MAC Protocols and Fairness Control in WDM Multirings with Tunable Transmitters and Fixed Receivers

Marco Ajmone Marsan, *Senior Member, IEEE,* Andrea Bianco, Emilio Leonardi,
Michela Meo, and Fabio Neri

***Abstract*— The paper illustrates novel proposals for medium access control protocols in all-optical packet networks based on WDM multichannel ring topologies where nodes are equipped with one fixed-wavelength receiver and one wavelength-tunable transmitter. Such networks provide separate slotted channels for disjoint subsets of destination nodes. Three access protocols based on local status information are described. A channel inspection capability is assumed to be available for the implementation of the access protocols. Global fairness control algorithms derived from those adopted in the Metaring high-speed metropolitan area network are also proposed. Access delays and throughputs are taken as performance indices for a simulation-based comparison of the proposed protocols, in the case of a 16-node multiring with either balanced or unbalanced traffic. Simulation results show that the throughput limitations and the fairness problems inherent in the network topology can be overcome with relatively simple protocols.**

I. INTRODUCTION

IN ALL-OPTICAL networks, data generated in the electronic domain by the source user are converted into the optical domain at the user/network interface, and transported to the destination with no conversion into an electronic format before the network/user interface.

Most proposals of all-optical networks partition the optical bandwidth into a number of channels whose rates match the speeds of electronic interfaces. This is normally obtained with wavelength division multiplexing (WDM) [1].

Several of the recently proposed WDM network architectures require as many wavelengths as the number of network nodes (see [2] for a survey). This results in a dedicated channel for any node for either transmission or reception. Since the state-of-the-art optical technology allows the separation of a number of WDM channels that by itself is not sufficient for the implementation of large networks, WDM is often combined with either space or time diversity, or both (see, e.g., [3]).

In this paper we consider all-optical networks where nodes are equipped with one tunable transmitter and one fixed receiver. The recent rapid evolution of fast tunable optical transmitters gives particular relevance to this architectural solution.

With our all-optical network design, separate slotted channels are associated with disjoint subsets of destination nodes (from one channel for all destinations up to one channel for each destination). These logical channels can in general be obtained by a partition of time, frequency, space, code, or any combination of them. In our case, logical channels are assumed to be obtained with a combination of WDM and space diversity, since in this case the feasibility of the node hardware was verified considering state-of-the-art optical components. The proposed protocols remain valid also for other environments, but the technological feasibility of the required node hardware in other cases should be investigated.

We focus on the particular case of ring topologies, which have become an attractive solution for all-optical LAN's and MAN's [4]–[9], mainly thanks to the successful progress achieved in optical amplifiers to compensate for insertion losses at intermediate nodes. Rings allow for slot synchronization even at extremely high data rates, hence they offer an efficient and flexible use of the available optical bandwidth for packet communications.

The multichannel ring topology jointly exploits space and wavelength diversity in order to provide a number of channels. The number of wavelengths required in the WDM comb equals the number of nodes divided by the product of the number of parallel fibers in the multiring times the number of destinations in each subset. We always assume that the transmitted information is removed from the channel by the destination node.

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.

This paper proposes collision-free access strategies and protocols for all-optical packet networks built over multichannel ring topologies providing one slotted transmission channel for a subset of destination nodes.

The linearity of the ring topology naturally leads to the design of collision-free access protocols. These protocols will be based on the assumption that a channel inspection capability is available at each node. The constraints of the optical

---

Manuscript received April 17, 1995. This work was supported in part by the Italian National Research Council, and the Italian Ministry for University and Scientific Research.

The authors are with the Dipartimento di Elettronica, Politecnico di Torino, Torino 10129, Italy.

Publisher Item Identifier S 0733-8724(96)04905-5.

technology available today are considered in the design. In [10] and [11], for example, all-optical node architectures are described that allow the implementation of the MAC protocols proposed in this paper in the special case in which one channel is devoted to each destination. Prototypes are being built in research laboratories according to these proposals.

The proposed access schemes are compared by simulation, considering a 16-node network with either a uniform traffic distribution, or a hot-spot traffic scenario where one server interacts with several clients. Access delays are plotted when the input traffic does not overload the network, while throughputs are shown in overload conditions. These plots are provided for each source/destination pair, so that fairness issues can be discussed.

The paper is organized as follows. The next section provides more details on the considered networks, and describes three access strategies in the special case in which one channel is devoted to each destination. Section III presents simulation results for these access strategies, showing that fairness problems may arise unless a global fairness control algorithm is adopted. Section IV proposes some fairness control algorithms derived from those adopted in the Metaring [12] metropolitan area network (MAN). Section V demonstrates the effectiveness of the fairness control algorithms on the 16-node network used as a simulation scenario. Up to this point, results always consider the special case in which one channel is devoted to each destination. Section VI generalizes the access strategies to the case in which each slotted channel is shared by a number of destinations, and Section VII presents and discusses simulation results for this more general case. Finally, Section VIII concludes the paper. The Appendix contains an approximate analytical model for the validation of the simulation results presented elsewhere in the paper.

II. ACCESS PROTOCOLS

For the sake of simplicity, we initially consider multiring networks providing one logical channel for transmissions to each destination. Denoting by M the number of nodes in the network, and by W the number of logical channels, we have $W = M$.

All nodes that need to transmit to the same destination share the same logical channel.

Fig. 1 depicts the logical network topology for the case $M = W = 4$. The M logical channels run in parallel; they are slotted and synchronized, so that M slots (one for each destination node) simultaneously arrive at a node every slot time. Fixed-length data packets are transmitted, and the slot size is such that one packet exactly fits into one slot.

In order to avoid collisions (i.e., simultaneous accesses to the same slot) among packets directed to the same destination, nodes exploit a channel inspection capability on each logical channel. Thus the source node, prior to transmitting a new packet with a given destination, senses the activity on the channel associated with the particular destination. If an empty slot is sensed, the packet can be transmitted, and it will reach its destination without collision; otherwise, the transmission attempt is delayed.

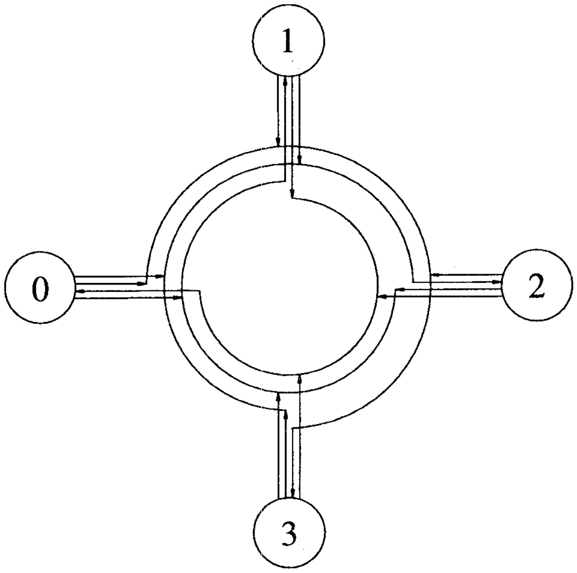

Fig. 1. The considered logical network topology.

It can be observed that, due to 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 other destinations. If we assume node numbers to be increasing in the transmission direction, the channel on which node i, $i = 0, 1, \cdots, M-1$, has the best access chance is the one leading to node[1] $|i-1|_M$, 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 $|i+1|_M$, since, to access this channel, node i must defer to the possible transmissions of $M-2$ upstream nodes. We assume nodes not to transmit to themselves.

In Fig. 1, 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.

We say that the traffic directed to node j from node i has lower priority than traffic directed to j from nodes $|j+1|_M, |j+2|_M, \cdots, |i-2|_M, |i-1|_M$. In particular, we say that node i has priority $|i-j|_M$ when transmitting to node j, 1 being the highest priority, and $M-1$ the lowest. When a packet is ready for transmission from node i to node j, the access must be delayed if the channel leading to j is already occupied with a transmission by nodes with higher priority with respect to destination j.

If propagation delays are neglected, an abstract model for the considered networks can be a $M \times M$ packet switch with input-buffering, and cyclic priorities.

The dependence of the access opportunity on the position along the multiring obviously raises fairness problems. One node can grab all the slots leading to a given destination, thus possibly starving downstream nodes competing for access to that channel.

---

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

Access protocols can be introduced to alleviate fairness problems and to guarantee a good utilization of the transmission resources. The slotted multiring easily allows the implementation of collision-free access protocols based on the channel inspection capability. Three alternatives are described below.

A. The FIFO MAC Protocol

The first MAC protocol that we propose for our slotted multiring is based on the assumption that only one logical first-in–first-out (FIFO) queue exists in each node for packets awaiting transmission. The packet at the head of the queue is considered for transmission at each slot. If it can be transmitted without collision, the packet is sent in the appropriate channel. Otherwise, the packet transmission attempt is delayed to the next slot. Packets are handled in a FIFO order: a packet is considered for transmission only after the successful transmission of previous packets.

This MAC protocol is particularly simple because no distributed synchronization algorithm is required. Indeed, each node simply monitors one channel at the beginning of the slot, thus using locally available information.

This FIFO MAC protocol does not affect the basic priority mechanisms of the considered networks. Again each node has a better-than-average access to some channels, and a worse-than-average access to other channels. Overall, under the assumption of uniform traffic destination, the protocol is fair, in the sense that it does not favor any particular source. However, fairness is not guaranteed if the traffic is not balanced.

This FIFO MAC protocol is an unconstrained slotted access to empty slots, with head-of-the-line blocking. The input queuing with blocking approach severely limits the achievable throughput to a fraction of the available network capacity. In order to quantify the network efficiency, it is possible to adapt the analytical model derived for the analysis of an input-buffering, space-division packet switch presented in [13]; the inspection of those results reveals that the fraction of the network capacity effectively used by the FIFO MAC protocol in uniform traffic conditions decreases when the number of nodes M increases. For $M \to \infty$, this fraction is about 58%.

