

Federated Learning over Next-Generation Ethernet Passive Optical Networks

Oscar J. Ciceri, Carlos A. Astudillo, Zuqing Zhu, and Nelson L. S. da Fonseca

ABSTRACT

Federated Learning (FL) is a distributed machine learning type of processing that preserves the privacy of user data, sharing only the parameters of ML models with a common server. The processing of FL requires specific latency and bandwidth demands that must be fulfilled by the operation of the communication network. This article introduces two Dynamic Wavelength and Bandwidth Allocation algorithms for TWDM-PONs: one based on bandwidth reservation and the other on statistical multiplexing for the Quality of Service provisioning for FL traffic over 50 Gb/s Ethernet Passive Optical Networks.

INTRODUCTION

The traditional use of Machine Learning (ML) models relies on batch processing in a central server in conjunction with the employment of datasets containing user data. With the worldwide adoption of data protection and privacy legislation, the creation of datasets and applications based on ML has been considerably limited. One way of coping with such restrictions is the adoption of Federated Learning (FL), which is a distributed way of processing machine learning algorithms that does not disclose private data. In FL, clients train a local ML model using a private dataset, and the parameters of these local models are then sent to a central server. The server then produces a global model on the basis of the numerous parameter values received and distributes this global model to the clients for further training. This round of processing is repeated until the global model produces results with an acceptable level of accuracy. In this way, user privacy is preserved. The most common approaches for the consolidation of the parameters sent by the clients to produce the global model rely on the assumption that clients are synchronized and that local datasets are independent and identically distributed [1].

FL can be classified either as cross-device or cross-silo [2]. While cross-device FL involves thousands of clients with limited computational capacity devices, such as smartphones, cross-silo FL involves hundreds of clients with devices with a larger computational capacity, such as edge devices. This article focuses on cross-silo FL with edge devices connected to a passive optical network.

The processing of FL models has brought several challenges to communication networks. FL clients may, for example, produce highly bursty

traffic when uploading their model parameters to the server. Clients training a Convolutional Neural Networks (CNNs) model with few convolution layers and thousands of parameters may need to send hundreds of megabytes. When millions of parameters are involved, the number of bytes sent can be of the order of gigabytes [3]. Moreover, in FL, data transmissions are usually encrypted [4]. Even for small models, the amount of transmitted data can be huge, due to a 100 to 1000 times increase in data caused by homogeneous encryption. When, for example, a 10 MB model is trained, the amount of data actually transmitted after homogeneous encryption can expand to 1–10 GB.

Moreover, FL may impose stringent communication delays for the uploading of client parameters to enhance fast convergence to the global model, especially when the federation involves numerous clients. To cope with diverse communication delays, the server may either wait for the arrival of the local parameters from all the clients, increasing convergence time, or exclude the late arriving data from the parameter consolidation step, which reduces the accuracy of the model [5]. Moreover, FL may also require a very large number of training rounds to produce accurate global models [6]. These challenges call for efficient resource allocation mechanisms to meet the FL requirements.

Passive Optical Network (PON) is a cost-efficient access network technology for delivering broadband services [7]. Operators have already deployed 10 Gb/s Time Division Multiplexing (TDM) PONs during the past two decades. In recent years, the ITU and IEEE standardization groups have proposed next-generation PONs based on Time and Wavelength Division Multiplexing (TWDM) to increase the network capacity for supporting demanding applications and services. TWDM allows allocation in various wavelength channels of 25 Gb/s (50G-EPON) and 10 Gb/s (40G-XPON) [8].

A few approaches have been proposed to deal with FL processing over PONs ([5, 9]). An architecture for scalable FL involving two-step of aggregation was introduced in [9]. The parameters of local models are first aggregated at clients connected to an Optical Network Unit (ONU) and then aggregated on a server connected to the Optical Line Terminator (OLT). As a consequence, the amount of upstream traffic remains relatively constant, regardless of the number of clients in the federation. However, this approach

does not define a Dynamic Bandwidth Allocation (DBA) algorithm required to handle the demands of both FL clients and conventional PON clients.

A seminal approach to support FL over TDM-PONs, called Bandwidth Slicing (BS), was introduced in [5]. It reserves a certain bandwidth, known as a slice (a fraction of the PON capacity), for FL clients [5]. Although the bandwidth reserved can change from round to round in FL processing, the bandwidth needed to cope with additional demands may not be available, since it may be allocated to other clients in a PON. When this happens, clients will need to send the model parameters to the server in several PON cycles, which increases the number of clients who are stragglers, and, consequently, the overall processing time of FL applications. Moreover, the BS approach does not specify which client is financially responsible for the reservation, which is not aligned with traditional PON business models based on charging according to what has been agreed in a Service Level Agreement (SLA). Moreover, BS was not developed for TWDM-PON and cannot be applied to Ethernet-PONs since it requires knowledge of the cycle length in advance.

