

Copyright 2002 Society of Photo-Optical Instrumentation Engineers. This paper was (will be) published in OptiComm 2002 and is made available as an electronic reprint (preprint) with permission of SPIE. One print or electronic copy may be made for personal use only. Systematic or multiple reproduction, distribution to multiple locations via electronic or other means, duplication of any material in this paper for a fee or for commercial purposes, or modification of the content of the paper are prohibited.

Performance of Converter Pools for Contention Resolution in Optical Burst Switching

Christoph M. Gauger

University of Stuttgart, Institute of Communication Networks and Computer Engineering

ABSTRACT

Optical burst switching (OBS) has recently attracted increasing interest as a network architecture for the future optical Internet. In OBS, efficient contention resolution is a key issue. This paper investigates the performance of shared converter pools for contention resolution in OBS. First, key design parameters for contention resolution in OBS nodes which employ wavelength converters and simple fiber delay line (FDL) buffers are discussed. Then, the performance of a converter pool is evaluated for a Poisson and a self-similar traffic model. Depending on load and node dimensioning, the number of wavelength converters can be reduced by 50–75%. Finally, different converter usage strategies for the combination of a converter pool and an FDL buffer are presented and compared. A strategy that prefers FDL's over converters for contention resolution can reduce burst loss probability for converter pools with a small number of converters.

Keywords: WDM, Optical Burst Switching, Contention Resolution, Wavelength Conversion

1. INTRODUCTION

Optical Burst Switching has been proposed as a new switching paradigm¹ for optical networks and has emerged as a candidate for the optical transport layer of the next generation Internet. Bandwidth and quality of service (QoS) requirements of future Internet services, e. g., voice-over-IP (VoIP) or distributed multimedia applications, as well as the currently exploding number of high-speed Internet access lines motivate a scenario in which optical transport capacity can be provisioned on the time-scale of IP-layer dynamics. In OBS networks (Figure 1), these dynamics can be supported by edge nodes that aggregate traffic and assemble it into variable length optical bursts as well as core nodes that switch these bursts in optics.

OBS is a fast circuit switching technique (FCS) that provides a granularity in between wavelengths and packets but neither mandates the use of optical header processing nor optical buffering, i. e., it requires less complex technology than optical packet switching. A key characteristic is the hybrid approach, in which control information is signaled out of band and processed electronically while data stay in the optical domain at all times. Another key concept of OBS is one-pass reservation, i. e., transmission of a burst is not delayed until an acknowledgment of successful end-to-end path setup is received but is initiated as soon as the burst has been assembled.

Efficient contention resolution in core nodes is essential in order to achieve a low burst blocking probability despite one-pass reservation strategy and statistical multiplexing. So far, all proposals and investigations assumed contention resolution by *full* wavelength conversion. This has been shown to provide low loss probabilities because all wavelength channels of an output line can be shared among all bursts directed towards this output line.^{2–6} This scheme has also been complemented by providing a number of fiber delay lines in an FDL buffer. Even FDL buffers with simple functionality and low technological requirements can improve OBS performance significantly.^{2,7,8} However, as wavelength converters are technologically complex and expensive, their usage should be minimized without significant degradation of performance in an OBS core node. Therefore, this paper concentrates on the performance of an OBS node which employs only *partial* wavelength conversion, i. e., wavelength converters are only available for a limited number of bursts at a time. In this scenario, tunable wavelength converters (TWC) are shared in a converter pool and can be assigned to a wavelength channel in

Correspondence: E-mail: gauger@ind.uni-stuttgart.de; Web: www.ind.uni-stuttgart.de/~gauger, Telephone: +49 711 685 7990, Fax: +49 711 685 7983, Address: IND, Pfaffenwaldring 47, 70569 Stuttgart, Germany

