On the use of prediction in Passive Optical LANs for healthcare latency-stringent applications

Georgios A. Tziroglou
Aristotle University of Thessaloniki
Department of Informatics
Thessaloniki, Greece
tziroglou@csd.auth.gr

Anastasios Valkanis
Aristotle University of Thessaloniki
Department of Informatics
Thessaloniki, Greece
tasosvalk@gmail.com

Constantine A.Kyriakopoulos
Aristotle University of Thessaloniki
Department of Informatics
Thessaloniki, Greece
kyriak@csd.auth.gr

Georgios I. Papadimitriou
Aristotle University of Thessaloniki
Department of Informatics
Thessaloniki, Greece
gp@csd.auth.gr

Petros Nicopolitidis
Aristotle University of Thessaloniki
Department of Informatics
Thessaloniki, Greece
petros@csd.auth.gr

***Abstract*— Passive Optical Networks (PONs) have been established as the most successful, viable and economic solution for the access part of an optical network. During the last two decades, a great proportion of the academic research in networking has focused on improving the bandwidth allocation algorithms for PONs so as to provide better user services. Recently, PONs have been suggested as the means for providing ultra-low latency for QoS stringent applications, such as the Tactile Internet, for healthcare applications in a hospital environment. Despite their potential, little attention has been given so far to prediction methods for allocating bandwidth in advance so as to decrease end-to-end latency in such environments. This paper aims to address this gap, by proposing a novel framework for Passive Optical LANs. By utilizing a prediction scheme and a Dynamic Bandwidth Allocation (DBA) algorithm that allocates bandwidth based on network feedback information, the proposes approach is shown to significantly contribute to reduction of latency.**

Keywords—Passive Optical Networks, Prediction, Healthcare applications, Knapsack Problem, Dynamic Bandwidth Allocation, Low latency.

I. Introduction

The amount of traffic needed to be handled significantly increases each year and estimates for and forthcoming years show that this trend will continue to hold. Cisco [1] has presented monthly traffic estimations per year, revealing that between the years of 2010-2014, Internet traffic statistics have shown exponential increase. Furthermore, not only does worldwide flowing traffic grow, but also new applications with special Quality of Service (QoS) requirements arise. For instance, Tactile Internet (TI) [2] is the new vision for future, allowing humans and machines to interact in real time. This new era of applications however, sets stringent requirements in data transfer, such as ultra-low latency, high availability and reliability, integration with existing technologies and security. 5G and Internet of Things (IoT) have accelerated advancements on the above areas and it is of crucial importance that optical networking will be able to facilitate such demands.

Research activity brings to the surface a large number of applications, which could be included in the group of tactile applications. In the past years, robotic surgery has become very popular in sensitive medical procedures, but physical appearance of a doctor is essential. The future is expected to bring remote surgery to life via use of haptic devices and augmented reality on the side of the surgeon and high precision machines on the side of the patient. In order to achieve such an application, an ultra-low latency networking infrastructure has to be available, so that surgical procedures carried out remotely can be implemented in real-time. Otherwise, a phenomenon called cyber sickness [3] may occur, in which the time lag between an input and its expected output is very large and the motions performed by the communicating endpoints are not synchronized. A 1ms latency is required to avoid this, resulting in a maximum of 500μs round-trip time for the network [4].

The presence of networking infrastructures in medial environments, such as hospitals, will also benefit mobile, networked Triage systems. Triage [5] is a prioritization set of methods for patients' classification, according to the severity of their status. They are mainly utilized in Emergency Medical Services (EMS) in order to assess possible risks and make accurate decisions concerning further treatment of waiting patients. Vital signs, clinical appearance and feedback from patients are monitored and transmitted by wearable devices which are inputs for the triage algorithms. However, the status of a patient may deteriorate during his sojourn in EMS, although his appearance may not reveal the ongoing situation. Consequently, medical staff may not acquire crucial information in time, resulting in dangerous situations. Frykberg [6] stresses the need for accurate Triage systems by observing a strong relation between misjudgment of critical events and mortality rates. For example, in a terrorist event in Bologna in 1980, 73.5% of the incoming events were categorized as critically injured when only 22% really needed that kind of care.

