A FAIRNESS ENFORCEMENT PROTOCOL FOR INTERCONNECTED WDM RINGS

A. Bianco, J.M. Finochietto, F. Neri and V. Sarra
Politecnico di Torino
corso Duca degli Abruzzi 24, 10129 Torino, Italy
{bianco, finochietto, neri}@polito.it

G. Galante
Istituto Superiore Mario Boella
Via Pier Carlo Boggio 61 - 10138 Torino, Italy
galante@ismb.it

Abstract The paper focuses on the Metro (metropolitan area) network defined in the IST European project DAVID. The DAVID Metro network is based on sets of optical rings interconnected by passive memoryless devices called Hubs. Ring access is regulated by a dynamic time-division multiple-access scheme allocating fixed size slots in sets of wavelengths that provide multi-channel pipes among ring pairs. Resource sharing among nodes is granted by a centralized scheduling algorithm running at the Hub and by an empty-slot media access control protocol implemented at the nodes. Node position along the ring has a significant impact on access opportunities; thus fairness issues become fundamental in this architecture. This paper proposes a new fairness enforcement algorithm for interconnected multi-channel rings and evaluates its performance by simulation.

Keywords: WDM, optical ring network, interconnected rings, fairness enforcement

1. Introduction

The DAVID (Data And Voice Integration over Dense-wavelength division multiplexing) project is part of the Information Society Technology (IST) Program sponsored by the European Union. Its aim is the design of an optical packet-switched network for the transport of IP traffic over metropolitan, national and international distances. The DAVID network [1] has a two-level hierarchical architecture where several optical-ring metropolitan area networks (Metros) are interconnected by a Wide Area Network (WAN) backbone. As shown in Figure 1, the backbone network consists of optical packet routers in-

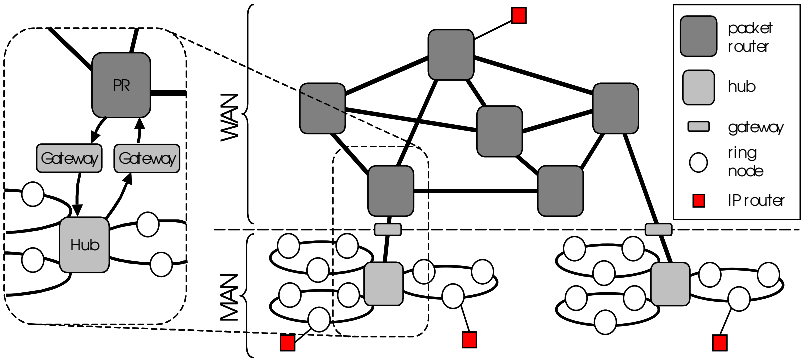

Figure 1. General overview of the DAVID network.

terconnected by a mesh network, while each Metro network comprises one or more rings interconnected through a bufferless Hub. Each ring collects traffic from several nodes; each Hub interconnects a number of rings, and is connected to an optical packet router in the WAN through a gateway. Access points to the network are provided both in DAVID Metro and backbone, and the traffic is collected by IP routers and switches connected to local area networks.

In packet-switched networks buffering inside routers is needed to solve contentions arising among packets arriving in a given node and headed to the same output port. In the DAVID WAN, optical packet routers provide buffering in the optical domain by means of fiber delay lines. No packet buffering in the optical domain is instead performed for packets flowing among ring nodes in the Metro network. Indeed, packets are buffered in ring nodes in the electrical domain, and are sent on the Metro network only when there are enough free resources to travel from source to destination through the Hub without being stored at any intermediate node. Thus, buffers are pushed towards the edge of the Metro network.

