

Statistical Analysis of end-to-end Route Feasibility in Converged Metro–Access Optical Networks

Original

Statistical Analysis of end-to-end Route Feasibility in Converged Metro–Access Optical Networks / Ali, A., Masood, M.U., Rosso, A., Malik, G., Pollone, M., Galardini, A., Curri, V.. - In: IEEE PHOTONICS TECHNOLOGY LETTERS. - ISSN 1041-1135. - (2026). [10.1109/LPT.2026.3653136]

Availability:

This version is available at: 11583/3006690 since: 2026-01-19T11:01:53Z

Publisher:

IEEE

Published

DOI:10.1109/LPT.2026.3653136

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)

Statistical Analysis of end-to-end Route Feasibility in Converged Metro–Access Optical Networks

Ahtisham Ali , *member, IEEE*, Muhammad Umar Masood , *member, IEEE*, Andrea Rosso , *member, IEEE*, Gulmina Malik , *member, IEEE*, Michela Pollone, Alessandro Galardini, Vittorio Curri , *Fellow, IEEE*,

Abstract—Disaggregation of the Radio Access Network (RAN) imposes stringent fronthaul latency and physical-layer performance requirements on supporting the optical transport infrastructure. This paper evaluates the feasibility of Access–DU–Access (A–DU–A) routes in a converged metro–access optical network, jointly considering Bit Error Rate (BER) and the fronthaul latency constraint of 250 μ s, considering functional split 7.2 defined by 3GPP. Two deployment scenarios are examined: (i) DU/CU hosting at metropolitan Points of Presence (PoPs) and (ii) DU/CU hosting at external data centers (DCs). Using Cassini and Phoenix whiteboxes hosting pluggable DCO transceivers (TRxs), the analysis quantifies the impact of modulation format, BER thresholds, and end-to-end propagation latency on service feasibility. Results show that metro-based hosting supports 87.9% latency-compliant routes compared to 47.6% under data-center hosting, and that Cassini consistently provides higher BER feasibility across modulation formats. The findings demonstrate that physical-layer impairments and functional placement must be jointly optimized to support latency-constrained, virtualized RAN transport in converged metro–access networks.

Index Terms—metro-access optical networks, RAN, QoS, BER analysis

I. INTRODUCTION

The evolution from 5G towards 6G is redefining the architectural and performance requirements of transport networks, demanding ultra-low latency, high capacity, and flexible resource allocation. The disaggregation of the RAN into Distributed Units (DUs), Central Units (CUs), and Radio Units (RUs) enhances architectural flexibility and enables dynamic placement of processing entities. However, it also introduces stringent constraints on the underlying optical transport layer, which must simultaneously ensure high signal quality, meet strict end-to-end latency budgets, and support heterogeneous traffic profiles [1, 2].

In this context, optical fiber remains the most suitable, reliable, and scalable medium for supporting latency critical and high capacity interconnections [3]. The converged metro access optical network offers an opportunity to exploit existing infrastructure for RAN transport, reducing capital expenditures (CAPEX) and enabling resource sharing between mobile and fixed services [4, 5]. Such convergence is increasingly recognized as a key enabler for flexible, virtualized, and service-oriented RAN architectures in 5G and beyond [6–9].

Ahtisham Ali (ahtisham.ali@polito.it), Muhammad Umar Masood, Andrea Rosso, Gulmina Malik, and Vittorio Curri are with Politecnico di Torino, Italy. Ahtisham Ali, Michela Pollone, and Alessandro Galardini are with Consorzio TOP-IX, Italy. This work has received funding from the EU Horizon 2020 MSCA project EWOC (101073265), ALLEGRO (101092766), the project PNRR-NGEU (MUR–DM117/2023), and the NG EU program “RESTART” (PE00000001).

Manuscript received December xx, 2025; revised December xx, 2025.

Within such converged network environments, optical transport networks are expected to meet the latency and quality-of-transmission requirements defined by 3GPP.

According to functional split (FS) 7.2 [10], the one-way midhaul latency between the CU and DU may reach up to 10 ms, whereas the fronthaul between the DU and RU is limited to 250 μ s. These constraints directly influence the feasible geographical distribution of RAN functions and restrict the set of physical routes that can satisfy both latency and BER limits. Beyond latency compliance, the route feasibility is ultimately bounded by physical-layer impairments, including amplified spontaneous emission (ASE) noise, nonlinear interference (NLI), and implementation penalties. Thus, evaluating the joint BER–latency feasibility of optical paths is essential to ensure reliable and flexible service delivery in converged optical–wireless networks.