What is envisioned is an optical network capable of facilitating a machine-assisted Triage, which can automate part of the procedure. Every patient entering the EMS will wear a special mobile/wearable device, able to obtain several

978-1-5386-9301-8/19/$31.00 ©2019 IEEE

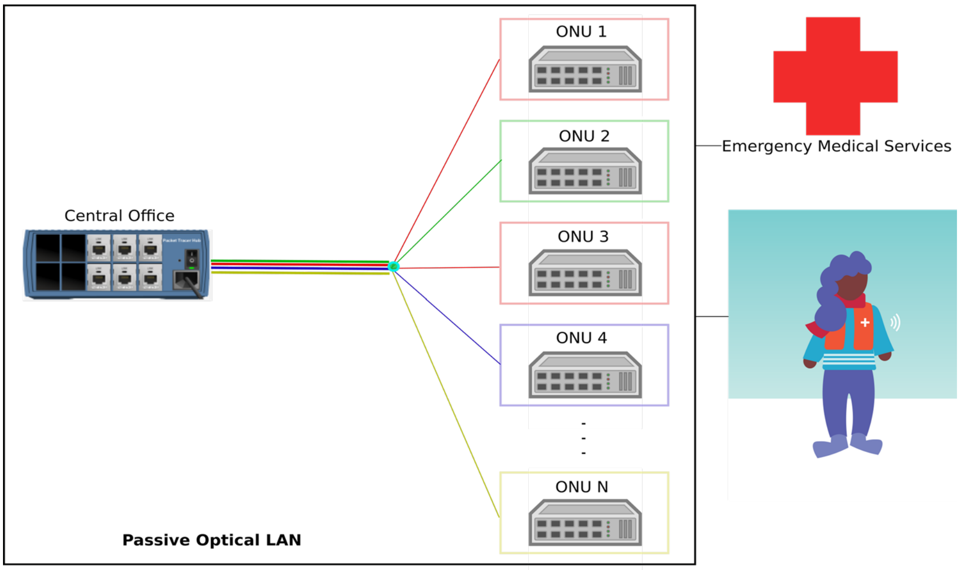

Fig. 1 Passive Optical LAN delivering healthcare services

measurements, such as heart rate, breathing rate, oxygen saturation etc. Periodically, these can be sent over a wireless-optical infrastructure in a centralized system, which will be responsible for prioritizing every patient, generating alarms, informing the medical staff or even steering and controlling remote devices. Consequently, a networking infrastructure able to provide low latencies is necessary for providing such services in a hospital environment.

Passive Optical Networks (PONs) have long been proposed as the viable, scalable and cheap solution for optical, access networks. Telecommunication vendors have long ago started to install such infrastructure, fulfiling the very old promise of Fiber to the Premises/Home. Therefore, new, broadband services can be delievered, such as live-streaming, video on demand, 4K etc. The main question that arises is why this promising tehnology has not yet been used in small, independent networks that require its beneficial properties.

Older PON standards used a TDM scheme to provide transmission opportunities for clients. Furthermore, their maximum speed was limited to 1 Gbps or less, shared by any unit wishing to transmit. In 2015, ITU-T presented its G.989 recommendation [7] of NG-PON2, which utilized both Wavelength and Time Division Multiplexing (TWDM), allowing total throughput for upstream and downstream directions of up to 40Gbps. By employing a smart DBA algorithm [8], [9] the network manager would be able to assign in a more efficient and dynamic process for usage with the available resources, enabling PONs to facilitate even the strictest demands. The need for highspeed, local and QoS aware [10] PONs could now become a reality.