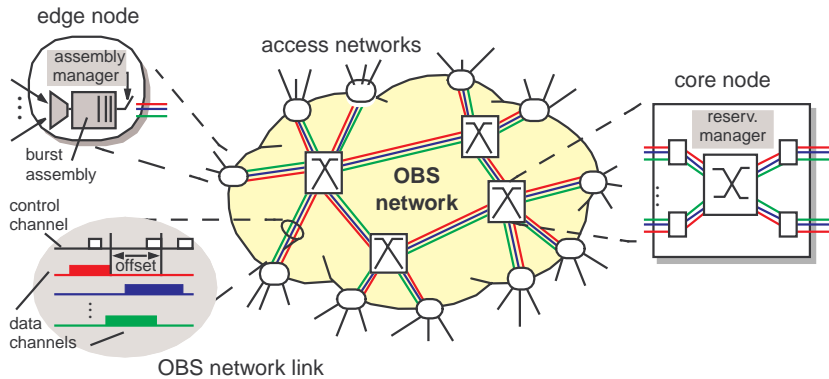


Figure 1: Network scenario for optical burst switching

case of contention. Partial wavelength conversion has been investigated for bufferless optical packet switches under the assumption that all packets have identical length^{9,10} and in the context of optical cross-connects for wavelength routed optical networks.^{11,12} It is reported that depending on node dimensioning, traffic model and converter assignment strategy a substantial number of converters can be saved. However, performance of partial wavelength conversion has not been evaluated under self-similar traffic and in combination with simple FDL buffers which both are interesting for OBS.

The remainder of this paper is structured as follows: In Section 2, options and key design parameters for contention resolution in OBS nodes which employ wavelength converters and simple FDL buffers are discussed. Section 3 describes the node and traffic model used for the performance evaluation. Section 4.1 evaluates the impact of the number of available converters on burst loss probability while Section 4.2 focuses on the impact of traffic models. Finally, the combination of a converter pool and an FDL buffer is studied in Section 4.3.

2. CONTENTION RESOLUTION IN OPTICAL BURST SWITCHING

In an all-optical burst switch, incoming bursts should primarily be sent on the same wavelength to their designated output line. In case of a reservation conflict, i. e., the wavelength on this output line is already reserved, one or a combination of the following three major options for contention resolution can be applied.

- *Wavelength domain:* By means of wavelength conversion, a burst can be sent on a different wavelength channel of the designated output line.
- *Time domain:* By applying an FDL buffer, a burst can be delayed until the contention situation is resolved. In contrast to buffers in the electronic domain, FDL's only provide a fixed delay and data leave the FDL in the same order in which they entered.
- *Space domain:* In deflection routing, a burst is sent to a different output line of the node and consequently on a different route towards its destination node. Space domain can be exploited differently in case several fibers are attached to an output line. In this case, a burst can also be transmitted on a different fiber of the designated output line without wavelength conversion.

In contrast to using a wavelength converter, an FDL buffer or a different parallel fiber deflection routing does not resolve contention locally in a single node but reroutes over-load traffic to neighboring nodes. Therefore, it depends heavily on network topology and routing strategy. Investigations comparing the efficiency of different contention resolution strategies have shown that it results in only limited improvement.³ As design and dimensioning of FDL buffers for OBS nodes have been investigated in previous work⁷ this paper focuses on partial wavelength conversion and combination of converters and FDL buffers.

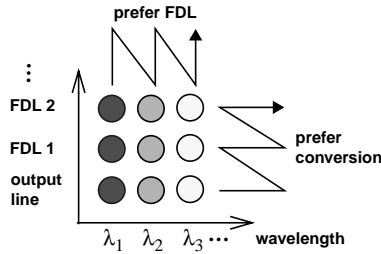


Figure 2: Usage strategies for wavelength converters and FDL's.

Performance of contention resolution in an OBS node is influenced by several parameters. Following list provides an overview of the key design parameters and introduces strategies which will be investigated in Section 4:

- *Node functionality:* This relates to the availability and architecture of wavelength converters and FDL buffers in the node. Nodes can provide no, full or partial wavelength conversion. In the latter case, converters are organized in a converter pool and can be used in case of a reservation conflict. Converters can be either placed at the input or at the output of the switch. A node can employ no buffers at all or FDL buffers of different complexity and architecture, e. g., feedback or feed-forward structure. FDL buffers can be shared per input line or per output line respectively or be shared per node.

Node functionality is closely related to the physical realization of the node, e. g., some node architectures trade-off wavelength conversion by the number of gates in the switch matrix and vice versa.^{10,13} Therefore, in order to assess fundamental contention resolution strategies in OBS nodes from a performance point of view, an abstract node model is chosen (c. f. Section 3).

- *Dimensioning and physical constraints:* In case of conversion, dimensioning relates to the number of converters in a converter pool. Wavelength converters can either convert from any input wavelength to any output wavelength or can only convert within a limited conversion range.³ FDL buffers are characterized by the number and individual delays of the FDL's in the buffer as well as by the number of wavelengths per FDL.
- *Usage strategy:* If wavelength conversion and FDL buffers are applied in combination an additional degree of freedom exists. The arrows in Figure 2 illustrate the order in which the two strategies seek reservation. In the *prefer conversion* strategy wavelength converters are applied preferably and only if they cannot resolve the conflict FDL buffers are used. In the *prefer FDL* strategy FDL buffers are applied preferably and only if they cannot resolve the conflict wavelength converters are used. While the former strategy extensively uses wavelength converters the latter can be expected to save converters.
- *Reservation strategy:* In case a burst uses a wavelength converter or an FDL buffer for contention resolution these resources also have to be reserved. Several strategies which have been proposed for output channel reservation can be applied.⁴ Regarding FDL buffers, an additional degree of freedom exists during the reservation process. The output channel for the time when the burst will leave the FDL buffer can be either reserved before entering the FDL or when leaving the FDL.⁷

