Sustainable 6G Optical Any-Haul with Coherent Pluggables in Converged Metro-Access

Sustainable 6G Optical Any-Haul with Coherent Pluggables in Converged Metro-Access

Ahtisham Ali*†, Sanwal Zeb†, Muhammad Umar Masood†, Andrea Rosso†, Gulmina Malik†, Renato Ambrosone†, Riccardo Schips†, Michela Pollone*, Stefano Straullu‡, Francesco Aquilino‡, Joao Pedro§, Antonio Napoli¶, Alessandro Galardini*, Vittorio Curri†
*Consorzio TOP-IX, Torino, Italy
†Dipartimento di Elettronica e Telecomunicazioni, Politecnico di Torino, Torino, Italy
‡Fondazione LINKS, Torino, Italy
§Nokia, Optical Networks, Carnaxide, Portugal
¶Nokia, Optical Networks, Munich, Germany
Email: ahtisham.ali@top-ix.org

***Abstract*—Future 6G networks will require transport solutions capable of meeting the stringent latency and bandwidth demands of any-haul traffic. This paper investigates the potential of converged metro-access networks offered as a service, within an Optical Network as-a-Service (ONaaS) framework, to address these requirements. Leveraging a physical-layer abstraction of a digital subcarrier multiplexing (DSCM)-based transceiver combined with statistical path analysis, we evaluate three key aspects: (i) route feasibility under split 7.2 latency and BER constraints for different Distributed Units (DUs) placement strategies; (ii) the resource footprint of Point-to-Point (P2P) versus Point-to-multipoint (P2MP) optical aggregation, accounting for transceivers, routers, and port occupancy at DU sites; and (iii) power consumption at both DU and network scale across varying spectral loads. Results show that increasing DU density significantly reduces the number of non-feasible routes, while P2MP architectures lower DU-site transceiver, router, and power requirements compared to P2P. These findings provide quantitative and qualitative gains for a scalable and energy-efficient any-haul design in 6G converged metro-access networks enabled by ONaaS.**

***Index Terms*—x-haul, metro-access, disaggregated RAN, network as-a-service, ONaaS, P2P, P2MP, Digital Subcarrier Multiplexing**

I. Introduction

The transition to 6G is anticipated to trigger unprecedented growth in mobile traffic alongside increasingly heterogeneous service demands, imposing strict requirements on capacity, latency, and power efficiency. According to the Ericsson Mobility Report [1], global mobile traffic is projected to grow at a compound annual growth rate (CAGR) ranging from 18.7% in 2025 to 15.4% in 2030. For long-term modeling, we adopt the equivalent average CAGR of 17.2% to extend this trend over a ten-year projection horizon. Such sustained growth places an unsustainable burden on conventional transport architectures, particularly P2P fiber deployments, which scale inefficiently in dense radio access networks (RAN). To mitigate these scalability limitations, converged metro-access optical networks offered as-a-Service (ONaaS) have emerged as a promising alternative, leveraging reconfigurable optical infrastructures and advanced pluggable transceivers (TRx)s to enable flexible and cost-efficient transport [2]. In particular, P2MP optical aggregation supported by DSCM provides statistical multiplexing gains, reducing hardware footprint and power consumption [3], while ensuring strict quality-of-transmission (QoT) guarantees [4].

In addition to the projected traffic surge, 6G will also introduce novel service paradigms, including immersive extended reality (XR), holographic communications, and ultra-reliable low-latency communications (URLLC), further intensifying the capacity and latency requirements on RAN transport [5]. These services impose highly diverse any-haul requirements across fronthaul, midhaul, and backhaul segments.

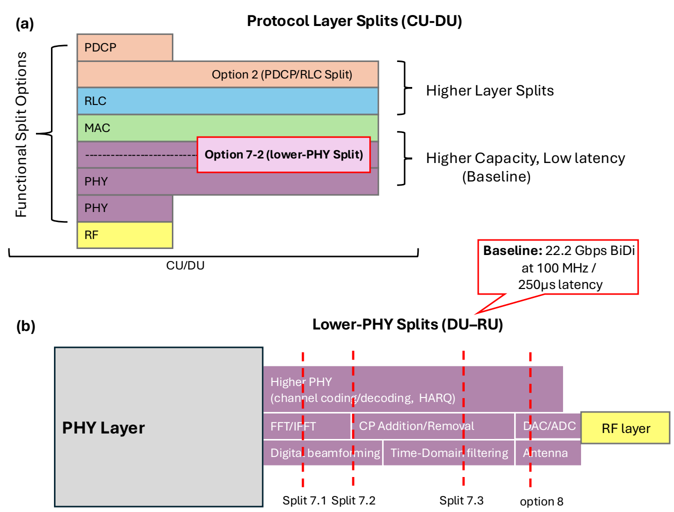

Fig. 1: Functional split options in the RAN: (a) CU–DU protocol layer splits highlighting Option 7.2 (baseline), (b) DU–RU lower-PHY split details including Splits 7.1, 7.2, 7.3, and Option 8.