This article introduces a novel algorithm for bandwidth slicing in TWDM PONs, called Multi-wavelength BS algorithm (MW-BS), as well as three variations of this algorithm each with different allocation policies. The BS algorithm [5] was adapted to employ multiple wavelengths, as well as an adaptive polling cycle as required in 50G-EPON networks with dynamic resource allocation. However, the MW-BS algorithm does not overcome the other mentioned limitations of the BS algorithm.

This article also introduces another novel Dynamic Wavelength and Bandwidth Allocation (DWBA) algorithm for 50G-EPONs based on DiffServ-like traffic prioritization. FL traffic is prioritized to support the demands of FL processing and communications, while maintaining the traditional guaranteed bandwidth scheme for all PON customers. Two variations of the DWBA algorithm using different prioritization policies for PON traffic are also proposed to reduce the delay of FL traffic and delay-critical applications in 50G-EPONs. In the first, the intra-ONU scheduler strictly prioritizes the FL traffic over that from other types of applications (*FL-first policy*). In the second, the delay-critical traffic is prioritized over the FL traffic (*DC-first policy*). These two policies differ from BS by dynamically allocating bandwidth without prior knowledge of the cycle length.

In summary, this article introduces the following original DWBA algorithms for TWDM-PONs:

- A DWBA algorithm for bandwidth slicing based on reservation and three variations of the algorithm employing different allocation policies.
- A DWBA algorithm based on traffic prioritization and two variations of this algorithm.

RESOURCE ALLOCATION IN PASSIVE OPTICAL NETWORKS

PON is a network access technology that offers larger capacity, greater cost-efficiency, and more energy savings than do other network access technologies. There are two main PON standards: Ethernet PON and Gigabit Capable PON, with EPON being less expensive. GPON transmission

The 50 Gb/s optical access network standardized in IEEE 50G-EPON 802.3ca-2020 [10] is a promising technology for adoption by Infrastructure Service Providers to support emerging services with stringent latency and bandwidth requirements.

system employs synchronous frames issued every 125 μ s, while those of EPON use Ethernet frames asynchronously for transmissions based on granting cycles of variable duration. While traditional PON standards allow bit rates of 1 Gb/s and 10 Gb/s, the next-generation PON standards allow those of 40 Gb/s to 100 Gb/s.

The 50 Gb/s optical access network standardized in IEEE 50G-EPON 802.3ca-2020 [10] is a promising technology for adoption by Infrastructure Service Providers to support emerging services with stringent latency and bandwidth requirements. This 50G-EPON technology employs the Time and Wavelength Division Multiple Access (TWDM) technique for controlling uplink transmissions between the ONUs and the OLT. There are three main TWDM-PON-based access architectures for the connectivity between the OLT and ONUs: Multiple-Scheduling Domain (MSD), Single-Scheduling Domain (SSD), and Wavelength Agile (WA). In the first, ONUs transmit on a single wavelength at a time. In the second, ONUs can transmit simultaneously on all wavelengths, and in the third, more than one wavelength can be granted to a single ONU.

In this technology, the signaling protocol Multipoint Control Protocol (MPCP) is employed for resource allocation, using Report and Gate messages for this purpose. Report messages are sent on the upstream to the OLT by the ONUs to request bandwidth, while Gate messages are sent on the downstream by the OLT to the ONUs to inform the granted wavelength(s) and transmission windows, as well as the starting time of the next transmission window. Resource allocation is carried out in two steps, one for wavelength allocation, and the other for bandwidth allocation. The use of different schemes for transmission on multiple wavelengths can be defined on the basis of conventional DBA algorithms for TDM-PONs.

For dynamic bandwidth allocation over EPONs, the Interleaved Polling with Adaptive Cycle Time (IPACT) algorithm has been adopted to complement the MPCP protocol. The IPACT algorithm employs an interleaved polling and statistical multiplexing technique that leads to efficient upstream channel usage. The Limited policy has been used to assure bandwidth to ONUs according to pre-defined Service Level Agreements. Moreover, the original IPACT algorithm employs a single wavelength channel for scheduling. It has been modified to operate with multiple wavelengths in [11–13]. The modified IPACT algorithm was proposed for the SSD and MSD architectures [11]. Additional algorithms have also been proposed: the Water-Fill (WF) [12] to promote fairness in the wavelength utilization and First-Fit (FF) [13] to provide less delay. Moreover, when there is no scheduler for Quality of Service (QoS) provisioning in the PON, the First-Come-First-Served (FCFS) queuing policy is employed. However, this strategy does not consider the priority or required bandwidth/delay of the applications. The performance of diverse appli-

