Providing Differentiated Service Categories in Optical Packet Networks[1]

J. Masip-Torné[a], J. Solé-Pareta[a], S. Borgione[a], B. Bostica[b], M. Burzio[b]

[a] Universitat Politècnica de Catalunya, Jordi Girona 1-3, Mòdul D6 (Campus Nord), 08034 Barcelona, Catalunya, Spain.
[b] CSELT-Centro Studi E Laboratori Telecomunicazioni, Via Guglielmo Reiss Romoli, 274, 10148 Torino, Italy.

Abstract. Advanced network technologies are evolving towards Optical Packet Networks (OPNs) as the means of matching fast packet switching and all optical network capabilities. However, basic OPN performance is highly limited by the small optical buffers available. Optical switch performance can be enhanced by introducing differentiated services techniques in the optical buffer scheduling and by using a hybrid optical-electrical buffer architecture. This paper presents and evaluates an innovative approach for such architecture, which is able to differentiate three categories of traffic, namely non-critical, time-critical and quality-critical. Two different topologies for the hybrid optical-electrical buffer are proposed, and evaluated by simulation under severe traffic conditions.

1 Introduction

Fast packet switching nodes have been typically implemented using electronic technologies that allow the use of high dimensional buffers (thousands of packets) required by present service traffic profiles. Furthermore, switching core functions can be implemented with low cost and low power electronics (CMOS) which, in turn, suffers from limited line speed and line driving capability. Data parallelization enables high throughputs, but the derived packaging and wiring complexity limit the overall capacity of electronic fast packet switches to a few hundred Gb/s [1].

The introduction of simple optical functions within the core of switching nodes has been proposed through several studies and test beds in order to obtain higher overall switching capacities [2,3]. In most of the studies optical technology is used to

[1] This work has been supported in part by a CIRIT (Generalitat de Catalunya) grant (No. 1997FI-00725), and in part by the CICYT (Spanish Ministry of Education) under contract TIC96-1127. This work has been possible thanks to the participation in a project included in a collaboration framework between the CSELT and the UPC.

Digitized by Google

exploit its intrinsic high bandwidth capability in terms of both bit rate and wavelength multiplexing (WDM). WDM gives an additional dimension with respect to the conventional electronic solutions based on time and space; this provides the possibility of carrying out high throughput nodes with moderate complexity. However, optical buffers can only be set up at present by using arrays of fiber optic loops, which limit the actual buffering capacity to some tens of packets. According to the previous considerations, optical technology has the potentiality for realizing Tbit/s range switching nodes. However, if huge buffering is required, the use of electronic technology is still necessary.

The use of optical switches can be conceived in both All Optical Network (AON) and Fast Packet Switching environments. In AON environments, data is optically transferred and switched in the optical domain [4]. An example is given by IP over DWDM networks where optical channels provide direct transfer capabilities to IP together with flexible transport network reconfiguration. Current PDH/SDH transmission hierarchies are progressively substituted by WDM networks as a physical transport layer. A further improvement can be achieved in terms of bandwidth granularity and resource sharing capabilities introducing optical fast packet switches [5]. As a result, Optical Packet Network (OPN) environments integrate both WDM (Layer 1) and Fast Packet Switching (Layer 2) functionality in a single layer (Layer 1+Layer 2). However, the limited optical buffer depth creates severe drawbacks in terms of packet loss in the presence of bursty and unbalanced data traffic coming from its client layer (Layers 3 or 4). Since this is a major impediment towards building a backbone that can efficiently cope with IP traffic growth and QoS requirements, it is of great interest to assess the performance of QoS provisioning strategies over optical buffers.

In this paper, it is shown that optical buffers can provide non-critical and time-critical traffic categories which combine a random-early-discard policy with non-consecutive buffer delay lines and that this is not sufficient to provide a quality-critical category. Subsequently, it is shown that, by adding a small electrical buffer, a limited amount of quality-critical traffic can be supported.

This paper is organized as follows. Section 2 discusses the possible QoS strategies for OPN nodes and describes two hybrid optical-electrical buffer topologies that are subsequently evaluated by simulation. The traffic and the hybrid buffer models to be used in the simulations are presented in Section 3. In Section 4 the performance of the single optical buffer structure and of the two hybrid optical-electrical buffer topologies described in Section 2 are discussed. Finally, conclusions are drawn in Section 5.

2 Providing QoS in OPN environments

The Internet Engineering Task Force has proposed two methods of improving the unpredictable best-effort service on the Internet, namely the Integrated Services (IS) approach [6] and the Differentiated Services (DS) approach [7]. The IS objective is to provide QoS to the users according to their specific application flow requirements.

Digitized by Google

Services are provided differentiating application flows at each network node and providing the committed resources. However, the IS model has scalability limitations due to the high number of flow states that the network has to be able to manage. For its part, the DS model provides QoS to aggregated application flows both in terms of delay/jitter priority and in terms of drop priority. DS packets are marked at the network edges by setting service bits according to the traffic category to which they belong. These bits will determine how packets are treated in the network nodes. DS is intended to provide scalable service discrimination without requiring per-application flow states and signalling at each network node.