Functional split 7.2, for instance, demands end-to-end latency below 250 µs, while simultaneously sustaining multi-gigabit throughput per radio unit (RU) [6, 7]. When combined with wider radio bandwidths, larger multiple-input multiple-output (MIMO) configurations, and increasingly dense small-

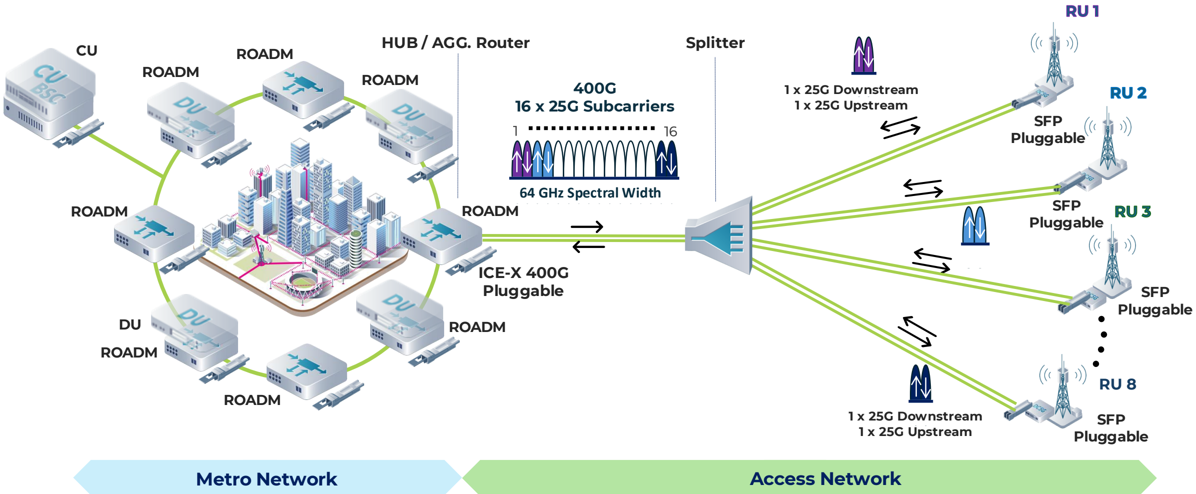

Fig. 2: Converged metro-access optical network as-a-service using P2MP TRx

cell deployments, these constraints result in huge capacity and latency demands.

Converged metro-access networks offered as-a-Service provide an attractive alternative, enabling operators to leverage existing fiber assets for RAN connectivity [8]. This approach can lower capital and operational expenditures, improve resource utilization, and facilitate dynamic capacity allocation [9, 10]. In particular, P2MP optical aggregation supported by advanced TRx technologies, such as DSCM TRx, allows statistical multiplexing of traffic from multiple RUs while maintaining strict QoT guarantees [11].

In this context, the present work addresses two complementary research questions. First, we evaluate the feasibility of optical paths within a converged metro-access topology, quantifying the impact of DU density on the share of feasible routes under split 7.2 latency and bit-error-rate (BER) constraints. Second, motivated by the projected traffic growth, we extend the analysis to hardware and power consumption aspects, comparing P2P and P2MP optical aggregation in terms of TRx counts at DU sites and network-wide, router and port requirements, and total power usage under varying spectral loads. By jointly considering route feasibility, hardware footprint, and power efficiency, this paper provides quantitative guidance for planning scalable, cost-efficient, and sustainable any-haul solutions in future 6G converged metro-access networks.

