Analysis of an Optical Burst Switching Router with Tunable Multiwavelength Recirculating Buffers

K. Merchant*, J. McGeehan[+], *Member, IEEE*, Student Member, OSA, A. E. Willner[+], *Fellow, IEEE, Fellow OSA*, S. Ovadia[⊥], *Senior Member, IEEE,* P. Kamath[A], *Student Member, IEEE,* J. Touch[A], *Senior Member, IEEE* and J. Bannister[A], *Senior Member, IEEE*

*Department of Computer Science
[+]Department of Electrical Engineering – Systems
University of Southern California
Los Angeles, California 90089-2565
Tel: (213) 740-4671, Fax: (213) 740-8729, kkmercha@usc.edu

[⊥]Intel Corporation
2200 Mission College Blvd
Santa Clara, California 95054-1537

[A]USC Information Sciences Institute
4676 Admiralty Way
Marina Del Rey, California 90292-6695

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 (FDLs) in an optical burst switching network. The router incorporates tunable FDLs that can vary their size to fit the burst being buffered. Tunable FDLs can be emulated using a set of static FDLs of unequal sizes. For this the size of static FDL set is monotonically increased, in step size increments, from minimum burst size till the throughput increase is equal to corresponding tunable FDL configuration. Tunable delays achieve up to 20% higher throughput than static 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 a 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. By varying the recirculation limit or number of FDLs, tunable buffering can match the gain achieved by wavelength conversion, possibly at lower hardware cost.

Index Terms: Optical burst switching, recirculation, fiber delay lines, optical buffering, wavelength conversion.

INTRODUCTION

Optical burst switching (OBS) supports high-speed, bursty 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) by eliminating 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 (FDLs) or by employing wavelength conversion. 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 buffers for the

PreRes (output port reserved before the burst enters the FDL) scheme and feedback 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 over many time scales. Hence, to obtain a lower bound analysis using an asynchronous model would be more helpful. Tančevski et al. [13] have shown that inter-burst 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. It is the time an output port is free because a burst was switched to a FDL having length longer than the transmission time of the previous burst at the contending 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 currently 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. This paper 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 and hence reduce the time for which the output port is free after transmission of the previous burst

(void). Recently there has been significant interest in the analysis and implementation of tunable FDLs. For example Liu et. al. [15] have demonstrated the implementation of all-optical delay by means of a recirculating loop controlled using optical processing technology. Sakamoto et al. [16] have demonstrated variable optical delay circuit using highly nonlinear fiber parametric wavelength converters.

Ramamirtham et al. [17] have been demonstrated that wavelength conversion can help improve the performance of routers in an OBS environment. Full wavelength conversion, however, is still considered too expensive and complex to be implemented practically. It has been shown in [18] that a router with a small range of wavelength conversion capabilities can achieve approximately the same improvement in performance as compared to a router with full wavelength conversion capabilities. This paper incorporates this assumption for the analysis of full wavelength conversion scenario. Also a similar analysis in [19] for optical packet switching concluded that for a small number of wavelengths it may be preferable to use optical buffers while in systems with large number of wavelengths, full wavelength conversion should be the preferred contention resolution scheme.

The proposed model presented herein does not suffer from the disadvantages of [10] as it evaluates the performance of the router using asynchronous bursty traffic. Also, as opposed to [9], router performance is examined for a wide range of recirculations (up to 1000) and the trade-off between the increase in throughput and accompanied increase in average latency is evaluated as well. Although, it is technologically infeasible to recirculate a burst more than a few times (the signal degrades without regeneration,

which is limited in all-optical networks) – the experiments with large recirculation limits are intended to serve as limiting cases.

First the advantages of assuming a tunable FDL architecture over a statically sized FDL architecture are demonstrated. Tunable FDLs can provide up to 20% increase in throughput as compared to a static configuration for a 32-port router with 256 FDLs and a recirculation limit of 16. Next the feasibility of assuming a tunable FDL architecture with is demonstrated using static FDLs having equal step size increments in length. By varying the maximum FDL size the set of static FDLs can achieve the same increase in throughput as the corresponding set of tunable FDLs. This work shows 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.

Next full wavelength conversion is compared with tunable buffering. With 8 wavelengths, wavelength conversion improves performance by about 28% at 100% load. To match this increase in throughput with 256 tunable FDLs, a recirculation limit of 4 is required. If the recirculation limit is raised to 16, for low loads 64 tunable FDLs can provide the same improvement as the 8-wavelength case while for the high load case 90 FDLs are required.

The rest of this paper is organized as follows. Section I discusses the router architecture and traffic characteristics. Simulation setup and parameter definitions and values are explained in Section II. Simulation results for the single wavelength and multiwavelength model in Section III and IV, respectively. Section V concludes the paper.

I. ARCHITECTURE

A router in the OBS core is modeled and its performance evaluated under varying buffering and wavelength conversion parameters. As shown in Fig. 1, the router can be thought of as an optical space switch with buffers/wavelength converters.

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 [20] for the LER model in an OBS environment. This algorithm (algorithm 2 in [20]) sets a timer as soon as a packet reaches the LER and when the timer elapses it then sends the burst into the core. (Note here that a burst refers to a data burst and in some previous works a control packet is called a control burst.)

In the core router, functions of switching, buffering and wavelength conversion are handled by the control unit possessing the control packets. The control packets contain information about the burst such as the input port, output port, burst length burst wavelength and expected burst arrival time at the node. Control packets undergo O-E conversion and configure the router based on the delayed reservation scheme of the just-enough-time (JET) protocol. This model assumes the JET protocol, as it has been shown to be a more efficient OBS control scheduling algorithm then other protocols [21].

Multiwavelength operation is likely to be the common mode of operation for most OBS systems of the future. Hence the later half of the analysis focuses on a multi-wavelength model and compares the advantages of wavelength conversion to multi-wavelength buffering. Full wavelength conversion is assumed, in which a burst contending on its destination output port wavelength can be converted to any available wavelength on the destination output port. Although it would be a difficult to implement full wavelength conversion among an arbitrarily large set of wavelengths, we believe it is a reasonable assumption given that we restrict the number of wavelengths to 8. For multiwavelength buffering a FDL is modeled such that it can simultaneously buffer bursts on different wavelengths contending for the same or different destination ports.

Tunable FDLs are a key feature of the model. A tunable FDL changes its size to just fit the size of the burst that is to be buffered in it, based on the information provided by the burst's control packet. We demonstrate that a set of static FDLs having step size increments can be used to model reasonably well a tunable FDL buffering system. Without loss of generality this analysis assumes that FDLs sizes are monotonically increasing and based on the following equation:

$$S = L + F * \xi$$

where

S is the FDL size, F is the FDL number which varies from 1 to R, R is the total number of FDLs, L is equal to the minimum size of the burst, ξ is the increment in the FDL size given by $[(U\text{-}L)/R]$ and U is the maximum size of the FDL and is increased till the throughput of the static set of FDLs approaches the throughput of set of tunable

configurations and. Fig. 5 explains this concept; a set of tunable FDLs are emulated using a set of static FDLs with step size increments.

We also model burst recirculation and study the effects on performance as maximum number of allowable recirculations varies. 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 may degrade 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 and recirculation features.

A tunable FDL can only change its length when there are no bursts present on any wavelengths in the FDL. A multiwavelength-tunable FDL is tuned to the length the first burst that enters it on a particular wavelength. Bursts entering later on other wavelengths must be shorter in length than the tuned size of the FDL. Only when it becomes completely empty can the FDL size be retuned.

Wavelength conversion is employed only for contention resolution; entering bursts are routed on the same wavelength on which they arrive if there is no contention. The delay for wavelength conversion is assumed negligible.

Traffic is modeled with an exponential interarrival burst time distribution and a heavy-tail Pareto burst size distribution. The interarrival time is varied to adjust the load on the router. The combination of exponential interarrival times and heavy-tail burst length probability distributions is known to result in self-similar traffic [22]. The destination wavelength and port distributions of the burst are uniformly random.

Although traffic on a local area network may follow a more complex model, the destination wavelength and port distributions are modeled as uniformly random because it is the simplest traffic model but can still provide important testing results for the switch.

