A-POSTERIORI ACCESS STRATEGIES IN ALL-OPTICAL SLOTTED WDM RINGS*

A. Bianco[a], V. Distefano[a], A. Fumagalli[b], E. Leonardi[a], F. Neri[a]
E-mail: {bianco,leonardi,neri}@polito.it, andreaf@utdallas.edu
(a)– Dipartimento di Elettronica, Politecnico di Torino, Italy
(b)– E. Jonsson School of Engineering and Computer Science, University of Texas at Dallas

The paper compares a number of strategies devised to gain access in slotted multi-channel all-optical ring networks whose nodes have a distinct transmission queues for each channel and limited transmission/reception capabilities, i.e., one tunable transmitter and one fixed receiver. The considered access strategies can be subdivided into two groups, depending on the information used to select the next packet for transmission. "A-priori" strategies select the packet only according to the knowledge of the state of transmission queues. "A-posteriori" strategies rely on the knowledge of the state of both transmission queues and optical channels. A number of strategies belonging to the former group were recently proposed by the authors and offer an intrinsic hardware and algorithmic simplicity. Two novel strategies belonging to the latter group are proposed in this paper with the aim of assessing the trade-offs between complexity and overall system performance, including achievable throughput and queuing delay. Simulation results obtained under several traffic scenarios show that the increased complexity of a-posteriori strategies pays off in terms of performance.

1 Introduction

All-optical networks are particularly attractive due to their capability of providing access to the huge bandwidth (several THz) available in the fiber without incurring in the so called "electronic bottleneck" due to multiple electro/optical conversions. A practical approach for exploiting the fiber bandwidth is based on Wavelength Division Multiplexing (WDM) that partitions it into a set of channels whose individual transmission rate matches the speed of the electronic interfaces at the User/Network Interface (UNI). WDM is often combined with either space (i.e., multiple parallel fibers) or time diversity, or both [1].

Depending on the number of available logical or physical channels, and the number of nodes in the network, several system architectures have been proposed (see [2] for a survey), some dedicating each channel entirely to a single node's transmission (e.g., [3]), some dedicating each channel entirely to a single node's reception (e.g., [4]), and some others providing a number of channels which are dynamically shared by a subset of nodes, or by all nodes, for either transmission or reception. Fiber based ring topologies have the flexibility to realize any of the above multiplexing architectures [4, 5, 6, 7, 8] and offer practical survivable techniques against faults [9]. In addition, it is possible to divide time into slots of fixed length and nodes' transmissions (in the slots) need not be scheduled in advance since both packet collisions on the channels and receiver contention at the destination can be avoided by monitoring the channels prior to packet transmission [4].

Without loosing generality, in the remainder of this paper we consider a ring network with a number of channels that is equal to the number of nodes. Each channel is associated with a destination node; thus, a source node willing to transmit a packet to a given destination must use a slot in the corresponding channel. Slot synchronization in the system is such that slots on different channels are perfectly aligned. In short, in each time slot a node is traversed by a set of slots, i.e., by a slot on each channel (which means a slot for each destination), and may decide to use one of the empty (unused) slots for transmission towards the corresponding destination node. This approach has the advantage that each node makes use of only a single fixed receiver and a fast tunable transmitter, optical components with proven technology [4, 5], and it does not originate packet collisions on the channels. In addition, receiver contention, i.e., packets transmitted on distinct wavelengths that simultaneously arrive to the same destination node, cannot occur.

In such a multi-channel ring network, since all source nodes that need to transmit to the same destination must share the corresponding channel, medium access control (MAC) protocols are required to arbitrate the access to the shared channels in each slot time. The access strategy is the part of the MAC protocol responsible for selecting the packet, i.e., the channel, for the node's next transmission. It determines network performance, level of fairness and complexity of the electronic control at the UNI.

Recently, the authors proposed a number of access strategies [4, 10, 11] that select the next packet transmission using only information local to the node, i.e., the destinations of the packets stored in the transmission queues. The advantage of this approach is the simplicity of the electronic control of the UNI as the packet is selected prior to monitoring the channels. Once the packet is selected, only the channel cor-

---

*This work was supported in part by the Italian National Research Council, and by the Italian Ministry for University and Scientific Research.

