Access-Aware Launch Power Optimization for RAN Fronthaul in Converged Metro-Access Optical Networks

Original
Access-Aware Launch Power Optimization for RAN Fronthaul in Converged Metro-Access Optical Networks / Ali, A., Masood, M.U., Zeb, S., Rosso, A., Malik, G., Galardini, A., Curri, V.. - (2026). (2026 Italian Conference on Optics and Photonics (ICOP) L'Aquila (Ita) 15 – 17 June 2026).

Availability:
This version is available at: 11583/3015550 since: 2026-09-15T22:16:27Z

Publisher:
IEEE

Published
DOI:

Access-Aware Launch Power Optimization for RAN Fronthaul in Converged Metro-Access Optical Networks

Ahtisham Ali*†, Muhammad Umar Masood†, Sanwal Zeb†, Andrea Rosso†, Gulmina Malik†,
Alessandro Galardini*, Vittorio Curri†
*Consorzio TOP-IX, Torino, Italy
†Dipartimento di Elettronica e Telecomunicazioni, Politecnico di Torino, Italy
Email: ahtisham.ali@top-ix.org

***Abstract*—Converged metro–access optical networks for 6G x-haul exhibit per-leaf access-loss variations among radio units (RUs), causing non-uniform received optical power at the splitter under conventional uniform launch power provisioning. This work investigates access-aware Tx-power optimization for RAN fronthaul over a realistic converged metro–access topology with multiple RUs, using an experimentally characterized 400 Gb/s DP-16QAM Phoenix DCO whitebox transceiver. Three launch-power provisioning strategies are evaluated under joint BER and 3GPP functional split 7.2 latency constraints: Uniform Launch Power (ULP), Per-leaf Adaptive (PLA), and Conservative Provisioning (CP). GSNR analysis identifies an optimum launch power of ∼+1 dBm. While ULP provides the highest GSNR, PLA achieves a uniform splitter optical power across leaves at the cost of ∼0.7 dB GSNR penalty caused by increased NLI. These results highlight the trade-off between maximizing physical-layer margin and achieving a uniform splitter power profile.**

***Index Terms*—Converged metro-access networks, RAN fronthaul, launch power optimization, physical-layer feasibility, GNPy**

I. Introduction

The evolution from 5G toward 6G is reshaping the requirements imposed on transport networks, where ultra-low latency, deterministic performance, and fine-grained service flexibility have become first-order design constraints. Global mobile traffic is projected to continue its rapid growth, with a compound annual growth rate (CAGR) of 15.4% by 2030, driven by emerging applications such as immersive extended reality, distributed AI workloads, and ultra-reliable low-latency communications [1–3]. To support such diverse and demanding services, the radio access network (RAN) is increasingly disaggregated into radio units (RUs), distributed units (DUs), and central units (CUs), enabling flexible function placement and virtualized deployments across geographically distributed sites [4]. This architectural shift, however, transfers the complexity from the radio domain to the underlying optical transport layer, which must now jointly satisfy strict latency budgets, quality-of-transmission (QoT) guarantees, and capacity scalability across the fronthaul (FH, RU–DU) and midhaul (MH, DU–CU) segments [5]. Considering functional split 7.2 as defined by 3GPP, the FH segment is constrained to a one-way latency budget below 250 µs, together with strict bit-error-rate (BER) requirements imposed by the adopted modulation format [5, 6].

In this context, converged metro–access optical networks have emerged as a key enabler for 6G x-haul, leveraging shared optical infrastructure to aggregate access traffic while providing the reach and capacity required for disaggregated RAN transport [7–9]. The convergence of metro and access domains, however, introduces a heterogeneous physical-layer environment in which access fiber spans of different lengths feed into a common metro segment through splitters, circulators, and reconfigurable optical add-drop multiplexers (ROADMs). As a consequence, each RU experiences a per-leaf access loss that varies across the network in realistic urban deployments. Under conventional uniform per-RU launch-power provisioning, this loss imbalance directly translates into a non-uniform received optical power at the splitter (ROPS), which in turn affects nonlinear interference (NLI) accumulation, ROADM passband penalty, and downstream amplifier and receiver dimensioning. This effect, hereafter referred to as the access-induced power imbalance, becomes a structural feature of converged metro–access x-haul that any practical transport strategy must explicitly address.

