

Assessing Techno-Economic Gains of Next-Generation Coherent Point-to-Multipoint Data Center Interconnects

Original

Assessing Techno-Economic Gains of Next-Generation Coherent Point-to-Multipoint Data Center Interconnects / Zeb, S., Masood, M.U., Ali, A., Schips, R., Ambrosone, R., Dipto, I.C., Pedro, J., Napoli, A., Curri, V.. - (2026). (2026 17th International Conference on Network of the Future (NoF) Roma (Ita) 29 September - 02 October 2026).

Availability:

This version is available at: 11583/3015547 since: 2026-09-15T21:41:56Z

Publisher:

IEEE

Published

DOI:

Terms of use:

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

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

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

(Article begins on next page)

Assessing Techno-Economic Gains of Next-Generation Coherent Point-to-Multipoint Data Center Interconnects

Sanwal Zeb*, Muhammad Umar Masood*, Ahtisham Ali*[†], Riccardo Schips*,
Renato Ambrosone*, Imran Chowdhury Dipto*, João Pedro^{‡§}, Antonio Napoli[¶], Vittorio Curri*

**Dipartimento di Elettronica e Telecomunicazioni, Politecnico di Torino, Italy*

[†]Consorzio TOP-IX, Torino, Italy

[‡]Nokia, Optical Networks, Carnaxide, Portugal

[§]Instituto de Telecomunicações, IST, Portugal

[¶]Nokia, Optical Networks, Munich, Germany

Email: sanwal.zeb@polito.it

Abstract—Rapid growth in AI-driven applications and 5G services is increasing traffic demand between Internet Service Providers (ISPs) and Data Centers (DCs), stressing the scalability of conventional point-to-point (P2P) optical architectures. Regional broadband traffic has sustained approximately 14% CAGR over the last decade, driving continued expansion of optical infrastructure and increasing CAPEX and OPEX. To improve scalability and reduce infrastructure expenditure, this paper investigates an alternate solution of coherent point-to-multipoint (P2MP) pluggables enabled by Digital Sub-Carrier Multiplexing (DSCM). The evaluation is based on real ISP-DC interconnect deployments and traffic trends over 10-years traffic projection. By allowing multiple ISP interconnects to share coherent transceivers through subcarrier-level resource allocation, P2MP reduces the need for proportional expansion of transceivers and physical links. Results show that P2MP reduces total transceiver count by up to 84% and aggregation-side transceivers by up to 93% over the 10-year projection period. Interconnect density is reduced by up to 69%, while normalized CAPEX decreases by 51-64% depending on the P2MP transceiver cost. P2MP also improves energy efficiency by up to 35% in terms of power per transported Gbps. These results demonstrate that coherent P2MP architectures can substantially reduce infrastructure scaling requirements in heterogeneous multi-operator ISP-DC environments.

Index Terms—Point-to-Point (P2P), Digital Subcarrier Multiplexing (DSCM), Point-to-multipoint (P2MP)

I. INTRODUCTION

Global internet traffic is increasing rapidly due to AI, cloud, and 5G-driven applications. Regional broadband traffic is projected to exceed 523 EB per year by 2030, corresponding to nearly a $3\times$ increase [1]. Over the last decade, traffic demand has sustained approximately 14% compound annual growth rate (CAGR), requiring ISP-DC interconnects to scale toward multi-Tb/s capacities while maintaining sustainable CAPEX and OPEX.

Conventional Point-to-Point (P2P) optical interconnects scale inefficiently under these conditions. Each new ISP

connection requires dedicated transceivers, wavelengths, and switching resources at the DC aggregation layer. As a result, transceiver count and interconnect density increases with traffic demand. In existing multi-operator deployments, the coexistence of heterogeneous ports (10G, 40G, 100G, ...), together with redundant links and Link Aggregation Group (LAG) configurations, further increases hardware cost, power consumption, and operational complexity [2]. These effects become particularly pronounced in ISP-DC interconnects, where traffic growth is highly asymmetric across operators and aggregation resources are frequently underutilized.

