Electronic architectures of optical slot switching nodes

Nihel Benzaoui, Yvan Pointurier, Jean-Christophe Antona

Electronic architectures of optical slot switching nodes

Nihel Benzaoui, Yvan Pointurier and Jean-Christophe Antona
Alcatel-Lucent Bell Labs, Nozay, France

Abstract— **An optical slot switching node network called POADM (packet optical add-drop multiplexers) has formerly been proposed as a flexible solution for metropolitan ring networks to carry data traffic with a sub-wavelength switching granularity and with a good energy efficiency, which is enabled by optical transparency.**

In this paper, we propose several architectures for the electronic side of optical slot switching nodes. Those architectures increase the flexibility with which a client, attached to a node, can access the transport medium, through the addition of electronic switches, working either at client frame granularity or at slot granularity (a slot encapsulates several client frames); such electronic switches can be located at either transmitter, receiver, or both sides of a node, thereby decreasing traffic latency, at the expense of increased node cost and/or energy consumption. This paper focuses on the latency aspect only. We assess and compare the latency of these node architectures with simulations; some results are also backed analytically. The utilization of an (electronic) slot switch enables load balancing across the transport channels, while the node architecture with an (electronic) client frame switch additionally permits flow aggregation, resulting in lower queuing delay.

The lowest queuing delay is achieved by the architecture embedding client frame switches at both transmitter (Tx) and receiver (Rx) sides, while the absence of electrical switches leads to the worst performance, but presents the advantage of lower cost and consumption with respect to the other architectures. The utilization of client frame switches at Tx or Rx presents intermediate performances.

Index Terms— **Optical packet switching, Network performance.**

I. Introduction

While traffic was mainly voice with little data, before the 2000's, making it predictable and amenable to circuit switching, trends are now reversed [1]. Most traffic is now data, leading to a modification of traffic characteristics to unpredictable and bursty, especially in the access and metro segments. In this context, an optical slot switching technology called "POADM" (packet optical add-drop multiplexers) [2] was proposed for metro ring networks, to provide sub-wavelength switching flexibility (down to the) in the optical domain.

A POADM ring is a Wavelength-Division Multiplexed (WDM) time-slotted ring. It is composed of several (e.g., 40) data channels, and of one additional "control channel" which carries all slot headers and is (electronically) processed at each node. Each node is connected to one or several clients (e.g., a Passive Optical Network (PON), an Ethernet switch, etc.). Data coming from/going to the clients is encapsulated/decapsulated within fixed-size packets called slots (of few μs) that can transit transparently, i.e. they are not converted to the electronic domain at the intermediate nodes [2]. Each node decides whether to drop or let transit each slot on each wavelength according to information carried by the control channel. This transparency helps to reduce energy consumption [3]. Each POADM node includes burst-mode fixed-wavelength receivers (Rx), making the nodes able to receive data on predefined wavelengths, and fast wavelength-tunable transmitters (Tx), which enable each node to communicate with every other node by tuning to the right wavelength on a per-slot basis [4].

The POADM technology has received much attention in the past: regarding network dimensioning [5][6], energy consumption [3], technological issues [4], and MAC, performance and QoS management [7] [8] [9] [10] [11] [12] [13] [14]. However, mapping issues between clients and nodes were never considered. In a plain, basic POADM node, clients are statically mapped to transponders (TRX; a transponder is a Tx and an Rx) or line cards. However, allowing a client to use any Tx or Rx at a given node provides more flexibility to send/receive data, and hence has an impact on network performance, measured for instance in terms of latency. The contribution of this paper is two-fold. First, we propose several optical slot switching node architectures enabling dynamic client/WDM TRX mapping, at several granularities (client frame or optical slot). Second, we assess the performance of each node architecture, essentially in terms of queuing delay. Some of the results are backed analytically. Note that this paper applies to any WDM slotted ring where nodes are equipped with tunable transmitters and fixed-wavelength receivers; POADM is one possible implementation of this concept.

This paper is organized as follows. In Section II, we present the basic and alternate architectures of our optical slot switching node. Latencies for some of those architectures are derived analytically in Section III and compared using simulations in Section IV. We give in Section V the most important conclusions of the performances studies.

