

Dynamic Bandwidth Allocation for PON Slicing with Performance-Guaranteed Online Convex Optimization

Genya Ishigaki, Siddhartha Devic, Riti Gour, and Jason P. Jue

Department of Computer Science, The University of Texas at Dallas, Richardson, Texas 75080, USA

Email: {gishigaki, sid.devic, rgour, jjue}@utdallas.edu

Abstract—The emergence of diverse network applications demands more flexible and responsive resource allocation for networks. Network slicing is a key enabling technology that provides each network service with a tailored set of network resources to satisfy specific service requirements. The focus of this paper is the network slicing of access networks realized by Passive Optical Networks (PONs). This paper proposes a learning-based Dynamic Bandwidth Allocation (DBA) algorithm for PON access networks, considering slice-awareness, demand-responsiveness, and allocation fairness. Our online convex optimization-based algorithm learns the implicit traffic trend over time and determines the most robust window allocation that reduces the average latency. Our simulation results indicate that the proposed algorithm reduces the average latency by prioritizing delay-sensitive and heavily-loaded ONUs while guaranteeing a minimal window allocation to all ONUs.

Index Terms—Dynamic Bandwidth Allocation (DBA), Online Convex Optimization (OCO), Passive Optical Networks (PONs)

I. INTRODUCTION

Emerging 5G networks are being designed to support a diverse array of network applications. Enhanced Mobile Broadband (eMBB) targets up to 10 Gbit/s download speeds for mobile devices, while massive Machine Type Communications (mMTC) handles dense connections for IoT and M2M applications, and Ultra Reliable Low Latency Communications (URLLC) involves strict latency constraints [1]. An approach for accommodating massive and diverse connections in 5G access networks is to implement network slicing in Passive Optical Networks (PONs) [2]. Optical access networks have attracted many industrial attentions, as they are expected to reduce operational expenditure (OPEX) [3]. In particular, the current development of C-RAN technology for mobile networks accelerates the benefit of such optical access networks [4], [5].

A PON consists of an Optical Line Terminal (OLT) and multiple Optical Network Units (ONUs) as depicted in Figure 1. The OLT and ONUs may be logically sliced so that one physical network can host multiple network infrastructure slices that are tailored for various types of applications. In this slicing scenario, additional consideration should be made to prioritize delay-critical network slices. For example, if a network slice hosts a URLLC network service, the ONUs on the slice should receive more bandwidth to meet the strict delay requirement on the order of 1 ms for both upstream and downstream traffic [6].




Fig. 1. PON Access Network Architecture: A physical network is sliced to support different application requirements.

Time Division Multiplexing (TDM) for the upstream traffic from ONUs to an OLT is facilitated by a *Dynamic Bandwidth Allocation (DBA)* algorithm that adjusts time windows for each ONU to minimize the latency of upstream traffic [7]. Conventional DBA algorithms periodically solve a static optimization problem to minimize the traffic latency by choosing an optimal bandwidth allocation given traffic information of the ONUs [8], [9], [10], [11]. The traffic information can be an estimation of the future traffic amount at the time when the ONU starts sending traffic to the OLT. The estimation is computed based on historical data, the amount of remaining data in queues, and/or neural network-based traffic prediction. Nevertheless, errors in such estimations may lead to suboptimal solutions for the bandwidth allocation problem. The suboptimality ironically originates from the fact that such estimation-based algorithms try to achieve the optimum window allocation in each static optimization problem, which is defined based on the estimated traffic demand. Since the allocation is optimum only for the estimated demand, the discrepancy between the estimation and the real traffic patterns is directly reflected in the performance of the algorithms.