We focus on the Metro network in this paper. The Metro network exploits synchronous operation on both data and control plane: data are fixed-size packets, because this greatly simplifies high-speed operations in both the optical and the electronic domains. Obviously, variable-size IP-like packets must be fragmented by the sender before being transmitted on the fiber, and reassembled at the receiver before being delivered to the intended recipient. Bandwidth partitioning in the Metro is obtained by a superposition of a Wavelength Division Multiplexing (WDM) and a Time Division Multiplexing (TDM) access protocol: each fiber carries several data wavelength channels (in the prototype running at 10 Gbit/s), shared in the time domain to provide connectivity among rings. Resource sharing among nodes is granted by a centralized scheduling algorithm running at the Hub and by an empty-slot MAC (Media Access Control) protocol implemented at the nodes. The scheduling algorithm aim is to provide bandwidth pipes to ring-to-ring pairs, whereas the MAC protocol regulates resource sharing among nodes belonging to the same ring. In [2] and [3] scheduling algorithms and MAC protocols for transporting on the DAVID Metro best-effort and guaranteed traffic were presented and analyzed.

Due to the empty-slot protocol and to the Metro architecture, nodes experience different access opportunities when trying to transmit data. Nodes in upstream position, i.e., closer to the Hub, have better access chances than downstream nodes. This is a well-known problem in ring architectures: recently, the IEEE 802.17 Resilient Packet Ring Working Group is discussing fairness issues on this optical Metro access technology. MetaRing [4] was proposed to solve fairness problems in electronic rings; later, MetaRing extensions were also studied in a single-ring multi-channel environment; see [5] for a detailed discussion on several possible extensions of the MetaRing protocol to rings using WDM technology. However, the DAVID Metro architecture introduces some new constraints with respect to single ring multi-channel architectures. Thus, when simply using previous proposals, significant throughput limitations can be observed. This paper describes new fairness issues arising in DAVID and proposes extensions to MetaRing to ensure fairness among nodes and overcome throughput limitations arising in the DAVID Metro.

The remainder of the paper is organized as follows. In Section 1.2 we describe in more detail the Metro operations. In Section 1.3 we first define the scheduling and access schemes, and later we propose a new fairness enforcement scheme. In Section 1.4 we present selected simulation results to assess the performance of the proposed scheme. We conclude the paper in Section 1.5, where we give future research directions.

2. Metro network architecture

In the Metro demonstrator, each fiber conveys $W = 32$ wavelengths at 10 Gbit/s called *data channels,* plus wavelengths at 2.5 Gbit/s for signaling purposes dubbed *signaling* or *control channels*. The separation induced by such an out-of-band signaling scheme allows to keep bulk data in the optical domain as long as possible and to process them electronically only when they are received. The price that has to be paid is that the control channel must be terminated, electronically processed and retransmitted at each node, and this justifies its lower bit rate.

The Metro consists of several uni-directional optical rings interconnected by a bufferless Hub. Each node in the Metro is connected to a *physical ring* formed by one or more fibers running in parallel; yet, in general, a node can only access a subset of contiguous wavelengths, called *waveband*, selected among those circulating on a fiber. This allows to partition a physical ring into several *logical rings* containing only some of the nodes connected to the physical ring. Each logical ring comprises a control channel serving its data channels, and we assume for simplicity that nodes connected to different logical rings cannot forward packets directly among them; therefore, packets are routed among different logical rings at the Hub. Note that, in the remainder of the paper, we use the term ring to identify a logical ring for simplicity; any reference to physical rings will be explicit.

The number of data channels in a ring is typically limited to 4 or 8, because nodes are equipped with integrated optical modules with limited tunability comprising one tunable *data transmitter* and one tunable *data receiver*. In addition, to simplify the access to the control channel, each node needs a dedicated *signaling transceiver* which is always tuned on the control channel independently of the data transceivers behavior.

The DAVID metro operates in a synchronous, time-slotted fashion. Time slots, whose typical duration is 500 ns, are the minimum granularity in resource allocation, and typically rule the dynamics of network control, in the sense that both node access decisions and Hub switching happen between slot boundaries. Time slots are aligned on all wavelengths of the same ring, so that a *multi-slot* (a slot on each wavelength available in the waveband) is available to each node in each time slot.

The time dimension is used to set up at the Hub ring-to-ring bandwidth pipes among different rings using a TDM access scheme. In DAVID demonstrators, the Hub is implemented with a broadcast-and-select structure [6] using Semiconductor Optical Amplifiers (SOAs) as a switching technology and operates on fixed-size slots.

