Throughput-Delay Trade-Offs in Slotted WDM Ring Networks

Thomas Bonald, Raluca-Maria Indre, Sara Oueslati, Chloé Rolland

► To cite this version:

Thomas Bonald, Raluca-Maria Indre, Sara Oueslati, Chloé Rolland. Throughput-Delay Trade-Offs in Slotted WDM Ring Networks. BROADNETS, 2010, Athènes, Greece. 10.1007/978-3-642-30376-0_21 . hal-01244071

HAL Id: hal-01244071

<https://hal.inria.fr/hal-01244071>

Submitted on 15 Dec 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Throughput-Delay Trade-Offs in Slotted WDM Ring Networks

T. Bonald¹, R.-M. Indre², S. Oueslati², and C. Rolland²

¹ Telecom ParisTech, France

² Orange Labs, France

Abstract. We analyse the throughput-delay trade-offs that arise in an optical burst-switched slotted WDM ring, where each node can transmit and receive on a subset of the available wavelengths. Specifically, we compare SWING, an access control scheme that combines opportunistic transmission and dynamic reservations, with a purely opportunistic access scheme. By means of analysis, we highlight the shortcomings of the opportunistic scheme in terms of load balancing and fairness. We then evaluate the performance of both schemes by simulation under several traffic scenarios and show that SWING yields a good throughput-delay trade-off.

Key words: Optical Burst Switching, WDM ring, MAC protocol, dynamic reservation, throughput, delay

1 Introduction

Increasing internet traffic volumes and the growing ubiquity of broadband access solutions (FTTx) challenge the simple use of over-provisioning in future metropolitan area networks (MAN). Most of today's MANs employ circuit-switching at the wavelength level [14] and have limited flexibility in handling traffic variability. An alternate switching paradigm, Optical Burst Switching (OBS), allows statistical multiplexing of data at the burst level, thus improving bandwidth utilization and network scalability.

This paper addresses the issue of resource access control in a burst-switched WDM metro ring. In such a network, a simple manner of providing resource access consists in allowing nodes to seize any unoccupied slot in order to send a waiting optical burst. We refer to this scheme as *purely opportunistic access*. Due to its simplicity, the opportunistic scheme suffers from poor load balancing and lack of fairness which leads to potentially high queuing delays. In the following, we compare the performance of the opportunistic scheme to that of SWING (Simple Wdm rING) [6], a distributed reservation protocol that aims at seizing all available transmission opportunities, while ensuring fair resource allocation.

The topic of resource access control in slotted WDM rings has drawn considerable research attention in recent years. Most of the existing proposals aim at providing a certain degree of fairness among the network nodes. Some of them are based on the idea that each station is allowed to emit up to a certain quota

of packets in a dynamically or statically determined cycle [3, 4, 9, 10]. The main drawback of these schemes is that the quota has to be large enough to ensure an efficient utilization of the ring, which implies a large interval between the transmission periods for each station and leads to potentially unacceptable packet delays. Other proposals rely on a connection-oriented reservation strategy [1, 2], where the ring is divided in multiple reservation frames in which consecutive slots can be used to transmit long packets without segmentation. Finally, a solution based on virtual circuit allocation has been investigated in [7], where a centralized reservation scheme is deployed.

In this paper, we analyze SWING, a novel MAC protocol that combines opportunistic access and dynamic slot reservation. The underlying reservation mechanism is original in many ways. Firstly, no reservation quota is imposed, which makes it more flexible. Secondly, the reservations are not connection-oriented, allowing a station to use a reserved slot for any destination. Lastly, it allows spatial reuse of reserved slots, thus enhancing bandwidth utilization.

This work builds on a previous paper [6], where SWING is presented in detail and its stability condition and throughput efficiency evaluated under the restrictive assumption of a single wavelength. The present study extends [6] by considering the impact of tunable transmitters and fixed receivers on network stability and performance. It turns out that the property of throughput optimality established for the opportunistic transmission scheme in [6, 12] is lost in the more realistic context of multiple wavelength channels. A thorough performance evaluation is carried out by means of analysis and simulations, the latter being realized in a multiservice context, with two service classes. The results highlight the interesting throughput-delay trade-off achieved by SWING, compared to the opportunistic scheme.