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

This paper uses the convention '*xy*' port to represent the port (input or output) wavelength combination, where *x* is the port (input or output) number and *y* is the wavelength number, where both are single digits (for discussion).

Consider a scenario in which bursts arrive on input ports 11, 21 and 41 and are all addressed to destination port 21 and bursts on input ports 22 and 32 are bound for destination port 32. This results in a destination wavelength contention, as shown in Fig. 2. Without loss of generality assuming that the bursts on port 11 and 22 arrive first, they will be routed to output port 21 and 32, respectively. In a bufferless switch, this output port contention would result in the dropping of bursts on input ports 21, 32 and 41 as no contention resolution scheme is available.

In the FDL contention resolution scheme, bursts are routed to FDL ports if there are any buffers available. Fig. 3 shows this type of contention resolution – the bursts from input ports 11 and 22 are routed to output ports 21 and 32 respectively, while the bursts from input ports 21 and 32 are directed to FDL 1 and the burst on port 41 is routed to FDL 2. FDL 1 output is scheduled to connect to output ports 21 and 32 when the corresponding bursts emerge. 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.

In the wavelength conversion scenario, as shown in Fig. 4, again assuming without loss of generality, the burst on input port 11 gets routed to 21 and the burst on port 21 gets converted to wavelength 2 and gets routed to port 22. In the case of a two-wavelength system the burst on port 41 has to be dropped as both wavelengths on destination port 2 are now occupied. Bursts on ports 22 and 32 can both be switched almost simultaneously to destination port 3 by converting the one from input port 32 to port 31.

II. SIMULATION SETUP

a) Design

The model is analyzed using a custom C++ discrete event-driven simulator. The simulator is divided into different components such as input port, router, output port and delay line to model different components of the router. Synthetic bursts are generated using a separate model (based on an algorithm in [20]) that uses the simulation parameters to create traffic having output wavelength and port distribution as uniformly random.

b) Parameter definitions

1. Normalized load per wavelength per port: This is the ratio of the total number of bits per second that enter a router input port on a wavelength to the bit rate. It is averaged over all the input wavelengths and ports.

2. Normalized throughput per wavelength per port: This is the ratio of the total number of bits per second that are routed through the output port on a wavelength of a router to the bit rate. It is averaged over all the output wavelengths and ports.

3. Latency: This is the time taken by a burst to route through the router. This includes the transmission time and buffering delay, if the burst is routed through an FDL port. For wavelength conversion the delay assumed to be is zero.

4. Average Latency: This is 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 FDLs and dropped after exceeding the recirculation limit due to contention at the output port are not considered for the latency calculation.

5. Throughput delta (Δ): This is the difference between the throughput achieved by the set of tunable FDLs and the set of static FDLs with step size increments. It is normalized to 100% load i.e., Δ value of 1 indicates that the tunable FDL configuration has 100% throughput while the set of static FDL configuration has 0% throughput. It is used as a measure to demonstrate the feasibility of tunable FDLs using a set of static FDLs (with step size increments).

c) **Parameter values**

The simulations were run with the following model parameter settings: line speed of 10 Gb/s, control packet processing time of 1 μs. The output wavelength and port of a burst is based on the uniform distribution and burst size is based on the 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   10 kB was set to 99.9999%. (Also note that

although these results are for maximum burst size of 10 kB, higher burst size limit will only increase the average latency and the FDL sizes correspondingly. The following analysis is still valid). 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)

4. The number of wavelengths (W)

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. For most results the confidence interval bars are too small and hence not visible. Simulations were performed for a router with N equal to 32 ports. The average degree of core routers is 3. Also some core routers have 100's of data ports, hence we believe that 32 is a reasonable compromise between these values. M was varied from 0 (bufferless) to $N/2$ (16), N (32), $2N$ (64), $4N$ (128), and $8N$ (256). K values were 1, 8, 16 and 1000 (effectively infinite). W is one for the single wavelength results (first five results) and for multiwavelength analysis W is varied between 4 and 8. Each single wavelength and multiwavelength runs were with about 32 and 256 thousand bursts respectively.

III. SINGLE WAVELENGTH ANAYLYSIS

a) **Comparing the effects of tunable and static variation of FDL sizes with constant number of FDLs and number of recirculations.**
This result we demonstrates the advantage of our architectural assumption of using tunable FDLs. Fig. 6 shows the simulation result of normalized throughput vs. load

per wavelength per port for tunable 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 $0.3L$ ($L = 10$ 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 (tunable), 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 tunable configuration is because the length of the voids at the output port is reduced.

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 highly loaded case bursts are emerging from the full set of FDLs as rapidly as possible; thus the bursts are able to be re-sampled for a free output port at the fastest possible rate.

The throughput of the static 99% curve remains constant at about 67% after 70% load. This could be because packets that contend before buffering will also contend after buffering (because buffers are the same size). So at higher loads, buffering helps, but eventually contention catches up. The limit may tend to be around 65%. We plan to investigate this issue further. Note here that this result evaluated set of static FDL that are all of the equal sizes while the following results evaluate a set of static FDLs with step size increments. (of unequal sizes)

b) Tunable FDLs vs a set of static FDLs with step size increments

Given the benefits of assuming a tunable FDL architecture, it is useful to consider the feasibility of implementing tunable FDLs using a set of static FDLs sizes with step size increments based on the equation defined in the architectural assumptions. Fig. 7 shows the simulation result of throughput delta (Δ) vs. maximum size of FDLs with 32, 64, 128 and 256 FDLs and maximum allowable recirculations restricted to 16. The value of the maximum size of the FDL (x-axis) is varied from 1 kB to 2 kB as 1 kB is the minimum size of the burst and by 2 kB all the configurations have Δ value of 0.

As shown in this figure for the 32- and 64-FDL cases the curves decrease exponentially and at 1.2 kB and 1.25 kB (respectively) Δ reaches 0 and continues to stay at that value as the maximum size of FDLs increases. For the 128- and 256-FDL case the Δ value also decreases exponentially, albeit at a slower rate, with the increase in the maximum FDL size and reaches 0 at about 1.35 kB and 1.65 kB respectively. Thus we claim that by adjusting the maximum value of the FDLs a static set of FDLs can be used to implement tunable FDLs. Based on this analysis the rest of the results are evaluated using a tunable FDL architecture.

Fig. 8 shows demonstrates that after setting the size of set of static FDLs based on the previous analysis at 100% load (Fig. 7), throughput delta (Δ) value varies less than 10% for the 32- and 256-FDL configurations at all loads. This further reinforces the claim that we can emulate a tunable set of FDLs using a set of static FDLs with step size increments. Note here that the plots for 64 and 128 FDLs are similar to 32 and 256 FDLs (Δ varies less than 10%) however they not been shown in Fig. 8 for clarity.

c) Constant number of recirculations and variable number of FDLs

Fig. 9 shows the simulation result of normalized throughput vs. load per wavelength 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 about 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.

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 to $K = 16$, as shown in the following Fig. 10.

Fig. 10 shows the simulation result of normalized throughput vs. load per wavelength per port with 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 256-FDL cases have an increase in throughputs of about 20% and 25% respectively over 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.

The simulations indicated that for the 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.

d) Constant number of FDLs and variable number of recirculations

Fig. 11 shows the simulation result of normalized throughput vs. load per wavelength 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. Beyond 8, however, 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.

Fig. 12 shows the simulation result of normalized throughput vs. load per wavelength 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. 11 for the same number of allowed recirculations.

Most of the increase is provided by the single-recirculation case and higher numbers of recirculations provide diminishing returns. The 8, 16 and 1000 recirculation cases provide an increase of about 32%, 36% and 37%repectively over the bufferless case. The 16 and 1000 recirculation curves almost overlap. Thus, allowing a maximum of 16 recirculations seems ideal, as further increases result in little change in throughput

at a cost of significantly increased burst attenuation. Next the effects of multiple recirculations on average latency are considered.

Fig. 13 shows the simulation result of average latency vs. normalized load per wavelength 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 µs while providing a negligible increase in throughput as compared to the 16 recirculations case. Thus, for this configuration, the 16 recirculations case is preferred. (The 8 recirculation case provides about a 5% lower increase in throughput and may be preferred incase maximum throughput is not the main objective.)