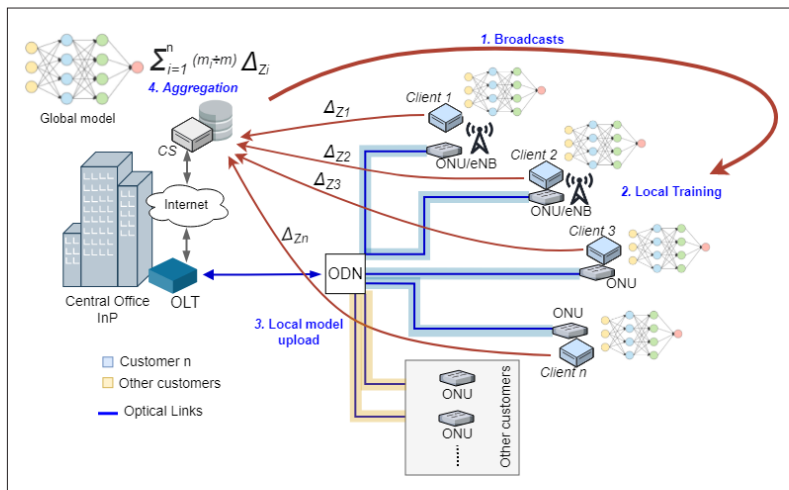


FIGURE 1. An overview of federated learning over passive optical networks.

cations in a PON is ensured by the QoS mechanism adopted, which controls the way frames are queued, prioritized, and scheduled. Such assurance of QoS can be provided by either the ONU or OLT. In the single-level architecture, the ONUs report individual queue sizes, while the OLT distributes the bandwidth for each type of traffic. In the hierarchical architecture, the OLT allocates bandwidth for each ONU, and the ONUs manage the amount of bandwidth to be allocated for each type of traffic.

The most straightforward method of facilitating QoS provisioning is the Differentiated Service approach. It classifies network traffic and delivers different services to different applications. The simplest way to implement this approach is to employ strict priority scheduling. The ONUs categorize the incoming traffic and put it in the buffer, imposing the prioritization of the different traffic classes. With the employment of Differentiated Services, however, the PON can support packetized voice and video with strict bandwidth and latency constraints, as well as Best Effort (BE) traffic [14]. However, none of the exiting mechanisms have been specially designed to support the QoS requirements of FL applications.

RESOURCE ALLOCATION FOR FEDERATED LEARNING

Figure 1 illustrates the scenario of FL processing over a PON. FL clients are connected to the ONUs. The FL server is remotely located in the Internet, and access to the server is provided by the OLT connected to the Internet.

The FL training process can be either asynchronous or synchronous. In the former, the global model parameters are computed as soon as the server receives updates of the parameters of the local models from a certain number of clients. In the latter, the server aggregates the local parameters that arrive in a period of constant duration and excludes the parameters from the late arriving stragglers. Such exclusion can, however, reduce the accuracy of the model and increase the time necessary to obtain a final global model.

The synchronization time per round includes the downstream, computing, network, and aggregation delays. The downstream delay includes the propagation and transmission delays of the parameters of the global model from the server to the clients. The computing time is the time taken

to train the local model at the client in each round and depends on the capacity of the client and the size of the training dataset. The network delay is the time spent in communicating the local model parameters from the clients to the server, including both transmission and propagation delays, and depends on the network load and the mechanism for allocation of bandwidth and wavelength(s) to the ONUs. Long network delays may increase the number of straggler clients, decreasing the model accuracy and increasing the time to reach a final global model. The aggregation delay is the processing time of the aggregation algorithm.

The time taken to transmit the local model parameters to the server depends on the bandwidth allocated to the FL traffic. PON customers receive a portion of the total available bandwidth in the PON due to the shared nature of the upstream channel. Residential and business customers usually have a guaranteed bandwidth of tens to hundreds of Mb/s, while other PON customers can require on demand up to tens of Gb/s. However, the large size of the local model parameters, which may be in the order of gigabytes, may demand several seconds to be fully transmitted, even with guaranteed bandwidth in the order of Gb/s.

The unique characteristics of FL processing, such as desired synchronization of clients in FL rounds and assurance of bandwidth, introduce challenges for the management of the network bandwidth in scenarios with limited bandwidth and a diversity of customers, such as that in commercial PONs.

DWBA SCHEMES FOR SUPPORTING FL TRAFFIC OVER 50G-EPON NETWORKS

This section introduces new DBA algorithm for the support of FL over 50G-EPON networks. These algorithms are based either on bandwidth reservation (BS) or statistical multiplexing.

BANDWIDTH SLICING DWBA FOR TWDM-EPON

The seminal BS approach to support FL over PONs was designed to operate in TDM-PONs [5]. It reserves bandwidth, known as slices, for FL clients. The reserved bandwidth can be allocated according to the ascending order in downstream client delay and computing time.

BS calculates the number of cycles an FL client requires to be completely served based on the required bandwidth and the fixed polling cycle length of 125 μ s employed in the GPON technology. However, next-generation PONs employ multiple wavelengths and the cycle duration is unknown a priori when an adaptive polling cycle mechanism is employed, such as the EPON technology. Therefore, the proposed BS algorithm cannot be directly used in TWDM-EPON networks.