Since the Hub is a bufferless all-optical TDM switch, switching is performed with a waveband granularity, implementing *permutations* between rings: no two input (or output) rings can be connected to the same output (or input) ring in a given slot. Thus, all packets conveyed by a multi-slot traveling on an input ring are transparently transferred to the same destination ring as depicted in Figure 2. This operation assumes that all wavebands comprise the same number of wavelengths, and we make this assumption in the sequel of the paper for simplicity.

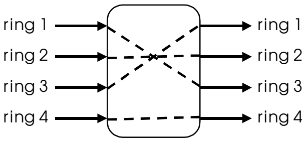

Figure 2. Hub performing a switching permutation between input and output rings.

Several alternatives were proposed in the DAVID project for the Metro node architecture. We consider here the scheme described in [7], in which nodes are attached to data channels on the rings with simple passive couplers. This makes the node much cheaper and permits to cascade a larger number of nodes (easily more than 16 nodes) on a ring. A single-channel receiver and a single-channel transmitter can be tuned to any wavelength in the waveband, in order to receive and transmit at most one packet in a given multi-slot.

Since packets cannot be erased from the channel, two separate wavebands are used: one for transmission and one for reception, requiring to double the number of wavelengths on the rings (without however requiring any increase in the amount of information that is switched at the Hub and at nodes). As a consequence, switching at the Hub must shift received information from the transmission waveband of the input ring to the reception waveband of the reception ring. This implies that no space reuse is available on the rings, which, for all practical effects, behave as a folded bus.

Packets traveling on the data channels are delayed with a fiber delay line so as to allow the signaling receiver to terminate the control channel and extract state information for the current multi-slot, to give the node CPU enough time to schedule packet transmissions in the current slot, and to regenerate the updated information in the signaling channel on the outgoing fiber.

3. Resource allocation

In this section, we describe the resource allocation problem in the Metro, by defining first the scheduling algorithm and later the access protocol. Finally, fairness issues are described and a suitable algorithm to ensure fairness is presented.

3.1 Hub Scheduling

A scheduling algorithm drives the Hub through a sequence of input/output configurations. To make the scheduling algorithm simpler, we assume, without loss of generality, that the propagation delay along all the physical (and therefore logical) rings is equivalent to the same integer number of slots. This constraint can be removed either making all the rings the same length by adding fiber delay lines, or by accounting in the scheduling for the propagation delay on each ring.

As shown in Figure 3, a sequence of F slots is organized into a fixed-length frame, where F is chosen according to the desired ring-to-ring connection granularity. The scheduling is driven by a traffic matrix, and can be computed at the Hub by using standard techniques based on iterated applications of modifications of the maximum size/maximum weight matching algorithms (see for instance, [8–9]). For simplicity, in this paper we assume that

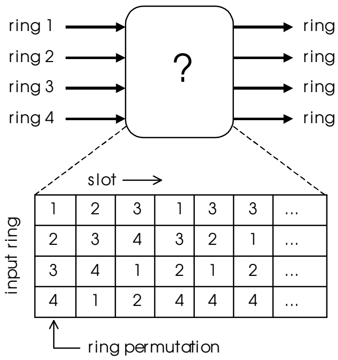

1	2	3	1	3	3	...
2	3	4	3	2	1	...
3	4	1	2	1	2	...
4	1	2	4	4	4	...

Figure 3. Hub scheduling.

- the traffic is described through a ring-to-ring traffic matrix,
- the ring-to-ring traffic matrix is known in advance; if this is not the case, it can be estimated with a measurement algorithm, as discussed in [2].

Note that with these assumptions the scheduling algorithm scales well with the network size, since its complexity depends on the number of rings (and not of nodes) in the Metro.

The slot allocation is distributed to nodes on the control channels as follows: the first few bits transmitted on each control channel at the beginning of a new slot denote the destination ring to which the data in the corresponding multi-slot will be forwarded upon reaching the Hub, as shown in Figure 4.

