

VARIABLE-SIZE PACKETS IN SLOTTED WDM RING NETWORKS*

Andrea Bianco, Massimiliano Bonsignori, Emilio Leonardi, Fabio Neri

Dipartimento di Elettronica

{bianco, leonardi, neri}@polito.it, bonsignori@mail.tlc.polito.it

Abstract Slotted WDM rings are considered to implement high-speed all-optical metro networks. The usual approach to transfer variable-size data units, such as IP datagrams, is to segment them into cells, and transfer each cell independently in the WDM ring. This approach entails complexity due to segmentation/reassembly procedures. We investigate in this paper architectures and control strategies to guarantee that all segments of a variable-size packet are sent contiguously on the ring, thereby drastically reducing reassembly complexity. Our proposal is based on register insertion techniques exploiting optical delay lines. A number of schemes are proposed, and compared by means of simulation, showing how performance impairments can be traded with hardware complexity.

1. Introduction

All-optical networking appears today as the only viable solution to keep pace with the constantly increasing need for bandwidth, mainly because the photonic domain allows bit rates extremely higher than those affordable in the electronic domain. Several architectures of all-optical networks have been studied and developed in the past (see [18, 19]); among these, all-optical rings with multiple WDM channels, also called *Multirings*, have been proposed in recent years to implement high-capacity, single-hop packet networks for the local and metro environments. In these systems, access is controlled by some distributed algorithm, and it is slotted, to ease the allocation of transmission resources and synchronization. Given the slotted operation, variable-size packets (such as IP datagrams), called generically *packets* in the paper, are segmented into fixed-size data units (often called *cells*, using the ATM

*This work was supported by the Italian Ministry for University and Scientific Research (MURST) under the RINGO project.

jargon) fitting into a slot, which are independently transmitted in the ring, and must be re-assembled at the destination. Segmentation and re-assembly of packets into cells can be very costly, especially at very high data rates, when the duration of a cell can be well below a μs , and in the case of simultaneous reception of several packets from many sources, as for example in the case of servers.

The aim of the work presented in this paper is to simplify costly packet re-assembling procedures by keeping all the cells belonging to the same packet contiguous on the ring. Packet contiguity can be achieved either by reservation of a set of contiguous slots, which implies a centralized access control, hence extra costs in terms of delay and control complexity, or by adapting to the WDM features and constraints the well known register insertion techniques, proposed in the past for single-channel, electrical ring-based networks [7, 15, 16]. In this work we investigate the latter approach, considering the possibility to perform optical buffering by means of delay lines [11, 13, 14].

2. Slotted WDM Rings

Multirings or *Multichannel rings* are considered an attractive architecture for all-optical LANs and MANs [8, 10, 12, 17, 20, 21]. The available transmission resources are partitioned in a number of channels operated in a slotted fashion. These logical channels can in general be obtained by a partition of time, frequency, space, code, or any combination of them. Currently, the state-of-the-art optical technology allows the separation of a number of WDM channels of the order of few dozens (usual numbers range between 8 and 32; experiments were run with up to 100 wavelengths), that by itself may not be sufficient for the realization of *large networks*.

Since each WDM channel is operated at a speed which is the largest affordable speed in the electronic domain, each network node can transmit and receive at most an average amount of information equal to one WDM channel. This means that nodes are equipped with transceivers that can receive at most from a channel, and can transmit at most in a channel, in each time slot. To achieve full connectivity among nodes, transceivers must however be tunable, i.e., they must be able to select different WDM channels for transmission and reception in the different time slots. While it can be conceived to have tunability on both the transmission and the reception side, it can be shown that having tunability on one side only is sufficient.

Due to the recent rapid evolution of fast tunable optical transmitters, the architectural solution considered in this paper refers to networks in

which nodes (see Figure 1) are equipped with one tunable transmitter and one fixed receiver (TT-FR systems). This means that WDM channels are associated with disjoint subsets of destination nodes (ranging from one channel for each destination to one channel for all destinations).