We propose a DWBA algorithm based on the BS approach [5] for TWDM-EPONs called multi-wavelength BS algorithm (*MW-BS*) which deals with multiple wavelengths and employs an adaptive polling cycle for dynamic resource allocation. A portion of the PON capacity (bandwidth slice) is still reserved for the FL traffic in each scheduling cycle, but instead of using the polling cycle information to grant the transmission windows for FL traffic and then share the remaining slice capacity with other traffic types, *MW-BS*

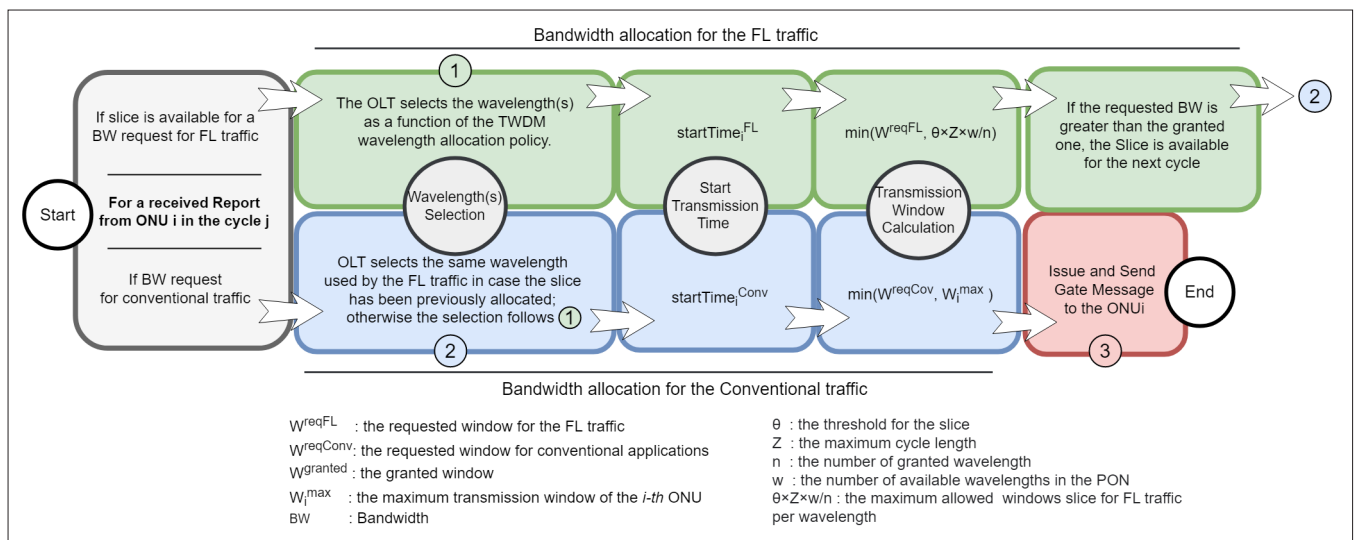


FIGURE 2. Flow chart of the bandwidth slicing algorithm adapted for TWDM-EPONs.

reserves the total slice for the FL traffic as long as a bandwidth request from any FL client exists. The use of a dynamic polling cycle reduces the FL traffic delays and avoids the need for information about the duration of the unknown upcoming cycles. Three variations of this MW-BS algorithm are proposed for different TWDM wavelength allocation policies, namely *MW-BS-SSD*, *MW-BS-MSD*, and *MW-BS-FF*. The flow diagram in Fig. 2 summarizes the proposed DWBA scheme residing in the OLT. The ONUs send Report messages requesting bandwidth for FL as well as for conventional applications. When a Report message arrives from an ONU containing a bandwidth request for FL traffic, the OLT first grants the bandwidth from the reserved slice, if available; otherwise, the OLT allocates bandwidth for the conventional traffic.

The OLT also reserves bandwidth for the ONU for upcoming cycles. It selects the wavelength(s) as a function of the TWDM wavelength allocation policy, and calculates the next starting time for the FL transmission. For the SSD policy, the OLT grants all wavelengths. For the MSD policy, the OLT grants a predetermined-fixed wavelength. For the FF policy, the OLT grants the first available wavelength, and then calculates the transmission window to be granted for each allocated wavelength, as a function of the number of wavelengths and the portion of the PON capacity designed for FL use. If the granted window is equal to the requested one, the FL traffic will be fully served, and the OLT will make the bandwidth slice available for the next cycle.

The OLT also calculates the granted bandwidth for conventional applications. If the OLT has previously allocated the bandwidth for a slice, it selects these wavelength(s) for the FL traffic. Otherwise, the OLT selects the wavelength(s) for the FL traffic and calculates the next starting time for that FL transmission as a function of the TWDM wavelength allocation policy involved. The transmission window for the conventional applications is calculated according to the limited policy. Finally, the OLT issues and sends a Gate message with the granted bandwidths for both FL and conventional applications.