Digital Sub-Carrier Multiplexing (DSCM) enables a more flexible scaling model. A single 400G DSCM transceiver synthesizes up to 16 Nyquist Sub-Carriers (SCs) within a 64 GHz optical channel, where each SC is capable of carrying upto 25 Gbps using DP-16QAM modulation. Through passive optical splitting, a hub transceiver can distribute SCs dynamically across multiple ISP endpoints, forming a coherent Point-to-Multipoint (P2MP) architecture [3]. This enables SC level granular bandwidth allocation and reduces the need for proportional expansion of physical interconnects and transceivers as traffic increases. However, dynamic SC allocation introduces additional control-plane complexity, requiring Software-Defined Networking (SDN) for automated provisioning and resource orchestration [4].

Prior work on DSCM-based P2MP architectures has demonstrated techno-economic benefits in metro aggregation and converged access networks. In [5], P2MP has been evaluated using synthetic nationwide topologies, reporting up to 76% total cost of ownership (TCO) reduction under high traffic growth scenarios. Other studies investigated P2MP in 5G optical metro networks using normalized techno-economic models and synthetic traffic assumptions [6], while experimental demonstrations validated coherent DSCM transmission feasibility [3]. More recently, DSCM-driven P2MP architec-

tures have been studied for 6G X-haul converged metro-access networks, reporting significant transceiver and energy savings under long-term traffic projections [7]. Complementary work has also established normalized techno-economic modeling as an effective framework for evaluating optical network scaling strategies [8].

However, existing studies primarily rely on synthetic traffic models and homogeneous metro aggregation assumptions. Real ISP-DC interconnect deployments exhibit different operational characteristics, including fragmented port utilization, mixed-rate interface coexistence, asymmetric traffic scaling, and per-operator redundancy requirements. These characteristics create inefficient hardware scaling behaviors. Furthermore, no prior work has evaluated coherent P2MP interconnects using multi-operator ISP-DC deployments combined with long-term traffic projections.

To evaluate the techno-economic impact of coherent P2MP in operational interconnects, this paper investigates DSCM-driven coherent P2MP interconnects using operational multi-operator deployment data under a 10-year traffic growth projection. The evaluated deployment comprises 13 service providers across Tier 1 and Tier 2 peering structures with heterogeneous 10G and 100G interconnects and approximately 1 Tbps baseline aggregate traffic. A techno-economic framework compares conventional P2P and coherent P2MP architectures in terms of transceiver scalability, required fiber links, CAPEX, and energy efficiency in terms of transceiver power consumption. The study also highlights how fragmented utilization, heterogeneous interconnects, and asymmetric traffic growth make operational interconnect deployments well suited to subcarrier-level coherent aggregation.

The remainder of this paper is organized as follows. Section II discusses the limitations of legacy P2P interconnects. Section III presents the P2P and P2MP architectures. Section IV describes the operator deployment records and techno-economic framework. Section V presents the comparative results, and Section VI concludes the paper.

II. CHALLENGES AND LIMITATIONS

Conventional ISP-DC interconnects are primarily built using P2P optical links. Each ISP connects to the DC aggregation layer through dedicated transceivers and wavelength channels operating at fixed line rates. While operationally simple, this one-to-one connectivity model scales inefficiently under sustained traffic growth [2]. Every new ISP connection requires additional transceivers, switching ports, and optical resources, causing hardware growth to scale nearly linearly with traffic demand.

These inefficiencies become more pronounced in operational multi-operator deployments. ISP-DC interconnects commonly include heterogeneous 10G, and 100G interfaces with highly asymmetric traffic profiles across operators. The DC aggregation hub must therefore provision high-capacity ports to serve many low-rate ISP endpoints simultaneously. This creates substantial stranded aggregation capacity and poor port utilization. For example, a 400G aggregation port serving a

100G demand operates at only 25% utilization [6]. Backup links and LAG configurations further increase transceiver count, switching complexity, and power consumption.