The paper is organized as follows. Section 2 presents the network architecture and the access protocols. The traffic model and performance metrics are described in Section 3. Section 4 presents the throughput analysis of the opportunistic scheme, focusing on load balancing and fairness issues. Section 5 is devoted to the simulation results, that show the throughput-delay trade-offs achieved by the opportunistic and dynamic reservation schemes. Section 6 concludes the paper.

2 System

This section presents the network architecture, the scheduling mechanism as well as the considered access protocols.

2.1 A slotted WDM ring

The network is a slotted WDM ring, which actually consists of two counter rotating fiber rings, one of which is used for backup in case of failure. We consider here the operation of a single ring. The ring consists of $N + 1$ nodes linked with

optical fibers, as illustrated by Figure 1 for $N = 6$. A particular station, say station 0, corresponds to the *hub* that connects the ring to the Internet. The other N nodes are *access stations* connecting end-users to the ring. Each node is an optical packet add/drop multiplexer (OPADM) capable of transmitting and receiving data on W wavelengths. Control information (occupancy of each wavelength, destination of each packet, reservation state when applicable) is sent via control packets on an out-of-band channel. One control packet is associated to every time slot across all wavelengths.

Each node i is equipped with t_i tunable transmitters, $t_i \in \{1, \dots, W\}$, and with a set of fixed receivers, $r_i \subset \{1, \dots, W\}$; it can thus simultaneously transmit data on any t_i distinct wavelengths and receive data on a fixed subset r_i of the wavelengths. Throughout the paper, we assume that in each time slot, the hub can transmit and receive on all wavelengths simultaneously i.e., $t_0 = W$, $r_0 = \{1, \dots, W\}$ while all other nodes can transmit and receive on a single wavelength (i.e., $t_i = |r_i| = 1$ for all $i \neq 0$). The received wavelengths are allocated in a cyclic manner from node 1 to node N , as shown in Figure 1. In the numerical examples, we consider either $W = 2$ or $W = 4$, which is compliant with the current economical and technological constraints [17].

Time is slotted so that there are S slots circulating on the ring at the speed of light, visiting nodes in a cyclic manner in the order $1, \dots, N$. Slots have a fixed duration, $10 \mu\text{s}$, which is taken as the time unit throughout the paper. Each slot on each wavelength channel may carry at most one optical burst. The size of an optical burst is thus given by the slot duration times the transmission speed on each wavelength, that is 12.5 KB for a transmission speed of 10 Gbit/s. It is equivalent to 8 Ethernet Maximum Transmission Units (MTU).

When a slot arrives at node i , all packets contained in this slot on the subset of wavelengths r_i and destined to node i are received and extracted from the corresponding wavelengths. In addition, node i can transmit up to t_i packets in this slot, depending on the queued packets and the occupancy state of the wavelengths, as described below.

2.2 Scheduling

Each node i maintains a queue of optical bursts per destination. At the arrival of each slot, if several wavelengths are unoccupied, the next wavelength to be used for transmission is selected using the reverse round robin policy, where wavelengths are scanned cyclically in the order $W, \dots, 1$. Among the destinations that can receive on the selected wavelength, the next queue to be served is also selected using reverse round robin scheduling (this reverse order tends to favor longer paths, see [6] for details). This process stops once t_i queues have been served or no other queue can be served, and resumes at the arrival of the next slot.

Fig. 1. Ring of $N = 6$ nodes with $W = 4$ wavelengths and fixed receivers $r_0 = \{1, 2, 3, 4\}$, $r_1 = \{1\}$, $r_2 = \{2\}$, $r_3 = \{3\}$, $r_4 = \{4\}$, $r_5 = \{1\}$, $r_6 = \{2\}$.

2.3 Access protocols

In the described network architecture, all nodes implement two MAC sub-layers: adaptation and transport. The adaptation sublayer is in charge of assembling variable-size data packets into fixed-size optical bursts. This is done by means of a timer-based burst-assembly mechanism proposed in [5].

The transport sub-layer is responsible for providing resource access control for the network nodes. In the following, we consider two different schemes at the transport sub-layer: a purely opportunistic scheme and SWING, a dynamic reservation scheme that allows on-demand slot reservation. In the opportunistic scheme, nodes seize every opportunity to use an idle slot in order to send a waiting optical burst. In such a scheme, nodes compete for access to the ring without any control beyond the packet scheduling policy described above.