Even though the BS approach reduces the latency for FL applications in relation to the traditional First Come First Served approach, bandwidth reservation prevents the BS and MW-BS algorithms from dealing with the requirements of bandwidth allocation to FL traffic. More specifically:

- The reserved bandwidth for the FL slice is adjusted according to the number of clients in different FL rounds. However, this bandwidth may not be available since InPs provide bandwidth guarantees for other PON customers according to agreements in the Service Level Agreements (SLAs).
- There might be not enough reserved bandwidth to cope with peaks of the bandwidth demands of FL clients, especially for the transport of large packets in synchronized FL rounds, unless the slice bandwidth is reserved for peak demands. Clients may have to send only part of the local model parameters per cycle, which increases the number of stragglers clients, and, consequently, the overall processing time of FL applications.
- The bandwidth of a slice is shared among all FL clients. This can cause bandwidth starvation for some clients.
- The BS approach is not compatible with the traditional PON business model, in which customers rent portions of the PON capacity from the InP to support their applications. In the BS approach, it is not specified who pays for the shared slice in the federation.

DYNAMIC BANDWIDTH ALLOCATION FOR FEDERATED LEARNING

Bandwidth reservation brings numerous limitations to the processing of FL over PONs. To address these limitations, we propose a DWBA algorithm based on statistical multiplexing rather than on bandwidth reservation. This provides flexibility to cope with dynamic bandwidth demands of FL clients. To guarantee the requirements of FL clients we introduce a DiffSev-like static prioritization for bandwidth allocation, as well as a DWBA that supports QoS provisioning for FL applications while meeting the requirements of delay-critical applications in TWDM-EPON networks. The algorithm is called DWBA for Federated Learning (DWBA-FL).

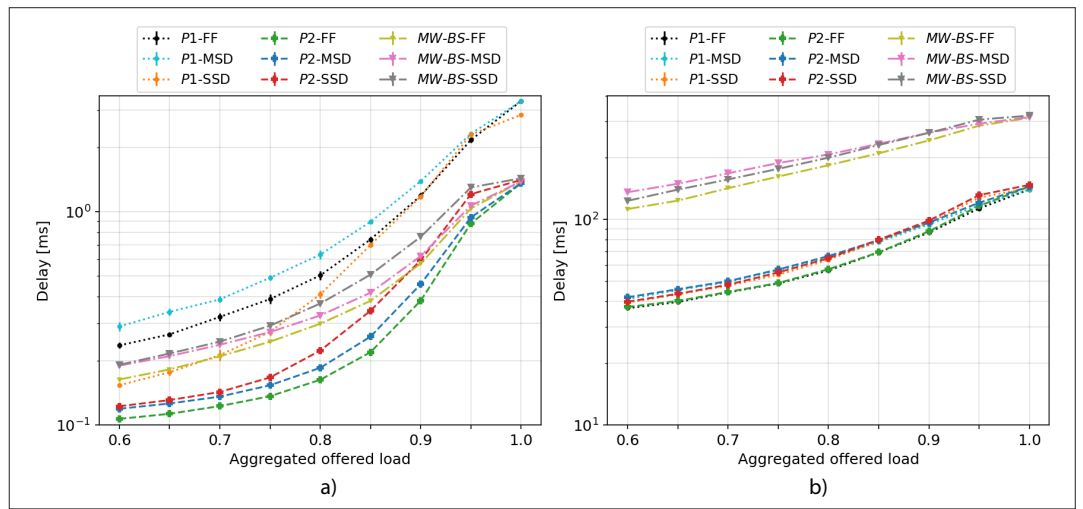


FIGURE 3. Average delay produced by the delay-critical and FL applications. $P1$ and $P2$ mean, respectively, DWBA-FL algorithm with FL -first policy and DC -first policy: a) Delay-critical application; b) Federated Learning application.

The idea behind our proposal is to allow PON customers to employ their guaranteed bandwidth for the scheduling of the FL application, but without jeopardizing the QoS provisioning of other delay-critical applications. To achieve this, the proposed mechanism adopts the widely-used Differentiated Service approach to tackle the QoS provisioning problem of FL applications over Ethernet PON. Just mapping FL traffic into a DiffServ per-hop behavior (PHB) does not provide the bandwidth guarantee needed for FL applications since the FL traffic would compete with traffic of other type of clients in the same PHB. By creating a PHB exclusive for FL, it is possible to treat FL traffic differently from the traffic of other PON clients, thus, allowing bandwidth allocation mechanisms to guarantee the required bandwidth. The prioritization of FL traffic complies with the traditional business model, as well as to improve statistical multiplexing gain.

The proposal employs an intra-ONU scheduler with a strict priority queuing policy for the ONU queues, and the ONUs can arbitrate the transmission demands of the different applications. Upon the arrival of a Report message, the OLT calculates the transmission window according to the conventional Limited policy and selects the wavelength(s) on the basis of the TWDM wavelength allocation policies. The OLT then sends a Gate message containing the resource allocation decision. Upon the receipt of that Gate message, the intra-ONU scheduler distributes the received bandwidth among the queues in the ONU. In our model, traffic is classified as FL, delay-critical, delay-sensitive, or BE. The ONUs maintains four different queues for buffering frames for these types of traffic.