0-7803-4984-9/98/$10.00

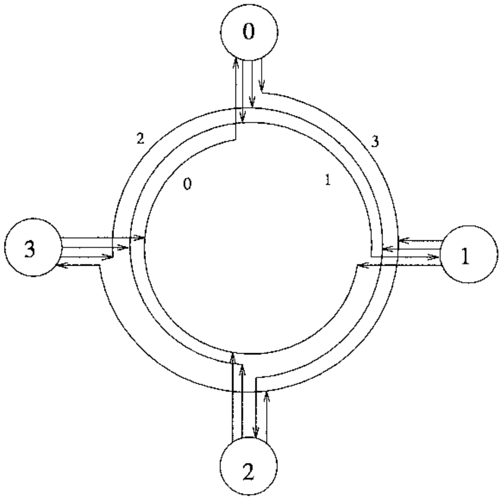

Figure 1: The logical network topology when $M = 4$

responding to the packet's destination is monitored by the UNI in order to determine whether the packet can be actually transmitted without generating collision. This group of strategies is termed "a-priori". In [10, 11, 12] it is shown that a-priori strategies which make use of a distinct queue associated with each channel, well exploit the system bandwidth asymptotically reaching 100% efficiency and guaranteeing a reasonable level of fairness when the distribution of the offered traffic is balanced, i.e., traffic patterns are uniform. Among the proposed access strategies, SRR (Synchronous Round Robin) [10, 12] leads to the best performance. However, the presence of unbalanced traffic patterns in the system does not allow to maintain a perfect SRR access and leads to some bandwidth waste. This limitation is due to the non-null probability that the a-priori strategy selects for transmission a channel whose slot is already busy.

To circumvent this drawback of a-priori strategies this paper proposes two novel strategies, the Sequential Priority Choice (SPC) and the Longest Queue Choice (LQC), belonging to the group of "a-posteriori" strategies. With these access strategies, the UNI first monitors the whole set of channels to establish which slots are empty, then it selects for transmission one packet among those whose corresponding channel is free. This requires an increase in the algorithmic and hardware complexity of the MAC protocol, but it guarantees that it is always possible to select (if any) a packet that will be successfully transmitted.

This study aims at establishing the trade-offs between the additional complexity required by the a-posteriori strategies when compared to the a-priori strategies and the potential gains in terms of system capacity and level of fairness.

2 System Description

We consider ring networks providing one dedicated logical channel for transmissions directed to each destination. If we denote by M the number of nodes, a total of M logical channels is available in the network. All nodes that need to transmit to the same destination share the same logical channel.

Fig. 1 depicts the logical network architecture for case $M = 4$. The M logical channels run in parallel; they are slotted and synchronized, so that M slots (one for each destination node) simultaneously reach a node every time slot. Full connectivity is achieved by tuning the transmitter at the source node on a packet-by-packet basis. Fixed-length data packets are transmitted, and the slot size is such that one packet transmission time and the transmitter tuning time fit into one slot. Each channel is cleared up at the corresponding destination node, to prevent optical signal re-circulation in the ring. This means that each logical ring is interrupted at a (different) node, and also that the multi-channel ring can be viewed as a set of staggered logical busses folded in a physical ring layout. We assume that nodes do not transmit to themselves.

The channel inspection capability required on every channel at each node to avoid collisions among packets directed to the same destination is obtained by monitoring the optical power in the channels. If contention arise, priority is always given to in-transit packets: as a consequence, once a packet is transmitted, it will reach its destination without collision.

It can be observed that, due to the ring symmetries, 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 some other destinations. If we increasingly number from $i = 0$ to $i = M-1$ the nodes along the propagation direction of the optical signal in the fiber, the channel on which node i, $i = 0, 1, \ldots, M-1$, gains the easiest access is the one leading to node[1] $|i-1|_M$. This occurs because, on the channel associated with node $|i-1|_M$, node i doesn't need to defer its transmission to any other node' transmission. Conversely, the channel on which node i gains the most difficult access is the one leading to node $|i+1|_M$, as, in order to access this channel, node i must defer its transmission to all transmissions of the $M-2$ upstream nodes. In general, to distinguish among the various transmissions, we define the *access priority* of a packet as the value $|j-i|_M$, where i is the source node and j is the destination node of the packet. According to this notation, 1 and $M-1$ are the lowest and highest access priorities, respectively.

Each node stores the generated packets in $M-1$ distinct (FIFO) transmission queues, one associated with each destination/channel. Since each queue stores packets with a specific access priority, for simplicity, we refer to each queue using the access priority of the corresponding stored packets.