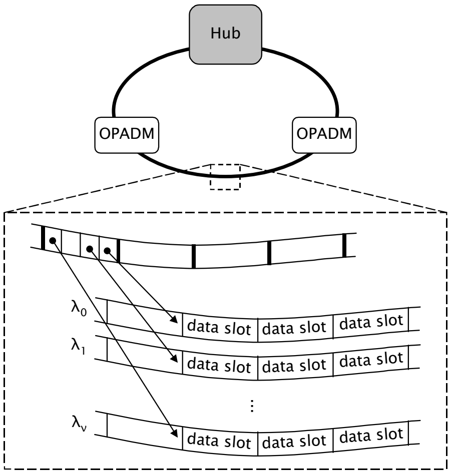

Figure 4. Control channel and data channels.

3.2 MAC protocol

Since at any given time each node can receive data on at most one wavelength in the same waveband, source nodes must make sure that there is no time overlapping among packets sent to the same destination node, by refraining from transmitting whenever this may happen.

This constraint is enforced locally at each node by a media access protocol that works as follows. A packet headed to node d can be sent on the data channels in the current multi-slot only when both the following conditions are met:

- it is possible to find an empty slot on the data channels,
- no other packet in the multi-slot is addressed to node d.

It is easy to observe that the Metro behaves as a distributed input-queued packet switch where buffers are located at node data transmitters; therefore it suffers from the well-known head-of-line (HOL) blocking phenomenon [10], which, however, can be easily overcome by providing virtual output queues (VOQ) in the electronic interfaces driving the data transmitters. Note that the HOL blocking can be completely removed from the Metro only when adopting a per-destination-node VOQ policy; yet, according to the scheduling algorithm which operates with a ring-to-ring granularity, we preferred to opt for a per-destination-ring VOQ policy to improve network scalability. Simulation results show that the performance degradation for using such non-optimal queuing architecture is minimal.

3.3 Fairness issues

The DAVID Metro network exhibits fairness problems induced by the ring topology, because upstream nodes have better access chances than downstream nodes. Credit-based schemes, such as MetaRing [4] proposed by Ofek for single ring electronic networks can enforce throughput fairness.

In MetaRing a control signal or message, called SAT, is circulated from node to node along the ring. A node forwarding the SAT is granted a transmission quota Q: the node can transmit up to Q packets before the next SAT reception. When a node receives the SAT, it immediately forwards the SAT to the next node on the ring if it is *satisfied* (hence the name SAT), i.e., if

- no packets are waiting for transmission on the ring, or
- Q packets were transmitted since the previous SAT reception.

If the node is not satisfied, the SAT is kept at the node until one of the two conditions above are met. Thus, SATs are delayed by nodes suffering throughput limitations, and SAT rotation times increase with the network load.

To be able to provide the full bandwidth to a single node, the quota Q must be at least equal to the number of data slots contained in the ring, i.e., proportional to the ring latency (propagation delay) measured in slot times. In overload, each node sends exactly Q packets per SAT rotation time.

In the case of the DAVID MAN, several rings exist, and multi-wavelength bandwidth pipes going through the Hub connect pairs of rings. Therefore, as the most straightforward extension to MetaRing, named SSRP-MetaRing (Single-Sat per Ring Pair MetaRing) in the paper, we may use a SAT for each ring pair (upstream ring, downstream ring): see [5] for more details on MetaRing extensions to multichannel architectures.

Note that each SAT signal can be represented by a few bits on the control wavelength. Each node inspects the incoming control slot, and operates on SATs as described above for the single ring case. Transmission opportunities are regulated by a MetaRing quota Q that may be different for each queue (recall we have a queue for each destination ring, i.e., each queue is controlled by a SAT). However, as previously stated, the quota Q must be greater than or equal to the ring latency; thus, since here we assume that all ring latencies are equal, we use the same value of the quota Q for all queues.

In the WDM single-ring architecture studied in [11], the rotational symmetry of the topology is such that each node is in a better position for the access to some channels, and in a worse position for the access to other channels. In the DAVID MAN, instead, due to the presence of the Hub, nodes closely following the Hub are in a better position in the access to all channels. This is even more true when transmission and reception resources are kept separate, since the first nodes observe all empty slots.

