

Performance Analysis of a Hybrid Optical–Electronic Packet Switch Supporting Different Service Classes

Wiem Samoud, Cédric Ware, and Mounia Lourdiane

Abstract—As a solution to the high energy consumption caused by numerous optical–electrical–optical (O-E-O) conversions in electronic switches, and the poor contention handling of all-optical switches, we investigated a hybrid switch that supplements optical switching with an electronic buffer. Our study takes into account reliable, fast, and default packets that have different requirements of quality-of-service performance criteria. We show, by simulations, that with only a few electronic ports to the buffer, the hybrid switch significantly improves the packet loss rates and the sustainable system load compared to an all-optical bufferless switch and meets the different packet classes' requirements. In addition, we quantified the considerable decrease of O-E-O conversions as well as the switching latency achieved by the hybrid switch compared to an electronic switch.

Index Terms—Contention resolution; Optical communication equipment; Optical packet switching; Optoelectronic devices.

I. INTRODUCTION

With the exponential increase of traffic and the need to support new services, many challenges must be met in designing future optical networks, notably curbing their energy consumption. The fact is that, although most traffic is carried as optical signals, at a relatively low energy per bit, packet routing cannot currently be performed in the optical domain. All traffic transmitted on optical fibers must be converted to the electrical domain at the level of routers and switches, and then reconverted to optical signals for further transmission. Reamplifying, reshaping, and retiming (3R) functions are performed electronically. Consequently, numerous optical–electrical–optical (O-E-O) conversions are required and result in a rapid growth of energy consumption [1]. We note that the network electricity consumption is growing fast, at a rate of 10% per year [2]. A way to address this problem is to think about different switching technologies that may provide economical energy consumption.

Manuscript received April 3, 2015; revised July 7, 2015; accepted August 10, 2015; published August 31, 2015 (Doc. ID 236657).

W. Samoud (e-mail: wiem.samoud@telecom-paristech.fr) and C. Ware are with Institut Mines-Télécom, Télécom ParisTech, CNRS LTCI, Paris, France.

M. Lourdiane is with Institut Mines-Télécom, Télécom SudParis, CNRS SAMOVAR, Évry, France.

<http://dx.doi.org/10.1364/JOCN.7.000952>

All-optical packet switching would do away with O-E-O conversions and reduce energy consumption. However, as all-optical buffering solutions are still impractical, those switches are extremely vulnerable to contention in which two or more coincident packets are to be forwarded to the same output port. This leads to notable packet loss rates (PLRs) even at unrealistically low loads [3]. Thus, except for quasi-static switching, all-optical switching solutions have not yet succeeded outside research laboratories.

A smarter approach could be combining the benefits of both optics and electronics. Therefore, a hybrid optoelectronic switch was proposed [4] and demonstrated [5,6]. According to abilities, roles are assigned between optics and electronics: the hybrid switch uses an optical switching matrix to route packets rapidly at a low energy cost if possible, or stores them in a shared electronic buffer if not to avoid their loss due to contention.

An analysis of a hybrid optical–electronic switch [7] shows performance improvements in terms of PLR and sustainable load compared to the all-optical case, with relatively few buffer ports. This was confirmed by simulations and an Engset-type analytical model. Another analysis consolidates these results while taking into account different classes of service [8]. Packets' priorities were defined only in terms of the PLR, and based on that, different techniques of buffer input ports' access were investigated. The study [9] integrates the hybrid switch in an intra-data-center network but without priority classification and with a different switching strategy than in the previously cited references; it employs deflection routing and fiber delay lines (FDLs) in addition to the electrical buffer. Considering the PLR as a performance criterion, the use of FDLs is discussed depending on the system load. In our studies, only a shared electrical buffer supplements the optical switching matrix to avoid contention.

In this article, we analyze by simulations the performance of a hybrid switch that supports different priority packets. Compared to former publications, besides the PLR and the sustainable system load, we take into account the latency (the delay) as a performance criterion. Thus, the packet classification is fairly realistic: reliable, fast, and default packets.