Access control assumes a slotted operation, so that the granularity in resource allocation is one time slot in one WDM channel. The transmitter is assumed to be tunable on a slot-by-slot basis. Slots are removed at the destination node, thereby allowing spatial reuse. Providing one channel for a subset of destinations means that all source nodes that need to transmit to a destination in the subset must share the corresponding channel: Medium Access Control (MAC) protocols are thus needed to arbitrate the access to the shared channels. The linearity of the ring topology, together with the assumption that a channel inspection capability is available at each node, ease the design of collision-free protocols. The channel inspection capability can be implemented in a number of ways; one possibility is using the Multiple Sub-carrier Signalling (MSS) technique [4, 5], consisting in multiplexing both data (payload) and control information (header) on the same optical carrier.

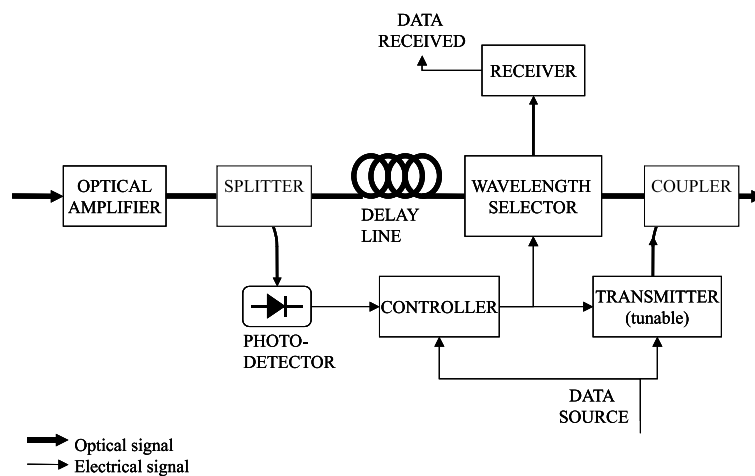


Figure 1. Architecture for a node of a WDM ring network.

2.1. Network Topology

The *Multiring* is an all-optical ring network, with N nodes and W logical channels. By “all-optical” we mean that no conversion to the electronic domain is performed on the data path from the source to the destination along the ring. Channels run in parallel and they are slotted and synchronized, so that a set of W simultaneous slots (a slot for each




Figure 2. The network topology when $N = W = 4$

channel), which we call *multislot*, arrives at a node every slot-time. The overall network latency is S slots.

W logical channels can be obtained purely through wavelength division, if the number of nodes is smaller than the number of wavelengths in the WDM comb, or through a combination of space and wavelength division, by exploiting a number of parallel fiber rings. The number of wavelengths required in the WDM comb equals the number of nodes divided by the product of the number of parallel fibers times the number of destinations in each subset of nodes.

For the sake of simplicity, we consider only the case $N = W$, i.e., multiring networks providing one logical channel for transmissions to each destination. In this case all the nodes that need to transmit to the same destination share the same logical channel. Figure 2 depicts such a case with $N = W = 4$.

It can be observed that each node has a better-than-average access to the channels leading to some destinations, and a worse-than-average access to other channels, leading to other destinations. If we assume node numbers to be increasing in the transmission direction, the channel on which node i , with $0 \leq i \leq N - 1$, has the best access chance is the one leading to node $|i - 1|_N$; indeed, since to access this channel node i needs not defer to transmissions of any other node. The channel on which node i has the worst access chance is the one leading to node $|i + 1|_N$, since, to access this channel, node i must defer to the possible transmissions of $N - 2$ up-stream nodes. In Figure 2, node 0 has the best access opportunity on channel 3 (leading to node 3), and the worst access opportunity on channel 1, where it must defer to transmissions of nodes 2 and 3.

2.2. Fixed and Variable Length Packets

Previous works on slotted WDM multiring networks (see [2, 3, 6, 9, 12, 20]) treated the case of fixed-length packets, whose size exactly matches that of a slot. In such a situation, only when the channel is sensed idle, i.e., the current slot is observed empty, a packet can be transmitted, thus avoiding collision with other packets and giving higher priority to in-transit packets. It is possible to define different MAC protocols to select packets, depending on the node queue structure, i.e., if there is only one queue for all destinations, or if there are several queues, as many as the number of channels, or as many as the number of destinations. We proposed in previous papers [2, 3] a number of access protocols for WDM rings. Distributed access control protocols can be divided into a-posteriori [6] and a-priori protocols, depending on the exploitation of information either only local to the node, or related to the channel state too.

