POLITECNICO DI TORINO
Repository ISTITUZIONALE

Stability-Aware Subcarrier Allocation in Filterless P2MP Metro-Access Networks

Original
Stability-Aware Subcarrier Allocation in Filterless P2MP Metro-Access Networks / Zeb, S., Masood, M.U., Ali, A., Dipto, I.C., Schips, R., Pedro, J., Napoli, A., Curri, V.. - (2026). (2026 International Conference on Photonics in Switching and Computing (PSC) Valencia (Spa) 15-17 September 2026).

Availability:
This version is available at: 11583/3015546 since: 2026-09-15T21:35:55Z

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)

30 October 2026

Stability-Aware Subcarrier Allocation in Filterless P2MP Metro-Access Networks

Sanwal Zeb[1], Muhammad Umar Masood[1], Ahtisham Ali[1,2], Imran Chowdhury Dipto[1], Riccardo Schips[1], João Pedro[3], Antonio Napoli[4], Vittorio Curri[1]

[1]Department of Electronics and Telecommunications, Politecnico di Torino, Torino, Italy
[2]Consorzio TOP-IX, Torino, Italy [3]Nokia, Optical Networks, Carnaxide, Portugal [4]Nokia, Optical Networks, Munich, Germany

Author e-mail address: sanwal.zeb@polito.it

Abstract:

We propose stability-aware subcarrier allocation for DSCM-based P2MP coherent metro-access networks. Three policies are evaluated under RAN, enterprise, and PON traffic with differentiated SLAs. The proposed SAA matches minimum blocking while reducing reconfigurations by 25–50×.

Keywords: *Coherent optical systems, Software-defined optical networking (SDON), Optical transceivers*

I. Introduction

Optical access networks are evolving from single-service residential broadband toward converged multi-service architectures. This is driven by factors such as B5G/6G small cell densification, enterprise cloud connectivity, and increasing residential bandwidth demands [1–3]. Access nodes are now required to simultaneously aggregate radio access network (RAN) fronthaul, enterprise (ENT.), and passive optical network (PON) traffic [4], each with distinct service level agreement (SLA) requirements [5, 6]. While widely deployed, legacy TDM-PON systems exhibit limitations in converged scenarios, as their uniform scheduling provides limited flexibility across service types, and dynamic bandwidth assignment may introduce delays that can affect latency-sensitive traffic such as RAN fronthaul [7–9].

Coherent point-to-multipoint (P2MP) pluggable transceivers based on digital subcarrier multiplexing (DSCM) address these limitations by serving multiple leaf nodes from a single hub using independently addressable subcarriers (SCs), each supporting 25 G with 16-QAM modulation [2]. Similar to PONs they leverage a filterless architecture with passive splitter/combiners rather than ROADMs. These transceivers are managed through software-defined networking (SDN) [10–12], where the controller performs network monitoring, resource management, and configuration [13, 14].

Prior work on DSCM-based P2MP has addressed filterless network design [15], dynamic SC activation for energy efficiency [3], and centralized ILP-based SC allocation for single-class traffic [16]. More recent contributions include QoS-aware MILP formulations for multi-class scheduling [17] and distributed multi-agent approaches for per-interval throughput optimization [18], both focusing on instantaneous demand adaptation. However, these approaches treat each time interval independently, without considering the cost of reassigning SCs across consecutive intervals. Achieving low blocking under bursty traffic requires frequent SC reassignments, increasing control-plane reconfiguration overhead [14] due to repeated southbound configuration updates at both hub and leaf transceivers, which introduces controller-induced latency. This exposes a trade-off between blocking performance and SC reconfigurations, which remains largely unexplored.