Albeit MetaRing provides throughput fairness in a single WDM ring, when using the SSRP-MetaRing extension previously described in the DAVID Metro, the total network throughput is lower than when fairness is not enforced. This is due to the fact that each node is equipped with only one data transceiver, and is therefore not able to fill completely by itself a ring-to-ring bandwidth pipe which contains $W > 1$ wavelength channels. Under heavy offered load we can assume that the node queues are always almost full. In such a situation, the first nodes on a ring have better access chances to transmission resources, so they exhaust their quota and idle much earlier than the last 2–3 nodes at the end of the same ring, which instead grab and retain the SAT for most of its rotation time. As a consequence, during each SAT rotation, there are time intervals when only a few (typically less than W) nodes at the end of the ring are unsatisfied (i.e., have quota greater than 0), and are therefore enabled to transmit. Whenever the number of unsatisfied nodes drops below W, the network utilization becomes less than 1, since each unsatisfied node can fill at most one wavelength in a multi-slot.

When the last node on the ring finally releases the SAT, the SAT goes quickly through the first nodes on the ring renewing their quota and re-enabling transmissions so that, by the time it reaches the last node (i.e., after one ring propagation time), it has still almost all if its original quota, and grabs the SAT, restarting the cycle.

This problem can be overcome by using W SATs (one for each wavelength) for each source/destination ring pair, and keeping as before one quota counter for each destination ring at each node. We name this extension SSW-MetaRing (Single-SAT per Wavelength MetaRing). The protocol is then modified introducing the additional constraint that no node can grab more than one SAT for the same source/destination ring pair, to avoid to behave as in the SSRP-MetaRing version. When a node is holding a SAT for a given source/destination/wavelength tuple and receives a SAT for the same source/ destination pair on a different wavelength, it adds Q to the quota counter for the corresponding ring and immediately releases the incoming SAT.

For example, in a network with $W = 4$ wavelengths per ring and 2 rings, each ring must convey $4 \times 2 = 8$ SATs controlling node transmissions on each wavelength, for each source/destination ring pair. Each node is then equipped with 2 per-destination-ring VOQs, each controlled by a different quota counter.

The SAT replication on each source/destination ring pair makes sure that less nodes refrain from sending packets on a ring as a consequence of quota exhaustion because it provides more opportunities to re-enable transmissions by restoring the nodes' quota counters to their original value upon SAT release. Therefore, it is more difficult that nodes on a ring reach a starvation situation, where most of them have traffic to send but are blocked by the SAT mechanism, thus guaranteeing higher throughput.

However, even the SSW-MetaRing does not work well as we would expect. The problem is in keeping a single quota per destination ring (i.e., a single quota for all wavelengths). Upstream nodes start transmitting as soon as they receive any SAT that control the given destination ring, regardless of the wavelength, since the quota is renewed, thus throttling downstream nodes that struggle transmitting since they can hold a single SAT. Severe unfairness can occur as shown in Section 1.4. The solution to this problem is to use separate quotas, one for each wavelength; downstream nodes can thus efficiently throttle upstream nodes on a specific wavelength, thus increasing overall throughput while guaranteeing fairness. This last version is called SSWQW-MetaRing (Single-Sat per Wavelength with Quota per Wavelength).

4. Simulation study

4.1 Network setup

The simulated network consists of 4 rings each connected to 15 nodes; thus the Hub has 4 ports. Each ring conveys $8+1$ wavelengths: 4 for transmission, 4 for reception and 1 for signaling purposes. Each node is equipped with 1 tunable data transceiver for accessing the $4+4$ data channels and 1 fixed signaling transceiver for accessing the control channel, and each VOQ can store up to 100 000 packets. The TDM frame length is set to $F = 400$ and one time slot lasts $T = 500$ ns. Each ring contains 8192 slots corresponding to ≈ 160 km, and the distance between the nodes has been fixed to 512 slots (≈ 10 km).