We propose two prioritization policies. The FL -first policy defines the FL traffic as that of the highest priority and the delay-critical, delay-sensitive, and BE traffic being of high, medium, and low priority, respectively. On the one hand, this strict prioritization of FL frames can reduce synchronization time for FL processing. It can also increase the delay of delay-critical application because the FL traffic requires a large bandwidth per cycle. To help alleviate this problem, we propose the DC -first policy, assigning highest priority and reducing FL

traffic to only the high priority. This proposal has been defined for all TWDM architectures.

The proposed algorithm processes each Report message from each ONU once per cycle, thus the allocation is performed with a computational complexity of $O(n)$, where n is the number of ONUs in the PON.

PERFORMANCE EVALUATION

The performance of the proposed DWBA scheme was evaluated using an EPON simulator (EPON-Sim), previously validated in [15]. This extension supports the three architectures, SSD, MSD and WA, proposed for 50G-EPON networks. Moreover, our proposed DWBA algorithms were introduced in the simulator.

SIMULATION MODEL AND SETUP

The simulation scenarios include a 50G-EPON network with 1 OLT serving 32 ONUs on an optical distribution network with a tree topology. Two wavelength channels of 25 Gb/s were employed for upstream transmission, giving a total capacity of 50 Gb/s. The total available bandwidth in the PON was equally distributed among the ONUs, so that each ONU has the same guaranteed bandwidth b_i , while the aggregated offered load per ONU l_i varied from $0.6 \cdot b_i$ to $1.0 \cdot b_i$ (for the sake of clearness and brevity, herein after, b_i is omitted from the offered load values of ONU i).

The aggregated load included the traffic of the four different types of application: FL, delay-critical, delay-sensitive, and BE. The benchmarking framework for learning in federated settings LEAF [6] was used to generate the FL traffic. The FEMNIST dataset and CNN with two 5×5 convolution layers were used for model training, while the FedAvg algorithm was employed to aggregate the local parameters in the server. Other configurations for the learning process, such as learning rate and batch size, followed the settings defined in [9]. FL clients generated 26.4 MBytes of data in each round of training. Moreover, the ONUs put the local parameters into frames according to the Ethernet protocol, which has a Maximum Transmission Unit of 1500 bytes and a header field for signaling (preamble) of 20 bytes.

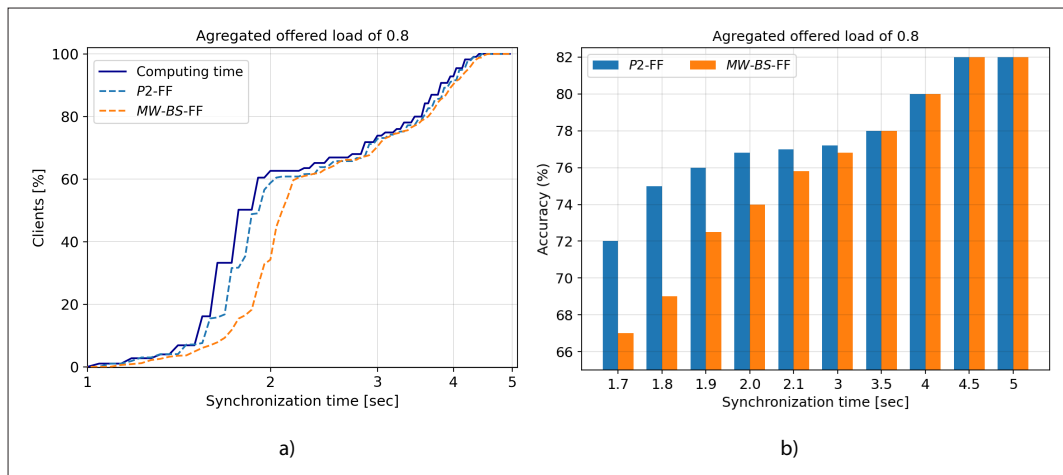


FIGURE 4. Percentage of the learning accuracy and involved clients depending on the synchronization time. P2 means DWBA-FL algorithm with DC-first policy: a) Involved clients vs. synchronization time; b) Learning accuracy (with 2000 rounds) vs. synchronization time.

The delay-critical applications were modeled employing a Constant Bit Rate (CBR) flow. It was coded with a fixed-size packet of 70 bytes and an inter-arrival time of 12.5 μ s, which produces an offered load of 44.8 Mb/s. The rest of the offered load I_i was evenly distributed between delay-sensitive and BE traffic. The traffic streams were generated employing Pareto ON-OFF sources. The ON period time and packet-burst size followed a Pareto and Bounded Pareto distributions, respectively. The aggregated traffic at the ONU had a Hurst parameter of 0.8. Moreover, the packet lengths were uniformly distributed between 64 and 1518 bytes.