II. Capacity Requirements from RAN

The evolution toward 6G will impose significantly higher bandwidth demands and stricter latency budgets on any-haul networks (fronthaul, midhaul, backhaul) [5]. One of the key contributing factors is the functional split between Central Unit (CU) and DU, as specified by 3GPP [6] and summarized in Fig. 1(a). Higher-layer splits reduce transport load on transport links but increase functional complexity at antenna sites. In contrast, lower-layer splits such as split 7.2 impose a strict latency budget of $250\,\mu s$ one-way, yet enable efficient coordination across the full protocol stack [6, 7]. For example, a split 7.2 configuration with 100 MHz bandwidth, 64-QAM modulation, and 32 antenna ports requires approximately 22.2 Gbps of bi-directional capacity per RU [12]. Looking ahead to 6G, scaling radio bandwidth to 400 MHz and 1 GHz increases the per-RU demand to about 117 Gbps and 293 Gbps, respectively [5].

The anticipated growth evolution highlights that while dedicated P2P fiber per RU could in principle accommodate higher traffic, it does so at the expense of proportional increases in TRxs, router ports, and associated power consumption. In dense small-cell deployments [13], this linear scaling rapidly becomes economically and operationally unsustainable. By contrast, shared metro-access infrastructures, such as P2MP optical aggregation, exploit statistical multiplexing to aggregate multiple RU traffic streams onto shared high-capacity optical channels. This enables the network to scale more efficiently with traffic growth while still meeting stringent any-haul requirements. In this work, we therefore adopt split 7.2 with a 22.2 Gbps per-RU demand as the baseline traffic model and build upon it to evaluate the feasibility, hardware requirements, and power consumption of P2P and P2MP optical aggregation.

A. ONaaS Architecture using P2MP Coherent Pluggable

P2MP TRxs enable simultaneous communication with multiple nodes. In DSCM enabled P2MP TRxs, the hub node modulates data onto optical subcarriers (SCs), which are broadcast to all leaf nodes. Each leaf extracts only its designated SCs through Nyquist multiplexing, while upstream transmission is achieved by reusing the same SCs to send data back to the hub or the aggregation site from a specific leaf [14].

Fig. 2 illustrates a converged metro-access network architecture enabled by DSCM TRxs. At the aggregation site, a

400G P2MP TRx generates 16 × 25 Gb/s Nyquist SCs using DP-16QAM modulation, where each SC occupies a spectral width of 4 GHz, providing an overall 64 GHz channel [15, 16]. These SCs are optically split and routed toward multiple leaf nodes using splitters/combiners. Each leaf node, corresponding to an RU, supports traffic up to 22.2 Gb/s in compliance with the functional split 7.2. Traffic from individual RUs is mapped onto 25 Gb/s downstream and upstream SCs and subsequently aggregated at the DU located in the metro node. This configuration enables cost efficient metro-access aggregation of RAN traffic without requiring intermediate regeneration. This approach integrates RAN fronthaul and metro-access transport while maintaining high spectral efficiency.

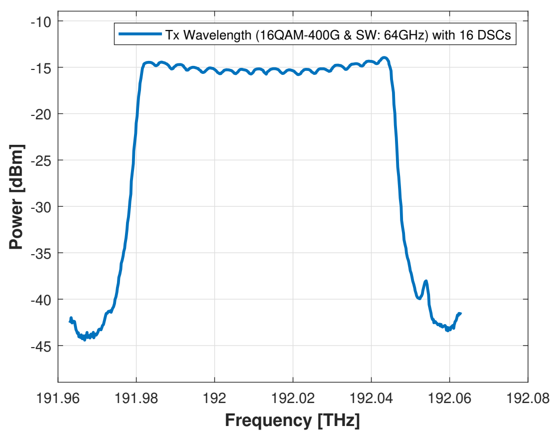

Fig. 3: OSA trace of the DP-16QAM signal having 16 Nyquist SCs

Fig. 3 shows the optical spectrum analyzer (OSA) trace of the DP-16QAM 400 Gb/s optical signal obtained during experimental back-to-back (B2B) characterization of the DSCM TRx. The signal consists of 16 Nyquist-shaped subcarriers (SCs), each occupying approximately 4 GHz, resulting in an overall spectral occupancy of about 64 GHz. During experimental B2B characterization, $\mathrm{SNR}_{\mathrm{TRX}}$, is extracted by varying the transmitted optical power using a variable optical attenuator (VOA). In parallel, OSA-based measurements are used to estimate the ASE noise power within the signal bandwidth, from which $\mathrm{SNR}_{\mathrm{ASE}}$ is obtained. This approach allows the combined impact of TRx impairments and ASE noise on the total received SNR, and consequently on the resulting BER, to be accurately assessed.