Previous works on converged metro–access x-haul have mainly focused on route-level feasibility under joint BER and latency constraints, considering different transceiver classes, modulation formats, and DU placement strategies [9–13], including hybrid scenarios introducing hollow-core fiber [9, 14]. In parallel, DSCM-based point-to-multipoint architectures have been proposed to improve spectral utilization at the DU side [15]. However, in all these studies, the per-RU launch power has been treated as a fixed system-level parameter, leaving the impact of the launch-power provisioning strategy on QoT and on the uniformity of the splitter power profile largely unexplored.

In this work, we address this gap by investigating access-aware Tx-power optimization for RAN fronthaul over a realistic urban topology comprising nine access clusters with eight RU-serving leaves each (72 RUs in total), using an experimentally characterized Phoenix DCO whitebox transceiver at 400 Gb/s with DP-16QAM and 64 Gbaud per channel [16, 17]. Under joint BER and FS 7.2 fronthaul latency constraints,

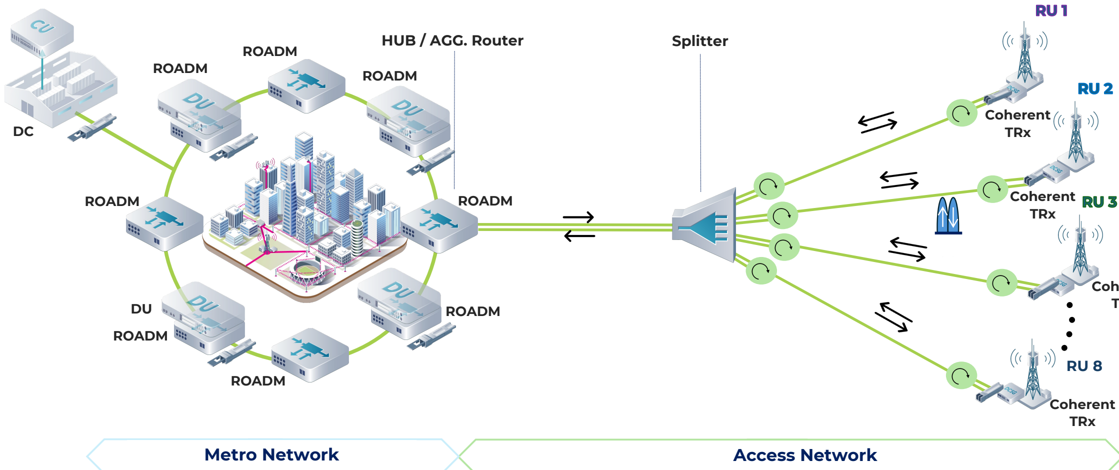

Fig. 1: Converged metro–access optical network architecture.

three per-RU Tx-power strategies are compared: Uniform Launch Power (ULP), Conservative Provisioning (CP), and Per-leaf Adaptive (PLA), formally defined in Section II. The strategies are characterized in terms of per-leaf GSNR margin, ASE/NLI decomposition, and splitter power uniformity, providing a physical-layer-grounded comparison rather than a purely connectivity-based feasibility assessment.

II. Modeling and Simulation Framework

The reference architecture, illustrated in Fig. 1, represents a converged metro–access optical infrastructure delivered as Optical Network-as-a-Service (ONaaS) for 6G x-haul. DU functions are hosted at metro Points-of-Presence (PoPs), while CUs are centralized at one or more data centers (DCs) interconnected to the metro segment. Each RU is connected to its serving metro node through an access leaf composed of a short-reach SSMF span, a splitter/MUX stage, and a circulator enabling bidirectional (BiDi) transmission over a single fiber.