Apart from the architecture, functionality and dimensioning of converters and FDL buffers, node dimensioning as well as traffic behavior have a significant impact on performance. Concerning node dimensioning, the number of input and output lines as well as the number of wavelength channels per output line are key parameters. Regarding traffic, the offered load, the arrival process of bursts as well as the distribution of burst lengths have to be considered.

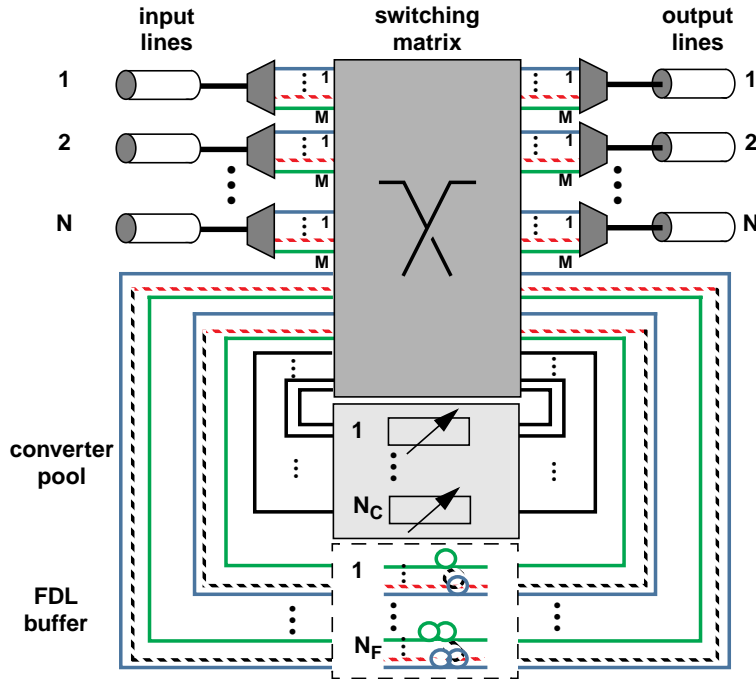


Figure 3: Model of an OBS node with a shared converter pool and an (optional) feedback FDL buffer

3. MODELING APPROACH

Performance of converter pools for OBS nodes is evaluated by simulation studies that are based on an abstract node model. This model is depicted in Figure 3 for an OBS node with a converter pool and an optional FDL buffer. The node has a non-blocking switching matrix, N single fiber input and output lines and M wavelength channels per fiber. The converter pool comprises N_C tunable wavelength converters which can be shared among all wavelength channels of this node and which can convert over the entire range of wavelengths. In order to compare performance of nodes with different dimensioning, the figure *conversion ratio*, r_c , is introduced: It is defined as the ratio of converters in the pool, N_C , and the total number of *output* wavelength channels that share this converter pool, $r_c = N_C / (N \cdot M)$. The FDL feedback buffer has N_F fiber delay lines with basic delay b and linearly increasing FDL delays, i. e., FDL i has delay ib , $i = 1, \dots, N_F$. FDL's employ WDM with M wavelength channels per FDL. No further architectural details of the optical switching matrix, the converter pool or the FDL buffer have been included in the model, because they can be very different for specific node realizations.

Wavelength channels on an output line and in an FDL as well as converters are reserved according to the just-enough-time (JET) reservation strategy.^{8,14} It considers the exact predetermined start and end times of each burst for reservation, which allows most efficient utilization of resources. In case an FDL buffer is applied, resources for buffering the burst in the FDL, for later transmission on the output line and potentially for wavelength conversion have to be reserved in a coordinated way. Here, it is assumed that a burst uses the same wavelength in the FDL buffer and on the output line, i. e., a potential wavelength conversion has to take place before buffering. Also, all resources are reserved at the same time, which was characterized as *PreRes* policy before.⁷ In JET, introducing an additional quality of service (QoS) offset time between reservation request and data burst for a prioritized service class leads to an early reservation of resources and consequently to a higher probability of successful reservation. Therefore, reserving the output line according to the *PreRes*, i. e., before entering the FDL, results in a prioritization and a high probability of successful reservation on the output line.