In this paper we consider an a-posteriori access protocol: knowing the state of the channel, the oldest transmissible packet is selected for transmission at each node. Moreover, we focus on variable-length data packets, i.e., the case in which packets occupy an integer number of slots: packet lengths can vary between 1 and MPS (Maximum Packet Size) slots. To cope with the problem of the transmission of packets larger than a single slot, a widely used technique is *interleaving* or *interleaved transmission* (INT): variable-size packets (such as IP datagrams) are segmented into fixed-size data units, one slot long, which are independently transmitted in the ring, and must be re-assembled at the destination.

To avoid reassembling procedures at the receiver, it is worth having the transmission of packets in their entirety, that is, all fragments (cells) of the packet are sent contiguously, such that no other cells can flow amongst those belonging to the same packet. This *continuous packet transmission* relieves the destination node of the complexity of interleaving, because a single packet is received at a time: the storage structure to reconstruct a packet in its integrity can simply be a buffer, whose maximum dimension is as large as the maximum packet size. Moreover, no controls are needed on the received cell, that must be simply stored in the buffer. This kind of transmission strategy can be achieved in several ways; one possibility is to adapt the well known [7, 15, 16] register insertion techniques to the WDM features and constraints.

This work investigates the feasibility and the performances of strategies based upon register insertion, in comparison with the interleaved solution.

3. Continuous Packet Transmission

To accomplish the task of continuous packet transmission, an improved node structure (Figure 3) must be considered with respect to the one proposed in Figure 1: after the wavelength extractor, and before the coupler for the transmitted signals, a new block must be inserted, which consists in a sort of “optical buffer”. First of all, each logical channel is separated by a demultiplexer: thus, each channel could be treated independently, and it is possible to focus the study on a single channel only. Then, every channel, enters a block, which provides the buffering mechanism exploiting delay lines. After the optical buffer (OB) block, which is composed of a switchable set of delay lines, all channels are multiplexed together. This structure allows to simplify the analysis of possible solutions, considering the behavior of a single OB as if the network had a single channel.

Note that the design suffers from the lack in the optical domain of an equivalent of electronic memories, capable of storing information for a variable amount of time. We approximate this behavior by using delay lines, as already proposed in the past by a number of researchers [11, 13, 14] to control contentions in all-optical networks. In the following subsections, several new schema that exploit delay lines to provide continuous packet transmission are presented.

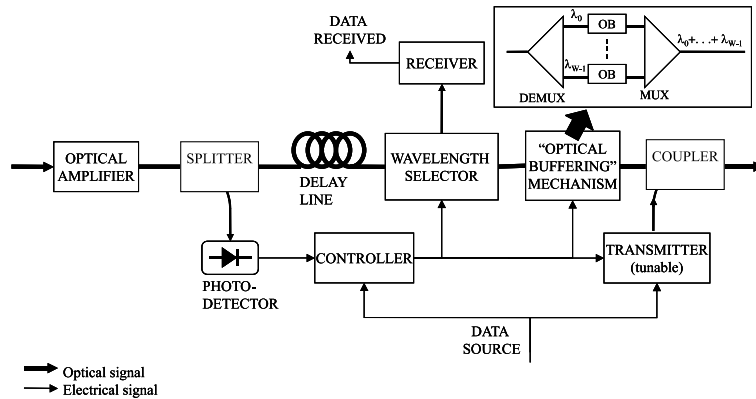


Figure 3. Improved node architecture for continuous packet transmission by delay lines

3.1. Simple Delay Line Scheme

The structure of the buffer must be able to offer continuous transmission for minimum size packets as well as for maximum size packets;

therefore, in the worst case it must provide a buffer whose dimension is the maximum packet size (MPS). However, it is important that the system does not favor packets with particular dimensions, i.e., that short and long packets are treated fairly.