First, the hybrid switch architecture and the switching policy are described in Section II. The switching policy has been established in order to find a trade-off between the

PLRs and the latencies and to satisfy the constraints of each packet class, mainly null PLR for reliable packets and low latency for fast packets. Second, we present, in Section III, the performance improvements in terms of the PLR and the sustainable load obtained by the hybrid switch compared to an all-optical one. Our simulations show that compared to the total number of optical ports, very few electronic buffer ports are sufficient to fulfill the requirements of the different packet classes. The switching latency is also acceptable, even for fast packets. Since we consider different parameters such as the number of buffer ports and the number of channels per connected azimuth, our simulation results help dimension the switch by giving the intervals of those parameters' values that lead to a good trade-off between performance improvements and energy savings. Third, a comparison with commercial (electrical) switches is given in Section IV, especially in terms of switching latency, which is reduced because only a fraction of the packets are converted to the electronic domain to be buffered. In addition, the reduction of O-E-O conversions is quantified and indicates that the hybrid switch may be a potential solution to the energy consumption problem. Section V concludes the paper and presents some of our perspectives.

II. ARCHITECTURE AND SWITCHING POLICY

The hybrid switch architecture is presented in Fig. 1. It is composed of an optical switching matrix supplemented with an electronic shared buffer that stores packets in case of contention. The main parameters to dimension the switch are the number of connected azimuths (n_a) that are supposed to be bidirectional, the number of supported channels per azimuth in each direction (n_c), and the number of electronic input ports as well as output ports (n_e). The required packet label processing and control unit for the switching matrix is supposed to be generic and is not presented in Fig. 1; our study assumes that label processing does not require O-E conversion of the whole packet, e.g., by sending labels out-of-band [10], or as a reduced-bit-rate header. We supposed that the channels are independent; an azimuth may receive up to n_c packets simultaneously. Channels are also supposed to be interchangeable: an ingress packet can use any available channel of its egress azimuth. This assumption works with space division multiplexing (SDM) non-wavelength-specific channels, such as parallel optical fibers in the same

cable or different cores in a multi-core fiber. Wavelength division multiplexing (WDM) channels, although more commonly used, add the wavelength constraint and affect the switch performance [11]; the present study is not directly applicable to them, at least not without liberal use of wavelength converters, which themselves consume power and may negate the energy savings achieved by the hybrid switch compared to an electronic one.

We consider three different service classes: reliable (R), fast (F), and default (D) packets, respectively, make up 10%, 40%, and 50% of the global traffic. These percentages are adapted from real data of circulating packets on metro and core networks [12]. Reliable packets (R)—which may refer to digital data and file transfer packets—must reach their destinations without loss, but they are the lowest priority packets in terms of delay. Fast packets (F)—which could refer to voice and interactive video packets—have the highest priority regarding latency, but they are more tolerant than R packets with regard to the PLR. Default packets (D)—which represent other types of packets—are the least restrictive with respect to both PLR and latency.

We considered a fixed packet duration of $\sigma = 10 \mu\text{s}$, which represents about 100 kbits for standard 10 Gbit/s systems. It may correspond to a jumbo Ethernet frame or an aggregation of several IP packets [7]. Thus, the system load ρ will depend only on the mean idle time per source τ , which is simply the average time interval separating two consecutive packets arriving from the same channel of a given azimuth:

$$\rho = \frac{\sigma}{\tau + \sigma}. \quad (1)$$

In our simulations, τ is generated randomly. The system is called “fully loaded” ($\rho = 1$) when packets are sent one after another unceasingly ($\tau = 0$).

The hybrid switch works in asynchronous mode: packets can arrive at any instant. Given this assumption, our study is equivalent to the case of having variable σ . Figure 2 describes the switching strategy. At the reception of a packet, the switch checks whether there is an available channel to its egress azimuth. If yes, the packet is directly sent on its way over this channel. Otherwise, if an electronic input port is available, the packet is buffered and then reemitted whenever a channel is released and an output electronic port is available. The first in first out (FIFO) technique is applied at the level of the buffer output ports: the first buffered packet is the first one to be reemitted; in addition, the reemission of buffered packets has priority over incoming packets for a given destination. Otherwise, depending on the packet class, there is a preemption policy: if the newly arrived packet is of type R , the switch may interrupt the transmission through the optical switching matrix of the last (preferably D , or F) packet being sent to the buffer or the last (preferably D , or F) packet being sent to the same egress azimuth and send the R packet preferentially. Otherwise, if the newly arrived packet is of type F , the switch checks whether there is a D packet being sent to the same egress azimuth to preempt it and send the F




Fig. 1. Hybrid switch architecture.