Building on our previous work [11], which provided a statistical evaluation of only BER-driven feasibility across access-to-access routes using a Cassini whitebox hosting pluggable DCO transceiver (TRx), this study extends the analysis by introducing three key advancements: (i) modeling a realistic metro–access network topology representative of operator deployments and extending the analysis to A–DU–A connectivity under DU/CU placement scenarios to enable the optical networking as-a-Service (ONaaS), (ii) incorporation of propagation-based latency evaluation, allowing a joint assessment of physical-layer signal quality and latency compliance under 3GPP FS 7.2, (iii) comparative analysis of two whiteboxes, Cassini and Phoenix, hosting pluggable DCO TRxs, to quantify the impact of BER margins for different modulation formats under realistic metro-scale path diversity. These contributions offer a more comprehensive physical-layer perspective on feasibility for virtualized RAN transport in converged metro–access infrastructures and clarify the role of functional placement in determining service compliance. The A–DU–A configuration reflects a typical operational architecture in fronthaul and aggregation segments, where traffic from distributed access nodes is routed through intermediate DUs to facilitate coordination and edge-cloud processing tasks. In this work, two representative DU/CU placement scenarios are examined: (i) metro-hosted functions at selected PoPs, and (ii) centralized hosting at external data centers. These scenarios, described in detail in Section II, enable evaluation of end-to-end A–DU–A feasibility under 3GPP BER–latency constraints.

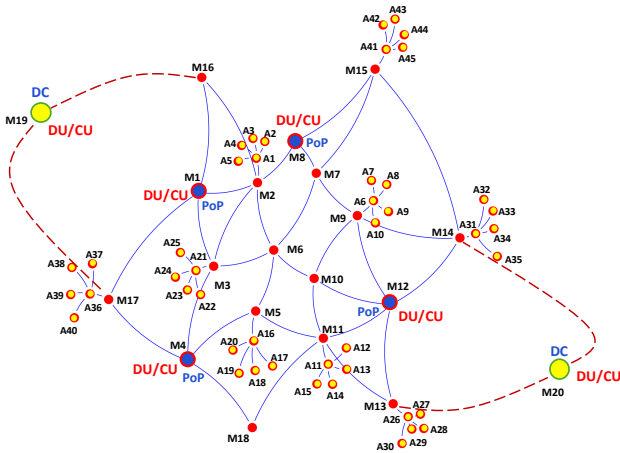


Fig. 1: Converged Metro Access Network Topology

II. SIMULATION SCENARIO

The performance evaluation is carried out on a fully converged metro access optical transport network, shown in Fig. 1. The overall network comprises 18 metro nodes (M1–M18) interconnected through bidirectional optical line systems with average inter-node distances of 4–5 km and a maximum span length of 12 km. The access domain includes 45 access nodes (A1–A45), each connected to its serving metro node via dedicated short-reach fiber spans of 1–3 km, emulating the optical connections between RUs and the metro PoPs. All fiber spans are modeled as Standard Single-Mode Fiber (SSMF) with an attenuation coefficient of 0.25 dB/km. Erbium-Doped Fiber Amplifiers (EDFAs) are deployed where required, assuming a nominal noise figure of 6 dB. The physical-layer model accounts for fiber attenuation, ASE noise accumulation, and implementation penalties associated with the coherent TRxs, consistent with metro-level optical systems. This topology is chosen as a representative metro–access configuration commonly adopted by regional operators, with span lengths and access loops matching European deployments. It provides a justified and realistic basis for assessing physical-layer BER–latency feasibility, independent of traffic profiles.

Two representative DU/CU placement scenarios are evaluated in this study. Scenario I models an edge-cloud configuration in which DU/CU functions are hosted within the metropolitan domain at selected PoPs. This reflects operator-grade metro–access deployments designed to minimize fronthaul propagation distance and support tight 3GPP latency budgets. Scenario II represents a centralized cloud-RAN configuration in which DU/CU functions are relocated to large external data centers interconnected with the metro backbone through additional long-haul spans. This placement increases the end-to-end transport distance and introduces wider propagation-delay variations. Evaluating both scenarios over the same physical topology enables direct comparison of how functional placement influences the feasibility of A–DU–A routes under joint BER–latency constraints.

