

State-Aware Subcarrier Allocation for SDN-Controlled Filterless Coherent P2MP Network

Original

State-Aware Subcarrier Allocation for SDN-Controlled Filterless Coherent P2MP Network / Zeb, S., Masood, M.U., Ali, A., Dipto, I.C., Schips, R., Pedro, J., Napoli, A., Curri, V.. - (2026). (GLOBECOM 2026 - 2026 IEEE Global Communications Conference Macau (Chi) 7 - 11 December 2026).

Availability:

This version is available at: 11583/3015549 since: 2026-09-15T22:08:26Z

Publisher:

IEEE

Published

DOI:

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

State-Aware Subcarrier Allocation for SDN-Controlled Filterless Coherent P2MP Network

Sanwal Zeb

Dept. of Elec. & Telecom.
Politecnico di Torino, Italy
sanwal.zeb@polito.it

Muhammad Umar Masood

Dept. of Elec. & Telecom.
Politecnico di Torino, Italy
muhammad.masood@polito.it

Ahtisham Ali

Consorzio Top-IX
Politecnico di Torino, Italy
ahtisham.ali@polito.it

Imran Chowdhury Dipto

Dept. of Elec. & Telecom.
Politecnico di Torino, Italy
imran.dipto@polito.it

Riccardo Schips

Dept. of Elec. & Telecom.
Politecnico di Torino, Italy
riccardo.schips@polito.it

João Pedro

Nokia Optical Network
Carnaxide, Portugal
joao.pedro@nokia.com

Antonio Napoli

Nokia Optical Network
Munich, Germany
antonio.napoli@nokia.com

Vittorio Curri

Dept. of Elec. & Telecom.
Politecnico di Torino, Italy
vittorio.curri@polito.it

Abstract—This paper presents a state-aware subcarrier allocation framework for SDN-controlled coherent P2MP networks under time-varying traffic conditions. In contrast to conventional approaches that perform resource allocation independently at each scheduling interval, the proposed framework incorporates backlog evolution and subcarrier reconfiguration cost, enabling time-coupled allocation decisions. The system considers heterogeneous service classes with distinct delay requirements and traffic dynamics. Two policies are evaluated: Reactive State-Aware (RS-SA), which prioritizes instantaneous demand, and Backlog-Stable State-Aware (BS-SA), which emphasizes backlog stability and reduced reconfiguration. Results show that RS-SA achieves higher throughput at the expense of increased reconfiguration overhead, whereas BS-SA reduces control activity and improves stability under high load, with moderate performance degradation. These results highlight the importance of state-awareness in balancing service performance and operational efficiency.

Index Terms—Point-to-Multipoint (P2MP), Digital Subcarrier Multiplexing (DSCM), Software-Defined Networking (SDN), Service-Level Agreements (SLA)

I. INTRODUCTION

Global Internet ecosystems are undergoing sustained and rapid growth, driven by the convergence of broadband connectivity, cloud computing, and AI-enabled applications. Both consumer and enterprise traffic continue to expand, imposing stringent requirements on transport infrastructures in terms of capacity, scalability, and energy efficiency. Recent industry analysis projects that global Internet traffic will increase from approximately 664 EB/month in 2023 to around 3–6 ZB/month by 2033, corresponding to a compound annual growth rate (CAGR) of 18%–27% [1]. This growth is further accelerated by the proliferation of 5G networks and the anticipated evolution toward 6G, alongside increasingly data-intensive applications.

To accommodate these demands, access network architectures are evolving toward converged metro-access infrastructures capable of supporting both mobile transport and fixed broadband services. The densification of 5G deployments significantly increases transport requirements in fronthaul and midhaul segments, where strict latency and syn-

chronization constraints must be satisfied [2–4]. In parallel, optical access technologies based on Passive Optical Networks (PONs) continue to scale through successive standardization efforts, including 50G-PON and beyond, enabling higher per-wavelength data rates and improved spectral efficiency [5, 6]. As a result, metro-access nodes are required to aggregate heterogeneous traffic classes while maintaining high efficiency and flexibility.