Moreover, the FIFO protocol leaves room for an unfair use of the transmission resource, as we shall see.

B. The ARR MAC Protocol

Better performance and fairness can be obtained with a different MAC protocol, termed "asynchronous round robin" (ARR). Under balanced heavy load conditions this protocol forces the network behavior to a TDMA operation on every channel, thus succeeding in exploiting most of the available capacity, and providing a fair access.

In order to implement the ARR protocol, nodes must keep one separate logical queue for every possible destination. Thus, $M - 1$ packet queues are necessary at each node. Every queue handles packets in a FIFO order. The queues are cyclically scanned by the node controller, looking for a packet to transmit. At node i, after the successful transmission of a packet directed to j, hence taken from queue q_j (queues are indexed with the destination node identifiers), the next packet to be transmitted is chosen as the first packet in the first next nonempty queue. This means that, in the selection of the next packet to transmit, node i first considers queue $q_{|j+1|_M}$. If this queue is not empty, its first packet is selected for transmission; otherwise, queue $q_{|j+2|_M}$ is considered, and so on, until either a packet is selected for transmission in the next slot, or all queues are found to be empty. Once the packet is selected, its transmission is repeatedly attempted until a collision-free slot is found, thus introducing a form of head-of-the-line blocking.

Note that this ARR MAC protocol requires nodes to select the next packet to be transmitted only according to an empty/nonempty information for each queue (thus the decision is only based on an array of $M - 1$ boolean variables). This can be considered to be a marginal increase in complexity with respect to the former FIFO MAC protocol.

If all queues at all nodes are nonempty, i.e., if the network is under heavy load conditions, the transmission attempts operated by the nodes tend to a "round robin" state, and the overall system behavior tends to TDMA due to a self synchronization among the nodes' activities. At low loads, instead, no coordination among nodes occurs, and packets are handled almost in a FIFO order.

The ARR MAC protocol performance is expected to be superior to that of the FIFO protocol, in terms of both efficiency and fairness. The performance loss of this MAC protocol with respect to an ideal access protocol is mainly due to the loss of synchronization caused by skipping a queue in the round robin scanning when the network is under heavy load conditions. This event occurs either when an empty queue is found, or when node i transmits a packet directed to node $|i - 1|_M$, and schedules for the next slot the transmission of a packet directed to node $|i + 1|_M$, since nodes do not transmit to themselves. The MAC protocol that we describe next overcomes this drawback.

C. The SRR MAC Protocol

The goal of the third MAC protocol is the same as that of the ARR MAC protocol: under heavy load conditions it forces the network behavior to TDMA. Due to its tighter enforcement of the synchronization of the node operations this MAC protocol is termed "synchronous round robin" (SRR) MAC protocol.

Like for ARR, nodes must keep one separate packet queue for every possible destination. Thus, $M - 1$ FIFO queues exist at each node. Queues are cyclically scanned by the node, looking for a packet to transmit. The difference with respect to the ARR MAC protocol is that now in an arbitrary time slot identified by the label s, node i deterministically schedules for transmission a packet from a specific queue. If we introduce the integer variable $k = |s|_{M-1}$, that cycles over the range $0, 1, 2, \cdots, M - 2$, the queue with packets directed to destination $|i + k + 1|_M$ is selected for transmission in the slot labeled s. If such a queue is empty, the transmission of the first packet from the longest queue is attempted in the slot labeled s. If more than one longest queue exists (a tie for the longest queue occurs) the lowest priority longest queue is selected. In any case, if transmission in slot s is not possible because it would generate a collision, in the following slot (labeled $s+1$) a *new* packet is selected for the transmission attempt according to the SRR algorithm.

By so doing, the SRR MAC protocol maintains the schedule among transmissions, and its behavior under heavy load conditions is closer to TDMA than that of the ARR MAC protocol.

Note that in this case the selection of the next packet to be transmitted requires the information about the lengths of all queues, as well as a common knowledge by all nodes about the slot label s, i.e., a global (network-wide) synchronization on the slot sequence. Moreover, a selected packet not successfully transmitted in the current slot is sent back to its queue, while in the other two MAC protocols, once selected, a packet is certainly transmitted.

Performance of this SRR MAC protocol is expected to be superior to that of the two previous protocols, in terms of both efficiency and fairness.

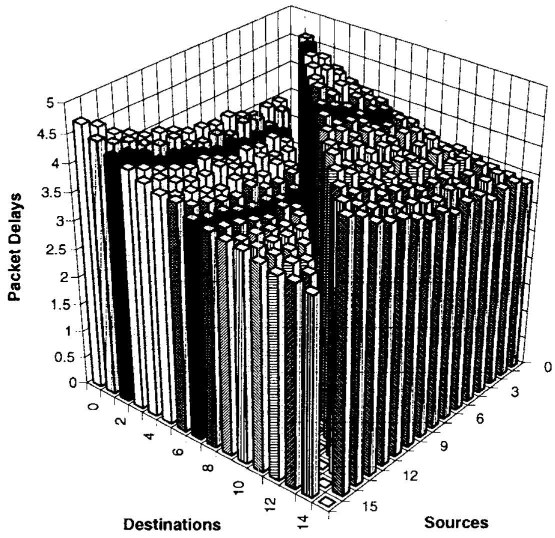

Fig. 2. Packet delays for the FIFO protocol with balanced traffic ($\Lambda = 8$).

III. Simulation Results—I

This section presents numerical results for a 16-node network where nodes have a large buffer to store packets awaiting transmission. Numerical results for the throughput and the packet delay distribution are obtained through simulation, referring to individual node-to-node traffic relations. The packet delay is the delay encountered by a packet prior to accessing the network, i.e., the time interval from the packet generation until the beginning of its successful transmission. We normalize delays to slot times, and express arrival rates and throughputs in packets per slot time.

Packets are assumed to always fit into one slot, and to be generated according to a Poisson process with fixed rate λ_i at node i. The total traffic in the network is $\Lambda = \sum_{i=0}^{M-1} \lambda_i$. The probability that a packet originated at node i is directed to node j is p_{ij}. We assume $p_{ii} = 0$.

We focus on very simple network scenarios, with the objective of providing an indication of the performance and fairness achievable with the proposed access protocols. Two traffic conditions are considered.

Balanced Traffic: All source nodes generate the same amount of traffic ($\lambda_i = \Lambda/M$, $\forall i$), and destinations are equally likely ($p_{ij} = 1/(M-1)$, $\forall i, j \colon i \neq j$).

Unbalanced Traffic: One "server" node (say node 0) is present in the network, the remaining $M-1$ "client" nodes direct half of their traffic to the server node, while the other half of the traffic is uniformly distributed among client nodes. The server generates an amount of traffic equal to half the total traffic generated by clients, and evenly distributes it to client nodes. In this case we use the arrival rates $\lambda_0 = \Lambda/3$ and $\lambda_i = (2/3)\Lambda/(M-1)$ $\forall i \neq 0$, while $p_{0j} = 1/(M-1)$ $\forall j \neq 0$, $p_{ij} = 1/[2(M-2)]$ $\forall i, j \neq 0$, $i \neq j$, and $p_{i0} = 1/2$ $\forall i \neq 0$.

We consider both cases where no channel is overloaded (thus no packets are lost due to overflows of node buffers), and where one or more channels are overloaded. In the first case, packet delays are mainly plotted; in the second case we focus on throughputs.

A. No Overloaded Channels

In the case of balanced traffic, the three MAC protocols are fair (in the sense that the rotational symmetry of the network guarantees the same performance at all nodes). Packet delays in the case of balanced traffic with a channel load equal to 50% ($\lambda_i = 0.5$ and $\Lambda = 8$) of the available capacity are shown in the 3-D graph of Fig. 2 for the FIFO MAC protocol. Delays are shown for every source/destination pair, using the same shading for delays observed on the same channel. On every channel the access delay increases for decreasing priority: this leads to largest and lowest delays near to the diagonal of the base matrix (due to rotational symmetries in the priorities). The average delays on each channel (i.e., for each destination) and for each source node are approximately the same.

Similar results (not shown here) were obtained with the other MAC protocols. ARR provides lower delays than the FIFO protocol, but causes larger differences between highest and lowest priority packets. SRR provides lowest delays, but largest differences. It is however important to mention that the lowest priority packets with the SRR protocol were observed to achieve better performance than the highest priority packets with the ARR protocol. The same holds true when FIFO and ARR are compared. This means that the ARR protocol brings an improvement over the FIFO protocol for all nodes, and that the SRR protocol brings and improvement over the ARR protocol for all nodes. Unfortunately, the improvement in performance is larger for some nodes and smaller for others, so that some unfairness still exists.

We now consider the unbalanced traffic case, with node 0 acting as a server. Nodes 1 through 15 direct half of their traffic to node 0, while the other half of the traffic is evenly distributed among the fifteen client nodes. The traffic generated by the server is evenly distributed among all client nodes. The total traffic in the network is taken to be $\Lambda = 2.7$, resulting in a

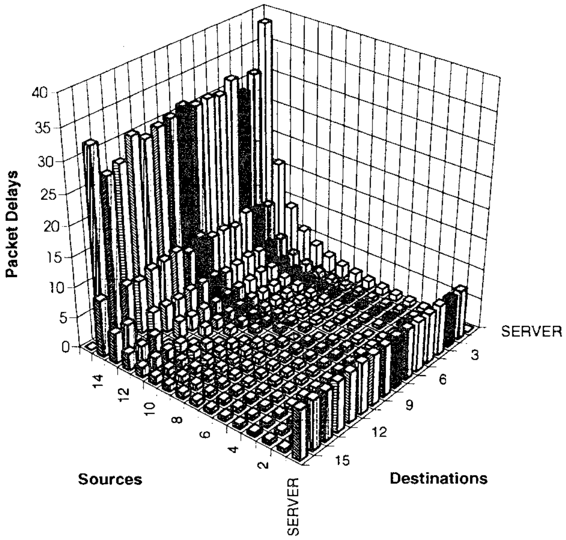

Fig. 3. Packet delays for the FIFO protocol with unbalanced traffic ($\Lambda = 2.7$).