Fig. 2. Switching strategy.

one instead. The preempted packet, which was still being transmitted through the optical switching matrix to its destination port or to the buffer input port when receiving the new higher priority packet, is dropped and taken into account in calculating the PLRs. In the worst case, in the absence of any of the possibilities listed above, the incoming packet is dropped.

This switching strategy, and especially the preemption policy, is chosen in order to meet each service class constraint of both PLR and latency. In fact, we simulate different policies, which are not detailed in this article, by changing the order of preemptions, and based on the PLR and latency of each class of service, we opt for the policy described above since it leads to the best trade-off between the PLR and the latency for all the packet classes.

Given the hybrid switch architecture and its switching policy, we present in the following section its performance considering different criteria.

III. HYBRID SWITCH PERFORMANCE

We implemented a dedicated simulator in C++. It takes as input parameters the hybrid switch dimensions: n_a , n_c , and n_e . Given these parameters, we start by simulating a fully loaded system ($\rho = 1$), launch packets, and calculate the PLR and the latency for each service class. The packets' destinations are chosen randomly. The traffic is symmetric: ρ is equal for all the channels of the connected azimuths.

Because of the simulations' time cost, we run simulations until having an average PLR of all classes that is less than or equal to 10^{-7} . We vary the system load (ρ) by decreasing it by 5% on each step and calculate the PLRs and the latencies. If the mean PLR is less than 10^{-4} , we decrease ρ with just 1.25% instead of 5% to be more accurate. For a given ρ , we start counting the PLRs and the latencies after the first drop of a packet. Each simulation ends if either

- 100 reliable packets are dropped, ensuring the accuracy of the resulting calculated PLR to 10%
- or the number of all switched packets is $n_p \geq 4 \times 10^8$, corresponding to $n_p(R) = 4 \times 10^7$ reliable packets, which is a large enough value to measure PLR_R down to about 10^{-7} with 95% confidence.

This lower bound stems from the fact that, at a given $\text{PLR} = p$, the probability of transmitting N packets without dropping any is $(1-p)^N \simeq \exp(-Np)$; therefore, if we lose no reliable packet at all (which is indeed the case in most of our simulations) after transmitting $n_p(R)$ of them, there is a less than 5% chance of PLR_R being actually higher than $-\ln(5\%)/n_p(R) \simeq 0.75 \times 10^{-7}$ [13].

We considered a degree-8 switch ($n_a = 8$). n_c takes the values of 1, 4, 8, 20, and 30 channels per azimuth, while n_e takes different values from 0 (which corresponds to an all-optical bufferless switch) to the number of optical links $n_a \times n_c$ (which corresponds to an all-electrical switch). This permits us to know when the hybrid switch is of interest and to determine the minimum value of n_e that is sufficient to have acceptable performance.

In particular, we will focus on the case of $n_a = 8$ and $n_c = 8$ because we want to compare the performance of the hybrid switch with the commercial Cisco Nexus 3064PQ electrical switch that has 64 ports [14]. For this example, n_e takes the values of 0, 3, 5, 10, 20, 30, 40, and 64 electronic ports.

In the present investigation, we did not focus on dimensioning the shared buffer capacity, which we assumed to be infinite. However, to ensure that the memory requirements remain reasonable, we performed sampling simulations in which we observe the number of packets that are currently in the buffer or being sent to the buffer (n_{buff}). Our simulations show that this number is always of the same order of magnitude as $2 \times n_e$ even for a fully loaded system. For example, Fig. 3 presents the histogram of the percentage of occurrences of each value of n_{buff} , where $n_a = 8$, $n_c = 8$, $n_e = 12$ or 20, and $\rho = 1$. Sampling is performed after every 10^6 circulated packets. n_{buff} is respectively around 24 and 40 when $n_e = 12$ and 20. Other simulations with different values of the switch dimensions n_a , n_c , and n_e show the same result: the demand on the buffer capacity remains reasonably limited.