IV. MULTIWAVELENGTH ANALYSIS

Next the performance in a multiwavelength scenario with up to 8 wavelengths is evaluated. It has been previously demonstrated [16] that wavelength conversion can significantly improve the performance of a router. The next section demonstrates this increase in throughput achieved with wavelength conversion and then discusses the trade-offs associated with achieving a similar increase using tunable buffering.

a) Increase in throughput with wavelength conversion

Fig. 14 shows the simulation result of normalized throughput vs. load per wavelength per port for a) 8 wavelengths with no FDLs and no wavelength conversion

and b) 4 and 8 wavelengths with wavelength conversion. As shown in the figure, incorporating wavelength conversion for 4 wavelengths provides an increase of about 20%. When the number of converters and wavelengths increases to 8 the increase in throughput is about 28%. Thus increasing the number of wavelengths and converters does increase the throughput but the returns are diminishing while it adds to the cost of hardware. Next we investigate that to achieve the same increase in throughput (28%) as the 8 wavelength conversion case what are the values of 1) the recirculation limit if router has a large number of FDLs (say $M = 8N$) and 2) the number of FDLs if the recirculation limit is large (say up to 16).

b) Comparison of wavelength conversion with buffering

1. With 256 FDLs and varying the number of recircualtions

Fig. 15 shows the simulation result of normalized throughput vs. load per wavelength per port for the curves with 8 wavelengths and a) no FDLs and no wavelength conversion, b) wavelength conversion and c) 256-FDL buffers with the maximum allowable recirculations restricted up to 4. The 256-FDL configuration with up to 4 recirculations achieves a similar increase in throughput (about 28%) as the 8 wavelength case with conversion. The nearly linear increase in throughput is because the large number of FDLs in this configuration provides sufficient buffering capacity to handle all the contending packets. (This is not the always case as we show in the next configuration – Fig. 17 and 18.)

Fig. 16 shows the simulation result of average latency vs. normalized load per wavelength per port for the curves with 8 wavelengths and a) no FDLs and no wavelength conversion, b) wavelength conversion, and c) 256-FDL buffers with the maximum

allowable recirculations restricted to 4. The wavelength conversion case demonstrates no increase in average latency (latency remains constant and equal to the single wavelength no FDL case). The configuration with 256 FDLs demonstrates an increase in average latency as the load increases. The 4 recirculation limit case demonstrates maximum latency of about 2.8 μs at 100% load.

From Figs. 15 and 16 it may appear that wavelength conversion is a better choice to increase the throughput as it does not increase the average latency as opposed to buffering which does increase the average latency. However the wavelength conversion configuration assumes that there is one converter per wavelength per port and for the 32-port router with 8 wavelengths 256 converters are needed. 256 converters are likely to be more expensive as compared to FDLs and thus the 256-FDL configuration with a slightly higher average latency may be the more efficient solution in this scenario.

A more optimal arrangement (in terms of hardware cost) can be found by increasing the recirculation limit (e.g. to 16). Given this higher recirculation limit, the number of FDLs required to provide the same increase in throughput as the 8-wavelength case with conversion can be determined.

2. With up to 16 recirculations and varying the number of FDLs

Fig. 17 shows the simulation result of normalized throughput vs. load per wavelength per port for the curves with 8 wavelengths and a) with no FDLs and no wavelength conversion, b) with wavelength conversion and c) 64- and 90-FDL buffers with the maximum allowable recirculations restricted up to 16. As shown in the figure up to 70% load, the 64-FDL case can achieve similar increase in throughput as compared to the 8-wavelength conversion case. Beyond 70% load about 90 FDLs are required to

provide increase comparable to the 8-wavelength conversion case. The non linear increase in throughput (the curve flattens out at about 70%) for the 64-FDL curve is because at high loads all the available FDLs are recirculating the bursts for long periods of time and the new contending bursts cannot be buffered and are thus dropped.

Fig. 18 shows the simulation result of average latency vs. normalized load per wavelength per port for the curves with 8 wavelengths and a) no FDLs and no wavelength conversion, b) wavelength conversion and c) 64- and 90-FDL buffers with the maximum allowable recirculations restricted up to 16. Similar to that shown in Fig. 16 the FDL scenario show an increase in average latency while the wavelength conversion case has no increase in average latency. The reason the FDL curves flatten out at about 65% load is because all the FDLs are filled up by 70% load. (same as above).

Figs. 17 and 18 show that by increasing the recirculation limit to about 16 we can achieve the same increase in throughput as the 8-wavelength conversion case with lower number of FDLs and thus lower hardware cost. However, increasing the number of recirculations may lead to increased attenuation and other physical layer problems such as dispersion and polarization mode dispersion (PMD)

c) Router Scaling

Fig. 19 shows the simulation result of number of FDLs vs. number of input/output ports for 35% increase in throughout at 100% load with recirculation limit of 16. The curve is plotted on a logarithmic scale as the number of input/output ports is increased exponentially from 4 to 64. The result demonstrates that the buffered router model grows linearly in terms of the number of FDLs required to achieve about the same increase in the throughput. For an increase of about 35% in throughput, *M* should be equal to about 8

times N. This result is limited up to $64N$; extending N further will need a similar increase in M.

V. CONCLUSION

This paper presents a performance analysis of an optical router with FDLs including recirculation for an optical burst switching environment. The analysis demonstrates the advantage of having a tunable FDL architecture in providing considerable increase in throughput compared to a static FDL setting. One way to implement tunable FDLs was demonstrated using a set of static FDLs with step size increments (by varying the maximum FDL size). 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.

Wavelength conversion was evaluated as an alternative contention resolution scheme. The comparison of wavelength conversion to buffering demonstrates that any increase in performance achieved using wavelength conversion can be matched by either varying the recirculation limit or the number of FDLs. Employing a larger FDL set (in terms of number of FDLs) with variation in the recirculation limit can provide a near-linear increase in performance but may lead to increased hardware cost. This increase in

hardware cost can be lowered by increasing the recirculation limit. However, that may lead to increased signal attenuation.

This work has evaluated tunable buffering and wavelength conversion individually and then compared them. A possible extension to this work could address the issue of combining these schemes and evaluating the parameter settings required to achieve an optimal increase in performance. Another extension could be to evaluate the performance after modifying the architecture to tune the FDL size equal to the time the output port is busy (instead of the size of the burst). Also analyzing the relation between the number input/output ports, FDL ports and recirculations to maximize the throughput and minimize the average latency could be an interesting future direction.

References

[1] S. Amstutz, "Burst switching – An update", IEEE Commun. Mag., Sept. 1989, pp. 50–57.

[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. 1,1999, pp.69-84.

[3] J.S. Turner, "Terabit burst switching", J. High Speed Networks, Vol. 8, No. 1, Jan. 1999, pp. 3-6.

[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., "Just-in-time signaling for WDM optical burst switching networks", IEEE/OSA J. Lightwave Tech., No. 18, 2000, pp. 2019-37.

[6] J.S. Turner, "WDM Burst Switching for Petabit Data Networks," OFC Technical Digest, 2000.

[7] M. Düser and P. Bayvel, "Analysis of a Dynamically Wavelength-Routed Optical Burst Switched Network Architecture", IEEE/OSA J. Lightwave Tech, No. 20, 2002, pp. 564-85.

[8] S. Ovadia, C. Maciocco, M. Paniccia, R. Rajaduray, "Photonic Burst-Switching (PBS) Architecture for Hop and Span Constrained Optical Networks", IEEE Optical Commun. Mag., Vol. 41, No. 11, 2003, pp. S24-S33.

[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/OSA J. Lightwave Tech, Vol.21, No. 3, Mar. 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, Vol. 2, No. 1, Feb. 1994, pp.1-15.

[12] M.E. Crovella and A. Bestavros, "Self-similarity in World Wide Web traffic: evidence and possible causes", IEEE/ACM Trans. Networking, Vol. 5, No. 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 Feb 1999. Technical Digest Vol. 3, pp. 21-26.

[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, Apr. 2004, pp. 282-286.

