

Performance Evaluation of a Router with Tunable Recirculating Buffers in an Optical Burst Switching Environment

K. Merchant, J. McGeehan, A. Willner
University of Southern California
{kkmercha, mcgeehan, willner}@usc.edu

S. Oviaia
Intel Corporation
shlomo.ovadia@intel.com

P. Kamath, J. Touch, J. Bannister
USC – ISI
{pkamath, touch, joseph}@isi.edu

Abstract

Optical burst switching presents challenges to the design of optical routers. This paper considers how to dimension a router of N input data ports with an additional M fiber delay lines ($N+M$ internal ports) in a hop-and-span constrained network. The router incorporates tunable FDLs that can vary their size to fit the burst being buffered. Tunable fiber delays achieve up to 20% higher throughput than static fiber delays at high input port load. Multiple recirculations are a critical requirement; when packets can circulate only once through the buffer, no measurable improvement is achieved after the number of as FDLs becomes equal to the number of data ports. When recirculation is permitted, throughput increases by up to 40%, depending on the combination of the number of FDLs added and the recirculation limit, which must increase in tandem. For a given number of FDLs, there is an optimal recirculation limit beyond which there is no measurable throughput benefit.

Keywords – optical burst switching, recirculation, fiber delay lines, optical buffering.

1. Introduction

Optical Burst Switching (OBS) schemes support high-speed bursty data traffic over wavelength-division-multiplexed (WDM) optical networks [1–3]. The OBS scheme offers a practical compromise between current optical circuit-switching and emerging all-optical packet switching technologies. In addition, the OBS scheme achieves high bandwidth utilization and quality of service (QoS) via the elimination of electronic bottlenecks and by using a one-way end-to-end bandwidth reservation scheme with variable time slot duration provisioning. Optical switching fabrics are

attractive because they offer at least one or more orders of magnitude lower power consumption with a smaller form factor compared to O-E-O (optics-electronics-optics) switches. Most of the recently published work on OBS networks focuses on next-generation backbone data networks (i.e. metropolitan or Internet-wide networks) using high-capacity (i.e. 1 Tb/s) WDM switch fabrics [4–7]. It has been previously suggested that the OBS scheme can be adapted to future high-speed enterprise networks in order to meet the growing demand for high bandwidth applications such as multimedia multicasting at a low cost [8].

One way to achieve some of these goals is by enhancing the performance of a core router using fiber delay lines (FDL). There have been several studies in the past that have tried to evaluate optical router performance with FDLs for burst and packet switching networks. Gauger's work [9] includes the evaluation of different buffering architectures for a wide-area OBS environment. This study compares simulation results for dimensioning of feed-forward (FF) buffers for the PreRes (output port reserved before the burst enters the FDL) scheme and feedback (FB) buffers for the PostRes (output port reserved after the burst enters the FDL) scheme. However, the study restricts its evaluation up to 4 recirculations and concludes that increasing the number of recirculations helps to improve the performance. Singh et al. [10] have analyzed the performance of a router using synchronous traffic and have provided exact and approximate models for throughput and blocking loss characteristics. Analysis of a synchronous model can help to provide an upper bound on the performance. However, as shown in [11, 12], the traffic on Ethernet and wide area networks tends to be bursty in some or all time scales. Hence, to obtain a lower bound analysis using an asynchronous model would be more helpful. Tančevski et al. [13] have shown that voids created by asynchronous traffic can significantly degrade the performance of an optical router having FDLs. A void is a gap in the output port packet distribution that is the time an output port is free

because a burst was switched to a FDL having a length longer than the time the previous burst took for transmission at that the output port. In [14] void filling has been proposed as an alternative to expensive synchronizing hardware. However, this process of void filling is too complicated and computationally intensive to be realized in high-speed optical networks. Also in [14] the authors show that the performance of a router with feedback FDLs with asynchronous traffic depends on the number of recirculated ports and the recirculation limit. Our work investigates these two parameters for an optical router with tunable FDLs capable of burst recirculation. We assume tunable FDLs to reduce the deleterious effects of voids. A tunable FDL can change its size to fit a buffered burst hence reducing the time for which the output port is free after transmission of the previous burst (void).