Fig. 3. Histogram: percentage of occurrence of the number of buffered and being buffered packets, $n_a = 8$ and $n_c = 8$. Left, $n_e = 12$; right, $n_e = 20$.

A. Performance in Terms of the PLR

Figures 4 and 5 show, respectively, the evolution of PLR_D and PLR_F as functions of the system load for eight azimuths, eight channels per azimuth, and different values of n_e . We considered a system load of 60% as a minimum acceptable operating point. It is a widespread reference value taken into account in several articles [7,15]. Other references consider that the system is heavily loaded when $\rho > 0.7$ [9].

We note that with the inclusion of just a few electronic ports ($n_e = 5$, for example), PLRs decrease significantly compared to an all-optical switch. At $\rho = 0.6$, PLR_D is reduced by a factor of 5 from 10^{-1} to 2×10^{-2} , while PLR_F is reduced from 2×10^{-3} to 1.4×10^{-4} , which means a reduction by better than an order of magnitude. The more ports the buffer has, the greater the decrease of PLR_D and PLR_F . For a 60% loaded system ($\rho = 0.6$), and with only 20 electronic ports, PLR_D is around 10^{-7} , while PLR_F is around 10^{-8} . In our simulations, thanks to the proposed switching strategy, no R packets were lost using the hybrid switch. In fact, R packets have the favor of preemption at the level of




Fig. 4. PLR_D versus system load ($n_a = 8$, $n_c = 8$).




Fig. 5. PLR_F versus system load ($n_a = 8$, $n_c = 8$).

the buffer access and also at the level of the optical ports. Thus the PLR constraint of the reliable class service is satisfied. We may say that PLR_R is under the sensitivity of our simulations, or at least it is always less than $-\ln(5\%)/n_p(R) \simeq 10^{-7}$, where $n_p(R)$ is the number of all switched R packets.

Thus, thanks to our switching strategy, our hybrid switch satisfies all the service classes' requests in terms of PLR, especially the R packets. In addition, it makes a great improvement compared to an all-optical switch, and it is therefore a good solution for the contention issue.

B. Performance in Terms of Sustainable System Load

The sustainable system load at a given value of PLR is the maximum system load for which the PLR is less than or equal to the given value. For a degree-8 switch, we plot in Figs. 6 and 7 the evolution of the sustainable system load at $PLR_D = 10^{-4}$ and at $PLR_F = 10^{-4}$, respectively, as a function of the ratio between the number of electronic ports and the optical links ($n_e/(n_a \times n_c)$). This ratio refers to the reduction of the number of electronic ports by the hybrid switch (n_e ports) compared to an electrical switch ($n_a \times n_c$ ports). The reduction of the electronic ports leads to a reduction of the energy consumption. We note that for interactive video packets (F packets), it is recommended that the PLR must be lower than 10^{-2} across the network [16]. Assuming paths can cross up to 100 nodes, we take as a reference a single node $PLR_F = 10^{-4}$. We also impose the same constraint to PLR_D even though D packets are more tolerant to the PLR.

The hybrid switch is considered of interest when it satisfies two conditions: first, since we considered a system load of 60% as a minimum acceptable operating point, the sustainable system load must be ≥ 0.6 . Second, the buffer must incur significantly fewer O-E-O conversions than an all-electronic switch of the same size, which we choose to express as the condition $n_e \leq (n_a \times n_c)/2$. The area where




Fig. 6. Sustainable system load versus $n_e/(n_a \times n_c)$ at $PLR_D = 10^{-4}$ ($n_a = 8$).




Fig. 7. Sustainable system load versus $n_e/(n_a \times n_c)$ at $PLR_F = 10^{-4}$ ($n_a = 8$).

these two conditions are fulfilled is presented by the rectangles in Figs. 6 and 7 and permits us to find a trade-off between the performance improvement and the energy savings.

The sustainable load increases with n_e and reaches 1 for $n_e = n_a \times n_c$, where an ingress packet can always be collected by the buffer if it cannot be directly switched. Figures 6 and 7 show that whenever the degree-8 hybrid switch is of interest, n_c and n_e have to be chosen among the values located inside the rectangles.