P2P architectures also require intermediate electronic aggregation equipment to connect low-speed ISP endpoints to high-capacity aggregation ports. This introduces Optical-Electrical-Optical (OEO) conversions at multiple nodes, increasing both infrastructure cost and operational power consumption. In addition, each P2P interconnect reserves dedicated wavelength capacity regardless of actual utilization, leading to fragmented spectrum usage and inefficient resource allocation across the aggregation layer [5].





Fig. 1: Overview of (a) 400G single-wavelength P2P and (b) DSCM-enabled 16×25 Gbps coherent P2MP architecture [3].

DSCM enables a more flexible hub-and-spoke interconnect architecture. As shown in Fig. 1, a single 400G DSCM transceiver synthesizes 16 Nyquist subcarriers (SCs) within a 64 GHz optical channel, where each SC carries 25 Gbps using DP-16QAM modulation. Through passive optical splitting, the aggregation hub distributes SCs dynamically across multiple ISP endpoints, forming a coherent P2MP architecture [3]. Each ISP terminates only its assigned SCs using lower-rate coherent pluggables [7].


This subcarrier level granular allocation improves aggregation efficiency by allowing multiple low-rate ISP demands to share high-capacity coherent transceivers. Passive optical splitting also reduces the need for intermediate electronic aggregation devices and associated OEO conversions. Traffic growth can therefore be accommodated more efficiently through dynamic SC allocation within existing coherent resources, reducing the need for proportional expansion of transceivers and physical interconnects. However, the techno-economic performance of coherent P2MP under operational multi-operator ISP-to-DC interconnect conditions remains largely unexplored. This paper addresses that gap using real deployment records and long-term traffic projections.

III. ISP-DC INTERCONNECT ARCHITECTURES

This section describes the two ISP-DC interconnect architectures evaluated in this techno-economic study. Fig. 2 compares the legacy P2P architecture with the DSCM-driven P2MP architecture for multi-operator ISP-DC connectivity. In the legacy P2P architecture shown in Fig. 2a, each ISP connects to the DC aggregation layer through a dedicated optical link. Every ISP connection requires a separate transceiver pair



(a) Legacy Point-to-Point (P2P)



(b) DSCM-driven Point-to-Multipoint (P2MP)

Fig. 2: ISP-DC interconnect architectures.

and wavelength channel [2]. Additional capacity is provided by adding more transceiver pairs in parallel, either as independent links or as part of a LAG. Active-backup link pairs are used to provide resilience. As traffic demand increases, the number of transceivers, switching ports, and optical links also increases. Service growth therefore remains tightly coupled to physical hardware expansion.

The DSCM transceiver divides a 400G optical channel into 16 independently addressable 25 Gbps SCs using a single wavelength [9]. This allows bandwidth to be allocated at subcarrier granularity instead of using a single fixed-rate wavelength. Multiple ISP connections can therefore share the same coherent transceiver while using only the SC capacity required for their traffic demand.

In the P2MP architecture shown in Fig. 2b, a single 400G DSCM transceiver at the DC hub broadcasts all SCs downstream through a passive optical splitter. Each ISP site uses a lower-rate coherent pluggable that terminates only its assigned SCs [7]. A single hub transceiver can therefore serve multiple ISPs simultaneously over a shared optical channel. As traffic demand changes, SCs can be reassigned dynamically across ISP endpoints without requiring dedicated wavelength additions for every traffic increment. Higher aggregate demand is handled by activating additional hub transceivers when needed.

This shared aggregation model improves utilization of high-capacity coherent transceivers and reduces stranded bandwidth at the aggregation layer [6]. Passive optical splitting also reduces the need for intermediate electronic aggregation devices and associated Optical-Electrical-Optical (OEO) conversions [5]. Compared to P2P interconnects, P2MP can support traffic growth more efficiently while reducing transceiver count, interconnect density, and infrastructure overhead.

