AI-driven converged metro-access optical network-as-a-service with point-to-multipoint coherent optics for 6G X-Hauling

Computer Networks 284 (2026) 112338

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

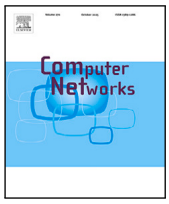

AI-driven converged metro-access optical network-as-a-service with point-to-multipoint coherent optics for 6G X-Hauling☆

Sanwal Zeb [a,*], Ahtisham Ali [a,b], Imran Chowdhury Dipto [a], Andrea Rosso [a], Muhammad Umar Masood [a], Riccardo Schips [a], Renato Ambrosone [a], Stefano Straullu [c], Francesco Aquilino [c], Antonino Nespola [c], João Pedro [d,f], Antonio Napoli [e], Alessandro Galardini [b], Vittorio Curri [a]

[a] Department of Electronics & Telecommunication (DET), Politecnico di Torino, Torino, Italy
[b] Consorzio TOP-IX, Torino, Italy
[c] LINKS Foundation, Torino, Italy
[d] Nokia, Optical Networks, Carnaxide, Portugal
[e] Nokia, Optical Networks, Munich, Germany
[f] Instituto de Telecomunicações, Instituto Superior Técnico, Lisboa, Portugal

ARTICLE INFO

Keywords:
Software Defined Network (SDN)
Network Function Virtualization (NFV)
Digital Subcarrier Multiplexing (DSCM)
Point-to-multipoint (P2MP)
Converged metro-access
Optical network-as-a-Service (ONaaS)
Coherent pluggable optics

ABSTRACT

Future 6G X-haul networks must satisfy strict latency and service reliability requirements, placing significant pressure on metro-access transport architectures. As deployments become denser, longer and more heterogeneous routes intensify physical-layer impairments, making feasibility assurance and Quality-of-Transport (QoT) evaluation increasingly complex. To address these challenges, this work proposes an AI-driven converged metro-access Optical Network-as-a-Service (ONaaS) architecture based on coherent Point-to-Multipoint (P2MP) transmission using Digital Subcarrier Multiplexing (DSCM). An experimentally characterized transceiver impairment model is embedded into a network-level simulator to perform end-to-end feasibility analysis under strict latency and BER constraints. The results show that connectivity is primarily bounded by accumulated impairments, while Distributed Unit (DU) densification improves performance mainly by shortening path lengths, with limited benefit beyond moderate routing depth. To enable scalable operation, a lightweight machine-learning-based BER estimator is developed for rapid QoT prediction. Trained on a minimal deployment scenario, the Random Forest model generalizes across DU densities and topologies with $R^2 > 0.98$, reducing evaluation time by several orders of magnitude. A techno-economic assessment further indicates up to 75% reduction in DU-site transceivers and 27%–30% energy savings compared to Point-to-Point (P2P) provisioning, demonstrating the efficiency and scalability of AI-enabled P2MP metro-access convergence for 6G.

1. Introduction

Future 6G networks will require highly flexible optical-wireless integration to support diverse and latency-sensitive applications. Transparent optical networks are expected to form the backbone of 6G transport by providing high-capacity, ultra-low-latency, and energy-efficient connectivity between distributed Radio Units (RUs), edge computing nodes, and centralized cloud infrastructures. Data-intensive applications such as extended reality, distributed AI processing, industrial automation, and autonomous systems impose strict Service-Level-Agreements (SLAs) in terms of latency, reliability, and throughput. Technologies including Wavelength-Division Multiplexing (WDM), advanced coherent modulation formats, photonic integration, and Software Defined Network (SDN) are therefore expected to enable terabit-per-second (Tbps) and ultra-reliable connectivity for such services [1, 2].

As mobile access networks evolve toward 6G, transport infrastructures must support tighter latency and throughput requirements, particularly under lower-layer functional splits such as split 7-2, while maintaining scalability and cost efficiency [3]. The adoption of Terahertz (THz) and sub-THz wireless interfaces further intensifies these demands by enabling high-capacity radio links that require equally capable optical aggregation networks [4]. In this context, converged

☆ This article is part of a Special issue entitled: 'Reconfigurable Networks' published in Computer Networks.

* Correspondence to: Politecnico di Torino, Corsa Duca Degli Abruzi 29, 10129, Torino, Italy.
E-mail address: sanwal.zeb@polito.it (S. Zeb).

https://doi.org/10.1016/j.comnet.2026.112338
Received 7 October 2025; Received in revised form 4 March 2026; Accepted 23 April 2026
Available online 27 April 2026

metro-access architectures are emerging as a practical solution for integrating Radio Access Network (RAN) deployments.

Within such architectures, optical transport plays a central role in interconnecting distributed RUs, edge processing nodes, and Central Units (CUs). Traditional P2P optical interconnect architectures allocate dedicated wavelengths and interfaces for each connection. Although operationally straightforward, this approach scales linearly with network densification, leading to increased TRX counts, fragmented spectrum utilization, and limited flexibility in handling asymmetric traffic patterns [5]. In particular, dedicated wavelength provisioning often results in spectral under-utilization when traffic demands are granular or dynamically varying, thereby reducing overall spectral efficiency in metro-access segments. These limitations become critical in 6G scenarios, where strict latency and throughput requirements must be satisfied under rapidly evolving and heterogeneous traffic conditions.

Building upon these limitations, P2MP coherent transmission offers a flexible and spectrally efficient alternative. By generating multiple digitally independent subcarriers within a single optical wavelength, it enables fine-grained bandwidth allocation, improves spectral utilization, and efficiently supports hub-and-spoke traffic aggregation in converged metro-access networks [6]. The feasibility of coherent P2MP transmission using digital subcarriers has been experimentally demonstrated in DSCM-based configurations [7]. However, these studies primarily focus on advanced DSP techniques and physical-layer transmission performance, and do not evaluate deployment feasibility in metro-access networks or consider higher-layer transport constraints such as latency and BER requirements.

Furthermore, a reconfigurable converged metro-access architecture based on coherent DSCM technology, referred to as the "MAIN" framework, has been proposed to enable flexible optical-wireless integration [8]. While this work outlines the architectural concept and highlights its techno-economic potential, it does not provide an end-to-end feasibility evaluation under strict functional split requirements. In particular, the impact of DU placement, routing diversity, and accumulated physical-layer impairments on scalable metro-access deployment is not systematically investigated.

While subcarrier-level reconfigurability significantly improves bandwidth utilization and spectral efficiency, it also increases control and orchestration complexity. Each subcarrier may experience distinct QoT conditions depending on path length, accumulated impairments, and network load. Consequently, parameters such as modulation format, launch power, and bandwidth allocation may require independent configuration per subcarrier through a centralized controller. To manage this complexity, network softwarization becomes essential. SDN enables centralized and programmable control of subcarrier-level resources, supporting automated provisioning and dynamic bandwidth adaptation [9]. Network Function Virtualization (NFV) further complements this approach by allowing control and monitoring functions to be deployed as virtualized services on commodity hardware, facilitating flexible and vendor-agnostic operation [10]. However, SDN/NFV-based reconfigurability also increases the need for rapid and scalable QoT estimation mechanisms capable of supporting real-time decision making within an ONaaS architecture.

Building upon the limitations described, this work investigates the feasibility and scalability of converged metro-access networks supporting 6G X-haul under stringent split 7-2 latency and BER constraints. An experimentally characterized DSCM-enabled coherent transceiver impairment model is integrated into a network-level simulation. This enables QoT evaluation and analysis of the impact of DU density and placement, routing feasibility, path-length distribution, and impairment accumulation across structurally distinct topologies. To enable scalable feasibility evaluation while reducing computational complexity, a lightweight ML-based QoT estimator is developed for efficient BER prediction, demonstrating cross-topology generalization when trained on minimal network dataset. Finally, a techno-economic assessment compares coherent P2MP aggregation with conventional P2P under long-term traffic growth, highlighting improvements in spectral efficiency, number of hardware required, port utilization, and operational efficiency.

The remainder of this paper is organized as follows. Section 2 reviews RAN architectures and the role of DSCM-based TRX in converged metro-access networks. Section 3 describes the experimental characterization, simulation framework, and ML-based QoT estimation model. Section 4 presents the numerical results and feasibility analysis. Finally, Section 5 concludes the work and outlines future directions.

2. Metro-access convergence

5G mobile networks come with numerous advances aimed at supporting ultra-reliable low-latency communication (uRLLC), enhanced mobile broadband, and massive machine-type communication. One of the most highlighted innovations is the disaggregation and virtualization of the RAN, which allows a flexible architecture composed of CUs, DUs, and RUs [11]. These functional components can be geographically separated and interconnected through high capacity transport segments, namely fronthaul (RU-DU) and midhaul (DU-CU), which increasingly rely on optical fiber networks due to their ability to meet high capacity and latency requirements [12]. This virtualized and disaggregated approach enables operators to dynamically allocate processing tasks across the network, optimize resource usage, and support multi-vendor interoperability. Metro networks will play a key role in enabling high-capacity, low-latency RAN by providing the critical fiber-based backbone that connects distributed edge computing nodes and centralized cloud resources. This ability to support high-bandwidth, low-latency transport, coupled with technologies such as network slicing, WDM, and software-defined networking, will ensure seamless front-haul and mid-haul connectivity for 5G/6G deployments. By integrating with cloud-native architectures and edge nodes, metro networks can help meet the stringent demands of emerging applications such as augmented reality, industrial automation, and reliable communications [13].