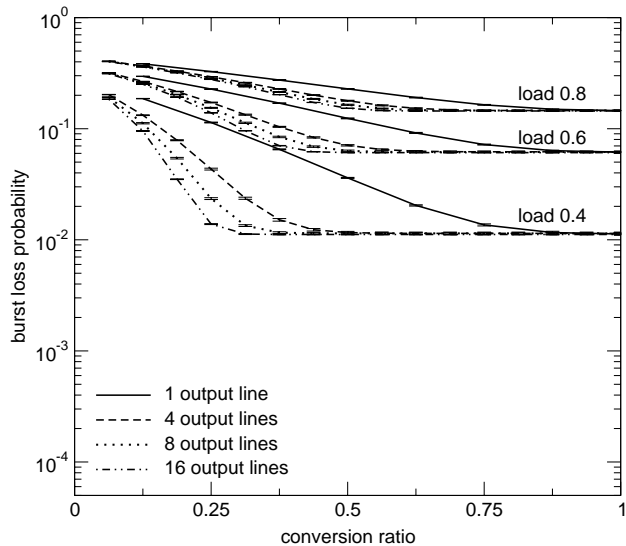


Figure 4. Burst loss probability vs. conversion ratio for Poisson traffic ($M = 8$)

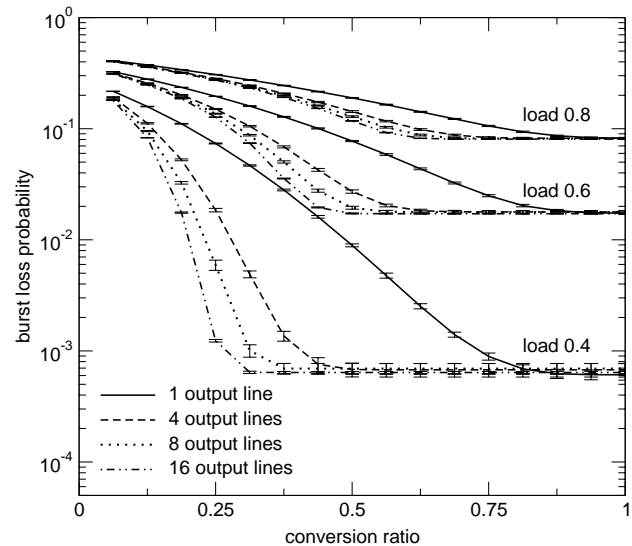


Figure 5. Burst loss probability vs. conversion ratio for Poisson traffic ($M = 16$)

For the studies in this paper, two models for the arrival process of bursts are applied. In the first model, bursts arrive according to a Poisson process. The second model is a *self-similar* traffic model and motivated by the rapid growth and dominance of Internet traffic which has been shown to exhibit self-similar behavior. In the future, when rather static traffic management functionality provided by ATM or SDH today is no longer used but IP is transported directly over WDM with only a thin adaptation layer in between, dynamics and long-range dependent characteristics of Internet traffic¹⁵ will eventually reach the optical layer. As illustrated in Figure 1, traffic originates from a large number of sources in metro and access networks. Superposition of a large number of ON/OFF sources with heavy-tailed ON or OFF phases has been proposed as a model for long-range dependent network traffic.¹⁶ An M/Pareto traffic model is obtained for an infinite number of sources with Pareto distributed ON phases.¹⁷ In an OBS edge node, incoming traffic is aggregated in the electronic domain and bursts are assembled dynamically, e. g., based on size or time strategies.^{18,19} Large traffic volumes directed to the same egress edge node will form trains of bursts. In the context of OBS, the M/Pareto model is used to generate trains of optical bursts with mean train interarrival time $E[T_A]$. The number of optical bursts, X_B , in a train of bursts follows a Pareto distribution with mean $E[X_B]$ and Hurst parameter, H . During a train of bursts, optical bursts are generated with constant interarrival times, d , i. e. the burst arrival rate is $1/d$. Instead of $E[T_A]$ and d , the model can be characterized by the mean burst arrival rate $m = E[X_B]/E[T_A]$ and the ratio of mean burst arrival rate and burst arrival rate during a train of bursts: $\beta = m \cdot d$.