Our proposed model does not suffer from the disadvantages of [10] because it evaluates the performance of the router using asynchronous bursty traffic. Also, as opposed to [9], we look at the router performance for a wide range (up to 1000) of recirculations and evaluate the trade-off between the increase in throughput and accompanied increase in average latency. Naturally, it is technologically infeasible to recirculate burst more than a few times – the experiments with large recirculation limits are intended to serve as limiting cases.

First we demonstrate the advantages of assuming a tunable FDL architecture over a statically sized FDL architecture. Dynamic FDLs can provide up to a 20% higher increase in throughput as compared to a static configuration for 32 port router with 256 FDLs and a recirculation limit of 16. Next we demonstrate that for a 32 port optical router with 32 tunable FDLs, a single recirculation provides about a 10% increase in throughput over the bufferless router. Increasing the number of FDLs beyond 32 for this configuration does not help. When the number of FDLs is increased to 256, up to 16 recirculations provide an improvement of 37% over the bufferless router. However increasing the number of recirculations beyond that provides a very small improvement (~2%) only at high loads. Also, with the maximum number of recirculations fixed at 16, the nonlinear throughput vs. load curve for 32 FDLs moves to a linear curve for 256 FDLs.

The rest of this paper is organized as follows. Section 2 discusses the router architecture and traffic characteristics. Simulation setup and parameter definitions and values are explained in Section 3. We explain our simulation results in Section 4. Section 5 concludes the paper.

2. Architecture

This work models a single router in a hop-and-span constrained OBS core [8] that routes bursts to the relevant output port based on their destination address. The label edge router (LER) function of aggregating packets into bursts is assumed to have been completed by using the burst assembly algorithm proposed in [15] for the LER model in an OBS environment. This algorithm (algorithm 2 in [15]) sets a timer as soon as a packet reaches the LER and when the timer elapses it then sends the burst into the core.

The control signaling is done out of band and there is a separate control channel for each data port. The control packet is sent ahead of the burst and contains information about the burst such as the input port, output port, burst length and expected burst arrival time at the node. It undergoes O-E conversion and configures the router based on the expected burst arrival time (the delayed reservation scheme of the JET protocol [16]). The control packet processing is implemented by maintaining a priority queue based on the expected burst arrival time so that the router can be configured for the next expected burst (as opposed to configuring based on the order of arrival of control packets). Should multiple bursts destined for the same output port be scheduled to arrive simultaneously, one is arbitrarily chosen to be routed to the appropriate port, while the others are buffered or dropped.

Tunable FDLs are a key feature of our model. A tunable FDL changes its size to fit the size of the burst that is buffered in it, based on the information provided by the burst's control packet. A practical router with tunable FDLs can be thought of as having a large set of static length FDLs and having control processing that can route a buffered burst to the FDL having size closest to it. Although it would be a difficult to implement tunable FDLs we believe it is a reasonable assumption to analyze the performance of the router.

We also model burst recirculation and study the effects on the performance with changing the maximum number of allowable recirculations. A burst needs to be recirculated if its intended output port is busy when it emerges from the FDL. However, recirculating the burst too many times affects switching performance and degrades the signal unacceptably.

Once a burst enters a FDL, it can only recirculate within that FDL (it cannot shift to some other FDL). If the output port is busy up to the maximum number of allowable recirculations then the burst is dropped. There can be only one burst at a time in the FDL due to the tunable FDL and recirculation features.

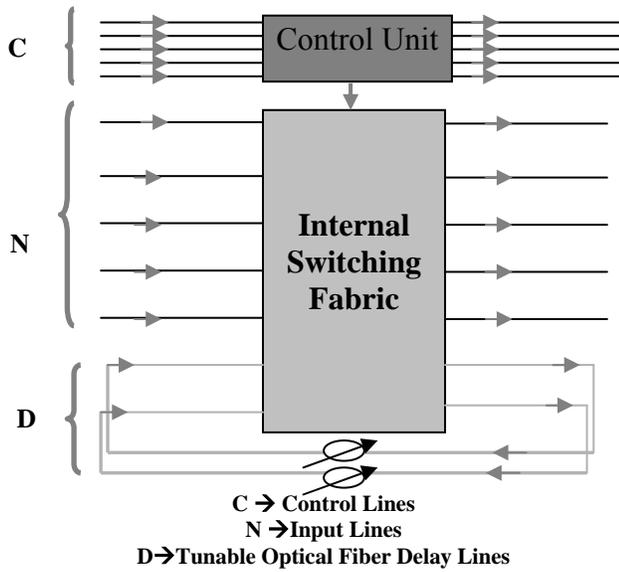


Fig. 1. Core router Architecture.