The introduction of one of the above QoS approaches in OPN environments is highly desirable because it would ensure compatibility with the QoS management of the Internet client layer and the associated scheduling mechanisms can be used to improve the raw OPN performance. Nevertheless, the IS approach is not a good choice because OPN nodes should provide QoS to large aggregated bandwidths of application flows. Thus, DS is the remaining candidate in scaling up to an OPN environment.

In this paper we address the problem of differentiating several categories of traffic at the OPN nodes. We deal with the following three categories of traffic: 1) *Non-Critical traffic* (NC) which has neither packet loss nor time delay requirements. 2) *Time-Critical traffic* (TC), which has strong time delay requirements, and moderate packet loss requirements. 3) *Quality-Critical traffic* (QC), which has strong packet loss requirements (even zero-packet loss requirements), and moderate time delay requirements.

Within OPN environments delay and jitter are not critical issues due to the short optical buffers. Then, a DS QoS policy based on introducing a traffic-drop priority can be enough for a simple and effective implementation of two service categories to differentiate TC traffic from NC traffic. That is, an *assured quality* service to support the TC traffic and a *best effort* service to deal with the NC traffic.

This traffic-drop priority scheme works as follows: if the network is not congested there is no difference in how the nodes treat each packet (e.g. no packet discard); and the packets are forwarded in the same order they are received. When congestion occurs best effort packets are discarded first, mitigating the congestion and leaving resources for the assured service packets. An example of how this strategy can be implemented is given in Figure 1.a. When a packet is received, the output port of the OPN node (and the associated queue) is determined. If the corresponding queue depth exceeds a given threshold T1 (congestion is imminent), best effort packets are discarded while assured packets are correctly sent into the queue.

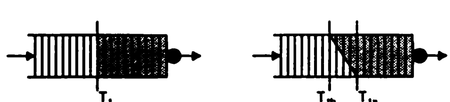

Figure 1: Output buffer with drop priority based on: (a) a threshold, (b) a RED policy.

Digitized by Google

This threshold mechanism can also be associated with a Random Early Discard (RED) strategy, providing a real advantage in the presence of optical packets carrying TCP/IP datagrams (Figure 1.b). When the buffer occupancy reaches the threshold T1a, the discard probability for a best effort packet begins to be more than zero and its value is increased to 1 when the buffer occupancy reaches the T1b value. In excess of this value only the assured quality packets can be buffered and all the best effort packets are discarded. It is worth noticing that in both cases (plain drop priority and RED) packets are discarded at the input of the OPN node: once packets have been entered the node, they undergo no scheduling modifications. This fits with the particular structure of optical buffers at OPN nodes, in which deleting a packet that has already entered the buffer is difficult and inadvisable.

2.1 Enhanced QoS policy for OPN environments

The model described above leads to a QoS scheme in which non-critical and time-critical traffic categories can be supported. Nevertheless, a real quality critical traffic category with zero packet loss ratio (PLR) or so (PLR less than 10^{-8}) is not supported. The QC traffic category is desirable because TCP throughput performance is highly degraded in the presence of packet loss, although the reserved bandwidth may be high. In this case, larger buffers will be required to provide a zero packet loss performance because the OTP the network does not include a congestion control mechanism to push congestion to the network edges [8]. As has been mentioned above, large buffers can only be implemented electronically.

We propose a hybrid optical-electrical output buffer architecture to support the NC, TC and QC traffic categories. Figure 2 shows the general structure of such architecture. It is composed of an optical switching core, an optical buffer devoted to NC and TC traffic, and an electrical buffer devoted to QC traffic. The RED policy is implemented at the optical buffer input to allow service differentiation between NC and TC traffic. The electrical buffer is served through the optical buffer. If the electrical buffer is directly served to the output of the optical buffer then its service probability decreases as the optical buffer load grows. Under extreme traffic load conditions, the electrical buffer would be never served and the electrical buffer size required for zero packet loss would tend to infinite.

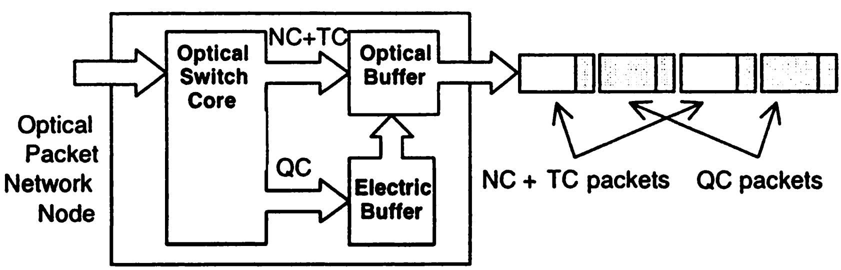

Figure 2: Hybrid optical-electrical buffer architecture.

Digitized by Google

Since the optical buffer is implemented as a set of fiber delay lines of different lengths, once a packet has entered in a fiber delay line of the optical buffer, it is both difficult and inadvisable to delete it and the packet output time cannot be modified at all.

Achieving low electrical buffer requirements to support a QC traffic category is crucial because very high-speed memories are involved and the switch complexity has to be kept under reasonable bounds. Consequently, the electrical technology used to implement the electrical buffer will establish the maximum buffer depth and this will bound the amount of QC traffic that an OPN are able to handle for a certain electrical buffer scheduling mechanism. Below, two different topologies for the hybrid optical-electrical buffer architecture are described, namely the Single Delay Line Topology, and the Double Delay Line Topology.