In the work of Hu [11], an Integer Linear Problem is formulated to assign both wavelengths and time slots amongst the Optical Networking Units (ONUs), where the clients reside. However, as the problem is NP-hard, two heuristics are proposed and tested, but latency constraints are not considered. In [12] and [13] multiwavelength, concurrent transmissions are considered and DBA algorithms are formulated to synchronize upstream allocations. Although, no QoS is taken into consideration, the results are very promising, since for low to medium loads, the delay is below 500μs. Nevertheless, installing such an infrastructure for a single organization, for instance a hospital, would not be economically viable. Neaime and Dhaini [14] introduce a DWBA scheme that addresses the challenging nature of Tactile Internet with simultaneous transmissions over multiple wavelengths. For network loads up to 24 Gbps, QoS is proven to be preserved with their extensive simulations. However, this solution is very complex for the purpose it is needed.

A promising solution to cope with the stringent end-to-end delays was proposed by Wong et al. [15]. In their work, prediction mechanisms are introduced to estimate the total expected traffic, while bandwidth and time allocation is performed in advance. In this fashion, average packet delay is reduced to approximately half a polling cycle, in contrast with classic PON packets which have to wait one and a half. Their DBA algorithm firstly designates one wavelength for all ONUs, while it transfers the furthest to the next one if bandwidth capacity does not suffice. Additionally, wavelength assignment is tightly connected to the distances of ONUs, although their simulations have shown that distance is not an important factor in end-to-end delay.

In this work, an efficient algorithm is proposed for bandwidth and wavelength allocation based on the very well known knapsack 0/1 problem [16]. Bandwidth predictions are quantized and transformed into an appropriate set of values, since the knapsack problem is NP-hard. Afterwards, the knapsack algorithm is performed for each wavelength, until all wavelengths or unassigned ONUs are exhausted.

Until 2014, Tactile Internet was not really known and was out of the scope of research activity. Resource allocation algorithms did consider Quality of Service as a basic parameter, but did not need to achieve such stringent constraints. This framework proposes a novel DWBA (Dynamic Wavelength and Bandwidth Allocation) algorithm in conjuction with traffic prediction information collected by all ONUs. So far, low latency was not in top priority [11]-[13], or was achieved by exploiting complex, but more expensive, functionalities of PONs [14]. "Knapsacking" is an improvement of [15], which distributes resources based on traffic volume. In addition, each ONU can transmit on single wavelengths per cycle, reducing the complexity of the system. Consequently, only the fundamental equipment of an NG-PON2 is required, and the need for extra hardware , e.g.. lasers or transmitters; therefore minimizing the cost. Thus, a small or medium healthcare enterprise could enchance the supplied services.

The rest of the paper is structured as follows. Section II, presents our proposed framework and our DBA algorithm. Section III, discusses the simulation results and section IV concludes our work.

II. Proposed Framework

TWDM optical access networks are the promising means to deliver services for stringent QoS data traffic (Figure 1). Correct estimation of forthcoming packets is necessary to perform efficiently; thus decreasing the end-to-end delay. Each ONU is given a designated wavelength and a time window in which it can transmit its buffered data. At the end of its transmission, it will send a REPORT message, containing helpful information that will enable the Central Office to predict future traffic and assign appropriate resources for each ONU.

To fulfill the time constraint of 500μs delay for real-time applications, a polling cycle of

$$T_{poll} = 0.95 \cdot 2\left(L_{cons} - T_{proc} - T_{rtt\max}\right), \quad (1)$$

is selected for all wavelengths. L_{cons} is set to 500μs, T_{proc} is the processing time needed for each packet and is set to 1μs [15], while $T_{rtt\max}$ is the round trip time of the furthest ONU.

Algorithm 1: DWBA algorithm performed by the Central Office

Inputs: Wavelength capacity, predicted traffic per ONU, number of ONUs

Outputs: Scheduling timetable per Wavelength

1. Wavelength_Capacity=100 (10Gbps case);
2. unassigned_Onus=numberOfOnus;
3. i=total_Number_of_Wavelengths-1
4. While (1){
5. Wait for all Reports to arrive
6. Transform the predictions according to the formula $\text{int}\left(prediction \cdot 10^{-8}\right)+1$
7. While(i>0 and unassignedOnus>0){
8. solveKnapsack(Wavelength_Capacity, transformed_predictions,unassigned_Onus)
   Update unassigned_Onus, i=i-1
   Remove elements corresponding to assigned Onus}
9. }
10. If(unassigned_Onus>0){
    Assign them to the last wavelength}
11. Create ONU scheduling slicing T_{poll} proportionally
12. **Return schedule**

A guardband factor multiplier of 0.95 is used in (1), in order to provide a safety margin and sustain the ultra-low latency for various loads of incoming traffic.