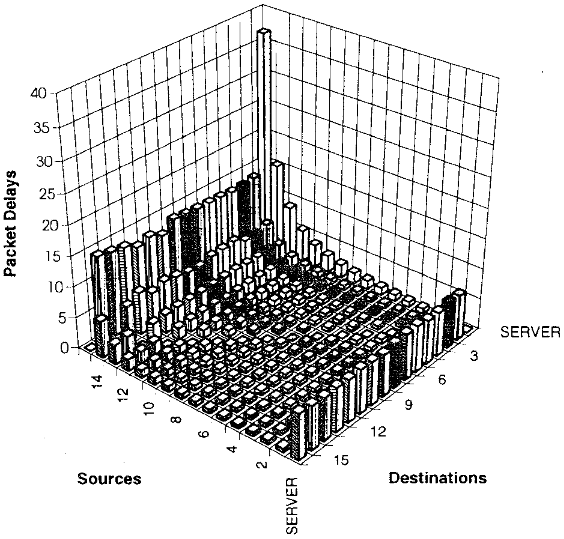

Fig. 4. Packet delays for the ARR protocol with unbalanced traffic ($\Lambda = 2.7$).

highly loaded (0.9) channel leading to the server. Client-bound channels are instead lightly loaded (0.12).

Figs. 3, 4, and 5, show results for the average packet delay for all source–destination pairs, with the FIFO, ARR, and SRR protocols, respectively.

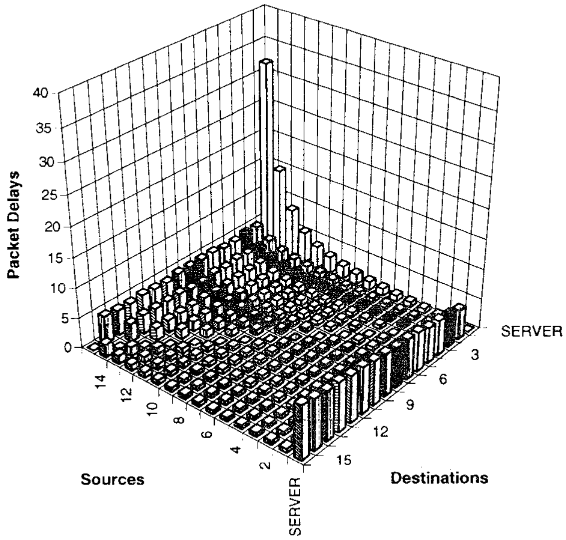

Fig. 5. Packet delays for the SRR protocol with unbalanced traffic ($\Lambda = 2.7$).

It can be easily seen that the FIFO protocol provides poor performance for all packets transmitted by nodes with high index, i.e., those nodes that have low priority when transmitting to the server. Note that the queuing delay is the same for all packets transmitted by the same source, hence packet delays differ only for the access delay (which in Fig. 3 is largest in the channel leading to the server); thus, for the FIFO protocol at high loads, all packets from the same source would suffer similar packet delays. The ARR protocol improves this situation, which is made even better by the SRR protocol, with hardly any delay penalty at other source nodes. However, the improvement is not significant in the channel leading to the server where the high load implies a long delay for the high index source nodes. On the other hand, it must be considered that with the considered load, the adoption of a fixed TDMA protocol in the channel leading to the server would result in delays of the order of 70 slots, i.e., about twice the largest delay obtained with the proposed approaches.

Note that the server faces an almost constant delay when transmitting (to clients), since the channels leading to the clients are lightly loaded. In the case of SRR, the server suffers larger delays at higher priority queues, since the access protocols favors lower priority queues when several longest queues are present.

B. Overloaded Channels

At low loads, the throughput is equal to the offered load (no packet losses are experienced). Fig. 6 shows the throughputs observed for unbalanced traffic with $\Lambda = 2.7$: a large amount of traffic is exchanged between the server and each client, and viceversa, while a smaller amount of traffic is exchanged among clients. Nodes do not transmit to themselves, thus no traffic is observed from i to i.

If we increase the traffic in the network, channels become overloaded and packets begin to be lost due to the finiteness of node buffers. Thus, while at low loads the throughput is equal to the offered load, it tends to a horizontal asymptote for increasing loads. For the FIFO protocol in balanced traffic, due to the finite number of nodes, the asymptote is slightly above 60% of the available capacity (it must tend to 0.58 for $M \to \infty$). For the ARR protocol the asymptote is close to 15/16, which shows the effect of synchronization loss whenever node i goes from destination $|i-1|_M$ to destination $|i+1|_M$. For the SRR protocol the asymptote is very close

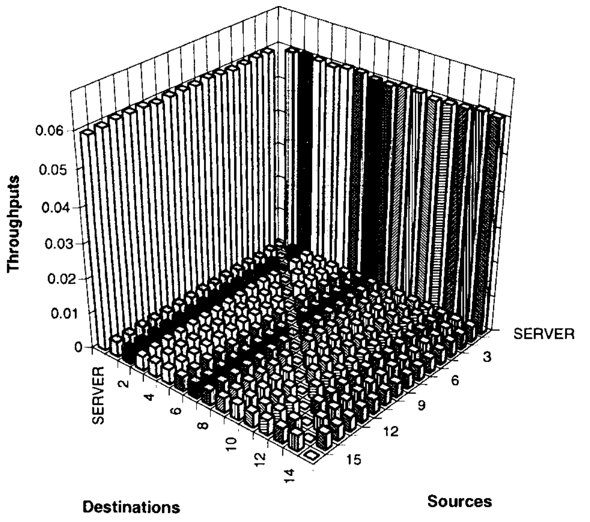

Fig. 6. Throughputs for the FIFO protocol with unbalanced traffic ($\Lambda = 2.7$).

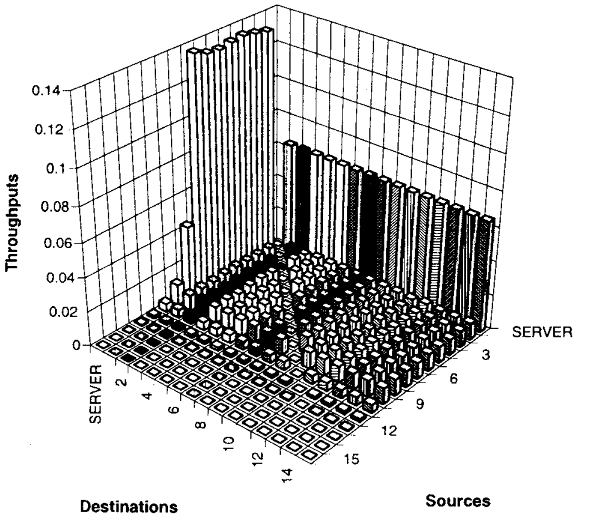

Fig. 7. Throughputs for the ARR protocol with unbalanced traffic ($\Lambda = 6$).

to 1. These limit throughputs (i.e., the protocol capacities) are shown in the second column of Table I.

Overloaded channels can show remarkable fairness problems. Figs. 7 and 8 show throughputs in the case of unbalanced traffic with overloaded server ($\lambda_0 = 2$ and $\Lambda = 6$) for the ARR and SRR MAC protocols. In either case the (overloaded) channel leading to the server (the white channel in the 3-D plots) shows starvation of lower priority nodes. Indeed, the nodes with the seven highest priorities use all the capacity available on the channel, and the other nodes are forced to remain silent. In the case of the ARR protocol, high index nodes obtain less throughput also on the other channels (which are not overloaded), due to the head-of-the-line blocking when attempting transmission on the server channel.

Little can be done against these unfair behaviors, unless some global fairness control algorithms are introduced. This is the topic of the next section.

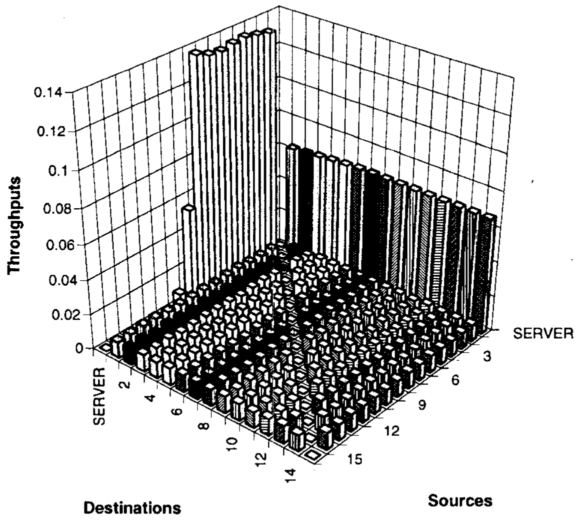

Fig. 8. Throughputs for the SRR protocol with unbalanced traffic ($\Lambda = 6$).

IV. FAIRNESS CONTROL

In order to avoid starvation of low priority nodes in overloaded channels, some global fairness control algorithm must be introduced. In this paper we adopt a mechanism originally proposed for the Metaring MAN.

Metaring [12], [14] is a MAN for the electronic domain, based on a dual ring topology, adopting a slotted access scheme and allowing slot reuse. Fairness in Metaring is granted by the circulation on the ring of a control message, named SAT, which normally rotates in the opposite direction (i.e., on the other ring) with respect to the data traffic it is regulating. Nodes are assigned a maximum number of packets to be transmitted between two SAT visits. Each node normally forward the SAT message on the ring with no delay, unless it is not SATisfied, in the sense that it has not transmitted the permitted number of packets since the last time it forwarded the SAT. The SAT is delayed at unSATisfied nodes until SATisfaction is obtained, in the sense that either the node packet buffer is empty or the number of permitted packet transmissions is achieved. The reader is referred to [12] for more details.