- ***The Single Delay Line Topology*** (SDLT) consists of an electrical buffer directly connected to the optical buffer through one of its delay lines (hereafter, the delay line X) as shown in Figure 3. In this topology, the electrical buffer is served when there is no optical packet at the X delay line input (NC and TC packets have priority over QC packets), and there are no packets in other delay lines with the same output time as the X delay line. Therefore, the service probability of the electrical buffer depends on the X delay line length, because packets inserted in a higher optical buffer position cannot be deleted and they will go down the buffer, reducing the service probability of the electrical buffer. In fact, a long X delay line will give a strong priority to QC traffic to the detriment of NC+TC traffic, while a short X delay line will give a weak priority to QC traffic to the benefit of NC+TC traffic.

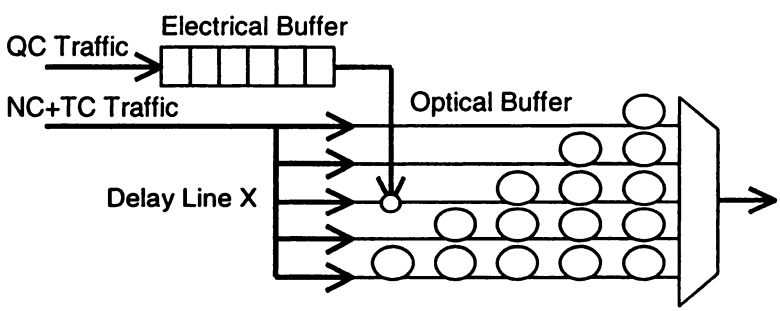

Figure 3: Single Delay Line Topology (SDLT).

- ***The Double Delay Line Topology*** (DDLT) combines the optical and the electrical buffers following the DQLT (Dual Queue Length Threshold) scheduling policy proposed for ATM multiplexers [9]. This scheduling policy uses a small buffer to serve real-time traffic (optical buffer for NC+TC traffic in our case), and a large buffer to serve non-real-time traffic (electrical buffer for QC traffic in our case). The DQLT policy achieves reasonable performance with low control complexity and it provides a threshold mechanism that is used to control the non-real-time

Digitized by Google

buffer load in the presence of severe congestion. The DQLT policy adapted to the hybrid buffer structure works as shown in Figure 4. When the electrical buffer occupancy is higher than the threshold T, the electrical buffer is served through the longest (high-priority) delay line (case B > T in Figure 4), otherwise the electrical buffer is served through the shortest (low-priority) delay line (case B < T in Figure 4).

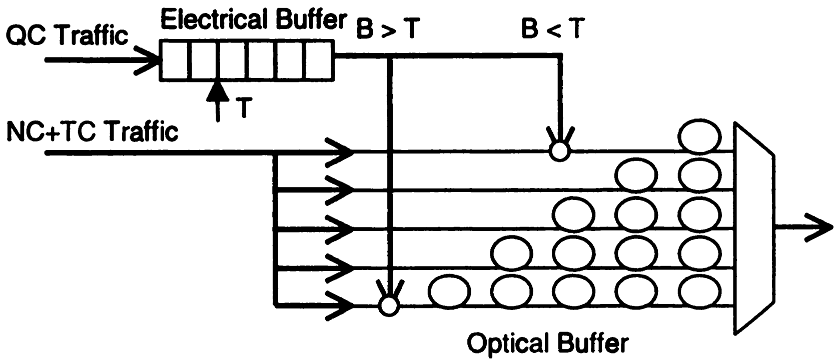

Figure 4: Double Delay Line Topology (DDLT).

SDLT and DDLT performance will depend strongly on its respective configuration parameters. The length of the delay line X where the electrical buffer is connected will determine the performance of the SDLT topology. This is a hardware parameter (it cannot be changed on-line). In contrast, the DDLT performance will be determined by the threshold value, which is a software parameter. In a wide sense, DDLT can be seen as the superposition of two SDLTs, and the threshold value controls their relative behaviour, thus reaching the average performance of a SDLT where the delay line to which the electrical buffer is served could be dynamically modified.

From all the above, we assume that SDLT and DDLT are good candidates for supporting NC, TC and QC traffic categories better than using a single optical buffer, specially under severe traffic conditions, and with small electrical buffer requirements. Below, we present the simulation environment and discuss the results of a set of experiments we carried out to verify this assumption.

3 Simulation environment

The simulation environment is shown in Figure 5. It is composed of a traffic generator and an output buffer. The models used in this simulation environment for the traffic generator and for the buffer are described in the following sub-sections.

Digitized by Google

3.1 The traffic generator model

The traffic generator can be set either to a single Poisson source, or to a set of N On-Off traffic sources. The Poisson model is used as the equivalent of a set of N Bernoulli traffic sources when N is large. On-Off traffic is used to test the performance of the hybrid buffer under unbalanced traffic conditions closer to real traffic behavior. Each On-Off traffic source generates a constant packet rate during burst periods. Both the burst length and the idle time periods are exponentially distributed. Burstiness is defined as the ratio of the burst rate and the average source rate. Irrespective of the traffic source type, packets are labeled according to a fixed rate of NC, TC and QC traffic categories. Packet labeling is performed randomly because it is assumed that the different traffic categories are highly multiplexed in the backbone nodes.