For a 60% loaded system, with four or more channels per azimuth, just $(0.4 \times n_a \times n_c)$ electronic ports are sufficient to have $PLR_D \leq 10^{-4}$, while PLR_F is less than or equal to 10^{-4} for $(0.3 \times n_a \times n_c)$ electronic ports. Thus, the hybrid switch leads to an acceptable sustainable system load with the number of electrical ports lower than half that of optical links, or even fewer.

C. Performance in Terms of the Delay

Another important performance criterion is the latency, also called the delay. It is the additional time it takes for a packet to arrive at its destination if it could not be switched directly to its egress, in the optical domain, because of contention. By this definition, the latency of a packet switched directly in the optical domain is null. As for packets that transit through the buffer, the three delays to consider are the duration of the packet's reception into the buffer, equal to σ ; the time spent in the buffer; and the packet's reemission time out of the buffer, also σ . Out of these delays, the first one does not count since, assuming that the packet is directly committed to memory as it comes, it is equal to the transit time through the optical switching matrix and we defined latency to be the additional time. Thus,

$$\text{mean latency} = \frac{\sum_{\text{packets}} (t_{\text{in the buffer}} + t_{\text{reemission}})}{\text{nbr of buffered then sent packets}}. \quad (2)$$

Figure 8 shows the evolution of the average delay for each class of service versus the system load, for different

values of n_a , n_c , and n_e . Curves presenting Delay_F and Delay_D as functions of the system load are increasing then decreasing curves. Delays have nonzero values for a fully loaded system ($\rho = 1$). Considering as an example $(n_a, n_c, n_e) = (8, 8, 10)$, in Fig. 8(a), $\text{Delay}_D = 1.4 \mu\text{s}$. Then, delays increase as ρ decreases; for the same example, Delay_D has its maximum value ($1.68 \mu\text{s}$) at $\rho = 0.8$. But after that, delays decrease until reaching a null value. In fact, for a heavily loaded system, F and D packets are often preempted by R packets when they are being transmitted to the buffer. They are less-commonly preempted when they are being sent directly to their destinations or during their retransmissions. So, F and D packets have little chance to be buffered; if they arrive at their destinations, in the majority of cases, they were sent directly. When ρ decreases, R packets have more chance to have free channels to their destinations or available buffer input ports, and so F and D packets are more likely to reach the buffer without being preempted. Therefore Delay_F and Delay_D increase (and of course their PLRs decrease). If ρ decreases more, it is even more common that F and D packets are sent directly to their destinations, and thus their delays decrease.

To further argue our interpretations, we present as an example for $(n_a, n_c, n_e) = (8, 8, 12)$, in Fig. 9, the evolution of the number of packets sent directly in the optical field to their destinations, and the number of packets that were buffered before being successfully sent to their final egress, as functions of ρ . Indeed, the curves presenting the “buffered then successfully sent” F and D packets have the same shape as those of Delay_F and Delay_D .




Fig. 8. Delays of R , F , and D packets versus system load. (a) $n_a = 8$ and $n_c = 8$, (b) $n_a = 10$ and $n_c = 10$, and (c) $n_a = 10$ and $n_c = 20$.




Fig. 9. Number of packets successfully sent to their destinations, directly or after bufferization, versus system load ($n_a = 8$, $n_c = 8$, and $n_e = 12$).

Besides, as seen in Fig. 8(a), for a low number of electronic ports ($n_e \leq 10$), Delay_D is lower than Delay_F . This is due to the numerous preemptions of D packets that are, consequently, rarely buffered. However, for a higher number of electronic ports ($n_e \geq 10$), Delay_F is lower.

When $n_a = 8$ and $n_c = 8$, even with 40 electronic ports, switching delays are less than 10 μs . Considering other switch dimensions, we present as examples in Figs. 8(b) and 8(c) the evolution of the delays as a function of ρ for $n_a = 10$ and $n_c = 10$ or 20. Delays are respectively less than 6 and 3.5 μs when ($n_c = 10$ and $n_e = 40 = 0.4 \times n_a \times n_c$) and when ($n_c = 20$ and $n_e = 40$).

