulated.




Figure 8. Performance of large switch simulated with self-similar trains. $F=8$, $W=32$.

5. CONCLUSION

Three different feedback buffer architectures have been examined:

1. Fixed delay feedback, which needs a very simple control for the switch fabric.
2. Incremental delay feedback, that by just using more fibre allows for improvement of the performance.

3. Switchable delay lines, dynamic devices that allow a finer control over the routing process.

Two main scenarios were simulated: a medium and a large switch. The results obtained show that for the medium sized switch and self-similar traffic, the switch with SDL feedback and using the *minBusy* strategy obtains better performance, especially for low loads. For the large switch the results show that incremental structures behave better than those based on SDL. The reason for this is that the SDL is not a very dynamic system; it switches the whole fibre and not each single wavelength and thus all packets that come into the SDL at the same time slot will be assigned the same delay. Therefore the bigger W is, the less dynamic the SDL will be and then the worse its performance will be.

ACKNOWLEDGMENTS

This work was done as part of the European IST-project STOLAS.

REFERENCES

- [1] De Zhong W., Tucker R. S., "Wavelength Routing-Based Photonic Packet Buffers and Their Applications in Photonic Packet Switching Systems", *IEEE Journal of Lightwave Technology*, Vol. 16, No. 10, October 1998.
- [2] R. Langenhorst et al., "Fiber Loop Optical Buffer", *IEEE Journal of Lightwave Technology*, Vol. 14, pp. 324-335, March 1996.
- [3] M. T. Hill, A. Srivatsa, N. Calabretta, Y. Liu, H. de Waardt, G.D. Khoe, H.J.S. Dorren, "1 x 2 optical packet switch using all-optical header processing", *IEEE Electronic Letters*, vol. 37, pp. 774-775, 2001.
- [4] C. Davelder, M. Pickavet, P. Demeester, "Choosing an appropriate buffer strategy for an optical packet switch with a feed-back FDL buffer", in *Proc. ECOC*, 2002, Sep. 2002, Vol. 3, pp. 8.5.4.
- [5] V. Paxson and S. Floyd, "Wide Area Traffic: The Failure of Poisson Modeling", *IEEE/ACM Transactions on Networking*, vol. 3, pp. 226-244, June 1995.
- [6] M.E. Crovella and A. Bestavros, "Self-Similarity In World Wide Web Traffic: Evidence And Possible Causes", *IEEE/ACM Trans. Networking*, vol. 5, Dec. 1997.
- [7] Chris Davelder, Mario Pickavet and Piet Demeester, "Strategies for an FDL Based Feed-Back Buffer for an Optical Packet Switch with QoS Differentiation", *COIN 2002*, Cheju, Korea, July 2002.
- [8] R. Jain and S. A. Routhier, "Packet Trains – Measurements and a New Model for Computer Network Traffic", *IEEE Journal on Selected Areas in Communications*, Vol. SAC-4, No. 6, pp. 986-995. September 1986.
- [9] Chris Davelder, Mario Pickavet, Piet Demeester, "On trains and wagons: switching variable length packets in a slotted OPS network", *COIN 2003*.