In-transit packets must not be interleaved, hence transmission is possible only in empty slots between busy periods. Consider a single channel for simplicity. The node, at the beginning of each new time slot, senses the channel to probe its state. Suppose that for some slots the channel has been sensed busy: the controller cannot know if the occupied slots belong to the same packet or to different packets; in other words, it cannot distinguish the end of a packet from the beginning of another. However, as soon as it senses an idle slot, it is sure that no interleaving will occur if the node starts to transmit its packet. The problem is that the node has no knowledge about the next incoming slots: they could be all empty, or they could be full. To guarantee continuous transmission while avoiding collisions, the node has to delay possibly incoming busy slots, at least for the time needed to end its transmission. Since it just transmitted one slot, at most it has still to send $\text{MPS} - 1$ slots; thus, $\text{MPS} - 1$ slots is the size of the delay line needed. After the transmission of a packet, up to $\text{MPS} - 1$ cells are held in a delay line which was switched into the data path in the ring.

Until now, it has been seen how the node can start to transmit, waiting for an empty slot between busy periods and setting up a delay line of $\text{MPS} - 1$ slots to hold the incoming traffic for the time needed by the local transmission. However, before attempting a new transmission, the node is supposed to free and to remove the delay line from the ring. Since slots in the delay line cannot be lost, any free slot incoming from the ring can be used to reduce the number of cells in the delay line, hence to reduce the length of the delay line. Only when the delay line is completely empty, it can be switched off from the ring. We call *normal node state* the situation in which no delay line is active and no delay is introduced for in-transit packets.

A number of strategies can be devised to control the access to the ring, and to manage the delay lines. The first transmission strategy, called *Simple Delay Line* (SDL) scheme, works as follows.

- 1 A node can transmit only when it is in normal state and it senses an empty slot.
- 2 If the packet to transmit is 1 slot long, no delay line is activated, because no collisions will occur.
- 3 If the packet is longer, a $(\text{MPS} - 1)$ -slot delay line is activated to hold incoming traffic to avoid collisions.

- 4 After the end of the transmission, at each new slot time, the node controls if the delay line is completely empty; if yes, the delay line is removed from the ring, and the node returns to the normal state.

This scheme becomes inefficient when the MPS increases, since, after a transmission, a node needs a long idle interval (i.e., a long sequence of consecutive empty slots) before returning to the normal state and attempting to transmit again. However, the optical buffer (OB) structure is simple in terms of both hardware components and control logic, since it must only provide a no-delay path, and a maximum-size delay line. If more delay lines with several dimensions were available, it would be possible to progressively reduce the size of the delay line until the OB is de-activated. This requires a good coordination in the switching operation among delay lines; by doing so, the delay suffered by incoming traffic would be decreased, because the buffer could be tailored to the size of the packet to transmit, and to the number of cells in the delay line.

Therefore, consider an OB composed of several parallel spans of fiber, each one with a different length, corresponding to D delay lines and to the normal (non-delayed) connection (Figure 4). The same behavior can be obtained by a switched cascade of D equally sized blocks. For both structures, it is a controller task to coordinate all the switches, so as to properly interconnect the delay lines, and to obtain the correct behavior.




Figure 4. A possible structure for the optical buffer.

Exploiting the improved hardware structure of the buffer, it is possible to define an improved SDL transmission scheme, where D is the granularity degree (i.e., there are D parallel delay lines, according to Figure 4).

- 1 A node can transmit only when it is in normal state and it senses an empty slot.
- 2 If the packet to transmit is 1 slot long, no delay line is activated, because no collisions will occur.
- 3 If the packet is longer (X slots), a delay line is activated to hold incoming traffic to avoid collisions. The dimension of the delay

line is $\frac{i}{D}\text{MPS} - 1$ slots, with $i = \left\lceil \frac{X D}{\text{MPS}} \right\rceil$, such that no information traversing the node is lost.