In this article, we address this trade-off by proposing three SC allocation policies: (i) Priority-based Class Allocation (PCA) as an SLA-aware baseline, (ii) Minimum-Blocking Allocation (MBA) as a blocking-optimal reference, and (iii) Stability-Aware Allocation (SAA), which penalizes unnecessary SC reassignments across consecutive intervals. Under bursty RAN, ENT., and PON traffic across four access/leaf nodes, SAA achieves blocking performance comparable to MBA (∼9% at 95 G offered load) while reducing SC reconfigurations.

II. Methodology

Fig. 1 shows the considered DSCM-based P2MP metro-access architecture, organized in three planes. At the data plane, a 1×400 G P2MP hub transceiver at the aggregation node serves $N = 4$ access/leaf nodes, each equipped with a 100 G leaf transceiver, interconnected through a passive splitter/combiner in a filterless configuration [3, 15]. The hub supports $K = 16$ SCs with 16-QAM modulation, each with $B = 25$ G capacity per SC, with $K_s = 4$ number of SCs per leaf [19]. A centralized network controller manages per-SC configuration of hub and leaf transceivers, while the proposed allocation policies (PCA, MBA, and SAA) run as controller applications that determine SC assignment.

At each time interval t, traffic demands are collected via the northbound interface, and SC-to-class assignments are computed and enforced via southbound configuration. Reassigning SCs between consecutive intervals triggers reconfiguration at both hub and leaf transceivers, incurring control-plane latency [14]. Each access node aggregates traffic from three classes $c \in \mathscr{C} = \{\text{RAN}, \text{ENT.}, \text{PON}\}$. The total offered load per leaf at interval t is given by $T_{\text{tot}}(s,t) = T_{\text{RAN}}(s,t) + T_{\text{ENT.}}(s,t) + T_{\text{PON}}(s,t)$, where $T_c(s,t)$ denotes the offered traffic of class c at leaf node s and interval t. Traffic shares are set to 40% (RAN), 35% (ENT.), and 25% (PON), representing a fronthaul-dominant scenario. SLA weights are

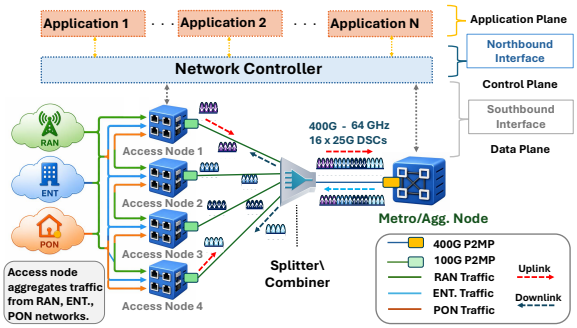


Fig. 1. Filterless P2MP converged metro-access network

defined as $w_{\text{ENT.}} = 1.3$, $w_{\text{RAN}} = 1.1$, and $w_{\text{PON}} = 0.7$, reflecting stricter requirements for RAN and enterprise traffic than for best-effort PON traffic.

Traffic varies over time due to the bursty nature of access traffic. At each interval, $T_c(s,t)$ follows a uniform distribution around a base load $T_c^{\text{base}}(s)$ with class-dependent burst amplitudes $a_{\text{ENT.}} = 5\%$, $a_{\text{RAN}} = 10\%$, and $a_{\text{PON}} = 20\%$. These values show that enterprise traffic is stable, while PON traffic is highly variable, with RAN in between. An additional global variation of $\pm 10\%$ is applied to the total traffic T_{tot}. The per-leaf base load ranges from 50 G to 100 G. Physical-layer feasibility, including OSNR margins and SC power equalization, are validated according to [15]. SC allocation is implemented at the controller using the following policies for multi-service traffic classes (e.g., RAN, PON, and ENT.).

