

Optimizing RAN X-Haul Performance through Targeted Hollow-Core Fiber Deployment in Converged Metro-Access Networks

Original

Optimizing RAN X-Haul Performance through Targeted Hollow-Core Fiber Deployment in Converged Metro-Access Networks / Ali, Ahtisham; Zeb, Sanwal; Rosso, Andrea; Masood, Muhammad Umar; Malik, Gulmina; Ambrosone, Renato; Schips, Riccardo; Correia, Bruno; Straullu, Stefano; Aquilino, Francesco; Pedro, João; Napoli, Antonio; Galardini, Alessandro; Curri, Vittorio. - (2025), pp. 1-6. (2025 International Conference on Software, Telecommunications and Computer Networks (SoftCOM) Split (Cro) 18-20 September 2025)
[10.23919/SoftCOM66362.2025.11197410].

Availability:

This version is available at: 11583/3006011 since: 2025-12-18T22:00:21Z

Publisher:

IEEE

Published

DOI:10.23919/SoftCOM66362.2025.11197410

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

©2025 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)

Optimizing RAN X-Haul Performance through Targeted Hollow-Core Fiber Deployment in Converged Metro-Access Networks

Ahtisham Ali
Consorzio Top-IX
DET, Politecnico di Torino
Torino, Italy
ahtisham.ali@top-ix.org

Sanwal Zeb
DET, Politecnico di Torino
Torino, Italy
sanwal.zeb@polito.it

Andrea Rosso
DET, Politecnico di Torino
Torino, Italy
andrea.rosso@polito.it

Muhammad Umar Masood
DET, Politecnico di Torino
Torino, Italy
muhammad.masood@polito.it

Gulmina Malik
DET, Politecnico di Torino, Italy
gulmina.malik@polito.it

Renato Ambrosone
DET, Politecnico di Torino, Italy
renato.ambrosone@polito.it

Riccardo Schips
DET, Politecnico di Torino, Italy
riccardo.schips@polito.it

Bruno Correia
Nokia, Carnaxide, Portugal
bruno.correia@nokia.com

Stefano Straullu
Links Foundation, Torino, Italy
stefano.straullu@linksfoundation.com

Francesco Aquilino
Links Foundation, Torino, Italy
francesco.aquilino@linksfoundation.com

João Pedro
Nokia, Carnaxide, Portugal
joao.pedro@nokia.com

Antonio Napoli
Nokia, Munich, Germany
antonio.napoli@nokia.com

Alessandro Galardini
Consorzio Top-IX, Torino, Italy
alessandro.galardini@top-ix.org

Vittorio Curri
DET, Politecnico di Torino, Italy
vittorio.curri@polito.it

Abstract—The growing trend of disaggregated and cloud-native RAN architectures in 5G deployments and future 6G networks imposes stringent latency and capacity requirements on optical transport networks. This paper evaluates the feasibility of using standard single-mode fiber (SSMF) and hollow-core fiber (HCF) for supporting x-haul transport across converged metro-access networks. Using a converged metro-access network topology and experimentally modeled performance of a NOKIA ICE-X multi-carrier transceiver, we analyze end-to-end RAN connection from Radio Units (RUs) to Central Units (CUs) via Distributed Units (DUs). Results show that HCF significantly outperforms SSMF in satisfying BER and latency constraints over long distances, particularly in midhaul-dominated scenarios. Case studies with varying DU–CU link lengths further demonstrate HCF’s potential to enhance service coverage for future latency-critical and high-capacity deployments. These findings position HCF as a strong candidate for enabling scalable and constraint-compliant optical transport in next-generation RAN infrastructures.

Index Terms—5G, 6G, X-haul, optical transport networks, hollow-core fiber (HCF), radio access networks, metro-access networks, network-as-a-service

I. INTRODUCTION

The global deployment of 5G and the evolution towards 6G are revolutionizing multiple sectors by enabling sustainable, resilient, and immersive services. To support such innovation, 5G networks embrace disaggregation of radio access networks (RAN), allowing flexible placement of network functions and efficient resource utilization. In particular, the evolution

towards Open Radio Access Networks (O-RAN) [1] introduces standardized, vendor-neutral interfaces, enabling multi-vendor ecosystems and dynamic RAN deployments. These trends are accompanied by stringent requirements for the underlying transport networks in terms of latency, capacity, and service granularity, pushing optical infrastructure to the forefront of 5G enablement [2].