III. Simulation Scenario

The simulation framework is based on an experimental characterization of commercially available P2P and P2MP TRxs, respectively, from which we derived a performance model relating Received Optical Power (ROP) and TRx Signal-to-Noise Ratio (SNR). This abstraction enables statistical evaluation of a converged metro–access optical network under joint BER and latency constraints. The network under study, depicted in Fig. 4, consists of 44 optical nodes partitioned into metro (M) and access (A) segments. The simulated metro network

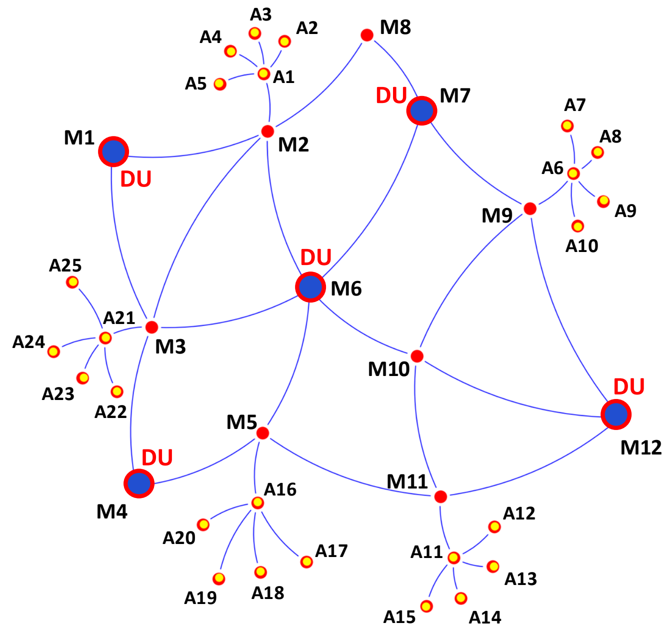

Fig. 4: Simulated network topology

is composed of reconfigurable optical add–drop multiplexers (ROADMs), interconnected by 19 fiber pairs with in-line amplification, and equipped with remotely controlled, interoperable DSCM pluggable TRxs that enable ONaaS capabilities in a converged metro-access environment. This results in an average metro node degree of 3.42, an average span length of 6.9 km, and a maximum span length of 12 km. Each access cluster includes four access nodes, each serving two RUs. Fiber spans are modeled using standard single-mode fiber (SSMF) with 0.25 dB/km attenuation, while ROADM nodes are assumed with an insertion loss of 10 dB, and EDFAs are employed where necessary to compensate span and node losses, assuming a 6 dB noise figure [17].

The topology captures dense metro-access deployment commonly adopted by regional operators, with European-style span lengths and access loops, and is suitable for assessing route feasibility for RAN transport [17]. The five considered DU placement configurations are summarized in Table I.

TABLE I: DU placement scenarios considered

DU placement (Metro nodes)	Number of DUs
M1	1
M1, M4	2
M1, M4, M12	3
M1, M4, M6, M12	4
M1, M4, M6, M7, M12	5

The TRx was characterized following the proposed method in [18, 19], based on simple curve fitting over experimental measurements using the relation between $\mathrm{SNR}_{\mathrm{TRx}}$ and ROP. A random statistical loss distribution was applied on the access segment to emulate variability due to connectors, splitters, and splicing [20, 21]. To capture the variability of statistically modeled random access losses, 1000 independent Monte Carlo simulations were run. For every route, the end-to-end BER was estimated according to

$$\mathrm{BER}_{16\mathrm{QAM}} = \frac{3}{8}\,\mathrm{erfc}\left(\sqrt{\frac{\mathrm{SNR}}{10}}\right), \tag{1}$$

where the effective SNR accounts for both optical channel and TRx penalties and is given as:

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

$$\mathrm{GSNR} = \frac{P_i}{P_{\mathrm{ASE}} + P_{\mathrm{NLI}}} \tag{3}$$

Here, GSNR is derived from accumulated fiber and amplifier impairments, including amplified spontaneous emission (ASE) and nonlinear interference (NLI) while TRx contribution is given as:

$$\mathrm{SNR}_{\mathrm{TRx}} = \frac{N \cdot \mathrm{ROP}_{\mathrm{mW}}}{\mathrm{ROP}_{\mathrm{mW}} + D}, \quad [18,\ 19] \tag{4}$$

where N and D are fitting constants and depend upon the type of TRx with the corresponding modulation format used.