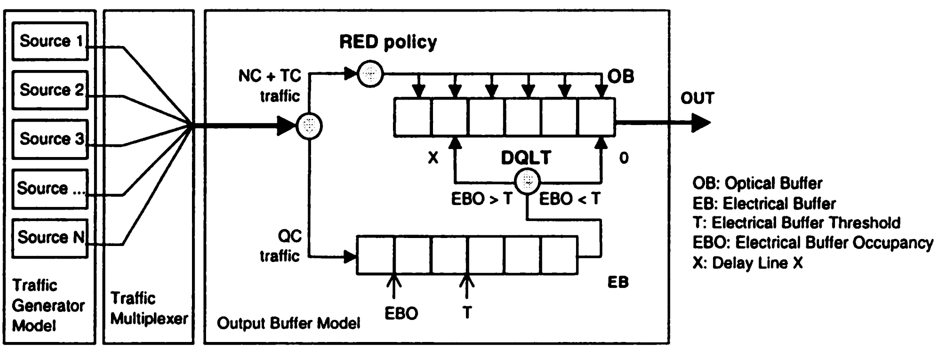

Figure 5: Simulation environment.

3.2 The output buffer model

The output buffer model is composed of an optical buffer model (OB), a FIFO electrical buffer model (EB) and three filters (the gray bullets), as shown in Figure 5. The filter at the hybrid buffer input separates QC and NC+TC traffic. Then the optical buffer accepts packets according to the RED policy filter implemented at the optical buffer input. Finally, the electrical buffer is served through two service delay lines (0 and X) depending on the electrical buffer occupancy (EBO) with respect to the threshold value (T). As can be seen in Figure 5, the DQLT filter selects the appropriate service delay line at the electrical buffer output.

3.3 The optical buffer model

Optical buffers have a highly limited capacity (tens of packets, as has been already mentioned) when compared with electrical buffers. In contrast, they have an

Digitized by Google

important implementation property, which is not available in the case of electrical buffers. Optical buffers consist of an array of delay transmission lines, and can be easily loaded in parallel, while the electrical buffers can only be serially loaded. Due to this fact, optical buffers can be modeled as packet shift registers, in which the input packets are loaded in parallel on the smallest delay positions that are available (empty) [10]. Figure 6.a shows an optical buffer with capacity for 5 optical packets, and the optical model using a 5-position packet-shift register. Note that the number of parallel accesses determines the complexity of the shift register and not the longest delay line. According to this model, optical buffers with the same complexity can have different access distributions, i.e., all the positions (consecutive delay lines) or for only some of them (non-consecutive delay lines). Figure 6.b shows three possible configurations of an optical buffer, the consecutive delay lines configuration, and two non-consecutive delay lines configurations: alternate accesses and accesses with increasing separation.

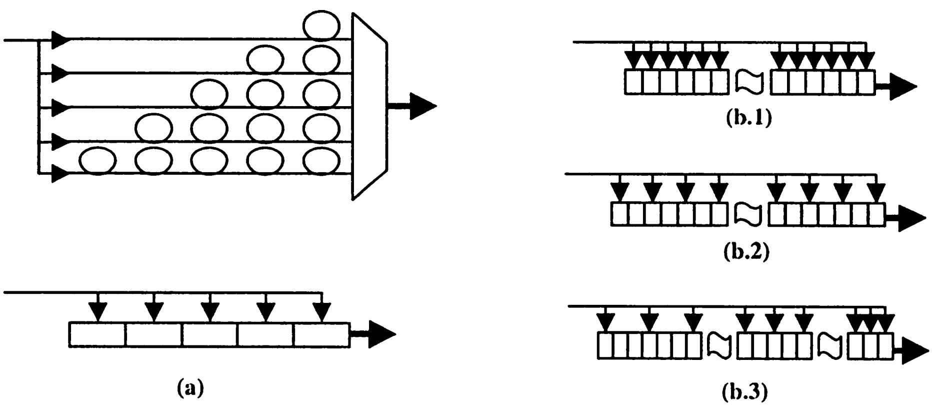

Figure 6: (a) Shift register model for a delay lines optical buffer. (b) Three possible configurations for the optical buffer shift register model.

The performance of a particular case of optical buffer with non-consecutive delay lines was first analyzed in [10] under conditions of regular and balanced traffic. It was demonstrated that non-consecutive delay lines optical buffers can have better performance in terms of PLR than consecutive-delay-line optical buffers. In contrast, non-consecutive-delay-lines optical buffers are not able to maintain the packet sequence integrity.

4 Experiments and results

Using the simulation environment described above, we ran a set of three experiments. The aim of these experiments was to evaluate and compare the performance of the

Digitized by Google

single optical buffer architecture and the hybrid optical-electrical buffer architecture with an SDLT topology and with a DDLT topology. Except where stated otherwise, the experiments were carried out with the following common conditions.