Within this evolving landscape, data centers (DCs) are critical for hosting centralized and virtualized network functions, including components of the RAN and core network. Efficient interconnection between Radio Units (RUs), Distributed Units (DUs), and Central Units (CUs) across the fronthaul (FH), midhaul (MH), and backhaul segments is essential to meet stringent 5G performance requirements. The fronthaul segment, linking RUs to DUs, typically carries digitized IQ samples and requires ultra-low latency ($< 100\mu\text{s}$) and high bandwidth (10–100 Gbps per sector) [2]. Midhaul, connecting DUs to CUs, supports control and coordination functions and has less stringent latency requirements (1 ms to 10 ms), but still requires high capacity to manage aggregated traffic from multiple RUs [3]. To address these challenges, technologies such as enhanced Common Public Radio Interface (eCPRI), Time-Sensitive Networking (TSN), and high-speed optical coherent links have been proposed and evaluated as potential enablers [2]. Digital Subcarrier Multiplexing (DSCM) has emerged as a promising transmission technique, offering higher spectral efficiency and hardware simplicity by generat-

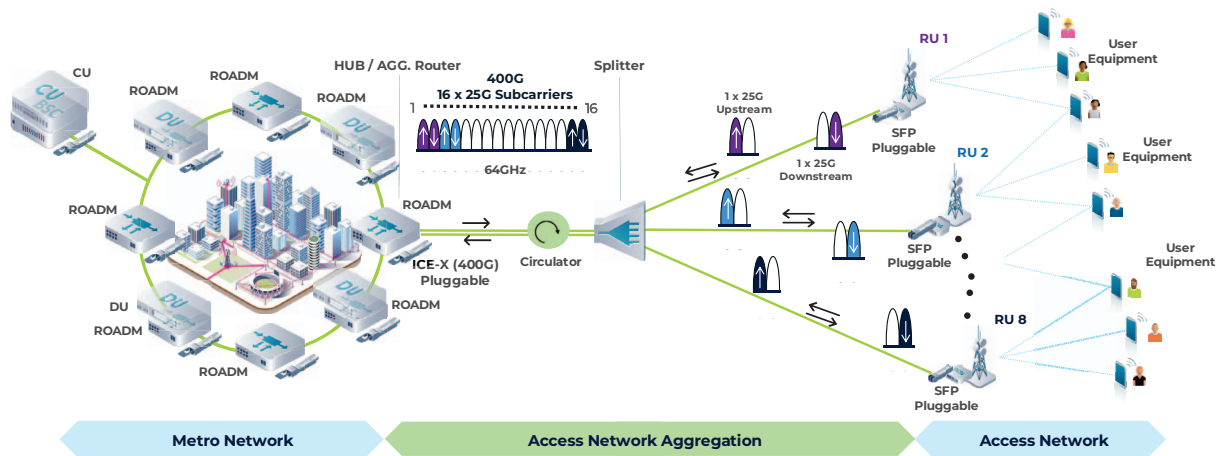


Fig. 1: P2P and P2MP Converged Metro Access Network Architecture using DSCM Pluggable TRx

ing multiple non-orthogonal subcarriers from a single laser source [4]. These subcarriers, processed independently via advanced Digital Signal Processing (DSP), allow dynamic routing and support rates up to 800 Gb/s per wavelength in metro and core deployments [5]. Equipment such as Nokia’s multicarrier ICE-X transceivers, supporting 400 Gbps with 16 subcarriers using DP-16QAM, can efficiently connect distributed units (DUs) to centralized units (CUs) over a shared optical infrastructure, enabling multi-RU aggregation. In addition to regular growth in consumer and enterprise traffic as forecasted by recent internet traffic reports [6], emerging AI workloads [7] and upcoming 6G deployments [8] are expected to intensify demands on transport network capacity. While these advances in modulation and transceiver design address capacity needs, reducing end-to-end latency remains a critical bottleneck. Standard Single-Mode Fiber (SSMF), while mature and widely deployed, faces intrinsic limitations in latency, nonlinear performance, and achievable transmission bandwidth, especially over long distances. In response, Hollow Core Fiber (HCF) has emerged as a promising alternative, offering distinct physical advantages by guiding light through an air-filled core instead of solid silica [9]. These properties can potentially enable lower propagation delay, reduced nonlinear effects, and higher usable bandwidth [10–12]. However, the integration of HCF into real-world networks remains an area of active research and evaluation, particularly regarding its manufacturability, interconnection challenges, and long-term performance stability. This air-guided transmission significantly reduces the refractive index, enabling lower latency and reduced nonlinear impairments compared to SSMF. State-of-the-art designs such as nested antiresonant nodeless fibers (NANF) have demonstrated 30% lower latency and losses as low as 0.1 dB/km [13], making them attractive candidates for metro and long-haul deployment [14]. Leading industry players, notably Microsoft, have accelerated this momentum by acquiring HCF startups and announcing large-scale HCF deployments—over 15,000 km by 2026 [15].