We adopt the fairness control mechanism of Metaring, with the modifications required by the different network topology and architecture. We call MMR (Multi-MetaRing) the modified algorithm. We call K the transmission quota assigned to each node between two SAT visits.[2] Each node can transmit at most K packets since the last time the SAT was forwarded; when this quota is reached, the node must refrain from transmission. The SAT is forwarded if either the node buffer is empty or K packets were transmitted since the last SAT visit.

Since the network architectures that we are considering in this paper only provide co-directional rings, the SAT message must propagate in the same direction of the regulated data. If

---

[2] The original Metaring proposal [12] defines two parameters: a node must refrain from transmitting after the usage of K slots if it is not holding the SAT, while it can transmit up to L packets while retaining the SAT. We consider the case $K = L$, since this guarantees better fairness (see [16]).

TABLE I
ACHIEVABLE THROUGHPUTS IN BALANCED TRAFFIC FOR THE PROPOSED MAC PROTOCOLS

	no global fairness	MMR-MS		MMR-SS	
		$K = 80$	$K = 1000$	$K = 1000$	$K = 5000$
FIFO	**0.605**	—	—	—	—
ARR	**0.937**	0.704	0.743	0.783	0.805
SRR	**0.999**	0.928	0.976	0.994	0.997

the SAT is forwarded to the upstream node, the forwarding node is in a good shape as far as access opportunity, since it is at the highest priority on the channel leading to the intended destination (i.e., to the upstream node); unluckily, SAT propagation delays are very large, since the SAT message must traverse almost the entire network.

Of course, it is possible to consider also the case where the SAT is co-directional with respect to the data it regulates, i.e., it is forwarded to the downstream node. In this case we would have more access delays and less propagation delays.

One further possibility is to devote one logical channel to SAT messages, in order to avoid or to reduce contention in SAT transmissions. This dedicated channel can be codirectional or counterrotating with respect to data propagation.

We do not take into consideration these options in this paper, restricting our attention to the case where SAT messages are forwarded to upstream nodes using the data channels, thus at high priority but with large propagation delays.

We assume that one slot is devoted to the transmission of the SAT message, i.e., some bandwidth is wasted for the propagation of control information. Since our networks provide several transmission channels, we can use either one "global" SAT message, or one separate SAT for each channel. In the first case, we have the MMR–SS (MMR Single SAT) protocol; in the second case we have the MMR–MS (MMR Multiple SAT) protocol.

As regards MMR–SS, we have two further options. At each node, we could associate the parameter K to all transmissions, disregarding destinations, or consider a separate quota for each channel.

As regards MMR–MS, there is a correlation between the forwarding of the SAT on one channel and the transmission of packets on other channels, in a way that is dependent on the basic access scheme (FIFO, ARR, or SRR). Indeed, while one node is delaying the SAT for one particular channel (because the quota K was not reached for that channel and more packets are waiting in the transmission queues), the access scheme may force the node to attempt transmissions on other channels (where some quota is still available). Several variations of the basic algorithms can be introduced at this point (for example, packets to be transmitted on those channels for which the SAT is being delayed could be given special attention), but we do not consider them here.

Note that the large SAT propagation delays prevent us from using low values for K, since the minimum SAT circulation time in slots on the rings is very large, and nodes with large transmission needs could be needlessly throttled. This means that fairness can be enforced only on a medium or large time scale. As a consequence, fairness can be obtained on throughputs, but not on delays, as we shall see in the next section.

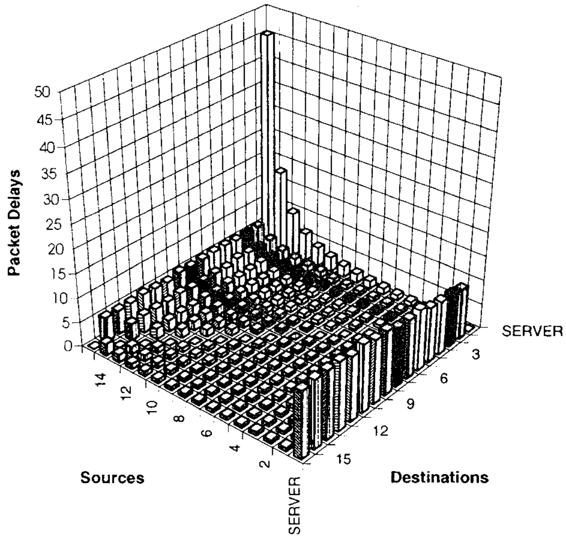

Fig. 9. Packet delays for the SRR+MMR–MS protocol with unbalanced traffic ($\Lambda = 2.7$).

V. SIMULATION RESULTS—II

Simulation results are provided in this section to study the superposition of the MMR algorithms onto the basic access protocols. An analytical model for a simple version of MMR is provided in Appendix A; it permitted the validation of the simulative approach.

In the case of balanced traffic, the limit throughputs obtained by simulation for ARR and SRR combined with MMR are shown in Table I. We did not consider FIFO in conjunction with MMR, mainly because MMR–MS requires a separation of packets to be transmitted according to destinations, while FIFO assumes one single packet queue. Note that MMR worsens the capacity of ARR, mainly due to the fact that SAT transmissions break the automatic synchronization mechanisms of ARR. The (marginal) capacity decrease due to the bandwidth devoted to SAT transmissions can also be observed from the table considering the SRR protocol; as expected, MMR–SS devotes less bandwidth to control.

In the case of no overloaded channels, the effects of MMR are negligible. As an example, Fig. 9 shows packet delays when the MMR–MS fairness control algorithm is added to the SRR access protocol, for the same scenario of Fig. 5. The MMR algorithm acts on a time scale which is much larger than access delays, so that no significant difference can be observed between the two figures (except for a marginal increase in delays due to the traffic increase because of SAT messages).

When we consider channel overloading, the fairness control algorithms come into play. Fig. 10 shows throughputs when MMR–MS is adopted in the same scenario of Fig. 8. Note that starvation is avoided, and bandwidth allocation is fair. Similar observations hold for the ARR MAC protocol, as shown in Fig. 11 (note again the slightly larger throughputs permitted by SRR).

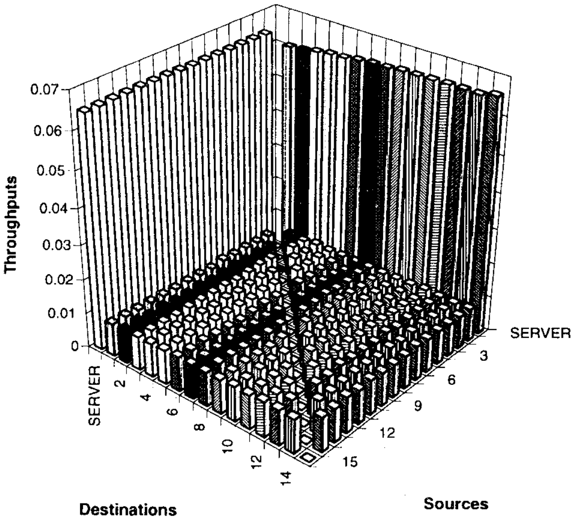

Fig. 10. Throughputs for the SRR+MMR-MS protocol ($K = 80$) with unbalanced traffic and $\Lambda = 6$.

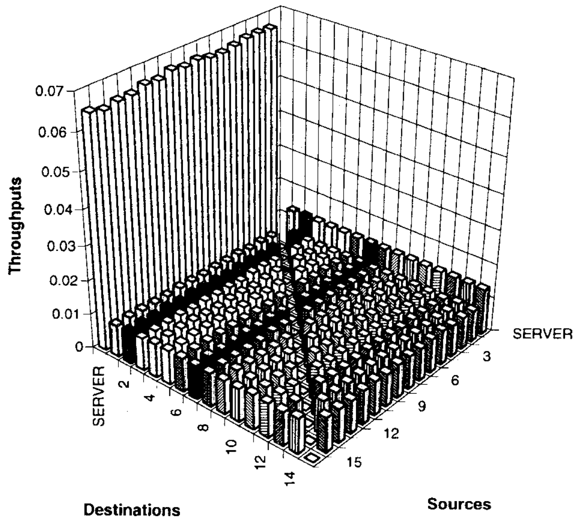

Fig. 12. Throughputs for the SRR+MMR-SS protocol ($K = 1000$ cumulative) with unbalanced traffic and $\Lambda = 6$.

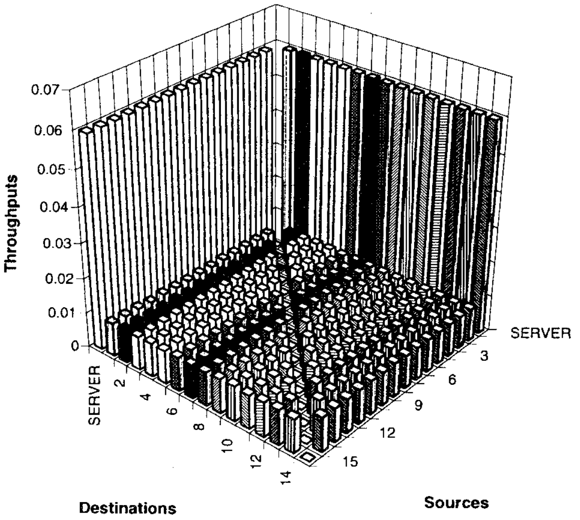

Fig. 11. Throughputs for the ARR+MMR-MS protocol ($K = 80$) with unbalanced traffic and $\Lambda = 6$.