In practice, voice and interactive video (fast packets) require 150 ms one-way, end-to-end delay, while streaming video has much laxer requirements because of a high amount of buffering that has been built into the applications [16]. Our simulation results show that switching delays are less than 10 μs , with four orders of magnitude below acceptable limits. Therefore, many hybrid switches may be used on a network while respecting the delay condition, especially for fast packets.

In addition, we tried another buffer output ports access technique. Rather than the FIFO technique, we respect the latency constraints: we give F packets the priority to be reemitted first. D packets are reemitted second, and finally R packets. We call this technique “reemission prioritization.” Figure 10 shows the comparison between the two techniques “FIFO” and “reemission prioritization” for $n_a = 8$ and $n_c = 8$.

The difference appears only if $n_e > 12$: when the buffer consists of only a few ports, R packets with no available channel or buffer port can preempt F packets (at the second level after D packets) that are being stored into the buffer, so there will not be many buffered F packets and the technique used for the reemission will not have influence. Otherwise, when the buffer has more electrical ports, there will be fewer preempted packets and therefore more F stored packets. Delay_F decreases when F packets have the priority to be reemitted before the other packets. Taking




Fig. 10. Delay versus system load. Technique of reemission: reemission prioritization (thin curves) versus FIFO (bold curves), $n_a = 8$ and $n_c = 8$.

the example of $n_e = 40$, Delay_F decreases by $\sim 2 \mu\text{s}$, at the expense of the increase of Delay_R by $\sim 7 \mu\text{s}$. Considering the percentages of each class relative to the global traffic, the average delay of all the packets does not change.

We note that for buffer input ports also, we previously studied different access techniques such as buffer input ports *partitioning*—where a specific number of ports are exclusively dedicated to a certain service class, as if we had a buffer for each class—buffer input ports *sharing*—where in addition to the partitioning, a packet of a certain class can access the ports that are dedicated to another class or classes—and FIFO access [8]. Considering the performance obtained by each of these techniques, we opted for the FIFO buffer input ports access.

In all cases, the switching delays are in the order of magnitude of some microseconds, which satisfies even the fast packets constraint.

IV. COMPARISON BETWEEN HYBRID SWITCH AND ELECTRICAL SWITCH

In this section, we focus on comparing our hybrid switch with commercial off-the-shelf switches. We begin by comparing the switching delay of the hybrid switch to a commercial electrical switch. Then, we will focus on the reduction of O-E-O conversions achieved by the hybrid switch compared to an electrical switch, which is another indicator of reduced energy consumption. More precisely, assuming that electronic switches consume energy mostly

through O-E-O conversions, energy savings are between the $n_e/[n_a \times n_c]$ ratio that we presented earlier (assuming that the n_e electronic ports are always on) and the actual reduction in O-E-O converted packets (best case, with electronic ports capable of sleep mode when not in use).

A. Delay Comparison

In order to compare the hybrid switch performance in terms of delay with electronic switches currently existing on the market, we considered as an example the Cisco Nexus 3064PQ Switch with sixty-four 10 gigabit Ethernet ports [14]. We fixed the parameters of our hybrid switch to $n_a = 8$ and $n_c = 8$, thus to have 64 packet emitters. The last in last out (LILO) latency for the Cisco Cut-Through switch is equal to 1.34 μ s, while we consider that the latency of our switch is null for packets switched directly on the optical field since they are not buffered and they are not subject to the scheduling algorithms, memory writing, and reading operations.

The LILO latency for the Cisco Store and Forward switch is about 11 μ s for a 100 kbit packet. For a fully loaded system, our hybrid switch leads to an average delay of 1.7 μ s for buffered packets when $n_e = 10$ and 8.5 μ s when $n_e = 40$. So, the delays obtained with the hybrid switch are in the same order of magnitude or even lower than those obtained with a commercial electrical one. The latency reduction may be explained by the reduction of O-E-O conversions that depends linearly on the number of buffered packets and is quantified by the following.

B. Reduction of O-E-O Conversions

The benefits of the hybrid switch compared to a bufferless all-optical switch are quantified in terms of PLR and sustainable system load. In this part, we aim to quantify the efficiency of the hybrid switch towards an electrical switch, through the reduction of O-E-O conversions.