- *Percentage of traffic categories:* 40% for the NC traffic, 30% for the TC traffic and 30% for the QC traffic.
- *On-Off source parameters:* the burstiness (peak rate/average rate) was set to b = 10, and the mean burst length (BL) was set either to BL = 10 packet times or to BL = 30 packet times. Since OPN multiplexes a large number of flows coming from its client layer, an aggregated flows with b = 10 and BL = 30 can be considered as a very high bursty traffic. The number of On-Off traffic sources in the experiments is N = 16.
- *Buffer configuration:* a FIFO infinite queue for the electrical buffer, and an optical buffer composed of 30 delay lines of independent length. In the case of the optical buffer, two structures were considered, the consecutive delay lines structure and a non-consecutive delay lines structure. As shown in Figure 7.a, the consecutive delay lines structure has an incremental step of one packet time for all the delay lines. Figure 7.b shows the non consecutive delay lines structure used in our experiments, note that it consists of groups of ten delay lines which have incremental steps of one, two and three packet times respectively. This leads to the longest delay line to have a length of 1x10+2x10+3x10 = 60 packet times.

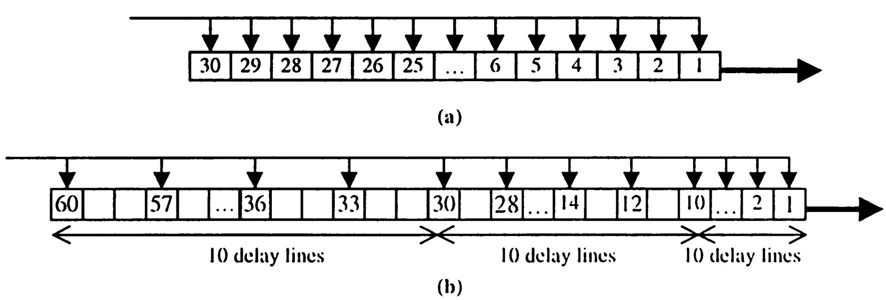

Figure 7: (a) Consecutive delay lines structure. (b) Non consecutive delay lines structure.

- *RED policy applied at the optical buffer:* the optical equivalent of the electrical buffer occupancy is the number of delay lines that cannot be accessed. Thus, the random discarding area, i.e. the values of the parameters T1a and T1b in Figure 2.a, were set to T1a = 10 and T1b = 20. Note that the RED policy only affects the NC traffic, i. e., the packets of the TC traffic flow are only discarded because of buffer overflow.

Digitized by Google

4.1 Experiment I: Optical buffer performance

In this experiment we tested the performance of three different configurations of the single buffer architecture, namely the consecutive delay lines structure and the non-consecutive delay lines structure without applying any drop priority policy, and the non-consecutive delay lines structure applying the above described RED policy. Figures 8, 9 and 10 depict the results obtained in the set of simulations of this experiment.

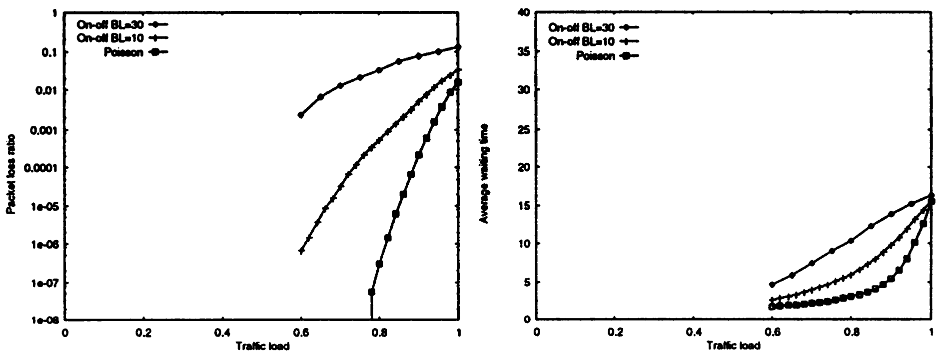

Figure 8: Single optical buffer performance: consecutive delay lines structure.

Figure 8 shows that, as expected, the packet loss ratio (PLR) highly degrades in presence of bursty traffic (On-Off with BL = 10 and BL = 30). Within the moderate traffic load conditions region (ranging from ρ = 0.6 to ρ = 0.8), the performance for the TC and QC traffic categories is only acceptable for non-bursty traffic (Poisson). Also as we expected the average waiting time in the buffer is very low, even under heavy traffic conditions (from ρ = 0.8 to ρ = 1).

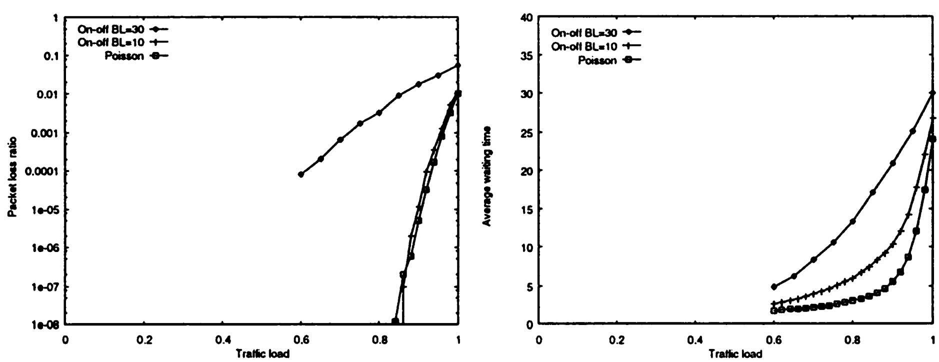

Figure 9: Single optical buffer performance: non-consecutive delay lines structure.