2.1. Functional split

Standardized functional splits, such as Option 2 (CU-DU split) defined by 3GPP and Option 7-2x (DU-RU split) promoted by the O-RAN Alliance, play a key role in determining transport capacity and latency requirements [12,14]. Higher-layer splits reduce transport bandwidth demand and relax latency constraints but require increased processing at the antenna site. In contrast, lower-layer splits enable tighter coordination across the protocol stack to support advanced RAN features, at the cost of significantly higher bandwidth consumption and stringent latency requirements.

According to 3GPP specifications [15], splits in the range of Option 8 to Option 6 require end-to-end latency below $250\,\mu s$, whereas Option 2 can tolerate latencies between $1.5\,ms$ and $10\,ms$. Split 8 maximizes centralization but imposes extremely demanding transport requirements. Split 7-2 slightly increases RU complexity while significantly reducing the required fronthaul bandwidth, thereby emerging as the most practical solution for modern deployments. The required any-haul bandwidth further depends on radio bandwidth, OFDM configuration, Multiple Input Multiple Output (MIMO) layers, antenna ports, modulation and coding schemes, and peak air-interface data rates. Considering a downlink functional split 7-2 with 100 MHz radio bandwidth, 64-QAM modulation, 32 antenna ports, 8 MIMO layers, 713.9 Mbps scheduling/control signaling, IQ bitwidth of 2×16 bits, 14 OFDM symbols, and 1200×5 radio subcarriers, the peak fronthaul bandwidth is approximately 22.2 Gb/s [12,14]. This capacity requirement can be effectively fulfilled by leveraging the fine granularity of subcarrier-based transport, where a single subcarrier provides up to 25 Gb/s capacity.

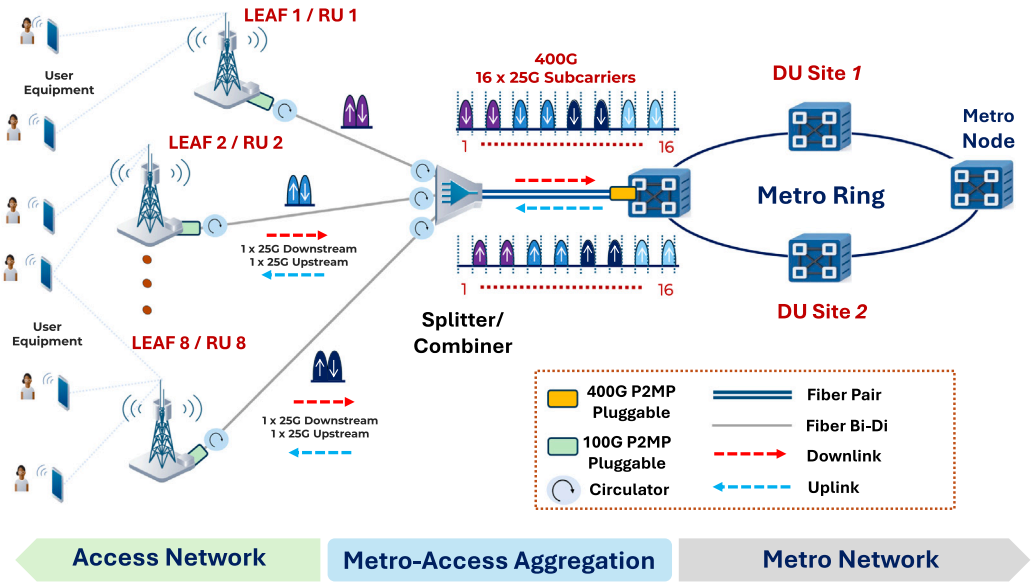

Fig. 1. P2MP converged metro-access network using DSCM coherent pluggable.

Future 6G networks are expected to further extend RAN disaggregation, increasing fronthaul bandwidth demand and transport heterogeneity. Dynamic functional splits, diverse radio configurations, and heterogeneous cell deployments-ranging from multi-sector macro cells to small cells with varying bandwidth and split options significantly complicate any-haul provisioning. Under such conditions, conventional P2P fiber provisioning becomes economically and operationally inefficient in dense B5G and 6G scenarios due to dedicated wavelength allocation and limited spectral flexibility. These challenges motivate the adoption of shared and dynamically managed transport infrastructures, where DSCM-based coherent P2MP aggregation enables spectrally efficient, granular, and scalable X-haul transport aligned with split 7-2 bandwidth requirements.

2.2. Converged metro-access using DSCM -based coherent pluggable

The rapid evolution of technologies such as IoT, edge computing, B5G/6G is driving new demands for underlying optical infrastructure. Service providers now require innovative solutions that not only deliver high capacity but also reduce power consumption, latency, and CAPEX, while enhancing scalability and service flexibility and efficiency. This is critical in fast growing metro/access and mobile front-haul networks, where Intensity Modulation Direct-Detection (IM-DD) and coherent technologies are the main options. While coherent offers significantly greater capacity, reach, and scalability, both systems are typically deployed in P2P architectures. These P2P architectures are inefficient, requiring multiple aggregation stages and struggling to handle the large networks whereas DSCM enables a coherent P2MP solution.

In Fig. 1, a coherent P2MP optical network architecture is depicted, leveraging DSCs for effective high-bandwidth distribution from a central Hub (i.e., 400G) to numerous leaf nodes (i.e., 25G each). The setup assumes to employ 16 DSCs at the hub with 16QAM-400G dual polarization (DP) modulation format [16,17]. Each DSC can be individually modulated and assigned to different leaf nodes, enabling dynamic bandwidth allocation [6,7]. The Splitter/Combiner acts as a passive optical node, eliminating the need for active electronics in the distribution segment, which divides the 400G optical signal into 16 × 25G subchannels, allocating 2 × 25G DSCs to each leaf or RU for bidirectional (Bi-Di) traffic. A circulator enables Bi-Di communication over a single fiber, separating upstream and downstream signals from the respective remote leaf to allow full-duplex transmission. Using higher modulation, it maximizes fiber utilization by packing more data into the same spectrum while enabling finer granularity in resource allocation. In the considered implementation, a single 400G wavelength is divided into 16 digital subcarriers using DP-16QAM modulation format, each supporting up to 25 Gb/s. Considering the 3GPP split 7-2, the required fronthaul capacity per RU is approximately 22.2 Gb/s, which closely matches the payload capacity of one DSC, allowing the maximum spectral utilization. This enables an efficient one-to-one mapping between fronthaul traffic demands and individual subcarriers, maximizing spectral utilization while avoiding bandwidth fragmentation or over-provisioning within the 400G optical transport window.

3. Methodology

3.1. Experimental testbed

Fig. 2 illustrates the B2B experimental setup used to characterize the intrinsic performance of the DSCM-enabled coherent P2MP TRX under controlled conditions. A Layer-2 switch serves as the central management and control entity, interfacing with hardware and controller via DCN-GE links. Real-time telemetry enables continuous monitoring of signal integrity, while an Optical Spectrum Analyzer (OSA) evaluates spectral characteristics and assesses the combined impact of transceiver impairments and noise on Signal-to-Noise Ratio (SNR) and BER. A Variable Optical Attenuator (VOA) is employed to emulate different Received Optical Power (ROP) levels [18,19].

The optical path passes through a network demarcation unit (NDU), hosting a 400G DSCM coherent pluggable P2MP transmission for both Tx and Rx operation. The B2B configuration eliminates line impairments, enabling isolation of the intrinsic transceiver noise contribution. Accurate network-level performance prediction therefore requires explicit characterization of transceiver-induced degradation, which becomes particularly relevant at low ROP conditions typical of access segments. The transceiver contribution is quantified through the definition of SNR_{TRX}, which depends on the device under test and the received power. In practice, BER is measured across varying ROP levels and converted to SNR by inverting the analytical relation $BER = k_1 \cdot \mathrm{erfc}\left(\sqrt{k_2 \cdot SNR}\right)$, where k_1 and k_2 depend on the modulation format [20,21]. Under B2B conditions, the measured SNR corresponds exclusively to SNR_{TRX}, since fiber impairments are negligible.

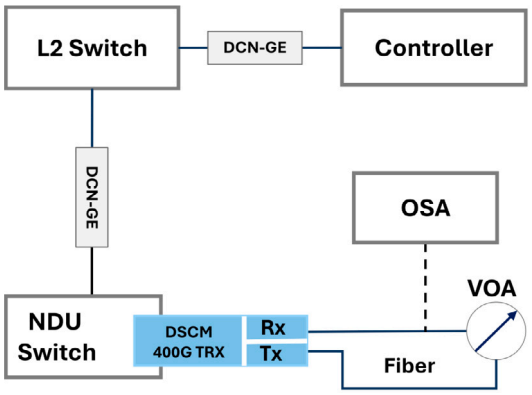

Fig. 2. Experimental setup in back-to-back (B2B)

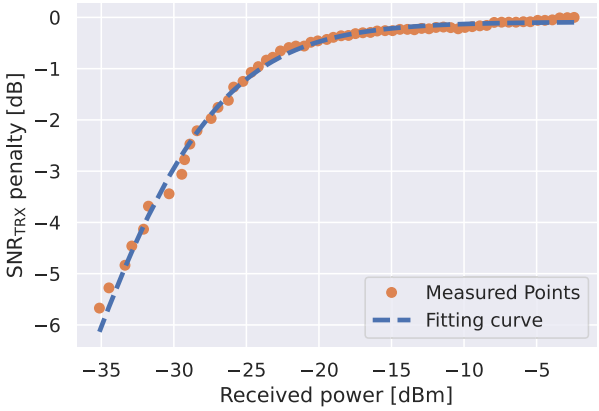

Fig. 3. B2B Characterization of TRX.