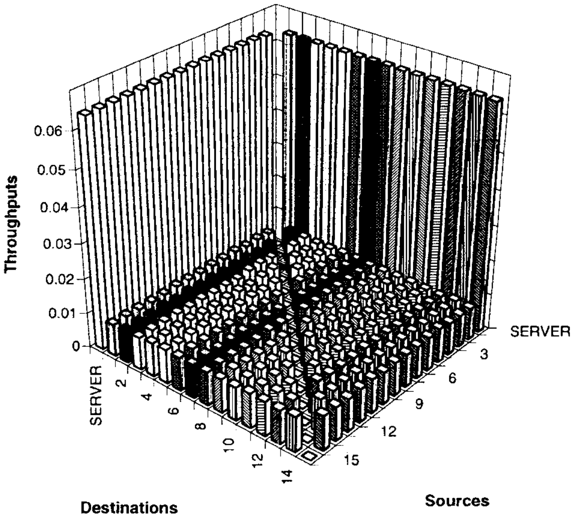

Fig. 13. Throughputs for the SRR+MMR-SS protocol ($K = 1000$ subdivided) with unbalanced traffic and $\Lambda = 6$.

Fig. 12 shows similar results when a single SAT is used, and the quota K is cumulative for the node. In this case the overall throughput of the server is forced to be equal to the overall throughput of each client, thus throttling the server also on channels that are not overloaded. This bandwidth waste is avoided in the single SAT case if the quota is partitioned between transmission channels. Fig. 13 shows throughputs in the same scenario, when one single SAT is propagated, but the quota K is partitioned among the $M - 1$ channels (nodes are SATisfied when, for all channels, either no packets are awaiting transmission, or $\lfloor K/15 \rfloor$ packets were transmitted since the last SAT release).

VI. ACCESS PROTOCOLS FOR $W < M$

While the previous sections considered only the case where one logical channel is devoted to each destination, the assumption that nodes are equipped with one fixed-wavelength receiver and one wavelength-tunable transmitter permits nodes to use logical channels to reach *disjoint subsets* of destinations.

This section considers the case $W < M$, i.e., the case where several destinations are reached through the same channel. Since we have a subdivision of the logical channels in space (through separate fibers) and frequency (through separate wavelengths), choosing $W < M$ allows a reduction in the number of fibers and/or in the number of wavelengths in each fiber with respect to the case $W = M$. Thus, choosing $W < M$ entails a reduction in the complexity of the system.

We always assume that the transmitted information is removed from the channel by the destination node, hence slots can be reused several times as they propagate along the channel. This slot reuse capability may entail additional complexity (slots need to be marked as free by the destination node), but we do not tackle these issues here.

The release of slots at their destination partially compensates for the resource reduction due to the decreased number of channels, since the available bandwidth on each channel is larger than the raw data rate. On the other hand, slot reuse raises fairness control problems. It is well understood that the additional capacity due to slot reuse is difficult to control: a node to which a large amount of information is directed generates a large amount of free slots and immediately downstream nodes are in a favorable position with respect to other nodes.

To alleviate fairness problems in uniform traffic conditions, we assume that the destination nodes of one logical channel are equally spaced (in terms of number of nodes) along the multiring. This means that, if $W = M/D$, the D destinations sharing the i-th channel are in positions $|i + dW|_M$, with $0 \leq d < D$.

The MAC protocols proposed in Section II can also be used in the case $W < M$ with minor modifications. Due to space limitations, in this section we only consider the SRR access protocol. More details on $W < M$ can be found in [15].

In the case $W = M$, SRR assumes at each node separate queues for each destination, i.e., for each logical channel. When $W < M$, we might ask either for separate queues for each destination, or separate queues for each logical channel. The latter choice is probably more consistent with the case $W < M$, since the complexity of the memory management inside each node is proportional to the number of logical channels and not to the number of nodes in the system.

In the case $W = M$, each node executing the SRR protocol cycles through the possible destinations, looking for a packet to transmit. When $W < M$, the round robin cycle can span over either the destinations (i.e., the cycle runs modulo $M - 1$), or the available channels (i.e., it runs modulo W). We chose the first option, i.e., to keep the round robin cycle on the destinations, similarly to $W = M$, mainly for fairness considerations.

This means that each node has a number of scheduled accesses to each logical channel equal to the number of destinations reached by that channel (one less for the channel where the node itself receives). If for example $W = M/2$, and each channel reaches two destinations ($D = 2$), each node has in each round robin cycle two scheduled accesses to each channel, except for the channel where it receives information, where it has only one scheduled access.

We consider the following three possible ways of adapting SRR to the case $W < M$.

SRR–FIFO: Nodes keep one packet queue for each logical channel (i.e., W queues). The round robin cycle of SRR selects one among the possible $M - 1$ destinations, which corresponds to one logical channel and one queue, from which the first packet is considered for transmission. Note that the packet considered for transmission might be directed to one of the other $D - 1$ destinations (this is true in the case of a channel for which the source node is not also a receiver, otherwise the other possible destinations are only $D - 2$; in order to simplify the description we refer to channels different from the one on which the source node is tuned for reception) that share the same channel and the same queue with the destination selected by the round robin schedule. If the selected queue is empty, the first packet of the longest queue is considered for transmission. If several longest queues exist, the first packet from the lowest priority longest queue (according to the definition of priorities for $W = M$) is selected.

SRR–RND: Nodes keep one packet queue for each destination (i.e., $M - 1$ queues). The round robin cycle of SRR selects one destination. This selected destination corresponds to one logical channel that, as usual, is shared with $D-1$ other destinations. This logical channel hence receives packets from D queues (from all the queues corresponding to the D destinations reached by the channel). The packet to be transmitted is chosen at random among the heads of these D queues. Again, the packet considered for transmission might be directed to one of the other $D-1$ destinations that share the same channel with the destination selected by the round robin schedule. If no packet can be selected according to the procedure described above, a packet is selected from the longest queue. If several longest queues exist, a packet from the lowest priority longest queue (according to the definition of priorities for $W = M$) is considered for transmission.

SRR–SRR: Nodes keep one packet queue for each destination (i.e., $M - 1$ queues). The same protocol as in the case $W = M$ is executed. This means that, on one logical channel, the same set of nodes has always a simultaneous scheduled access in a given slot of the round robin cycle.

The MMR fairness control can be adopted in the case $W < M$. We could consider MMR–MS and MMR–SS, but in the next section we restrict our attention only to the case MMR–SS, with a subdivided quota K. The subdivision of K can be either per logical channel (this is consistent with the SRR–FIFO protocol), or per destination.

VII. SIMULATION RESULTS—III

To study the case $W < M$, we report simulation results for a system with $M = 16$ nodes, and with $W = 8, 4, 2$ logical channels, i.e., respectively $D = 2, 4, 8$ destinations per logical channel. As already mentioned, we assume that the receivers on the same channel are evenly spaced along the multiring (this means that, with $W = M/2$, nodes 0 and 8 receive from channel 0; nodes 1 and 9 receive from channel 1, and so on).

In the case of balanced traffic, the limit throughput achievable by an ideal access protocol capable of granting fairness on a multiring can be analytically computed using the model presented in [15]. For $M = 16$, this limit throughput is equal to 1 when $W = 16$; it is 0.690 when $W = 8$, 0.425 when $W = 4$, and 0.239 when $W = 2$. Note that slot reuse partially compensates for the reduced transmission capacity of the system. Indeed, with $W = 8$ the system capacity is halved, but the throughput decreases by less than 40%; with $W = 2$ the capacity is divided by 8, while the throughput is reduced by a factor smaller than 5.

Simulation results show that SRR–FIFO and SRR–RND approach very closely this ideal behavior. It can be further noted that SRR–FIFO and SRR–RND in uniform overload

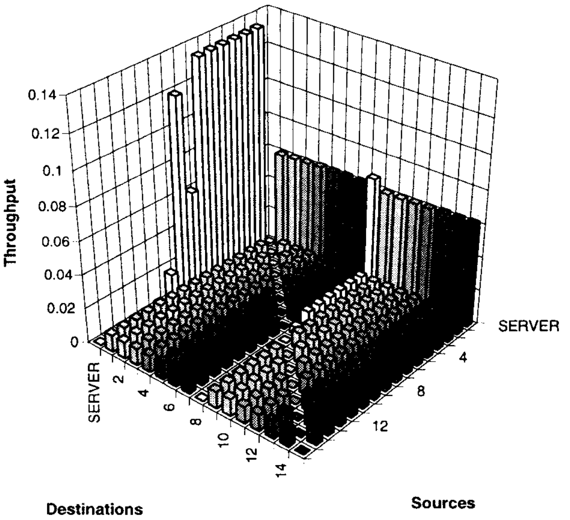

Fig. 14. Throughputs for the SRR–FIFO protocol with unbalanced traffic in the case $W = 8$.

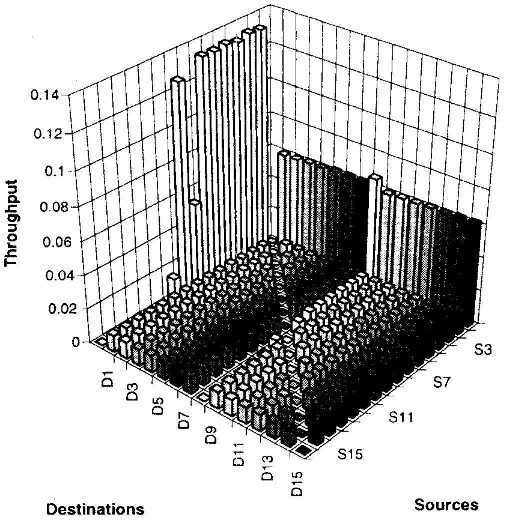

Fig. 15. Throughputs for the SRR–RND protocol with unbalanced traffic in the case $W = 8$.

behave very similarly when the MMR fairness control is not used. For SRR–SRR in sustained overload, larger throughputs can be obtained at the cost of severe unfairness. To explain this behavior, consider for example the case $W = M/2$, in which two receivers are present in each logical channel. Always the same two transmitters compete for the access to the slots in the same position of the SRR cycle. Spatial reuse is possible only when the two transmitters must reach the closest receiver. Otherwise, only one of the two transmitters can use the slot. With SRR–FIFO and SRR–RND, each source can independently address one of the two receivers with the same probability, hence space reuse is possible with probability 0.25, and no reuse is possible with probability 0.75. With SRR–SRR, the transmitters have a synchronized access, hence either both address the nearer receivers (with probability 0.5), or both address the farther receivers (again with probability 0.5). Thus more spatial reuse, and larger throughputs are observed. The cost of the larger throughput is that one of the two transmitters deterministically wins (i.e., uses the slot) when no reuse is possible, while the other transmitter is completely starved for transmission to the particular destination. This is true in sustained overload; more fairness is obtained when the SRR scheduling is broken due to empty queues.