- 4 When there are active delay lines, a node always controls if it is possible to scale down, decreasing the delay suffered by in transit packets: if the active delay line is not the smallest one and if an idle interval of MPS/D slots has passed, then the next smaller delay line is immediately connected to the input of the buffer, and it is connected to the output as soon as the previous one becomes empty (because all its stored information has been forwarded).
- 5 After the end of the transmission, at each new slot time, the node controls if it is possible to tear down the delay line and to return to the normal state: if the active delay line in the OB is the smallest one ($\frac{\text{MPS}}{D} - 1$ slots), and if it is completely empty, then the node returns to the normal state.

3.2. Complex Delay Line Scheme

Besides reducing the length of the idle intervals needed to reduce the size of the delay line, another way to improve the expected performance of the SDL scheme is to release the constraint that the node must return to the normal state before it is allowed to transmit again: it is possible to enable the transmission of a local packet also when a delay line is active, but it is possible to store both the residual cells in the delay line *and* the cells generated by the local packet. This new scheme is called *Complex Delay Line* (CDL). It must be noted that the OB structure for CDL is the same as the one for the SDL scheme. The difference lies only on the control algorithm. CDL works as follows:

- 1 A node can transmit a new packet only when it has completely transmitted the previous packet.
- 2 If the delay line is inactive and a non-empty slot is sensed, the node is not allowed to transmit.
- 3 If the delay line is inactive and an empty slot is sensed, the node is allowed to transmit a maximum size packet.
- 4 If the delay line is active, the node is allowed to transmit packets smaller than or equal to the number of empty consecutive slots at the end of the delay line: this operation is called *insertion*.
- 5 A node always attempts to return to the normal state when there are active delay lines: this operation is the scaling down followed

by the tearing down (i.e., by the removal of the delay line from the ring). During transmission, the tearing down is disabled, while the scaling down is allowed.

3.3. Fixed Delay Line Scheme

Both SDL and CDL schemes tend to minimize the delay suffered by in transit traffic, due to the priority given to scaling down the delay line. A different transmission scheme, that does not attempt delay minimization, considers a buffer simply comprising a $(MPS - 1)$ -slots delay line that is never removed from the ring data path. We call this scheme *Fixed Delay Line* (FDL). FDL is similar to CDL but it considers an always active delay line (note that the normal state does not make sense anymore), and each packet traversing the node suffers from an extra fixed delay. So, the structure simplicity trades off the larger average delays.

- 1 If both the delay line and the incoming slot are completely empty, then the node can transmit maximum size packets.
- 2 Otherwise, the node is allowed to transmit packets smaller than or equal to the number of empty consecutive slots at the end of the delay line.

The FDL mechanism requires a very simple structure: besides demultiplexer and demultiplexer, it needs only W fixed size delay lines and no switches. For some node architecture designs, one single delay line, with no multiplexers and demultiplexers, where all WDM channels are simultaneously delayed, can do the job, with evident gains in terms of implementation complexity. This simplicity trades with the larger average delay suffered by packets due to the larger propagation delay with respect to SDL and CDL (for which delay lines can also be inactive). Moreover, both the variance of delays and the average round trip time (RTT) increase. Therefore this scheme seems not to be valid for RTT-sensitive network control algorithms and applications.

4. Experimental set-up

A simulator was built to investigate the behavior of the WDM ring network, with different transmission schemes, with different medium access protocols, and under several traffic conditions and patterns. The results reported here consider the following scenario:

- the network has 16 nodes and 16 channels (each channel is devoted to a different destination);
- the ring latency is 240 slots, with 15 slots between adjacent nodes;

- the packet size distribution is uniform between 1 and MPS slots;
- the packet generation process follows a Poisson random process;
- at each node, a FIFO queue is maintained for each ring channel (i.e., for each destination node);
- the adopted MAC protocol is a-posteriori: in each multislot, all incoming slots are inspected, and, among packets at the head of local packet queues, one packet is selected among those that can start transmission in the current multislot; if several such packets exist, the *oldest* packet is selected for transmission;
- the total traffic in the network (ρ) is normalized to the aggregate network bit rate, i.e., to the sum of the bit rates of all nodes;
- the traffic in the network due to the node n_i is ρq_i , where q_i is the percentage of traffic originated by the node itself;
- the traffic in the network due to communications from node n_i to node n_j is $\rho q_i p_{ij}$, where p_{ij} is the probability that a packet originated at node n_i is directed to node n_j ($p_{ii} = 0$);
- the considered traffic pattern is uniform: i.e., all source nodes generate the same amount of traffic ($q_i = 1/N, \forall i$), and destinations are equally likely ($p_{ij} = 1/(N - 1), \forall i, j : i \neq j$);