Fig. 1 shows a core router architecture with 5 control inputs, 5 input/output ports and 2 FDL ports. FDL ports are shared between all the output ports. Contention resolution at any output port can be handled by using any of the available FDL ports. All bursts are assumed to be on the same wavelength, and this model does not incorporate wavelength conversion.

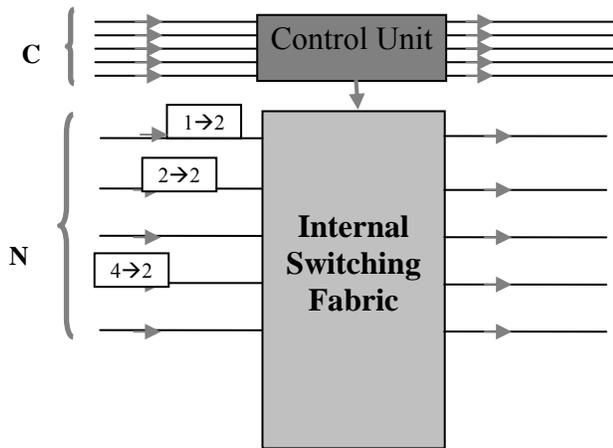


Fig. 2. Burst contention scenario – Bursts on input ports 1, 2 and 4 are all contending for output port 2.

Traffic is modeled with exponential distribution for the inter-arrival times of the bursts and a heavy-tail Pareto distribution for burst sizes. The interarrival time

is varied to adjust the load on the router. The output port distribution of the bursts is uniform. The combination of exponential inter-arrival times and heavy-tail burst length probability distributions is known to result in self-similar traffic [17].

We consider the burst propagation delay in the router to be negligible. Throughout the rest of the paper we will assume that limit of 1 recirculation means that a burst can pass through the only FDL once.

When bursts arrive on input ports 1, 2 and 4 that are all addressed to output port 2 there will be output port contention, as shown in Fig. 2. Without loss of generality assuming that the burst on port 1 arrives first, in that case it will be routed to output port 2. In a bufferless switch, this output port contention would result and the bursts on input ports 2 and 4 will be dropped from the network as no contention resolution scheme is available.

In the FDL contention resolution scheme, the burst is routed to an FDL port if there are any buffers available. Fig. 3 shows this type of contention resolution – the burst from input port 1 is routed to output port 2, while the bursts from input ports 2 and 4 are directed to FDL ports 1 and 2, respectively. Each FDL output is scheduled to connect to output port 2 when the burst emerges. If the output port is busy when a burst exits the FDL then the burst is recirculated up to a maximum number of allowable recirculations. If the output port is busy after the maximum number of recirculations are completed then the burst is dropped.

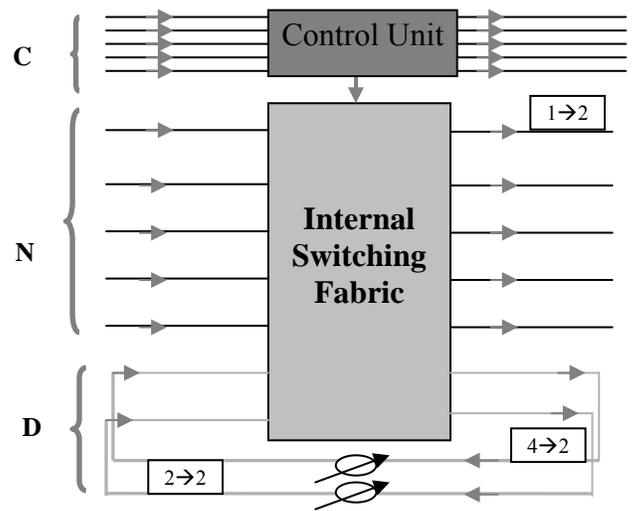


Fig. 3. FDLs for contention resolution – contending bursts are diverted to FDLs, where they can, upon exiting, be routed through the appropriate output ports.