3 A-Posteriori versus a-Priori Access Strategies

The dependence of the access priority on the position of the node along the channel intuitively raises fairness issues. An extreme case occurs when a node grabs all the slots on a channel possibly starving downstream nodes competing for the access to the same resource. The MAC protocol of the slotted multi-channel ring must therefore arbitrate the nodes' access on each channel in order to assure that fair and efficient utilization of every channel band-

---

[1] The notation $|\cdot|_M$ indicates the modulo M operator.

width is achieved in the system. The key factor of the MAC protocol that is responsible for achieving the above objective is the **access strategy**, adopted by the node to select the next packet for transmission. We recall that the node makes use of a single tunable transmitter which allows the transmission of at most one packet per time slot.

We focus our attention on access strategies that offer distributed access control and do not require nodes' coordination prior to transmission. However, some strategies require to label all the slots travelling in the multi-ring with their ordinal number (i.e., a global slot identifier is needed). Access strategies can be subdivided into two families, namely a-priori and a-posteriori access strategies.

A-priori strategies select the packet to be transmitted only on the basis of the distribution of the packets in the transmission queues of the node, without the knowledge of the channels state. Once the packet has been selected, the channel associated with the packet's destination is monitored: if the corresponding slot is free, the packet is successfully transmitted, otherwise, if the slot is busy, the node refrains from transmission to avoid collision. In each time slot this procedure is executed anew.

A-priori strategies are appealing due to their simple control. They do not require the inspection of all channels in each time slot as only one channel must be monitored in order to decide whether the selected packet can be successfully transmitted or not. The drawback is that not always the a-priory choice succeeds in selecting a packet whose destination channel is free, thus some bandwidth may be wasted. In fact, since nodes using a-priori strategies select a packet for a tentative transmission independently from the channel state, they can select packets directed to destinations whose channel is busy, resulting in contention, even when other packets that could have been transmitted without contention are available. A non optimal utilization of the network capacity is therefore expected. This problem was addressed in previous works that proposed some access strategies that mitigate this drawback by reducing contentions among packets whose transmission is attempted in the same slot to the same receiver.

A-posteriori strategies rely on the knowledge of both the distribution of the packets in the transmission queues of the node and the current availability of the channels. First, the node monitors the channels and detects those with an empty slot. After this preliminary detection, a packet (if any) chosen from the queues associated with the free channels is successfully transmitted. We say that non-empty transmission queues corresponding to destinations whose slots are sensed free are **enabled;** a-posteriori strategies thus select one among the packets at the head of enabled queues.

With an a-posteriori strategy a node is always able to transmit a packet if it has at least one enabled transmission queue. Thus, the bandwidth efficiency issue of a-priori strategies is taken care of at the cost of some additional complexity of the MAC protocol. In the optical layer this complexity consists of an array of sensors that continuously monitor all the channels. At the electronic level this complexity implies that the node must process the information derived from every channel in parallel, thus, requiring additional fast electronics when compared to the a-priori approach. In terms of processing, the selection of the packet must be done only *after* that the state of all channels has been detected, i.e., the detection and selection operations must be serialized, a problem that does not arise in a-priori strategies where the selection of the packet and the channel inspection can run in parallel (see [5]). This serialization leads to an increased signal latency at the node.

Although straightforward, a-posteriori strategies must pay particular attention to the fairness issue. Contrary to a-priori strategies where the selection of the packet is totally independent from the availability of the channels, in a-posteriori strategies the selection of the packet is affected by the traffic originated by upstream nodes. Intuitively, this mechanism tends to concentrate the packet selection on those channels that are more often free. As a result, packets with high access priorities ($M-1$, $M-2$, ...) are more likely to be transmitted than others, thus exacerbating the unfairness problem of slotted ring network. In addition, it must be noted that the node does not have any information regarding the activity of downstream nodes. As a consequence, each node can statistically improve its packet transmissions chance in subsequent slots, for example by maximizing the number of transmission queues that are left non-empty after each transmission, but it cannot maximize the transmission probabilities of the other (downstream) nodes.

Several a-priori strategies were proposed by the authors in previous papers (see [10, 11]); we choose to compare in this paper two novel a-posteriori strategies with the a-priori strategy called SRR (Synchronous Round Robin), which allows an optimal utilization of the transmission bandwidth under uniform traffic, as shown in [11, 12]. We expect to overtake some of the SRR limitations by using a-posteriori strategies, namely its reduced efficiency under non-uniform traffic patterns, and the tight requirements imposed on the ring latency in order to allow the SRR scheduling to work properly.

The two a-posteriori access strategies are the Sequential Priority Choice (SPC) and the Longest Queue Choice (LQC). SPC is a direct extension of SRR to the a-posteriori case, and like SRR, it performs very well under uniform traffic conditions. LQC instead privileges transmissions on the more loaded channel, with the aim of balancing the occupancy of the transmission queues, thereby improving the overall network performance.