DSCM-driven P2MP architectures introduce additional operational and physical layer considerations. Each SC can experience different Quality of Transmission (QoT) conditions depending on frequency allocation, launch power, and propagation characteristics. SC power levels, modulation formats, and bandwidth allocations must therefore be managed inde-

pendently for each ISP connection. SDN enables centralized control of SC-level resources by automating provisioning, monitoring QoT, and dynamically reallocating bandwidth as traffic demand changes [4]. In addition, the passive optical splitter introduces insertion loss, reducing the optical power budget available at each far end leafs or spoke sites [7]. Higher split ratios further reduce received optical power and can limit transmission reach. Upstream wavelengths from different ISP leaves must also be managed carefully to avoid interference. QoT variation across SCs therefore requires continuous monitoring and careful link budget planning to maintain stable transmission performance across all ISP connections.

IV. TECHNO-ECONOMIC MODEL

The deployment data used in this study is sourced from an operator (FT) comprising 13 independent service providers connected to a centralized DC aggregation hub using legacy P2P optical interconnects. The dataset includes traffic demand, transceiver interface type, occupied aggregation-side transceivers, and port utilization information for each ISP client. In total, the deployment contains 56 physical P2P interconnects between ISPs and the DC aggregation layer.




Fig. 3: Interconnects deployment summary (Baseline Y1)

Fig 3 summarizes the baseline interconnect deployment characteristics for multi-tier service providers. The deployment exhibits several common characteristics. DC-side transceiver interfaces range from 10G to 100G, while traffic demand varies significantly across ISP clients. ISPs are classified into two tiers based on aggregate interconnect traffic, with five Tier 1 ISPs operating above 100 Gbps and eight Tier 2 ISPs operating below 100 Gbps. The baseline aggregate traffic across all ISP clients is approximately 1.06 Tbps. Port utilization also varies considerably across the deployment. Several low-capacity interconnects operate near 50–60% utilization, whereas multiple high-capacity exceed 90% utilization. This creates fragmented aggregation capacity and inefficient resource utilization at the DC hub, motivating more flexible subcarrier-granular resource allocation.

The 14% CAGR used in this study is supported by national internet exchange traffic records. Internet exchange traffic increased from 3.5 Tbps in 2020 to 6.8 Tbps in 2025, with annual growth rates consistently between 13% and 16%, averaging approximately 14% over this period. This is consistent with regional broadband growth projections reported in [10]–[12]. Applying this CAGR over 10 years projects the baseline aggregate interconnect traffic from 1,062 Gbps to approximately 3,453 Gbps, corresponding to a $3.25\times$ increase.

Traffic growth is modeled proportionally across ISP clients based on their baseline carried traffic. For each client i with baseline bandwidth B_i , the projected traffic at year t is given by:

$$B_i(t) = B_i \cdot (1 + r)^t \quad (1)$$

where $r = 0.14$ represents the annual traffic growth rate.

P2P interface count. For each ISP client, the number of required P2P transceiver pairs at year t is calculated as the ceiling of the projected client bandwidth divided by the interface rate. A factor of two is applied to account for primary and backup links. ISP-7 operates at 100% port utilization with no backup headroom at Year 1, reflecting a fully-loaded legacy deployment. The redundancy factor is applied uniformly in the model; operational deployments may defer backup provisioning until port capacity is upgraded:

$$N_i^{\text{P2P}}(t) = 2 \cdot \left\lceil \frac{B_i(t)}{C_{\text{port}}} \right\rceil \quad (2)$$

where C_{port} represents the interface capacity, equal to either 10 Gbps or 100 Gbps depending on the ISP client tier.

P2MP interface count. In the P2MP architecture, each ISP client is served using one or more 100G DSCM pluggables at the spoke end. Traffic is allocated proportionally using fixed 25 Gbps subcarriers without QoT-aware optimization. The number of required SCs for client i at year t is:

$$N_i^{\text{SC}}(t) = \left\lceil \frac{B_i(t)}{25} \right\rceil \quad (3)$$

Each 100G DSCM pluggable supports four SCs. The number of required 100G spoke pluggables per ISP client is therefore given by $\lceil N_i^{\text{SC}}/4 \rceil$. At the aggregation hub, each