4. PERFORMANCE EVALUATION AND RESULTS

Performance evaluation has been performed by discrete event simulation of the model described above. The tool is realized in C++ and based on a simulation library²⁰ that provides support for simulation control, traffic generation as well as statistical evaluation. In the simulation, bursts are generated according to either a Poisson or a self-similar model as discussed in the previous section. Burst length is assumed to be negative-exponentially distributed for the Poisson model and constant for the self-similar model with mean 100 kbits which leads to a mean transmission time $h = 10 \mu\text{s}$ on a 10 Gbps line. The term *load* refers to offered load per wavelength. Destination of bursts is uniformly distributed over all N output lines. Also, selection of the wavelength on which a burst arrives to the node follows a uniform distribution. Eight wavelength channels per output line or per FDL are used except for Figure 5 in which there are 16 wavelengths per output line. The burst loss probability is used as the key performance metric. Unless stated differently, graphs obtained for the Poisson arrival process include 95%-confidence intervals based on the batch simulation method. However, as fundamental assumptions

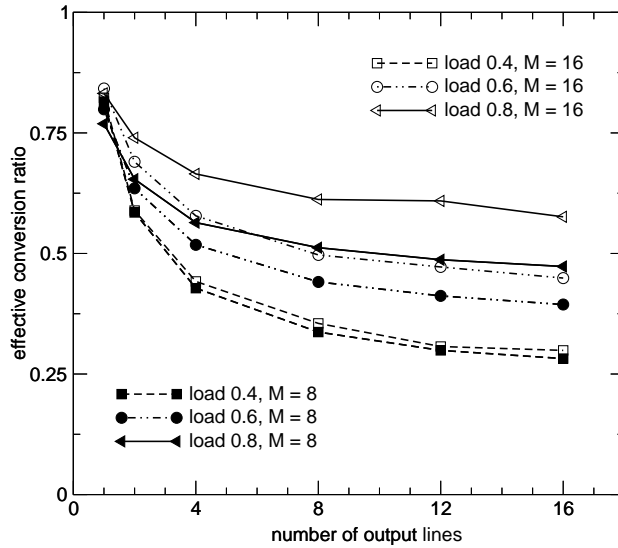


Figure 6: Effective conversion ratio vs. number of output lines

for the analysis of confidence intervals do not hold for long-range dependent traffic graphs for the self-similar arrival process show no confidence intervals. In this case, a larger number of bursts has been simulated.

4.1. Impact of Conversion Ratio

In Figures 4 and 5, burst loss probability is shown versus the conversion ratio, $r_c = N_C / (N \cdot M)$ for different number of output lines and values of load for scenarios with 8 and 16 wavelength channels. In both scenarios, loss probability decreases steadily with increasing number of converters until it reaches a lower boundary. In case of 8 wavelengths and load 0.4, this boundary is reached at conversion ratios of approximately 0.25–0.5, i. e., between 50% and 75% of the maximum number of converters can be saved. For greater values of load the lower boundary is only reached for greater conversion ratios. When comparing respective curves for 8 and 16 wavelength channels it can be seen that the lower boundary is reached at approximately identical conversion ratios. However, loss probability decreases more significantly for 16 wavelength channels when increasing the number of converters. This can be explained by the multiplexing gain of the larger bundle of channels. From Figures 4 and 5 it can be seen that for a larger number of output lines the lower boundary is reached for a smaller conversion ratio. Again, this can be explained by the multiplexing gain which results from the larger number of converters in the pool. The curve for a single output line is included for reference and corresponds to a converter pool which is only shared for a single output.

In order to quantify the dependence on the number of output lines and to obtain values for the number of converters needed in order not to exceed a certain penalty in loss probability, the *effective conversion ratio* is introduced. It is defined as the conversion ratio which results in a 10% higher loss probability than the case of full conversion. Figure 6 depicts the effective conversion ratio versus the number of output lines for different values of load and 8 as well as 16 wavelengths. No confidence intervals are provided as data have been extracted from graphs. The effective conversion ratio decreases fast for a smaller number of output lines and much slower for a larger number of output lines. From approximately 8 ports on and a load of up to 0.6, the effective conversion ratio is less than 0.5 and can be as low as 0.25 which results in substantial savings of wavelength converters. Again, the curves for different numbers of wavelengths have similar shape but for more wavelengths the effective conversion ratio is greater which is due to the higher carried traffic, i. e. occupancy of the output lines.