II. Node Architectures

A POADM node is able to encapsulate and decapsulate client traffic into optical fixed-duration optical packets or "slots". In addition, a POADM node can add, drop, and let optical slots transit on an optical ring. Hence, a POADM node consists of: client cards, for client traffic handling (essentially, en/decapsulation and, optionally, QoS management); an optical packet blocking fabric with the relevant controller, and line cards directly connected to the metro ring to transmit and receive optical slots. We focus here solely on the client and line cards of POADM nodes, and we propose several solutions

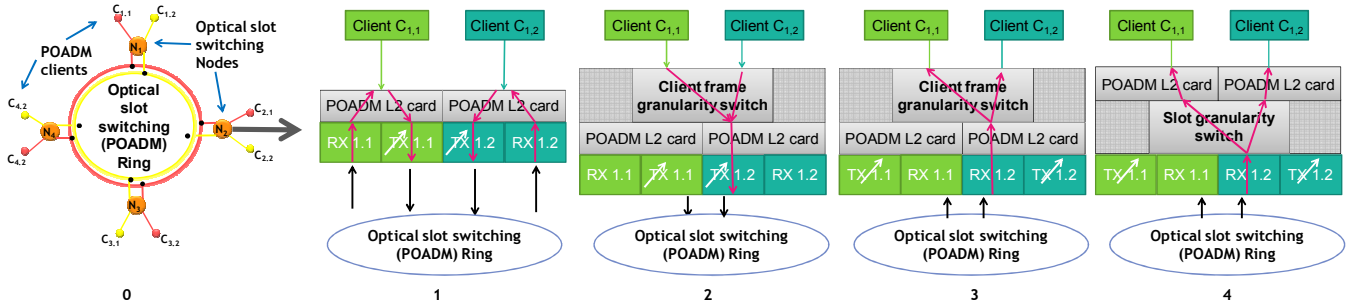

Figure 1:a) Optical slot switching (e.g., POADM) ring, 4 nodes and 2 clients per node; b) Basic node architecture; c) Node architecture with client granularity switch (c-switch) at transmitter. d) Node architecture with c-switch at receiver. e) Node architecture with slot granularity switch (s-switch) at receiver.

to electrically interconnect them within a node. Fig. 1 depicts the reference network for this study as well as the investigated node architectures. The network under study, as shown in Fig 1.a, is a ring with 4 POADM nodes (N_n, with node index n=1 … 4) and 2 clients per node ($C_{n,i}$ with node index *n* and client index *i*=1 or 2). For the sake of clarity, packet blockers located between the transmitters and the receivers are not depicted, see [2] for details. This ring supports here two wavelengths w_1 and w_2.

A. *Basic node architecture*

In the first architecture, each client is mapped, via a (L2) client card, to a TRX (line card) as illustrated in Fig. 1.b. Recall that, with the tunable Tx on the line side, a client can send its traffic on any wavelength, but the client is always mapped to a predefined wavelength for reception. In Fig. 1.a, for example, if client $C_{2,2}$ on node N_2 is mapped to wavelength w_2, then $C_{1,2}$ must use w_2 to reach client $C_{2,2}$ on N_2.

Since each pair (client card, line card) is independent, an optical slot can be formed only with frames from the same source client to the same destination client; thus in the previous example the considered optical slot only contains the frames from $C_{1,2}$ to $C_{2,2}$.

B. *Client frame electronic switch (c-switch)*

To enhance the flexibility of the previous architecture and remove the fixed mapping between client and wavelength of the basic architecture, we propose to include, either at transmitter or receiver (or even both) side, an electronic switch working at client frame granularity, further denoted by "c-switch". For instance, if the metro ring carries Ethernet traffic, those switches could be implemented with Ethernet switching fabrics. Note that this enhanced flexibility (and improved performance, as will be seen in Section IV), comes at the cost of adding at each POADM node one or two electronic switches working at the client frame granularity. These options are described as follows:

1) With client frame electronic switch at transmitter only:

When a client frame switch is used at the transmitter side (Fig. 1.c), the same line card will be shared between several clients (of the same node), such that frames coming from different clients can be encapsulated in the same optical slot, under the condition that these frames have the same (client and node) destination.

2) With client frame electronic switch at receiver only:

In this configuration, the electronic client frame switch is placed at receiver (Fig. 1.d), such that several clients can receive traffic encapsulated in the same optical slot. In this case, the optical slots are filled with frames coming from a single client on a source node, going to several clients on the destination node.

C. *Slot electronic switch (s-switch)*

We also propose to consider a less expensive, less energy consuming electronic switch working at slot granularity further denoted by "s-switch", rather than client frame granularity, between the client and the line cards, as illustrated in Fig. 1.e. Note that fast (slot-granularity) wavelength tunability, already present at the Tx side, provides the same functionality as an electronic slot switch, therefore the scenario with an electronic slot switch at Tx is not considered here, and only scenarios where this switch is located at the Rx side are considered. This switch enables the mapping of client cards to line cards at the slot (rather than client frame) granularity: it can be seen as an intermediate scenario between the no-switch case (Section II.A) and the client frame switch case (Section II.B).