Latency is computed by mapping the physical path length of the fronthaul segment to a propagation delay of $5\,\mu\mathrm{s/km}$, which is the dominant factor in latency measurement and is consistent with SSMF characteristics [22]. A route is considered feasible if it satisfies both a BER threshold of 10^{-3} and a 3GPP split 7.2 latency constraint of $250\,\mu\mathrm{s}$.

IV. Results and Discussion

Based on the simulation framework described earlier, the following section describes route feasibility, hardware scalability, and power efficiency for P2P and P2MP converged metro-access architectures. The route feasibility was assessed under BER and strict latency constraint using commercially available P2MP TRxs.

Fig. 5 presents the feasibility analysis with the maximum hop size limited to nine. As the number of DUs in the metro network increases, two main effects become evident. First, the average distance between access nodes and their potential destination DUs decreases, allowing more routes to satisfy the latency requirement. Second, a higher DU density naturally expands the set of available source-destination pairs, thereby increasing the total number of possible routes.

Consequently, both the total and feasible route counts rise with higher DU density, the share of feasible paths increases from 60.4% with one DU to 73.5% with five DUs. This improvement stems from shorter access–metro spans and a richer set of routing alternatives, which together enhance the likelihood of meeting both BER and latency requirements. To further assess the impact of hop count on route feasibility, Fig. 5a presents results for paths with hop size $\leq$ 5. In this case, all routes satisfy both BER and latency constraints, yielding 100% feasibility irrespective of DU placement. Conversely, Fig. 5b illustrates longer paths with hop sizes $>$ 5, where feasibility drops considerably from 57.6% with one DU

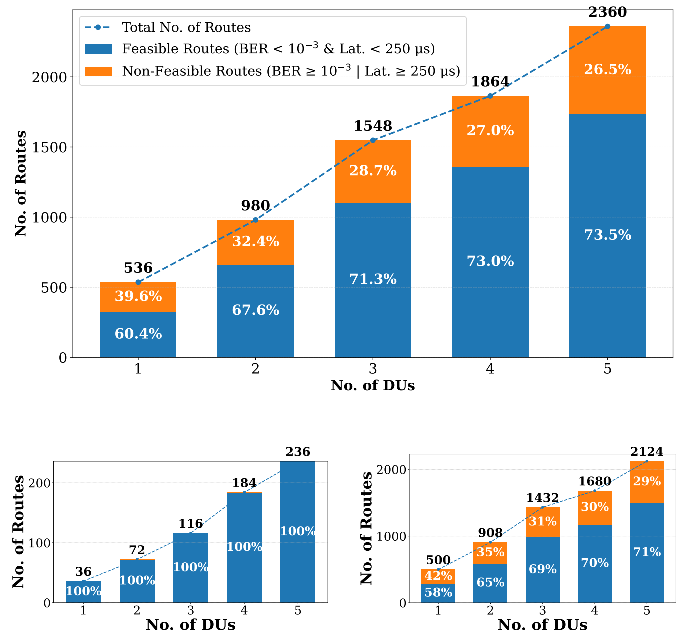

(a) (b)

Fig. 5: Feasibility analysis with DU densities: overall with hop size $\leq$ 9 (top), (a) hop size $\leq$ 5, and (b) hop size 6–9

to 70.5% with five DUs. These results highlight the strong dependence of end-to-end performance on optical path length, emphasizing that longer metro–access routes are more prone to QoT and latency degradation.

While the preceding feasibility analysis addressed physical-layer constraints such as BER and latency, it did not capture the network-level implications of deploying the P2MP architecture under realistic traffic conditions. To bridge this gap, the following analysis also evaluates the techno-economic advantages of P2MP over P2P, highlighting reductions in hardware requirements and power consumption as traffic demand grows. The study extends this assessment over a ten-year horizon using mobile traffic projections from the Ericsson Mobility Report [1], which forecasts annual growth rates of 18.7%, 18%, 17.3%, 16.5%, and 15.4% between 2025 and 2030. These correspond to an average compound growth rate of 17.2%. For the subsequent five years (shaded regions represented in Figs. 6 and 7), we extrapolate by applying this average rate, providing a ten-year projection horizon. Starting from a traffic load of 1 Tb/s in year 0, this model results in nearly a fivefold increase by year 10.