Priority-based Class Allocation (PCA) is an SLA-aware baseline that assigns SCs in a fixed priority order (ENT. $>$ RAN $>$ PON). At each interval t, SCs are sequentially allocated by fully serving higher-priority classes before considering lower-priority ones, subject to the per-leaf SC budget. PCA is memoryless, i.e., assignments at t are independent of $t-1$, and may starve low-priority traffic under high traffic load. **Minimum-Blocking Allocation (MBA)** minimizes blocking by assigning each SC to the class with the highest weighted residual demand. Let $d_{s,c}(t)$ denote the unmet demand of class c at leaf node s. For each SC k, MBA selects:

$$c^\star = \arg\max_{c \in \mathscr{C}} \ (w_c \cdot d_{s,c}(t),) \tag{1}$$

MBA adapts to instantaneous demand but is memoryless, allowing frequent SC reassignments when demand ranking changes across intervals.

Stability-Aware Allocation (SAA) extends MBA by incorporating reconfiguration awareness to jointly minimize blocking and SC reassignments. For each SC k, SAA selects:

$$c^\star = \arg\max_{c \in \mathscr{C}} \left(w_c \cdot d_{s,c}(t) - \beta \, [c \neq c_{\text{prev}}(s,k)] \right) \tag{2}$$

where $c_{\text{prev}}(s,k)$ denotes the class assigned at $t-1$, and $[c \neq c_{\text{prev}}(s,k)]$ equals 1 for SC reassignment and 0 otherwise. The penalty β acts as a threshold: reassignment occurs only if the gain in weighted demand exceeds β. When $\beta = 0$, SAA reduces to MBA, while large β suppresses reassignment, approaching PCA-like behavior. SAA follows MBA, but avoids changing an SC unless the gain exceeds the penalty β.

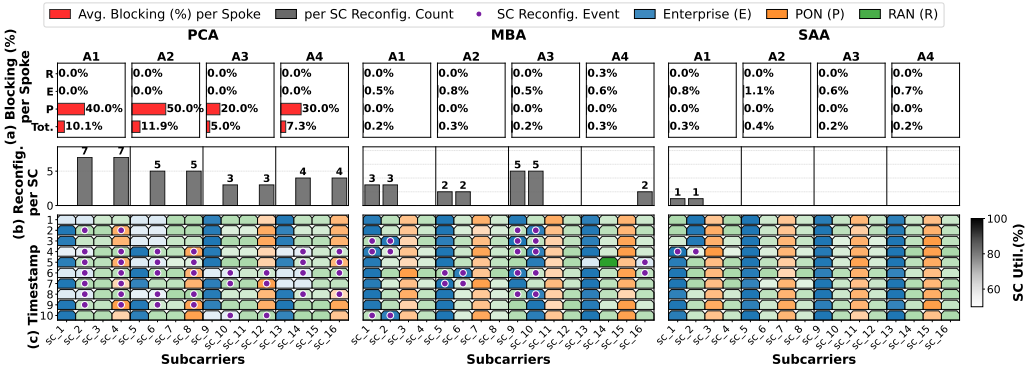


Fig. 2. SC allocation results at 70 G offered load: **row (a):** avg. blocking (%) per class per $(A_1 - A_4)$; **row (b):** tot. per-SC reconfig.; row **(c):** per interval SC allocation across $(A_1 - A_4)$ where color indicates traffic class, shading reflects per-SC capacity util. (%).

This introduces temporal consistency at the SC level by considering the previous assignment. As a result, small traffic variations do not trigger reassignment, since the gain in weighted demand remains below β. In contrast, when traffic changes are large enough, they exceed β and trigger SC reassignment to avoid blocking. In this work, β is set to $0.15B$, where B is the SC capacity i.e. 25 G, achieving MBA-like blocking while significantly reducing reconfiguration frequency.