Therefore, we propose a performance-guaranteed online DBA algorithm based on Online Convex Optimization (OCO), which does not involve static optimizations based on an unreliable estimation of traffic patterns. An OCO algorithm generates a sequence of decisions, observing the actual past traffic information rather than an estimation of future states. The decision sequence will be adjusted along its runs to minimize the *regret*, the loss incurred by diverting from a decision that is robust over the entire time horizon. Note that a truly robust




Fig. 2. Conventional interactions between OLT and ONUs: The OLT decides a bandwidth allocation $\mathbf{x}_t[i]$ to each ONU u_i based on the demand request included in a REPORT message.

decision can be defined only in retrospect. In the PON context, the OCO-based DBA algorithm decides a window allocation for the next time cycle based on the historical window sizes and *actual* traffic information observed in the previous cycles, so that the resulting allocation becomes robust against the fluctuating traffic. In other words, the algorithm implicitly learns a potential distribution of future traffic and provides a safe window allocation that can accommodate any traffic dynamics that may happen under the distribution. Furthermore, the performance of the algorithm can be theoretically bounded in terms of regret [12]. Our algorithm is more promising than DBA algorithms based on traffic estimation, since its allocation would realize lower latency within a potential traffic range, instead of being optimal only at one estimated traffic amount. Our DBA problem is formulated in such a way as to ensure fairness among ONUs, which guarantees that each ONU receives a minimal size of time window even when it is not heavily loaded. This formulation could be seen as an extension of the work in [13] to a dynamic setting. The simulation results show that our OCO-based DBA algorithm reduces the latency at all ONUs, demonstrating slice-awareness, demand-responsiveness, and allocation fairness.

II. PON ACCESS NETWORK MODEL

A Passive Optical Network (PON) consists of an Optical Line Terminal (OLT) and Optical Network Units (ONUs). Since all the ONUs share one optical cable connecting them to the OLT, they need to be coordinated to share the bandwidth for upstream traffic. Time Division Multiplexing (TDM) PON realizes the sharing by allocating a time window to each ONU in each logical time cycle.

Figure 2 illustrates conventional interactions of an OLT and an ONU $u_i \in U$ in a PON. An OLT has a Dynamic Bandwidth Allocation (DBA) algorithm that decides the allocation to ONUs. In EPON or 10G EPON standards, an ONU sends a REPORT message that includes a demand request for its next turn. The request is generated by a traffic estimation algorithm at ONUs, which may consider the amount of currently buffered traffic and the future traffic pattern. Receiving the request, the OLT decides the actual bandwidth allocation for

the ONU and informs the ONU of the assigned time windows through a GATE message.

However, it is pointed out that the DBA computation using the latest requests from ONUs causes an extra delay to grant the next TDM windows [14]. Therefore, some works even propose DBA algorithms that do not use the REPORT messages. For example, Cooperative DBA algorithms (CO-DBAs) [15] compute the future bandwidth allocation based on the mobile schedule information provided by a Base Band Unit (BBU), which is not an element of a PON itself, instead of waiting for the requests from ONUs. The allocation decision that does not rely on the ONU requests enables CO-DBAs to reduce the overall latency.

Based on the discussion about the negative effect of request messages, we also consider a bandwidth allocation problem in a PON model that is slightly different from the conventional OLT/ONU interactions. Our model does not have allocation requests from ONUs, which cause the extra delay and could include unreliable statistics such as outdated queue state or estimated future traffic demands. The REPORT messages are repurposed so that they report the actual traffic amount, which implies the discrepancy between the actual window size needed and the actual allocation executed in the previous time cycle. In other words, GATE messages to grant next upstream windows are generated based on not the ONU requests but the allocation decision computed from the history of the actual traffic amount buffered at each ONU.

In this paper, we assume that an OLT divides the time horizon into multiple logical time cycles $\{C_t\}$ ($t = 1, \dots, T$) and solves the bandwidth allocation for each time cycle. Note that each ONU receives its turn to send out upstream traffic within the time cycle. With the notation, the amount of bandwidth allocated to an ONU $u_i \in U$ at a cycle C_t is denoted as $\mathbf{x}_t[i]$.

III. PROBLEM STATEMENT

Our DBA problem is formulated to provide *slice-aware*, *demand-responsive*, and *fair* bandwidth allocation for ONUs. The slice awareness indicates the capability to satisfy additional delay requirements imposed by some slices. The demand responsiveness is defined as the ability to adjust window sizes based on the amount of traffic loads at ONUs. Furthermore, the bandwidth allocation is fair when some ONUs with huge traffic do not dominate a cycle; i.e., each ONU receives an opportunity to send a minimal amount of its queued traffic within each cycle.