TABLE I: Relative Cost and Power Factors

Parameter	Relative C.F. / P.F.
Existing Traffic at Agg.	≈ 1 Tbps
Traffic CAGR [1], [11]	14%
10G P2P (C.F.) [6]	x
100G P2P (C.F.) [6]	$1.43x$
400G P2P (C.F.) [6]	$3.43x$
100G P2MP (C.F.) [6]	$1.43x\alpha$, $\alpha \in \{1.5, 3.5\}$
400G P2MP (C.F.) [6]	$3.6x\alpha$, $\alpha \in \{1.5, 3.5\}$
10G P2P (P.F.) [13], [14]	y
100G P2P (P.F.) [13], [14]	$3.67y$
100G P2MP (P.F.) [13], [14]	$12y$
400G P2MP (P.F.) [13], [14]	$15y$

400G DSCM transceiver supports 16 SCs. The total number of required hub transceivers is determined by the aggregate SC demand across all ISPs.


Cost model. A techno-economic framework is used to compare infrastructure scaling trends independently of vendor-specific pricing [8]. Transceiver costs are relative to a 10G P2P baseline cost factor (C.F.) x . The relative C.F. used in the evaluation are summarized in Table I. P2MP transceivers have C.F. ($x\alpha$) where $\alpha \in \{1.5, 3.5\}$ to account for uncertainty in DSCM pluggable pricing [6]. Total CAPEX at year t is computed as the sum of transceiver counts multiplied by their respective normalized cost factors.

Power model. Transceiver power consumption is normalized relative to a 10G P2P baseline power factor (P.F.) of $y = 1.5$ W. The P.F. for each pluggable type are listed in Table I. Total aggregation-side transceiver power consumption at year t is calculated as the sum of active transceiver counts times their corresponding P.F.. Energy efficiency is expressed in W/Gbps and computed as the ratio of total transceiver power consumption to total carried traffic. The evaluation considers pluggable transceiver power only and excludes DC facility infrastructure, cooling systems, and switching equipment.


V. RESULTS AND DISCUSSION

This section compares P2P and P2MP ISP-DC interconnects over a 10-year traffic projection. The evaluation considers transceiver scalability, interconnect density, CAPEX, and energy efficiency, as summarized in Fig. 4.


Fig. 4a compares the transceiver scaling behavior of P2P and P2MP over the 10-year projection period. In the legacy P2P architecture, each new ISP demand and redundancy path requires additional dedicated transceivers. As a result, the total P2P transceiver count increases from 112 at Year 1 to 334 at Year 10, showing near-linear hardware growth with increasing traffic demand. In contrast, P2MP consolidates multiple ISP connections over shared coherent transceivers at the aggregation layer. Total P2MP transceiver count increases more gradually from 24 at Year 1 to 52 at Year 10. Overall transceiver savings improve from 78.6% at Year 1 to 84.4% at Year 10. At the aggregation layer specifically, coherent P2MP hub transceivers replace multiple dedicated P2P ports, maintaining 91–93% transceiver reduction throughout the projection period. At the ISP spoke side, savings increase from




(a) Transceiver scalability



(c) CAPEX variation



(b) Impact on interconnect density



(d) Power-driven OPEX comparison

Fig. 4: Techno-economic comparison of P2P and P2MP interconnect architectures over a 10-year traffic projection.

66.1% to 75.4% as multiple low-rate 10G interconnects are consolidated into higher-capacity DSCM pluggables. These results show that coherent P2MP significantly reduces hardware scaling requirements under sustained traffic growth.

Fig. 4b compares the required fiber links growth of P2P and P2MP architectures. In P2P, each ISP connection and redundancy path requires a dedicated physical link to the aggregation hub. Consequently, the total number of P2P interconnects increases from 56 at Year 1 to 167 at Year 10. P2MP reduces this growth by aggregating multiple ISP connections over shared coherent infrastructure through passive optical splitting. Total P2MP interconnect count increases from 19 to 41 over the same period, while the splitter-to-aggregation segment remains nearly constant throughout most of the projection horizon. Interconnect savings improve steadily from 57.1% at Year 1 to 68.9% by Year 10 as available subcarrier capacity becomes more efficiently utilized. This reduction in physical interconnect density lowers fiber usage, aggregation port count, and operational complexity at the DC hub.