[15] Y. Liu, M. Hill, R. Geldenhuys, H. de. Waardt, G.D. Khoe, H. J. S. Dorren, "Demonstration of an all-optical variable delay for re-circulating buffers", Proc. ECOC 2004, Vol. 4, pp. 892-893.

[16] T. Sakamoto, A. Okada, O. Moriwaki, M. Matsuoka, K. Kikuch, "Performance analysis of variable optical delay circuit using highly nonlinear fiber parametric wavelength converters", IEEE/OSA J. Lightwave Tech., Vol. 22 , No. 3 , Mar. 2004, pp. 874 – 881.

[17] J. Ramamirtham, J. Turner, J. Friedman, "Design of wavelength converting switches for optical burst switching", IEEE J. Selected Areas in Commun.**,** Vol. 21, No. 7, Sept. 2003, pp. 1122-1132.

[18] S. Gangxian, S. Bose, C. Hiang, L. Chao, C. Yoong, "Performance study on a WDM packet switch with limited-range wavelength converters", IEEE Commun. Lett., Vol. 5, No. 10, Oct. 2001, pp. 432-434.

[19] A. Bononi, G. Castanon, O. Tonguz, "Analysis of hot-potato optical networks with wavelength conversion"**,** IEEE/OSA J. Lightwave Tech., Vol. 17, No. 4, Apr. 1999 pp. 525-534.

[20] R. Rajaduray, S. Ovadia, D.J. Blumenthal, "Analysis of an edge router for span-constrained optical burst switched (OBS) networks", IEEE/OSA J. Lightwave Tech., Vol. 22, No. 11, Nov. 2004 pp. 2693-2705.

[21] 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.

[22] 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.

Figure Captions

Figure 1: Core router architecture

Figure 2: Burst contention scenario – Bursts on wavelength 1 on input ports 1, 2 and 4 are all contending for wavelength 1 on output port 2. Bursts on wavelength 2 on input ports 2 and 3 are contending for wavelength 2 on output port 3.

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

Figure 4: Wavelength conversion for contention resolution – contending bursts are converted on available wavelength on the output port. The burst on wavelength 1 at input port 4 has to be dropped as no wavelength is available for conversion at output port 2 (2-wavelength system).

Figure 5: Tunable FDLs emulated using a set of static FDLs with variable sizes. L is the lower limit of FDL size, ξ is the increment in the FDL size and n is number of FDLs –1.

Figure 6: Normalized throughput vs. load per wavelength per port for 32-port router with 256 FDLs and 16 recirculations for tuanble and static sized FDLs

Figure 7: Throughput delta (Δ) vs. maximum size of FDLs for 32-port router with 256, 128, 64 or 32 FDLs and number of recirculations restricted to 16

Figure 8: Throughput delta (Δ) vs. load for 32-port router with 256 FDLs or 32 FDLs and number of recirculations restricted to 16. Maximum size of FDL set at throughput delta value equal to 0. (Figure 7)

Figure 9: Normalized throughput vs. load per wavelength per port for 32-port router with up to 256 FDLs and number of recirculations restricted to 1

Figure 10: Normalized throughput vs. load per wavelength per port for 32-port router with up to 256 FDLs and number of recirculations restricted to 16

Figure 11: Normalized throughput vs. load per wavelength per port for 32-port router with 32 FDLs and number of recirculations restricted to up to 1000

Figure 12: Normalized throughput vs. load per wavelength per port for 32-port router with 256 FDLs and number of recirculations restricted to up to 1000

Figure 13: Average latency vs. normalized load per wavelength per port for 32-port router with 256 FDLs and number of recirculations restricted to up to 1000

Figure 14: Normalized throughput vs. load per wavelength per port for 32-port router with 4 and 8 wavelengths using wavelength conversion

Figure 15: Normalized throughput vs. normalized load per wavelength per port for 32-port router with 8 wavelengths, 256 FDLs and number of recirculations restricted to up to 8

Figure 16: Average latency vs. normalized load per wavelength per port for 32-port router with 8 wavelengths, 256 FDLs and number of recirculations restricted to up to 8

Figure 17: Normalized throughput vs. normalized load per wavelength per port for 32-port router with 8 wavelengths, up to 100 FDLs and number of recirculations restricted to up to 16

Figure 18: Average latency vs. normalized load per wavelength per port for 32-port router with 8 wavelengths, up to 100 FDLs and number of recirculations restricted to up to 16

Figure 19: Number of FDLs vs. number of input/output ports for 35% increase in throughout at 100% load with recirculation limit of 16

Figures

C – Control Line
N – Data Lines
M – Tunable Fiber Delay Lines
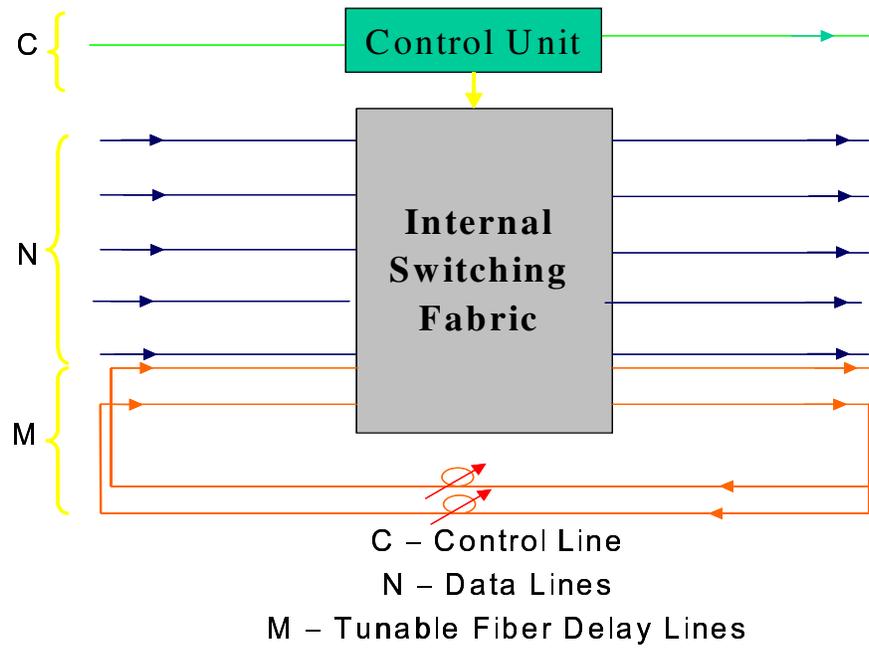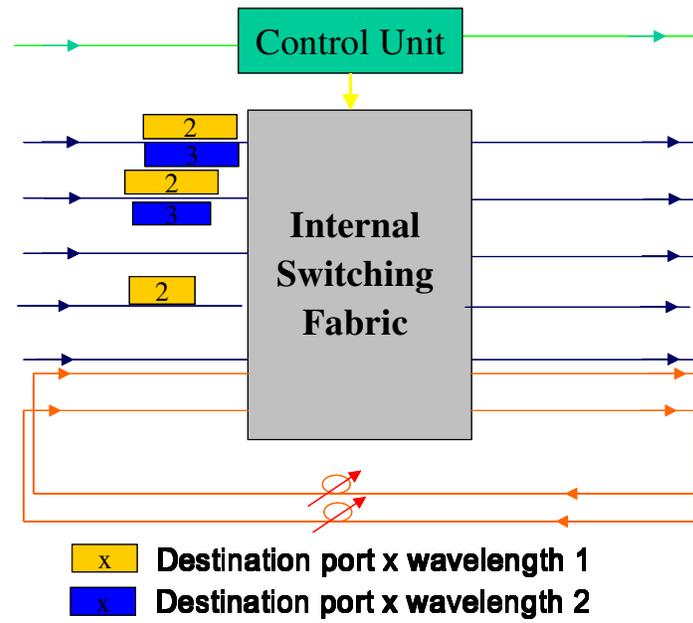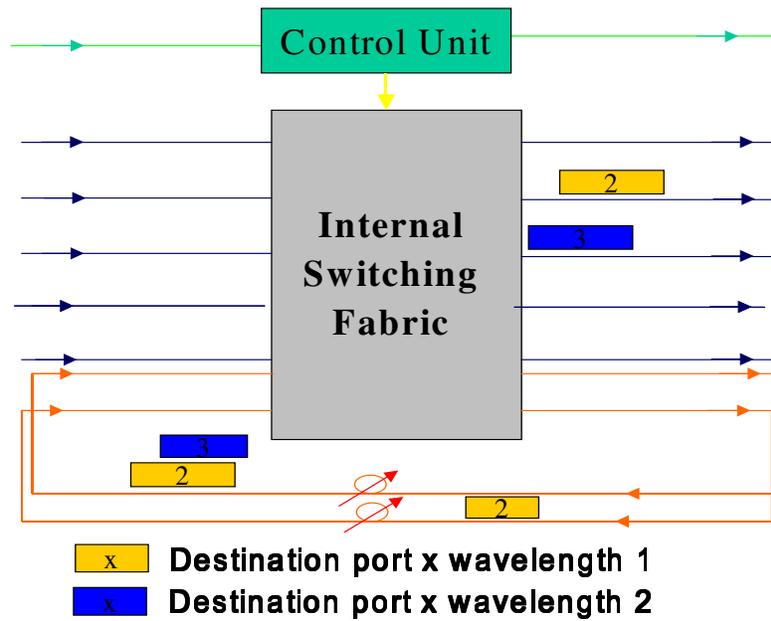
Figure 1