This model contains two timers, called Timer O and Timer B, which are individually for each burst. Timer O sets the output port state FREE based on the time the burst is expected to leave the output port. This timer value is based on the expected burst arrival time and the burst transmission time (T). Timer B is set to switch the FDL port output to the corresponding output port based on when a burst is expected to emerge from an FDL. These timers are required as there is no other easy way of knowing when a burst will leave the output port/FDL as it is routed without any O-E-O conversion.

3. Simulation setup

a) Design

The model is analyzed using a C++ discrete event-driven simulator. The simulator is divided into different components such as input port, router, output port, delay line, queue and scheduler as shown in Fig. 4. The input port, output port and delay line components perform the functions of the corresponding part of the router. The router component handles the control packet processing and the switch configuration functions. The queue component is purely a simulation component and functions as a FIFO queue while transferring bursts from one component to another. The scheduler component works as a clock for the entire model and handles event scheduling. Synthetic bursts are generated using a separate model (based on an algorithm in [15]) that uses the simulation parameters to create traffic having output port distribution as uniformly random.

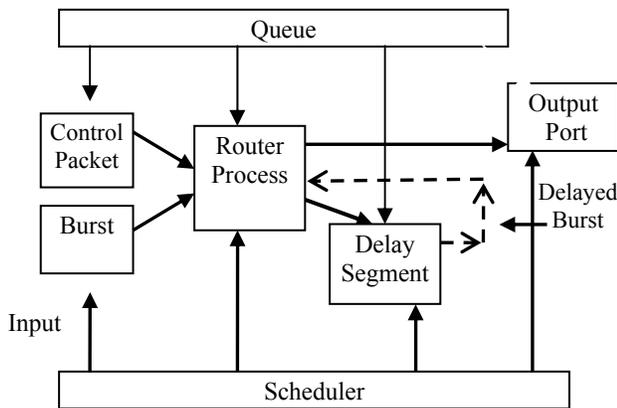


Fig. 4. Conceptual diagram of our simulator/model indicating the interconnection between simulator components.

Although traffic on a LAN may follow a more complex model, we have modeled the output port distribution as uniformly random as it is the simplest traffic model but can still provide important testing results for the switch.

b) Parameter Definitions

1. Normalized load per port: The ratio of the total number of bits that enter a router input port per second to the bit rate. It is averaged over all the input ports.
2. Normalized throughput per port: The ratio of the total number of bits that are routed through the output port of a router per second to the bit rate. It is averaged over all the output ports.
3. Latency: The time taken for the burst to route through the router. This includes the transmission time and buffering delay, if the burst is routed through an FDL port.
4. Average Latency: The ratio of the sum of the latency experienced by all the routed bursts to the number of bursts routed. The bursts that are diverted to the FDL(s) and dropped after exceeding the recirculation limit due to contention at the output port are not considered for the latency calculation

c) Parameter Values

The simulations were run with the following model parameter settings: line speed of 10 Gb/s, control packet processing time $1\mu\text{s}$. The output port of a burst is based on uniform distribution and burst size is based on a Pareto distribution with maximum burst size set to 10 kB. (Note here to generate synthetic bursts, a maximum value needs to be set for the Pareto distribution.) Hence the probability of generating a burst size of 10kB was set to 99.9999%. The following parameters are varied in the simulation

1. The number of router ports (N)
2. The number of FDL ports (M)
3. The number of maximum allowable recirculations (K)

The FDLs are tunable such that they just fit the burst being buffered. All the results are with 95% confidence intervals for five randomly seeded simulation runs. Simulations were performed for a router with N equal to 32 ports and M was varied from 0 (bufferless) to $N/2$ (16), N (32), $4N$ (128), and $8N$ (256). K values were 1, 8, 16 and 1000 (effectively infinite).

4. Simulation Results

a) Comparing the effects of dynamic and static variation of FDL sizes with constant number of FDLs and number of recirculations.