Results obtained for the INT scheme are considered as a reference; they represent the situation in which packets are segmented and then re-assembled, without any guarantee about their contiguity. All delay line schemes (S/C/F/DL) allow that packets are transmitted as a whole, avoiding that the receiver has to reconstruct several packets simultaneously. This simplification in the receiver hardware has a cost in terms of network capacity, i.e., the maximum throughput that the network can sustain.

4.1. Performance Results

Figures 5 and 6 show curves obtained by simulation for network throughput and packet average delay (total delay except constant propagation and transmission delays), pooling all the different schemes, so as to have a direct comparison for a given value of MPS=64 slots. Table 1 summarizes the network capacity (i.e., the maximum achievable throughput) under different schemes and packet sizes.

As expected (see [6]), the interleaved transmission scheme (INT), with an a-posteriori protocol, allows the full bandwidth exploitation, regardless of the MPS. This is due to the fact that packets are segmented in




Figure 5. Throughput results in uniform traffic for MPS = 64




Figure 6. Average delays in uniform traffic for MPS = 64

independently transmitted cells, so that all slots in the channel can be used. For what concerns average delays, they grow with the maximum packet size; this is obvious, since access and queueing delays are taken into account, and they increase proportionally with the transmission time. In fact, the MPS value is almost a perfect scaling factor: the average delays for the case $\text{MPS} = X$ (with X given) are about X times larger than the average delays in the case $\text{MPS} = 1$.

Simulation results obtained exploiting a SDL scheme with the minimum granularity ($D = 1$) show that for large packets performance is poor, as already mentioned in Subsection 3.1. Indeed, this scheme allows transmission only after emptying the delay line, whose size increases with the maximum packet size, and therefore at high load it is very difficult to access the ring.

By running simulation with different MPS values, and by looking at Table 1, it can however be observed that, when increasing the maximum packet size (i.e., for $\text{MPS} \rightarrow \infty$), the network capacity tends to an asymptotic value, which is around 0.6.

When $\text{MPS} = 64$, a reasonably large value, it is possible to see that performance improve for increasing granularity values: with $D = 4$ the maximum throughput increases by 35%. It must be remarked that with $D = \text{MPS}$, throughput reaches almost 1 and delays are very similar to those obtained for INT: in this case complexity has been removed from the receiver and has been completely moved to the transmitter. Almost no difference is observed on delays between the cases $D = 16$ and $D = 64$: this observation allows to reduce the complexity of optical buffers (OB) block without significant losses in performance.

Note in Figure 6 that the schemes enforcing continuous packet transmission may lead to packet delays slightly smaller than INT (at intermediate load values). This behavior was observed and analytically modeled in [1] for input-buffered switches.

The CDL scheme has been proposed as an improvement of the SDL one, allowing to exploit the same node structure. Indeed, in the basic scenario ($W = N$ and $D = 1$, with varying MPS), better throughput and delays are obtained: this gain can be related to the fact that transmission can happen also when the delay line is active, allowing to fill gaps otherwise left empty. As it was supposed, when MPS is small (e.g., 4), improvement is less pronounced, because the delay line does not need long idle intervals to be inactivated, whereas for large MPS values (e.g., 64) the insertion mechanism takes place more frequently. By simulation we obtained that the maximum throughput becomes 0.85 for $\text{MPS} = 4$ (from 0.83 of SDL), 0.73 for $\text{MPS} = 16$ (from 0.67 of SDL) and 0.71 for $\text{MPS} = 64$ (from 0.63 of SDL).

On the other hand, when the structure complexity increases ($W = N$ and $\text{MPS} = 64$, with varying D), the improvement with respect to SDL is less noticeable, because the scaling down operation prevents exploiting the insertion mechanism, except for a small fraction of packets, as explained in Subsection 3.2. Indeed, for $D = \text{MPS}$, results are identical between SDL and CDL.