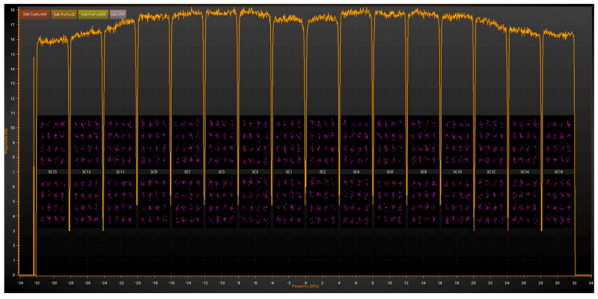

Fig. 4. Trace of Spectrum Occupancy for 16 DSCs.

The resulting characterization yields a functional relationship between ROP and SNR_{TRX}, as shown in Fig. 3. The measured SNR_{TRX} values are normalized to an arbitrary maximum of 0 dB to anonymize device-specific performance. The relationship is analytically modeled as

$$SNR_{TRX,Lin} = N\left(\frac{ROP_{TRX,Lin}}{ROP_{TRX,Lin} + D}\right), \quad (1)$$

where N and D are fitting parameters determined through numerical regression. The curve exhibits two distinct regions: a saturation region at high ROP, where performance variations are minimal, and a steep degradation region at low ROP, corresponding to receiver sensitivity limits.

Fig. 4 shows the spectrum trace at the aggregation side, representing the Fourier transform of the received electrical signal and confirming the presence of 16 digital subcarriers. The system operates using DP-16QAM modulation over a 64 GHz spectral window, achieving 400G aggregate capacity. Through advanced DSP, 16 × DSCs are multiplexed within a single wavelength, each with a symbol rate of 4 GBaud and supporting up to 25 Gb/s per subcarrier [22].

Importantly, the experimentally derived SNR_{TRX}- ROP model is directly embedded into the network-level simulator that enables impairment-aware route feasibility evaluation under realistic transceiver constraints.

3.2. Statistical simulation model

Fig. 5 illustrates two converged metro-access network topologies, denoted as N_1 (Fig. 5(a)) and N_2 (Fig. 5(b)), which serve as baseline architectures for evaluating route feasibility and QoT metrics under 6G fronthaul (DU-RU) constraints using DSCM-based coherent TRXs.

Network topology N_1 comprises 37 nodes, including 12 metro ROADMs (M-nodes) and 25 access nodes (A-nodes). Metro nodes are interconnected through 19 bidirectional fiber links, achieving an average node degree of 3.42. The average link length between metro nodes is 6.9 km, with a maximum span of 10 km, representing a compact metro ring. Each access cluster consists of four A-nodes serving two RUs, interconnected with the metro segment via a circulator to enable single fiber bi-directional (BiDi) transmission in the access segment and fiber-pair transmission in the metro segment. Up to five DUs are strategically deployed at metro nodes to provide distributed processing capability closer to the network edge.

In N_2, the network scale and connectivity are extended, comprising 63 nodes, of which 18 belong to the metro segment and 45 to the access segment. Metro nodes exhibit higher redundancy, resulting in an average node degree of 4.06 and a maximum fiber length of 12 km. The increased connectivity and structural diversity introduce a broader set of routing options, enhancing scalability and resilience. Similar to N_1, each access segment aggregates four A-nodes per RU and connects to the metro segment via circulators. Additionally, a random statistical power-loss distribution is applied to the access segment, following the approach in [23].

Topologies N_1 and N_2 represent structurally distinct converged metro-access networks, as summarized in Table 1(a). While N_1 exhibits moderate connectivity with shorter average path lengths and limited routing diversity, N_2 features higher node degree, larger scale, and greater path diversity. These structural characteristics directly influence the effectiveness of the DU densities listed in Table 1(b). In N_1, constrained routing diversity limits variations in path length across DU locations, resulting in modest placement sensitivity. In contrast, the higher connectivity of N_2 amplifies the impact of DU positioning on path-length distribution and impairment accumulation, leading to more pronounced feasibility differences. Although absolute feasibility values differ between the two networks, the dominant trends are governed by routing diversity, and their interaction with DU density and placement.

A Python-based simulator is developed to evaluate the proposed ONaaS architecture using DSCM-based coherent pluggable TRXs. The simulator imports the converged metro-access topologies (N_1 and N_2) and computes all candidate end-to-end routes by incorporating a detailed physical-layer budget. This budget accounts for fiber attenuation, ROADM insertion losses, passive component losses, amplifier noise accumulation, and experimentally characterized transceiver impairments derived from the B2B measurements described in Section 3.1.

For each candidate route, the received optical power (ROP) is computed by accounting for accumulated fiber attenuation, ROADM insertion losses, passive component losses, and amplifier noise contributions. The resulting SNR is evaluated using Eq. (2), where both generalized line-induced noise (GSNR) and transceiver-induced impairments from (1) are considered, and the corresponding BER is obtained using Eq. (3). Latency is modeled as the fiber propagation delay along the selected route, expressed as $t = \frac{d}{v}$, where d denotes the total fiber length and v is the propagation speed in optical fiber (with c representing the speed of light in vacuum). Since a fully transparent optical

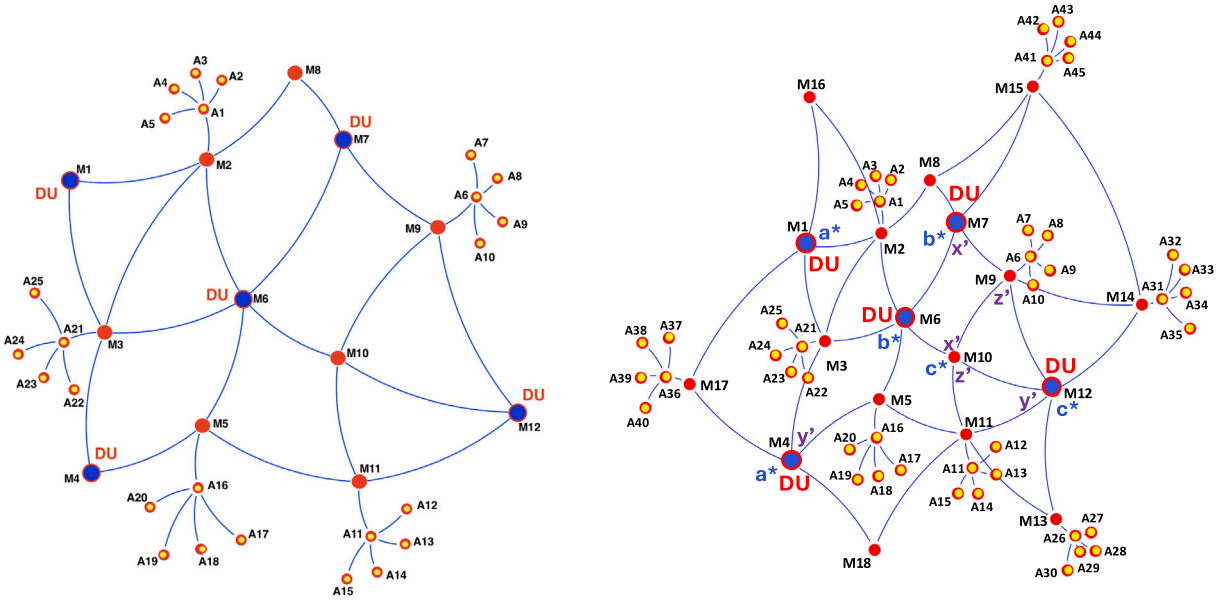

(a) Network topology (N_1) (b) Network topology (N_2)

Fig. 5. Simulated converged metro-access topologies.

Table 1
Topology characteristics and DUs configuration.

(a) Converged metro-access topologies

Parameters	N_1	N_2
Total Nodes	37	63
Metro Nodes (M)	12	18
Access Nodes (A)	25	45
Average Metro Node Degree	3.42	4.06
Average Link Length	6.9 km	7.31 km
Maximum Link Length	10 km	12 km
Access Segment Design	4 A-nodes, BiDi via circulator	
Reliability & Redundancy	Moderate	High
Scalability	Moderate	High
Dominant Risks	Limited fault tolerance	Higher latency & QoT degradation

(b) DU Densities in metro ring

Metro Nodes	No. of DU
M1	1
M1, M4	2
M1, M4, M12	3
M1, M4, M6, M12	4
M1, M4, M6, M7, M12	5

architecture is assumed, latency in the fronthaul section is dominated by propagation delay. Although the feasibility analysis is simulation-based, it integrates experimentally measured transceiver characteristics and realistic impairment modeling to ensure practical representation of converged metro-access network scenarios.

$$\mathrm{SNR}^{-1} = \mathrm{GSNR}^{-1} + \mathrm{SNR}_{\mathrm{TRX}}^{-1} \tag{2}$$

$$\mathrm{BER}_{16\mathrm{QAM}} = \frac{3}{8}\,\mathrm{erfc}\left(\sqrt{\frac{\mathrm{SNR}}{10}}\right) \tag{3}$$

Using a Monte Carlo approach, route feasibility is evaluated under varying hop limits and DU densities in a fully transparent converged metro-access scenario with an overall spectral occupation of 64 GHz [16,24]. Each DSC supports up to 25 Gb/s throughput. Accordingly, pairs of DSCs are allocated to serve the BiDi traffic demand of a single RU, consistent with the split 7-2 bandwidth requirement, as illustrated in Fig. 1. This approach enables a modeling of TRX characteristics into network-level feasibility evaluation.

3.3. ML-based QoT estimation

Fig. 6 illustrates the proposed machine-learning-based regression framework for QoT estimation in converged metro-access networks. While the feasibility analysis described in Section 3.2 provides accurate route evaluation, it incurs significant computational overhead as the network topology scales, hop limits and DU densities grow due to the large number of candidate routes that must be evaluated. To enable scalable and rapid decision-making within an ONaaS environment, a data-driven estimator is introduced to predict BER without repeatedly executing computationally intensive physical-layer simulations.