Between transmissions of upstream traffic, each ONU is responsible for maintaining metrics related to its buffered traffic. In our simulation two types of traffic were considered, Tactile Internet (TI) and Non Tactile Internet (NonTI). As our concern was mainly on the effectiveness of our DBA, TI was considered to follow a Poisson process and NonTI a periodic one. According to [15], human to human networks can be described by the aforementioned process, while sensing devices scenarios, such as the afore-mentioned ones, create a periodic traffic. When the REPORT message needs to be sent at the end of the transmission window, the ONU piggybacks this information.

A DWBA algorithm is performed once all REPORT messages are received by the CO. Special care has been taken, so that these messages will all arrive at the CO before the polling cycle ends. Then, the CO possesses all available information so as to calculate the proportion of bandwidth for each ONU for the forthcoming polling cycle. After that, and whenever a REPORT message arrives, the OLT will generate a GATE message, informing the respective ONU for the starting time and the allowed bandwidth allocated for it, for the next cycle.

The core of our approach is the algorithm that performs the bandwidth assignments. Inspired by algorithm design [11], knapsack 0/1 dynamic programming algorithm seems to fit best for our problem. What this problem states is that if a bag with maximum capacity *W* (in our case link capacity of 10 Gbps) and a set of valuable objects (ONUs) are possessed, each one with weight w_i (which in our case corresponds to estimated traffic) and value v_i (which in our case corresponds a priority and are all set to 1, since all ONUs are of equal importance), find a subset of these that fit the bag and concurrently maximizes the total value of it. In a more formalistic way the problem states the following:

$$\max_{x_i} \sum_{objects} x_i v_i \ , \quad (2)$$

$$\text{subject to} \sum_{objects} x_i w_i \le W \quad (3)$$

with $x_i = 1$, if ONU i is included and 0, otherwise.

Two implications should be also considered. Firstly, the optimal solution can be attained if the multi-knapsack is solved, which considers all knapsacks at once (in our case the wavelengths), but has been proven to be NP-Hard [17]. Instead, each knapsack is computed at a time and objects (unassigned) ONUs that were not chosen are left for the next wavelength. Secondly, an efficient, linear complexity algorithm exists if *W* is integer. In order to tackle with this, the predicted traffic is quantized in steps of 100 Mbits and multiply the results by 10^{-8}. As quantizing may result to imprecisions, number one is added to the result. Thus, the total amount of the wavelength capacity is 100, while the smallest amount of predicted traffic is 1.

The DWBA (Algorithm 1) firstly transforms the predictions, according to the methodology presented in the previous paragraph. Then, it performs sequentially a knapsack solution for each knapsack, until all the knapsacks but the last is totally filled or the number of unassigned ONUs is zero. If all wavelengths, except the last, are full and there exist ONUs remaining to be served, these are assigned to the last wavelength. Finally, for each wavelength, the total amount of traffic is calculated by the untransformed predictions of each ONU, a proportion of which is assigned according to its needs. Algorithm 1 presents a compact description of our scheme.

The time complexity of this algorithm is mainly dominated by the solution of the knapsack decision problem, which is proven to be:

$$O\left(Wavelength\_Capacity \cdot Number\_of\_Onus\right) \qquad (5)$$

Assuming that *Wavelength_Capacity* is predefined, the complexity of our algorithm is linear with respect to the number of ONUs. Scaling the network, by adding more ONUs can, therefore, be easily achieved.