In the unbalanced traffic scenario, with node 0 acting as a server, we report results for a total traffic in the network $\Lambda = 6$, i.e., for the overload situation considered in the previous results sections. Note that reducing W means offering more traffic to the channels: while the channel leading to the server has an offered traffic equal to 2 when $W = 16$, it has an offered load equal to 2.27 when $W = 8$, and equal to 3.87 when $W = 2$.

Figs. 14–16 show the throughputs for each source/destination pair in the case $W = 8$, for SRR–FIFO, SRR–RND, and SRR–SRR, respectively. We used the same shading for throughputs obtained on the same logical channel.

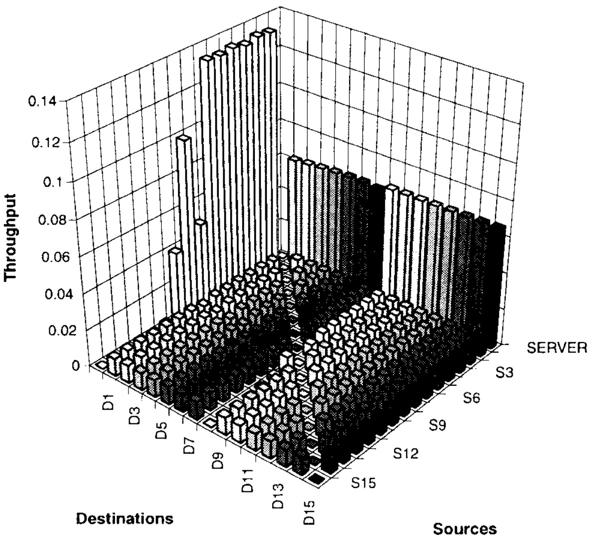

Fig. 16. Throughputs for the SRR–SRR protocol with unbalanced traffic in the case $W = 8$.

The channel where the server receives is highly overloaded. While nodes far from the server can use slots at will, nodes close to the server suffer starvation, as already noted in Fig. 8. Since we are considering the case $W = 8$, the destination node 8 shares the channel with the server. Starvation can thus be observed also for transmissions to 8 from nodes closer (i.e., at lower priorities for transmissions) to the server. The server, as a source of packets directed to node 8, takes advantage of slot reuse, hence gains a slightly higher throughput with respect to transmissions to the other clients.

The advantage of slot reuse is evident in the channel leading to the server and to node 8. Note indeed that node 8 has a good throughput on that channel, while node 7 is starving. This gain is dependent on the access protocol: SRR–FIFO and SRR–SRR allow node 8 to use less slots made available by slot reuse, and more throughput gain is observed by nodes 9 and 10.

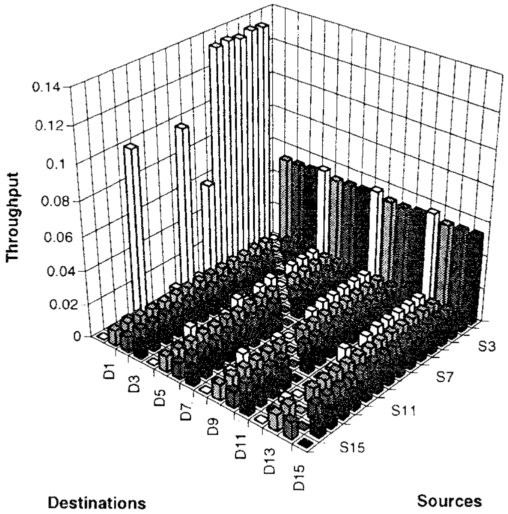

Fig. 17. Throughputs for the SRR–RND protocol with unbalanced traffic in the case $W = 4$.

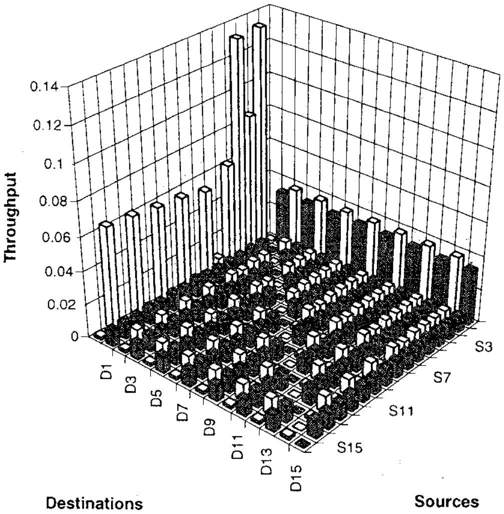

Fig. 18. Throughputs for the SRR–RND protocol with unbalanced traffic in the case $W = 2$.

A reason for the different exploitation of slot reuse is as follows. Let's focus on the channel leading to the server and to node 8. Node 8 on this channel only transmits to the server; it has a scheduled access only once per round robin cycle (while the other nodes have two scheduled accesses to this channel per cycle). All nodes (except the server) are in heavy load condition for what regards transmissions to the server (and possibly to node 8), while they have lightly loaded transmission queues for the other destinations. Thus, when access is scheduled from a lightly loaded queue, the queue might be found empty, and the access protocol asks the transmitter to select the packet from the longest queue, i.e., to transmit on the server channel. Therefore it quite often happens that transmitters break the round robin scheduling, thus "stealing" scheduled slots from downstream nodes. The impact of this behavior depends on the particular access protocol. For SRR–SRR, the round robin cycle is broken when a queue for one destination is found empty; this happens more often than for SRR–RND and SRR–FIFO, where the scheduling is broken when no packets are found for *one channel* (i.e., for D destination nodes). Hence SRR–SRR allows less throughput to node 8 on the server channel.

SRR–FIFO, on the other hand, maintains one single queue for packets directed to the server and to node 8. This queue is overloaded; this means that some packets with destination 8 are lost, and that the number of packets received by node 8 decreases with respect to SRR–RND. Moreover, when the SRR scheduling is broken due to an empty queue, the access protocol attempts transmission of a packet taken from the longest queue. The longest queue contains only packets directed to the server for SRR–RND, but also packets directed to node 8 for SRR–FIFO. Thus, some of the packets directed to 8 are not transmitted according to the SRR scheduling, and cannot be reused by node 8, which is scheduled for transmission to another destination in those slots. These are two reasons for the fact that SRR–FIFO shows slightly less reuse by node 8 on the server channel with respect to SRR–RND.

TABLE II
ACHIEVABLE THROUGHPUTS IN BALANCED TRAFFIC FOR SRR + MMR–SS WHEN $W < M$

	$W = 16$	$W = 8$	$W = 4$	$W = 2$
SRR–FIFO	0.999	0.650	0.403	0.231
SRR–RND	0.999	0.668	0.411	0.233
SRR–SRR	0.999	0.621	0.395	0.228
ideal ring	1.000	0.690	0.425	0.239

Figs. 17 and 18 show the throughputs for each source/destination pair for the SRR–RND protocol, in the cases $W = 4$, and $W = 2$, respectively.

For lower values of W, the phenomena observed for $W = 8$ become more evident. Remember that the load offered to the channels increases, since the traffic between each source/destination pair is kept constant, but the number of channels is reduced. Starvation appears for transmission to all destinations that share the channel with the receiving server. Slot reuse is effectively exploited by these destinations when transmitting on the most overloaded channel.

We now consider the use of MMR fairness control when $W < M$. In the case of balanced traffic, the throughputs achievable by the three variations of SRR in conjunction with MMR–SS with large values of K are shown in Table II. The table also shows the limit throughputs for the ideal access protocol, taken from [15]. Note that SRR–RND best approaches the ideal behavior. For SRR–SRR, and SRR–FIFO, enforcing fairness has a cost in terms of throughput.

The MMR fairness control with unbalanced traffic is studied in the case $W = 4$. Fig. 19 shows the throughputs for each source/destination pair for SRR–FIFO in the case $W = 4$, when MMR–SS with subdivided K controls fairness. The subdivision of the quota K is per channel consistently with the definition of the SRR–FIFO protocol, whose complexity is proportional to the number of channels.

The subdivision of the quota per logical channel prevents a fair access to the channels, although a remarkable improvement with respect to Fig. 17 is observed. Nodes sharing the

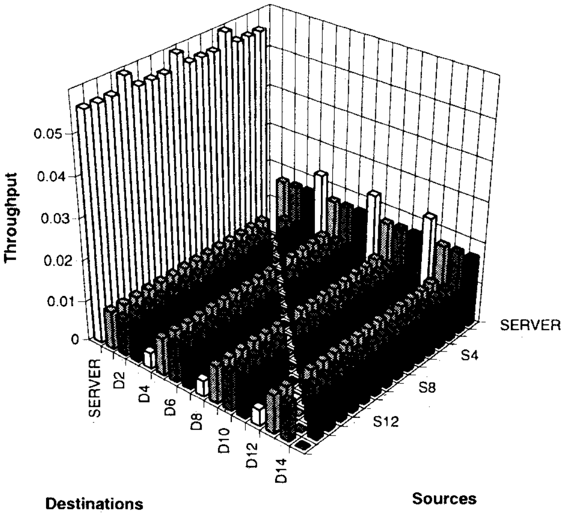

Fig. 19. Throughputs for the SRR–FIFO + MMR–SS protocol with unbalanced traffic in the case $W = 4$.