The full metro–access topology used as the reference deployment is shown in Fig. 2 and corresponds to a representative urban operator-grade scenario, consistent with our previous works on converged metro–access x-haul [9, 11], with link lengths ranging from $5\,\mathrm{km}$ to $20\,\mathrm{km}$ and a network diameter of 9 hops ($122\,\mathrm{km}$). In the present work, the analysis specifically focuses on the access segment connecting each RU to its serving metro DU node, analyzing the role of the per-leaf access loss and of the per-RU launch-power strategy on the optical signal delivered to the splitter. The physical-layer performance is evaluated using a Python-based simulator built on top of the GNPy [18, 19], accounting for fiber attenuation, ASE noise, and NLI. The transceiver contribution is captured through an experimentally characterized Phoenix DCO white-box operating at $400\,\mathrm{Gbit\,s^{-1}}$ with DP-16QAM at $64\,\mathrm{Gbaud}$ and $75\,\mathrm{GHz}$ channel spacing, following the methodology in [16, 17]. For each FH path, the end-to-end SNR combines optical-channel and transceiver contributions as $\mathrm{SNR}^{-1} = \mathrm{GSNR}^{-1} + \mathrm{SNR}_{\mathrm{TRx}}^{-1}$, and the pre-FEC BER is computed assuming DP-16QAM as $\mathrm{BER} = (3/8)\,\mathrm{erfc}(\sqrt{\mathrm{SNR}/10})$, with a feasibility threshold of 10^{-2} [9, 11]. Latency is evaluated from the physical path length assuming $5\,\mathrm{\mu s\,km^{-1}}$ for SSMF, and a route is feasible when both BER and the FS 7.2 fronthaul latency budget of $250\,\mathrm{\mu s}$ are satisfied.

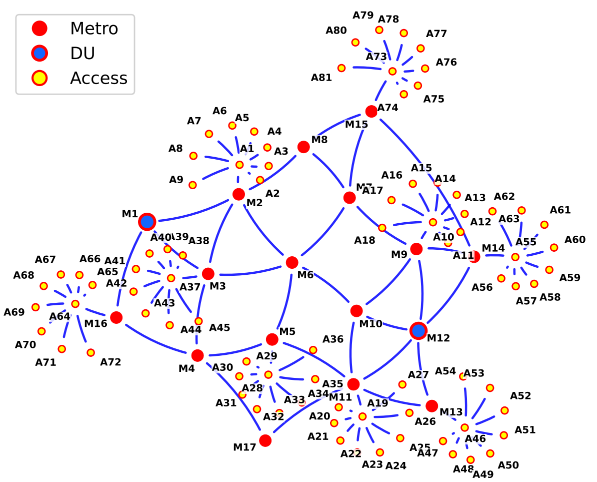

Fig. 2: Converged metro-access optical network topology

Across the 72 RUs of the considered topology, the per-leaf access loss L_i falls in the $11\,\mathrm{dB}$ to $13\,\mathrm{dB}$ range, depending on the access SSMF span and on circulator and splitter contributions. This variation translates, under conventional uniform provisioning, into a non-uniform splitter optical power, which is the access-induced power imbalance addressed in this work. Three per-RU Tx-power strategies are compared under joint BER and FS 7.2 latency constraints, formally defined in Algorithm 1: Uniform Launch Power (ULP), Per-leaf Adaptive (PLA), and Conservative Provisioning (CP). The strategies are evaluated in terms of per-leaf and per-cluster GSNR margin, ASE–NLI decomposition of the noise contributions versus

Algorithm 1 Per-RU launch-power assignment under the three Tx-power strategies

Require: Set of RUs $\mathcal{R} = \{1, \ldots, N\}$; per-leaf access loss L_i, $\forall i \in \mathcal{R}$; target Rx power $P_{\mathrm{Rx}}^{\mathrm{tgt}} = -10$ dBm; strategy $s \in \{\mathrm{ULP, PLA, CP}\}$
Ensure: Per-RU launch power $P_{\mathrm{Tx}}^{(i)}$, $\forall i \in \mathcal{R}$
$L_{\max} \leftarrow \max_{j \in \mathcal{R}} L_j$
Compute $P_{\mathrm{Tx}}^{(i)}$ **according to strategy** s**:**
if $s = \mathrm{ULP}$ **then**
  // Uniform Launch Power
  $P_{\mathrm{Tx}}^{(i)} \leftarrow 0$ dBm, $\forall i \in \mathcal{R}$
else if $s = \mathrm{PLA}$ **then**
  // Per-leaf Adaptive
  $P_{\mathrm{Tx}}^{(i)} \leftarrow P_{\mathrm{Rx}}^{\mathrm{tgt}} + L_i$, $\forall i \in \mathcal{R}$
else if $s = \mathrm{CP}$ **then**
  // Conservative Provisioning
  $P_{\mathrm{Tx}}^{(i)} \leftarrow P_{\mathrm{Rx}}^{\mathrm{tgt}} + L_{\max}$, $\forall i \in \mathcal{R}$
end if
return $\{P_{\mathrm{Tx}}^{(i)}\}_{i \in \mathcal{R}}$

launch power, and uniformity of the splitter optical power across the leaves of each cluster.