III. Results and Discussion

Figs. 2 and 3 show results at 70 G and 95 G offered load per leaf (A_1–A_4), comparing PCA, MBA, and SAA in terms of (a) per-class blocking, (b) per-SC reconfigurations, and (c) per-SC allocation across 10 intervals, where color denotes traffic class, shading indicates SC utilization (%), and markers represent reconfiguration events. At moderate-to-high load (70 G, Fig. 2), PCA blocks (20–50)% of PON traffic across leaf nodes by prioritizing ENT. and RAN, resulting in (5–12)% total blocking. MBA and SAA both keep total blocking below 0.4%, but MBA triggers 2–5 reconfigurations per leaf across 10 intervals as traffic variations alter class priorities, whereas SAA triggers at most one reconfiguration at a single leaf node. This is confirmed by the heatmaps in row (c), where MBA exhibits frequent reconfiguration events in contrast to stable SAA allocation.

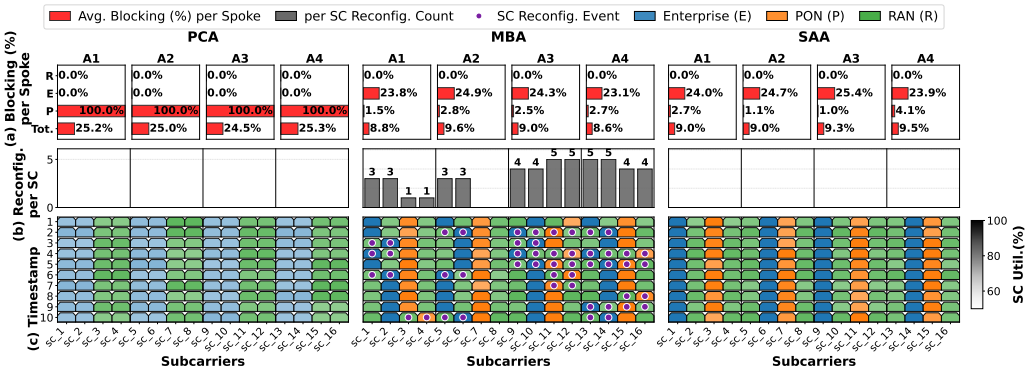


Fig. 3. SC allocation results at 95 G offered load: **row (a):** avg. blocking (%) per class per ($A_1 - A_4$); **row (b):** tot. per-SC reconfig.; row **(c):** per interval SC allocation across ($A_1 - A_4$) where color indicates traffic class, shading reflects per-SC capacity util. (%).

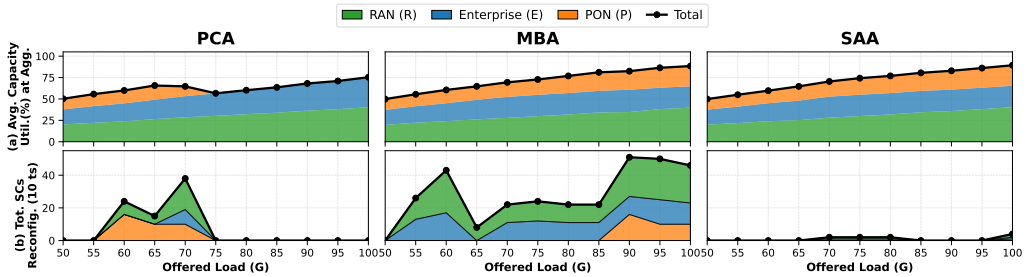


Fig. 4. Allocation policies performance at the hub/agg. node for different offered loads: **row (a):** avg. capacity utilization (%) across 10 timestamps; **row (b):** tot. SCs reconfigurability across all access nodes.

At high offered load (95 G, Fig. 3), PCA blocks 100% of PON traffic across all leaf nodes, resulting in (∼25%) total blocking, as prioritized ENT. and RAN demand nearly exhaust the available SC capacity. In contrast, MBA and SAA reduce total blocking to (∼9%), distributing it across ENT. (∼24%) and PON (∼1–4)% while maintaining 0% RAN blocking. However, MBA triggers 1–5 reconfigurations per SC across 10 intervals due to small variations in instantaneous demand. In contrast, SAA achieves zero reconfigurations across all leaf nodes (A_1–A_4). Under high load, the relative demand across traffic classes remains largely unchanged across intervals. As a result, the SC allocation stays consistent.