In SWING, each node reserves slots when necessary and preempts the reservations of other nodes in order to preserve fairness, as explained in the following. The packet scheduling policy is still the one described in §2.2. A wavelength k is considered free by node i in slot n if this slot carries no burst and is either:

- *unreserved* (opportunistic transmission);
- *reserved for node i* , which means the reservation made by node i in the previous cycle has not been preempted;
- *reserved for another node and spatially reusable*, i.e. node i can send an optical burst to node j on a reserved slot as long as node j is not downstream from the node that reserved the slot.

All reservations held by node i on slot n are then released.

The reservation scheme is fully distributed. Each node i maintains a counter $c_{ij}(k)$ of the number of reservations node j has made on each wavelength k over slots $1, \dots, n-1$, where n is the current slot. Note that these counters are local visions of the reservation status and, in general, are different from one station to another.

A maximum number t_i of new reservations are attempted by node i on slot n if the pending reservations are not sufficient to clear all queues of node i . To determine whether the number of pending reservations of node i , namely $\sum_{k=1}^W c_{ii}(k)$, is sufficient, the burst scheduling algorithm is *virtually* run on all

queues for each wavelength k , assuming that $c_{ii}(k)$ slots are available on wavelength k and no other slot is available. The algorithm stops either when there are no more pending optical bursts in the queues or when all reserved slots have been used. Reservations are attempted on wavelengths r_j associated with queues j that still have waiting bursts when the algorithm stops. Node i requests the reservations on unreserved wavelengths and preempts the reservations made by other nodes, if allowed; the preemption of a reservation held by node j on wavelength k is allowed if and only if node i has fewer pending reservations, that is if $c_{ii}(k) < c_{ij}(k)$.

3 Model

This section presents the traffic characteristics and the performance metrics used throughout the paper.

3.1 Traffic characteristics

Packets are assumed to arrive according to a Poisson process of intensity λ . Although this assumption is known to be incorrect in local area networks [16, 19], it is reasonable in the metropolitan and wide area networks we are interested in [8]. Traffic is indeed shaped by the transmission speed of end systems and access queues; typical peak rates of end-to-end data flows do not exceed a few hundred Mbit/s while the transmission speed on each wavelength of the ring is equal to 10 Gbit/s. The multiplexing of a large number of such flows typically leads to Poisson packet arrivals.

Let p_{ij} be the associated traffic matrix, with $p_{ii} = 0$ for all i : an incoming packet has source i and destination j with probability p_{ij} . We denote by $\lambda_{ij} = \lambda p_{ij}$ the packet arrival rate on the source-destination pair (i, j) and by $\lambda_i = \sum_j \lambda_{ij}$ the total packet arrival rate at node i .

3.2 Throughput metric

A key criterion for MAC efficiency is its ability to fully use the transmission and reception capabilities of each node. Formally, this can be expressed as the following stability issue, which is reminiscent of those arising in the performance analysis of ad-hoc networks [13] and switching fabrics [15].

Nodes are indexed by integers modulo N . We refer to link l as that connecting node l to node $l+1$, for all $l = 1, \dots, N$. We denote by $[i, j]$ the set of nodes between node i and node j , that is the set $\{i, \dots, j\}$ if $i \leq j$ and the set $\{i, \dots, N\} \cup \{1, \dots, j\}$ otherwise. For the network to be stable, in the sense that all queues (i, j) remain finite, it is necessary that traffic does not exceed the transmission and reception capabilities of each node, namely:

$$\forall i, \quad \lambda_i < t_i \tag{1}$$

and

$$\forall l, \forall r \subset \{1, \dots, W\}, \sum_{i,j: l \in [i,j], r_j \subset r} \lambda_{ij} < |r|. \quad (2)$$

We say that the network is *throughput-optimal* if these conditions are also sufficient for stability.

