

Techno-Economic and Feasibility Study of Point-to-Multipoint Communications in the Metro-Core

P. Soumplis^{1,2}, K. Christodoulopoulos^{1,3}, P. Kokkinos^{1,4}, A. Napoli⁵, M. Hosseini⁵, M. Quagliotti⁶, E. Riccardi⁶, A. Pagano⁶, K. Yiannopoulos^{1,7} and E. Varvarigos^{1,2}

¹*Institute of Communication and Computer Systems, Zografou, Greece*

²*School of Electrical and Computer Engineering, National Technical University of Athens, Zografou, Greece*

³*Department of Informatics and Telecommunications, National and Kapodistrian University of Athens, Athens, Greece*

⁴*Department of Digital Systems, University of the Peloponnese, Sparti, Greece*

⁵*Infinera, Munich, Germany*

⁶*TIM-Telecom Italia S.p.A., Turin, Italy*

⁷*Department of Informatics and Telecommunications, University of the Peloponnese, Tripoli, Greece*

*soumplis@mail.ntua.gr

Abstract— We investigate the introduction of point-to-multipoint (P2MP) transceivers in a realistic scenario inspired by the structure of real metro-regional DWDM networks. By establishing light-trees to aggregate traffic from multiple nodes towards the metro-core and backbone, we investigate how P2MP can complement existing point-to-point (P2P) architectures in reducing the number of required transceivers while optimizing the optical signal-to-noise ratio (OSNR) budget. Our analysis of light-tree construction in relation to the available OSNR budget of the traffic aggregation nodes demonstrates that the efficient deployment of P2MP transceivers can significantly reduce reliance on P2P connections and lower overall transceiver costs, particularly when the OSNR is sufficient to support light-tree lengths of approximately 120 km.

Keywords— Coherent transceivers, coherent point-to-multipoint, metro optical networks

I. INTRODUCTION

Optical coherent P2MP technologies present an alternative to P2P, offering benefits that include lower cost and enhanced flexibility and efficiency in managing network resources [1], [2]. Coherent P2MP transceivers utilize digital sub-carriers (DSCs) which can be allocated independently, and, as a result, a single high-capacity transceiver is capable of concurrently serving the traffic from multiple lower rate sources. This leads to an all-optical tree-like communication scheme, where the high-capacity transceiver serves as the “hub” and the lower capacity ones constitute the “spokes”. Hub-and-spoke communications are especially advantageous in the access-metro network, characterized by low traffic between spokes. It has been previously demonstrated that P2MP can reduce the number of required transceivers by approximately 50% compared to point-to-point (P2P) configurations. Furthermore, P2MP transceivers, through the utilization of DSCs, can eliminate the need for intermediate electrical packet aggregation points [3], resulting in significant advantages such as lower hardware requirements for transceivers and networking elements (switches and routers). Since the intermediate aggregation points typically involve O/E/O conversions and packet switching, both of which introduce latency, P2MP can also contribute to reduce the latency in the network. This could prove of key importance in

emerging applications in 5G (6G in perspective) era, such as tactile internet for remote operation, Digital Twins (DTs), Industrial IoT (IIoT) with cloudification, AR/VR and holographic type communication, to name a few [4], [5].

Currently, the all-optical hub-and-spoke paradigm is under investigation for use in the access-metro network segment, where its integration is particularly advantageous. This is largely because the metro aggregation nodes, which collect traffic from mobile radio sites and PON OLTs in a given access area, typically connect to metro-core nodes over horseshoes. Each horseshoe is served by two metro-core nodes, one at each endpoint, making it straightforward to install P2MP hubs at one or both of the metro-core nodes and P2MP spokes at the aggregators. All the traffic from an aggregation node, arranged in DSCs, is directed toward one of the two hubs of the horseshoe where it is optically terminated and electrically aggregated with other traffic. Aggregated traffic from this metro core is then optically transported P2P to one of the Service Points of Presence (S-POP) of the network for service delivery; these are typically located in so called National nodes. This approach constitutes a “baseline” architecture wherein the hubs are strictly associated with the horseshoes they terminate on.

The baseline architecture exhibits two main shortcomings: (i) traffic coming from aggregation nodes belonging to a given horseshoe is constrained to be terminated on hubs on metro core nodes belonging to the same horseshoes: this reduces the possibilities to optimize the use of DSCs on the hubs (i.e., use them at their full capacity, which means using all available DSCs). (ii) In addition, many P2P connections, at least one for each metro core node that hosts one or more hubs, are required to carry the traffic from metro core nodes to S-POP nodes.

The first problem can be tackled by disassociating the optical P2MP hubs from the horseshoe’s terminal metro-nodes, allowing spokes to connect to any optically reachable hub in a path toward the S-POP. This approach increases the multiplexing gain and fully utilizes the capacity of P2MP transceivers but also permits spoke assignment based on additional criteria, such as specific hubs that accommodate regional datacenters. The second challenge can be addressed by allowing hubs to migrate deeper in the metro-core towards the backbone, ideally placing hubs in S-POP nodes, which alleviates the need for P2P intermediate connections and multiplexes traffic from numerous horseshoes.

This work has been funded by the EU H2022 ALLEGRO project (Grant agreement ID: 101092766).

We propose to replace light-paths with P2MP light-trees [6] to obtain an advantage from the reduction of intermediate aggregation points and packet switching. Each light-tree uses a distinct set of contiguous and non-overlapping DSCs that are collectively routed in the metro-core network using currently available technologies and a modified ROADM architecture that was recently proposed [7].