D. *Combinations of c-switches and s-switches*

In addition to the aforementioned scenarios, we can also consider a combination of electronic switches at client frame granularity (at Tx, Fig. 1.c; and/or Rx side, Fig. 1.d) and at slot granularity (at Rx side, Fig. 1.e), resulting in 8 possible node architectures overall.

III. Model and Analytical Considerations

We derive here analytical results for the gain brought by the c-switches in terms of slot filling time, or, equivalently, client frame queuing delay assuming the channels are not congested, and the absence of a timer capping the maximum slot filling time. Those results are validated numerically in Section IV.

A. *Notation and notion of flow aggregation*

Let network load ($0 \le \text{load} \le 1$) be the *total load* (the total amount of traffic demand for the network), normalized by the *maximum supported load*: $\text{load} = \sum_{s,i,d,j} D_{s,i,d,j} / C_{\max}$, where $D_{s,i,d,j}$ is the traffic demand of (client-layer) data flow (s,i,d,j) from client *i* on node *s* to client *j* on node *d*, and $C_{\max} = 2WR$ [15] , where W is the number of wavelengths, each with capacity R.

The queuing delay for a client frame is the time it spends

Tx \ Rx		N_1		N_2		N_3		N_4	
		C_1 (w_1)	C_2 (w_2)	C_1 (w_1)	C_2 (w_2)	C_1 (w_1)	C_2 (w_2)	C_1 (w_1)	C_2 (w_2)
N_1	C_1			a	b	a	b	a	b
	C_2			c	d	c	d	c	d
N_2	C_1	a	b			a	b	a	b
	C_2	c	d			c	d	c	d
N_3	C_1	a	b	a	b			a	b
	C_2	c	d	c	d			c	d
N_4	C_1	a	b	a	b	a	b		
	C_2	c	d	c	d	c	d		

Figure 2: Clients/nodes/wavelengths mapping and traffic demand.

within the source node, i.e., the time between its emission by the client at the source node and its reception at the destination node's client card, assuming no propagation delay (propagation delays are deterministic and can easily be added, given physical link lengths). The optical slot is formed gradually by the arrival of client frames; once the maximum slot size or duration (fixed by the utilization of a Timer, see Section IV.B) is reached, the slot is buffered before it can be sent onto the fiber medium. Therefore, the queuing delay results from the waiting time spent in the node buffers and the optical slot filling time. The buffering time is directly related to the congestion of the channels.

The optical slot filling time is the time elapsed between the arrival time of the first frame in a slot, and the last frame forming this slot. If slots contain frames from a single flow, the corresponding slot filling time is a function of the arrival rate of client frames in this flow. When several flows are participating to the composition of an optical slot, their superposition is equivalent to a single flow with mean client frame arrival rate equal to the sum of the arrival rates from all participating flows, such that the optical slot filling time decreases when the client frame arrival rate increases, and depends on the number and sizes of the flows participating in the formation of the same optical slot. This process is called flow aggregation. Flow aggregation capability, and hence performance (queuing delay), strongly depends on the node architecture.

B. *Performance gains with c-switches*

In the following we derive the gain in latency seen by the client layer thanks to c-switches at Tx, Rx, or both, in the absence of congestion.

Consider client layer flow *(s,i,d,j)*. With no switch (neither at Tx nor Rx), the time to form an optical slot is $T_{s,i,d,j}$(no switch)$=K/D_{s,i,d,j}$. where constant K depends on the slot and (average) client frame sizes only. The average slot filling time between s and d (which, in absence of congestion, is also equal to the queuing delay of a client frame):

$$\overline{T_{s,d}^{\text{no switch}}} = \frac{1}{f_{s,d}} \sum_{i,j} \frac{K}{D_{s,i,d,j}} \qquad (1)$$

where $f_{s,d}$ is the number of flows from node *s* to node *d*.

With a c-switch at Tx on node *s*, i.e. when all flows from all clients of s to some client *j* of node *d* are aggregated into a single flow with intensity $A_{s,d,j} = \sum_i D_{s,i,d,j}$, hence the time to form an optical slot is: $T_{s,d,j}(\text{c} - \text{sw at Tx}) = K / A_{s,d,j}$. Let K_d be the number of aggregated flows from *s* to *d*, i.e., K_d is the number of clients on node *d*. The mean slot forming time between *s* and *d* is then:

$$\overline{T_{s,d}^{\text{Tx}}} = \frac{1}{K_d} \sum_j \frac{K}{A_{s,d,j}}. \qquad (2)$$