In this context, coherent point-to-multipoint (P2MP) optical architectures based on digital subcarrier multiplexing (DSCM) have emerged as a promising solution. DSCM enables a single optical carrier to be partitioned into multiple independently controllable subcarriers (SCs), each operating at fine capacity granularity (e.g., 25 Gb/s), which can be dynamically assigned across access nodes [7–9]. In a hub-leaf topology, SCs are provisioned in the frequency domain, where each subcarrier occupies a distinct spectral slot and can be flexibly mapped across nodes. Fig. 1 illustrates the considered P2MP architecture, where subcarriers generated at the hub are dynamically assigned to multiple leaf nodes. Slight frequency offsets between upstream and downstream SCs enable simultaneous bidirectional operation while avoiding spectral overlap. Advanced digital signal processing (DSP) enables precise subcarrier separation, frequency alignment, and bidirectional transmission through slight spectral offsets [10]. Moreover, DSCM-based P2MP systems can operate over filterless optical infrastructures, similar to PONs, enabling cost-efficient deployment while providing fine-grained resource allocation capabilities [11].

The programmability of such systems is further enhanced through software-defined networking (SDN), which provides centralized control, global network visibility, and dynamic reconfiguration of SC resources. SDN-controlled P2MP architectures enable service-aware allocation of subcarriers, allowing resources to be adapted to time-varying traffic demands across nodes [12, 13]. This capability transforms metro-access transport from a statically provisioned infrastructure into a




Fig. 1. DSCM enabled independent SCs on a single shared wavelength, where upstream and downstream transmissions are separated by slight offsets in central frequency to enable di-directional transmission.

programmable and adaptive system.

However, this flexibility introduces new operational challenges. SC reconfiguration is inherently non-instantaneous, involving both control-plane processes, such as telemetry collection, centralized decision-making, and configuration updates, and data-plane operations including transceiver activation, wavelength tuning, and signal stabilization [14, 15]. These processes introduce non-negligible latency and overhead. While frequent reconfiguration can improve short-term service efficiency, it also increases control-plane load and may degrade system stability, leading to a fundamental trade-off between performance and operational efficiency.

Prior research has extensively investigated SC allocation strategies in coherent P2MP systems. Early studies on dynamic SC allocation demonstrated that adapting SC assignments to instantaneous traffic conditions can significantly improve resource utilization [16]. Subsequent efforts explored centralized and distributed SDN control frameworks, enabling near-real-time adaptation of SC resources [13]. More recently, QoS-aware allocation schemes have been proposed to support heterogeneous service requirements by incorporating service differentiation and performance constraints into the allocation process [17]. In parallel, experimental and network-level studies have validated the feasibility of coherent P2MP systems and analyzed route feasibility under quality-of-transmission (QoT) constraints [18].

Despite these advances, existing approaches largely rely on stateless allocation models, where resource assignment is optimized independently at each scheduling interval. This assumption neglects the inherently temporal nature of metro-access traffic, where demand evolves over time and unmet traffic can persist across intervals as backlog. Furthermore, repeated SC reconfiguration decisions are typically not penalized, leading to excessive switching activity and increased control overhead. As a result, these approaches may achieve short-term performance gains while compromising long-term system stability and efficiency.

This limitation highlights a critical gap: effective SC allocation in coherent P2MP metro-access networks requires a time-coupled, state-aware framework that jointly captures (i) temporal traffic dynamics, including backlog evolution and non-stationary demand, (ii) heterogeneous service requirements with different delay sensitivities, and (iii) the operational cost

associated with SC reconfiguration.

Motivated by these observations, this paper presents a controller-driven, state-aware SC allocation framework for SDN-controlled P2MP metro-access networks under time-varying traffic conditions. The proposed approach leverages network state information, including backlog evolution and prior allocation decisions, to guide SC assignment over time while explicitly accounting for reconfiguration overhead. By balancing service performance and operational stability, the framework enables efficient and robust resource allocation under realistic traffic dynamics.