However, several technical challenges remain before HCF

can be adopted at scale. These include inter-modal interference (IMI), especially over long distances [9], and elevated insertion and splice losses between HCF and conventional fibers [16, 17]. Although recent studies have demonstrated reduced IMI and improved interconnection performance [18], uncertainties regarding manufacturing scalability and mechanical tolerances (e.g., bending resilience) still pose deployment risks [19, 20]. Despite these limitations, the benefits of HCF are compelling, especially for latency-sensitive services in metro-access optical networks. Prior studies show that HCF can significantly reduce the number of in-line amplifiers required and improve optical reach, simplifying network architectures [21]. The potential for latency reduction is particularly relevant for converged infrastructures offering metro-access network resources as-a-service to support RAN X-Haul [9].

In this work, we investigate the feasibility of using HCF as an alternative to SSMF in a converged metro-access optical transport network to support RAN FH and MH segments. Unlike prior studies such as [11, 21], which primarily assessed HCF benefits in terms of point-to-point latency reduction or amplifier savings, our paper presents three key differentiators: (i) modeling a realistic urban optical topology to enable the usage of Optical Network as-a-Service (ONaaS), (ii) employing experimentally characterized performance of NOKIA’s ICE-X multi-carrier transceiver to emulate metro-scale deployments supporting 400G transport, and (iii) performing a comparison of using HCF and SSMF, under Bit Error Rate (BER) and latency constraints as discussed in [22], for dedicated links between Metro (DUs) to data centers (CUs). Given the evolving nature of HCF technology, its standardized attenuation characteristics are yet to be conclusively established, but for simplicity, we use 0.11 dB/km as reported in [21]. This study provides new insights into the performance gains and challenges of targeted HCF deployment within the emerging paradigm of optical networks as-a-service.