4.2 Traffic scenarios

The traffic flowing between different Metro rings is described by a 4×4 *ring-to-ring* matrix $\mathbf{R}$, representing the probability that a packet generated by a node attached to ring i is sent to ring o. The normalized (i.e., ranging from 0 to 1) load offered among each ring pair is computed by multiplying matrix $\mathbf{R}$ by a constant Λ, representing the normalized network offered load. All source nodes generate the same amount of traffic, measured in terms of number of slots; packet inter arrival times are geometrically distributed. Packets addressed to ring o are uniformly distributed among all nodes connected to ring o.

We consider two traffic scenarios, named *uniform*, and *client-server*. For simplicity, we assume that the Hub scheduling is pre-computed and matched to the traffic pattern.

4.2.1 Uniform traffic.

$$\mathbf{R} = \begin{pmatrix} 0.25 & 0.25 & 0.25 & 0.25 \\ 0.25 & 0.25 & 0.25 & 0.25 \\ 0.25 & 0.25 & 0.25 & 0.25 \\ 0.25 & 0.25 & 0.25 & 0.25 \end{pmatrix}$$

4.2.2 Client-server traffic. This traffic matrix has been designed to emulate the traffic unbalance that is observed today on the Internet, where client nodes typically originate a small amount of traffic directed towards server nodes, which answer with a larger amount of data. This causes the load to be distributed asymmetrically among different nodes and rings.

Three different node types are defined in our model: *server* node, *medium-size* node and *small-size* node. Each node type is characterized by a transmission rate and by an array specifying how the traffic generated is partitioned among different node typologies. A medium-size node generates 10% of the traffic produced by a server node whereas one small-size node outputs 5% of the traffic generated by a server node. Both small-size and middle-size nodes direct 17% of their traffic towards small-size nodes, 33% towards middle-size nodes and the remaining 50% towards server nodes. Finally, server nodes send 53% of their traffic towards small-size nodes, 43% towards middle-size nodes and 4% towards server nodes. Rings differ for the number of the nodes of a given size they connect as shown in Table 1.

Table 1. Node assignment on the different Metro rings

Ring	*Server nodes*	*Medium-size nodes*	*Small-size nodes*
1	5	5	5
2	1	4	10
3	1	10	4
4	3	6	6

The resulting ring-to-ring traffic matrix **R** is as follows:

$$\mathbf{R} = \begin{pmatrix} 0.226071 & 0.271729 & 0.254892 & 0.247308 \\ 0.091678 & 0.074734 & 0.079137 & 0.084886 \\ 0.110047 & 0.086239 & 0.087108 & 0.099215 \\ 0.166084 & 0.176687 & 0.169862 & 0.165628 \end{pmatrix} \quad \begin{matrix} 1.000000 \\ 0.330435 \\ 0.382609 \\ 0.678261 \end{matrix}$$

$$\begin{matrix} 0.593880 & 0.964632 & 0.942145 & 0.948307 \end{matrix}$$

The numbers besides each row (column) represent the row (column) sums. The matrix is normalized in such a way that the maximum weight row (column) sums to 1. It is easy to see how ring 1, which contains 5 server nodes, generates the highest amount of traffic towards all other rings. Further, the amount of traffic received from rings 2, 3 and 4 is much higher than the traffic they generate to emulate a real-world client/server scenario.

4.3 Performance Results

Figure 5 shows, under uniform traffic, a comparison of the overall network throughput obtained by the different MetaRing versions and when no MetaRing scheme is enforced. The SSWQW-MetaRing allows a 5% increase in the overall network throughput to be obtained with respect to SSRP-MetaRing. Note that when using no fairness enforcement protocol and when using SSW-MetaRing the overall network throughput is maximized. However, this is obtained at the price of a significant unfairness in the throughput obtained by single nodes, as can be observed in the upper and lower left side of Figure 6 where the single node saturation throughput for ring 1 is plotted when the network offered load is

1.5. Note that a similar behavior is observed on any ring, due to the symmetric traffic considered. Thus, SSWQW-MetaRing permits to jointly achieve high overall throughput and good throughput fairness.