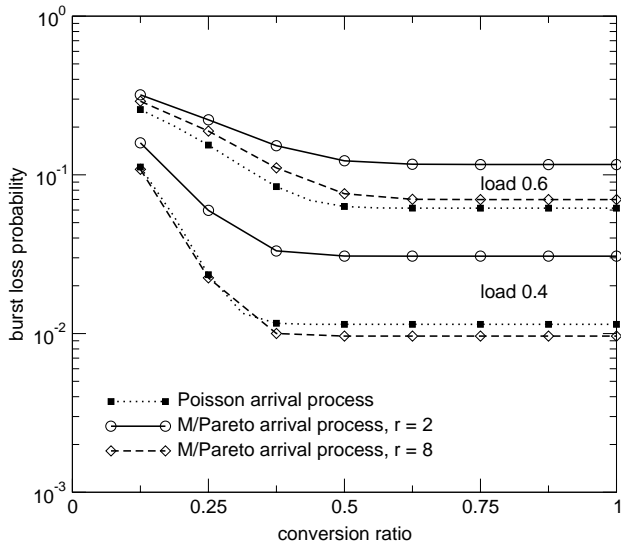


Figure 7. Burst loss probability vs. conversion ratio for self-similar traffic ($N = 8$, $M = 8$, $E[X_B] = 8$)

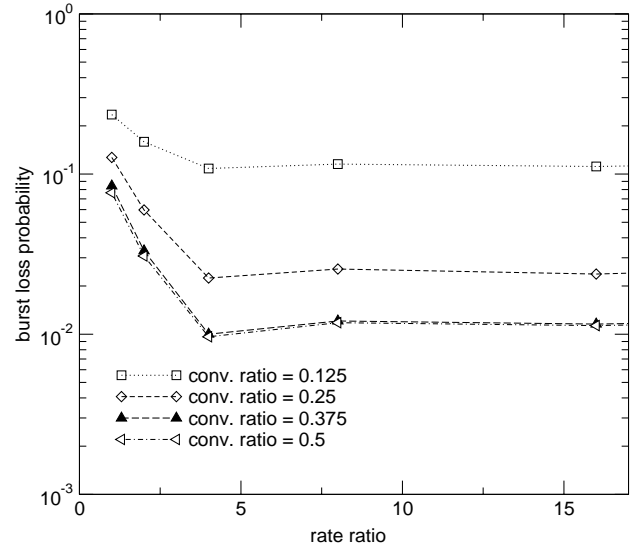


Figure 8. Impact of rate ratio on burst loss probability (load 0.4, $N = 8$, $M = 8$, $E[X_B] = 8$)

4.2. Impact of Traffic Characteristics

Figure 7 compares the burst loss probability for the Poisson traffic model and for the self-similar traffic model for two values of the rate ratio, β . The number of output lines is 8 and in case of the self-similar traffic model there are in average $E[X_B] = 8$ bursts in a train of bursts and the Hurst parameter is 0.6. The M/Pareto arrival process with rate ratio $\beta = 8$ results in an almost identical burst loss probability as the Poisson process. In case of rate ratio $\beta = 2$, the lower boundary is reached for the same conversion ratio but the burst loss probability is significantly higher. As highly dynamic traffic is a key motivation for OBS, the impact of parameters of the self-similar traffic model on burst loss probability is studied. Results which are not included in figures show that varying the Hurst parameter between 0.6 and 0.7 and the mean number of bursts in a train of bursts have hardly any impact.

For an offered load of 0.4, Figure 8 depicts this dependence of the burst loss probability on the rate ratio for different values of the conversion ratio. Higher conversion ratios do not yield better performance and are therefore not included in the graph. It can be seen that for a small rate ratio loss probability is significantly higher than for a large rate ratio. Also, the improvement of loss probability due to wavelength conversion is smaller for a small rate ratio. It can be concluded that the higher the rate of individual streams compared to the aggregate rate the higher the loss probability and the less efficient the wavelength converter pool.

4.3. Impact of Converter Usage Strategies

For a combination of a wavelength converter pool and an FDL buffer, Figure 9 compares the strategies *prefer conversion* and *prefer FDL* with respect to burst loss probability. Also, the case of no buffers is included for reference. A node with $N = 4$ output lines, an FDL buffer with $N_F = 2$ FDL's and $N_C = 8$ wavelengths per fiber and FDL are assumed. Note that if an FDL buffer is applied the maximum number of converters in the converter pool is the sum of wavelength and FDL channels. However, in order to receive comparable results, the conversion, r_c , is still calculated with respect to $N \cdot M$ and could therefore exceed 1.