III. Results and Discussion

Fig. 3 reports the per-leaf power at the splitter as a function of the per-RU launch power, for one representative access cluster of the considered topology. The three Tx-power strategies generate three distinct operational signatures. Under ULP, all leaves share the same launch power (0 dBm) but exhibit a ∼1 dB spread in ROPS, directly reflecting the per-leaf access-loss variation across the cluster. CP shifts the entire launch-power distribution up by $L_{\max}$, leaving the ROPS spread unchanged. PLA is the only strategy that delivers a uniform ROPS at the target $P_{\mathrm{Rx}}^{\mathrm{tgt}} = -10$ dBm across all the leaves, at the price of a per-leaf launch-power distribution that follows the access loss profile of the cluster ($1.9\,\mathrm{dB}$ to $3.15\,\mathrm{dB}$). ULP and CP propagate the loss imbalance to the splitter, while PLA absorbs it at the RU side.

To explain this behavior, Fig. 4 decomposes the inverse GSNR into its ASE and NLI contributions as a function of launch power, together with the operating points of the three provisioning strategies. The ASE scales as -1 dB/dB

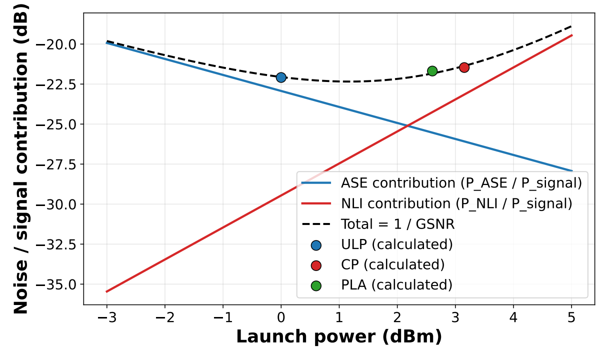

Fig. 4: ASE and NLI contributions to $1/\mathrm{GSNR}$ vs. launch power.

and the NLI term as $+2$ dB/dB, consistent with launch-power-independent ASE accumulation and with Kerr-induced P^2 scaling, respectively [19]. The GSNR is maximized at a per-route launch-power optimum of approximately $+1$ dBm, corresponding to the minimum of the total $1/\mathrm{GSNR}$ curve. Given the -1 dB/dB and $+2$ dB/dB slopes, this optimum is reached when $P_{\mathrm{ASE}} = 2\,P_{\mathrm{NLI}}$, and therefore sits ∼1 dB below the ASE/NLI crossing point (∼+2 dBm) rather than at it. ULP, with its 0 dBm setpoint, operates just below this optimum and therefore minimizes the NLI penalty. PLA and CP, with launch powers in the $1.9\,\mathrm{dBm}$ to $3.15\,\mathrm{dBm}$ range, sit above the optimum and pay a ∼0.7 dB NLI penalty with respect to ULP. The system is therefore operating in the NLI-limited regime for both PLA and CP, and the additional launch power used to equalize (PLA) or to over-provision (CP) the splitter ROPS is paid back by a corresponding NLI penalty.

The consequence at the physical-layer level is shown in Fig. 5, reporting the per-leaf GSNR versus launch power for the same cluster, and in Fig. 6, summarizing the per-leaf Tx and ROPS profiles as a function of the access leaf distance. ULP exhibits the highest GSNR margin in all leaves (with a minimum margin of ∼2.8 dB above the DP-16QAM feasibility threshold), followed by PLA and CP, which behave similarly within ∼0.1 dB and trail ULP by the expected ∼0.7 dB NLI penalty. The same picture is consistently ob-

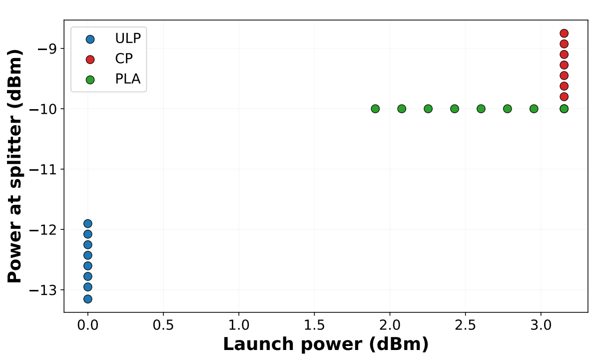

Fig. 3: Per-leaf optical power at the splitter vs. launch power.