Scenario I represents the case where DU/CU functionalities are hosted within the metropolitan zone. In this configuration, four PoPs located at nodes M1, M4, M8, and M12 serve as DU/CU hosting sites. This configuration reflects an edge-

cloud deployment model in which processing functions are positioned close to the access domain, thereby minimizing fronthaul propagation delay and supporting stringent 3GPP latency budgets.

Scenario II corresponds to the case where DU/CU functionalities are hosted at centralized large data centers situated outside the metropolitan region. These data centers, associated with nodes M19 and M20, are connected to the metro network through additional longer links of 25 km (M16–M19, M13–M20) and 30 km (M14–M20, M17–M19), respectively. With these extensions, the effective metro network expands to 20 metro nodes. This configuration represents a centralized cloud-RAN architecture with significantly longer transport distances and larger propagation-delay variations.

The physical-layer performance of each route is evaluated using a Python-based optical network simulation framework, developed following the modeling principles proposed in [11, 12]. In particular, the model incorporates both ASE noise generated by optical amplifiers and Non-Linear Interference (NLI) accumulated along the fiber spans. The Generalized Signal-to-Noise Ratio (GSNR) of the i^{th} channel is given as

$$\text{GSNR}_i = \frac{P_{S,i}}{P_{\text{ASE}}(f_i) + P_{\text{NLI},i}(f_i)} \quad (1)$$

where $P_{S,i}$ is the signal launch power, $P_{\text{ASE}}(f_i)$ is the accumulated ASE noise, and $P_{\text{NLI},i}(f_i)$ represents the NLI contribution. For every route, the effective Signal-to-Noise Ratio (SNR) and the associated BER for each route are obtained through

$$\text{SNR}^{-1} = \text{GSNR}^{-1} + \text{SNR}_{\text{TRX}}^{-1} \quad (2)$$

$$\text{BER} = k_1 \cdot \text{erfc}\left(\sqrt{k_2 \cdot \text{SNR}}\right) \quad (3)$$

where k_1 and k_2 are modulation-dependent constants defined in [13]. Two whiteboxes, Cassini and Phoenix, hosting coherent pluggable DCO TRxs [14, 15] are evaluated in this investigation. As discussed in [13, 16, 17], TRx contribution is modeled through the calibrated term SNR_{TRX} in the end-to-end SNR formulation shown in Eq. 2 based on experimental characterizations and modeling methodology.

Every A–DU–A route is assessed under a pre-FEC BER threshold of 10^{-2} and a fronthaul latency constraint of 250 μs (FS 7.2). Latency is computed by mapping the physical length of each RU–DU and DU–RU segment to the propagation delay $\approx 5 \mu\text{s}/\text{km}$ in case of SSMF [12]. This provides the deterministic propagation component of end-to-end delay within a transparent optical path, which is the focus of the physical-layer feasibility analysis.

To limit computational complexity, a maximum hop count of 10 was imposed on each RU–DU path evaluated. A route is classified as feasible only if both BER and latency constraints are satisfied. This modeling framework enables a direct comparison of Scenario I and Scenario II, quantifying the combined impact of DU/CU placement, physical-layer impairments, and TRx characteristics on service feasibility in converged metro access optical networks.




Fig. 2: Joint BER–latency distribution for Cassini and Phoenix whiteboxes under both deployment scenarios.

III. RESULTS AND DISCUSSION

This section presents a comprehensive assessment of A–DU–A route feasibility under the two DU/CU placement scenarios. The results are organized into three components: (i) the joint BER–latency behavior, (ii) statistical distribution of fronthaul latency, and (iii) the impact of BER thresholds on service feasibility across both TRxs and modulation formats.

Figure 2 illustrates the joint distribution of the BER and end-to-end latency for all evaluated end-to-end routes using Cassini and Phoenix whiteboxes under both scenarios. The red dashed lines indicate the fronthaul latency bound of $250\ \mu\text{s}$, while the shaded regions denote the feasible operating zones where both BER and latency constraints are satisfied. The elliptical contours indicate the modulation regimes corresponding to DP-QPSK and DP-16QAM modulation formats.

Under Scenario I, where DU/CU functionalities are hosted within the metropolitan domain, both TRxs exhibit compact BER clusters with limited dispersion. The lower-BER region from 10^{-9} to 10^{-7} corresponds to DP-QPSK, confirming its higher robustness and larger optical Signal-to-Noise Ratio (OSNR) margin. The higher-BER region from 10^{-4} to 10^{-2} corresponds to DP-16QAM, reflecting its expected sensitivity to accumulated impairments. Nevertheless, the reduced propagation distances in the metro domain enable DP-16QAM to remain within the BER threshold for a meaningful subset of routes, indicating that higher-order modulation is still viable when DU/CU functions are positioned close to the access segment. Phoenix exhibits a similar behaviour, albeit with broader dispersion due to larger implementation penalties.