3.1 The SRR access strategy

The SRR access strategy is based on a cyclic (preferential) scheduling of transmission queues. This scheduling is obtained by first numbering the slots in the ring with an incremental counter s. In time slot s, node i selects for a tentative transmission the oldest packet with destination $|i+k+1|_M$, with $k=|s|_{M-1}$ (where k is called SRR in-frame-slot-label for s). If the corresponding queue is empty, the transmission of the oldest packet from the longest queue is attempted. If more than one longest queue exists (a tie for the longest queue occurs) SRR selects the packet in the queue with the lowest access priority. In any case, if the transmission of the selected packet in not possible because it would generate collision, a new selection is made in the next slot (labeled $s+1$), according to the described algorithm.

SRR behaves almost like a single-queue FIFO strategy for low network loads, while its preferential scheduling deterministically orthogonalizes nodes' transmission attempts when all queues are always non-empty. More precisely, for network loads larger than the channel capacity, nodes behave like in a Time Division Multiple Access (TDMA): during a SRR frame, whose length is equal to $M-1$ slots, all nodes have exactly one access opportunity for transmissions toward each other node, and they exploit this access opportunity deterministically, thereby avoiding potential conflicts. In this sense SRR is optimal under uniform traffic conditions. A more detailed description of the SRR access strategy can be found in [11].

3.2 The SPC access strategy

According to the SPC strategy, and similarly to the SRR strategy, node i selects for transmission in slot s the oldest packet in the queue associated with destination $|i+k+1|_M$, with $k=|s|_{M-1}$, under the condition that the queue is enabled, i.e., it is non-empty and the corresponding channel is free. As an alternative, SPC selects the oldest packet in the longest enabled queue. If two or more queues satisfy the condition (a tie for the longest queue occurs) the lowest access priority among the longest enabled queues is selected.

3.3 The LQC access strategy

Nodes following the LQC access strategy always select for transmission the packet at the head of the longest enabled queue, i.e., among queues storing packets directed to destinations whose channel is sensed free. If more than one longest enabled queue exist (a tie for the longest queue occurs), the lowest priority among longest enabled queues is selected.

The rationale behind the LQC design is the attempt to maximize the number of future transmissions by maximizing the number of queues left non-empty after each transmission, thereby improving overall performance.

4 Performance Comparison

Performance comparison among the three access strategies is carried out via simulation using a network with 16 nodes, i.e., $M=16$. In the simulation model, packets whose fixed length fits into one slot, are generated according to a Poisson process whose rate is denoted by λ_i at node i. The total network traffic is thus $\Lambda=\sum_{i=0}^{M-1}\lambda_i$. The probability that a packet originated at node i is directed to node j is indicated by p_{ij}, with $p_{ii}=0$. The length of the multi-ring is equivalent to 30 slots (i.e., 6 Km at 1 Gb/s for packets of 1000 bits). We assume that each transmission queue has a finite capacity of 100 packets. Packets that cannot find room in the transmission queue are dropped.

4.1 Uniform traffic scenario

Fig. 2 shows the average queuing delay, measured in time slots, versus packet priority in a uniform traffic scenario, i.e., $\lambda_i=\Lambda/M\ \ \forall i$ and $p_{ij}=1/(M-1)\ \ \forall i\neq j$. Delays are averaged over all packets with the same access priority at all nodes. Four cases are shown, by varying the offered load: $\lambda_i=0.75, 0.88, 0.94$, and 0.97.

In addition to the three access strategies, the delays obtained by TDMA, where successive slots are statically devoted to transmissions to successive destinations, are shown for comparison.

It can be seen that SRR asymptotically approaches TDMA for increasing loads. A reduced queueing delay is achieved by the a-posteriori strategies. At heavy loads, the gain achieved with the a-posteriori strategies decreases since SRR is designed to be optimal in sustained overload.

Focusing our attention on the two a-posteriori strategies, it can be noticed that LQC has the tendency to reduce the worst queuing delay (corresponding to the packets with the lowest access priority) at the expenses of the queuing delay of the highest access priority packets. This result is confirmed by the intuition that by giving priority to the transmission of packets in the longest queues, LQC attempts to equalize the delay in the transmission queues. In addition, LQC achieves a slightly lower overall average queuing delay since its balancing effect on the average leads to a larger number of non-empty queues, which in turns increases the probability to find a packet for transmission at any given time slot.