The concept of *proportional fairness* [16] is introduced to formulate the DBA problem incorporating the three properties. A vector $\mathbf{x}^*$ is said to be $(\mathbf{w}, 1)$ -proportionally fair if the proportions of value differences in any other allocations $\mathbf{x} \in X$ sum to negative:

$$\sum_i \mathbf{w}_i \frac{\mathbf{x}[i] - \mathbf{x}^*[i]}{\mathbf{x}^*[i]} \leq 0 \quad (\forall \mathbf{x} \in X). \quad (1)$$

Here, the weight $\mathbf{w}_i$ indicates the relative importance of component i , which receives a resource quota.

It is known that a maximizer $\mathbf{x}^*$ for the objective function of the following form over a convex space $X = \{\mathbf{x}\}$ satisfies the proportional fairness [13].

$$\sum_i \mathbf{w}_i \log \mathbf{x}[i]. \quad (2)$$

The intuition behind the relationship between the objective function in Eq. (2) and the fairness concept comes from the decreasing increment of logarithmic functions. Let us consider two points x_1 and x_2 such that $x_1 < x_2$. When evaluating the increment of a logarithmic function h with the same increment δ from each point, $h(x_1 + \delta) > h(x_2 + \delta)$ always holds.

In our context, the increment of our objective function, which represents the utility of a specific bandwidth allocation, becomes larger when providing more bandwidth to an ONU that is assigned a smaller bandwidth window. This property enables us to avoid the domination of a time cycle by a small number of ONUs.

The DBA problem is to determine a time partition $\mathbf{x}_t = (\mathbf{x}_t[i])_{i=1,\dots,|U|}$ of a cycle C_t , where each element in the partition corresponds to a window size allocated to an ONU $u_i \in U$. Let $C \in \mathbb{R}_+$ denote the fixed cycle duration where every ONU receives a fraction of the cycle. An ONU slice $S_j \subseteq U$ is defined as a subset of ONUs that are running on a network slice j , and the prioritization weight p_j indicates the sensitivity of the slice S_j to delay.

Assuming the availability of accurate information regarding the amount $\mathbf{b}_t = (\mathbf{b}_t[i])_{i=1,\dots,|U|}$ of data queued in every ONU u_i at the starting time of its window in C_t , the optimization problem, which maximizes the total weighted utility of allocated windows, is defined as follows:

$$\begin{aligned} \text{Maximize } f_t(\mathbf{x}_t) &= \sum_i \mathbf{b}_t[i] p_{j:u_i \in S_j} \min\{\log(\mathbf{x}_t[i] + 1), \\ &\quad \log(\mathbf{b}_t[i] + 1)\} \\ \text{subject to } \mathbf{x}_t^\top \mathbf{1}_N &\leq C - \sum_i \mathbf{d}[i], \end{aligned} \quad (3)$$

where $\mathbf{d}$ is a vector of guard window sizes that prevent collisions when switching ONUs.

The min function prevents the utility function from increasing by overallocation. The maximum utility that an ONU u_i can experience at a cycle C_t must be bounded by the allocation that empties its entire buffer ($\mathbf{x}_t[i] = \mathbf{b}_t[i]$). When the allocation is less than the emptying allocation, the utility of the ONU should be represented as an increasing function over the allocation $\mathbf{x}$. To represent these property, the objective cuts off the increase in the utility by taking a minimum between two logarithm values.

Since the constraint in Eq. (4) defines a scaled simplex of $\mathbf{x}_t$, the decision space X is a convex set. The objective function f has the exact same form as the proportionally fair optimization discussed above, recognizing the pair of variables $\mathbf{b}_t[i]$ and $p_{j:u_i \in S_j}$ as a weight $\mathbf{w}_i$.

In addition, the fairness is defined with respect to two weights $\mathbf{b}$ and p that indicate the load on ONUs and the prioritization weight of slices, respectively. Therefore, the

maximization with these weights realizes the slice awareness and demand responsiveness, since ONUs with more traffic load and/or with more stringent delay requirements would receive more bandwidth allocations.