In this result we demonstrate the advantage of our architectural assumption of using tunable FDLs. Fig. 5 shows the simulation result of normalized throughput vs. load per port for dynamic and two configurations of statically sized 256 FDLs with up to 16 recirculations. The two configurations of the statically sized 256 FDLs, each of size $0.112L$ and $0.3L$ ($L = 10\text{KB}$), can accommodate about 50% and 99% of the bursts and show an increase of about 15% and 18%, respectively, over the bufferless case. The dynamically sized FDLs, which can tune their size to fit the burst being buffered according to information provided by the burst's control packet, show an improvement of about 37%. The reason for the higher increase with the dynamic configuration is because the length of the voids at the output port is reduced.

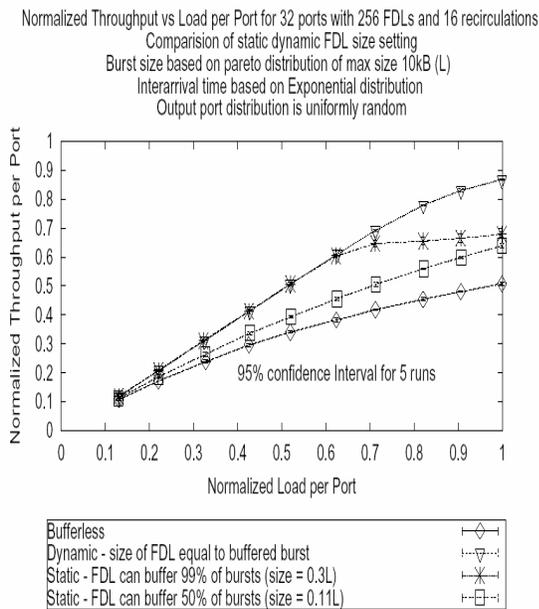


Fig. 5. Simulation result of Normalized Throughput vs. Load per Port for 32 port router with 256 FDLs and 16 recirculations for dynamic and static sized FDLs

The buffered burst has to wait the minimum possible amount of time (a FDL has to be as long as the buffered burst size to be able to buffer it) before it finds that its output port is free while in the static case a buffered burst that has yet to traverse the entire FDL might lose the output port to a new burst from an input port due to the excess length of the FDL. Essentially, fitting the burst exactly within a FDL means that in the heavily loaded case bursts are emerging from the full set of FDLs as rapidly as possible; thus the bursts are able to be resampled for a free output port at the fastest possible rate.

b) Constant number of recirculations and variable number of FDLs

1) 32 port router with maximum allowable recirculations restricted to 1.

Fig. 6 shows the simulation result of normalized throughput vs. load per port with 32 and 256 FDLs and maximum allowable recirculations restricted to 1. As shown in this figure, both the 32 and 256 FDL cases have nearly identical increases of $\sim 10\%$ over the bufferless router (the curves overlap). Curves for the 16 and 128 FDL cases would also overlap. This is because with recirculations restricted to 1 the FDLs are quickly freed and hence can buffer other bursts whose output ports are busy.

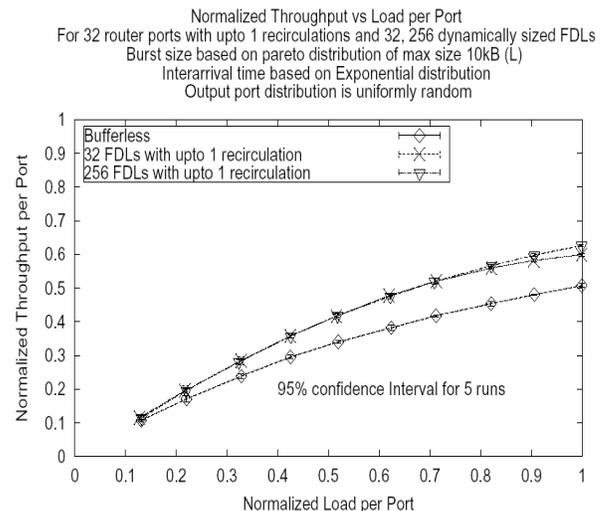


Fig. 6. Simulation result of Normalized Throughput vs. Load per Port for 32 port router with up to 256 FDLs and number of recirculations restricted to 1.