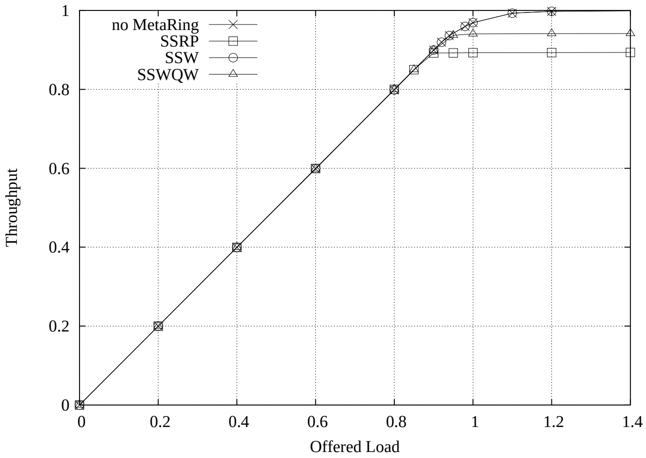

Figure 5. Performance comparison of fairness enforcement schemes under uniform traffic.

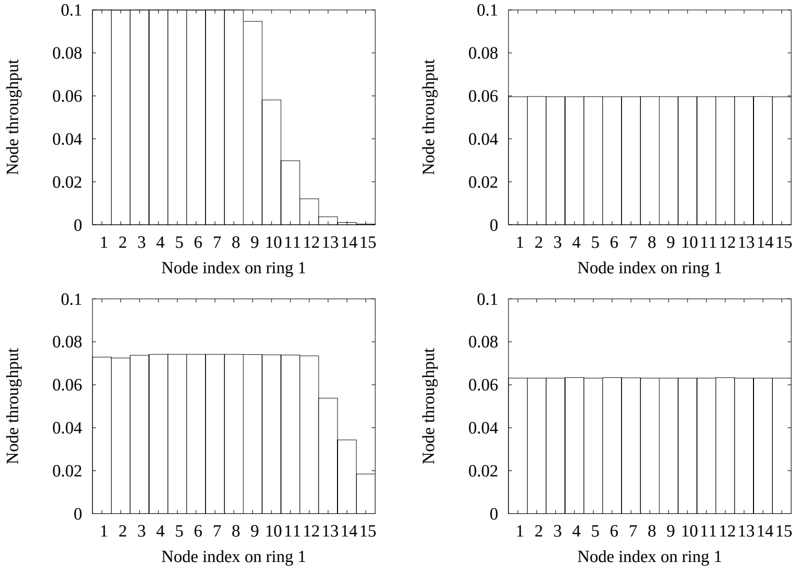

Figure 6. Saturation throughput obtained by single nodes on ring 1 when using no fairness enforcement algorithm (upper left), SSRP-MetaRing (upper right), SSW-MetaRing (lower left), and SSWQW-MetaRing (lower right) are used.

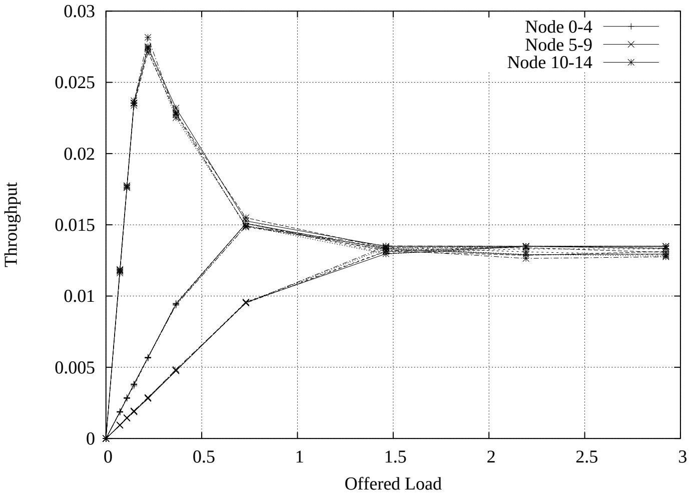

Figure 7. Comparison of the throughput obtained by different ring nodes and overall throughput for the traffic relation from ring 1 to ring 1 under client-server traffic.