IV. PERFORMANCE-GUARANTEED DYNAMIC BANDWIDTH ALLOCATION BASED ON OCO

A. OCO-based DBA Algorithm

With the assumption of an oracle to report the accurate amount of queued data $\mathbf{b}_t$, the defined DBA problem is easy to solve, as it has a concave objective function and a convex variable space. However, the practical difficulty of DBA problems arises from the fact that it is impossible to obtain the actual data amount $\mathbf{b}_t$ at the starting time of the window allocated to each ONU. This is because the demand request in a REPORT message is computed based on the data amount at the time when the message is issued. Furthermore, a prediction of future traffic often differs from the actual amount.

The unavailability of the actual traffic amount hinders an optimal bandwidth allocation. Conventional DBA algorithms try to solve the problem by estimating the future traffic amount with statistical or machine learning methods. Nevertheless, the static optimization relying on the such estimates may demonstrate high variance in its performance, as the estimation is not an easy task in general even with recent machine learning techniques. Furthermore, such estimation will be more complex when considering the high user mobility in access networks. Hence, it seems desirable to use an allocation algorithm that gradually adjusts its allocation strategy based on the past traffic patterns to obtain an optimal allocation robust against fluctuating traffic over time, instead of changing an allocation at every time cycle to make it the exact optimum for an estimated traffic amount. In this sense, we want to design an allocation algorithm that is numb to traffic changes over a short time span.

Our proposal is to compute the window size $\mathbf{x}_t$ from the actual traffic from C_0 to C_{t-1} by an Online Convex Optimization (OCO) algorithm, instead of solving the static optimization problem with the estimated amount $\hat{\mathbf{b}}_t$. Since the exact amount of queued data at the previous cycle C_{t-1} is observed after the window allocation $\mathbf{x}_{t-1}$, our algorithm can use the actual data $\mathbf{b}_{t-1}$ to decide an allocation for the next cycle C_t .

By defining a convex version of the objective, $g_t(\mathbf{x}_t) \triangleq -f_t(\mathbf{x}_t)$, the maximization problem of f_t is converted to a minimization problem of a convex function g_t over $\mathbf{x}_t$ with the same constraint in Eq. (4). We use Projected Gradient Descent (PGD) [12] to solve this problem. At C_{t-1} , the accurate amounts of queued data are collected at the OLT via REPORT messages (See Figure 2). Thus, the actual objective function at the previous cycle g_{t-1} becomes available (Eq. (4) with $\mathbf{b}_{t-1}$). Note that this approach is different from the traditional use of REPORT messages in that ONUs report the




Fig. 3. Average Latency at ONUs in respect to the cycle size $C = 1$.

actual traffic amount after the allocated window at C_{t-1} starts. PGD computes the next window size $\mathbf{x}_t$ by

$$\mathbf{x}_t \leftarrow \Pi_X(\mathbf{x}_{t-1} - \eta_{t-1} \nabla g_{t-1}(\mathbf{x}_{t-1})), \quad (5)$$

where $\Pi_X(\mathbf{x})$ is a projection bringing $\mathbf{x}_t$ back to the space X , and η_{t-1} is a diminishing learning step size. Since the allocation for the next round is defined as a function of the previous allocation, the recursive relation can be represented as $\mathbf{x}_t \leftarrow \text{PGD}(\mathbf{x}_0, \dots, \mathbf{x}_{t-1})$, which was initially desired.

B. OCO Performance Bound

We are interested in bounding the average performance of a DBA algorithm over time rather than obtaining the exact optimum allocations at each cycle, since the exact optimum solution is not feasible without an oracle reporting the exact future traffic amount $\mathbf{b}_t$. The *regret* of an algorithm A quantifies the discrepancy between a sequence of allocations by A and a hindsight optimum allocation (the most robust allocation over all cycles till C_T):

$$\text{regret}_T(A) = \sup_{\{g_1, \dots, g_T\} \subseteq G} \left\{ \sum_{t=1}^T g_t(\mathbf{x}_t) - \min_{\mathbf{x} \in X} \sum_{t=1}^T g_t(\mathbf{x}) \right\}.$$

Note that the supremum over all possible sequences of objective functions implies that the regret bound considers adversarial scenarios, which is, in our case, a situation where the actual traffic pattern is extremely irregular.