In addition, when the burst emerges after 1 recirculation, if the output port is busy, the burst will be dropped, which does not help in increasing the throughput. Simulations indicated that, given 32 router ports and recirculations restricted to 1, the maximum number of FDLs that are occupied is about 60 (even at 100% load). In other words, adding FDLs without increasing the recirculation limit in tandem will not yield any benefit. One way to increase the throughput for this configuration is by increasing the number of allowable recirculations as shown in the following Fig 7 to $K = 16$.

2) 32 port router with maximum allowable recirculations restricted to 16.

Fig. 7 shows the simulation result of normalized throughput vs. load per port with 16, 32, 128 and 256 FDLs and maximum allowable recirculations restricted to 16. As shown in this figure, increasing the recirculation limit up to 16 from 1 delivers a significantly enhanced performance for the same configuration. The 128 FDL and the 256 FDL cases have an increase in throughputs of about 20% and 25% respectively over and above the 1 recirculation limit case. This increase is because the buffered bursts have a higher probability of finding their output port free after recirculating more than once.

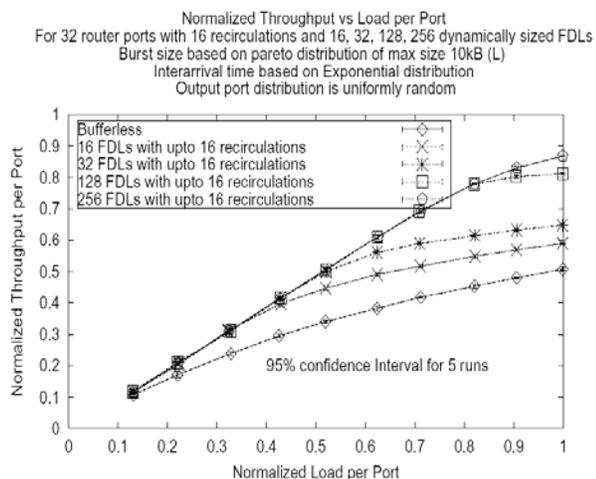


Fig. 7. Simulation result of Normalized Throughput vs. Load per Port for 32 port router with up to 256 FDLs and number of recirculations restricted to 16.

The simulation indicates that for 256 FDL configuration the average number of recirculations was about 7.5 and about 225 FDLs are used to buffer bursts.

The 128 and 256 FDL curves tend to overlap until high load cases as for low loads, less than 128 FDLs are needed to buffer all contending bursts.

c) Constant number of FDLs and variable number of recirculations

1) 32 port router with 32 FDLs.

Fig. 8 shows the simulation result of normalized throughput vs. load per port for the bufferless case, and for the 32 FDL buffered case with the maximum allowable recirculations set to 1, 8 and 1000. This figure also shows an increase in throughput of about 10% with 1 recirculation. Increasing the number of recirculations to 8 provides an increase of about 15% above the bufferless configuration, however, beyond 8 there is no increase in throughput. With 8 recirculations all the FDLs are filled up and the router starts to drop bursts that need to be buffered. Further increase in throughput requires corresponding increase in the number of FDLs.

2) 32 port router with 256 FDLs.

Fig. 9 shows the simulation result of normalized throughput vs. load per port for the bufferless case, and for the 256 FDL buffered case with the maximum allowable recirculations set to 1, 8, 16 and 1000. As shown in this figure, the increase in throughput obtained with 256 FDLs is higher compared to the 32 FDL case shown in Fig. 8 for the same number of allowed recirculations.

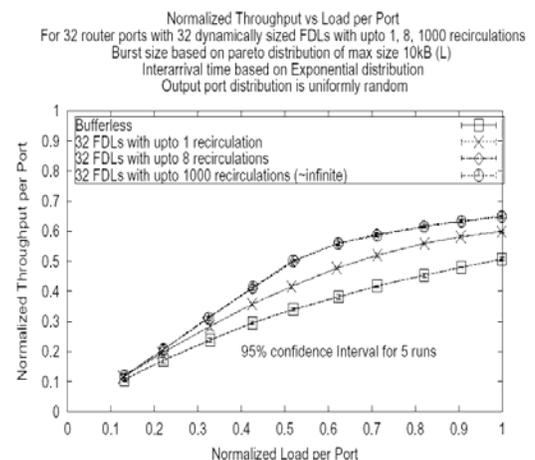


Fig. 8. Simulation result of Normalized Throughput vs. Load per Port for 32 port router with 32 FDLs and number of recirculations restricted to up to 1000.