The simulator generates a dataset by evaluating multiple end-to-end routes across varying DU densities and topological configurations. The dataset includes network-level parameters and corresponding QoT metrics, summarized in Table 2. To develop a lightweight estimator suitable for scalable deployment, feature selection excludes parameters directly derived from physical-layer calculations (e.g., GSNR or ROP), as these quantities are inherently correlated with BER and would require full physical-layer evaluation. Instead, structural features that can be obtained without detailed physical-layer computation are selected. Specifically, `distance` and `boundary_id` are used as input features, with BER as the target label. The `distance` feature captures cumulative impairment effects along optical paths, while `boundary_id` uniquely represents each source–destination pair. By grouping routes that share common endpoints, `boundary_id` captures routing diversity and network QoT conditions within source–destination pair.

As shown in Fig. 7, the left plot illustrates the relationship between transmission distance and BER, which generally increases with distance due to the cumulative effects of attenuation, amplifier noise,

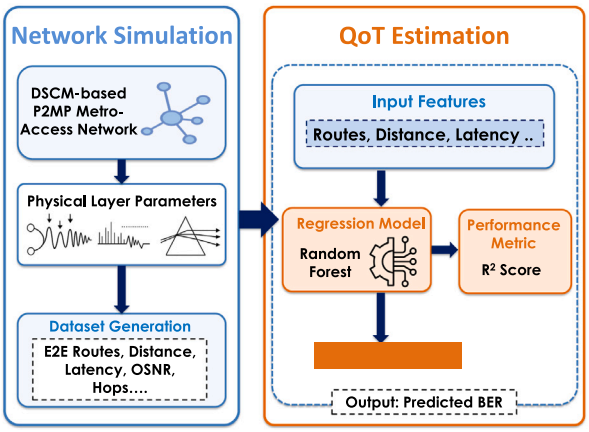

Fig. 6. ML-driven QoT estimation model.

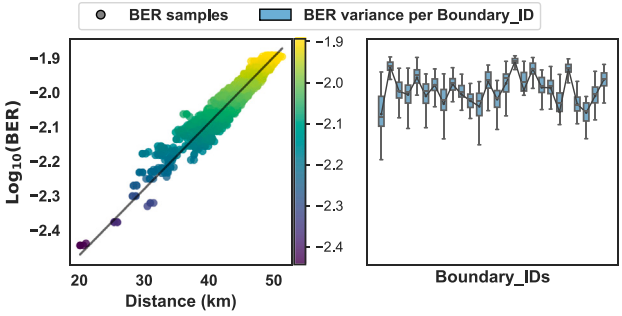

Fig. 7. Distance and boundary_id vs. BER (N_2-5DU)

and component losses along the optical path. However, the relationship is not strictly linear, as BER is also influenced by route-specific impairment accumulation. Consequently, multiple BER samples appear for similar distance values, corresponding to different end-to-end routes with similar path lengths but distinct physical-layer conditions. This dispersion reflects the impact of heterogeneous impairment accumulation along different routes within the network. The right plot presents the BER distribution across different `boundary_ids`. The error bars represent the dispersion of BER values associated with routes connecting the same boundary, reflecting variability introduced by routing diversity. The mean trend line represents the average BER per boundary, reflecting the typical QoT conditions for routes connecting a given source–destination pair. Differences across boundaries arise from variations in path-length distribution and topology-dependent impairments.

These observations indicate that `distance` and `boundary_id` jointly capture impairment accumulation and route variability, enabling accurate BER prediction without explicit physical-layer calculations.

3.3.1. Training and testing strategy

The training and testing strategy is designed to assess both scalability and cross-topology generalization. Model training is conducted using the simplest deployment configuration (1DU) of topology N_2, which exhibits the lowest routing complexity and shortest average path lengths. This minimal scenario serves as the base training configuration. The trained model is subsequently evaluated under two distinct conditions:

1. $N_2 \rightarrow N_2$: Testing on the same topology under increasing DU densities (1DU, 3DU, and 5DU) to assess scalability.

Table 2
Simulated network parameters.

Parameters	Description
boundaries	Src. and Dst. pair
path	End-to-end (E2E) route between Src. and Dst.
distance	Total length of an (E2E) path (km)
GSNR	Generalized SNR (dB)
ROP	Received optical power (dBm)
SNR	Signal-to-noise ratio (dB)
BER	Bit-error-rate
latency	Propagation delay (ms)
boundary_id	Unique identifier for boundaries
last_link_km	Final hop length of a route (km)

2. $N_2 \rightarrow N_1$: Testing on a structurally different topology to evaluate cross-topology transferability.

This strategy verifies whether the model captures general BER behavior governed by structural network features rather than memorizing topology-specific routes. To evaluate model performance under varying data availability, the training dataset size is progressively reduced from 100% to fractions of 60%, 50%, 30%, and 10%. To further assess robustness, additive Gaussian noise is introduced during testing with standard deviation levels

$$\sigma = [0.5,\ 0.4,\ 0.3,\ 0.2,\ 0.1,\ 0.0],$$

thereby emulating uncertainty and variability in simulated network conditions. This procedure ensures that the estimator learns transferable structural relationships between network parameters and BER, rather than relying on deterministic route-specific patterns. Once trained, the regression model predicts BER for unseen routes, enabling computationally efficient route feasibility assessment without executing full physical-layer exhaustive simulations.

3.3.2. Model selection and performance

To identify a suitable regression model for BER prediction, multiple machine learning models are evaluated. Specifically, Linear Regression (LR), Decision Tree (DT), and Random Forest (RF) models are trained and tested under identical conditions as described in Section 3.3.1 to ensure a consistent comparison.

The comparative evaluation is conducted across varying DU densities, noise levels, and training data fractions. Among the evaluated models, RF consistently achieves higher predictive accuracy and improved robustness compared to LR and DT. This performance advantage arises from the ensemble structure of RF, which aggregates multiple decorrelated trees to reduce overfitting while effectively capturing nonlinear relationships between network features and BER [25,26]. Consequently, RF is adopted as the primary regression model for QoT estimation. Model performance is quantified using the coefficient of determination (R^2), defined as

$$\mathrm{R}^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y})^2}, \tag{4}$$

where y_i denotes the observed BER, $\hat{y}_i$ the predicted value, and $\bar{y}$ the mean of the observed values. The R^2 coefficient measures the proportion of variance in the observed BER, with values indicating stronger predictive capability. This metric aligns with the objective of developing a lightweight and accurate estimator capable of generalizing BER behavior across topologies and varying DU densities while maintaining computational efficiency.

4. Results and discussion

This section presents the impairment-aware route feasibility analysis for topologies N_1 and N_2 under strict latency and BER constraints.

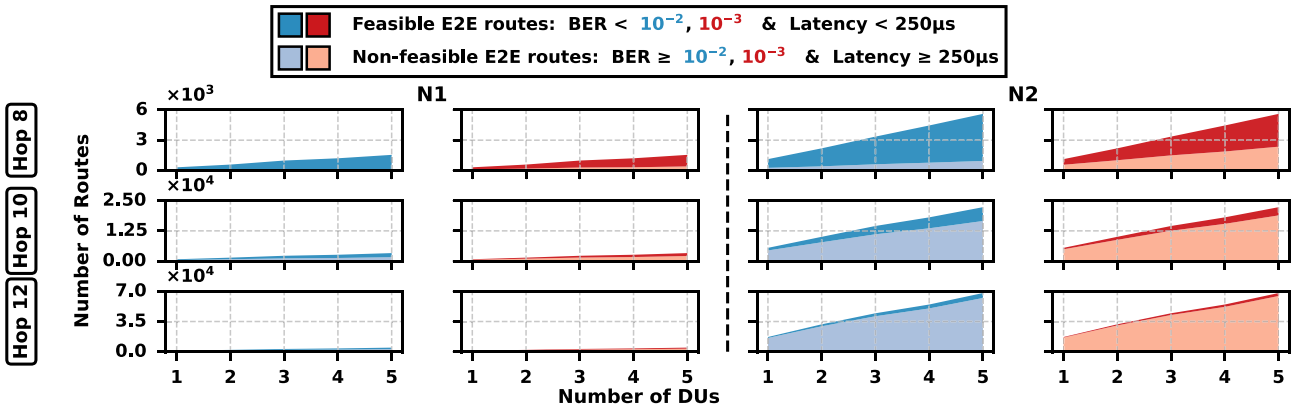

Fig. 8. Route feasibility under BER thresholds (10^{-2}, 10^{-3}) and latency < 250 µs in N_1 and N_2 across hop limits and DU densities.

The impact of DU placement on achievable routes and QoT performance is systematically examined. In addition, the effectiveness of the proposed ML-based QoT estimator is evaluated under computational complexity constraints, followed by a techno-economic comparison between coherent P2MP and P2P architecture.

4.1. Route feasibility analysis

Fig. 8 presents route feasibility for topologies N_1 and N_2 under BER thresholds of 10^{-2} and 10^{-3}, subject to a strict latency constraint of < 250 µs. The analysis considers varying hop limits (8, 10, and 12) and DU densities (1–5 DUs) in the metro segment, as defined in Table 1(b), where DUs are progressively deployed to evaluate the effect of proximity on routing performance.

Although increasing hop limits significantly increases the total number of routes, the proportion of feasible routes converges at moderate routing depth (around 8 hops). Extending the hop limit to 10 and 12 primarily introduces longer paths that are constrained by accumulated impairments rather than enhancing feasible connectivity. Consequently, a hop limit of 12 is adopted to represent a mid-to-high routing depth and to evaluate whether extended path diversity yields additional admissible routes.