Figure 2

Figure 3
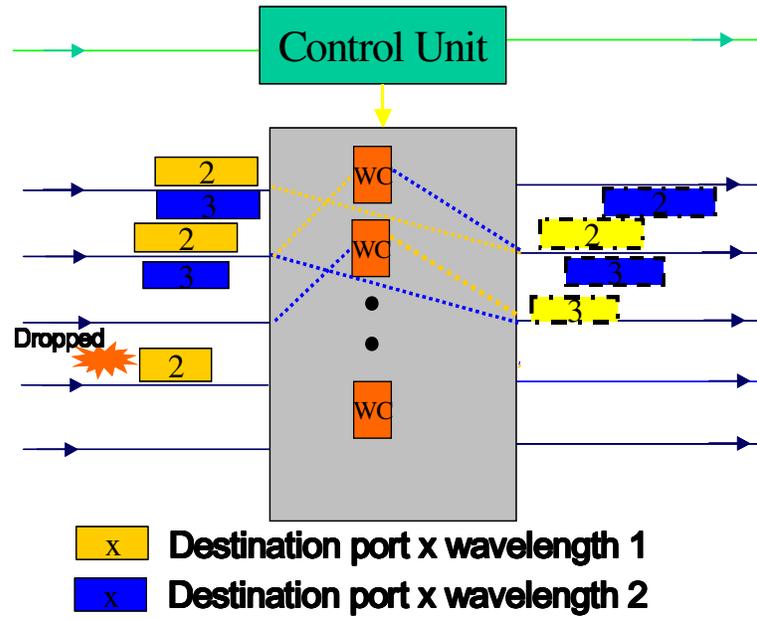
Figure 4
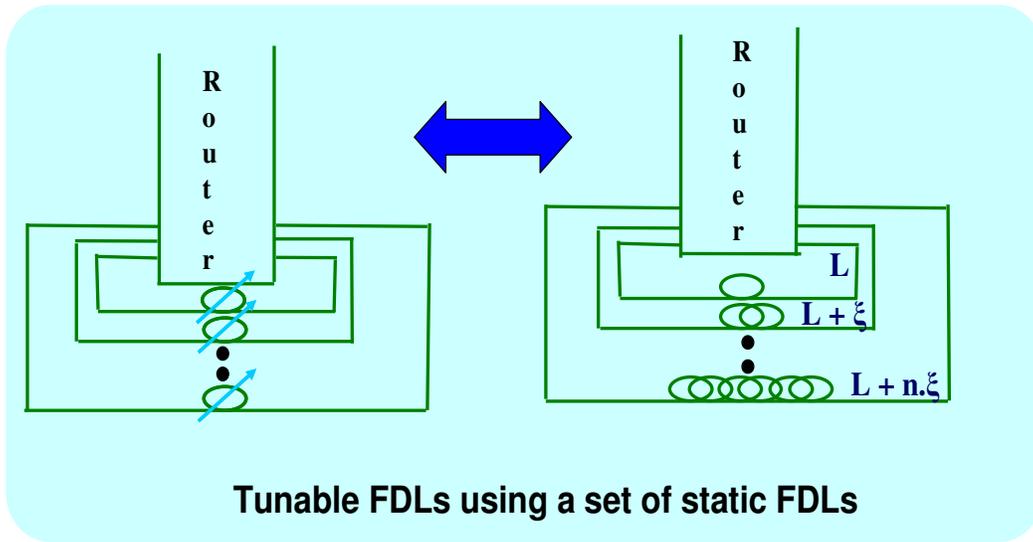
Figure 5