Hence the gain in queuing delay (during the slot formation process) for flow *(s,i,d,j)* brought by the presence of an e-switch at Tx side is:

$$G_{s,d}^{\text{Rx}} = \frac{\overline{T_{s,d}^{\text{no switch}}}}{\overline{T_{s,d}^{\text{Tx}}}} = \frac{K_d}{f_{s,d}} \frac{\sum_{i,j} 1/D_{s,i,d,j}}{\sum_j 1/A_{s,d,j}}. \qquad (3)$$

By symmetry, in the scenario with c-switch at Rx, where flows from one source client to several destination clients are aggregated within the same stream of optical slots, the gain is:

$$G_{s,d}^{\text{Tx}} = \frac{K_s}{f_{s,d}} \frac{\sum_{i,j} 1/D_{s,i,d,j}}{\sum_i 1/B_{s,i,d}}, \text{ where } B_{s,i,d} = \sum_j D_{s,i,d,j}. \qquad (4)$$

When both source node and destination node are equipped with a c-switch, flows from all clients of a node s to all clients of a node d are aggregated into a single flow with intensity $E_{s,d} = \sum_{i,j} D_{s,i,d,j}$, thus the time to form a slot between *s* and *d* is $T_{s,d}(\text{c - sw at Tx and Rx}) = K / E_{s,d}$. The gain in slot filling time brought by a c-switch at both Tx and Rx is thus:

$$G_{s,d}^{\text{Tx+Rx}} = \frac{E_{s,d}}{f_{s,d}} \sum_{i,j} 1/D_{s,i,d,j}. \qquad (5)$$

Note that s-switches have no impact on flow aggregation, and their impact cannot be measured in terms of latency when the network is not congested. The impact of the various node architectures on network performance are numerically investigated in the next section.

IV. Performance Studies

A. *Simulation description*

We evaluate the aforementioned node architectures for performance (latency) using an ns3-based network simulator.

The POADM unidirectional ring (a second direction could be used for protection) consists of 4 nodes, each connected to 2 clients, and 2 wavelengths, as seen in Fig. 1.0. Each node can transmit and receive on the two wavelengths via two transponders (fast tunable Tx + colored fixed receiver Rx); Fig. 2 shows the wavelengths/node/client mappings. Medium Access Control (MAC) issues are out of the scope of this paper, and an opportunistic MAC, which can use any available slot (without reservation) is assumed. Other MACs were proposed in [7]. Each flow (*s,i,d,j*) is assumed to made of fixed-size client frames of size f=558 bytes; client frames arrivals for each flow are Poisson (with average frame arrival rate $L_{s,i,d,j}=f/D_{s,i,d,j}$) and 100 ms-long simulations are run (Typically a simulation corresponds to the generation of a few several hundreds of thousands of client frames). The client-card and line-card rate is set to 10 Gb/s. The slot duration is set to roughly 8 µs (corresponding to 18 client frames). Slots can be electronically buffered once they are formed; at each client a buffer can contain up to 50 slots. The reported

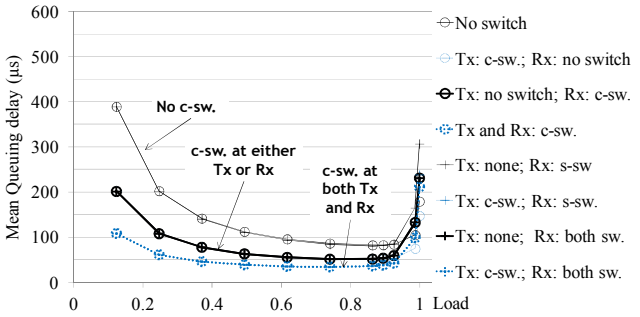

Figure 3: Mean queuing delay, without timer, a=b=c=d.

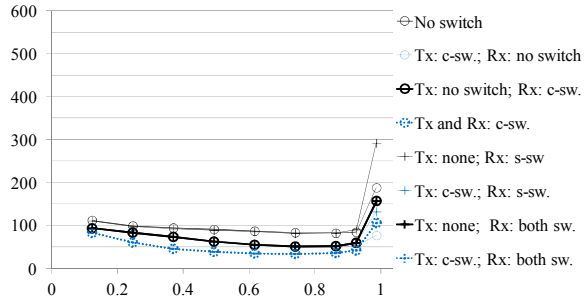

Figure 4: Impact of using a timer = 15 slots when filling an optical slot for uniform traffic (a=b=c=d).

queuing delays only account for frames successfully received and client frames lost due to congestion are ignored.

B. *Performance for uniform traffic*

We evaluate the 8 node architectures resulting from the combinations of a client frame switch (at transmitter and/or receiver) and/or an optical switch at receiver. We first assume traffic is uniform i.e. *a=b=c=d* in the demand matrix (Fig. 2).