Digitized by Google

Figure 9 shows that under heavy load traffic conditions, the non-consecutive delay lines based solution do not exhibit better performance than the consecutive delay lines structure. From $\rho = 0.8$ to $\rho = 1$, the resulting PLR is extremely high for the TC and QC traffic requirements. In contrast, under the moderate load traffic (from $\rho = 0.6$ to $\rho = 0.8$), low bursty traffic gives an acceptable performance for the TC and QC traffic categories. The performance is same for low bursty traffic (On-Off with BL = 10) than for non-bursty (Poisson). However, in this case, packet sequence integrity is not maintained.

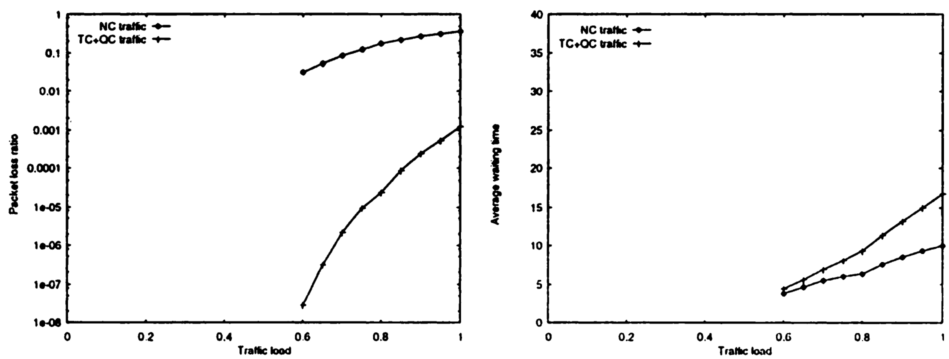

Figure 10: Single optical buffer performance: non-consecutive delay lines structure with RED policy applied.

Figure 10 corresponds to the non-consecutive delay lines buffer configuration when applying the two levels described in Figure 2.b. In this case only the performance for the most bursty traffic (On-Off with BL = 30) was tested. Figure 10 shows two clearly differentiated categories of services, namely best effort (for NC traffic) and assured quality (for TC and QC traffic together). The only acceptable PLR is that provided by the assured quality service under moderate load conditions, which could be acceptable for TC traffic, but not for QC traffic. In this case, even the NC traffic has unacceptable performance. Complementary simulations applying a three level RED policy shown that the differentiation of QC and TC traffic can only be achieved in detriment of the TC traffic performance and that no better performance for NC traffic was obtained. The average waiting times continue at very low levels.

As a general conclusion, from the experiment I it has to be pointed out that with a single optical buffer architecture it is difficult to differentiate the TC and QC traffic categories. Furthermore, a proper PLR for these categories of traffic can only be obtained in detriment of the NC traffic category.

Digitized by Google

4.2 Experiment II: SDLT hybrid buffer performance

In the Experiment II the SDLT topology for the hybrid optical-electrical architecture was tested. In the optical buffer a two level RED policy was applied, the electrical buffer size was set to infinite and, as in the case of the non-consecutive delay lines optical buffer configuration, the simulations were ran only for the most bursty traffic conditions (On-Off sources with BL = 30).

The obtained results have been summarized in Figure 11, and are given in terms of the delay line though which the electrical buffer was served (we scanned the performance from the delay line X = 0 to the delay line X = 60).