First, it can be seen that the combination of a shared converter pool and an FDL buffer has a significantly reduced burst loss probability compared to the case without an FDL buffer. For a conversion ratio less than 0.5 and for load 0.4, the strategy *prefer FDL* leads to a lower loss probability than the *prefer conversion* strategy. For a large number of converters or higher values of load the *prefer conversion* strategy performs equally good or better. It can be seen that the strategy *prefer FDL* which tries to minimize converter usage can better utilize

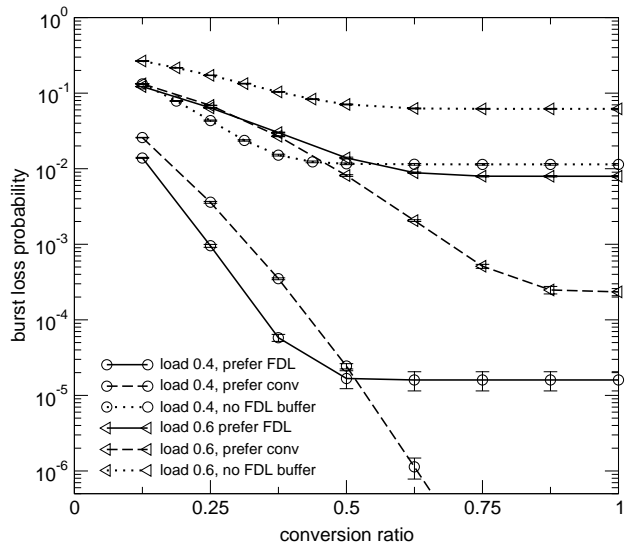


Figure 9. Impact of converter usage strategies on burst loss probability ($M = 8$, $N = 4$, 2 FDL's)

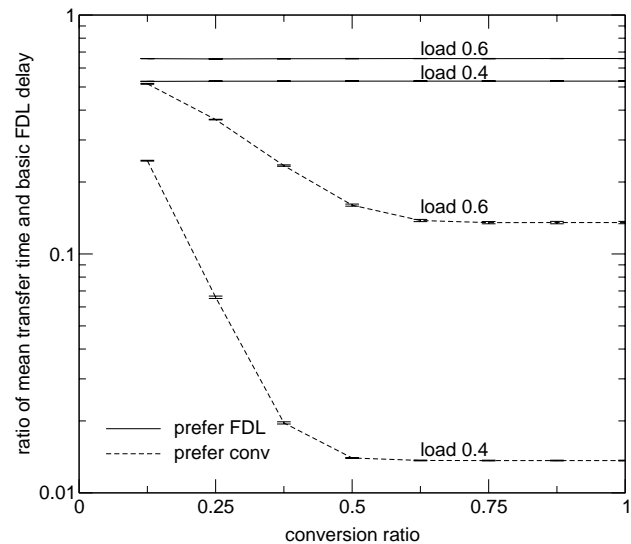


Figure 10. Impact of converter usage strategies on mean transfer time ($M = 8$, $N = 4$, 2 FDL's)

a reduced number of converters. *Prefer conversion* extensively uses converters which can be seen from the fact that the lower boundary is only reached for larger values of the conversion ratio.

The lower loss probability of the strategy *prefer FDL* comes at the cost of increased transfer delay. Figure 10 compares both strategies for the same set of parameters with respect to the transfer time normalized by the basic FDL delay which is chosen twice the mean burst transmission time. In case of the *prefer FDL* strategy, the transfer time stays constant while the transfer time in case of the *prefer conversion* strategy decreases when increasing the number of converters until a lower boundary is reached. This shows that for the *prefer conversion* strategy FDL usage decreases as the number of converters increases. In this scenario, the mean transfer delay is in the order of few tens of microseconds and can be neglected compared to propagation delay in long-haul networks or assembly delay in edge nodes. Increasing the FDL delay within the physical limits can further reduce the loss for low loads but also leads to increased transfer times.

5. CONCLUSION AND OUTLOOK

In this paper, key design parameters for application of converter pools and FDL buffers for contention resolution in optical burst switching nodes are discussed. The influence of number of converters in the converter pool, number of wavelengths per output line and number of output lines on burst loss probability is evaluated. Investigations are performed for a Poisson and a self-similar traffic model including a study on the impact of traffic characteristics of the latter model. Even for a rather small number of output lines a substantial reduction in the number of wavelength converters can be achieved. For the self-similar traffic model, a large ratio of mean burst arrival rate and burst arrival rate during a train of bursts is shown to yield a high burst loss probability. For the combination of a converter pool and an FDL buffer two usage strategies are introduced and investigated. A strategy that prefers FDL's over converters for contention resolution is shown to achieve improved loss performance for converter pools with a small number of converters. Further work could include the investigation of different reservation strategies for converter pools with or without FDL buffers. Also, an integrated investigation including realization issues for different node architectures could be performed.

ACKNOWLEDGMENTS

The author would like to thank Stefan Bodamer, Hao Feng, Erwin Patzak and Jürgen Saniter for invaluable discussions and all reviewers for their comments. This work was funded within the TransiNet project by the German Bundesministerium für Bildung und Forschung under contract No. 01AK020C (www.transinet.de).