The main contributions of this work are summarized as follows:

- A state-driven SC allocation framework that adapts resource assignment based on network conditions, including backlog evolution and prior allocation states.
- Explicit modeling of temporal traffic dynamics, enabling the system to capture carry-over demand and non-stationary traffic behavior across scheduling intervals.
- Incorporation of SC reconfiguration overhead into the allocation process, exposing the trade-off between service performance and operational stability.
- Stability-aware allocation policies that balance immediate demand servicing with long-term system efficiency and reduced reconfiguration activity.
- Evaluation under realistic time-varying traffic conditions, demonstrating service performance and stability under stressed load regimes.

II. SYSTEM MODEL AND STATE-AWARE ALLOCATION FRAMEWORK

To model dynamic resource allocation in coherent P2MP metro-access networks, we consider the SDN-controlled architecture illustrated in Fig. 2. The system follows a hub-leaf topology, where a central coherent transceiver at the aggregation node generates multiple digital SCs that are distributed across access nodes [18]. Each SC represents a fine-grained capacity unit that can be flexibly assigned to different service classes. Traffic at each access node is aggregated from heterogeneous sources and mapped to predefined service classes, enabling dynamic capacity sharing across access nodes and traffic services.

SC allocation is governed by a centralized SDN controller that maintains a global view of network conditions, including traffic demand, SC utilization, node-level states, and physical-layer QoT parameters obtained through real-time telemetry. Based on this information, the controller dynamically updates SC-to-service mappings at each access node through programmable southbound interfaces, such as NETCONF/YANG, enabling adaptive configuration of coherent pluggables and switching functions[19]. To reflect practical operation, reconfiguration overhead is incorporated into the model, capturing the impact of SC reassignment on system stability and control-plane efficiency.




Fig. 2. Filterless P2MP converged metro-access architecture.

A. System and Resource Model

We consider a coherent P2MP metro-access system with one hub and a set of access nodes $S = \{1, \dots, S\}$, with $S = 4$. The hub generates $N_{SC}^{\text{hub}} = 16$ Nyquist-shaped subcarriers (SCs), each with capacity $C_{SC} = 25$ Gb/s. Each leaf or access node s is provisioned with $N_{SC}^{(s)} = 4$, yielding:

$$C^{(s)} = N_{SC}^{(s)} \cdot C_{SC} = 100 \text{ Gb/s}. \quad (1)$$

At each time interval t , SCs are assigned to service types $c \in \mathcal{C}$:

$$\sum_c n_c^{(s)}(t) \leq N_{SC}^{(s)}. \quad (2)$$

Traffic from heterogeneous services is modeled at the level of aggregated flows rather than individual packets. Resource allocation decisions are performed at the control-plane time scale, reflecting SDN-based monitoring and reconfiguration processes [12, 19].

Time is divided into fixed scheduling intervals (epochs), indexed as $t = 1, 2, \dots, T$, where each slot corresponds to $\Delta t = 6$ minutes, reflecting traffic aggregation at different time index.

B. Traffic Model

Each access node aggregates traffic from:

$$\mathcal{C} = \{\text{mob}, \text{ent}, \text{res}\}, \quad (3)$$

representing mobile transport, enterprise connectivity, and residential broadband, respectively. The total offered demand at node s and time interval t is:

$$T_{\text{tot}}(s, t) = \sum_{c \in \mathcal{C}} T_c(s, t), \quad (4)$$

where $T_c(s, t)$ denotes the traffic demand of service type c .

The aggregate load evolves over time according to:

$$L_{\text{tot}}(t) = L_0 \cdot \phi(t), \quad (5)$$

where L_0 is the base (mean) load and $\phi(t)$ is a normalized diurnal profile satisfying $\mathbb{E}[\phi(t)] = 1$ and $\max \phi(t) \approx 1.2$ – 1.3 , capturing peak-to-mean traffic variation.

The demand of each service is expressed as:

$$T_c(s, t) = L_{\text{tot}}(t) \cdot x_c(t), \quad (6)$$




Fig. 3. Time-varying traffic demand per service at 100 Gb/s offered load.

where $x_c(t)$ represents the time-varying traffic share of a service.