For N_1 with 5DUs, feasibility under BER < 10^{-2} decreases from 93.19% (8 hops) to 51.64% (10 hops) and 35.83% (12 hops). Under the stricter BER < 10^{-3} threshold, feasibility reduces from 74.42% to 36.55% and 25.21%, respectively. A more pronounced degradation is observed in N_2, where feasibility with 5DUs under BER < 10^{-2} drops from 83.27% (8 hops) to 25.62% (10 hops) and 8.40% (12 hops), while under BER < 10^{-3} it declines from 58.00% to 15.21% and 5.00%.

The stronger sensitivity of N_2 reflects its larger scale and higher routing diversity, which introduce longer average path lengths and greater impairment accumulation. As hop count increases, fiber attenuation, ROADM insertion losses, and amplifier noise collectively degrade the received SNR, directly reducing route feasibility. Under the fully transparent architecture assumption, latency is dominated by propagation delay; therefore, longer multi-hop routes are also more likely to violate the 250 µs constraint.

Overall, the results indicate that under strict split 7-2 latency and BER constraints, route feasibility in converged metro-access networks is fundamentally bounded by physical-layer impairments. Beyond moderate routing depth, increasing path diversity yields diminishing gains in feasible connectivity.

4.2. Impact of DU placement

To evaluate the impact of DU placement on route feasibility, the analysis is conducted on topology N_2, which exhibits higher node degree and greater routing diversity compared to N_1. Owing to its dense

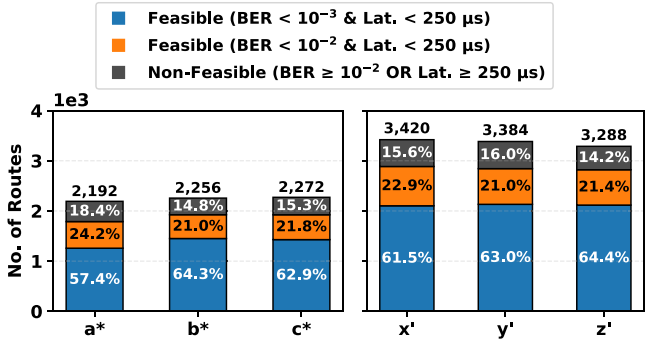

Fig. 9. Comparison of (a) 2× & (b) 3× DU placement scenarios in topology N_2 under 8 hops.

structural connectivity, N_2 is more sensitive to placement variations and therefore better highlights the influence of DU positioning on route feasibility. Since the route feasibility analysis in Section 4.1 shows that feasible connectivity stabilizes at eight hops; therefore, the hop limit is fixed to 8 to focus on placement effects while excluding predominantly non-feasible longer routes.

Fig. 9 compares different DU placement scenarios using stacked bars that classify routes into three categories: (1) feasible (BER < 10^{-3} and latency < 250 µs), (2) marginally feasible (BER < 10^{-2} and latency 250 µs), and (3) non-feasible. The numeric value above each bar denotes the total number of routes, while the colored segments indicate the proportion of routes in each feasibility category.

For the 2DU case shown in Fig. 9(a), the total number of routes remains comparable across all placements (≈ 2300). However, the feasibility distribution varies noticeably. Configuration a* (M1, M4) yields 57.4% feasibility (cat. 1) with 18.4% non-feasible routes, whereas configuration b* (M6, M7) achieves the highest cat. 1 feasibility (64.3%) and the lowest non-feasible fraction (14.8%). Configuration c* (M10, M12) provides intermediate performance (62.9% cat. 1 feasibility and 15.3% non-feasible). Despite the marginal variation in the total routes, cat. 1 feasibility differs by nearly 7%, demonstrating that DU placement directly influences QoT rather than simply increasing connectivity.

A similar trend is observed for the 3DU configurations shown in Fig. 9(b). Configuration x' (M3, M7, M10) achieves 61.5% cat. 1 feasibility, configuration y' (M4, M8, M12) increases to 63.0%, and configuration z' (M6, M9, M10) reaches the highest cat. 1 feasibility of 64.4% with the lowest non-feasible proportion (14.2%). Although route availability changes only marginally across evaluated configurations (≈3400), the improvement in cat. 1 feasibility results primarily from reduced impairment accumulation along the routes rather higher number of available routes.

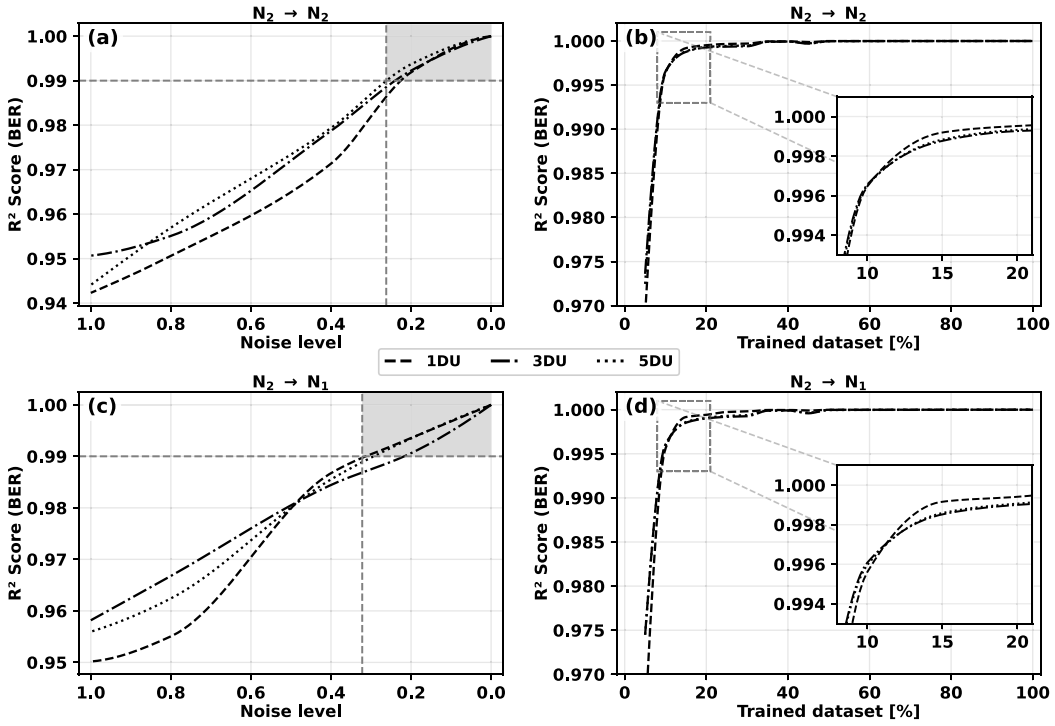

Fig. 10. BER prediction using ML-based Random Forest (RF) for 1DU, 3DU, 5DU scenario. (a) R^2 with varying noise level for N_2 (Training) → N_2 (Testing), (b) R^2 with varying trained dataset (%) for N_2 (Training) → N_2 (Testing), (c) R^2 with varying noise level for N_2 (Training) → N_1 (Testing), (d) R^2 with varying trained dataset (%) for N_2 (Training) → N_1 (Testing)

The observed behavior is primarily driven by physical-layer effects. Placing DUs at highly connected metro nodes reduces the average end-to-end path length and limits cumulative attenuation, ROADM insertion losses, and noise accumulation along optical routes. Consequently, a larger fraction of paths satisfies both BER and latency constraints. In contrast, suboptimal DU placement increases routing distance, resulting in stronger SNR degradation and a lower proportion of feasible routes even under the same hop depth.

The analysis therefore focuses on the 2DU and 3DU densities, where placement effects are most pronounced and routing diversity strongly influences impairment accumulation. At moderate (2–3 DUs) densities, appropriate positioning leads to consistent improvements in strict feasibility. From an operator perspective, increasing the number of deployed DUs expands the available fronthaul routing options, enabling more balanced traffic distribution, improving redundancy for X-haul operation.

4.3. QoT estimation with complexity constraints

Although individual route evaluation may appear computationally manageable for small network configurations, the combinatorial growth of routing paths with increasing hop limits, DU densities, and topology size leads to rapidly increasing computational complexity. Consequently, exhaustive QoT evaluation becomes impractical for large metro-access deployments. To address this challenge, the proposed ML-based estimator is evaluated across varying DU densities and network topologies, as shown in Fig. 10.

The model is trained using the simplest configuration (1x DU in N_2) with a hop limit of 12, corresponding to the lowest routing complexity w.r.t DU densities. This minimal training setup is selected to verify whether the estimator captures the underlying BER-distance relationship rather than memorizing topology-specific routing patterns. The trained model is subsequently evaluated under the testing scenarios described in Section 3.3.1.

Fig. 10(a) and (c) report the R^2 score as a function of noise level for intra ($N_2 \rightarrow N_2$) and cross-topology ($N_2 \rightarrow N_1$) evaluations, respectively. In both cases, prediction accuracy improves as noise decreases, indicating stable estimator behavior under cleaner input conditions. The gray shaded regions indicate the noise ranges where the RF estimator achieves consistently high R^2 score across DU configurations (1DU, 3DU, and 5DU). For the intra-topology case, convergence occurs at slightly lower noise levels compared to the cross-topology scenario, reflecting structural differences in routing diversity between the two evaluated topologies.

Fig. 10(b) and (d) analyze the impact of training dataset size on prediction accuracy. In both intra- and cross-topology scenarios, the R^2 score increases rapidly with the fraction of training data and reaches near saturation when approximately 10%–20% of the training dataset is used. The inset plots highlight this convergence region, indicating that only a small portion of the dataset is sufficient to achieve stable BER prediction. Beyond this point, additional training data yields only marginal improvements in accuracy while slightly enhancing prediction stability, particularly in the cross-topology case.