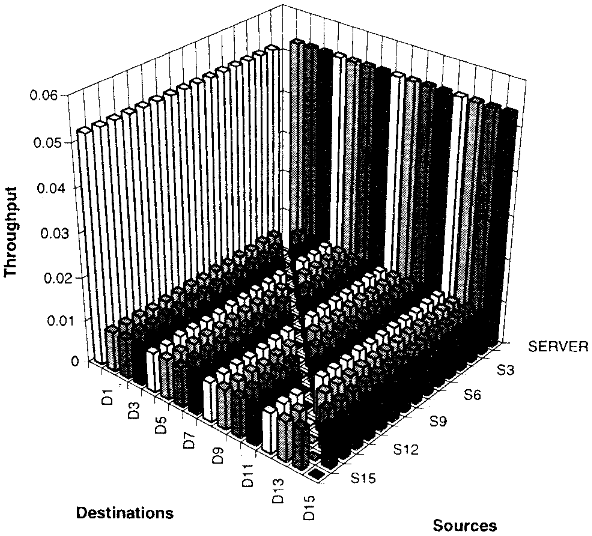

Fig. 20. Throughputs for the SRR–RND + MMR–SS protocol with unbalanced traffic in the case $W = 4$.

reception channel with the server obtain lower throughputs, due to a larger loss probability in the finite queue congested by packets directed to the server. Server transmissions are limited by a larger cycle time with respect to SRR-RND. This increase in the cycle time is due to the sharing of transmission quotas by all sources sharing the FIFO queue devoted to transmissions on the server channel, which results in more transmissions from client to server than from server to client. Note that nodes receiving on the server channel obtain larger throughputs on this channel, since they are granted the same quota but transmit to fewer (i.e., to $D - 1$) destinations.

Fig. 20 is similar to Fig. 19, but the SRR–RND access protocol is considered, and the quota K is subdivided per destination. The subdivision per destination successfully achieves fairness, and an ideal throughput partitioning is obtained. A similar behavior was observed also for SRR–SRR.

VIII. CONCLUSION

With the rapid success of optical amplifiers, WDM ring architectures are entering the arena of all-optical network architectures capable of providing transmission support to future bandwidth-greedy applications. As a direct consequence, the well-known and well-established properties of classic ring networks can be adapted for the benefit of all-optical networks.

This paper presented three access protocols for collision-free multichannel rings with slotted packet transmissions. The proposed access protocols were studied through simulation, showing that significant improvements of the throughput limitations deriving from head-of-the-line blocking, and of the unfairness generated by the network architecture can be obtained.

In the case of channel overloading, fairness can be enforced by suitable fairness control algorithms. We proposed the MMR family of algorithms, derived from the Metaring MAN. These global fairness control algorithms nicely fit the considered network architectures, and provide the desired throughput performance at a reduced cost. Several variations of the proposed fairness control algorithms are possible, and left for future investigation.

APPENDIX A
AN APPROXIMATE ANALYTICAL MODEL

An approximate analytical model was developed for a simple version of the access protocol and fairness control scheme.

As in previous sections, denote by M the number of nodes in the multiring and by W the number of channels. We restrict here our attention to the case $W = M$, i.e., one logical channel is available for each destination node, and to MMR–MS, i.e., separate SAT messages on each channel. Extensions to $W < M$ and MMR–SS are presented in [15]. Let K be the number of packets that each node is allowed to transmit by the MMR–MS fairness control scheme between two SAT departures on each channel.

Assume that each node has a separate packet queue for each destination. We associate with every queue a priority, equal to the priority of the source node in the transmission to the destination corresponding to the queue. A queue at a node is *blocked* if the node has already transmitted K packets from the queue since the last transmission of the SAT on the channel. We call Δ the fixed time required to pass the SAT message from node to node.

In order to simplify the analysis, assume that each node selects at random, according to a uniform distribution, the queue from which a packet transmission is attempted in the next slot (this case is simpler to analyze than the FIFO, ARR, or SRR protocols). The random selection considers nonempty queues that are not blocked due to the exhaustion of the credit of the MMR–MS scheme.

The aim of the analytical model is to evaluate the throughput achievable under uniform overload traffic conditions. Under these conditions every node exhibits the same overall average behavior, but the individual queue performances depend upon the queue priority, i.e., upon the position of the node on the multiring.

For any given node define the following quantities that refer to priority i, $i = 1, \cdots, M-1$:

- C_i is the average cycle time at priority i, that is the average time elapsing between two subsequent instants in which a node transmits the SAT on the channel used by the ith priority queue; in uniform overload conditions, the cycle time is independent of i;
- Pb_i is the probability that the ith priority queue of a node is blocked, equal to the average fraction of the average cycle time in which the queue is blocked;
- τ_i is the average activity period within a cycle for the ith priority queue of a node, that is the average interval between the instant in which the node transmits the SAT and the instant in which the i-priority queue credit at that node exhausts;
- R_i is the average rate of transmission of ith priority packets during the activity period of their queue: $R_i = K/\tau_i$ due to the overload assumption;
- X_i is the average throughput of ith priority queue packets, $X_i = K/C_i$. Note that $X_i \leq R_i$.

We introduce two simplifying assumptions in order to make our model analytically tractable.

- The state of each queue in a node is statistically independent from the states of the other queues in the same node.
- The state of each queue in a node is statistically independent from the states of the queues of other nodes containing packets directed to the same destination.

The second assumption in particular entails a drastic simplification, as known from the study of polling models [17]. We are losing the direction of SAT propagation: queue independence means random SAT cycles.

The approximate analysis technique is iterative. The algorithm to be followed for the computation of the quantities X_i is the following:

1. provide initial values to the probabilities Pb_i, $i = 1, \cdots, M-1$, and to the overall throughput $\bar{X}$;
2. evaluate the rates R_i, $i = 1, \cdots, M-1$;
3. evaluate X_i, $i = 1, \cdots, M-1$;
4. evaluate C_i, $i = 1, \cdots, M-1$;
5. check for convergence, and eventually stop;
6. compute the new values of the probabilities Pb_i, $i = 1, \cdots, M-1$, and return to Step 2.

In the following subsections we separately describe the key steps of our iterative approach.

A. Evaluation of R_i

Consider a given node and use the current values for the probabilities Pb_i. We want to evaluate R_i, the average transmission rate for ith priority packets in a node, given that the corresponding queue is not blocked. R_i is obtained as the product of the probability Ps_i that a packet from the ith priority queue is chosen for transmission and the probability Pt_i that the channel on which the packet should be transmitted is free

$$R_i = Ps_i \, Pt_i.$$

The value of Ps_i is obtained from the total probability theorem.

When no queue is blocked, Ps_i equals $1/(M-1)$ where $M-1$ is the number of queues in the node. The probability that no queue is blocked, given that queue i is not blocked, is

$$\begin{aligned}(1-Pb_1)\cdots(1-Pb_{i-1})&(1-Pb_{i+1})\cdots(1-Pb_{M-1})\\ &= \prod_{j=1, j\neq i}^{M-1} (1-Pb_j).\end{aligned}$$

When one of the queues (different from the one with ith priority packets) is blocked, Ps_i equals $1/(M-2)$, and the probability that one such queue is blocked is

$$\begin{aligned}&Pb_1(1-Pb_2)(1-Pb_3)\cdots(1-Pb_{i-1})(1-Pb_{i+1})\\ &\qquad\cdots(1-Pb_{M-1})+\\ &(1-Pb_1)Pb_2(1-Pb_3)\cdots(1-Pb_{i-1})(1-Pb_{i+1})\\ &\qquad\cdots(1-Pb_{M-1})+\\ &\vdots\\ &(1-Pb_1)(1-Pb_2)\cdots(1-Pb_{i-1})(1-Pb_{i+1})\\ &\qquad\cdots(1-Pb_{M-2})Pb_{M-1}\\ &= \sum_{k=1, k\neq i}^{M-1}\left[Pb_k \prod_{j=1, j\neq k, j\neq i}^{M-1} (1-Pb_j)\right].\end{aligned}$$

Similarly, when two queues are blocked, Ps_i equals $1/(M-3)$, and to evaluate the probability that two queues are blocked we add the probabilities that each one of all the possible pairs of queues is blocked.

We can proceed accordingly for all possible cases.

In order to write the final formula in a compact way we define the following shorthand notation

$$\beta_i^0 = Pb_i$$
$$\beta_i^1 = 1 - Pb_i$$

so that we can write

$$Ps_i = \sum_{a_1=0}^{1}\sum_{a_2=0}^{1}\cdots\sum_{a_{i-1}=0}^{1}\sum_{a_{i+1}=0}^{1}\cdots\sum_{a_{M-1}=0}^{1} \frac{\beta_1^{a_1}\beta_2^{a_2}\cdots\beta_{i-1}^{a_{i-1}}\beta_{i+1}^{a_{i+1}}\cdots\beta_{M-1}^{a_{M-1}}}{1+a_1+a_2+\cdots+a_{i-1}+a_{i+1}+\cdots+a_{M-1}}.$$

The value of Pt_i is easily obtained.

A packet chosen for transmission is actually transmitted only if the channel is free, i.e., if no other node with higher priority on the same channel has transmitted a packet. Pt_1 is 1 because the node has highest priority on this channel. Pt_2 equals $1 - X_1$, since the node can transmit only if the one at higher priority does not transmit, and so on. We obtain

$$Pt_1 = 1$$
$$Pt_i = 1 - \sum_{j=1}^{i-1} X_j \qquad \forall i \neq 1.$$

B. Evaluation of X_i and C_i

Since a node transmits packets from the ith priority queue with rate R_i during its activity period and it transmits no packet when the queue is blocked, the average throughput of the ith priority queue is given by

$$X_i = R_i\,(1 - Pb_i) \tag{1}$$

while the total node throughput is

$$X = \sum_{i=1}^{M-1} X_i.$$

Since in each cycle a node transmits K packets to each destination, we have $X_i = X/(M-1)$ for every i. This also implies that X_i does not depend on i, although the proposed iterative algorithm provides not identical X_i's before convergence. The aggregate value X must be therefore used in subsequent formulas.