The service traffic share is defined as:

$$x_c(t) = \frac{\tilde{x}_c(t)}{\sum_{c'} \tilde{x}_{c'}(t)}, \quad \tilde{x}_c(t) = b_c p_c(t) e_c(t) m_c(t), \quad (7)$$

where:

- b_c is the average share of traffic for service c over time,
- $p_c(t)$ captures regular daily patterns (e.g., day–night variations),
- $e_c(t)$ models predictable demand increases (e.g., peak hours or events),
- $m_c(t)$ captures random traffic fluctuations over time.

The stochastic component $m_c(t)$ models time-varying traffic fluctuations, preserving the average load while capturing burstiness and time-correlated traffic variations.

Fig. 3 represents the generated time-varying demand for the three service types at an offered load of 100 Gb/s. The figure highlights distinct diurnal patterns: enterprise connectivity peaks during business hours, residential broadband exhibits strong evening demand, and mobile transport shows relatively smoother variation with moderate peaks. The shaded regions indicate stochastic variability around the mean profile, capturing realistic traffic fluctuations over time.

C. Service, Backlog, and Optimization Model

The assumed P2MP metro-access supports heterogeneous services with distinct traffic characteristics and delay requirements. Such multi-service abstractions are commonly used in metro-access resource allocation to capture QoS differentiation across service domains [17].

TABLE I
HETEROGENEOUS SERVICES CLASS PARAMETERS

Service	Traffic share	Delay tolerance (τ_c)	Weight (w_c)
Mobile (mob)	34.4%	0	1.50
Enterprise (ent)	35%	1	1.15
Residential (res)	30.6%	2	1.00

Table I summarizes the service-specific parameters. The traffic share represents the relative proportion of each service in the aggregate demand. The delay tolerance τ_c defines the maximum number of time intervals that traffic of class c can be deferred before it is considered expired. The weight w_c reflects the relative priority of each service in the allocation

process, with higher values indicating stricter service requirements. At each interval t , the effective demand for service c includes both new arrivals and previously unserved traffic:

$$D_c^{\text{eff}}(t) = T_c(t) + B_c(t), \quad (8)$$

where $T_c(t)$ denotes newly generated demand and $B_c(t)$ represents backlog carried from previous intervals. This formulation explicitly captures temporal traffic dependencies, which are typically neglected in stateless allocation approaches that optimize each interval independently [13].

Given an SC allocation $n_c(t)$, the served traffic is:

$$S_c(t) = \min(D_c^{\text{eff}}(t), n_c(t) C_{\text{SC}}), \quad (9)$$

where C_{SC} is the capacity per SC i.e. 25G.

Backlog evolves as:

$$B_c(t+1) = B_c(t) - S_c^{\text{old}}(t) + C_c(t), \quad (10)$$

where $S_c^{\text{old}}(t)$ denotes the portion of served traffic originating from backlog, and $C_c(t)$ represents newly deferred traffic. Traffic exceeding the delay tolerance τ_c is dropped:

$$D_c^{\text{drop}}(t) = \sum_{a > \tau_c} B_c^{(a)}(t), \quad (11)$$

where $B_c^{(a)}(t)$ denotes backlog with age a .

The SC allocation problem aims to balance service performance with system stability. The optimal allocation is obtained by solving an optimization overall performance score J :

$$J = w_d J_d + w_b J_b + w_f J_f + w_u J_u - w_r J_r, \quad (12)$$

where:

$$\begin{aligned} J_d &= \sum_c w_c S_c \quad (\text{served traffic}), \\ J_b &= \sum_c w_c S_c^{\text{old}} \quad (\text{backlog reduction}), \\ J_f &= \sum_c A_c S_c \quad (\text{fairness}), \\ J_u &= \text{avg. SC utilization per access} \end{aligned}$$

The first four terms improve service performance by increasing served traffic, reducing backlog and deferred traffic, and ensuring fair resource allocation across heterogeneous traffic services. One of the key performance parameter is the reconfiguration cost:

$$J_r = \sum_{\text{SC}} \mathbb{I}\{\text{class}_{\text{SC}}(t) \neq \text{class}_{\text{SC}}(t-1)\}$$

which counts the number of SC reassignments between consecutive time intervals t and $t-1$. Frequent reconfiguration increases control overhead and may impact system stability. By penalizing J_r , the model avoids unnecessary changes and promotes stable operation. This formulation enables a trade-off between improving service performance (e.g., reducing blocking and deferred traffic) and limiting control-plane overhead.