Figure 6

Figure 7
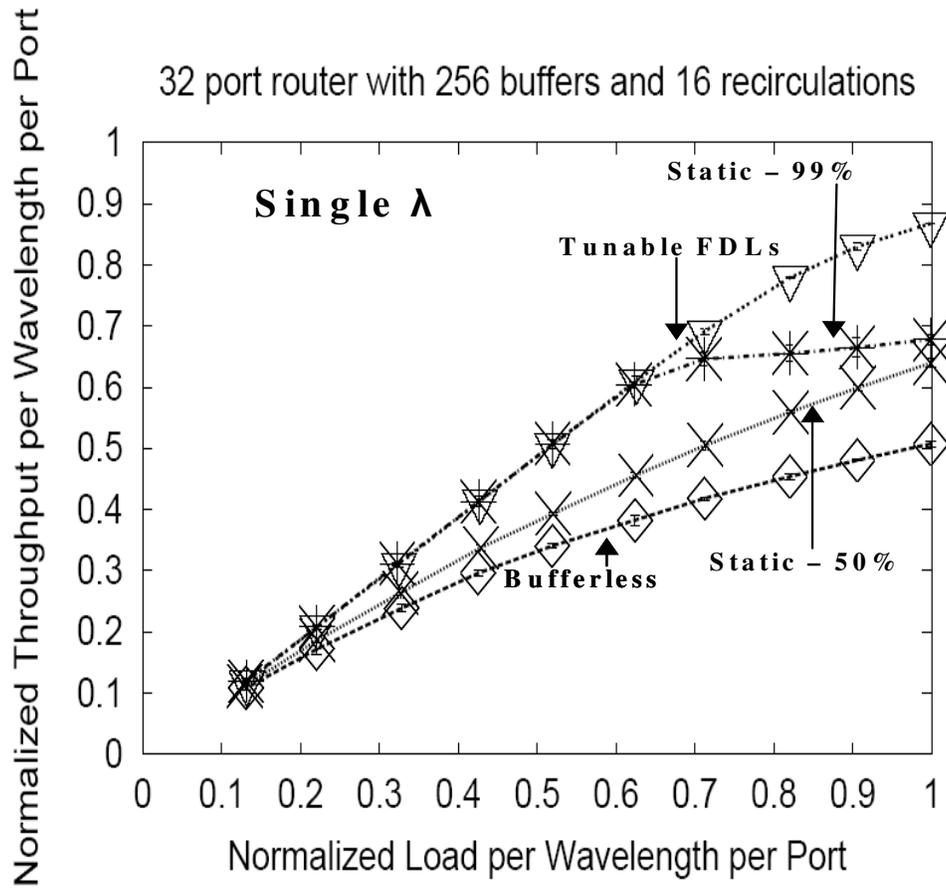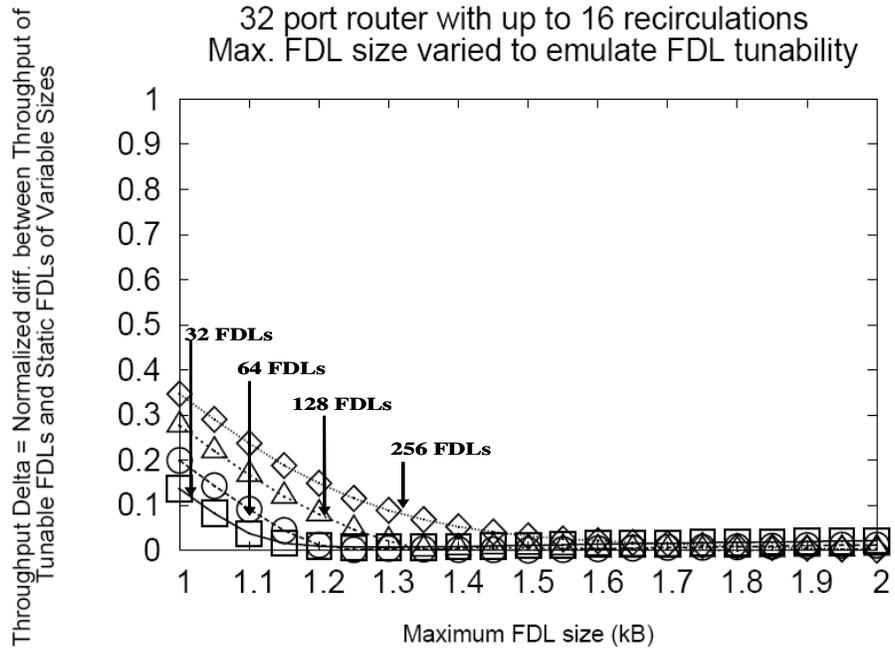
32 port router with up to 16 recirculations
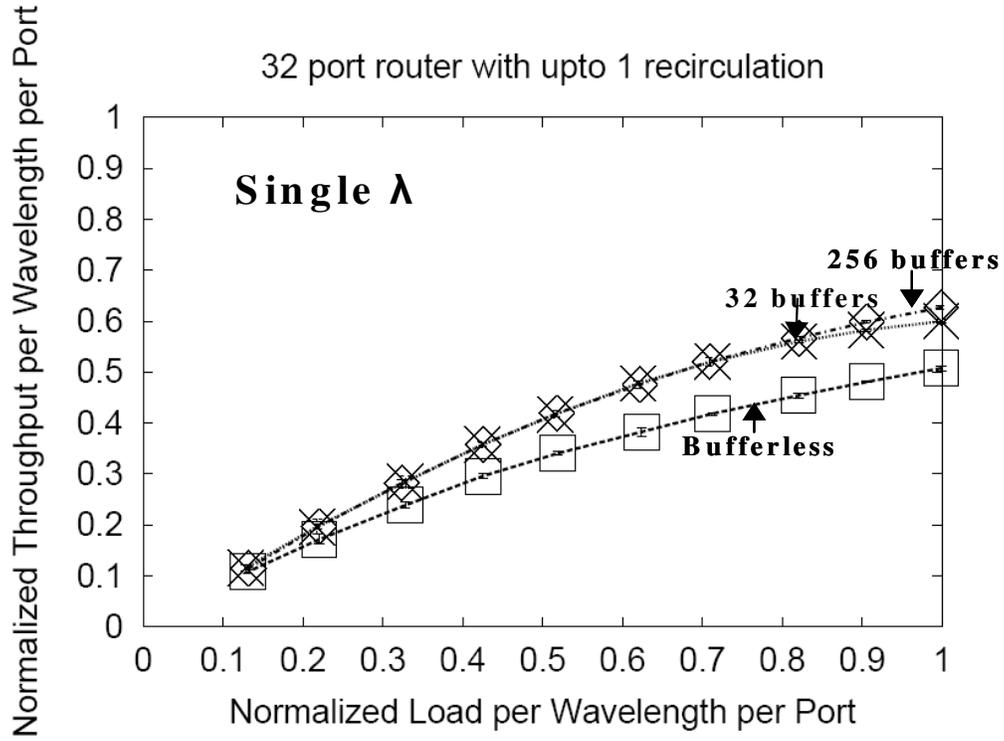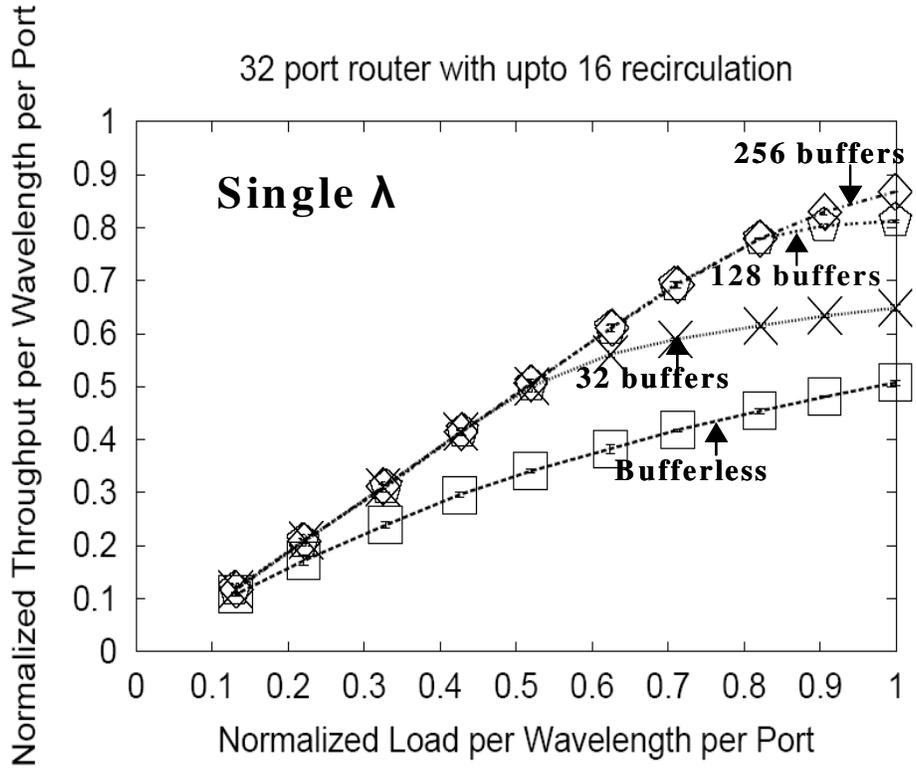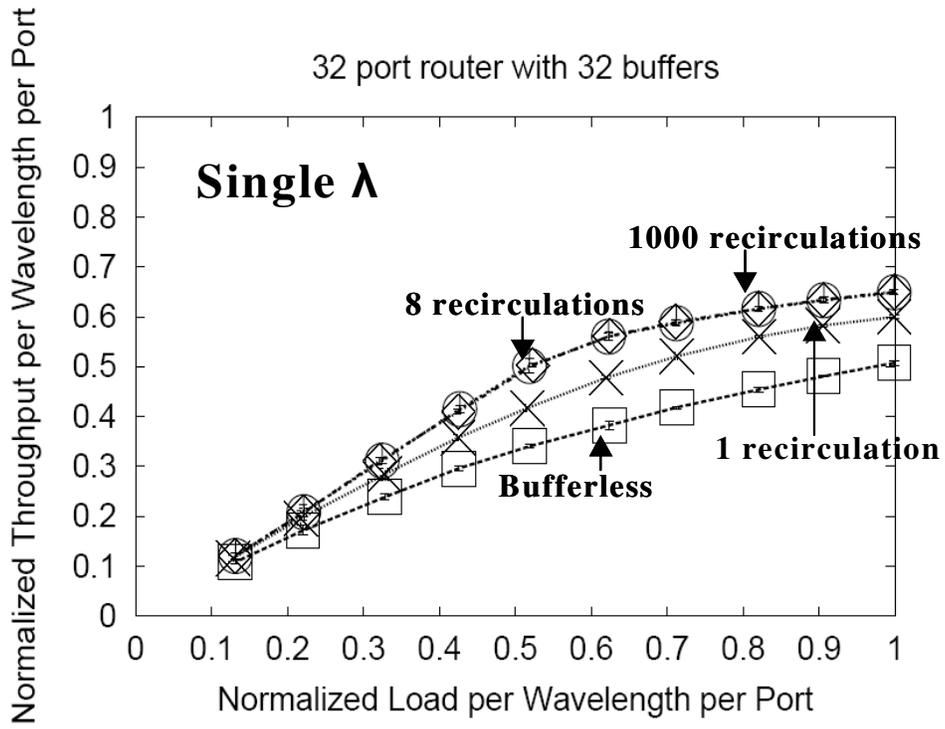