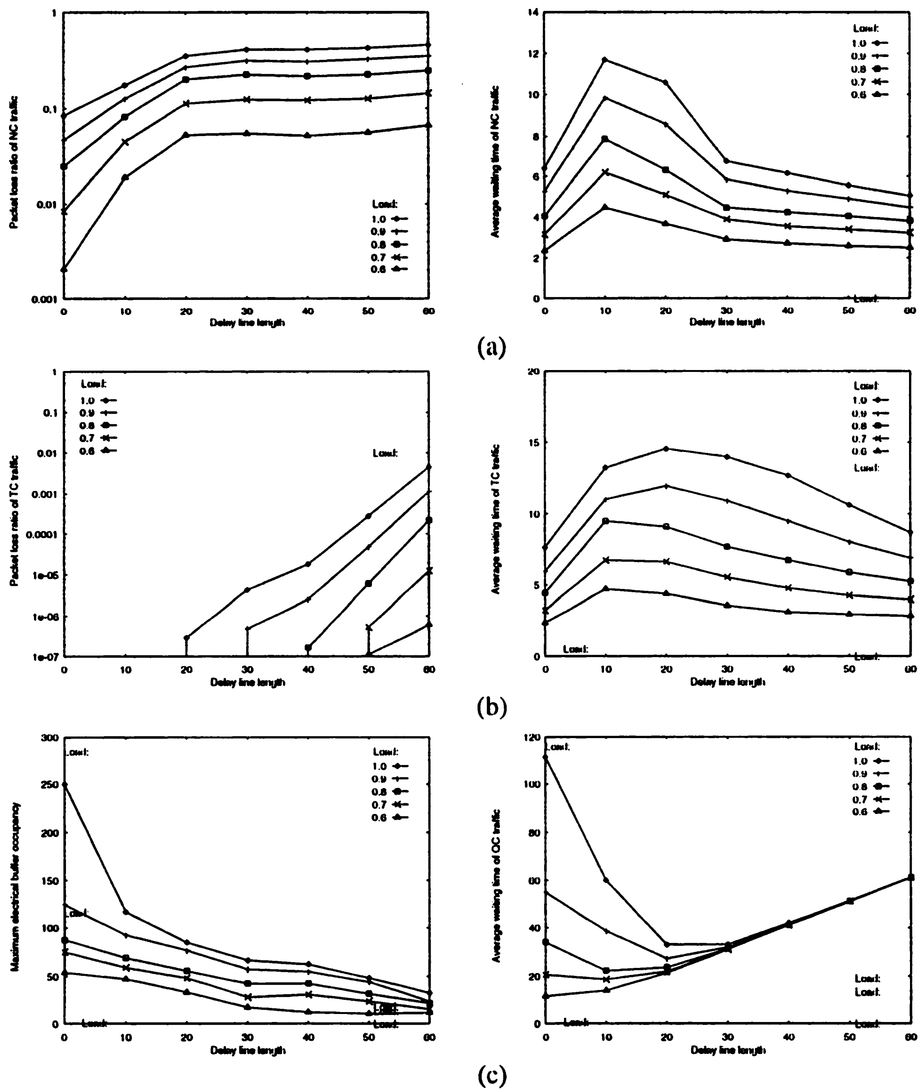

Figure 11: Hybrid SDLT buffer performance: (a) NC traffic, (b) TC traffic, (c) QC traffic.

Digitized by Google

As expected, Figure 11 show that the best performance in terms of PLR is obtained when the electrical buffer is served through the shortest delay lines. Between delay lines of delay 0 and 20 packet times, the PLR is acceptable for all types of traffic. Consider the case of serving the electrical buffer through the delay line X = 0 and extreme load traffic conditions (ρ = 1). In that case, NC traffic exhibits a PLR behind the 10% (see Figure 11.a), while in the case of a single optical buffer reached the 50% (see Figure 10). TC traffic has an acceptable PLR within the entire range of traffic load (see Figure 11.b), and no QC packets were lost during the simulation time, that means a PLR lower than 10^{-8}. Note that to achieve such a low PLR for the QC traffic, it is required a minimum electrical buffer size of 250 packets (see Figure 11.c). In what concern to the average waiting time, obviously grows for the QC traffic category, but it can be acceptable even for severe traffic load conditions (Fig. 11.c).

4.3 Experiment III: DDLT hybrid buffer performance

In the Experiment III the DDLT topology for the hybrid optical-electrical architecture was tested under the same conditions than the DDLT topology. The obtained results have been summarised in Figure 12, and are given in terms of the value of the threshold T, in a logarithmic scale (we scanned the performance from T = 1 to T = 300).

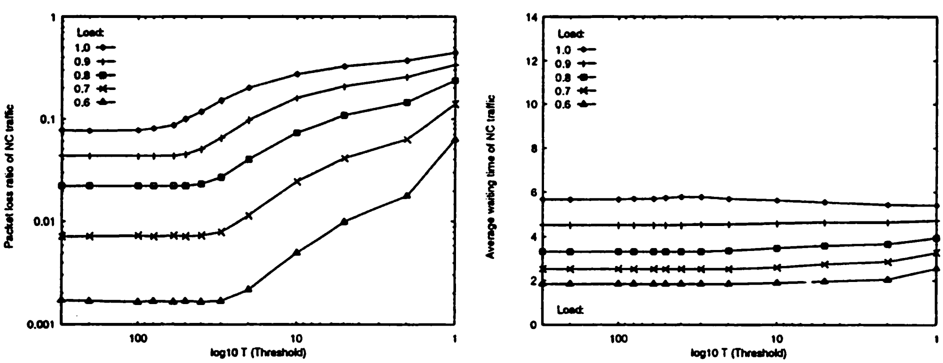

(a)

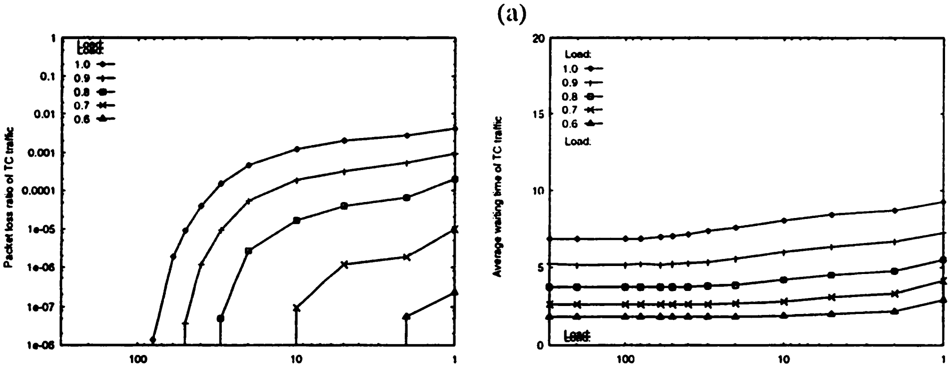

(b)