Figure 7, instead, shows a comparison among the throughput obtained by different ring nodes and the overall throughput for the traffic relation from ring 1 to ring 1 under client-server traffic when running SSWQW-MetaRing. It is easy to see that the fairness enforcement algorithm works pretty well also in this scenario, guaranteeing that all the nodes belonging to the same class experience the same throughput. Moreover, a desirable max-min fairness enforcement is observed. Indeed, for an offered load around 0.2, when the resources available for the traffic relation from ring 1 to ring 1 become scarce, SSWQW-MetaRing, first subtracts bandwidth from server nodes (occupying positions 1–5); then, as the offered load becomes greater than 0.7, also middle-size nodes (labeled as nodes 6–10) start decreasing their throughput until all nodes converge to a fair bandwidth share of 0.013, according to a max-min fairness criteria, since, when increasing the load, all nodes try to overload the channel.

5. Conclusions

We have highlighted the difficulties of obtaining good throughput fairness properties in the DAVID Metro network. Several non trivial extensions to the MetaRing protocol, conceived for single channel electronic rings, are described and discussed. Their performance is assessed by simulation. SSWQW-MetaRing shows very good fairness properties and allows a significant increase in overall network throughput to be obtained.

Acknowledgments

This work has been funded as part of the European Union IST DAVID project.

References

[1] C. Develder, M. Pickavet, N. Leligou, F. Callegati, F. Neri, “The European IST Project DAVID: a Viable Approach Towards Optical Packet Switching,” to appear in *Journal on Selected Areas in Communications*

[2] A. Bianco, G. Galante, E. Leonardi, and F. Neri, “Measurement Based Resource Allocation for Interconnected WDM Rings,” *Photonic Network Communications*, vol. 5, no. 1, pp. 5–22, Jan. 2003.

[3] A. Bianco, J. Finochietto, E. Leonardi, P. Mitton, F. Neri, L. Quarello, “Multiclass Resource Allocation in Interconnected WDM Rings,” *7th Working Conference on Optical Network Design and Modeling (ONDM 2003)*, Feb. 2003, Budapest, Hungary.

[4] I. Cidon, Y. Ofek, “MetaRing — a Full-Duplex Ring with Fairness and Spatial Reuse,” *IEEE Transactions on Communications*, Vol. 41, No. 1, Jan. 1993, pp. 110–120.

[5] M. Ajmone Marsan, A. Bianco, E. Leonardi, F. Neri, S. Toniolo, “MetaRing Fairness Control Schemes in All-Optical WDM Rings,” IEEE INFOCOM’97, Kobe, Japan, April 1997.

[6] D. Chiaroni, et al., “First demonstration of an asynchronous optical packet switching matrix prototype for Multi-Terabit-Class Routers/Switches,” *post-deadline paper at ECOC 2001*, Vol. 6, PD.A.1.8, pp. 60-61, 2001.

[7] A. Stavdas, “Architectures, Technology and Strategies for a Gracefully Evolving Optical Packet Switching Networks,” *Optical Networks Magazine*, Vol. 4, No. 3, May/June 2003.

[8] B. Hajek, T. Weller, “Scheduling Non-Uniform Traffic in a Packet-Switching System with Small Propagation Delay,” *IEEE Transactions on Networking*, Vol. 5, No. 6, Dec. 1997, pp. 813–823.

[9] C.S. Chang, W.J. Chen, H.Y. Huang, “Birkhoff-von Neumann Input Buffered Crossbar Switches,” *IEEE Conference on Computer Communications (INFOCOM 2000)*, Tel Aviv, Israel, Mar. 2000, pp. 1614–1623.

[10] M. Karol, M. Hluchyj, S. Morgan, “Input Versus Output Queuing on a Space Division Switch,” *IEEE Transactions on Communications*, Vol. 35, No. 12, Dec. 1987, pp. 1347–1356.

[11] M. Ajmone Marsan, A. Bianco, E. Leonardi, A. Morabito, F. Neri, “All-Optical WDM Multi-Rings with Differentiated QoS,” *IEEE Communications Magazine*, Vol. 37, No. 2, Feb. 1999, pp. 58–66.