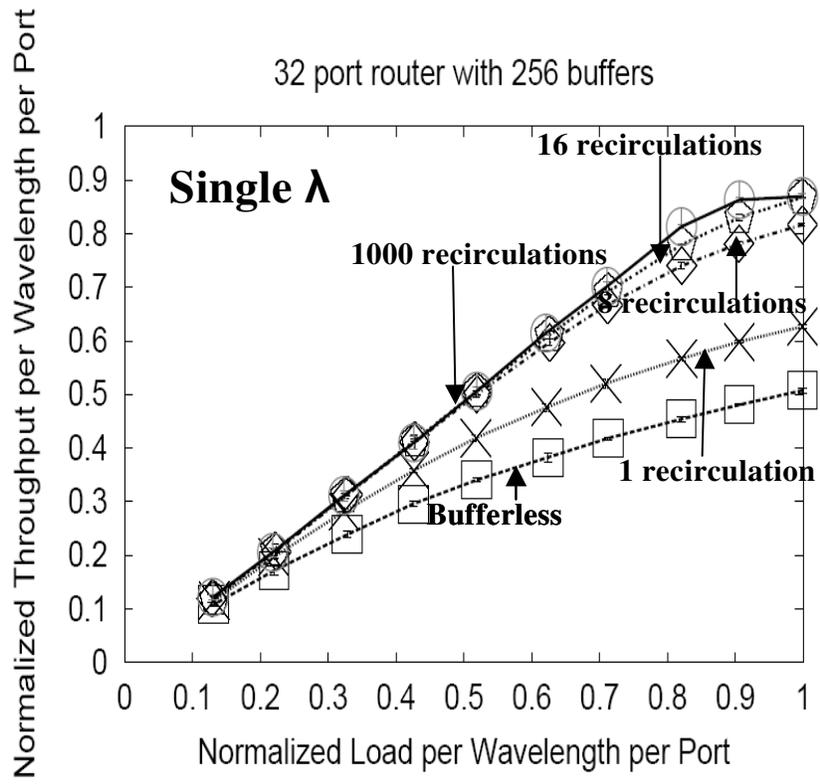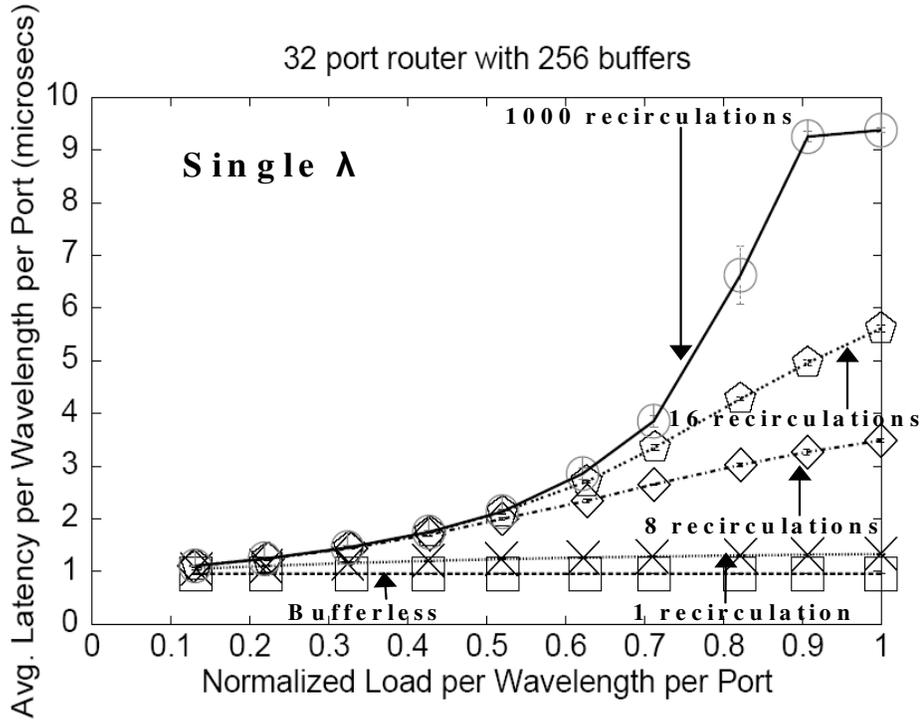Max. FDL size set at throughput delta value equal to 0 (Fig. 7)
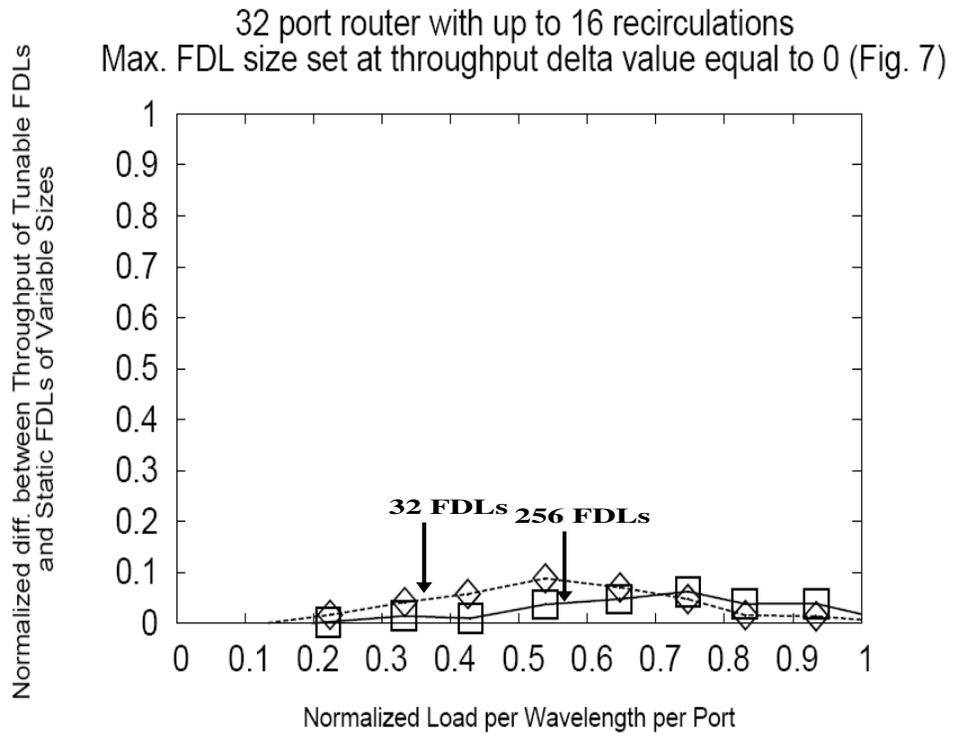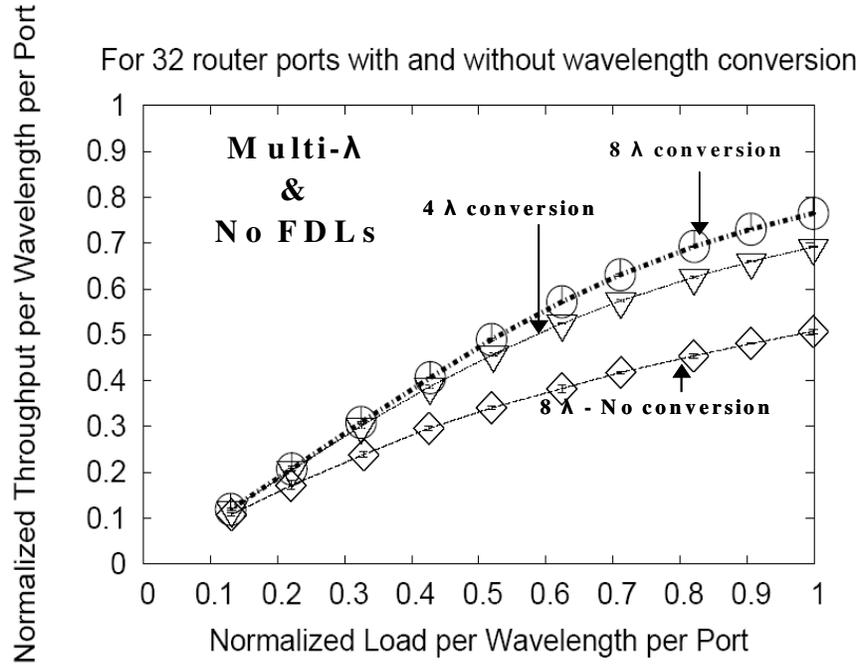