In view of (1) and (2), it is natural to define the *load* as:

$$\rho = \max \left(\max_i \frac{\lambda_i}{t_i}, \max_{l, r \subset \{1, \dots, W\}} \frac{1}{|r|} \sum_{i,j: l \in [i,j], r_j \subset r} \lambda_{ij} \right). \quad (3)$$

The optimal stability condition can then be simply written as $\rho < 1$. In general, the actual stability condition is more restrictive, say $\rho < \theta$ with $\theta \leq 1$. We refer to θ as the throughput efficiency, which depends both on the considered MAC and on the traffic matrix. Throughput optimality is equivalent to the equality $\theta = 1$.

3.3 Delay metric

Data packet delay is another key performance metric. We estimate the probability that the queuing delay of any packet exceeds some threshold, taken equal to two cycle times ($2S$ slots). This delay is supposed to be sufficient for the dynamic reservation scheme to work out since one cycle is needed before a node can benefit from its reserved slot.

4 Throughput Analysis of the Opportunistic Scheme

To analyse the performance of the opportunistic scheme, we consider in this section a simple burst-level model where bursts (instead of packets) arrive according to Poisson processes.

4.1 Load balancing impairment

The opportunistic scheme was proven to be throughput-optimal for a single wavelength [12, 6]. It turns out that this property is lost for multiple wavelengths, as shown by the example of Figure 2 with $N = 2$ nodes, $W = 2$ wavelengths. Recall that the hub station can simultaneously receive and transmit on all wavelengths, i.e. $r_0 = \{1, 2\}$ and $t_0 = 2$ when $W = 2$. The access nodes are able to transmit on any wavelength but only on one wavelength at a time ($t_1 = t_2 = 1$). They can receive on a single predefined wavelength ($r_1 = \{1\}$ and $r_2 = \{2\}$). We further assume that there is traffic only on source-destination pairs $(1, 0)$ and $(2, 1)$. In view of (1)-(2), the optimal stability condition is given by:

$$\lambda_1 < 1 \quad \text{and} \quad \lambda_2 < 1.$$

It is achieved by forwarding all traffic of the source-destination pair $(1, 0)$ on wavelength 2, so that the traffic of the source-destination pair $(2, 1)$ can fully use wavelength 1. Unfortunately, node 1 is not aware of the traffic matrix and thus cannot make the optimal decision; in view of the scheduling policy described in §2.2, it chooses wavelength 1 or 2 alternately, resulting in the actual stability condition:

$$\lambda_1 < 1 \quad \text{and} \quad \frac{\lambda_1}{2} + \lambda_2 < 1.$$

In the homogeneous case, i.e. $\lambda_1 = \lambda_2 = \lambda/2$, this reduces to the inequality $\lambda < 4/3$ whereas the optimal stability condition is $\lambda < 2$. The throughput efficiency θ of the opportunistic scheme is thus equal to $2/3$ in this case.

Fig. 2. A simple scenario where opportunistic transmission is suboptimal.

We generalize this toy example to the case $N \geq W$. We consider that there is traffic only on source-destination pairs $(N, 1)$ and $(i, 0), i \in \{1, \dots, N-1\}$. We consider the homogeneous case in which $\lambda_i = \lambda/N$ for all $i \in \{1, \dots, N\}$.

Nodes $1, \dots, N-1$ alternately choose one of the W wavelengths, thus limiting the transmission opportunities of node N on wavelength 1. Stability is defined by the capacity of wavelength 1 on link N . From (2) and (3) we obtain the throughput efficiency of the opportunistic scheme:

$$\theta = \frac{N}{N + W - 1}.$$

Table 1 shows the corresponding throughput efficiencies for different values of W and N . Since the traffic on source-destination pair $(N, 1)$ decreases with the number of nodes N , the fraction of slots required by station N also decreases, leading to a better throughput efficiency. Similarly, efficiency decreases with W . Under SWING access control, the dynamic reservation scheme forces nodes $1, \dots, N-1$ to use wavelengths $2, \dots, W$, yielding throughput-optimality (i.e., $\theta = 1$).

4.2 Fairness impairment

We study a scenario in which there is traffic only from the access nodes to the hub, namely $\lambda_{i0} = \lambda/N$ for all $i \neq 0$. We refer to this scenario as *hub uplink*.

N	4	6	8	10
$W = 2$	0.80	0.85	0.88	0.90
$W = 4$	0.57	0.67	0.72	0.77

Table 1. Throughput efficiency θ of the opportunistic scheme in the studied scenario.