In the case of FDL, the maximum normalized throughput is well below one. This is due to the fact that, with respect to the case of switched delay lines, no compression of used slots is performed on the different channels, and channel usage becomes irregular, leaving a large number of small gaps between packets, which cannot be used if packet contiguity is enforced. As already mentioned in Subsection 3.3, delays for FDL have a lower bound (observe the horizontal asymptote for decreasing offered loads in Figure 6), due to the fact that traveling along the ring, packets suffer a delay which increases linearly with the maximum packet size and with the number of traversed nodes. Under uniform traffic conditions, the average delay is biased by a constant extra delay:

$$\begin{aligned} \bar{d}_{\text{extra}} &= \sum_{i=1}^{N-1} \left[\frac{1}{N-1} (\text{MPS} - 1) (i - 1) \right] = \\ &= \frac{N-2}{2} (\text{MPS} - 1) \end{aligned}$$

because a packet transmitted, say from node 0, has to traverse $i-1$ delay lines to reach node i . For example, when $\text{MPS} = 4, 16, 64$, the average extra delay is respectively 21, 105, 441 slots in our simulation scenario.

With respect to CDL with $D = 1$, the FDL scheme adopts the insertion mechanism only, so it tries to fill all the gaps. From Table 1 it is possible to see that, for low MPS values, CDL performs better than FDL; on the contrary, for high MPS values, FDL has a better behavior. This difference is motivated only by the activation/deactivation operation of the delay line, because, when the line is active, the behavior of an OB block is the same in CDL and FDL. However, when the delay line is inactive, CDL can transmit packets in idle intervals smaller than the packet size, at the price of modifying the gap; instead, when the idle interval is larger than the packet size, the behavior is similar to FDL (insertion only) again.

5. Conclusions

This work investigates several strategies to realize the insertion buffer technique in the optical domain, to allow continuous transmission of variable-size packets. It seems that the task to alleviate the receiver

| SCHEME | | MAXIMUM PACKET SIZE | | | |
|--------|----------|---------------------|--------|--------|--------|
| | | 1 | 4 | 16 | 64 |
| INT | | 0.9999 | 0.9999 | 0.9990 | 0.9973 |
| SDL | $D = 1$ | — | 0.8404 | 0.6897 | 0.6553 |
| | $D = 4$ | — | 0.9999 | 0.9202 | 0.8796 |
| | $D = 16$ | — | — | 0.9994 | 0.9718 |
| | $D = 64$ | — | — | — | 0.9985 |
| CDL | $D = 1$ | — | 0.8675 | 0.7507 | 0.7389 |
| | $D = 4$ | — | 0.9999 | 0.9242 | 0.8845 |
| | $D = 16$ | — | — | 0.9994 | 0.9723 |
| | $D = 64$ | — | — | — | 0.9985 |
| FDL | | — | 0.8374 | 0.7876 | 0.7751 |

Table 1. Simulation results for the network capacity under uniform traffic

from heavy re-assembly operations can be reasonably afforded without high costs in the transmitter complexity: CDL with $D = 1$ and FDL provide interesting performances, exploiting only delay lines, and, for CDL, few switches. If it were economically worth, solutions with larger values of D would assure an almost full bandwidth utilization. The choice about CDL and FDL depends on the requirements about delays: if they are severe, the choice is for CDL because FDL implies too large constant delays, which increase with both the number of nodes and the maximum packet size; on the contrary if delays are less important, FDL allows a simpler structure.