Fig. 4 shows (a) hub capacity utilization (%) and (b) total SC reconfigurations over 50–100 G offered load across all time intervals. MBA and SAA increase utilization (%) from ∼50% to ∼90%, whereas PCA plateaus at ∼70% beyond 75 G due to PON starvation. Reconfiguration trends exhibit three load regimes: (i) low load (50–55 G), (ii) moderate load (60–75 G), and (iii) high load (>75 G). (i) At low load (50–55 G), all policies produce few reconfigurations as demand remains within capacity. (ii) At moderate load (60–75 G), PCA peaks at (∼40) near 70 G as demand variations push SCs across class boundaries, while MBA dips around (60–65 G) where capacity margins absorb bursts, before rising as the SC budget tightens. Overall, SAA achieves (25–50)× fewer SC reconfigurations than MBA while maintaining comparable blocking and utilization. These results demonstrate that SAA achieves near-MBA blocking performance while ensuring stable SC allocation across offered load regimes.

IV. Conclusion

We proposed three controller-driven SC allocation policies for filterless P2MP metro-access networks serving converged RAN, ENT., and PON traffic. The proposed SAA matches MBA blocking (∼9% at 95 G) and hub utilization (∼90%) while reducing SC reconfigurations by (25–50)× and lowering PON blocking from 100% under PCA to (1–4)%. These results show that penalizing SC reassignments effectively balances blocking and SC reconfiguration for multi-traffic allocation.

Acknowledgment

Funded by 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), and (ALLEGRO, G.A. 101092766).

REFERENCES

[1] D. Welch, A. Napoli, J. Bäck, W. Sande, J. Pedro, F. Masoud, A. Chase, T. Duthel, C. Fludger, H. Sun, S. J. Hand, T.-K. Chiang, M. Plantare, A. Mathur, T. A. Eriksson, M. Olson, S. Voll, and K.-T. Wu, "Point-to-multipoint optical networks using coherent digital subcarriers," *Journal of Lightwave Technology*, vol. 39, no. 16, pp. 5232–5247, 2021.

[2] D. Welch, A. Napoli, J. Bäck, S. Buggaveeti, C. Castro, A. Chase, X. Chen, V. Dominic, T. Duthel, T. A. Eriksson *et al.*, "Digital subcarrier multiplexing: Enabling software-configurable optical networks," *Journal of Lightwave Technology*, vol. 41, no. 4, pp. 1175–1191, 2023.

[3] A. Napoli, C. Castro, P. Torres-Ferrera, M. M. Hosseini, G. Parisi, D. Hillerkuss, J. Pedro, J. Bäck, A. Rashidinejad, F. Masoud, S. J. Hand, S. Erkilinc, T. Duthel, and C. Fludger, "Towards truly scalable sustainable flexible optical networks," in *Proc. European Conference on Optical Communications (ECOC)*, 2025.

[4] G. Simon, P. Chanclou, M. Wang, D. Abgrall, and D. Minodier, "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.

[5] S. Sarmiento, J. M. Delgado Mendinueta, J. A. Altabás, S. Spadaro, S. Shinada, H. Furukawa, J. J. Vegas Olmos, J. A. Lázaro, and N. Wada, "High capacity converged passive optical network and RoF-based 5G+ fronthaul using 4-PAM and NOMA-CAP signals," *Journal of Lightwave Technology*, vol. 39, no. 2, pp. 574–580, 2021.

[6] M. Wang, G. Simon, L. Anet Neto, I. Amigo, L. Nuaymi, and P. Chanclou, "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.

[7] D. Zhang, D. Liu, X. Wu, and D. Nesset, "Progress of itu-t higher speed passive optical network (50g-pon) standardization," *Journal of Optical Communications and Networking*, vol. 12, no. 10, pp. D99–D108, 2020.