D. State-Aware SC Allocation Policies

Algorithm 1 summarizes the controller-side execution of the proposed *Reactive State-aware SC Allocation (RS-SA)* and *Backlog-Stable State-aware SC Allocation (BS-SA)* policies. The controller evaluates feasible SC layouts, scores each candidate using the objective function, and applies the best allocation subject to reconfiguration constraints. The two policies differ in how the available capacity is assigned between backlog and newly arriving traffic.

Algorithm 1 State-Aware SC Allocations: (RS-SA, BS-SA)

Require: Demand $T_c(t)$, backlog $B_c(t)$, previous layout $\ell(t-1)$, policy π

Ensure: Updated layout $\ell(t)$ and backlog $B_c(t+1)$

- 1: Compute effective demand $D_c^{\text{eff}}(t)$ for all services.
- 2: Generate feasible SC layouts $\mathcal{L}$.
- 3: **for** each layout $\ell \in \mathcal{L}$ **do**
- 4: Compute allocated capacity and served traffic.
- 5: **if** $\pi = \text{RS-SA}$ **then**
- 6: Split served traffic proportionally between backlog and new demand.
- 7: **else if** $\pi = \text{BS-SA}$ **then**
- 8: Serve backlog first, then use remaining capacity for new demand.
- 9: **end if**
- 10: Compute objective score and reconfiguration cost.
- 11: Apply hysteresis, max-reconfiguration, and holding-time constraints.
- 12: **end for**
- 13: Select the best feasible layout.
- 14: Update backlog and drop traffic exceeding SLA limits.
- 15: **return** $\ell(t)$, $B_c(t+1)$

RS-SA policy prioritizes instantaneous demand by proportionally serving both backlog and new traffic, resulting in higher responsiveness and improved throughput. In contrast, the BS-SA policy prioritizes backlog clearance and limits frequent reconfiguration, leading to more stable operation with reduced control overhead.

III. RESULTS AND DISCUSSION

This section evaluates the proposed state-aware subcarrier allocation under time-varying traffic for heterogeneous services. Two policies are compared: RS-SA, which prioritizes instantaneous demand, and BS-SA, which emphasizes backlog stability. Performance is analyzed across increasing load using traffic fate and system behavior metrics.

Fig. 4 shows the normalized traffic fate as a function of offered load for enterprise, mobile, and residential services under RS-SA and BS-SA policies. At moderate load, most traffic is successfully served across all services. As load increases, served traffic decreases while deferred and dropped traffic increase, with different behaviors across service types. Mobile traffic shows no deferral due to strict delay constraints, and unmet demand is directly dropped as load increases. In




Fig. 4. Service outcome versus offered load for RS-SA and BS-SA.

contrast, enterprise and residential services allow buffering, leading to increasing deferred traffic under congestion. Under RS-SA, more traffic is served at moderate load due to aggressive reallocation, but this results in higher service-level-agreement (SLA) expiration at high load. BS-SA, instead, retains more traffic in backlog, reducing SLA violations while increasing deferred traffic. Overall, the figure highlights the trade-off between immediate service (RS-SA) and backlog stability (BS-SA) under increasing load.




Fig. 5. Backlog dynamics versus offered load: normalized backlog, clearance ratio, and blocking ratio.