4.2 Varying the ring latency

The performance of SRR is sensitive to the ring length. This fact can be explained considering that, in order to optimally apply the SRR packet selection strategy, slot s after one round trip in the ring must return to the node as slot $s+k(M-1)$, where k is an integer, in order to maintain the same SRR in-frame-slot-label during the successive tours of the multi-ring. Only under this condition the source node maintains a perfect round robin schedule over the other (destination) nodes. When this condition is not satisfied by the ring length, the TDMA-like orthogonalization of the SRR scheme under heavy load conditions cannot be achieved. As a consequence, the ring length, measured in slot times, must be an integer multiple of the SRR frame, (i.e. of $M-1$ slots): in Fig. 2 it was taken to be equal to $2\times(M-1)=30$ slots. When the ring length does not satisfy the above constraint, a strong performance degradation is experimented.

Fig. 3 reports the queuing delay measured at node 0 versus packet priority for each of the three strategies. Traffic is uniform, with $\lambda_i=0.88$ and the round trip delay of the ring is chosen to be different from an integer multiple of $M-1$ time slots. As expected SRR reveals to be unfair with a delay factor of 100 between the lowest and the highest priority packets. Conversely, despite its synchronous scheduling SPC achieves a queuing delay that is only marginally sensitive to the variation of the ring length. At higher loads SPC originates a more pronounced unfairness behavior, whereas LQC, due to its asynchronous access, is totally unaffected by the ring length at any load.

4.3 Client-server scenario

We consider next a client-server type of traffic scenario: one node ($i=15$) behaves like a server generating and attracting a large fraction of the network throughput; the remaining $M-1$ nodes behave like clients and direct half of their traffic to the server

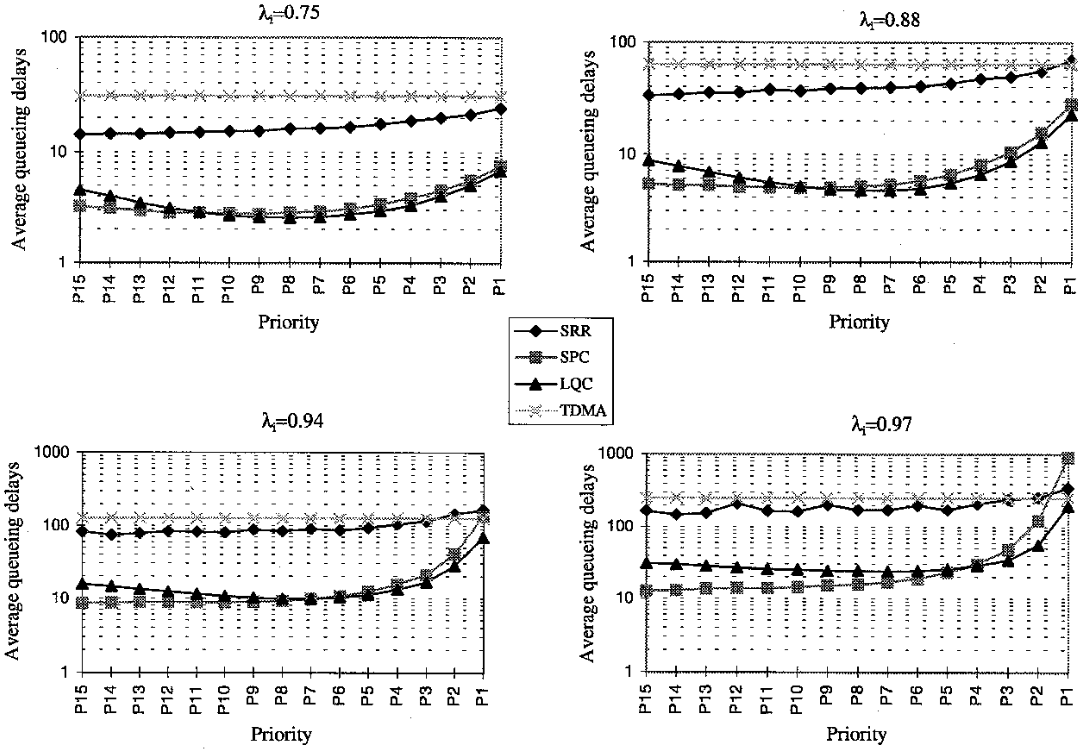

Figure 2: Queuing delays by access priorities for the three access strategies and the TDMA scheme for variable network loads