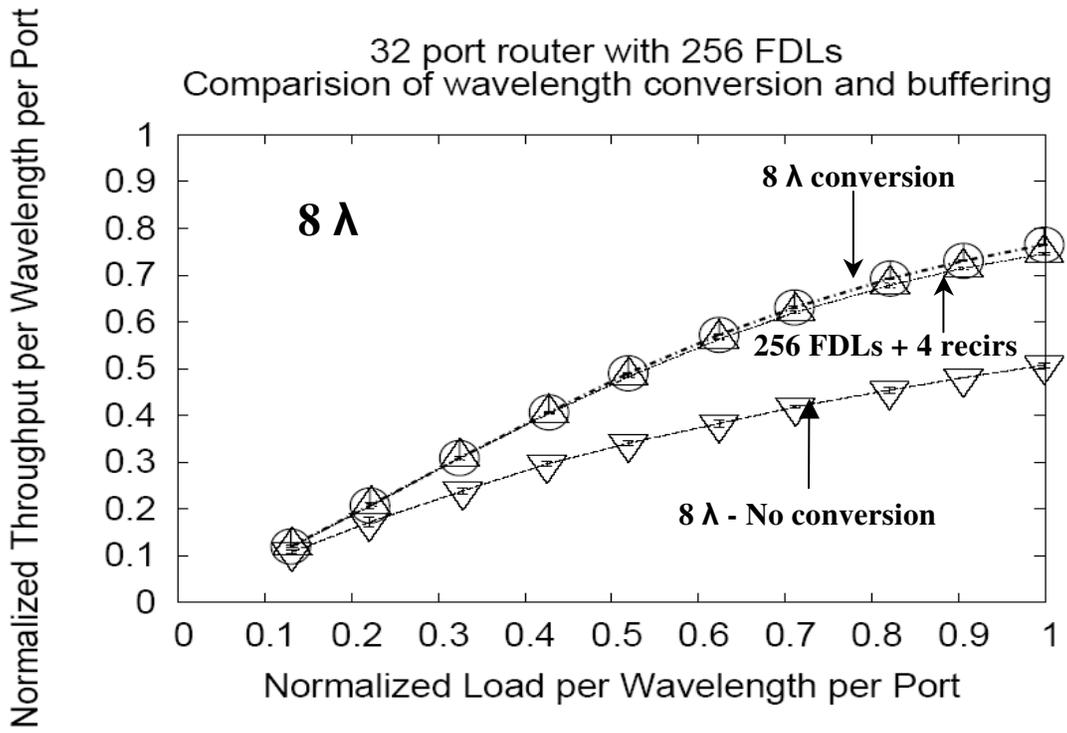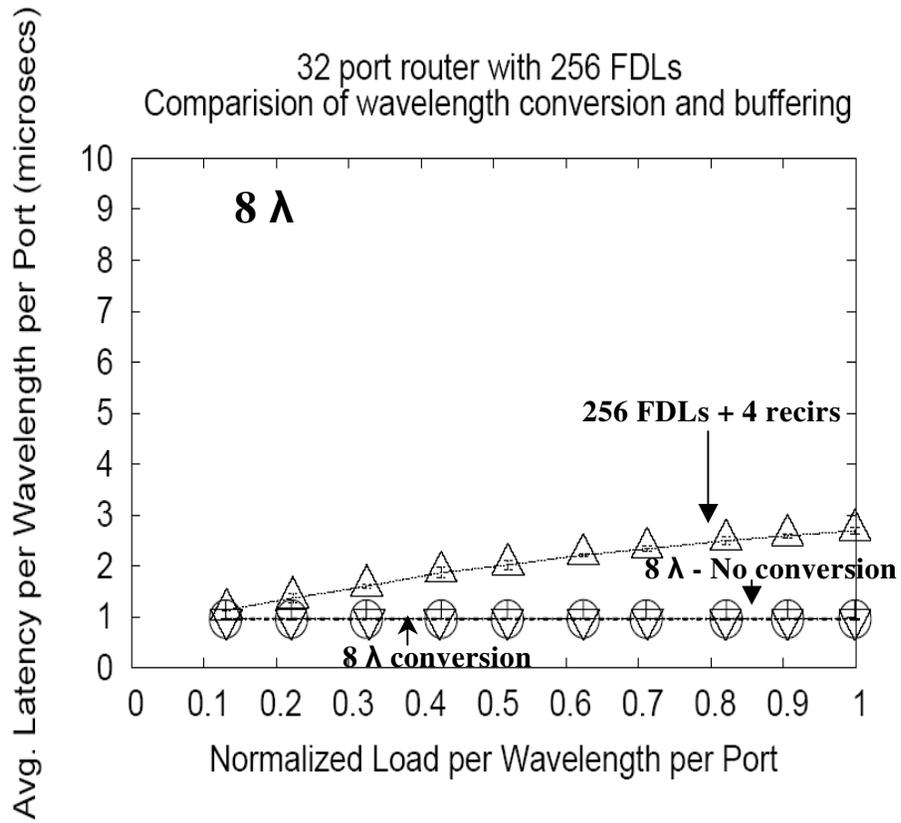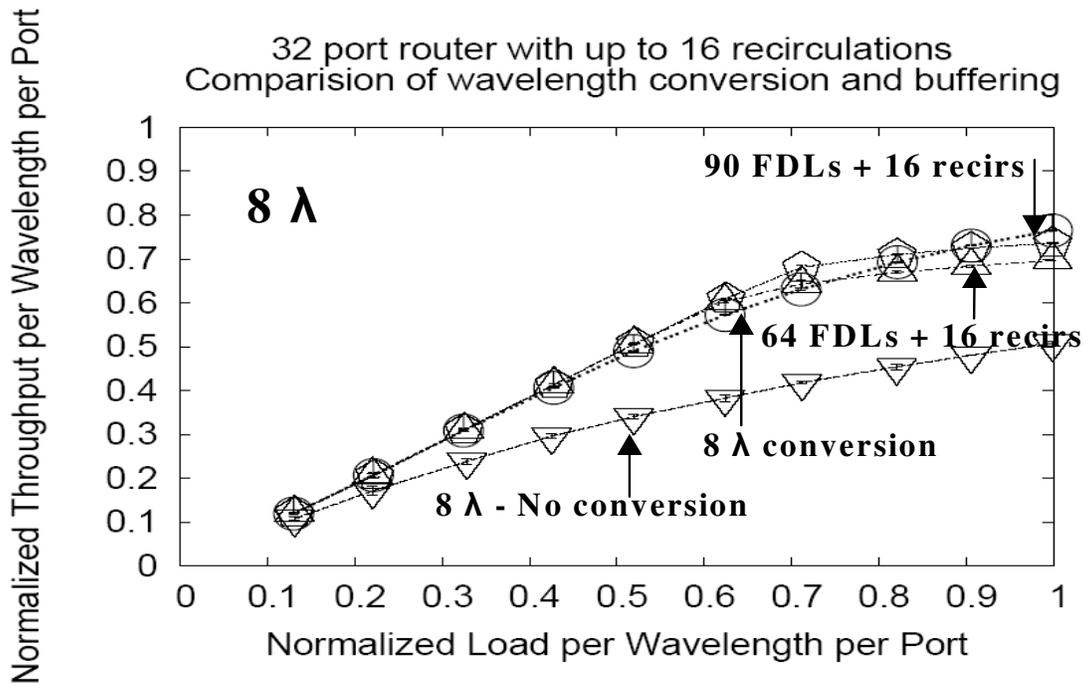
Figure 8

Figure 9

Figure 10

Figure 11

Figure 12

Figure 13

Figure 14

Figure 15

Figure 16

Figure 17

Figure 18

Figure 19