Most of the increase is provided by the single recirculation case and higher numbers of recirculations provide diminishing returns. The 16 and 1000 recirculation cases provide an increase of about 37% over the bufferless case, these curves almost overlap. Thus, allowing a maximum of 16 recirculations seems ideal as increasing beyond this results in little change in throughput at a cost of significantly increased burst attenuation. Next we consider the effects of multiple recirculations on average latency.

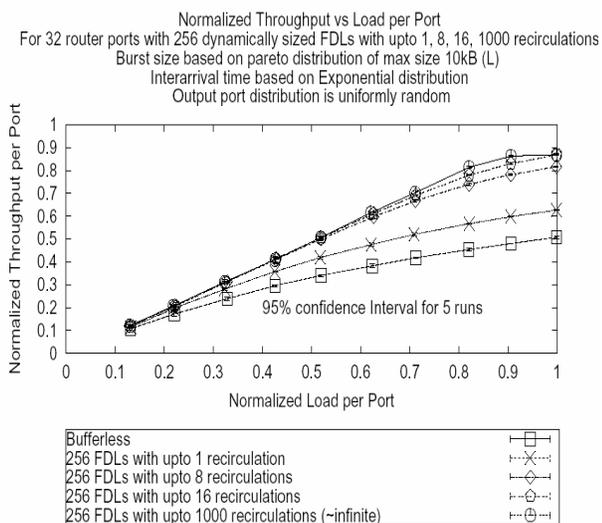


Fig. 9. Simulation result of Normalized Throughput vs. Load per Port for 32 port router with 256 FDLs and number of recirculations restricted to up to 1000.

Fig. 10 shows the simulation result of average latency vs. normalized load per port for the bufferless case, and for the 256 FDL buffered case with the maximum allowable recirculations set to 1, 8, 16 and 1000 (effectively infinite) recirculations. This figure demonstrates the trade-off associated with increasing the number the recirculations. Increasing the number of recirculations from 1 to 16 moderately increases the average latency by about 48% and provides a 20% increase in throughput. With up to 1000 recirculations the curve almost increases exponentially, reaching a near-maximum of about $9.6 \mu s$ while providing a negligible increase in throughput as compared to the 16 recirculations case. Thus, for this configuration, the 16 recirculations case is the preferred.

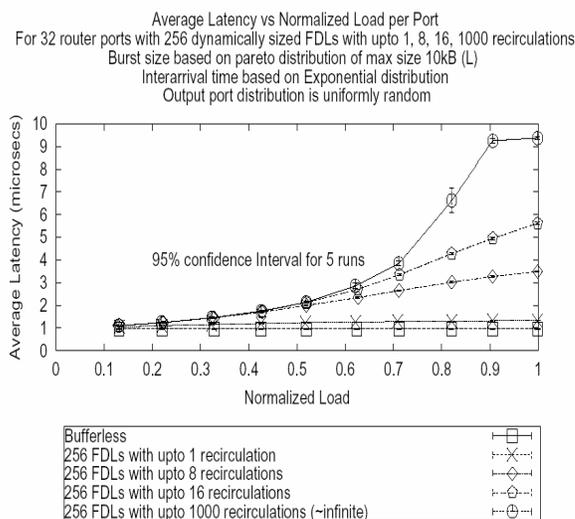


Fig. 10. Simulation result of Average Latency vs. Normalized Load per Port for 32 port router with 256 FDLs and number of recirculations restricted to up to 1000.

5. Conclusions

This paper presents a performance analysis of an optical router with FDLs including recirculation for a hop-and-span constrained environment. The analysis demonstrates the advantage of having a dynamic FDL architecture in providing considerable increase in throughput compared to a static FDL setting. Analysis of the dimensioning of the router as a function of the number of recirculations shows that although a single recirculation can help to increase the throughput, having multiple recirculations can provide a significant improvement. However, these results also showed that when the number of FDLs remains constant, increasing the number of recirculations beyond a threshold value provides diminishing returns, at a cost of increased attenuation and burst latency. When this threshold value is reached, to increase the throughput further we need to increase the number of FDLs.