Fig. 4c compares the normalized CAPEX trends of P2P and P2MP over the 10-year projection. In P2P, infrastructure C.F. increases almost linearly from 118.9 to 352.9 units, as each traffic increment requires additional dedicated hardware

resources. P2MP exhibits substantially slower C.F. growth due to shared aggregation and improved transceiver utilization. For a P2MP transceiver low C.F. of $\alpha = 1.5$, CAPEX reaches 127.3 units at Year 10, corresponding to 50.6% savings at Year 1 and 63.9% savings by Year 10 relative to P2P. Even under the higher C.F. of assumption of $\alpha = 3.5$, P2MP remains consistently lower cost than P2P throughout the projection period. The hub/aggregation capacity utilization gap between P2P and P2MP decreases from $\Delta 21.7\%$ at Year 1 to $\Delta 12.2\%$ at Year 10. Importantly, the remaining $\Delta 12.2\%$ at Year 10 represents additional free aggregation-side coherent capacity available in the P2MP architecture, demonstrating more efficient utilization of aggregation resources under increasing traffic demand. This residual capacity headroom can accommodate future traffic growth without proportional infrastructure expansion, thereby reducing long-term CAPEX in heterogeneous multi-operator interconnect deployments. These results show that coherent P2MP can substantially reduce long-term infrastructure expenditure in heterogeneous multi-operator interconnect deployments.

Fig. 4d compares normalized transceiver power consumption and energy efficiency for P2P and P2MP. P2MP transceiver power consumption remains higher throughout the

projection period because DSCM pluggables operate at higher line rates than conventional 10G conventional P2P. Total normalized P2MP power consumption increases from 0.45 to 1.0, while P2P increases from 0.2 to 0.65. However, energy efficiency measured in W/Gbps shows the opposite trend. P2MP energy efficiency improves from approximately 0.105 W/Gbps at Year 1 to 0.068 W/Gbps at Year 10 as SC utilization increases and coherent resources carry more aggregate traffic. In comparison, P2P energy efficiency improves more slowly from approximately 0.090 W/Gbps to 0.075 W/Gbps over the same period. The efficiency curves intersect around Year 5, after which P2MP becomes more energy efficient per transported bits. By Year 10, coherent P2MP achieves approximately 25–35% lower W/Gbps compared to P2P. This indicates that although coherent DSCM pluggables consume higher absolute power, improved aggregation efficiency enables better long-term energy utilization under sustained traffic growth.

VI. CONCLUSIONS

This study investigated the long-term infrastructure scalability and techno-economic gains of DSCM-driven coherent P2MP optical interconnects under projected traffic growth. The evaluation used real deployments comprising heterogeneous 10G and 100G interconnects with fragmented utilization and asymmetric traffic demand. Compared to conventional P2P architectures, coherent P2MP significantly reduced infrastructure scaling requirements over a 10-year traffic projection at 14% CAGR. Total transceiver count was reduced by up to 84%, while aggregation-side transceivers were reduced by up to 93% through coherent resource sharing at the DC hub. Fiber interconnect growth was also reduced by up to 69%, lowering aggregation complexity and fiber infrastructure requirements. In terms of techno-economic gains, coherent P2MP achieved 51–64% CAPEX reduction depending on the DSCM transceiver C.F.. Although coherent DSCM pluggables consume higher power, improved aggregation efficiency enabled 25–35% lower W/Gbps energy consumption under sustained traffic growth. These results demonstrate that fragmented utilization, mixed-rate interface coexistence, and asymmetric ISP traffic scaling make operational ISP-DC deployments well suited to SC level coherent aggregation architectures. Overall, coherent P2MP can substantially improve long-term infrastructure efficiency and reduce scaling overhead in heterogeneous multi-operator DC interconnects. Future work will investigate QoT-aware subcarrier allocation and bandwidth orchestration under multi-service SLA-driven traffic conditions.