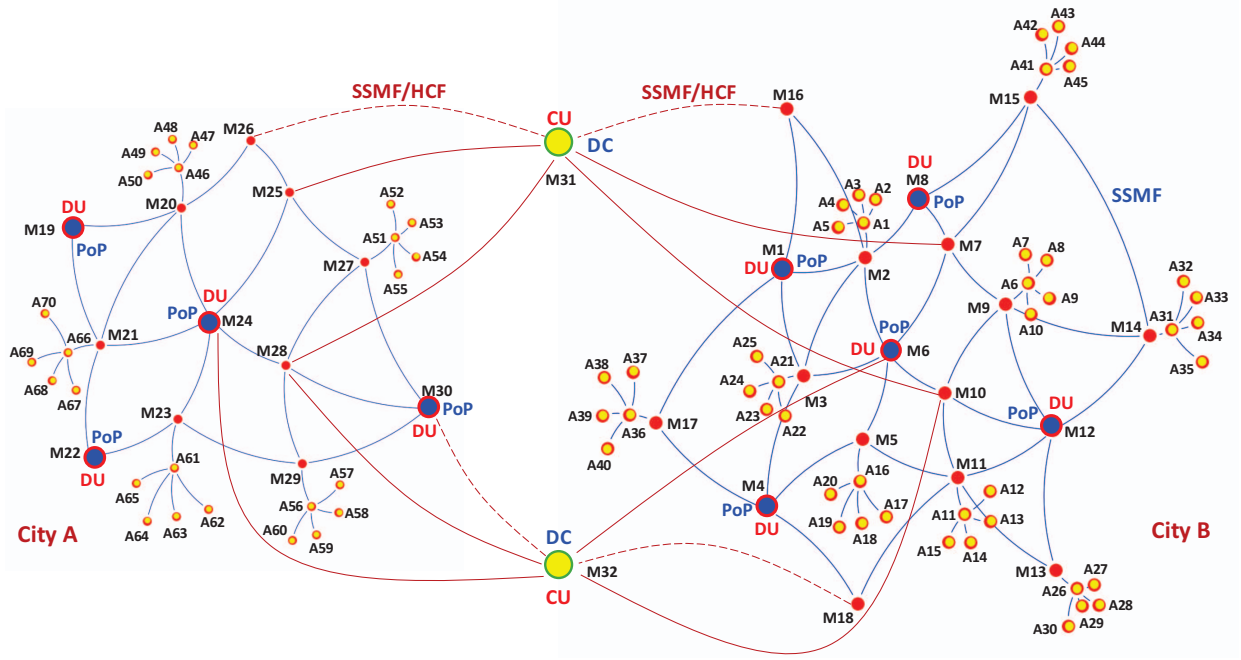


Fig. 2: Converged Metro Access Network Topology

II. SIMULATION SCENARIO

The simulation framework is based on an initial experimental characterization of the NOKIA ICE-X multi-carrier transceiver (TRx), from which we derived the key performance model relating Received Optical Power (ROP) and Signal-to-Noise Ratio (SNR). This model enabled statistical analysis of a fully converged metro-access optical network under both BER and latency constraints.

The simulated network includes two metropolitan areas—referred to as City A and City B—as depicted in Fig. 2. Each city features both metro (M) and access (A) segments. The access segment is structured in clusters, each composed of four access nodes, and each node serves two RUs. These access nodes are connected to the metro segment via Reconfigurable Optical Add-Drop Multiplexers (ROADMs) and circulators, enabling seamless integration between single-fiber bidirectional (BiDi) transmission in the access segment and fiber pairs in the metro segment. The network emulates a realistic urban deployment, making it suitable for assessing route feasibility for RAN transport. To limit computational complexity, a maximum hop count of 12 was imposed on each RU–CU path evaluated.

- City A: 44 optical nodes, including 12 metro nodes and 25 access nodes connected to RUs. The network has an average node degree of 3.42 and an average fiber link length of 6.9 km, with a maximum link length of 10 km.
- City B: 63 optical nodes, of which 18 belong to the metro segment and 45 to the access segment. The average node degree of the metro is 4.06, and the maximum fiber link length is 12 km.

We assume that DUs and CUs are being hosted at Points of Presence (PoPs) inside the metro network and Data Centers

(DCs) outside the metro network, respectively. Both cities, A and B, are connected through CUs located at M31 and M32. Each RU–CU route is evaluated on the basis of BER and latency constraints defined for FH and MH. The total SNR for each route accounts for the combined contributions of the optical channel (GSNR) and the transceiver noise (SNR_{TRX}), computed as:

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

As the NOKIA ICE-X multicarrier transceiver operates with 16 subcarriers (SCs) under DP-16QAM with an overall spectral occupation of 64 GHz, each SC carrying 25 Gbps [24]. In our analysis, we assume that a couple of SCs can serve the bidirectional (BiDi) traffic request of an RU, which aligns with the capacity and performance characteristics under the functional split 7-2, i.e., 22.2 Gbps per RU [22]. BER estimation is therefore performed assuming DP-16QAM modulation, consistent with the experimental characterization of the NOKIA ICE-X multi-carrier transceiver, which supports 400 Gbps transmission and can be potentially widely deployed in metro networks. The BER is then estimated using Eq. 2 for DP-16QAM:

$$\text{BER} = \frac{3}{8} \operatorname{erfc} \left(\sqrt{\frac{\text{SNR}}{10}} \right) \quad (2)$$

Latency has been calculated for both fronthaul (RU–DU) and midhaul (DU–CU) segments by converting physical distances into latency. In the case of SSMF, a propagation delay of $\approx 5 \mu\text{s}/\text{km}$ has been considered, while $\approx 3.45 \mu\text{s}/\text{km}$ considering 30% lower latency with a propagation loss of 0.1 dB/km in the case of HCF [13]. Given HCF's lower propagation delay and reduced attenuation, it is expected to

perform better compared to SSMF to enable BER compliance (i.e., 10^{-3} and 10^{-2}) over large distances.