III. Simulation Setup and Results

A C++ simulator was built to evaluate our approach. Table I shows the parameters used for our simulation study. A TWDM PON was studied having 15 ONUs and fibers of four wavelengths of 10 Gbps each. The distance (in meters) between the OLT and each respective ONU is set by the formula:

$$(i+1)\cdot 10, i = 0,\ldots, numberOfOnus - 1 \qquad (6)$$

Only upstream traffic was simulated, while tuning times were neglected as per [15]. TI traffic arrives with an exponential distribution, whereas NonTI is periodic. TI packets are 64 bytes as specified by IEEE 802.3av and NonTI traffic are Ethernet frames with mean length 791 bytes. Each ONU possesses a buffer of 1 MB, containing two queues, one for each type of traffic. At the beginning of the transmission window, an ONU calculates the total number of packets per queue. It first empties the recorded packets of TI and if further transmission time exists, it begins to empty the NonTI packets.

Two scenarios are evaluated to assess the performance of our algorithm. Scenario 1 splits the load of an ONU equally among TI and NonTI traffic, while Scenario 2 stresses our algorithm by adjusting the proportion of TI to NonTI traffic to 90:10 [14].

A feature of our algorithm is that it can manage ONUs with different loads of traffic. Therefore, to assess its capabilities, each ONU generates a different amount of traffic. The mean load per ONU is calculated by:

$$meanLoadPerOnu = \frac{totalLoad}{numberOfOnus}, \qquad (7)$$

while the actual load is distributed according to a uniform distribution with a mean value according to (7) and a maximum value of 5 Gbps. Table 1 summarizes the simulation setup.

Our performance results for the first scenario are shown in Figures 2-5. Figure 2 plots the mean end-to-end delay versus the normalized offered load, which is defined as the total offered load in Gbps divided by the fiber link capacity of 40 Gbps. It can easily observed that TI packets experience a delay lower than the limit of 500μs. When normalized load of 1 is implemented, delay slightly exceeds the 500μs goal. The result is expected, due to the fact that the channel limits are reached and the distribution of ONUs per wavelength may be imperfect due to the quantization introduced. It should also be noted that the above results were obtained, only when the exact expression of (1) was implemented. The factor 0.95 decreases the polling cycle time, leading to reduced delays. If this factor was set to 1, the delay experienced by TI packets would fail to achieve the defined goal. T_{poll} cannot be set as large or as small as desired. The maximum value is limited by the target delay, whereas its minimum is dictated by the tuning time of lasers, the distances between the ONU and the OLT and several other limiting factors imposed by current technology. Non TI packets, on the other hand, experience more delay than the TI ones. Nevertheless, up to a normalized load of 0.9 latency is still below 1 ms. As non-Ti packets are Best-Effort, higher latency may be in cases

Table I
Simulation Parameters

Simulation Parameters	Values
Number of ONUs	15
Distance OLT-ONU	10,20,…,150m
Latency Constraint L_{cons}	500μs
Scenarios (TI, NonTI traffic percentages)	(50, 50), (90, 10)
Number of Wavelengths	4
Link Capacity	10 Gbps
ONU buffer	1 MB
TI packet length	64bytes
Non TI packet length	791bytes

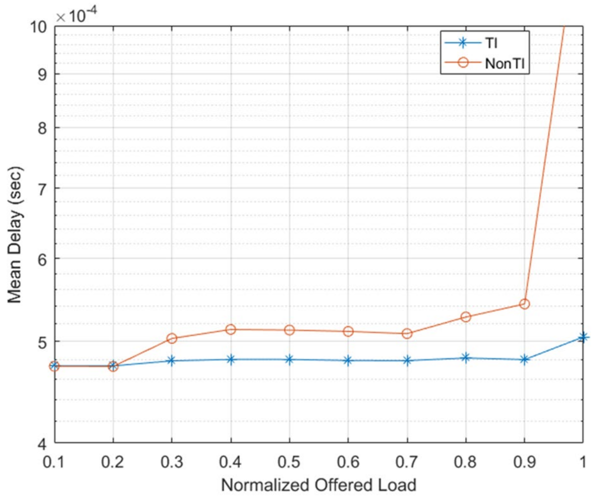

Fig 2. Mean Delay for TI and NonTI packets versus normalized offered load (Scenario 1)