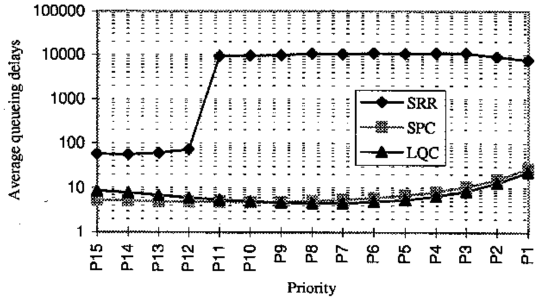

Figure 3: Queuing delays at node 0 when the ring length is equal to 34 slots

node, while the other half of the traffic is uniformly distributed among client nodes. The server generates an amount of traffic that is equal to half the total traffic generated by clients, and evenly distributes it to client nodes. The arrival rates are, therefore, $\lambda_{15} = \Lambda/3$ and $\lambda_i = (2/3)\Lambda/(M-1)\ \forall i \neq 15$, while $p_{15,j} = 1/(M-1)\ \forall j \neq 15$, $p_{ij} = 1/[2(M-2)]\ \forall i,j \neq 15$, $i \neq j$, and $p_{i,15} = 1/2\ \forall i \neq 15$.

An improvement in packet delay performance at the server node is achieved with respect to SRR by using either of the two a-posteriori access strategies, as it can be seen in Fig. 4, where the delay performance of SRR and SPC are reported for all source/destination pairs with a network load $\Lambda = 0.95$. A-posteriori strategies allow the server to optimize its transmission scheduling as a source node, while SRR forces the server to behave like in a TDMA scheme, which is optimal only under overload. Results not reported for LQC are very similar to SPC.

No significant differences are found on how the three strategies treat packets transmitted by the lowest access priority clients and directed to the server. The poor access performance of lowest priority clients on the channel leading to the server are mainly due to the lack of free slots on that channel (whose load is 0.9), therefore also a-posteriori strategies are not able to improve that situation.

A strong reduction in the delays at client nodes being at low access priority on the channel leading to the server when transmitting toward other client nodes is observed for the a-posteriori strategy. Under SRR, client nodes close to the server attempt to transmit packets to other clients only once for each SRR frame, when the corresponding transmission queues are selected according to the SRR preferential scheduling algorithm. This occurs because, every time the queue scheduled by SRR is found empty, a transmission is attempted on the server channel, where the transmission queue is very long. As a consequence, a packet destined to other clients must wait, on average, at least half the SRR frame before being successfully transmitted. Under SPC (and LQC), every time the server channel is found busy, a packet directed to other clients can be transmitted successfully. Thus, the waiting time of packets

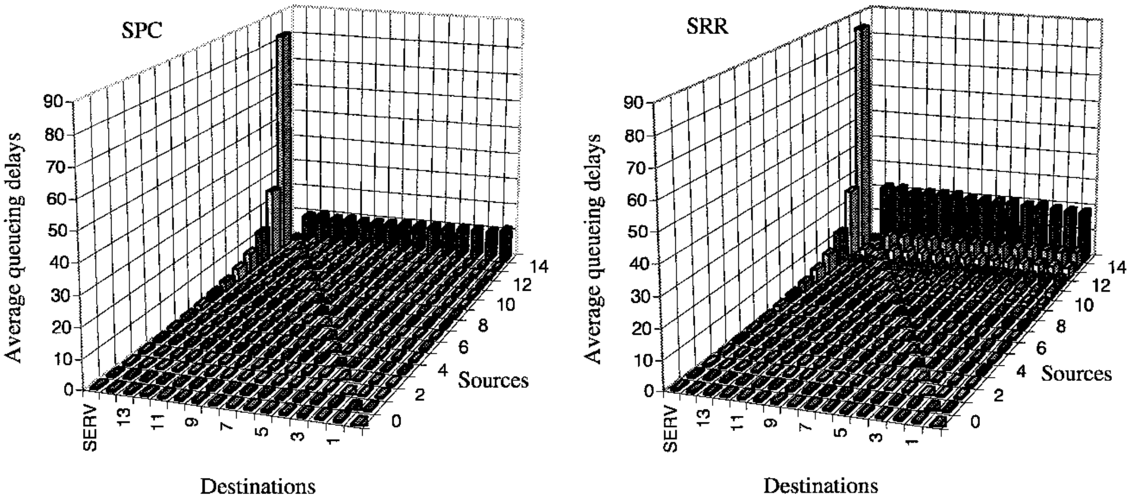

Figure 4: Queueing delays for SPC (left) and SRR (right) in a client-server traffic scenario

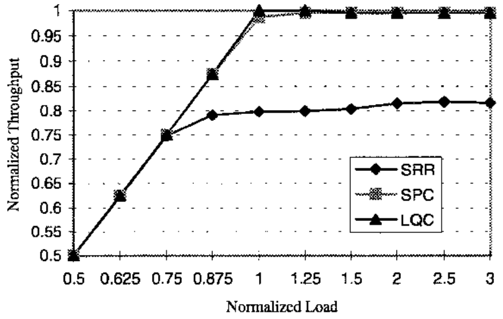

Figure 5: Throughput for SPC, LQC and SRR when only 8 nodes are active

directed to other clients is shorter.