Fig. 5 shows the evolution of normalized backlog, backlog clearance ratio, and current blocking ratio with increasing offered load. As load increases, backlog grows for both policies, with BS-SA exhibiting consistently higher backlog due to its backlog retention strategy. In contrast, RS-SA maintains lower backlog levels by prioritizing immediate service. This behavior is reflected in the backlog clearance ratio, where BS-SA achieves higher clearance across all load regimes, indicating more effective servicing of accumulated demand. However, this comes at the expense of higher blocking, where BS-SA exhibits a steeper increase in blocking ratio compared to RS-SA. Overall, the figure highlights the trade-off between backlog stability and immediate service: BS-SA improves backlog clearance and stability, while RS-SA reduces blocking by favoring instantaneous traffic servicing.

Fig. 6 shows SLA compliance as a function of offered load for enterprise, mobile, and residential services. At moderate load, all services achieve near-perfect compliance under both

policies. As load increases, SLA compliance degrades for enterprise and residential traffic due to backlog accumulation and delayed servicing, consistent with the increase in deferred traffic observed in Fig. 4. BS-SA maintains higher SLA compliance than RS-SA for enterprise traffic at high and saturation loads, as its backlog-aware strategy reduces delay violations. However, this improvement is less pronounced for residential traffic, where SLA compliance degrades similarly under both policies due to longer delay tolerance and sustained backlog accumulation. In contrast, RS-SA exhibits a sharper decline in enterprise SLA due to its emphasis on instantaneous demand. Mobile traffic maintains 100% SLA compliance across all load levels due to its hard-drop behavior, where unmet demand is immediately discarded rather than delayed.



Fig. 6. SLA compliance versus offered load for RS-SA and BS-SA across service classes.

Fig. 7 shows the average number of subcarrier reconfigurations per hour as a function of offered load. Reconfiguration activity increases with load for both policies, reflecting higher adaptation requirements under traffic dynamics. However, RS-SA consistently incurs higher reconfiguration rates, with a sharp increase in the high and near-saturation regimes due to its reactive adaptation to instantaneous demand. In contrast, BS-SA exhibits a more gradual increase by maintaining allocation consistency and limiting unnecessary changes, resulting in reduced control-plane overhead and improved operational stability.

Fig. 8 summarizes the trade-off between service performance and normalized reconfiguration rate, where the service




Fig. 7. Average subcarrier reconfiguration rate versus offered load for RS-SA and BS-SA.




Fig. 8. Service performance versus normalized reconfiguration rate for RS-SA and BS-SA.

performance score is computed using the objective function in Eq. 12. RS-SA operates in a higher performance regime but at the cost of significantly increased reconfiguration, particularly at high loads. BS-SA, on the other hand, achieves lower reconfiguration rates with faster performance degradation under congestion. The two policies define distinct operating regions rather than converging to a single optimal point, highlighting the inherent trade-off between performance and stability. Overall, these results confirm that incorporating backlog-awareness and reconfiguration cost enables controlled tuning between throughput, stability, and operational overhead in coherent P2MP metro-access networks.

IV. CONCLUSION

This paper presented a state-aware subcarrier allocation framework for SDN-controlled coherent P2MP metro-access networks under time-varying traffic. By incorporating backlog evolution and reconfiguration cost, the proposed approach captures the trade-off between service performance and operational stability. Results show that RS-SA achieves higher throughput but incurs up to $\sim 40\text{--}50\%$ higher reconfiguration rates at high load, while BS-SA reduces reconfiguration activity by $\sim 30\text{--}40\%$ and improves backlog clearance, at the

cost of $\sim 10\text{--}15\%$ performance degradation under congestion. These results highlight the importance of temporal awareness in resource allocation. Overall, the framework enables controlled tuning between performance and stability, making it suitable for practical deployments. Future work will consider larger-scale scenarios, adaptive learning-based policies, and integration of Agentic-AI for efficient resource allocation.

ACKNOWLEDGMENTS

Funded by the European Union's Horizon Europe research and innovation programme under the Marie Skłodowska-Curie Doctoral Network (MSCA-DN NESTOR, G.A. No. 101119983), (EWOC, G.A. No. 101073265), (AL-LEGRO, G.A. 101092766) in collaboration with Nokia Optical Networks, (Germany & Portugal).