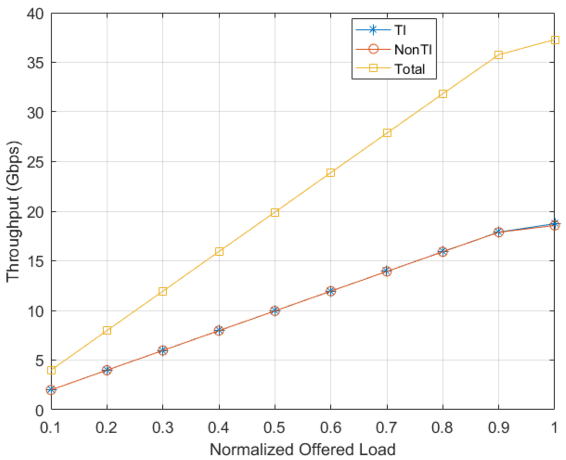

Fig 3. Throughput for TI, NonTI and Total traffic versus normalized offered load (Scenario 1)

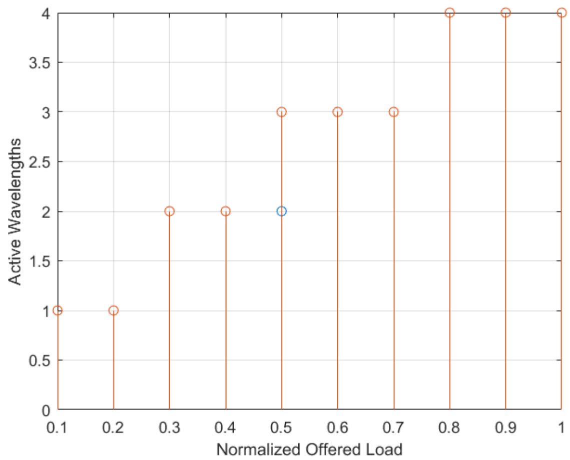

Fig 5. Active wavelengths versus normalized offered load (Scenario 1)

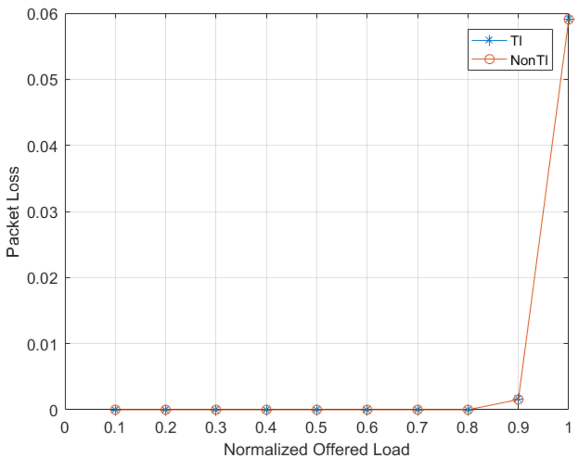

Fig 4. Packet Loss for TI and NonTI traffic versus normalized offered load (Scenario 1)

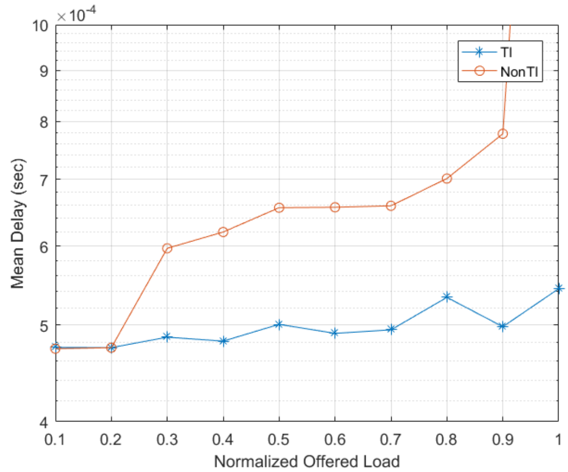

Fig 6. Mean Delay for TI and NonTI packets versus normalized offered load (Scenario 2)

acceptable for them.