1) *Load regimes:*

Fig. 3 shows the queuing delay for different loads and for each node architecture, *with no timer*. At low loads (e.g., load <0.4), the inter arrival time of frames to each node is large, leading to a large optical slot filling time, and hence relatively large queuing delay. At medium and high loads (e.g., 0.4 < load < 0.95), the optical slot filling time decreases, reducing the queuing delay. At very high loads (>0.95), unlike with the two first cases, buffering times are no longer negligible and queuing delays increase quickly due to the congestion of the network rather than increasing optical slot filling time (in fact the large number of arriving frames per second leads to fast formation of slots).

2) *Impact of s-switch*

In the uniform scenario, the load is balanced on the 2 wavelengths by construction and the addition of an s-switch at reception (which, as will be seen further, enables wavelength load balancing) brings no gain, hence in Fig. 3 the matching curves with s-switch and without s-switch are identical.

3) *Impact of c-switch at either Tx or Rx*

Architectures with a c-switch at either Tx or Rx aggregate, respectively, the flows from several clients (on the same source node) to the same destination client (c-sw. at Tx, Fig. 1.c), or the flows from one client to several clients (on the same destination node) (c-sw. at Rx; Fig. 1.d). Flows are uniformly distributed over all pairs of clients, hence the frame inter-arrival times for the aggregated flows are equal for both cases (c-switch at either Tx or Rx) resulting in identical queuing delays. In the regime where queuing delay is dominated by slot filling times (load < 0.95), when c-switch is used at either Tx or Rx, using (3) and (4) (for any (*s*,*d*) pair) the gain in queuing delay w.r.t the no-switch architecture is $G_{s,d}^{\mathrm{Tx}} = G_{s,d}^{\mathrm{Rx}} = 2$. This can be verified in Fig. 3.

4) *Impact of c-switch at both Tx and Rx*

The largest reduction in queuing delay is provided by the architecture with c-switches at both Tx and Rx. This is due to the larger number of aggregated flows, leading a higher reduction of the optical slots filling time, compared to the other architectures, and a better pooling of the network resources (Tx, Rx, wavelengths) across the clients. In this case the gain in latency w.r.t. the no-switch scenario is $G_{s,d}^{\mathrm{Tx+Rx}} = 4$ (see (5)). This can be verified in Fig. 3.

5) *Timer*

In order to bound the queuing delay at low loads, we force optical slots to be sent whether they are fully filled or not upon expiration of a (per-slot) timer initialized when the slot is created. This sets an upper bound on the slot filling time, at the expense of a potential channel capacity waste, since an optical slot uses an entire time slot, even if it may not be completely filled. This latency vs. channel capacity waste trade-off requires an adequate selection of the timer value.

Fig. 4 presents the results of the same previous simulation scenario, now with a timer value *Ts*=15 slots (i.e., 120 µs). We notice, compared with Fig. 3, that the queuing delay is effectively bounded by the slot filling time (=15 slots) for the lower loads, until congestion occurs at higher loads.

The utilization of a timer of 10 slots results in frame losses (not shown here); a timer of 15 slots is a good value since there is no loss for load < 0.95 (not shown here). Hence, in all further simulations, we use a timer value *Ts*=15 slots.

C. *Non-uniform case: a=c, b=d, a ≠ b*

In order to refine the node architectures comparison, we modify the traffic matrix; here we assume that *a=c* and *b=d* but $a \neq b$. In this case the traffic matrix is uniform at the node level (each node sends/receive the same quantity of traffic) but not at the client level, such that the two wavelengths carry a different amount of traffic (traffic sent by each node: 6*a* on w_1 and 6*b* on w_2). The behavior of the network differs depending on whether *a* is small or large, as seen below.

1) *'a' small*

Fig. 5 shows the queuing delay, for each node architecture, for *a* = 50 Mb/s and *b* varying between 50 and 1600 Mb/s.

a) *Impact of s-switch*

In the absence of a switch at the receiver side, there is no possibility to balance the load over the wavelengths, resulting in a rapid increase of the queuing delay and important losses at load > 0.5 (the plateau in Fig. 5 corresponds to congestion of w_2, buffers overflows, and important losses, while w_1 is lightly used), corresponding to the maximum capacity supported by a single wavelength. Using an s-switch at Rx enables load balancing and removes the plateau.