REFERENCES

- [1] Nokia, "EMEA broadband index report 2025," Nokia White Paper, <https://www.nokia.com/newsroom/>.
- [2] M. Wang *et al.*, "SDN East-West Cooperation in a Converged Fixed-Mobile Optical Access Network: Enabling 5G Slicing Capabilities," *Journal of Optical Communications and Networking*, vol. 14, no. 7, pp. 540–549, 2022.
- [3] G. Simon *et al.*, "Optical Access Evolutions Towards SDN and Disaggregated Hardware: An Operator Perspective," *Journal of Optical Communications and Networking*, vol. 14, no. 6, pp. C57–C69, 2022.
- [4] A. Ali *et al.*, "Statistical analysis of end-to-end route feasibility in converged metro-access optical networks," *IEEE Photonics Technology Letters*, pp. 1–1, 2026.
- [5] P. Torres-Ferrera *et al.*, "Overview of High-Speed TDM-PON Beyond 50 Gbps per Wavelength Using Digital Signal Processing," *Journal of Optical Communications and Networking*, vol. 14, no. 12, pp. 982–996, 2022.
- [6] D. Zhang, D. Liu, and D. Nessel, "Progress of Very High Speed PON in ITU-T," *Journal of Optical Communications and Networking*, vol. 18, no. 2, pp. A112–A122, 2025.
- [7] D. Welch, A. Napoli *et al.*, "Point-to-Multipoint Optical Networks Using Coherent Digital Subcarriers," *Journal of Lightwave Technology*, vol. 39, no. 16, pp. 5232–5247, 2021.
- [8] D. Welch, A. Napoli *et al.*, "Digital Subcarrier Multiplexing: Enabling Software-Configurable Optical Networks," *Journal of Lightwave Technology*, vol. 41, no. 4, pp. 1175–1191, 2023.
- [9] A. Ali *et al.*, "Evaluating metro-access optical networks as-a-service for flexible ran x-haul," in *2025 IEEE Conference on Standards for Communications and Networking (CSCN)*, 2025, pp. 1–6.
- [10] T. Duthel *et al.*, "DSP design for coherent optical point-to-multipoint transmission," *Journal of Lightwave Technology*, vol. 42, no. 3, pp. 1109–1118, 2023.
- [11] M. M. Hosseini *et al.*, "Optimized Design of Filterless Horseshoe Networks Exploiting P2MP Coherent Transceivers," *Journal of Optical Communications and Networking*, vol. 15, no. 9, pp. 569–578, 2023.
- [12] A. Sgambelluri *et al.*, "Fully Disaggregated Optical Network Control with ONOS," *Journal of Optical Communications and Networking*, vol. 13, no. 1, pp. A138–A147, 2021.
- [13] H. Shakespear-Miles *et al.*, "Centralized and Distributed Control of P2MP Systems," *Journal of Optical Communications and Networking*, vol. 16, no. 5, pp. 565–576, 2024.
- [14] R. Alvizu *et al.*, "Comprehensive Survey on Transport SDN," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 4, pp. 2232–2283, 2017.
- [15] M. Radovic *et al.*, "SDN-Controlled Open RAN X-Haul with P2MP Transceivers," *Journal of Optical Communications and Networking*, vol. 17, no. 11, pp. E109–E116, 2025.
- [16] H. Shakespear-Miles *et al.*, "Dynamic Subcarrier Allocation for Multipoint Optical Connectivity," in *OECN/PSC*, 2022.
- [17] P. Soumplis *et al.*, "QoS-aware Resource Allocation in P2MP Metro Networks," in *OFC*, 2025.
- [18] S. Zeb *et al.*, "AI-driven converged metro-access optical network-as-a-service with point-to-multipoint coherent optics for 6G X-Hauling," *Computer Networks*, vol. 284, p. 112338, 2026.
- [19] A. Rafel *et al.*, "Self-configuring metro-access network with OpenXR Pluggables and filter-less OADMs," *Journal of Optical Communications and Networking*, vol. 17, no. 6, pp. 543–554, 2025.