4.4 Other non-uniform traffic scenarios

Several other traffic scenarios were considered, in order to exhaustively investigate the differences between a-posteriori access strategies and SRR. We discuss some relevant results obtained under two traffic scenarios.

In the first scenario, only eight nodes (e.g., the nodes with even index) are active. Each active node generates the same amount of traffic that is evenly transmitted towards the other active nodes.

In this case, the a-posteriori access strategies manage to fully exploit the network capacity whereas SRR fails, as documented in Fig. 5, where node throughput vs. node load is plotted. In each node, once every two slots the transmission queue selected according to the SRR preferential scheduling is found empty, and the transmission of the packet at the head of the longest queue is attempted. As a consequence, also in heavily loaded conditions (i.e. when queues are almost never empty) contentions among nodes are likely to arise.

In the second traffic scenario only three nodes are active, i.e., nodes 0, 1, and 2. Node 0 transmits all its packets to node 3, node 1 transmits all its packets to node 4, and finally, node 2 transmits half of its packets to node 3 and half of its packets to node 4. Nodes 0 and 1 generate the same amount of traffic $\lambda_0 = \lambda_1 = 0.5$, while $\lambda_2 = 2\lambda_0 = 1.0$.

In this case, both SRR and a-posteriori strategies are not able to fully exploit the network capacity; to obtain this goal, all the slots on both channels leading to nodes 3 and 4 must be continuously filled by the three transmitting nodes. Nodes 0 and 1 should alternate their transmissions, so as to fill just one channel in each slot time, leaving the other channel empty for transmissions of node 2. In this case, node 2 would be able to transmit in each slot on the free channel.

However, either under SRR or any of the a-posteriori strategies, transmissions of nodes 0 and 1 are completely independent; thus, with probability 0.25 nodes 0 and 1 transmit in the same slot time in two different channels. Node 2 finds both channel busy, and can not transmit any packet; with probability 0.25, instead, both channels are left free from nodes 0 and 1, in the same slot time; since node 2 is able to fill only one slot at a time, being equipped with a single transmitter, this leads to an obvious throughput inefficiency. In addition, under SRR, node 2 is not always able to transmit even when a free channel is left by nodes 0 and 1, since it could make the wrong a-priori choice with probability 0.5. The effect of these phenomena is shown in Fig. 6, where throughputs experimented by node 2 on the channels leading to nodes 3 and 4 are reported for the three considered access strategies. SRR performs worse than both the a-posteriori strategies, but the maximum achievable throughput is never reached.

5 Conclusions

This paper studied performance versus complexity trade-offs in two families of access strategies used in MAC protocols for slotted multi-channel all-optical ring networks whose nodes have a distinct transmission queue for each channel and limited transmis-

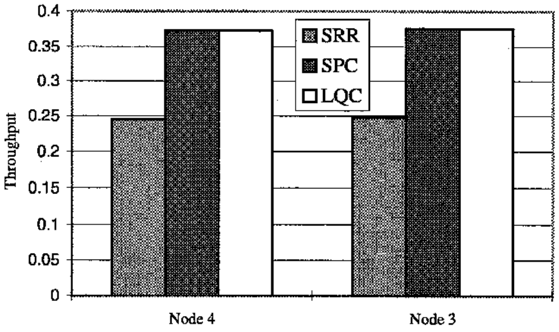

Figure 6: Throughput for SRR, SPC and LQC at node 2 on the two channels leading to nodes 3 and 4

sion capabilities, i.e., one tunable transmitter and one fixed receiver.

The first family, termed a-priori, comprises those strategies that select the packet for transmission based on the distribution of the packets in the transmission queues. The second family, termed a-posteriori, consists of strategies that select the packet for transmission after having sensed the channels. Both families operate solely on local information, i.e., network nodes are not aware of other nodes' state. Nodes have a "carrier sense" capability to inspect the channel availability in order to avoid packet collisions in the channels.

We presented two novel a-posteriori strategies, SPC and LQC, that require an increase in the hardware and algorithmic complexity with respect to a-priori strategies, but lead to significant performance improvements under several traffic patterns. Among the two a-posteriori strategies, LQC, being based on a simpler algorithm and leading essentially to the same performance of SPC, seems the more promising. Our conclusion is that the increase of complexity required by a-posteriori access strategies when compared to a-priori strategies is well justified when the network under design is expected to operate under highly unbalanced traffic scenarios.