This projection framework enables us to quantify how variable network elements, such as TRxs, router ports, and power consumption evolve under increasing demand, while fixed infrastructure is kept constant. The following subsections present a stepwise analysis: first, the scaling of TRxs at DU sites and across the network; second, the router and port requirements; and finally, the power consumption of P2P versus P2MP. The scaling of TRxs with projected traffic growth is shown in Fig. 6, comparing P2P and P2MP under

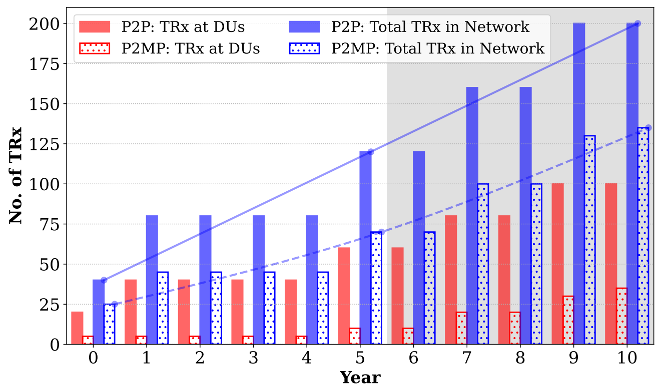

Fig. 6: TRx requirements at DU sites and network-wide for P2P and P2MP under 10-year traffic growth

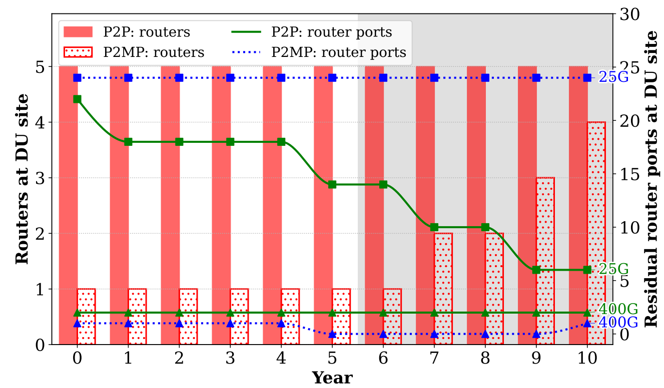

Fig. 7: Router demand and residual port availability at DU sites for P2P and P2MP under 10-year traffic growth

the ten-year forecast. For this analysis, only variable elements that evolve with traffic, namely TRxs are considered, while fixed infrastructure remains unchanged.

At year 0, P2P requires 4 TRx at each DU site and 40 across the network, whereas P2MP requires only 1 TRx per DU (5 total) and 25 network-wide. This efficiency gap widens in the early years: by year 4, P2P demands 8 TRx per DU, while P2MP still operates with only 1. As traffic continues to grow, the difference at the DU becomes striking. By year 10, P2P requires 27 TRx per DU (135 network-wide), whereas P2MP requires only 7 per DU (35 across all DUs), a nearly fourfold reduction at the DU level, even though both architectures converge to similar network-wide totals under very high traffic.

The curves exhibit a staircase behavior, reflecting the discrete granularity of DSCM TRxs: capacity is added in chunks rather than continuously. In P2P, these steps occur frequently and at low load thresholds, forcing rapid growth in DU-site TRx. In P2MP, SC multiplexing allows multiple RU flows to share the same DU-side TRx, so capacity additions occur less often and at higher aggregate loads. This results in significant savings at the DU even though total network-wide TRx eventually converges under very high traffic.

The key advantage of P2MP lies not only in reducing the total number of TRxs in the long run, but in optimizing their distribution. By lowering DU-site TRx counts by nearly a factor of four compared to P2P at year 10, P2MP alleviates the most constrained locations in the network, defers costly upgrades of routers and ports, and reduces the operational footprint. This makes P2MP a more sustainable and practical solution for scaling converged metro-access networks to support 6G any-haul. The router and port requirements at DU sites were evaluated under the same ten-year traffic projection. For this analysis, we assumed that each DU site is equipped with a router providing 28 ports: 24×25G, 2×100G, and 2×400G. In P2P case, the 25G ports are consumed to connect TRxs, leaving the 400G interfaces unused. In contrast, P2MP case leverages the 400G ports, enabling efficient aggregation of multiple RU flows over fewer high-capacity links.

Fig. 7 shows the number of routers required per DU site (bars, left axis) and the residual free ports per router (lines, right axis) as traffic scales. In the P2P scenario, router demand scales rapidly with traffic: by year 0, a single router is sufficient, but as TRx counts grow, additional routers must be added at DU sites to provide enough 25G interfaces. By year 10, several routers are needed at individual DUs to keep up with traffic, reflecting both port and slot pressure.