We present in Fig. 11 the evolution of the O-E-O conversions' reduction as a function of the system load for a degree-8 switch with eight channels per azimuth and for different values of n_e . The reduction in percentage is equal to

$$\begin{aligned} &\text{O-E-O reduction} \\ &= 100\% \times \left(1 - \frac{\text{nbr of buffered packets}}{\text{nbr of all switched packets}} \right). \end{aligned} \quad (3)$$

As expected, the curves have the same shape as those presenting the evolution of the number of buffered packets (Fig. 9). For a highly loaded system, the more electronic ports the hybrid switch has, the less important the O-E-O conversion reduction is. In this case the buffer is used often.

We note that the curves become superimposed when the system load decreases. For example, at a system load of 0.65, having 10 or more electronic ports leads to the same percentage of O-E-O reduction, which is 87%. However, having the same O-E-O diminution regardless of n_e does




Fig. 11. Reduction of O-E-O conversions versus system load ($n_a = 8$, $n_c = 8$).

not mean having the same PLRs. Indeed, for a given value of n_e , a certain proportion of packets are sent to the buffer and O-E-O conversions take place, but some D and F packets are then preempted when they are being sent to the buffer by R packets. If the switch has more electronic ports, the same number of O-E-O conversions will take place, but fewer packets will be preempted, and thus the PLRs will decrease.

The reduction of O-E-O conversions is higher than 40% even with 40 electronic ports and for a considerably loaded system ($\rho = 0.95$). At a system load of 0.75, the hybrid switch does away with more than 75% of O-E-O conversions compared to an all-electronic switch, whatever the number of buffer ports. Thus, compared to an all-electronic switch, the reduction of O-E-O conversions by the hybrid switch is important and leads to the reduction of power consumption if the electronic ports are capable of sleep mode while not in use.

V. CONCLUSION

Our investigation shows that the proposed hybrid switch is a good compromise between the all-optical bufferless switch and current commercial electrical switches. In fact, it leads to much better performance in terms of PLR and sustainable system load compared to all-optical switches and meets the requirements of the different packet classes (especially for reliable data packets) for a relatively low number of electronic ports to/from the shared buffer. In addition, compared to electrical switches, the hybrid switch reduces the switching latency to a value that is acceptable even for fast voice and interactive video packets.

To dimension the hybrid switch, we tried to find a trade-off between the performance improvements, notably in terms of PLR and sustainable load, and the energy savings. Indeed, compared to an all-electronic switch, the hybrid switch reduces the number of electronic ports from ($n_a \times n_c$) to only (n_e) ports. Furthermore, it significantly decreases the O-E-O conversions, which may reduce the

energy consumption still further if the buffer ports are capable of sleep mode when not in use. Thus, the hybrid switch is a promising solution to reduce power consumption. This is a great advantage since power consumption is becoming a very important issue due to the increasing amount of transmitted data in networks.

In future work, we will review the analytical model given in the previous analysis [7] and provide amelioration for an easier dimensioning of the hybrid switch and taking into consideration the different classes of service.

Another perspective is to consider WDM channels, which do not satisfy the interchangeability condition, but are the most widely used for multiplexing channels on a single fiber. In this case, wavelength converters would be required to improve the switch performance, mainly the PLRs [11]. However, since they are energetically costly, some studies propose to have shared converters [15]. We will investigate whether the shared buffer or the shared wavelength converters make a better compromise between the obtained performance and the energy consumption reduction. Further, we will pursue our investigation of importing the benefits of channels' interchangeability by combining the SDM and WDM techniques [17].