Digitized by Google

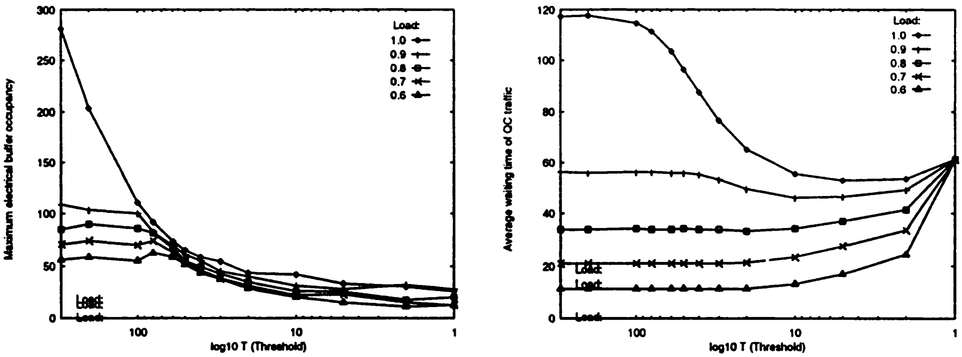

(c)

Figure 12: Hybrid DDLT buffer performance: (a) NC traffic, (b) TC traffic, (c) QC traffic.

From Figure 12 it can be seen that DDLT topology has similar performance than the SDLT topology. Actually DDLT topology exhibits a slight performance improvement compared with that of the SDLT topology, but what we want to emphasize is the fact that in the DDLT topology the tunable parameter (T) can be set up by software, either manually or dynamically. This gives to this topology a flexibility that can be very useful in the real case where the percentages of the different traffic are not fix like we considered in our simulations. Note, however, that the DDLT topology do not preserve the packet sequence integrity for any of the three traffic categories, nor for the QC traffic, which can be preserved by using the SDLT topology.

5 Conclusions

In this paper we validated that NC and TC traffic categories can be supported, even under severe traffic load conditions, by a non-consecutive delay lines optical buffer where a RED policy is applied and by admitting the end to end recovery of the packet sequence integrity. However, single optical buffers were not sufficient to provide the zero packet loss requirements of the QC traffic category. Thus, the introduction of a hybrid optical-electrical buffer architecture was proposed and its performance evaluated. Two hybrid optical-electrical buffer topologies were considered, namely the SDLT topology and the DDLT topology.

The simulation results shown that both SDLT and DDLT topologies provide acceptable QoS to NC, TC and QC traffic categories even under severe traffic load conditions. For the SDLT, the shorter the delay line of the optical buffer through which the electrical buffer is served, the better is the performance. Nevertheless, serving the electrical buffer through shorter delay lines, implies a significant increase of the average waiting time of the QC packets. The DDLT topology allows to achieving similar and even better performance than the SDLT topology by properly

Digitized by Google

setting up the value of the electrical buffer threshold T. For the DDLT, the highest the electrical buffer threshold value, the better is the performance.

The DDLT topology has the advantage that its configuration parameter (the electrical buffer threshold) can be dynamically modified. It is expected that the DDLT topology will be able to cope with variable traffic conditions in terms of the percentages of the different traffic categories by introducing an appropriate control algorithm. Nevertheless, the DDLT topology does not preserve the packet sequence integrity for any of the three traffic categories.

The DDLT topology behavior under variable traffic conditions in terms of the percentages of the different traffic categories, and the impact on the performance of recovering the packet sequence integrity has been let for further studies.

References

[1] J. Y. Hui and K. Cheung, "Optical Versus Electronic Switching for Broadband Netwoks", IEEE Network, November/December 1996.

[2] J. Masip-Torné, J. Sole-Pareta and G. Junyent-Giralt, "Electronic and Photonic Switching in the ATM Era", IEE International Broadcasting Convention 1997, Conference Publication, 1997.

[3] B. Bostica, P.Gambini and L.Licciardi, "The role of electronics and photonics in the architectural evolution of switching systems", ISS'97-XVI World Telecom Congress Proceedings, pp.565-571.

[4] P. E. Green et al., "Optical Networking Update", IEEE Journal on Selected Areas in Communications, June 1996.

[5] M. Renaud, F. Masetti, C. Guillemot and B.Bostica, "Network and System Concepts for Optical Packet Switching", IEEE Communications Magazine, April 1997.

[6] P.P. White, "RSVP and Integrated Services in the Internet: A Tutorial", IEEE Communications Magazine, May 1997.

[7] K. Nichols et al., "Differentiated Services Operational Model and Definitions", IETF Internet-Draft: draft-nichols-dsopdef-00.txt, February 1998.

[8] S. Kalyanaraman, R. Jain, S. Fahmy and R. Goyal, "Performance and Buffering Requirements of Internet Protocols over ATM and UBR Services", IEEE Communications Magazine, June 1998.

[9] J.S. Jang, S.H. Shim and B.C. Shin, "Analysis of DQLT Scheduling Policy for an ATM Multiplexer", IEEE Communications Letters, November 1997.

[10] P. Cadro, A. Gravey, C. Guillemot, "Performance Evaluation of an Optical Transparent Packet Switch", ITC 15, Elsevier, 1997.

Digitized by Google