Acknowledgment

The authors would sincerely like to thank S. M. Reza Motaghian Nezam for his valuable feedback after reviewing the paper.

References

- [1] S. Amstutz, "Burst switching – An update", IEEE Commun. Mag., pp. 50–57, Sept. 1989.
- [2] C. Qiao and M. Yoo, "Optical Burst Switching - A New Paradigm for an Optical Internet", J. High Speed Networks, Special Issue on Optical Networks, Vol. 8, No. 1, pp.69-84, 1999
- [3] J.S. Turner, "Terabit burst switching", J. High Speed Networks, Vol. 8, No. 1, January 1999, pp. 3–16.
- [4] C. Qiao, "Labeled Optical Burst Switching for IP-over-WDM Integration," IEEE Commun. Mag., vol. 38, no. 9, 2000, pp. 104–14.
- [5] J.Y. Wei and R.I. McFarland Jr., IEEE J. Lightwave Tech., no. 18, 2000, pp. 2019–37.
- [6] J.S. Turner, "WDM Burst Switching for Petabit Data Networks," Tech.Dig. OFC, 2000.
- [7] M. Düser and P. Bayvel, "Analysis of a Dynamically Wavelength-Routed Optical Burst Switched Network Architecture," IEEE J. Lightwave Tech., no. 20, 2002, pp. 564–85.
- [8] S. Ovadia, C. Maciocco, M. Paniccia "Photonic Burst Switching (PBS) Architecture for Hop and Span Constrained Optical Networks", IEEE Optical Commun., 2003
- [9] C. Gauger, "Dimensioning of FDL Buffers for Optical Burst Switching Nodes", Proc. Optical Network Design and Modeling (ONDM 2002), Torino, 2002
- [10] Y.N. Singh, A. Kushwaha, S.K. Bose, "Exact and approximate analytical modeling of an FLBM-based all-optical packet switch", IEEE J. Lightwave Tech., Volume: 21, Issue: 3, March 2003, pp. 719–726.
- [11] W.E. Leland, M.S. Taqqu, W. Willinger, D.V. Wilson, "On the self-similar nature of Ethernet traffic (extended version)", IEEE/ACM Trans. Networking, Volume: 2, Issue: 1, Feb. 1994, Pages:1–15
- [12] M.E. Crovella and A. Bestavros, "Self-similarity in World Wide Web traffic: evidence and possible causes", IEEE/ACM Trans. Networking, Volume: 5, Issue: 6, Dec. 1997, pp. 835–846
- [13] L. Tancčevski, A. Ge, G. Castanon, L. Tamil, "A new scheduling algorithm for asynchronous variable length IP traffic incorporating void filling", Proc. OFC/IOOC '99. Technical Digest Volume: 3, pp. 21–26 Feb. 1999
- [14] S. Shou-Kuo, T. Meng-Guang, T. Hen-Wai, P. Sreedevi, W. Jingshown, "Performance analysis of feedback type WDM optical routers under asynchronous and variable packet length self-similar traffic", Proc. High Performance Switching and Routing, 2004, April 2004, pp. 282–286.
- [15] R. Rajaduray, D.J. Blumenthal, S. Ovadia, "Impact of Burst Assembly Parameters on Edge Router Latency in an Optical Burst Switching Network", Proc. IEEE/LEOS Annual Meeting, Tucson, Oct. 2003.
- [16] M. Yoo and C. Qiao, "Just-Enough-Time (JET): A High Speed Protocol for Bursty Traffic in Optical Networks", Proc. IEEE/LEOS Conf. on Technologies for a Global Information Infrastructure, pp. 26–27, Aug. 1997.
- [17] A. Feldmann, A.C. Gilbert, W. Willinger, "Data networks as cascades: Investigating the multifractal nature of Internet WAN traffic", Proc. ACM Sigcomm '98, Vancouver, Sept. 1998, pp. 42–55.