Fig. 11 compares predicted and true BER values for intra- and cross-topology evaluations under varying noise levels and training dataset fractions. Subfigures (a)–(d) illustrate the impact of noise level for 1DU and 5DU deployments. In the intra-topology case, the predicted BER values closely follow the trend line, indicating accurate estimation. Increasing noise levels introduce visible dispersion around the trend line. Subfigures (e)–(h) analyze the influence of the training dataset fraction. Increasing the training data from 30% to 60% reduces the dispersion of BER prediction and improves their alignment with the trend line in both topology scenarios.

Fig. 12 illustrates the computational performance comparison between exhaustive simulation and ML-based QoT estimation for topology N_2. The horizontal axis represents the number of DUs, while the vertical axis shows execution time in seconds on a logarithmic scale. The bar chart on the left reports the total simulation time required to evaluate route feasibility for increasing DU densities under a hop

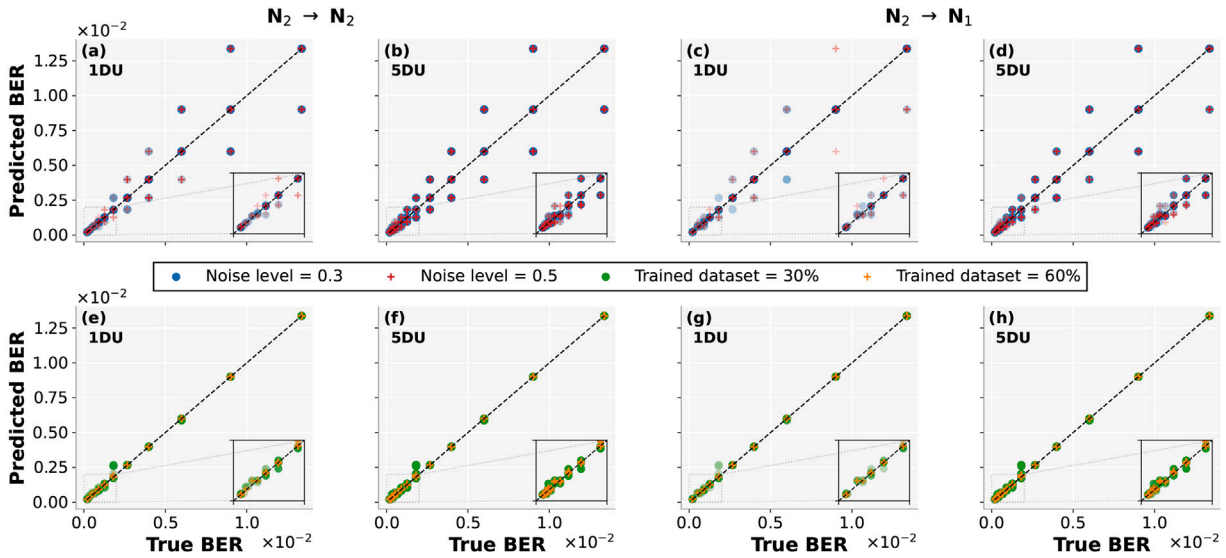

Fig. 11. Predicted vs. True BER comparison for 1DU, 5DU deployment scenarios. Subfigures (a)–(b) show $N_2 \rightarrow N_2$ and (c)–(d) $N_2 \rightarrow N_1$ with varying noise level, Subfigures (e)–(f) show $N_2 \rightarrow N_2$ and (g)–(h) $N_2 \rightarrow N_1$ with varying trained dataset(%)

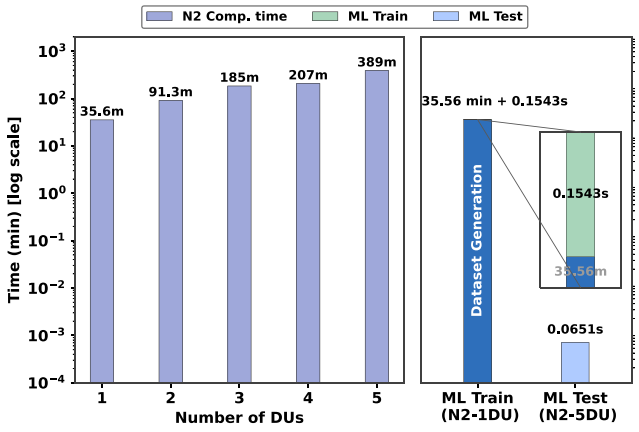

Fig. 12. Computational performance comparison between exhaustive simulation and the ML-model.

limit of 12. The panel on the right summarizes the ML processing time, showing the one-time dataset generation time, training, and testing time required for BER prediction.

As observed, simulation time increases sharply with DU density due to the rapid growth in the number of routes evaluated through full physical-layer computations. The execution time rises from (≈ 35.56 min) for the 1DU configuration to about (≈ 389 min) for the 5DU case. In contrast, the ML-based approach requires a one-time dataset generation step of (≈ 35.6 min) for the simplest configuration (1DU in N_2), followed by a model training time of only (≈ 0.1543 s). Once trained, BER prediction for higher DU configurations is completed in approximately (≈ 0.065 s) for the 5DU case, corresponding to an inference speedup of roughly five orders of magnitude compared with exhaustive simulation. This corresponds to an overall computational reduction of (≈ 11×), obtained by dividing the simulation time for 5DU by the total ML processing time (dataset generation, ML-training and testing). The dataset generation step corresponds to a single simulation of the minimal configuration and is performed only once, whereas exhaustive simulation must be repeated for every DU density and topology scenario. These results demonstrate that accurate QoT estimation under strict split 7-2 constraints can be achieved with significantly reduced computational complexity, thereby supporting scalable ONaaS deployment for future 6G X-haul networks.

4.4. Techno-economic gains of P2MP over P2P

To quantify the economic implications of metro-access convergence, a techno-economic assessment is performed comparing conventional P2P with DSCM-enabled P2MP optical transport architecture. The techno-economic assessment is conducted on topology N_1 with a 5-DU deployment under a 10-year traffic projection assuming a 17.2% compound annual growth rate (CAGR) derived from global mobile traffic forecasts [27], starting from an aggregated load of 1 Tb/s. The analysis compares conventional P2P optical transport based on dedicated 25G optical links per RU with coherent DSCM-enabled P2MP aggregation within a 400G coherent transport window. Both architectures are evaluated under identical topology, traffic evolution, and service constraints following the techno-economic modeling framework described in [28]. Table 3 summarizes the evolution of the key techno-economic indicators.

Due to dedicated wavelength provisioning, the P2P architecture exhibits near-linear scaling of transceivers and router ports as traffic grows, leading to continuous hardware expansion. In contrast, P2MP leverages subcarrier-level aggregation, allowing multiple RUs to share high-capacity coherent resources within a single 400G transport window. Consequently, capacity upgrades occur in aggregated increments rather than incremental expansions per RU, resulting more efficient resource utilization.

As shown in Table 3, P2MP significantly reduces hardware requirements across the evaluated period. At the DU site, the number of optical transceivers is reduced by approximately 75% compared with the P2P architecture, reflecting the aggregation of multiple RUs within a single coherent transport window. At the network level, transceiver reductions reach approximately 37.5% during early deployment stages and increase as traffic grows, indicating improved spectral utilization efficiency over time.

Router port scaling is similarly deferred due to more efficient traffic aggregation, reducing the rate at which switching infrastructure must expand. From an operational perspective, power consumption advantages become more pronounced at higher spectral loads. Using the coherent transceiver energy model described in [29], the evaluated scenario shows approximately 27%–30% lower power consumption for P2MP under high spectral utilization. These results demonstrate that subcarrier-based P2MP aggregation enables more efficient

Table 3
Techno-economic projection for topology N_1 with a 5DU deployment under a 10-year traffic growth (17.2% CAGR).

Metric	Year 0	Year 10	Techno-economic Implication
TRX per DU site (P2P)	4	27	Linear hardware scaling
TRX per DU site (P2MP)	1	7	Subcarrier-based aggregation leads to reduction in links
DU-site TRX reduction (P2MP)	75%	74.1%	Sustained CAPEX reduction
Network-wide TRX (P2P)	40	135	Growing infrastructure footprint
Network-wide TRX (P2MP)	25	35	Slower hardware expansion
Network-wide TRX reduction (P2MP)	37.5%	74.1%	CAPEX advantage increases with traffic
Router expansion (DU-site)	Low	High (P2P)	Earlier port saturation in P2P case
Power efficiency (P2MP, high spectral load)	Comparable	27%–30%	Reduction in OPEX

resource utilization, reducing hardware footprint, delaying infrastructure expansion, and improving energy efficiency under long-term traffic growth.

5. Conclusion and future work

This work presented a feasibility and scalability analysis of a converged metro-access network designed to support 6G X-haul under split 7-2 latency and BER constraints. Experimental characterization of a DSCM-enabled coherent transceiver was incorporated into a network-level simulation to enable QoT evaluation. The results show that increasing DU density improves feasible connectivity primarily by reducing average path length and latency, while topology expansion alone does not guarantee improved feasibility due to longer routes. Consequently, strategic DU placement plays a key role in improving strict route feasibility. To address the scalability limitations of exhaustive QoT evaluation, a lightweight RF-based estimator was developed and validated across varying DU densities and structurally different topologies. The model achieves high prediction accuracy with minimal input features and limited training data, remaining robust under varying noise levels while enabling rapid BER prediction. Compared with full physical-layer simulation, the ML-based approach reduces the overall computational load by ($\approx$ 11×), including dataset generation and model training and testing. Finally, a techno-economic assessment shows that coherent P2MP aggregation provides sustained CAPEX and OPEX advantages over conventional P2P architectures under long-term traffic growth. Overall, these findings support scalable and cost-efficient converged metro-access networks suitable for future 6G X-haul.