Under Scenario II, where DU/CU functionalities are centralized at external data centers, the latency distribution shifts significantly toward higher values. Although many routes

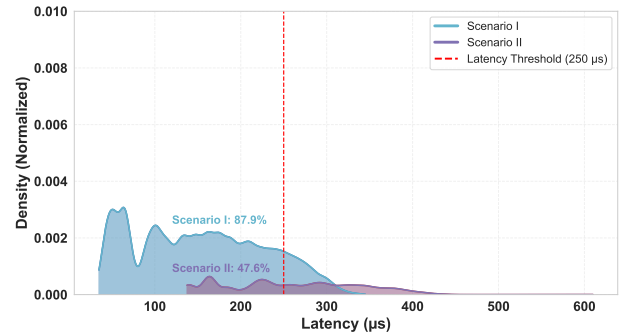


Fig. 3: Normalized latency distribution for A–DU–A routes under both deployment scenarios. The red dashed line marks the fronthaul latency threshold.

continue to achieve BER levels in the 10^{-4} to 10^{-2} range for DP-16QAM, but a substantial share of them exceed the $250\ \mu\text{s}$ latency limit due to longer propagation distances. Cassini retains a wider feasible region than Phoenix due to better OSNR margin; however, the feasibility reduction is primarily driven by propagation delay rather than signal quality. This demonstrates that Scenario II becomes latency-limited rather than BER-limited, even when OSNR performance is adequate. Overall, the joint BER–latency analysis confirms that Scenario I supports substantially more feasible routes overall, with DP-QPSK showing the broad feasibility region due to its higher OSNR tolerance, while Scenario II becomes predominantly latency-limited rather than BER-limited.

Figure 3 compares the fronthaul latency distribution in both scenarios, normalized with respect to the scenario containing the largest number of evaluated routes. In Scenario I, the distribution peaks well below the threshold, with 87.9% of routes satisfying the fronthaul constraint. This behavior is




Fig. 4: Percentage of feasible A–DU–A routes for Cassini and Phoenix whiteboxes under multiple BER thresholds.

attributed to the localized placement of DU/CU functions within the metro domain, resulting in short propagation distances and minimal latency variation. In contrast, Scenario II exhibits a markedly different distribution. The addition of 25–30 km spans to reach the external data centers increases the overall path lengths, causing the latency distribution to shift toward significantly higher values. As a result, only 47.6% of routes satisfy the 250 μ s requirement, and the spread of latency values becomes substantially wider due to increased topological diversity.

These results indicate that centralizing DU/CU functions outside the metro region yields an approximate 40% reduction in latency-compliant routes relative to metro-hosted deployments. Importantly, the underlying optical transmission remains physically feasible in most cases; functional placement becomes the dominant factor determining end-to-end latency compliance in converged metro–access transport networks.

Figure 4 summarizes the percentage of feasible end-to-end routes under three BER thresholds of 10^{-2} , 5×10^{-3} , and 10^{-3} . Shaded bars correspond to Cassini results, while the unshaded bars correspond to Phoenix. Each horizontal bar relates the feasibility metrics from Scenario II and Scenario I, demonstrating how variations in DU/CU placement translate into differences in route feasibility. At the relaxed threshold of 10^{-2} , both TRxs achieve high feasibility ($\geq 70\%$) across all modulation formats and deployment scenarios. When the threshold is tightened to 5×10^{-3} , feasibility decreases primarily for Phoenix DP-16QAM ($\approx 51\%$), demonstrating greater BER sensitivity in high-order modulation. Cassini maintains higher feasibility across all thresholds due to superior OSNR tolerance. At the strictest threshold of 10^{-3} , only DP-QPSK remains feasible, while DP-16QAM yields negligible or zero feasibility, in line with its higher susceptibility to accumulated noise and nonlinear impairments.

Across all thresholds, Scenario I consistently outperforms Scenario II, confirming that shorter propagation distances preserve both OSNR and latency, thereby sustaining BER-compliant operation. The gap between Cassini and Phoenix

further indicates the advantage of whiteboxes with stronger noise margins for metro-scale deployments.