In P2MP, a single router with 400G interfaces can sustain traffic growth far longer. Even as the number of TRxs increases, the higher-capacity ports delay the need for deploying additional routers. Moreover, the residual port curves highlight that P2MP routers maintain unused capacity across the decade, which provides headroom for supporting traffic surges beyond the projected growth without immediate hardware upgrades. In contrast, P2P routers quickly reach saturation of their 25G ports, leaving little spare capacity for unanticipated demand. This analysis underscores that P2MP not only minimizes the number of routers required at DU sites but also ensures a more favorable port utilization profile. The combination of fewer routers and sustained residual port availability translates into lower capital costs, reduced operational complexity, and greater flexibility in absorbing traffic uncertainties over time.

The power consumption of the network was analyzed by comparing P2P and P2MP under different levels of spectral utilization of the DSCM TRx. For P2P, each optical connection is assumed to use a 25G SFP module consuming 3 W per TRx. In contrast, P2MP employs DSCM-enabled CFP2 modules with a rated power consumption of approximately 22 W, consistent with commercially available implementations. In P2MP, the digital SC structure enables flexible allocation of capacity: 25% utilization corresponds to 4 out of 16 SCs active, 50% to 8 SCs, 75% to 12 SCs, and 100% to all 16 SCs. This allows us to model how power efficiency evolves as traffic load increases.

Fig. 8 shows normalized power consumption relative to the P2P case at full (100%) load. A clear trend emerges: At very low utilization (25%), P2MP consumes more power than P2P, since the overhead of maintaining a multi-carrier system outweighs the benefits of aggregation when few SCs

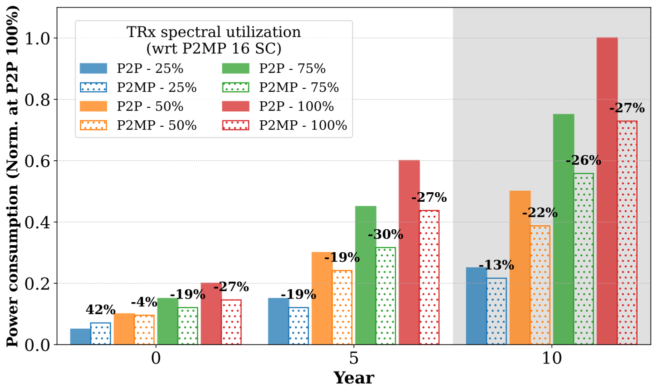

Fig. 8: Norm. power consumption of P2P and P2MP across different levels of spectral utilization of DSCM TRx

are active. However, as utilization increases, P2MP rapidly becomes more efficient. At 50% load, P2MP already reduces power consumption by nearly 20% compared to P2P, and this advantage grows at higher loads: up to 27-30% lower consumption at 75-100% capacity utilization.

This behavior reflects the inherent trade-off of DSCM-based TRxs, where the baseline power draw remains constant regardless of the number of active SCs, and efficiency improves as utilization increases. In contrast, P2P scales linearly with traffic, since each additional flow requires a dedicated TRx pair. Consequently, P2MP provides substantial long-term power savings under sustained traffic growth-once utilization exceeds 50%, it consistently outperforms P2P by 20-30%. This establishes P2MP as a more power-efficient and sustainable architecture for converged metro-access networks supporting dense 6G RAN traffic.

V. Conclusion

This work presented a quantitative assessment of a converged metro–access network for 6G any-haul, comprising feasibility analysis for P2MP and a techno-economic evaluation showing the comparative advantage of P2MP over P2P architectures. The results confirm that higher DU densities enhance end-to-end route feasibility by better satisfying functional split 7.2 latency and BER constraints. Moreover, while P2P optical aggregation ensures direct connectivity, P2MP architectures provide superior scalability by reducing the number of TRxs, router ports, and associated power footprint at DU sites and across the network. These efficiency gains become increasingly significant under higher spectral loads, underscoring their role as a key enabler of network scalability and interoperability. Overall, the findings demonstrate that P2MP enabled converged metro–access ONaaS offers a practical pathway towards power efficient, cost effective, and latency compliant any-haul solutions in future 6G deployments.

Acknowledgments

Ahtisham Ali's work is fully funded by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie project EWOC under GA No. 101073265. This work has also received support from the Marie Skłodowska-Curie NESTOR project under GA No. 101119983, the EU Horizon Project ALLEGRO GA No. 101092766, the project PNRR-NGEU (MUR–DM117/2023), and the EU Next Gen. EU partnership on 'Telecommunications of the Future' (PE00000001 – program 'RESTART').