BER is estimated from the accumulated optical impairments along the path utilizing the transceiver model, while latency is derived from the actual fiber length and its propagation speed. This approach ensures that feasibility reflects realistic network conditions for both SSMF and HCF. A route is considered feasible if conditions of BER, FH latency, and MH latency (i.e., 10^{-3} , $250 \mu\text{s}$, and $1000 \mu\text{s}$ (1 ms) respectively) are met [22, 23]. This dual-criteria assessment provides a comprehensive indication of route viability for supporting RAN transport over the converged optical network.

III. RESULTS AND DISCUSSION

We performed a feasibility analysis while comparing using SSMF and HCF for X-Haul under different conditions. A complete end-to-end connection from each RU to CU has been considered by combining paths from RU→DU (FH) and from DU→CU (MH). To evaluate the impact of MH distance on RAN transport feasibility, we performed a comparative analysis of SSMF and HCF under increasing MH distances, as shown in 3. Here, we focused on only selected variable inter-city links ranging from 10 km to 250 km, namely M26↔M31↔M16 and M30↔M32↔M18, as depicted with red dotted lines in Fig. 2.

We assessed the feasibility of the RU–CU route using SSMF and HCF by checking compliance with BER threshold of $< 10^{-3}$ and latency constraints of $250 \mu\text{s}$ and $1000 \mu\text{s}$ for FH and MH, respectively. Different constraint combinations, i.e., (i) all constraints involving both BER and latency; (ii) BER-only; and (iii) latency-only, where FH latency and MH latency were also evaluated separately. This layered evaluation provides a detailed understanding of how each constraint affects feasible routes using both fiber technologies. As shown in Fig. 3, the most stringent case, where both BER and latency constraints must be satisfied simultaneously, reveals a clear distinction between the performance of SSMF and HCF. For SSMF, feasibility begins to drop sharply after 100 km and falls to 0% at 110 km. In contrast, HCF maintains over 60% feasibility up to 190 km and then eventually falls to 0% 240 km. This can be attributed to HCF’s lower propagation delay and better signal quality, which together help satisfy both latency and optical signal-to-noise requirements more consistently across longer distances.

To further analyze this behavior, we examined the individual effect of each constraint. When only the BER constraint is applied (ignoring latency), both fiber types maintain high feasibility upto a distance of 100 km. However, HCF outperforms SSMF beyond 100 km and performs till 230 km due to its improved tolerance to nonlinear effects, which results in better SNR and lower BER over long distances. Latency-only evaluations highlight the benefits of HCF’s reduced propagation delay. When applying constraints of fronthaul latency (FL) and midhaul latency (ML) independently, HCF maintains 70% feasibility up to 190 km and remains more than 50% upto 220 km. In comparison, SSMF feasibility starts to




Fig. 3: Feasibility percentage versus variable link distance considering SSMF and HCF

decrease after 130 km and declines significantly beyond 180 km and drops to 0% at 200 km. This indicates that latency is the dominant constraint at longer distances. When only MH latency constraint ($< 1000 \mu\text{s}$) is considered, feasibility remains relatively stable for both SSMF and HCF across the entire range of 130 km, then SSMF starts decreasing and eventually drops to zero at 200 km, while HCF starts decreasing at 200 km and remains above 50% upto a distance of 250 km. The results clearly indicate that in deployments constrained by fronthaul latency and signal degradation, HCF offers a viable pathway to extend transport reach and emerges as a key enabler for future-proofing X-Haul infrastructure in latency-critical and high-capacity RAN environments.

To further investigate the impact of the HCF link on MH latency, we analyze a specific RU→CU path while incrementally increasing the length of a selected variable link. The end-to-end route under consideration starts from access node A32 and terminates at CU node M31, passing through M4. The complete path includes the following sequence of nodes: A32(RU)→A31→M14→M9→M10→M11→M5→M4(DU)