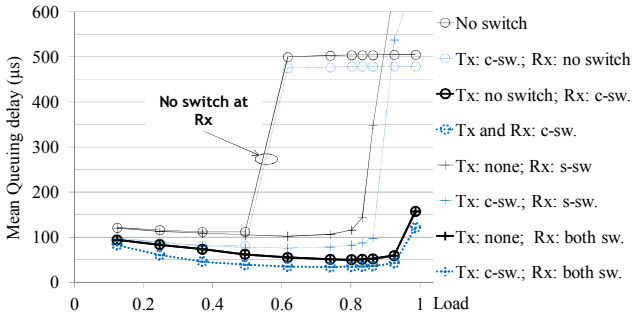

Figure 5: Mean queuing delay, using timer = 15 slots, a=c=50Mb/s, b=d= [50, 1600] Mb/s.

b) c-switch at Tx only; no s-switch

Since the traffic is not symmetric, the slot filling time varies depending on whether aggregation of the client flows is enabled by a c-switch at Tx or at Rx. When queuing delay is dominated by slot filling times (0.4 < load < 0.95), when c-switch is used at Tx only, using (3) (for any (*s*,*d*) pair) the gain in queuing delay w.r.t the no-switch architecture is $G_{s,d}^{\mathrm{Tx}}=2$, as verified in Fig. 5. However, the absence of a switch at Rx results in a plateau, as explained above.

c) c-switch at Rx only

Using a c-switch at Rx removes the latency plateau (see Section IV.C.1)a)). When queuing delay is dominated by slot filling times, the queuing delay gain with c-sw. at Rx only is larger than with a c-sw. at Tx only: $G_{s,d}^{\mathrm{Rx}}=(a+b)^2/(2ab)>2$.

d) c-switch at both Tx and Rx

When an c-switch is used at both Tx and Rx the gain is even larger, and $G_{s,d}^{\mathrm{Tx+Rx}}=(a+b)^2/(ab)$.

e) Maximum supported load

The maximum load supported by the network strongly depends on the node architecture. Fig. 5 shows that the highest supported load (above which congestion occurs and queuing delay diverges) is below 0.86 using the node architecture without c-switches, 0.89 using the c-switch at Tx, 0.95 using c-switch at Rx and 0.98 with switches at both Tx and Rx.

For certain node architectures, enumerated below, queuing delay diverges because the offered channel capacity is exceeded. Indeed, with a timer (with value *Ts*) the capacity required by any flow is *S*/*Ts* b/s (*S* is the slot size in bits) even if the corresponding demand is lower. With the architectures with no c-switch, client flows cannot be aggregated and it can easily be shown that a capacity of 650 Mb/s is used by any client flow with demand below 650 Mb/s for *Ts*=15 slots; in Fig. 5 some client demands are much lower (*a*=50 Mb/s, 50<*b*<1600 Mb/s), leading to a maximum supported load (=0.86) below the theoretical limit (=1).

In the presence of c-switches at both Tx and Rx, the slot filling time is always below Ts and optical slots are always completely filled, ensuring that all the channels capacity can be used, leading to low queuing delays even for loads very close to 1. For the architectures with c-switches at either Tx or Rx, the slot filling time is below Ts but larger than in the case with c-switch at both Tx and Rx. This allows higher channel utilization than with the basic architecture, but lower than with the architecture with both c-switches.

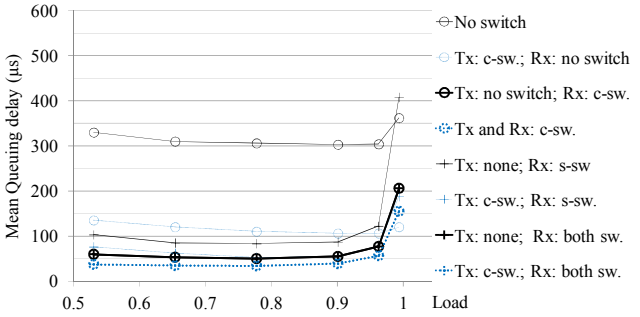

Figure 6: Mean queuing delay, using timer = 15 slots, a=c=810Mb/s, b=d= [50, 800] Mb/s.

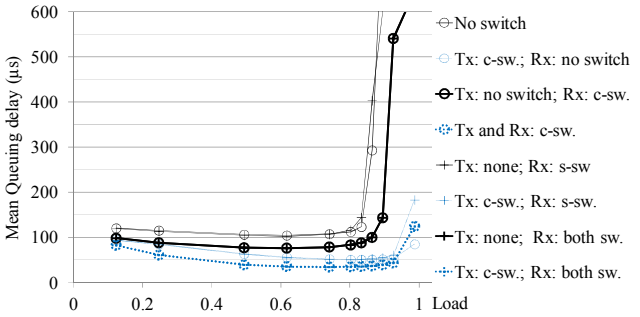

Figure 7: Mean queuing delay, using timer = 15 slots, a=b=50Mb/s, c=d= [50, 1600] Mb/s.