We consider there are W wavelengths in the network. When no access control is implemented, the amount of free slots available to each node depends on the node's position on the ring. Indeed, the stations preceding the hub are more likely to suffer from slot starvation, particularly station N . It is this station that will drive the stability of the network. In the following, we analyze the characteristics of the waiting queue associated to node N .

We model our ring network as a system of two waiting queues. The first queue is an aggregation of the waiting queues at stations $1, \dots, N-1$. The second queue simply corresponds to station N 's waiting queue. At each queue, clients arrive according to a Poisson process of rate $\lambda_a = \sum_{i=1}^{N-1} \lambda_{i0}$ and $\lambda_b = \lambda_{N0}$, respectively. For both queues, the service rate μ corresponds to the arrival rate of free time slots. Our network can thus be reduced to a single-server queuing system as shown in Figure 3. Under opportunistic access, the stations succeeding the hub will have privileged access to the free slots meaning that the first queue has priority over the second.

Fig. 3. Single-server priority queuing system.

Fig. 4. Markov chain for a single-server preemptive priority queuing system.

We consider the worst case scenario in which the first queue has *preemptive priority* over the second one (Figure 4). Using the generating function of the stationary distribution of the second queue, we express the tail asymptotics of the *number of clients* in the second queue (see [18] for more details):

$$P(Q_b = n) \sim \frac{1 - \rho}{\rho_b} (\rho^2 - \rho_a) \rho^{n-1}, \quad (4)$$

where $\rho_a = \lambda_a/\mu$, $\rho_b = \lambda_b/\mu$ and $\rho = \rho_a + \rho_b$.

If we limit the size of the second queue to 25 clients¹ and we consider the network to be stable as long as the packet loss rate is less than 2%, then (4) can be used to compute the maximum sustainable load. Table 2 gives the maximum sustainable load for different values of N and for $W = 4$. We note that the lack of fairness of the opportunistic scheme can lead to efficiency losses as high as 10%.

N	4	6	8	10
Analysis	1.00	0.88	0.90	0.92
Simulation	0.99	0.90	0.91	0.93

Table 2. Throughput efficiency θ of the opportunistic scheme in the hub uplink scenario.

5 Simulation Results

This section presents the results of our data packet-level simulations which were conducted on a multiservice network supporting two classes of service: real-time and best-effort.

5.1 Simulation parameters

The number of slots S is taken equal to 100, corresponding to a cycle time of 1 ms for a typical slot duration of 10 μ s. Each simulation run lasts 10^6 cycles. For throughput performance, we need to decide if the network is stable or unstable at the end of the simulation, for any given load ρ . To this end, we chose to limit each queue size to 100 MTU packets (each node has 2 queues, one for each class of service). We consider the network to be stable as long as the packet loss rate is less than 2% at each queue. We have verified that the results are not very sensitive to these parameters.

We assume real-time packets to have absolute priority over best-effort packets (see [5] for details). Real-time packets arrive according to a Poisson process, while

¹ This corresponds to a total queue size of 200 MTU packets (cf §2.1) which is also the value considered in the simulations of §5.

best-effort packets arrive according to a *batch* Poisson process, with a batch size uniformly distributed between 1 and 16 packets. We consider 10% of the traffic to be real-time traffic. The packet size distribution is given by 60% of 40B packets (minimum transmission unit) and 40% of 1500B packets (maximum transmission unit), for both real-time and best-effort packets.

We simulate the access protocols described in §2.3 where data packets are aggregated into optical bursts using a timer-based burst assembly mechanism.

5.2 Traffic scenarios

In order to evaluate throughput-delay performance in typical traffic conditions, we consider the following three traffic scenarios, depicted by Figure 5:

Local scenario: All traffic is local in the sense that there is no traffic originating from or going to the hub, yielding:

$$\forall i, j \neq 0, i \neq j, \quad \lambda_{ij} = \frac{\lambda}{N(N-1)}$$

Hub scenario: All traffic originating from node $i \neq 0$ is destined to the hub. Referring to “upstream” as the traffic destined to the hub and to “downstream” as the traffic originating from the hub, the traffic matrix is characterized by a single parameter α representing the fraction of upstream traffic:

$$\forall i \neq 0, \quad \lambda_{0i} = \frac{\lambda(1-\alpha)}{N}, \quad \lambda_{i0} = \frac{\lambda\alpha}{N}.$$