Figure 3 presents the throughput achieved for both type of traffic as well as for the total traffic. As TI and NonTI traffic are of equal load and are successfully transmitted from the respective ONU, both traffic types show the same throughput. It can also be seen from Figure 2 that throughput increase is limited for when load approaches to 1, as in this case, the total bandwidth reaches the maximum capacity the network can accommodate. Loads per ONU are uniformly distributed, creating the possibility that a wavelength may not be fully utilized. These mismatches lead to an overloaded last wavelength, rendering it incapable of successfully accommodating some requests. This is the main reason the total throughput drops to approximately 37 Gbps instead of 40.

Figure 4 depicts the packet drop with respect to the total amount of created traffic for this type of packet. No dropping is observed until 0.8 of normalized offered load. When the latter is reached, a drop rate of below 1 percent is attained, which can be acceptable for the Tactile traffic [3]. Increasing the load leads to severe data loss and can be attributed to the fact stated in the previous paragraph.

Figure 5 considers Scenario 1 and reveals the possible number of active wavelengths per load. Simple arithmetic is in line with the results obtained by simulations. Notably, the normalized load of 0.5 (20 Gbps) presents a peculiar behavior. In many cases two wavelengths are assigned, but sometimes three wavelengths can also be activated. Normally, 20 Gbps fit exactly in two wavelengths; nevertheless, the estimation of traffic may not be perfectly accurate, sometimes leading to a higher estimate, enabling three wavelengths, while other times an underestimation occurs and only two wavelengths are used for that polling cycle.

Finally, Figure 6 assesses the performance of our algorithm in a more demanding scenario (Scenario 2) with TI packets being 90% of the total traffic. It can be seen that our algorithm indeed enables TI traffic to preserve its low latency for loads up to 0.8, with a corresponding increase in NonTI packet latencies. This increase of NonTI traffic delay can be easily explained due to the fact the traffic is dominated by TI packets and only when the ONU is completely empty of them, will it be able to transmit the NonTI ones. Furthermore, a great proportion of these packets may have to wait more than one polling cycle to be transmitted, leading to severe latencies in contrast with those experienced in Scenario 1. The best-effort nature of these packets, nevertheless, can justify higher end-to-end transmission times.

The different traffic proportions, configured in scenarios 1 and 2, have stressed the validity of the knapsack DWBA algorithm. Up to normalized loads of 0.8, TI packets suffer latencies below 500 μs, as desired, independent of their proportion. When exceeding 0.8, time constraints are not fulfilled, packet drops become significant and throughput is much lower than maximum. NonTI packets, on the other hand, experience higher delays, due to their lower transmission priority. Since they can be considered as best effort traffic, their quality metrics can vary significantly without degrading the quality of service of applications linked to them.

IV. Conclusion

This paper proposed a prediction-assisted DBA for LAN PONs that can satisfy stringent QoS requirements, such as low latency and high reliability that are envisioned for medical environments. A knapsack-based algorithm was presented, that can fit as many ONUs as possible per wavelength. In general, the knapsack problem is NP-Hard and to reduce its complexity, a transformation of the traffic predictions was introduced, rendering the decision problem linear and capable of being completed timely. A set of simulations proved the efficiency of our framework and has shown that Tactile traffic was successfully delivered below the required time constraint of 500 μs for the vast majority of tested network loads. Therefore, the use of prediction can lead to a new generation of LAN PONs which are capable of supporting healthcare applications, remote controlling/steering etc.

Acknowledgment

This research has been co-financed by the European Union and Greek national funds through the Operational Program Competitiveness, Entrepreneurship and Innovation, under the call RESEARCH – CREATE – INNOVATE (project code:T1EDK-02489).

References

[1] A. Sumits, "The History and Future of Internet Traffic", blogs@Cisco - Cisco Blogs, 2019. [Online]. Available: https://blogs.cisco.com/sp/the-history-and-future-of-internet-traffic. [Accessed: 18- Jan- 2019].