The node activity period can be computed as $(M-1)K/X$. For the computation of C_i, we have a further constraint on the SAT propagation time over the ring $M\Delta$

$$C_i = \max\left[\frac{(M-1)\,K}{X}, M\Delta\right].$$

C. Evaluation of Pb_i

Pb_i can be obtained by dividing the average period during which the ith priority queue is blocked by the average cycle time.

$$Pb_i = \frac{C_i - \tau_i}{C_i} = 1 - \frac{K}{C_i R_i}.$$

When K is large, the probability Pb_i must tend to zero for the lowest priority queue.

D. Validation of the Model

The iterative algorithm described in this appendix showed reasonably fast convergence. The probabilities Pb_i were initialized to 0, while the overall throughput X was initialized to $M/2$. The check for convergence was based on the difference between the values taken by X_{ij} in subsequent iterations.

Table III compares, for a system with $M = W = 16$ and node-to-node distances equal to 15 slot times, the limit throughputs obtained by simulation and with the analytical models for varying values of the quota K. The SAT is supposed to be transmitted to the upstream node with no access delays using a dedicated control channel, hence the propagation time Δ of the SAT on the ring is 15 × 16 slot times.

A threshold behavior is observed for varying values of K: when K is small, each queue exhausts its quota before the SAT returns, some capacity is lost, and the network behaves like a token protocol. When K is above a threshold value, the full utilization of the ring is obtained, and practically no cost has to be paid for fairness control. The threshold value depends on propagation delays, and on the number of available channels.

TABLE III
COMPARISON BETWEEN SIMULATION AND ANALYTICAL RESULTS FOR MMR–MS, IN THE CASE $M = W = 16$

K	simulation	analysis
2	0.1181	0.1172
4	0.2354	0.2343
8	0.4666	0.4687
10	0.5253	0.5851
15	0.5611	0.6457
20	0.5872	0.6457
50	0.6007	0.6457
100	0.6039	0.6457
200	0.6056	0.6457
500	0.6087	0.6457
1000	0.6147	0.6457

The analytical results are in a reasonable agreement with simulation results, although the latter show a smoother transition around the threshold. This provides confidence on the correct simulator behavior in this case as well as in the case of more complex protocols that could be studied only through simulation.

REFERENCES

[1] C. A. Brackett, "Dense wavelength division multiplexing networks: Principles and applications," *IEEE J. Select. Areas Commun.*, vol. 8, pp. 948–964, Aug. 1990.

[2] B. Mukherjee, "WDM-based local lightwave networks—Part I: Single-hop systems," *IEEE Network*, vol. 6, pp. 12–27, May 1992.

[3] C. B. Brackett *et al.*, "A scalable multiwavelength multihop optical network: A proposal for research on all-optical networks," *J. Lightwave Technol.*, vol. 11, pp. 736–753, May/June 1993.

[4] W. I. Way *et al.*, "Self-routing WDM high-capacity SONET ring network," in *OFC'92*, San Jose, CA, Feb. 1992.

[5] A. Hamel *et al.*, "Multilayer add-drop multiplexers in a self-healing WDM ring network," in *OFC'95*, San Diego, CA, Feb. 1995.

[6] M. J. Chawki *et al.*, "Demonstration of a WDM survivable open ring network using reconfigurable dropping receivers," in *ECOC'94*, Firenze, Italy, Sept. 1994.

[7] J. Pietzsch, "Optical access networks with HDWM-race 2070 MUNDI," in *OPNET'95*, Paris, France, Jan. 1995.

[8] Y. Cai *et al.*, "Demonstration of photonic packet-switched ring networks with optically transparent nodes," *IEEE Photon. Technol. Lett.*, vol. 6, pp. 1139–1141, May 1994.

[9] M. J. Chawki *et al.*, "Wavelength reuse scheme in a WDM unidirectional ring network using a proper fiber grating add/drop multiplexer," in *OFC'95*, paper ThI4, San Diego, CA, Feb. 1995.

[10] M. Ajmone Marsan, A. Fumagalli, E. Leonardi, and F. Neri, "R-daisy: An all-optical packet network," in *EUROPTO European Symp. Advanced Networks Services*, Amsterdam, The Netherlands, Mar. 1995.

[11] I. Chlamtac, A. Fumagalli, L. G. Kazovsky, and P. T. Poggiolini, "A Multi-Gbit/s WDM optical packet network with physical ring topology and multi-subcarrier header encoding," in *ECOC'93*, Montreux, Switzerland, Sept. 1993.

[12] I. Cidon and Y. Ofek, "MetaRing—A full-duplex ring with fairness and spatial reuse," *IEEE Trans. Commun.*, vol. 41, pp. 110–120, Jan. 1993.

[13] M. J. Karol, M. G. Hluchyj, and S. P. Morgan, "Input versus output queueing on space-division packet switch," *IEEE Trans. Commun.*, vol. COM-35, pp. 1347–1356, Dec. 1987.

[14] J. Chen, I. Cidon, and Y. Ofek, "A local fairness algorithm for the metaring and its perfomance study," *IEEE J. Select. Areas Commun.*, vol. 11, pp. 1183–1192, Oct. 1993.

[15] M. Ajmone Marsan, A. Bianco, E. Leonardi, M. Meo, and F. Neri, "On the capacity of MAC protocols for all-optical WDM multi-rings," in *INFOCOM'96*, San Francisco, CA, Mar. 1996.

[16] G. Anastasi, L. Lenzini, and P. Motta, "Performance evaluation of a metaring MAC protocol integrating video and data traffics in an interconnected environment," in *INFOCOM'95*, Boston, MA, Apr. 1995.

[17] H. Takagi, *Analysis of Polling Systems*. Cambridge, MA: M.I.T. Press, 1986.

Marco Ajmone Marsan (SM'86) was born in Torino, Italy, in 1951. He received the Dr. Ing. degree in electronic engineering from Politecnico di Torino, and the Master of Science degree from the University of California, Los Angeles.

He is a Full Professor at the Dipartimento di Elettronica of Politecnico di Torino, in Italy. From November 1975 to October 1987, he was at the Dipartimento di Elettronica of Politecnico di Torino, first as a Researcher, then as an Associate Professor. From November 1987 to October 1990, he was a Full Professor at the Computer Science Department of the University of Milan, in Italy. During the summers of 1980 and 1981, he was with the Research in Distributed Processing Group, Computer Science Department, UCLA. He has coauthored the two books "Performance Models of Multiprocessor Systems," (Cambridge, MA: M.I.T. Press) and "Modelling with Generalized Stochastic Petri Nets," (New York: Wiley). His current interests are in the fields of performance evaluation of data communication and computer systems, communication networks, and queueing theory.

Dr. Ajmone Marsan received the Best Paper Award at the Third International Conference on Distributed Computing Systems in Miami, FL, in 1982. His email address is ajmone@polito.it

Andrea Bianco received the Dr.Ing degree in electronic engineering and the Ph.D. degree in telecommunications engineering, both from Politecnico di Torino.

He is a Research Assistant at the Dipartimento di Elettronica of Politecnico di Torino. Since May 1991 to December 1994, he was with the Dipartimento di Elettronica of Politecnico di Torino as a Ph.D. student. In 1993, he visited Hewlett–Packard Laboratories, Palo Alto, CA, doing research on the performance analysis of TCP protocol over ATM networks. From January to November 1995, he was with the Dipartimento di Sistemi di Produzione ed Economia dell'Azienda of Politecnico di Torino. Since November 1995, he has been with the Dipartimento di Elettronica of Politecnico di Torino. His current research interests are in the field of access protocols for all-optical networks, performance analysis of ATM networks, and formal description techniques. His email address is: bianco@polito.it

Emilio Leonardi received the Dr.Ing. degree in electronic engineering and is currently working toward the Ph.D. degree in the Dipartimento di Elettronica of Politecnico di Torino, Turin, Italy.

Since September 1994, he spent one year at the Computer Science Department of the University of California at Los Angeles (UCLA), working on deadlock-free routing and flow control techniques for asynchronous high speed wormhole networks. His current research interests are in the fields of performance evaluation of high-speed and all-optical networks, and of queuing theory. His email address is: leonardi@polito.it

Michela Meo received the Dr.Ing. degree in electronic engineering from Politecnico di Torino in July 1993. Since November 1993, she has been working toward the Ph.D. degree at the Dipartimento di Elettronica of Politecnico di Torino.

In 1994, she visited the University of California at Santa Barbara doing research on performance analysis of packet radio networks. Her current research interests are in the fields of access protocols for all-optical networks, algorithms for the efficient solution of Markov chains and formal description techniques. Her email address is: michela@polito.it

Fabio Neri received the Dr. Ing. and Ph.D. degrees in electronic engineering from Politecnico di Torino in 1981 and 1987, respectively.

He is currently an Associate Professor at the Dipartimento di Elettronica of Politecnico di Torino, Turin, Italy. He was born in Novara, Italy, in 1958. From 1991 to 1992, he was with the Information Engineering Department at University of Parma, Parma, Italy, as an Associate Professor. From 1982 to 1983, he was a visiting scholar at George Washington University in Washington, DC. In the summer of 1995, he was a visiting researcher at the Computer Science Department of the University of California at Los Angeles (UCLA). His research interests are in the fields of performance evaluation of communication networks, high-speed and all-optical networks, discrete event simulation, and queuing theory. His email address is: neri@polito.it