A threshold value of $\theta = 0.015$ was employed in the MW-BS algorithm, as in [5]. This algorithm reduces the bandwidth for each ONU since it reserves bandwidth for the slice. Moreover, the same aggregated offered load was employed in the simulated algorithms to make a fair comparison. The duration of the guard period was set to 0.624 μ s, with a maximum cycle length of 1 ms. Each simulation scenario lasted 100 s and was replicated 10 times.

SIMULATION RESULTS AND DISCUSSION

Mean delay values obtained by the DWBA-FL were less than 80 ms and 150 ms for the FL traffic in both underloaded and overloaded conditions, respectively. The delay values given by BS were at least twice as large as those given by our proposal (Fig. 3b). This improvement is a consequence of the large windows allocated for transmissions of FL traffic when our proposal is employed.

Moreover, the use of the DC-first policy produced lower, delay values for the delay-critical traffic lower than those given by either the FL-first policy or the BS algorithm (Fig. 3a). This result is due to the static allocation of bandwidth slice for the FL traffic. Furthermore, the strict prioritization of FL traffic employing the FL-first policy and the huge amount of traffic produced by the FL application leads to bandwidth starvation for delay-critical application. The mean delay of the delay-critical traffic produced by the FL-first policy was from 200 μ s to 1000 μ s, greater than that produced by FL-first policy. Thus, the DWBA-FL

with DC-first policy produces mean delay values for the FL and delay-critical applications lower than those for the other algorithms.

Furthermore, the FF policy produces a slight decrease in delay values for both type of traffic in relation to the other wavelength allocation policies (i.e., SSD and MSD). These results are a consequence of the waste of bandwidth due to the excessive uses of guard periods and poor multiplexing gain when employing the SSD and MSD, respectively.

In Fig. 4a, the blue curve shows the proportion of clients involved as a function of the computing time. It shows the minimal synchronization time per round without any communication delay. The MW-BS algorithm requires a longer synchronization time per round to produce the same percentage of the involved clients than is required by the proposed scheme with DC-first policy. For example, synchronization times of 1.9 s and 2.1 s were required to produce a percentage of involved clients of 50 percent with our proposal and the MW-BS algorithm, respectively.

In the simulations, the target accuracy was obtained after 2000 rounds of training (Fig. 4b), that is, the verified convergence time in rounds. To achieve a training accuracy of 76 percent, the proposed scheme can reduce 9.5 percent of the training time required by the BS algorithm (i.e., 0.2 s less for a synchronization time of 2.1 s), when the total traffic load is 0.8.

Figure 5 shows the network delay as a function of the ONU offered load. The MW-BS produces delay values greater than 300 ms, whereas, with the DWBA-FA algorithm, these values are reduced to less than 150 ms. Moreover, for 80 percent of the clients, which is the typical percentage of clients that lead to an accuracy greater than 75 percent (Fig. 4), the MW-BS scheme imposes a network delay greater than 200 ms, while the DWBA-FA imposes delay values less than 100 ms, under underloaded conditions (i.e., load < 0.85). In summary, the DWBA-FL algorithm reduces the network delay in relation to the MW-BS scheme. This reduction in delay may decrease the number of stragglers, which in the end leads to faster convergence and greater model accuracy.




FIGURE 5. Network delay performance for the FL application. Each boxplot shows 10th, 30th, 50th, 80th, and 100th percentile of network delay for FL clients; P2 means DWBA-FL algorithm with DC-first policy.

CONCLUSION

This article has introduced two DWBA algorithms for the support of FL applications over TWDM-EPON networks. A DWBA algorithm based on bandwidth reservation, as well as three different variations of this algorithm, have been introduced. Moreover, a DWBA algorithm that employing static prioritization of FL traffic, with two variations proposed. The latter includes a strict prioritization for FL and delay-critical traffic.

Results show that the DWBA-FL algorithm with DC-first policy increases the FL model accuracy and reduces the delay of federated learning and delay-critical applications when compared to the BS approach and the FL-first policy.

Future research directions are envisioned as follows. Mechanisms are needed to address the QoS provisioning appropriately for diverse FL applications co-existing in the same PON. These schemes may schedule the FL traffics based on required bandwidth but also consider the number of stragglers, diverse FL packet sizes, and synchronization time.

ACKNOWLEDGMENTS

This work was partially sponsored by grant #15/24494-8, São Paulo Research Foundation (FAPESP), and CNPq.