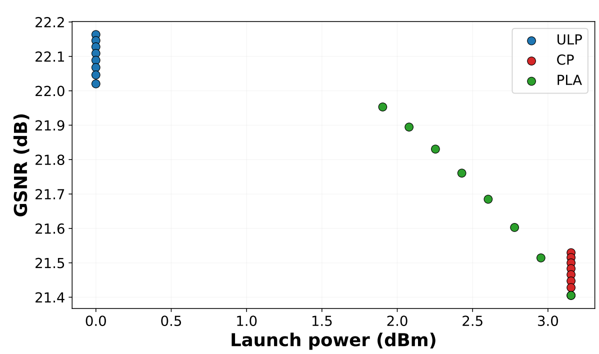

Fig. 5: Per-leaf GSNR vs. launch power under the three Tx-power strategies for the same representative cluster.

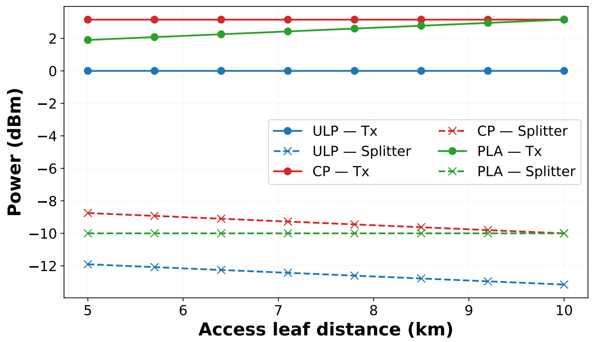

Fig. 6: Per-leaf Tx and splitter optical power as a function of the access leaf distance, under the three strategies.

served across the other clusters of the topology (not shown for brevity), confirming that the launch-power vs. received-power-uniformity trade-off is the key design consideration of the converged metro–access access segment, rather than a cluster-specific effect. From a deployment perspective, the choice reflects a structural launch-power vs. received-power-uniformity trade-off: ULP maximizes physical-layer margin, CP guarantees the worst-ROPS for the leafs at the price of over-provisioning, and PLA uniquely delivers a flat splitter power profile at the cost of a ∼0.7 dB NLI penalty relative to ULP. This uniformity property, while not quantified in the access-only analysis presented here, has direct implications for downstream metro-segment behavior, including DSCM-based per-subcarrier power balancing, ROADM passband behavior, and downstream amplifier and receiver dimensioning, which will be the subject of future work.

IV. Concluding Remarks

This work investigated access-aware Tx-power optimization for RAN fronthaul in a converged metro–access optical network, addressing the received power imbalance caused by non-uniform access losses across RUs. Three per-RU launch-power provisioning strategies (ULP, PLA, CP) were compared on an experimentally characterized 400 Gb/s DP-16QAM Phoenix DCO whitebox transceiver under joint BER and FS 7.2 fronthaul latency constraints. A GSNR decomposition into ASE and NLI contributions identified a per-route launch-power optimum of ∼+1 dBm, with the system operating in the NLI-limited regime above it. ULP maximizes the GSNR margin but propagates the access-loss imbalance to the splitter. Among the evaluated strategies, PLA achieves a uniform splitter optical power across all access leaves, paying a ∼0.7 dB NLI penalty matched by CP. The choice is therefore a structural launch-power vs. received-power-uniformity trade-off, rather than a connectivity-driven feasibility one.

Future work will extend the analysis from the access segment to the full converged metro–access infrastructure, characterizing the end-to-end RU–DU–CU behavior under the three strategies and quantifying the impact of the splitter power profile on ROADM passband penalty, NLI accumulation along the metro spans, and downstream amplifier and receiver dimensioning. The uniformity property delivered by PLA is also a natural enabler for DSCM-based deployments, where per-subcarrier power balancing at the splitter directly affects the optical comb shape and the per-SC received power across the cluster; the integration of access-aware per-RU control with DSCM per-SC provisioning will be the subject of a dedicated study.