References

- [1] M. Ajmone Marsan, A. Bianco, P. Giaccone, E. Leonardi, and F. Neri. Packet scheduling in input-queued cell-based switches. *IEEE INFOCOM 2001*, April 2001.
- [2] M. Ajmone Marsan, A. Bianco, E. Leonardi, M. Meo, and F. Neri. MAC protocols and fairness control in WDM multirings with tunable transmitters and fixed receivers. *IEEE/OSA Journal of Lightwave Technology*, 14(6):1230–1244, June 1996.
- [3] M. Ajmone Marsan, A. Bianco, E. Leonardi, A. Morabito, and F. Neri. All-optical WDM multi-rings with differentiated QoS. *IEEE Communications Magazine*, 2(37):58–66, February 1999.
- [4] M. Ajmone Marsan, A. Fumagalli, E. Leonardi, F. Neri, and P. Poggiolini. Daisy: a scalable all-optical packet network with multifiber ring topology. *Computer Networks & ISDN Systems*, 30:1065–1082, 1998.
- [5] S. Benedetto and P.T. Poggiolini. Theory of subcarrier encoding of packet headers in quasi-all-optical broadband WDM networks. *IEEE/OSA Journal of Lightwave Technology*, 12(10):1869–1881, October 1994.

- [6] A. Bianco, V. Distefano, A. Fumagalli, E. Leonardi, and F. Neri. A-posteriori versus a-priori access strategies in slotted all-optical WDM ring. *Computer Networks and ISDN Systems*, 32:617–631, 2000.
- [7] W. Bux and M. Schlatter. An approximate method for the performance analysis of buffer insertion rings. *IEEE Transactions on Communications*, COM-31(1):50–55, January 1983.
- [8] Y. Cai, R.M. Fortenberry, and R.S Tucker. Demonstration of photonic packet-switched ring networks with optically transparent nodes. *IEEE Photonic Technology Letters*, 6(9):1139–1141, May 1994.
- [9] A. Carena, M. Vaughn, R. Gaudino, M. Shell, and Blumenthal D.J. Opera: an optical packet experimental routing architecture with label swapping capabilities. *IEEE/OSA Journal of Lightwave Technology*, 16(12):2135–2145, December 1998.
- [10] M.J. Chawki, V. Tholey, E. Delevaque, S. Boj, and E. Gay. Wavelength reuse scheme in a WDM unidirectional ring network using a proper fiber grating add/drop multiplexer. *OFC'95*, paper ThI3, February 1995. San Diego (CA).
- [11] I. Chlamtac et al. CORD: Contention resolution by delay lines. *IEEE Journal on Selected Areas in Communications*, SAC-14(5):1014–1029, June 1996.
- [12] I. Chlamtac, A. Fumagalli, Kazovsky L., and Poggiolini P. A contention and collision free WDM ring network for multi-gbit/s packet switched communication. *Journal of High Speed Networks*, 1(4):1–19, 1995.
- [13] I. Chlamtac, A. Fumagalli, and C.-J. Suh. A delay line receiver architecture for all-optical networks. *IEEE INFOCOM '96*, 2:419–426, 1996.
- [14] I. Chlamtac, A. Fumagalli, and Chang-Jin Suh. Multibuffer delay line architectures for efficient contention resolution in optical switching nodes. *IEEE Transactions on Communications*, COM-48(12):2089–2098, December 2000.
- [15] I. Cidon and Y. Ofek. MetaRing – A full-duplex ring with fairness and spatial reuse. *IEEE Transactions on Communications*, COM-41(1):110–120, January 1993.
- [16] D.E. Hubber, W. Steinlin, and P.J. Wild. SILK: An implementation of a buffer insertion ring. *IEEE Journal on Selected Areas in Communications*, SAC-1(5):1214–1223, November 1983.
- [17] A. Jourdan et al. The perspective of optical packet switching in ip-dominant backbone and metropolitan networks. *IEEE Communications Magazine*, 39(3), March 2001.
- [18] B. Mukherjee. *Optical Communication Networks*. McGraw-Hill, 1997.
- [19] R. Ramaswami and K.N. Sivarajan. *Optical Networks: A Pratical Perspective*. Morgan Kaufmann Publishers, 1998.
- [20] K.V. Shrikhande et al. Hornet: a packet-over-WDM multiple access metropolitan area ring network. *IEEE Journal of Selected Areas in Communications*, 18(10):2004–2016, October 2000.
- [21] V. Tholey, M.J. Chawki, L. Berthou, I. Le Gac, E. Gay, and A. Poudoulec. Demonstration of a WDM survivable open ring network using reconfigurable dropping receivers. *ECOC'94*, September 1994. Firenze, Italy.