REFERENCES

- [1] D. C. Kilper, G. Atkinson, S. K. Korotky, S. Goyal, P. Vetter, D. Suvakovic, and O. Blume, "Power trends in communication networks," *IEEE J. Sel. Top. Quantum Electron.*, vol. 17, no. 2, pp. 275–284, Mar./Apr. 2011.
- [2] S. Lambert, W. V. Heddeghem, W. Vereecken, L. Bart, D. Colle, and M. Pickavet, "Worldwide electricity consumption of communication networks," *Opt. Express*, vol. 20, no. 26, pp. B513–B524, 2012.
- [3] Z. Rosberg, H. Le Vu, M. Zukerman, and J. White, "Performance analyses of optical burst-switching networks," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 7, pp. 1187–1197, 2003.
- [4] R. Takahashi, T. Nakahara, K. Takahata, H. Takenouchi, T. Yasui, N. Kondo, and H. Suzuki, "Ultrafast optoelectronic packet processing for asynchronous, optical-packet-switched networks [Invited]," *J. Opt. Netw.*, vol. 3, no. 12, pp. 914–930, 2004.
- [5] D. Chiaroni, R. Urata, J. Gripp, J.-E. Simsarian, G. Austin, S. Etienne, T. Segawa, Y. Pointurier, C. Simonneau, Y. Suzuki, T. Nakahara, M. Thottan, A. Adamiecki, D. Neilson, J.-C. Antona, S. Bigo, R. Takahashi, and V. Radoaca, "Demonstration of the interconnection of two optical packet rings with a hybrid optoelectronic packet router," in *European Conf. on Optical Communication (ECOC)*, Torino, Italy, Sept. 2010, postdeadline paper.
- [6] R. Takahashi, T. Nakahara, Y. Suzuki, T. Segawa, H. Ishikawa, and S. Ibrahim, "Recent progress on the hybrid optoelectronic router," in *Photonics in Switching (PS)*, Ajaccio, Corsica Island, France, Sept. 2012.
- [7] P. Delesques, T. Bonald, G. Froc, P. Ciblat, and C. Ware, "Enhancement of an optical burst switch with shared electronic buffers," in *Int. Conf. on Optical Networking Design and Modeling (ONDM)*, Brest, France, Apr. 2013, pp. 136–141, poster.
- [8] W. Samoud, C. Ware, and M. Lourdiane, "Investigation of a hybrid optical-electronic switch supporting different service classes," in *Photonics North (PN)*, Montreal, Canada, May 2014, paper 9288.
- [9] K.-I. Kitayama, Y.-C. Huang, Y. Yoshida, R. Takahashi, and M. Hayashitani, "Optical packet and path switching intra-data center network: enabling technologies and network performance with intelligent flow control," in *European Conf. on Optical Communication (ECOC)*, Cannes, France, Sept. 2014, paper Tu.1.6.1.
- [10] Z. Pan, H. Yang, J. Yang, J. Hu, Z. Zhu, J. Cao, K. Okamoto, S. Yamano, V. Akella, and S.-J. B. Yoo, "Advanced optical-label routing system supporting multicast, optical TTL, and multimedia applications," *J. Lightwave Technol.*, vol. 23, no. 10, pp. 3270–3281, Oct. 2005.
- [11] W. Samoud, M. Lourdiane, and C. Ware, "Comparison between interchangeable and WDM channels for hybrid optoelectronic switch," in *Photonics North (PN)*, Ottawa, Canada, accepted for oral presentation in June 2015.
- [12] M. Gharaei, H. Castel, and M. Lourdiane, "100 Gb/s Réseau Internet Adaptative (100GRIA) FUI9 project," Tech. Rep., Dec. 2012.
- [13] M. Justin Redd, "Calculating statistical confidence levels for error-probability estimates," *Lightwave Mag.*, vol. 21, no. 5, pp. 110–114, Apr. 2000.
- [14] Cisco, "Understanding switch latency," Tech. Rep. [Online]. Available: http://www.cisco.com/c/en/us/products/collateral/switches/nexus-3000-series-switches/white_paper_c11-661939.html#_Toc327313472.
- [15] C. Okonkwo, C. Almeida, Jr., and R.-E. Martin, "Performance analysis of an optical packet switch with shared parametric wavelength converters," *IEEE Commun. Lett.*, vol. 12, no. 8, pp. 596–598, Aug. 2008.
- [16] T. Szigeti and C. Hattingh, "Quality of service design overview," Cisco, Dec. 2004 [Online]. Available: <http://www.ciscopress.com/articles/article.asp?p=357102>
- [17] W. Samoud, M. Lourdiane, and C. Ware, "Performance improvements of hybrid opto-electronic packet switch using SDM in addition to WDM," in *European Conf. on Optical Communication (ECOC)*, Valencia, Spain, accepted for poster presentation in Sept. 2015.