Acknowledgments: Ahtisham Ali's work is fully funded by the EU's Horizon 2020 research and innovation programme under the MSCA project EWOC under GA No. 101073265. This work has also received support from the MSCA project NESTOR (101119983), the EU Horizon Project ALLEGRO (101092766), and the project PNRR-NGEU (MUR–DM117/2023).

References

[1] Ericsson, "Mobile Traffic Forecasts, Ericsson Mobility Report," Tech. Rep., 2025, accessed: 2025-08-29. [Online]. Available: https://www.ericsson.com/en/reports-and-papers/mobility-report/dataforecasts/mobile-traffic-forecast
[2] C.-X. Wang *et al.*, "On the Road to 6G: Visions, Requirements, Key Technologies, and Testbeds," *IEEE Commun. Surv. & Tutor.*, vol. 25, no. 2, pp. 905–974, 2023.
[3] P. A. Baziana, "Optical Data Center Networking: A Comprehensive Review on Traffic, Switching, Bandwidth Allocation, and Challenges," *IEEE Access*, vol. 12, pp. 186 413–186 444, 2024.
[4] A. Tzanakaki *et al.*, "Optical Networking in Support of 5G and 6G Infrastructures," in *Proc. Int. Conf. Opt. Netw. Design and Model. (ONDM)*, 2025, pp. 1–3.
[5] 3GPP, "Study on CU-DU lower layer split for NR; (Release 15)," 3GPP, Tech. Rep. TR 38.816 V15.0.0 (2017-12), 2017.
[6] FLEX-SCALE Consortium, "Deliverable D2.1 – 6G network requirements," Tech. Rep., 2023. [Online]. Available: https://6g-flexscale.eu/en/outcomes
[7] M. Klinkowski, "Latency-aware DU/CU placement in convergent packet-based 5G fronthaul transport networks," *Applied Sciences*, vol. 10, no. 21, p. 7429, 2020.
[8] G. Rizzelli *et al.*, "Analysis and experimental demonstration of possible architectures for future coherent metro+PON converged networks [Invited]," *J. Opt. Commun. Netw.*, vol. 17, no. 2, pp. A142–A154, 2025.
[9] A. Ali *et al.*, "Optimizing RAN X-Haul Performance Through Targeted Hollow-Core Fiber Deployment in Converged Metro-Access Networks," in *Proc. Int. Conf. on Softw., Telecommun. and Comput. Netw. (SoftCOM)*, 2025, pp. 1–6.
[10] A. Ali *et al.*, "Evaluating Metro-Access Optical Networks as-a-Service for Flexible RAN X-haul," in *Proc. IEEE Conf. Standards for Commun. and Netw. (CSCN)*, 2025, pp. 1–6.
[11] A. Ali *et al.*, "Statistical Analysis of end-to-end Route Feasibility in Converged Metro–Access Optical Networks," *IEEE Photon. Technol. Lett.*, pp. 1–1, 2026.
[12] A. Ali *et al.*, "Optimized Distributed Unit (DU) Placement under Physical Constraints for 6G x-Haul," in *Proc. Int. Conf. Transparent Opt. Netw. (ICTON)*, 2026, pp. 1–4.
[13] 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.
[14] A. Ali *et al.*, "Benefits of Selective Hollow-Core Fiber Deployment for 6G X-Haul in Converged Metro-Access Networks," in *Proc. Int. Conf. Opt. Netw. Design and Model. (ONDM)*, 2026.
[15] D. Welch *et al.*, "Digital subcarrier multiplexing: Enabling software-configurable optical networks," *J. Lightwave Technol.*, vol. 41, no. 4, pp. 1175–1191, 2023.
[16] A. Rosso, "Modeling and controlling optical transponder white-boxes based on the physical layer digital twin," Master's thesis, PoliTo, 2024. [Online]. Available: https://webthesis.biblio.polito.it/id/eprint/31825
[17] T. Mano *et al.*, "Modeling the Input Power Dependency of Transceiver BER-OSNR for QoT Estimation," in *Proc. Opt. Fiber Commun. Conf. (OFC)*, 2024, p. M1H.4.
[18] TIP, "GNPy: Open Optical Performance Modeling Tool," [Online]. Available: https://github.com/Telecominfraproject/oopt-gnpy, 2019.
[19] V. Curri, "GNPy model of the physical layer for open and disaggregated optical networking [Invited]," *J. Opt. Commun. Netw.*, vol. 14, no. 6, pp. C92–C104, 2022.