P2P scenario: There is some traffic between any pair of nodes (i, j) and traffic is symmetric in the sense that $\lambda_{ij} = \lambda_{ji}$ for all $i, j, i \neq j$, which is representative of peer-to-peer traffic. The traffic matrix is characterized by a single parameter β representing for any node $i \neq 0$ the fraction of local traffic, that is the fraction of traffic originating from all nodes $j \neq 0$:

$$\forall i \neq 0, \quad \lambda_{0i} = \lambda_{i0} = \frac{\lambda(1-\beta)}{2N},$$

$$\forall i, j \neq 0, i \neq j, \quad \lambda_{ij} = \frac{\lambda\beta}{2N(N-1)}.$$

Unless otherwise specified, we take $\alpha = \beta = 20\%$.

Fig. 5. Considered traffic scenarios (from left to right): local, hub and P2P (hub station represented in black)

5.3 Results

Throughput Figure 6 gives the throughput efficiency with respect to the number of nodes N (other than the hub) in the three traffic scenarios of §5.2, for $W = 4$ wavelengths, for both SWING and the opportunistic scheme. The throughput efficiencies of the two schemes are fairly similar and are higher than 0.75, regardless of the considered scenario. Note that efficiency is slightly higher for $N = 6$ for both schemes. In this case, wavelengths 1, 2 carry more traffic than wavelengths 3, 4 (see Figure 1), which reduces the number of limiting resources and increases the maximum throughput.

As shown in [6], in the case of a single wavelength, the opportunistic scheme is throughput optimal while SWING experiences efficiency losses due to the reservations which limit possibilities for opportunistic transmission. In the case of multiple wavelengths where each node can transmit and receive on a subset of the available wavelengths, the efficiency of the opportunistic scheme decreases because of its poor load balancing over the wavelengths and lack of fairness (cf. Section 4). Despite the efficiency loss due to reservations, SWING is able to achieve the same level of throughput efficiency as the opportunistic scheme, while ensuring fairness.

Delay Figure 7 shows the probability that the queuing delay of a data packet exceeds 2 ms, for both SWING and the opportunistic scheme for $N = 8$ nodes plus the hub. Delays of best effort traffic are low as long as the load is less than the maximum sustainable load, for both access protocols in all traffic scenarios.

We further notice that the opportunistic scheme experiences low delays for real-time traffic in both local and hub scenarios. However, this is not the case in the P2P scenario, where delays of real-time traffic are significant for $\rho > 0.8$. This is a direct consequence of the opportunistic scheme's absence of access control. As seen in §4.2, depending on its position on the ring, a node may be starved from free slots on a specific wavelength that serves a given destination. The node is then unable to satisfy the corresponding queue, which leads to high delays for both real-time and best effort traffic. It turns out that such situations occur more often in the P2P scenario due to the joint effect of high traffic coming from and going to the hub and local traffic exchanged among access nodes.

We notice that SWING is able to guarantee low delays for real-time traffic even at high network loads in any of the considered traffic scenarios. Indeed, the overall probability that real-time packets wait longer than 2 ms is close to 0. For instance, only 5 out of 10^6 real-time packets wait longer than 2 ms at load $\rho = 1$ in the P2P scenario. As for the opportunistic scheme, over 7% of real-time packets will wait longer than 2 ms under the same traffic conditions.

Fig. 6. Throughput efficiency of SWING (top) and opportunistic scheme (bottom) against the number of nodes N in the local, hub and P2P scenarios.

6 Conclusion

In this paper we have analysed the throughput-delay trade-offs that arise in slotted WDM rings for a purely opportunistic transmission scheme and for SWING, a dynamic slot reservation protocol. We have highlighted the throughput suboptimality of the purely opportunistic scheme as well as its fairness impairments in the context of multiple wavelength channels. The results of our simulations have shown that the dynamic reservation scheme of SWING ensures fairness while maintaining high throughput efficiency. Unlike the opportunistic scheme, SWING is able to guarantee low delays for high priority traffic under any of the considered traffic scenarios.