The gradient descent-based solution (PGD), which we use, is known to have $O(\log T)$ -bound for the regret with the convergence rate of $1/\sqrt{T}$ for a convex objective function over a convex variable range [12].

V. EXPERIMENT AND DISCUSSION

A. Network Settings

Simulations are conducted in a network with three slices: S_0, S_1 , and S_2 . The network hosts ten ONUs, and each slice has a different number of ONUs as follows: $\{u_0, u_1, \dots, u_4\} \in S_0$, $\{u_5, u_6, u_7\} \in S_1$, and $\{u_8, u_9\} \in S_2$. Among the ten ONUs, ONUs u_0, u_3, u_6 , and u_9 are heavily loaded with a traffic generation function based on the Poisson distribution $\text{Pois}(\lambda = 10)$, while the traffic for the other ONUs is generated

based on $\text{Pois}(\lambda = 1)$. Unless otherwise specified, the slice prioritization weight p_j is set to 1.0 for all slices. It is only changed when evaluating the slice-awareness.

B. Performance Comparison: Typical DBA Algorithms

The proposed method is compared with typical generic DBA algorithms: MAXWIN and AVGPRED.

1) *MAXWIN*: allocates either the predefined maximum window size m or the traffic amount queued in ONUs, which is reported in the previous REPORT message. Thus, this algorithm is quite responsive to the demand from each ONU. While both $m = 0.2$ and 0.4 are tested in the simulations, only the result with $m = 0.2$ is discussed below, since their performances are similar.

2) *AVGPRED*: assigns the average of the actual traffic queued at each ONU in past cycles. When computing the average at round C_t , the algorithm uses the past traffic amounts from C_{t-1} to C_{t-h} , where h is a given time horizon. The performance is shown with a horizon h that provided the best result among different h 's ($h \in \{10, 100, 1000\}$). This algorithm stays optimum for every static problem defined with a predicated traffic amount. Note that an allocation vector can be projected back to the solution space to make it a feasible solution, preserving the relative allocation ratio among ONUs.

C. Experimental Results

We discuss our simulation results from the three aspects desired for a DBA algorithm; namely, demand-responsiveness, slice-awareness, and allocation fairness. In addition, an interpretation of the OCO performance guarantee in our problem is briefly summarized at the end.

1) *Demand-responsiveness*: Figure 3 illustrates the average latency versus the delta value (Δt) depicted in Figure 2, during which additional traffic unknown to a DBA arrives. This value implies the amount of traffic that was not reported to the OLT and can be added to the queues of ONUs before the next round. The y-axis is the average latency of each traffic unit from 50 simulation runs. Δt and the average latency are represented with respect to the cycle size $C = 1$.

The proposed OCO-based algorithm realizes the lowest queuing delay, on average, among the three methods. This result indicates that our solution reduces the average latency even though it does not provide the optimum for every single static optimization. This is because the OCO-based approach gradually adjusts its solution to an allocation that is robust against underlying fluctuations of traffic.

The general decreasing trend of MAXWIN and AVGPRED in the average latency could be explained by excessive adjustment towards traffic amount in each time cycle. As Δt increases, the total amount of unreported traffic becomes more random. Therefore, the negative impact of an unmatched bandwidth allocation will be smaller in scale *on average*. In contrast, the impact stays notable when Δt is smaller. It is worth mentioning that this behavior does not imply the usefulness of MAXWIN and AVGPRED with a larger Δt . This is because it is reasonable to assume the duration between a




Fig. 4. Transition of the Allocation to ONU u_0 : The proposed DBA algorithm adjusts its allocation over time to provide a robust bandwidth allocation, considering the past traffic pattern.




Fig. 5. Allocation Depending on the Traffic Load at C_{10000} : A heavily-loaded ONU receives more bandwidth.

GATE message and the previous REPORT message (Δt) is within one time cycle C , and our OCO approach performs better in the range.