IV. CONCLUSION

This work presented a detailed feasibility analysis of A–DU–A routes in a converged metro–access optical network under two deployment architectures for distributed and centralized RAN functions. By jointly evaluating BER and fronthaul latency for Cassini and Phoenix whiteboxes hosting pluggable DCO TRxs across different modulation formats, the study quantified the impact of physical-layer impairments, TRx characteristics, and functional placement on end-to-end service feasibility.

The results demonstrate that Scenario I, where DU/CU functions are hosted within the metropolitan domain, enables the majority of routes to satisfy both BER and the 250 μ s fronthaul latency requirement. In Scenario II, where DU/CU functions are centralized at external DCs, the extended propagation distances reduce the share of latency-compliant routes by approximately 40%, while BER performance remains largely unchanged. This confirms that latency, not optical signal degradation, is the dominant limiting factor in centralized cloud-RAN deployments.

Cassini consistently outperforms Phoenix due to its higher OSNR margin, particularly under strict BER thresholds, highlighting the importance of TRx selection for metro-scale converged infrastructures. The comparative analysis demonstrates that DU/CU placement is a primary determinant of transport feasibility and demonstrates the advantages of metro-hosted architectures for latency-sensitive RAN services within 3GPP requirements. Future work will extend this analysis by integrating dynamic orchestration and predictive control to enable real-time physical-layer feasibility management and adaptive resource allocation in converged optical–wireless networks.

REFERENCES

- [1] F. Musumeci *et al.*, “Optimal bbu placement for 5g c-ran deployment over wdm aggregation networks,” *JLT*, vol. 34, no. 8, pp. 1963–1970, 2016.
- [2] J.-i. Kani *et al.*, “Disaggregation and virtualization for future access and metro networks [invited tutorial],” *JOCN*, vol. 17, no. 1, pp. A1–A12, 2024.
- [3] A. Ali *et al.*, “Capacity assessment in converged metro-access optical networks for end-to-end ran fronthaul,” in *2025 25th ICTON*. IEEE, 2025, pp. 1–4.
- [4] F. Cavaliere *et al.*, “Will a metro-access optical continuum ever fly? access network evolution trends and enabling technologies,” *JOCN*, vol. 17, pp. A134–A141, 2025.
- [5] M. Klinkowski, “Latency-aware DU/CU placement in convergent packet-based 5G fronthaul transport networks,” *Applied Sciences*, vol. 10, no. 21, p. 7429, 2020.
- [6] F. B. Saghezchi *et al.*, “Drive towards 6g,” in *Enabling 6G Mobile Networks*. Springer, 2021, pp. 3–35.
- [7] J. Costa-Requena *et al.*, “Sdn-enabled thz wireless x-haul for b5g,” in *2021 Joint EuCNC/6G Summit*. IEEE, 2021, pp. 211–216.
- [8] C. Ranaweera *et al.*, “Design and deployment of optical x-haul for 5g, 6g, and beyond: progress and challenges,” *JOCN*, vol. 15, no. 9, pp. D56–D66, 2023.
- [9] Z. Vujicic *et al.*, “Towards virtualized optical-wireless heterogeneous networks,” *IEEE access*, 2024.
- [10] 3GPP, “Study on CU-DU lower layer split for NR; (release 15),” 3GPP, Tech. Rep. TR 38.816 V15.0.0 (2017-12), 2017.
- [11] A. Ali *et al.*, “Statistical evaluation of lightpath feasibility in converged optical metro-access networks,” in *2025 IEEE Photonics Society SUM*, 2025, pp. 1–2.
- [12] —, “Optimizing ran x-haul performance through targeted hollow-core fiber deployment in converged metro-access networks,” in *SoftCOM*, 2025, pp. 1–6.
- [13] A. Rosso, “Modeling and controlling optical transponder white-boxes based on the physical layer digital twin,” Master’s thesis, Polito, 2024.
- [14] H. Nishizawa *et al.*, “Open whitebox architecture for smart integration of optical networking and data center technology,” *JOCN*, vol. 13, no. 1, pp. A78–A87, 2020.
- [15] Wistron Corporation, “Galileo flex t: Tip oopt phoenix 400g disaggregated muxponder,” 2023. [Online]. Available: <https://telecominfraproject.com/exchange/offering/?id=635018c718709a0019224033>
- [16] 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.
- [17] —, “Modeling transceiver ber-onsr characteristic for qot estimation in short-reach systems,” in *2023 ONDM*, 2023, pp. 1–3.