[2] M. Maier, M. Chowdhury, B. Rimal and D. Van, "The tactile internet: vision, recent progress, and open challenges", IEEE Communications Magazine, vol. 54, no. 5, pp. 138-145, 2016. Available: 10.1109/mcom.2016.7470948.

[3] G. Fettweis, "The Tactile Internet: Applications and Challenges", IEEE Vehicular Technology Magazine, vol. 9, no. 1, pp. 64-70, 2014. Available: 10.1109/mvt.2013.2295069.

[4] M. Simsek, A. Aijaz, M. Dohler, J. Sachs, and G. Fettweis, "5G-Enabled tactile internet," IEEE J. Sel. Areas Commun., vol. 34, no. 3, pp. 460–473,Mar. 2016.

[5] I. Lidal, H. Holte and G. Vist, "Triage systems for pre-hospital emergency medical services - a systematic review", Scandinavian Journal of Trauma, Resuscitation and Emergency Medicine, vol. 21, no. 1, p. 28, 2013. Available: 10.1186/1757-7241-21-28.

[6] E. Frykberg, "Triage: Principles and Practice", Scandinavian Journal of Surgery, vol. 94, no. 4, pp. 272-278, 2005. Available: 10.1177/145749690509400405.

[7] ITU-T, NG-PON2 Standard, Available: https://www.itu.int/rec/T-REC-G.989/en.

[8] Sarigiannidis, P., Papadimitriou, G., Nicopolitidis, P.,Varvarigos, E. and Yiannopoulos, K., 2014, April. Towards afair and efficient downlink bandwidth distribution in xg-pon frameworks. In Electrotechnical Conference (MELECON),2014 17th IEEE Mediterranean (pp. 49-53). IEEE.

[9] Kyriakopoulos, C.A. and Papadimitriou, G.I., 2017. Predicting and allocating bandwidth in the optical access architecture XG-PON. Optical Switching and Networking, 25, pp.91-99.

[10] Sarigiannidis, P., Gkaliouris, A., Kakali, V., Louta, M., Papadimitriou, G., Nicopolitidis, P. and Obaidat, M., 2014, December. On forecasting the ONU sleep period in XG PON systems using exponential smoothing techniques. In Global Communications Conference (GLOBECOM), 2014 IEEE (pp. 2580-2585). IEEE.

[11] X. Hu, X. Chen, Z. Zhang and J. Bei, "Dynamic Wavelength and Bandwidth Allocation in Flexible TWDM Optical Access Network", IEEE Communications Letters, vol. 18, no. 12, pp. 2113-2116, 2014. Available: 10.1109/lcomm.2014.2367119.

[12] S. Hussain, W. Hu, H. Xin, A. Mikaeil and A. Sultan, "Flexible Wavelength and Dynamic Bandwidth Allocation for NG-EPONs", Journal of Optical Communications and Networking, vol. 10, no. 6, p. 643, 2018. Available: 10.1364/jocn.10.000643.

[13] L. Wang et al., "Dynamic Bandwidth and Wavelength Allocation Scheme for Next-Generation Wavelength-Agile EPON", Journal of Optical Communications and Networking, vol. 9, no. 3, p. B33, 2017. Available: 10.1364/jocn.9.000b33.

[14] J. Neaime and A. Dhaini, "Resource Management in Cloud and Tactile-Capable Next-Generation Optical Access Networks", Journal of Optical Communications and Networking, vol. 10, no. 11, p. 902, 2018. Available: 10.1364/jocn.10.000902.

[15] E. Wong, M. Pubudini Imali Dias and L. Ruan, "Predictive Resource Allocation for Tactile Internet Capable Passive Optical LANs", Journal of Lightwave Technology, vol. 35, no. 13, pp. 2629-2641, 2017. Available: 10.1109/jlt.2017.2654365.

[16] J. Kleinberg and E. Tardos, Algorithm Design. Harlow: Pearson Education, 2014.

[17] G. Gens and E. Levner, "Complexity of approximation algorithms for combinatorial problems", ACM SIGACT News, vol. 12, no. 3, p. 52, 1980. Available: 10.1145/1008861.1008867