REFERENCES

1. C. Qiao and M. Yoo, "Optical burst switching (OBS)—a new paradigm for an optical Internet," *Journal of High Speed Networks* **8**, pp. 69–84, January 1999.
2. M. Yoo and C. Qiao, "QoS performance in IP over WDM networks," *IEEE Journal of Selected Areas in Communications* **18**, pp. 2062–2071, October 2000.
3. S. Yao, B. Mukherjee, S. J. B. Yoo, and S. Dixit, "All-optical packet-switched networks: a study of contention-resolution schemes in an irregular mesh network with variable-sized packets," in *Proceedings of OptiComm 2000*, (Dallas), October 2000.
4. K. Dolzer, C. Gauger, J. Späth, and S. Bodamer, "Evaluation of reservation mechanisms for optical burst switching," *AEÜ International Journal of Electronics and Communications* **55**, pp. 18–26, January 2001.
5. K. Dolzer and C. Gauger, "On burst assembly in optical burst switching networks—a performance evaluation of just-enough-time," in *Proceedings of the 17th International Teletraffic Congress (ITC 17)*, pp. 149–160, (Salvador, Brazil), December 2001.
6. J. S. Turner, "Terabit burst switching," *Journal of High Speed Networks* **8**, pp. 3–16, January 1999.
7. C. M. Gauger, "Dimensioning of FDL buffers for optical burst switching nodes," in *Proceedings of the 5th IFIP Optical Network Design and Modeling Conference (ONDM 2002)*, (Torino), February 2002.
8. Y. Xiong, M. Vanderhoute, and C. C. Cankaya, "Control architecture in optical burst-switched WDM networks," *IEEE Journal of Selected Areas in Communications* **18**, pp. 1838–1851, October 2000.
9. V. Eramo and M. Listanti, "Packet loss in a bufferless optical WDM switch employing shared tunable wavelength converters," *Journal of Lightwave Technology* **18**, pp. 1818–1833, December 2000.
10. S. L. Danielsen, B. Mikkelsen, C. Joergensen, and T. Durhuus, "WDM packet switch architectures and analysis of the influence of tuneable wavelength converters on the performance," *Journal of Lightwave Technology* **15**, pp. 219–227, February 1997.
11. J. Späth and S. Bodamer, "Routing of dynamic poisson and non-poisson traffic in WDM networks with limited wavelength conversion," in *Proceedings of the 24th European Conference on Optical Communication (ECOC 1998): Regular and Invited Papers*, (Madrid), September 1998.
12. J. Späth, "Resource allocation for dynamic routing in WDM networks," in *Proceedings of OptiComm 1999*, (Boston), September 1999.
13. H. Feng, E. Patzak, and J. Saniter, "Physikalische Grenzen von "Broadcast and Select"-Schaltknoten für Optical Burst Switching (Physical limitations of broadcast and select switching nodes for optical burst switching)," in *Proceedings of the 3rd ITG Workshop on Photonic Networks*, (Leipzig), April 2002.
14. M. Yoo, M. Jeong, and C. Qiao, "A high speed protocol for bursty traffic in optical networks," in *3rd SPIE Conference on All-Optical Communication Systems*, pp. 79–90, (Dallas), November 1997.
15. M. E. Crovella and A. Bestavros, "Self-similarity in World Wide Web traffic: evidence and possible causes," *IEEE/ACM Transactions on Networking* **5**, pp. 835–846, December 1997.
16. W. Willinger, M. S. Taqqu, R. Sherman, and D. V. Wilson, "Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level," *IEEE/ACM Transactions on Networking* **5**, pp. 71–86, February 1997.
17. T. D. Neame, M. Zukerman, and R. G. Addie, "A practical approach for multimedia traffic modeling," in *Proceedings of the 5th IFIP TC6 International Conference on Broadband Communications (BC 1999)*, (Hong Kong), November 1999.
18. A. Ge, F. Callegati, and L. S. Tamil, "On optical burst switching and self-similar traffic," *IEEE Communications Letter* **4**, March 2000.
19. K. Dolzer, "Assured horizon—a new combined framework for burst assembly and reservation in optical burst switched networks," in *Proceedings of 7th European Conference on Networks and Optical Communications (NOC 2002)*, (Darmstadt, Germany), June 2002.
20. Institute of Communication Networks and Computer Engineering, University of Stuttgart, "The IND simulation library," 2002. www.ind.uni-stuttgart.de/INDSimLib.