[8] M. Ruffini, A. Ahmad, S. Zeb, N. Afraz, and F. Slyne, "Virtual DBA: virtualizing passive optical networks to enable multi-service operation in true multi-tenant environments," *Journal of Optical Communications and Networking*, vol. 12, no. 4, pp. B63–B73, 2020.

[9] P. Torres-Ferrera, F. Effenberger, M. S. Faruk, S. J. Savory, and R. Gaudino, "Overview of high-speed TDM-PON beyond 50 gbps per wavelength using digital signal processing [invited tutorial]," *Journal of Optical Communications and Networking*, vol. 14, no. 12, pp. 982–996, 2022.

[10] M. Channegowda, R. Nejabati, and D. Simeonidou, "Software-defined optical networks technology and infrastructure: Enabling software-defined optical network operations [invited]," *Journal of Optical Communications and Networking*, vol. 5, no. 10, pp. A274–A282, 2013.

[11] D. Kreutz, F. M. V. Ramos, P. E. Veríssimo, C. E. Rothenberg, S. Azodolmolky, and S. Uhlig, "Software-defined networking: A comprehensive survey," *Proceedings of the IEEE*, vol. 103, no. 1, pp. 14–76, 2015.

[12] R. Alvizu, G. Maier, N. Kukreja, A. Pattavina, R. Morro, A. Capello, and C. Cavazzoni, "Comprehensive survey on T-SDN: Software-defined networking for transport networks," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 4, pp. 2232–2283, 2017.

[13] R. Muñoz, L. Gifre, W. Akbar, R. Vilalta, R. Martínez, and R. Casellas, "Dynamic SDN Control of Pluggable Transceivers in disaggregated packet-optical networks for 6G Xhaul," in *2025 25th Anniversary International Conference on Transparent Optical Networks (ICTON)*. IEEE, 2025, pp. 1–4.

[14] M. Radovic, A. Sgambelluri, F. Cugini, R. Kapuscinski, M. Hosseini, S. Parker, R. Berozashvili, F. Paolucci, A. Napoli, and N. Sambo, "SDN-controlled open RAN X-haul with point-to-multipoint transceivers on a horseshoe network," *Journal of Optical Communications and Networking*, vol. 17, no. 11, pp. E109–E116, 2025.

[15] M. M. Hosseini, J. Pedro, A. Napoli, N. Costa, J. E. Prilepsky, and S. K. Turitsyn, "Optimized design of filterless horseshoe networks exploiting point-to-multipoint coherent transceivers," *Journal of Optical Communications and Networking*, vol. 15, no. 9, pp. 569–578, 2023.

[16] H. Shakespear-Miles, M. Ruiz, A. Napoli, and L. Velasco, "Dynamic subcarrier allocation for multipoint-to-point optical connectivity," in *Proc. OptoElectronics and Communications Conference / International Conference on Photonics in Switching and Computing (OECC/PSC)*, 2022, pp. 1–4.

[17] P. Soumplis, K. Christodoulopoulos, K. Yiannopoulos, and E. Varvarigos, "Qos-aware resource allocation in point to multipoint (p2mp) metro aggregation networks," in *2025 Optical Fiber Communications Conference and Exhibition (OFC)*, 2025, pp. 1–3.

[18] H. Shakespear-Miles, Q. Lin, S. Barzegar, M. Ruiz, X. Chen, and L. Velasco, "Centralized and distributed approaches to control optical point-to-multipoint systems near-real-time," *Journal of Optical Communications and Networking*, vol. 16, no. 5, pp. 565–576, 2024.

[19] S. Zeb, A. Ali, I. C. Dipto, A. Rosso, M. U. Masood, R. Schips, R. Ambrosone, S. Straullu, F. Aquilino, A. Nespola, J. Pedro, A. Napoli, A. Galardini, and V. Curri, "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.