Figure 4 shows the learning process of the allocation to ONU u_0 by the proposed algorithm along with the observed actual traffic pattern (grey bars) and the allocations of the other algorithms. The result indicates that the OCO-based algorithm quickly adjusts its allocation to the traffic pattern at earlier cycles and continues fine-tuning at later cycles. The result also indicates the convergence of our OCO-based allocation over time.

Figure 5 illustrates the converged allocation to ONU u_0 at C_{10000} by each method. It indicates that the proposed algorithm allocates slightly more resources to normally loaded ONUs. While the difference in the allocation among the three methods is quite small (on a scale of 10^{-3}), the results discussed above show that the small difference in allocation succeeds in reducing the average delay. Since the heavily

TABLE I
ALLOCATED WINDOW SIZES WITH DIFFERENT SLICE PRIORITIZATION WEIGHTS: THE SLICE WEIGHT p_j ALLOWS DELAY-SENSITIVE SLICES TO BE ALLOCATED LARGER WINDOW SIZES. (UNIT: CYCLE)

	Latency-sensitive Slice ($p_j = 1.2$)	Standard Slice ($p_j = 1.0$)
Heavily-loaded ONU	0.2642	0.1657
Normally-loaded ONU	0.0569	0.0169

TABLE II
AVERAGE LATENCY AT EACH ONU AND STANDARD DEVIATION σ_U OF THE DELAYS IN ALL ONUS IN U . (UNIT: CYCLE)

	ONU 0	ONU 1	ONU 5	ONU 6	σ_U
OCO	0.0351	0.0172	0.0169	0.0352	0.0094
MAXWIN	0.0230	0.0304	0.0302	0.0231	0.0039
AVGPRED	0.0231	0.0306	0.0310	0.0230	0.0039

loaded ONUs aggressively request the resources, the other algorithms, which are relatively myopic, tend to allocate more resources towards the ONUs with heavier loads. However, the overallocation could be alleviated in the proposed algorithm due to its gradual allocation updates towards a robust solution.

2) *Slice-awareness*: Table I indicates the difference in the allocated window sizes depending on the difference among network slices with different delay-sensitivity. The slice prioritization weight p_j was adjusted to represent the specific delay requirements. In particular, assuming that slice S_3 hosts a delay-sensitive network service, the weight p_3 of slice S_3 was set to 1.2, while the weights for the other slices are 1.0. It is observed that the ONUs on S_3 receive more bandwidth. For example, the heavily-loaded ONU u_9 on S_3 is allocated 0.2642 cycles at round C_{10000} , while the heavily-loaded ONUs on S_1 and S_2 receive 0.1657 cycles. This result empirically verifies the possibility to accommodate different types of network slices with diverse requirements through the tuning of the parameter p_j .

3) *Fairness*: Table II summarizes the average latency in representative ONUs (u_0, u_1, u_5 , and u_6) and the standard deviation σ_U of the delays across all ONUs in U . The standard deviation of the proposed OCO-based algorithm is relatively higher than the other methods. While this implies the prioritization of some ONUs, the standard deviation stays in the same order (10^{-3}) as the other algorithms. This fact empirically shows that the proposed algorithm maintains the fairness in allocation in terms of the latency, even though it allocates more resources to some ONUs depending on the weight metrics. In other words, the algorithm fairly allocates the resources so that the latency levels stay within a certain range across all ONUs.

VI. CONCLUSION

This paper proposes an Online Convex Optimization (OCO) based solution to the Dynamic Bandwidth Allocation (DBA)

problem. Our algorithm is aimed at realizing an bandwidth allocation that is slice-aware, demand-responsive, and fair among ONUs in PON access networks, formulating the DBA problem based on the concept of proportional fairness with appropriate weight parameters. The use of the OCO scheme enables our allocation solution to be robust against fluctuations of traffic, which eventually results in the reduction of the average latency over time. The simulation results indicate that the proposed solution mitigates the latency compared to other typical allocation approaches. Furthermore, the results infer the effectiveness of a robust allocation over time for access networks with dynamic traffic patterns, in contrast to the optimum solutions at static optimizations formulated with an estimated traffic pattern.