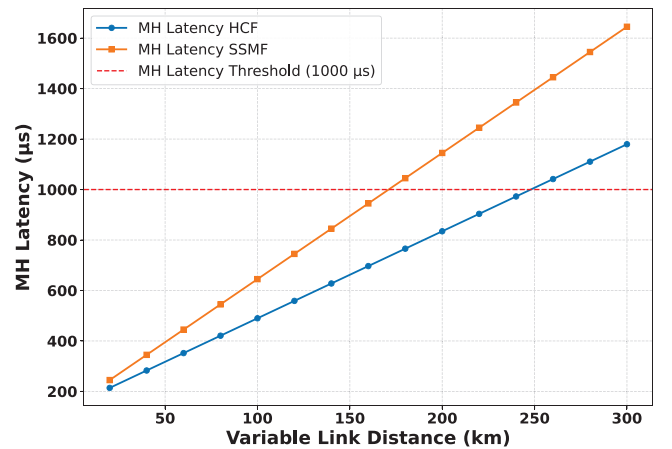


Fig. 4: MH Latency vs MH Distance considering route A32-M4-M31

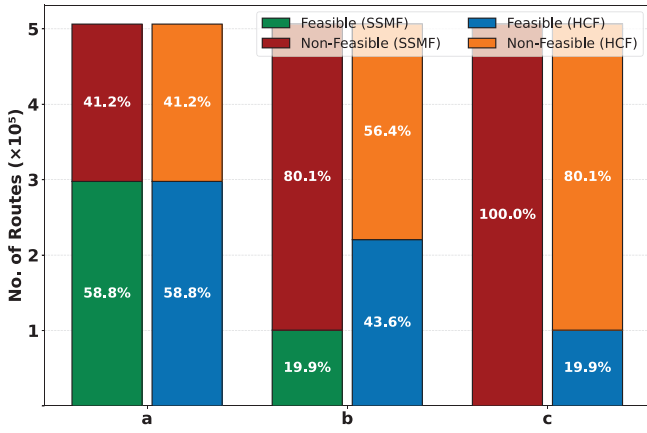


Fig. 5: Feasibility Comparison using SSMF and HCF

→M3→M6→M2→M16→M31(CU), where M4 acts as the DU and M31 as the CU. The key MH segment under study is the link M16→M31, with a variable link distance to simulate different regional deployment scenarios. In this analysis, we focus exclusively on midhaul latency (ML), which is computed by adding up the propagation delays across all optical links from the DU→CU. Propagation delay is modeled as $\approx 5 \mu\text{s}/\text{km}$ for SSMF and approximately $\approx 3.45 \mu\text{s}/\text{km}$ for HCF, considering a 30% lower latency. As shown in Fig. 4, MH latency increases linearly with the variable link length for both fiber types. However, HCF consistently exhibits lower latency across the entire distance range. For SSMF, the latency crosses the $1000 \mu\text{s}$ threshold at approximately 200–220 km of variable link length, indicating that the RAN MH constraint is no longer satisfied beyond this point. In contrast, HCF remains within the $1000 \mu\text{s}$ bound up to roughly 280–300 km, extending the feasible range by nearly 80 km. This result quantitatively confirms that the use of HCF in latency-sensitive MH segments can significantly expand the allowable transport distance without exceeding functional split constraints.

As shown in the network topology 2, each Central Unit (CU), located at nodes M31 and M32, is connected to both City A and City B through three distinct fiber links indicated with color. These CU–DU connections span a range of physical distances, which were grouped into three cases to model increasingly challenging MH scenarios, as summarized in Table I. Each case aggregates all RU–CU paths that traverse through the corresponding FH and MH segments, enabling a comparative feasibility analysis across fiber types.

Fig. 5 summarizes the feasibility comparison between

Links → Cases ↓	<i>M25–M31–M7</i> <i>M30–M32–M18</i> (km)	<i>M28–M31–M10</i> <i>M28–M32–M10</i> (km)	<i>M24–M32–M6</i> <i>M26–M31–M16</i> (km)
a	30	40	50
b	60	80	100
c	120	140	160

TABLE I: Link Distances considered for feasibility analysis

SSMF and HCF across three representative distance scenarios. In Case (a), representing the shortest distances, both SSMF and HCF exhibit similar performance, with 58.8% of routes feasible and 41.2% non-feasible. This confirms that for short FH and particularly MH distances, both fiber types can satisfy the combined BER and latency requirements of the RAN architecture. However, in Case (b), feasibility diverges significantly. With increasing DU–CU distances, the feasibility for SSMF drops to just 19.9%, while HCF maintains 43.6% feasibility. This highlights the advantage of HCF’s lower propagation delay and improved signal quality in longer metro or regional MH deployments. The gap becomes even more pronounced in Case (c), where links (in red color) reach 120–160 km. Here, HCF enables 19.9% of the routes to remain feasible, while SSMF fails to satisfy the combined BER and latency requirements for any route. This sharp drop-off in SSMF feasibility is primarily due to cumulative propagation delay and optical impairments, which exceed allowable thresholds for both FH and MH segments in functionally split RAN architectures. It clearly demonstrates that HCF substantially increases the feasible coverage area for RAN transport in converged metro-access networks. Its ability to sustain long-distance connectivity under strict performance constraints makes it a promising candidate for future-proof, performance critical, and capacity demanding X-Haul deployments.