REFERENCES

- [1] B. McMahan et al., "Communication-Efficient Learning of Deep Networks From Decentralized Data," *Artificial Intelligence and Statistics, Proc. PMLR*, 2017, pp. 1273–82.
- [2] P. Kairouz et al., "Advances and Open Problems in Federated Learning," 2021.
- [3] K. He et al., "Deep Residual Learning for Image Recognition," *Proc. 2016 IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–78.
- [4] J. Ma et al., "Privacy-Preserving Federated Learning Based on Multi-Key Homomorphic Encryption," *CoRR*, vol. abs/2104.06824, 2021; available: <https://arxiv.org/abs/2104.06824>.
- [5] J. Li et al., "Bandwidth Slicing to Boost Federated Learning Over Passive Optical Networks," *IEEE Commun. Letters*, vol. 24, no. 7, 2020, pp. 1492–95.
- [6] S. Caldas et al., "Leaf: A Benchmark for Federated Settings," *arXiv preprint arXiv:1812.01097*, 2018.
- [7] F. de Melo Pereira, N. L. S. da Fonseca, and D. S. Arantes, "A Fair Scheduling Discipline for Ethernet Passive Optical Networks," *Comput. Networks*, vol. 53, no. 11, 2009, pp. 1859–78.
- [8] J. S. Wey, "The Outlook for Pon Standardization: A Tutorial," *J. Lightwave Technology*, vol. 38, no. 1, 2020, pp. 31–42.
- [9] J. Li, L. Chen, and J. Chen, "Scalable Federated Learning Over Passive Optical Networks," *arXiv preprint arXiv:2010.15454*, 2020.
- [10] "IEEE Standard for Ethernet Amendment 9: Physical Layer Specifications and Management Parameters for 25 Gb/S and 50 Gb/S Passive Optical Networks," *IEEE Std 802.3ca-2020 (Amendment to IEEE Std 802.3-2018 as amended by IEEE*

- 802.3cb-2018, IEEE 802.3bt-2018, IEEE 802.3cd-2018, IEEE 802.3cn-2019, IEEE 802.3cg-2019, IEEE 802.3cq-2020, IEEE 802.3cm-2020, and IEEE 802.3ch-2020)*, 2020, pp. 1–267.
- [11] M. P. McGarry, M. Reisslein, and M. Maier, "WDM Ethernet Passive Optical Networks," *IEEE Commun. Mag.*, vol. 44, no. 2, 2006, pp. 15–22.
- [12] L. Wang et al., "Dynamic Bandwidth and Wavelength Allocation Scheme for Next-Generation Wavelength-Agile EPON," *J. Optical Commun. Networking*, vol. 9, no. 3, 2017, pp. B33–42.
- [13] S. B. Hussain et al., "Low-Latency Dynamic Wavelength and Bandwidth Allocation Algorithm for NG-EPON," *J. Optical Commun. Networking*, vol. 9, no. 12, 2017, pp. 1108–15.
- [14] M. P. McGarry, M. Reisslein, and M. Maier, "Ethernet Passive Optical Network Architectures and Dynamic Bandwidth Allocation Algorithms," *IEEE Commun. Surveys & Tutorials*, vol. 10, no. 3, 2008, pp. 46–60.
- [15] O. J. Ciceri et al., "Passive Optical Networking for 5g and Beyond 5g Low-Latency Mobile Fronthauling Services," *CoRR*, vol. abs/2109.08989, 2021; available: <https://arxiv.org/abs/2109.08989>.

BIOGRAPHIES

OSCAR CICERI (oscar@lrc.ic.unicamp.br) completed his five-year degree in Electronics and Telecommunications Engineering at Universidad del Cauca (UNICAUCA), Colombia, in 2015, and his M.S. degree in Computer Science at the University of Campinas (UNICAMP), Brazil, in 2019. He is a Ph.D. student and researcher in the Network Computing Laboratory (LRC) at that university. His research interests include Passive Optical Networks (PONs), 5G and 5G beyond networks, Machine Learning, and virtualization. His current work is on the low-latency PON technologies for 5G and 5G beyond mobile optical access systems.

CARLOS ASTUDILLO (castudillo@lrc.ic.unicamp.br) obtained his Ph.D. degree in Computer Science from the University of Campinas (UNICAMP), Brazil, in 2022. He received first place in the 29th Theses and Dissertations Contest (CTD) of the Brazilian Computer Society (SBC), the 2016 best master's thesis award from IC/UNICAMP, and two IEEE best Conf. paper awards. His research interests include radio resource allocation and quality of service provisioning for machine to machine communications and the Internet of things, fronthauling for 5G & beyond 5G networks, and machine learning in communications and networking.

ZUQING ZHU [M'07, SM'12] (zqzhu@ieee.org) received his Ph.D. degree from the Department of Electrical and Computer Engineering, University of California, Davis, in 2007. He is currently a full professor at USTC. He has published more than 200 papers in referred journals and Conf.s. He was an IEEE Communications Society Distinguished Lecturer (2018–2019) and a Senior Member of OSA.

NELSON L. S. DA FONSECA [M'88, SM'01] (nfonseca@ic.unicamp.br) obtained his Ph.D. degree from the University of Southern California in 1994. He is Full Professor at the Institute of Computing, State University of Campinas (UNICAMP). He published 400+ papers and supervised 70+ graduate students. He is Senior Editor for the *IEEE Commun. Mag.* and the *IEEE Systems Journal*. He is also an Associate Editor for *Peer-to-Peer Networking and Applications* and *Computer Networks*. He is past EIC of the *IEEE Communications Surveys and Tutorials*. He served as the *IEEE ComSoc VP Publications*, *VP Technical and Educational Activities*, and *VP Member Relations*.