ACKNOWLEDGEMENT

This work was supported in part by the National Science Foundation under Grant No. CNS-2008856.

REFERENCES

- [1] S. Zhang, "An overview of network slicing for 5G," *IEEE Wireless Communications*, vol. 26, no. 3, pp. 111–117, 2019.
- [2] H. Uzawa, K. Honda, H. Nakamura, Y. Hirano, K. Nakura, S. Kozaki, and J. Terada, "Dynamic bandwidth allocation scheme for network-slicing-based TDM-PON toward the beyond-5G era," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 12, no. 2, pp. A135–A143, 2020.
- [3] J. Ahmed, J. Chen, L. Wosinska, B. Chen, and B. Mukherjee, "Efficient inter-thread scheduling scheme for long-reach passive optical networks," *IEEE Communications Magazine*, vol. 51, no. 2, pp. S35–S43, 2013.
- [4] A. Checko, H. L. Christiansen, Y. Yan, L. Scolari, G. Kardaras, M. S. Berger, and L. Dittmann, "Cloud RAN for mobile networks—a technology overview," *IEEE Communications Surveys Tutorials*, vol. 17, no. 1, pp. 405–426, 2015.
- [5] G. O. Pérez, J. A. Hernández, and D. Larrabeiti, "Fronthaul network modeling and dimensioning meeting ultra-low latency requirements for 5G," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 10, no. 6, pp. 573–581, 2018.
- [6] J. Navarro-Ortiz, P. Romero-Díaz, S. Sendra, P. Ameigeiras, J. J. Ramos-Munoz, and J. M. Lopez-Soler, "A survey on 5G usage scenarios and traffic models," *IEEE Communications Surveys Tutorials*, vol. 22, no. 2, pp. 905–929, 2020.
- [7] B. Skubic, J. Chen, J. Ahmed, L. Wosinska, and B. Mukherjee, "A comparison of dynamic bandwidth allocation for EPON, GPON, and next-generation TDM PON," *IEEE Communications Magazine*, vol. 47, no. 3, pp. S40–S48, 2009.
- [8] A. Helmy and A. Nayak, "Energy-efficient decentralized framework for the integration of fog with optical access networks," *IEEE Transactions on Green Communications and Networking*, vol. 4, no. 3, pp. 927–938, 2020.
- [9] J. A. Hatem, A. R. Dhaini, and S. Elbassuoni, "Deep learning-based dynamic bandwidth allocation for future optical access networks," *IEEE Access*, vol. 7, pp. 97 307–97 318, 2019.
- [10] M. Xu, X. Liu, N. Chand, F. Effenberger, and G. Chang, "Flex-frame timing-critical passive optical networks for delay sensitive mobile and fixed access services," in *2017 Optical Fiber Communications Conference and Exhibition (OFC)*, 2017, pp. 1–3.
- [11] A. Helmy and A. Nayak, "Towards more dynamic energy-efficient bandwidth allocation in EPONs," in *2018 IEEE Global Communications Conference (GLOBECOM)*, 2018, pp. 1–6.
- [12] S. Shalev-Shwartz, "Online learning and online convex optimization," *DOI:10.1561/22000000018*, 2012.
- [13] N. Merayo, P. Pavon-Marino, J. C. Aguado, R. J. Durán, F. Burrull, and V. Bueno-Delgado, "Fair bandwidth allocation algorithm for PONS based on network utility maximization," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 9, no. 1, pp. 75–86, 2017.
- [14] D. Hisano and Y. Nakayama, "Two-stage optimization of uplink forwarding order with cooperative DBA to accommodate a TDM-PON-based fronthaul link," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 12, no. 5, pp. 109–119, 2020.
- [15] T. Tashiro, S. Kuwano, J. Terada, T. Kawamura, N. Tanaka, S. Shigematsu, and N. Yoshimoto, "A novel DBA scheme for TDM-PON based mobile fronthaul," in *OFC 2014*, 2014, pp. 1–3.
- [16] J. Mo and J. Walrand, "Fair end-to-end window-based congestion control," *IEEE/ACM Transactions on Networking*, vol. 8, no. 5, pp. 556–567, 2000.