IV. CONCLUSION

This paper presents a comprehensive feasibility assessment of using SSMF and HCF for RAN X-Haul transport within a converged metro-access optical network. By modeling realistic RU–DU–CU routes across a metropolitan topology, we evaluated how latency and BER constraints impact the viability of end-to-end connections. The analysis considered a range of deployment scenarios, including variable midhaul distances, single-path latency evolution, and grouped inter-city link cases with increasing fronthaul and midhaul lengths. Across all cases, HCF consistently showed better performance in satisfying stringent BER and latency requirements, particularly as link distances increased beyond 100 km. The feasibility improvements with HCF’s lower propagation delay and improved signal integrity enable end-to-end performance across longer metro-regional paths in compliance with stringent RAN functional split requirements.

The results indicate that while SSMF may suffice for short to medium X-Haul deployments, its feasibility sharply declines in long-reach scenarios where latency becomes a dominant constraint. In contrast, HCF significantly expands the range of feasible routes, even under the most restrictive joint BER and latency conditions. This positions HCF as a strong candidate for supporting future 5G-Advanced and 6G transport architectures, where capacity, latency, and scalability are paramount. As the industry transitions toward flexible and disaggregated RAN deployments, our results underscore the importance of fiber type selection in X-Haul infrastructure planning and highlight the role of HCF in enabling performance-critical, future-proof converged transport networks.

In future, we will extend this work in several directions. First, while the current analysis assumes DP-16QAM modulation format based on the NOKIA ICE-X transceiver specification, we plan to investigate adaptive modulation strategies exploiting the transceiver's full capabilities, including DP-QPSK and DP-8QAM formats. This will enable elastic sub-carrier allocation to optimize reach–capacity trade-offs (e.g., QPSK delivering 12.5 Gbps per subcarrier, requiring four subcarriers per RU for bidirectional X-Haul traffic). Second, we will conduct a techno-economic analysis of selective HCF deployment to quantify cost–performance trade-offs, which will become increasingly relevant as commercially mature HCF solutions become available. Third, the study will be extended to consider spatial-division multiplexing (SDM) technologies, such as multicore fibers (MCF) and ribbon fibers, as promising approaches to massively increase network capacity and resilience, in line with the trends highlighted in [25]. Finally, we will assess the applicability of HCF in point-to-multipoint (P2MP) architectures, such as Fiber-to-the-X (FTTx) scenarios, to explore its feasibility in supporting both mobile and fixed access networks under strict latency and capacity constraints.

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 grant agreement No. 101073265. This work has also received support from the Marie Skłodowska-Curie NESTOR project under grant agreement No. 101119983, the EU Horizon Project ALLEGRO grant agreement No. 101092766, the project PNRR-NGEU (MUR–DM117/2023), and the EU Next Generation EU partnership on 'Telecommunications of the Future' (PE00000001 – program 'RESTART').