2) 'a' large

Fig. 6 presents the results again for *a*=*c* and *b*=*d*, but now using *a* = 810Mb/s and *b*= [50,800] Mb/s, so that the capacity of w_1 is no longer exceeded, as was the case for `a' small. The high queuing delay for the basic architecture is due to the saturation of the wavelengths: w_1 by the large demands (*a* = 810Mb/s), and w_2 because the timer transforms demands *b*= [50, 800] Mb/s into demands of at least 650 Mb/s, as seen above. Each of c-switch or s-switch at Rx decreases the latency (compared with the basic, no-switch case); in addition, the queuing delay is lower with a c-switch than with an s-switch. However, the combination of both a c-switch and s-switch at Rx does not bring any gain in queuing delay with respect to using a c-switch only[1].

D. Non-uniform traffic, a=b, c=d, a ≠ c:

In this case the traffic matrix is again uniform at the node level but not at the client level; however the traffic demand is now uniformly distributed over the two wavelengths. As in Section IV.B, queuing delays relative to the basic architecture and the architecture with s-switch at Rx are the same.

Fig. 7 shows the queuing delay according to the load for *a* = 50 Mb/s and *b*= [50, 1600] Mb/s. Compared with Section IV.C, the delays of the median curves are reversed, i.e. the electronic switch at Tx side significantly reduces latency and decreases the capacity wasted through the utilization of the timer, w.r.t. the use of electronic switch at Rx side.

In the regime where queuing delay is dominated by slot filling times (0.4 < load < 0.95), when c-switch is used at Rx only, using (4) (for any (*s*,*d*) pair) the gain in queuing delay

[1] Note that the same kind of reasoning cannot be applied at transmitter side, since as Tx are tunable to ensure each nodes can communicate and Rx are fixed-wavelength, the case with s-switch only at Tx does not exist.

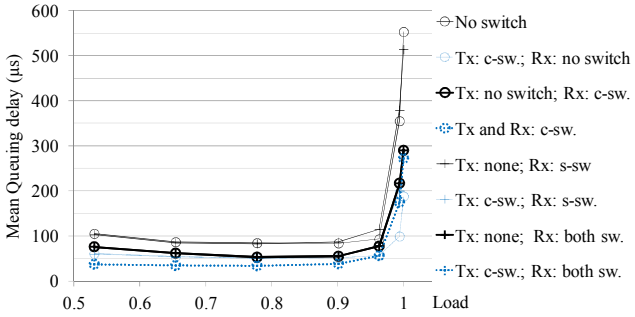

Figure 8: Mean queuing delay, using timer = 15 slots, a=b=810Mb/s, c=d= [50, 800] Mb/s.

w.r.t the no-switch architecture is $G_{s,d}^{\mathrm{Rx}}=2$, as verified in Fig. 7. The queuing delay gain is larger: $G_{s,d}^{\mathrm{Tx}}=(a+c)^2/(2ac) > 2$ when a c-switch is used at Tx only; when a c-switch is used at both Tx and Rx the gain is even larger: $G_{s,d}^{\mathrm{Tx+Rx}}=(a+c)^2/(ac)$. It is easily verified that the values for the uniform case are obtained for a=b in the expressions above.

Fig. 8 shows the queuing delay for a=810 Mb/s and b= [50,800] Mb/s. Compared with Fig. 6, the difference noticed in Section IV.C regarding the utilization of switches at Rx vanishes. This is because the traffic is uniformly distributed in Fig. 8, cf. Section IV.B for the full explanation.

V. Conclusion

In this paper we proposed and compared the performance (in terms of client frame queuing delay) different possible electrical architectures for optical slot switching nodes, embedding electrical switches at two different granularities: client frame and/or slot, at Tx and/or Rx.

Table 1 summarizes the maximum achievable load to ensure a queuing delay below 200 µs, for each node architecture. Architectures without any switch at Rx do not support some of the traffic demands (dark grey cells), because of their inability to balance the load, as seen above. The architectures with c-switch at either Tx or Rx permit to support a higher load (at least 80%), depending to the aggregated traffic. The architectures with c-switch at both sides sustain the highest load (more than 95%).

Although the basic node architecture with no electrical switch presents the worst queuing delay in all simulated scenarios, this worst case never exceeds a relatively low value in the order of magnitude of 100 µs at loads below 80%. This node architecture presents the advantage of lower cost and consumption with respect to the other architectures. However, with small flows, slots are filled slowly and we have to use a timer to bound the queuing delay. With the basic node architecture the timer results may in forming poorly filled slots and wastes channel capacity: resulting at high loads in an increase of queuing delays and in large frame losses.

The lowest queuing delay is achieved with client frame switches at both Tx and Rx sides. This architecture permits to reduce the time of optical slot formation while ensuring a high optical slot filling rate and hence better utilization of all network resources (Tx/Rx capacity, channel capacity).