Future work will focus on AI-driven orchestration including subcarrier-level resource allocation based on traffic load. The investigation of hybrid and distributed architectures will further examine trade-offs in performance, energy efficiency, and scalability. In addition, SDN-based control of individual subcarriers using the experimental testbed will be pursued to enable telemetry-driven and on-demand reconfigurable ONaaS.

CRediT authorship contribution statement

Sanwal Zeb: Writing – original draft, Visualization, Methodology, Investigation, Conceptualization. **Ahtisham Ali:** Resources, Formal analysis. **Imran Chowdhury Dipto:** Writing – review & editing, Software, Methodology, Data curation. **Andrea Rosso:** Formal analysis. **Muhammad Umar Masood:** Writing – review & editing, Visualization, Validation, Methodology, Conceptualization. **Riccardo Schips:** Writing – review & editing. **Renato Ambrosone:** Writing – review & editing. **Stefano Straullu:** Writing – review & editing, Resources. **Francesco Aquilino:** Writing – review & editing, Resources. **Antonino Nespola:** Project administration. **João Pedro:** Writing – review & editing, Supervision, Project administration. **Antonio Napoli:** Writing – review & editing, Supervision, Project administration. **Alessandro Galardini:** Project administration. **Vittorio Curri:** Writing – review & editing, Supervision, Project administration.

Acknowledgments

This work is funded by Horizon Europe which is the EU's main funding program for research and innovation, for project *"NESTOR"* GA No. *101119983*, project *"EWOC"*, Italy GA No. *101073265*, project *"ALLEGRO"*, Italy GA No. *101092766*, project *"RESTART"*, Italy with GA *PE00000001*, and also in collaboration with *"LINKS Foundation, Torino, Italy"* & *"NOKIA (Munich & Carnaxide)"*.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The authors do not have permission to share data.

References

[1] P. Schulz, M. Matthe, H. Klessig, M. Simsek, G. Fettweis, J. Ansari, S.A. Ashraf, B. Almeroth, J. Voigt, I. Riedel, et al., Latency critical IoT applications in 5G: Perspective on the design of radio interface and network architecture, IEEE Commun. Mag. 55 (2) (2017) 70–78.
[2] M. Nooruzzaman, X. Fernando, Cost dynamics of converged optical-wireless networks: enabling low-latency xRANs through a reconfigurable hybrid split, J. Opt. Commun. Netw. 16 (6) (2024) 659–669.
[3] W. Jiang, B. Han, M.A. Habibi, H.D. Schotten, The road towards 6G: A comprehensive survey, IEEE Open J. Commun. Soc. 2 (2021) 334–366.
[4] Z. Vujicic, M.C. Santos, R. Méndez, B. Klaiqi, J. Rodriguez, X. Gelabert, M.A. Rahman, R. Gaudino, Towards virtualized optical-wireless heterogeneous networks, IEEE Access (2024).
[5] N. Skorin-Kapov, F.J.M. Muro, M.-V.B. Delgado, P.P. Marino, Point-to-multipoint coherent optics for re-thinking the optical transport: case study in 5G optical metro networks, in: 2021 International Conference on Optical Network Design and Modeling, ONDM, IEEE, 2021, pp. 1–4.
[6] C. Castro, J. Sime, T. Eriksson, M.S. Erkilinc, M. Porrega, J. Pedro, M. Quagliotti, E. Riccardi, C. Fludger, A. Napoli, Power and spectral savings in metro-aggregation networks exploiting coherent point-to-multipoint transceivers, in: ECOC 2024; 50th European Conference on Optical Communication, VDE, 2024, pp. 519–522.
[7] D. Welch, A. Napoli, J. Bäck, W. Sande, J. Pedro, F. Masoud, C. Fludger, T. Duthel, H. Sun, S.J. Hand, et al., Point-to-multipoint optical networks using coherent digital subcarriers, J. Lightwave Technol. 39 (16) (2021) 5232–5247.
[8] Y. Hu, A. Yan, J. Zhao, S. Xing, C. Shen, Z. Li, Y. Zhou, J. Shi, Z. He, N. Chi, et al., All-optical metro-access integration network bidirectional transmission enabled by coherent digital subcarrier multiplexing, J. Opt. Commun. Netw. 17 (1) (2024) 58–70.
[9] M. Radovic, A. Sgambelluri, F. Cugini, R. Kapuscinski, M. Hosseini, S. Parker, R. Berozashvili, F. Paolucci, A. Napoli, N. Sambo, SDN-controlled open RAN X-haul with point-to-multipoint transceivers on a horseshoe network, J. Opt. Commun. Netw. 17 (11) (2025) E109–E116.
[10] A.S. Thyagaturu, A. Mercian, M.P. McGarry, M. Reisslein, W. Kellerer, Software defined optical networks (SDONs): A comprehensive survey, IEEE Commun. Surv. Tutor. 18 (4) (2016) 2738–2786.
[11] Cloud architecture and deployment scenarios for O-RAN virtualized RAN, 2023, O-RAN Working Group 6 (Cloudification and Orchestration).
[12] 3GPP, Radio Access Architecture and Interfaces (Release 14), Tech. Rep. 3GPP TR 38.801 V14.0.0 (2017-03), 3GPP, 2017.

[13] E. Kosmatos, C. Matrakidis, D. Uzunidis, A. Stavdas, S. Horlitz, T. Pfeiffer, A. Lord, E. Riccardi, Real-time orchestration of qos-aware end-to-end slices across a converged metro and access network exploiting burst-mode technology, J. Opt. Commun. Netw. 15 (1) (2022) 1–15.

[14] A. Larranaga, S. Lagén, J.M. Fàbrega, J.M. Rivas-Moscoso, J.P. Fernández-Palacios, I. Tomkos, R. Munoz, Fronthaul/midhaul networks: Capacity and latency requirements imposed by 6G disaggregated RANs, IEEE Commun. Mag. (2025).

[15] 3GPP, Study on CU-DU lower layer split for NR; (Release 15), Tech. Rep. TR 38.816 V15.0.0 (2017-12), 3GPP, 2017.

[16] 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, J. Lightwave Technol. 41 (4) (2023) 1175–1191.

[17] J. Pedro, M.M. Hosseini, A. Napoli, Extended network applications of coherent pluggable transceivers, J. Opt. Commun. Netw. 17 (2) (2025) A210–A223.

[18] T. Mano, A. D'Amico, Y.-K. Huang, G. Borraccini, H. Nishizawa, T. Wang, V. Curri, K. Takasugi, Measuring the transceiver's back-to-back BER-osnr characteristic using only a variable optical attenuator, in: ECOC 2024; 50th European Conference on Optical Communication, 2024, pp. 1499–1502.

[19] T. Mano, Y.-K. Huang, G. Borraccini, E. Ip, A. D'Amico, Z. Wang, H. Nishizawa, G. Zussman, T. Chen, T. Wang, K. Asahi, D. Kilper, V. Curri, K. Takasugi, Modeling the input power dependency of transceiver BER-ONSR for QoT estimation, in: 2024 Optical Fiber Communications Conference and Exhibition, OFC, 2024, pp. 1–3.

[20] A. Rosso, Modeling and Controlling Optical Transponder White-Boxes Based on the Physical Layer Digital Twin (Master's thesis), Politecnico di Torino, 2024.

[21] A. Ali, M.U. Masood, A. Rosso, G. Malik, M. Pollone, A. Galardini, V. Curri, Statistical analysis of end-to-end route feasibility in converged metro–access optical networks, IEEE Photonics Technol. Lett. (2026) http://dx.doi.org/10.1109/LPT.2026.3653136, 1–1.

[22] A. Ali, S. Zeb, A. Rosso, M.U. Masood, G. Malik, R. Ambrosone, R. Schips, S. Straullu, F. Aquilino, J. Pedro, A. Napoli, A. Galardini, V. Curri, 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.

[23] G. Simon, P. Chanclou, F. Saliou, J. Potet, M. Wang, Clustering G-PON field data to improve flexibility in next generation PON systems, in: 2021 European Conference on Optical Communication, ECOC, IEEE, 2021, pp. 1–4.

[24] A. Ali, S. Zeb, A. Rosso, M.U. Masood, G. Malik, R. Ambrosone, R. Schips, B. Correia, S. Straullu, F. Aquilino, J. Pedro, A. Napoli, A. Galardini, V. Curri, Optimizing RAN X-Haul performance through targeted hollow-core fiber deployment in converged metro-access networks, in: 2025 International Conference on Software, Telecommunications and Computer Networks, SoftCOM, 2025, pp. 1–6.

[25] L. Breiman, Random forests, Mach. Learn. 45 (1) (2001) 5–32.

[26] A. Liaw, M. Wiener, et al., Classification and regression by randomforest, R News 2 (3) (2002) 18–22.

[27] Ericsson, Mobile traffic forecasts, ericsson mobility report, 2025, https://www.ericsson.com/en/reports-and-papers/mobility-report/dataforecasts/mobile-traffic-forecast.

[28] A. Marotta, C. Centofanti, F. Graziosi, M. Quagliotti, M. Agus, On the techno-economic viability of coherent P2MP DSCM for mobile network fronthaul, in: Photonic Networks and Devices, Optica Publishing Group, 2025, pp. NeW4B–3.