REFERENCES

- [1] S. Marinova and A. Leon-Garcia, "Intelligent O-RAN beyond 5G: Architecture, use cases, challenges, and opportunities," *IEEE Access*, vol. 12, pp. 27088–27114, 2024.
- [2] A. Tzanakaki, M. Anastasopoulos, V. Alevizaki, and A. Manolopoulos, "Optical networking in support of 5G and 6G infrastructures," in *Proc. Int. Conf. Optical Network Design and Modeling (ONDM)*, 2025, pp. 1–3.
- [3] A. Tzanakaki, M. Anastasopoulos, and V.-M. Alevizaki, "Optical transport networks converging edge compute and central cloud: An enabler for 6G services," in *Proc. Optical Fiber Communication Conf. (OFC)*, 2024, pp. 1–3.
- [4] R. Hui, B. Zhu, R. Huang, C. T. Allen, K. R. Demarest, and D. Richards, "Subcarrier multiplexing for high-speed optical transmission," *J. Lightwave Technol.*, vol. 20, no. 3, pp. 417–, 2002.
- [5] H. Sun *et al.*, "800G DSP ASIC design using probabilistic shaping and digital sub-carrier multiplexing," *J. Lightwave Technol.*, vol. 38, no. 17, pp. 4744–4756, 2020.
- [6] CISCO, "Cisco Annual Internet Report (2018–2023)," 2020. [Online]. Available: <https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-cll-741490.html>
- [7] P. A. Baziana, "Optical data center networking: A comprehensive review on traffic, switching, bandwidth allocation, and challenges," *IEEE Access*, 2024.
- [8] C.-X. Wang *et al.*, "On the road to 6G: Visions, requirements, key technologies, and testbeds," *IEEE Commun. Surveys Tuts.*, vol. 25, no. 2, pp. 905–974, 2023.
- [9] P. Poggiolini and F. Poletti, "Opportunities and challenges for long-distance transmission in hollow-core fibres," *J. Lightwave Technol.*, vol. 40, no. 6, pp. 1605–1616, 2022.
- [10] K. Borzycki and T. Osuch, "Hollow-core optical fibers for telecommunications and data transmission," *Applied Sciences*, vol. 13, no. 19, p. 10699, 2023.
- [11] M. Ibrahim, G. S. Sticca, F. Musumeci, and M. Tornatore, "On the benefits of hollow-core fiber in next-generation optical networks," in *Proc. Int. Conf. Optical Network Design and Modeling (ONDM)*, 2025, pp. 1–5.
- [12] G. S. Sticca, M. Ibrahim, N. Di Cicco, F. Musumeci, and M. Tornatore, "Hollow core fiber as a long-term solution for capacity scaling in optical networks," in *Proc. Optical Fiber Commun. Conf. Exhib. (OFC)*, 2025, pp. 1–3.
- [13] Y. Chen *et al.*, "Hollow core DNANF optical fiber with <0.11 dB/km loss," in *Proc. Optical Fiber Communication Conf. (OFC)*, 2024, pp. Th4A–8.
- [14] B. Correia, J. Pedro, and N. Costa, "Hollow-Core Fiber Specifications for Competitive Deployment in Regio/Long-Haul Optical Networks," in *Proc. OFC*, 2025, pp. Tu3L3.
- [15] Microsoft, "Satya Nadella's keynote at Microsoft Ignite 2024," 2024. [Online]. Available: <https://news.microsoft.com/ignite-2024/>
- [16] D. Suslov *et al.*, "Low loss and high performance interconnection between standard single-mode fiber and antiresonant hollow-core fiber," *Sci. Rep.*, vol. 11, p. 8799, 2021.
- [17] A. Zhong *et al.*, "Connecting Hollow-Core and Standard Single-Mode Fibers With Perfect Mode-Field Size Adaptation," *J. Lightwave Technol.*, vol. 42, no. 6, pp. 2124–2130, 2024.
- [18] A. Nespolo *et al.*, "Ultra-Long-Haul WDM Transmission in a Reduced Inter-Modal Interference NANF Hollow-Core Fiber," in *Proc. OFC*, 2021, pp. F3B.5.
- [19] G. T. Jasion *et al.*, "Hollow Core NANF with 0.28 dB/km Attenuation in the C and L Bands," in *Proc. OFC*, 2020, pp. Th4B.4.
- [20] G. T. Jasion *et al.*, "0.174 dB/km Hollow Core Double Nested Antiresonant Nodeless Fiber (DNANF)," in *Proc. OFC*, 2022, pp. Th4C.7.
- [21] B. Correia and J. Pedro, "Potential of Hollow-Core Fibers to Reduce Number of in-line Amplifiers in Optical Transport Networks," in *Proc. ONDM*, 2025, pp. 1–6.
- [22] FLEX-SCALE Project Consortium, "Deliverable D2.1 – 6G Network Requirements," Horizon 2020 Project, 2023. [Online]. Available: <https://6g-flexscale.eu/en/outcomes>
- [23] M. Klinkowski, "Latency-aware DU/CU placement in convergent packet-based 5G fronthaul transport networks," *Applied Sciences*, vol. 10, no. 21, p. 7429, 2020.
- [24] 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.*, vol. 41, no. 4, pp. 1175–1191, 2023.
- [25] P. J. Winzer, "The future of communications is massively parallel," *J. Opt. Commun. Netw.*, vol. 15, no. 10, pp. 783–787, Oct. 2023.