References

[1] Ericsson, "Mobile Traffic Forecasts, Ericsson Mobility Report," Ericsson, Tech. Rep., 2025, accessed: 2025-08-29. [Online]. Available: https://www.ericsson.com/en/reports-and-papers/mobility-report/dataforecasts/mobile-traffic-forecast

[2] A. Ali *et al.*, "Evaluating Metro-Access Optical Networks as-a-Service for Flexible RAN X-haul," in *2025 IEEE Conference on Standards for Communications and Networking (CSCN)*, 2025, pp. 1–6.

[3] J. Bäck *et al.*, "CAPEX Savings Enabled by Point-to-Multipoint Coherent Pluggable Optics Using Digital Subcarrier Multiplexing in Metro Aggregation Networks," in *2020 European Conference on Optical Communications (ECOC)*, 2020, pp. 1–4.

[4] 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.

[5] A. Larranaga *et al.*, "Fronthaul/midhaul networks: Capacity and latency requirements imposed by 6G disaggregated RANs," *IEEE Communications Magazine*, pp. 1–8, 2025.

[6] 3GPP, "Study on CU-DU lower layer split for NR; (release 15)," 3GPP, Tech. Rep. TR 38.816 V15.0.0 (2017-12), 2017.

[7] J. Pérez-Romero *et al.*, "A tutorial on the characterisation and modelling of low layer functional splits for flexible radio access networks in 5G and beyond," *IEEE Communications Surveys & Tutorials*, 2023.

[8] A. Ali *et al.*, "Evaluating metro-access optical networks as-a-service for flexible ran x-haul," in *2025 IEEE Conference on Standards for Communications and Networking (CSCN)*, 2025, pp. 1–6.

[9] Z. Vujicic *et al.*, "Towards virtualized optical-wireless heterogeneous networks," *IEEE access*, vol. 12, pp. 87 776–87 806, 2024.

[10] G. Rizzelli *et al.*, "Analysis and experimental demonstration of possible architectures for future coherent metro+ PON converged networks," *Journal of Optical Communications and Networking*, vol. 17, no. 2, pp. A142–A154, 2025.

[11] D. Welch *et al.*, "Digital subcarrier multiplexing: Enabling software-configurable optical networks," *JLT*, vol. 41, no. 4, pp. 1175–1191, 2023.

[12] 3GPP, "Radio access architecture and interfaces (release 14)," 3GPP, Tech. Rep. 3GPP TR 38.801 V14.0.0 (2017-03), 2017.

[13] A. Rafel *et al.*, "Comparing optical transport technologies for x-hauling 5G small cells in the sub-6 GHz," *JOCN*, 2022.

[14] D. Welch *et al.*, "Point-to-multipoint optical networks using coherent digital subcarriers," *Journal of Lightwave Technology*, vol. 39, no. 16, pp. 5232–5247, 2021.

[15] J. Bäck *et al.*, "CAPEX savings enabled by point-to-multipoint coherent pluggable optics using digital subcarrier multiplexing in metro aggregation networks," in *2020 European Conference on Optical Communications (ECOC)*. IEEE, 2020, pp. 1–4.

[16] M. Mohammad Hosseini *et al.*, "Multi-period planning in metro-aggregation networks exploiting point-to-multipoint coherent transceivers," *Journal of Optical Communications and Networking*, vol. 15, no. 3, pp. 155–162, 2023.

[17] A. Ali *et al.*, "Statistical analysis of end-to-end route feasibility in converged metro–access optical networks," *IEEE Photonics Technology Letters*, 2026.

[18] A. Rosso, "Modeling and controlling optical transponder white-boxes based on the physical layer digital twin," Master's thesis, PoliTo, 2024.

[19] T. Mano *et al.*, "Modeling the input power dependency of transceiver ber-onsr for qot estimation," in *OFC*. Optica Publishing Group, 2024, pp. M1H–4.

[20] G. Simon *et al.*, "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.

[21] A. Ali *et al.*, "Capacity Assessment in Converged Metro-Access Optical Networks for end-to-end RAN Fronthaul," in *2025 25th Anniversary International Conference on Transparent Optical Networks (ICTON)*. IEEE, 2025, pp. 1–4.

[22] A. Ali *et al.*, "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)*. IEEE, 2025, pp. 1–6.