ACKNOWLEDGMENTS

Funded by the European Union’s Horizon Europe research and innovation programme under Marie Skłodowska-Curie project NESTOR (G.A. No. 101119983) in collaboration with EWOC (G.A. No. 101073265) and ALLEGRO (G.A. No. 101092766) in collaboration with NOKIA Optical Networks, (Germany & Portugal) & Future Telecom (FT).

REFERENCES

- [1] Nokia, “EMEA Broadband Index Report 2025,” Nokia White Paper.
- [2] D. Welch, A. Napoli, J. Bäck, S. Buggaveeti, C. Castro, A. Chase, X. Chen, V. Dominic, T. Duthel, T. A. Eriksson *et al.*, “Digital subcarrier multiplexing: Enabling software-configurable optical networks,” *Journal of Lightwave Technology*, vol. 41, no. 4, pp. 1175–1191, 2023.
- [3] 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,” *Journal of Lightwave Technology*, vol. 39, no. 16, pp. 5232–5247, 2021.
- [4] A. S. Thyagaturu, A. Mercian, M. P. McGarry, M. Reisslein, and W. Kellerer, “Software defined optical networks (SDONs): A comprehensive survey,” *IEEE Communications Surveys & Tutorials*, vol. 18, no. 4, pp. 2738–2786, 2016.
- [5] J. Bäck, P. Wright, J. Ambrose, A. Chase, M. Jary, F. Masoud, N. Sugden, G. Wardrop, A. Napoli, J. Pedro *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.
- [6] N. Skorin-Kapov, F. J. M. Muro, M.-V. B. Delgado, and 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.
- [7] S. Zeb, A. Ali, I. C. Dipto, A. Rosso, M. U. Masood, R. Schips, R. Ambrosone, S. Straullu, F. Aquilino, A. Nespola *et al.*, “AI-driven converged metro-access optical network-as-a-service with point-to-multipoint coherent optics for 6G X-Hauling,” *Computer Networks*, p. 112338, 2026.
- [8] M. U. Masood, S. Zeb, I. Khan, and V. Curri, “Evaluating Capacity Scaling Strategies for Cost-Efficient Optical Network Evolution,” in *Optical Fiber Communication Conference*. Optica Publishing Group, 2026, pp. W4A–1.
- [9] S. Zeb, I. C. Dipto, M. U. Masood, A. Ali, J. Pedro, A. Napoli, and V. Curri, “ML-driven QoT estimation for adaptive P2MP optical network-as-a-service in 6G X-Haul,” in *2026 IEEE International Conference on Machine Learning for Communication and Networking (ICMLCN)*, 2026, pp. 1–6.
- [10] E. C. Kirci, A. Mishra, and L. Vanbever, “Five Blind Men and the Internet: Towards an Understanding of Internet Traffic,” *arXiv preprint arXiv:2509.06515*, 2025.
- [11] Pakistan Telecommunication Authority, “Annual Report (2024–25),” Islamabad, Pakistan, Tech. Rep., 2025, <https://www.pta.gov.pk/assets/media/Annual%20Report%202025-25%20Nov%2025.pdf>.
- [12] A. Ali, S. Zeb, M. U. Masood, A. Rosso, G. Malik, R. Ambrosone, R. Schips, M. Pollone, S. Straullu, F. Aquilino, J. Pedro, A. Napoli, A. Galardini, and V. Curri, “Sustainable 6G optical any-haul with coherent pluggables in converged metro-access,” in *2026 IEEE 11th International Conference on Network Softwarization (NetSoft)*, 2026, pp. 1–6.
- [13] 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,” *Journal of Optical Communications and Networking*, vol. 17, no. 6, pp. 526–542, 2025.
- [14] J. A. Hernández, M. Quagliotti, L. Serra, L. Luque, R. L. da Silva, A. Rafel, Ó. G. de Dios, V. López, A. Eira, R. Casellas *et al.*, “Comprehensive model for technoeconomic studies of next-generation central offices for metro networks,” *Journal of Optical Communications and Networking*, vol. 12, no. 12, pp. 414–427, 2020.