Previous works [10, 13, 14] showed that the SRR access strategy, and in general any a-priori strategy, does not guarantee by itself the full exploitation of the network bandwidth when the load on some channel exceeds the channel capacity. To design a more effective MAC protocol, a fairness control scheme can be superposed to SRR to achieve a fair partition of the network resources among nodes [10, 13]. Moreover, a-priori strategies, even with fairness control, lack the ability to provide acceptable service to applications requiring a certain degree of QoS; to overcome this limitation, in [14] a reservation scheme is proposed, which permits to efficiently transport real time traffic in the slotted multi-ring.

Similar considerations apply to the a-posteriori strategies presented in this paper.

References

[1] C.A.Brackett *et al.*, "A Scalable Multiwavelength Multihop Optical Network: A Proposal for Research on All-Optical Networks", *IEEE/OSA Journal of Lightwave Technology,* v.11, n.5/6, pp.736-753, May/June 1993

[2] B.Mukherjee, "WDM-Based Local Lightwave Networks Part I: Single-Hop Systems", *IEEE Network,* v.6, pp.12-27, May 1992

[3] M.-S.Chen, N.R.Dono, R.Ramaswami, "A Media-Access Protocol for Packet-Switched Wavelength Division Multiaccess Metropolitan Area Networks," *IEEE Journal on Selected Areas in Communications,* v.8, n.6, pp.1048-1057, August 1990

[4] I.Chlamtac, A.Fumagalli, L.G.Kazovsky, P.T.Poggiolini, "A Multi-Gbit/s WDM Optical Packet Network with Physical Ring Topology and Multi-Subcarrier Header Encoding", *ECOC '93,* Montreux, Switzerland, Sept.1993, also in *Journal of High Speed Networks,* v.4, n.2, pp.201-219, 1995

[5] M.Ajmone Marsan, A.Fumagalli, E.Leonardi, F.Neri, "R-Daisy: an All-Optical Packet Network," *EUROPTO European Symposium on Advanced Networks and Services,* Amsterdam, The Netherlands, Mar.1995

[6] Y.Cai, *et al.*: "Demonstration of photonic packet-switched ring networks with optically transparent nodes", *IEEE Photonic Technology Letters,* v.6, n.9, pp.1139-1141, May 1994

[7] M.J.Chawki, *et al.*: "Wavelength reuse scheme in a WDM unidirectional ring network using a proper fiber grating add/drop multiplexer", *OFC '95,* San Diego, CA, Feb.1995

[8] H.Toba, K.Oda, K.Inoue, K.Nosu, T.Kitoh, "An Optical FDM-Based Self-Healing Ring Network Employing Arrayed Waveguide Grating Filters and EDFA's with Level Equalizers", *IEEE Journal on Selected Areas in Communications,* v.14, n.5, pp.800-813, June 1996

[9] Tsong-Ho Wu, "Fiber Network Survivability," *Artech House, Boston,* 1992.

[10] M.Ajmone Marsan, A.Bianco, E.Leonardi, M.Meo, F.Neri, "MAC Protocols and Fairness Control in WDM Multirings with Tunable Transmitters and Fixed Receivers", *IEEE/OSA Journal of Lightwave Technology,* v.14, n.6, pp.1230-1244, June 1996

[11] M.Ajmone Marsan, A.Bianco, E.Leonardi, M.Meo, F.Neri, "On the Capacity of MAC Protocols for All-Optical WDM Multi-Rings", *INFOCOM'96,* San Francisco, CA, Mar.1996

[12] M.Ajmone Marsan, A.Bianco, E.Leonardi, F.Neri, S.Toniolo "An Almost Optimal MAC Protocol for All-Optical WDM Multi-Rings with Tunable Transmitters and Fixed Receivers", *ICC'97,* Montreal, Canada, May 1997

[13] M.Ajmone Marsan, A.Bianco, E.Leonardi, F.Neri, S.Toniolo, "MetaRing Fairness Control Schemes in All-Optical WDM Rings", *INFOCOM'97,* Kobe, Japan, Apr.1997

[14] M.Ajmone Marsan, A.Bianco, E.Leonardi, A.Morabito, F.Neri, "SR^3: a Bandwidth-Reservation MAC Protocol for Multimedia Applications over All-Optical WDM Multi-Rings", *INFOCOM'97,* Kobe, Japan, Apr.1997