Future work is focused on analysing throughput-delay trade-offs arising in other network topologies. We also intend to investigate fairness in the context of WDM ring interconnection.

Fig. 7. Delay performance of SWING (top) and opportunistic scheme (bottom) in the local, hub and P2P scenarios in a ring of $N = 8$ nodes.

Acknowledgements

The authors would like to thank Alain Simonian for his contribution to the analytic part of this work, in particular on the considered priority queuing system.

References

1. M. Ajmone Marsan, A. Bianco, E. Leonardi, A. Morabito, F. Neri, *SR3: a Bandwidth-Reservation MAC Protocol for Multimedia Applications over All-Optical WDM Multi-Rings*, Proc. INFOCOM, 1997.
2. K. Bengi, H.R. van As, *QoS Support and Fairness Control in a Slotted Packet-Switched WDM Metro Ring Network*, Proc. GLOBECOM, 2001.
3. A. Bianco, D. Cuda, J. Finochietto, F. Neri, *Multi-MetaRing Protocol: Fairness in optical packet ring networks*, Proc. ICC, 2007.
4. A. Bianco, D. Cuda, J. M. Finochietto, F. Neri, C. Piglion, *Multi-Fasnet Protocol: Short-Term Fairness Control in WDM Slotted MANs*, Proc. GLOBECOM, 2006.
5. T. Bonald, R-M. Indre, S. Oueslati, C. Rolland, *On Virtual Optical Bursts for QoS Support in OBS Networks*, Proc. of ONDM, 2010.

6. T. Bonald, S. Oueslati, J. Roberts, C. Roger, *SWING: Traffic capacity of a simple WDM ring network*, Proc. of ITC 21, 2009.
7. C. Cadere, N. Izri, D. Barth, J. M. Fourneau, D. Marinca, S. Vial, *Virtual Circuit Allocation with QoS Guarantees in the ECOFRAME Optical Ring*, Proc. ONDM, 2010.
8. J. Cao, W.S. Cleveland, D. Lin, D.X. Sun, *Internet Traffic Tends Toward Poisson and Independent as the Load Increases*, in Nonlinear Estimation and Classification, Springer, New York, 2002.
9. A. Carena, V. De Feo, J. Finochietto, R. Gaudino, F. Neri, C. Piglion, P. Poggiolini, *RingO: An experimental WDM optical packet network for metro applications*, IEEE JSAC, Vol 22, No 8, Oct 2004.
10. I. Cidon, Y. Ofek, *Metaring: A full-Duplex Ring with Fairness and Spatial Reuse*, Proc. INFOCOM, 1990.
11. D. Chiaroni, *The French RNRT ECOFRAME project : packet technology perspectives in metro networks*, Invited presentation, APOC 2008.
12. J-F. Dantzer, V. Dumas, *Stability Analysis of the Cambridge Ring*, Queueing Systems: Theory and Applications, Vol 40, No 2, 2002.
13. A. Ephremides and L. Tassiulas, *Stability Properties of Constrained Queueing Systems and Scheduling Policies for Maximum Throughput in Multihop Radio Networks*, IEEE Trans. on Automatic Control, 1992.
14. M. Herzog, M. Maier, and M. Reisslein, *Metropolitan Area Packet-Switched WDM Networks: A Survey on Ring Systems*, IEEE Communications Surveys and Tutorials, Vol 6, No 2, pp. 2–20, 2004.
15. I. Keslassy and N. McKeown, *Analysis of Scheduling Algorithms That Provide 100% Throughput in Input-Queued Switches*, Proc. Allerton Conference on Communication, Control, and Computing, 2001.
16. W.E. Leland, M.S. Taqqu, W. Willinger, D.V. Wilson, *On the self-similar nature of Ethernet traffic*, IEEE/ACM Transactions on Networking, 1994.
17. Matisse Networks, *EtherBurst: Optical Packet Transport System*, White Paper, 2008.
18. R. G. Miller, *Priority Queues*, The Annals of Mathematical Statistics, Vol 31, No 1, 1960.
19. V. Paxson, S. Floyd, *Wide area traffic: the failure of Poisson modeling*, IEEE/ACM Transactions on Networking, 3-3, pp. 226–244, 1995.