The utilization of a client frame switch at Tx or Rx presents intermediate performances, which depend heavily on the number and sizes of the aggregated client flows.

	a=b=c=d (Fig. 4)	a=c, b=d, a small (Fig 5)	a=c, b=d, a large (Fig 6)	a=b, c=d, a small (Fig 7)	a=b, c=d, a large (Fig 8)
No switch	>0.95	0.5	0	0.86	>0.95
Tx: c-sw	>0.95	0.5	>0.95	>0.95	>0.95
Rx: c-sw	>0.95	>0.95	>0.95	0.89	>0.95
Tx : c-sw. Rx: c-sw.	>0.95	>0.95	>0.95	>0.95	>0.95
Rx: s-sw.	>0.95	0.86	>0.95	0.86	>0.95
Tx: c-sw.; Rx: s-sw.	>0.95	0.89	>0.95	0.86	>0.95
Rx: c-sw & s-sw.	>0.95	>0.95	>0.95	0.89	>0.95
Tx: c-sw.; Rx: c-sw & s-sw.	>0.95	>0.95	>0.95	>0.95	>0.95

Table 1: maximum achievable load to ensure a queuing delay below 200 µs.

The utilization of an electronic slot switch at Rx enables load balancing. The node architecture with this switch only does not permit flow aggregation and is therefore less efficient than those with client frame switches.

In this paper we focused solely on performance analysis and the quantified trade-off in terms of CAPEX (or energy) for the various architectures was not evaluated. Future work will consider joint CAPEX and performance analysis of optical slot switched networks.

Acknowledgment

Work partly supported by the CELTIC+ SASER project.

References

[1] S. Korotky, "Semi-empirical description and projections of internet traffic trends using a hyperbolic compound annual growth rate," Bell Laboratories, Technical Memorandum, 29 November 2012.

[2] D. Chiaroni, G. Buforn Santamaria, C. Simonneau, S. Etienne, J.- C. Antona, S. Bigo, and J. Simsarian, "Packet OADMs for the next generation of ring networks," Bell Labs Technical Journal, Winter 2010.

[3] Y. Pointurier, B. Uščumlić, I. Cerutti, A. Gravey and J.-C. Antona, "Dimensioning and energy efficiency of multi-rate metro networks", IEEE/OSA Journal of Lightwave Technology, Nov. 2012.

[4] J. Simsarian and L. Zhang, "Wavelength locking a fast-switching tunable laser," IEEE Photonics Technology Letters, Jul. 2004.

[5] B. Uščumlić, A. Gravey, P. Gravey and I. Cerutti, "Traffic Grooming in WDM Optical Packet Rings", in Proc. ITC 21, 2009.

[6] B. Uščumlić, I. Cerutti, A. Gravey, P. Gravey, D. Barth, M. Morvan, and P. Castoldi, "Optimal dimensioning of the WDM unidirectional ECOFRAME optical packet ring," Photonic Network Comm., 2011.

[7] T. Bonald, S. Oueslati, J. Roberts and C. Roger "SWING: Traffic Capacity of a Simple WDM Ring Network", in Proc. ITC 21, 2009.

[8] M.A. Marsan, A. Bianco, E. Leonardi, A. Morabito and F. Neri, "All-optical WDM multi-rings with differentiated QoS", IEEE Communications Magazine, Feb. 1999.

[9] F. Hacimeroglu and T. Atmaca, "Impacts of Packet Filling in an Optical Burst Switching Architecture", in Proc. IEEE AICT, 2005.

[10] H.-C. Leligou, K. Kanonakis, J. Angelopoulos, I. Pountourakis and T. Orphanoudakis, "Efficient burst aggregation for QoS-aware slotted OBS systems", European Transactions on Telecommunications, 2006.

[11] A. Pantaleo, M. Tornatore, A. Pattavina, C. Raffaelli, and F. Callegati, "Dynamic service differentiation in OBS networks", in Proc. Broadnets, 2007.

[12] V.M. Vokkarane, K. Haridoss and J.P. Jue, "Threshold-Based Burst Assembly Policies for QoS Support in Optical Burst-Switched Networks", in Proc. OptiComm 2002.

[13] X. Yu, Y. Chen and C. Qiao, "Performance evaluation of optical burst switching with assembled burst traffic input", IEEE Network, 2004.

[14] T. Eido, D.T. Nguyen and T. Atmaca, "Packet Filling Optimization in Multiservice Slotted Optical Burst Switching MAN Networks", in Proc. IEEE AICT, 2008.

[15] R.-M. Indre, "Performance Analysis in Subwavelength Switching Optical Networks," Ph.D. Thesis, Telecom ParisTech, 2012.