[29] C. Castro, P. Torres-Ferrera, M. Sezer Erkilinç, J. Sime, G. Parisi, J. Pedro, M. Quagliotti, M. Porrega, D. Hillerkuss, C. Fludger, et al., Power consumption considerations of coherent transceivers in filterless point-to-multipoint metro-aggregation networks with digital subcarrier multiplexing, J. Opt. Commun. Netw. 17 (6) (2025) 526–542.

Sanwal Zeb is a Doctoral candidate in Department of Electronics and Telecommunications at Politecnico di Torino, Italy, funded by the EU Horizon Europe Marie-Curie (MSCA) project "NESTOR". He is a member of the PLANET research group, where his work focuses on digital twin of physical layer for short reach scenarios Metro-Access convergence and Next Generation Optical Networks. He hold a Master's degree in Electrical Engineering from the National University of Sciences & Technology (NUST), Pakistan. He worked in China Mobile Pakistan (CMPak), a leading Telco, specializing in network planning and design. He also served as Assistant Manager for Cloud & Fixed Networks at Telenor Pakistan.

Ahtisham Ali is a PhD candidate in Electrical, Electronics, and Telecommunication Engineering at the Department of Electronics and Telecommunications, Politecnico di Torino, Italy, under the Marie Curie Project EWOC (Enabling Wireless and Optical Communication Converged). He is currently affiliated with the company Consorzio Top-IX and the PLANET Research Group at DET, Politecnico di Torino. He received his Bachelor's and Master's degrees in Physics from Government College University, Faisalabad, Pakistan. His research interests include converged communication systems, wireless communication, and optical networks.

Imran Chowdhury Dipto is currently working as an MSCA fellow under the NESTOR project as part of his PhD under the Supervision of Prof. Vittorio Curri at the PLANET lab in collaboration with NOKIA Portugal at Politecnico di Torino, Italy. Imran obtained his postgraduate degree from Manchester Metropolitan University, UK with an optional placement year. He received a Bachelor's Degree in Information Technology from the University of Derby, UK. His research interests include Machine Learning, Deep Learning and their real-world applications in Optical Communications.

Andrea Rosso (Member IEEE) is a Ph.D. candidate in Telecommunications Engineering at Politecnico di Torino, Italy, where he also received his M.Sc. degree with a thesis on transceiver modeling and characterization for optical communication systems. He is a member of the PLANET research group, where his work focuses on enhancing optical networks through the integration of open network control architectures with physical layer awareness. His research interests include optical channel modeling, device characterization, and the study of underlying physical phenomena affecting optical transmission.

Muhammad Umar Masood received his Ph.D. in Optical Communication from Politecnico di Torino, Italy, in 2024, where he is currently a Postdoctoral Research Fellow with the PLANET research team. His work focuses on machine learning for optical networks, including QoT estimation, fault detection, and predictive performance modeling. He also investigates multi-band photonic integrated switching for scalable optical transport, with applications in autonomous control, intelligent monitoring, and energy-efficient 5G/6G backbone networks. His research interests span optical transport systems, photonic integration, AI/ML for networking, and disaggregated infrastructures.

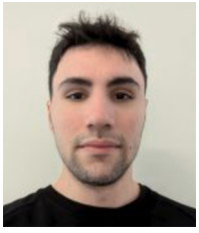

Riccardo Schips is a Ph.D. candidate at Politecnico di Torino and a member of the PLANET research group. His research, funded by the ALLEGRO project, focuses on the control and telemetry of disaggregated and partially disaggregated optical networks. In particular, he is exploring the integration of novel devices and software interfaces to enable seamless interoperability across all components and services within the optical infrastructure.

Renato Ambrosone obtained a bachelor's degree in Computer Science Engineering at the Università degli Studi della Campania in 2019. In 2021, he earned a master's degree in Computer Networks and Cloud Computing at Politecnico di Torino. He is currently pursuing a PhD at Politecnico di Torino within the Optical Communications (OptCom) research group. His research concentrates on the orchestration and control of disaggregated and partially disaggregated optical networks

Stefano Straullu received the Graduate degree in telecommunications engineering from Politecnico di Torino, Turin, Italy, in 2005, with a thesis on the realization and testing of opto-electronic subsystems for packet-switched optical networks, completed in the PhotonLab of Istituto Superiore Mario Boella, Turin; and the Ph.D. degree in electronics and communications engineering from Politecnico di Torino, Turin, with a thesis on next-generation passive optical networks. Since May 2009, he has been a Researcher at the Istituto Superiore Mario Boella (now LINKS Foundation) of Turin in the Optical Communication Unit. His main research interests are focused on the performance analysis and design of optical transmission systems and the application of digital signal processing techniques in optical links. He has published more than 120 journal and conference papers.

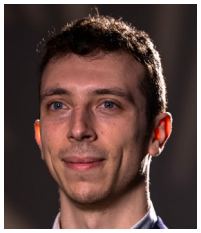

Francesco Aquilino received the Master Degree in Communication and Computer Network Engineering from Politecnico di Torino in 2019. He worked at Skylogic SPA (Eutelsat group) and in TIM. From March 2022 he works as a Researcher at LINKS Foundation and collaborates with POLITO and companies (Cisco Photonics, SM-Optics). His main research area are Optical Networks, Optical devices, Optical Communication and Fiber Sensing.

Antonino Nespola received the M.S. (1995) and Ph.D. (2000) in Electrical Engineering from Politecnico di Torino, Italy. From 1997 to 1998 he was a visiting researcher at the Photonics Laboratory, University of California, Los Angeles. From 1999 to 2003 he coordinated the Optical Systems Unit at Corning OTI, working on high-speed modulators. In 2003 he joined Pirelli Labs to develop nanotechnology-enabled integrated optical systems in collaboration with the MIT Microphotonics Center. Since 2004 he has led the Optical Communications Unit at the LINKS Foundation and serves as Adjunct Professor at Politecnico di Torino. In 2023 he joined the scientific coordination committee of the VCSELence Centre of Excellence. Since 2024 he has been Director of the CNIT Research Unit 'Fondazione LINKS', affiliated with CNIT (National Inter-University Consortium for Telecommunications). His interests include DSP for coherent optics, nonlinear optics, quantum communications, and modeling/experiments in optical systems. He has authored 190+ publications and holds five US/EU patents.

João Pedro holds a M.Sc. and Ph.D. degrees in Electrical and Computer Engineering from Instituto Superior Técnico (IST), University of Lisbon, Portugal. He was a research engineer and a system architect at Nokia Siemens Networks and Coriant and he is currently a senior principal engineer at Infinera, now part of Nokia. He leads the Optical Architecture Group in Lisbon and he is responsible for the design of capacity planning algorithms for next-generation optical and multi-layer networks and supporting performance and techno-economic investigations of future-looking network architectures. He holds 13 patent applications and has co-authored over 350 publications in international conferences and journals. He has also participated in several EU-funded projects, currently acting as the coordinator of Horizon NESTOR, has been a lecturer of courses on network planning and transport networks, and has co-supervised multiple PhD and MSc students. His current research interests include high-capacity and reliable optical networks, flexible metro-aggregation architectures, routing and spectrum assignment, multilayer optimization and machine learning applications. He is also a permanent staff member of Instituto de Telecomunicações, a senior member of the IEEE and a member of OPTICA.

Antonio Napoli (Senior Member, IEEE) received the M.Sc. degree in Telecommunication Engineering in 2002 and the Ph.D. degree in Electrical Engineering in 2006, both from the Politecnico di Torino, Italy. He is currently leading the Strategic Architecture Engineering Group at Nokia, where he is responsible for advanced research and system design for next-generation optical communication systems. From 2006 to 2025, he held research and leadership positions at Siemens, Nokia Siemens Networks, Coriant, and Infinera, focusing on DSP algorithms, coherent optical technologies, wideband optical transmission, machine learning for optics, and intelligent transceiver architectures from access to submarine applications. He has co-authored more than 260 peer-reviewed publications, holds seven patents, and has served as Guest Editor for the Journal of Optical Communications and Networking and the Journal of Lightwave Technology. He was recognized as an OSA/IEEE Top Reviewer in 2016 and 2017, and his contributions are widely acknowledged in both the industrial and academic optical communications communities.

Alessandro Galardini is a network Engineer at Consorzio Top-IX. He is passionate about technology, graduated in Telecommunication Engineering from Politechnico di Torino specializing in electromagnetic compatibility and radio systems. With Top-IX since 2009, his main role is in the design, operation, and monitoring of Top-IX network infrastructure at all levels (passive, optical, switch/routing, and DC infrastructure).

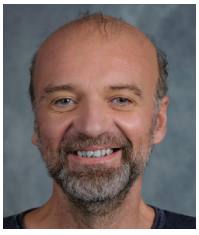

Vittorio Curri received the Laurea degree in electrical engineering cum laude in 1995, and the Scientific Doctoral degree in optical communications in 1999, both from Politecnico di Torino (PoliTO), Italy. He is a founder member of the OptCom group and PhotonLab at PoliTO and currently leads the Open PLANET Lab. He is a Full Professor with the Department of Electronics and Telecommunications, PoliTO. He has been a Visiting Scholar at the Stanford University and UC at Santa Barbara. His major research interests are in fiber transmission modeling, including nonlinearities, transmitter and receiver optimization, design strategies for optical links, Raman amplification and simulation and modeling of optical communication systems. Prof. Curri is one of the early promoters of using multi-band transmission in single-mode fibers and he is also active in the analysis of the impact of physical layer on networking for planning and control of open and disaggregated optical networks based on vendor-neutral digital-twins, including AI techniques. He is the Scientific Chair of the GNPy open source project within the industry consortium Telecom Infra Project. He is also an early promoter of using the optical network as-a-sensor. Prof. Curri is IEEE Fellow, OPTICA Fellow and AAIA Fellow, and is the co-author of more than 550 scientific publications.