The rest of this paper is organized as follows: Section II provides an overview of the topology considered in the study which is inspired by the real structure of the metro-regional network of TIM. Furthermore, Section II gives the details regarding the introduction of P2MP communications via light-trees. Section III elaborates on the light-tree formation algorithm and the parameters used in simulations. Section IV presents simulation results for the network topology under two different traffic scenarios. In a centralized architecture, S-POPs including network functions (e.g., UPF and BNG) and services (CDN caches, OTT content servers, etc.) are located exclusively in the National nodes of the metro-core network; the proposed light-tree based architecture demonstrates a significant cost reduction compared to the baseline architecture, assuming a sufficient OSNR budget at the horseshoes. Analyzing the selected network scenario, the number of P2MP transceivers is minimized when root nodes are positioned approximately 120 km from the horseshoes, requiring around 50 light-trees to adequately serve all leaf node traffic demands. Moreover, P2P connections are completely eliminated when root nodes are stationed 200 km from the horseshoes, the maximum distance between horseshoe hubs and theirs nearest National node. When traffic is directed towards a combination of destinations, such as a National node and a Datacenter, light-trees are similarly formed, yet P2P connections are still necessary to link National nodes with the Datacenter. It is also suggested that a full P2MP implementation won't be competitive unless P2MP transceivers become more affordable than P2P ones.

II. TOPOLOGY AND P2MP ARCHITECTURE

A. Reference Topology

The regional-metro network considered in this study has the structure of Fig. 1 and comprises an aggregation layer, a metro-core layer and interconnections with the National backbone. This network architecture features horseshoes that link aggregation nodes (in blue) with two metro-nodes at the ends, that hierarchically may be either Regional nodes (in red) or National nodes (in green). The aggregation nodes collect the traffic coming from a variable number of mobile sites (few

to many for each node) and fixed access (OLTs of PONs) of a given access area. For simplicity, the access layer is omitted from this representation as it falls outside the scope of our analysis. The number of aggregation nodes in a horseshoe can vary from one to eight, interconnected through the blue links. The metro-regional core mesh (red links) connects the Regional nodes (both hub, in red, or non-hub, in orange) and National nodes (in green), within a particular macro region. As for the aggregation nodes, each National or Regional node collects the fixed and mobile traffic of its own access area. Within this metro-regional core mesh, there is typically also a node, which plays the role of large Data Center (DC) for the entire macro region. The National nodes are interconnected to the National backbone (by orange links) and assure the exchange of traffic between the macro regions and outside the operators' network, i.e. peering and the Internet.




Fig. 2. The topology of the metro-regional core mesh.

As regards the network's structure used in numerical experiments, Fig. 2 shows the meshed topology of the metro-regional core mesh (horseshoes are not shown), emphasizing fiber optic links equipped with Dense Wavelength Division Multiplexing (DWDM) systems. This topology consists of 95 nodes in total, with 89 serving as Regional nodes—66 of these act as hubs, while the remaining 23 solely collect their traffic without being part of an aggregation horseshoe. Additionally, there are 5 National nodes and 1 Data Center. The core mesh features 139 links, with lengths varying from 1.4 to 110.2 km, averaging 26.8 km. The network accommodates 342 aggregation nodes, organized into 87 horseshoes anchored to pairs of Regional or National nodes, with an average of 3.9 Aggregation nodes per horseshoe (ranging from 1 to 8, Fig. 4(a)). The average link length within horseshoes stands at 21.4 km (from 0.6 to 166 km), with the total length of horseshoes averaging 95.6 km (from 15.2 to 95.6 km).

The metro-core topology shown in Fig. 2 and its aggregation horseshoes extension (not depicted) is representative of a big macro region of TIM's network. For confidential reasons it was obtained by modifying the real network, maintaining its statistical properties. Concerning the service architecture, it is assumed that the telco network functions (e.g., Border Network Gateway (BNG) and User Plane Function (UPF)) and services (content caches, compute, storage) are hosted in the National nodes that play the role of S-POP, in other words the architecture is centralized. Consequently, traffic adopts a hub-and-spoke pattern, with all traffic from access areas (collected at Aggregation, Regional, and National nodes) being directed to the nearest National node, where the telco network functions reside, before it is either terminated or relayed to its final destination.




Fig. 1. Hierarchical structure of the metro-regional network used in the study.

B. P2MP Architecture

The incorporation of P2MP transceivers in the network under study relies on the establishment of light-trees. In Fig. 3 both baseline and P2MP architecture are shown to give evidence to the differences between the two. In Fig 3 (a), which shows the baseline architecture, aggregation nodes terminate their traffic at hubs located within one of their horseshoe's metro core nodes; in this specific case they are the green, blue and orange hubs. P2P connections must then carry traffic from the hubs to the National node where the S-POP is located. In Fig 3 (b), which shows the P2MP architecture, light-trees are formed between the aggregation nodes, that reside in the horseshoes, and any of the metro-core nodes which satisfies optical feasibility constraints of the optical paths between the Aggregation nodes and the metro core node where the tree is originated. In the example tree 1 (green) is originated in a Regional node while tree 2 (orange) is originated in a National node. The aggregation nodes constitute the leaves of the light-tree and utilize lower capacity coherent transceivers to transmit their data in the form of DSCs. Each leaf node is assigned with DSCs at distinct wavelengths and modified ROADMs combine the DSCs from multiple leaves to form the light-tree. Modified ROADM merges the DSCs in the direction from leaf nodes to the hub along the whole tree up to the root node transceiver. In the opposite direction, the root node transceiver broadcasts the DSCs to all the leaf nodes that belong to the light-tree using the intermediate ROADMs. Broadcasting is required due to the tight guard-bands between the DSCs, which makes it infeasible to separate them using the currently available WSSs that are utilized in the ROADMs. The leaf nodes receive all the DSCs of the corresponding light-tree and process their own traffic via coherent demultiplexing.




Fig. 3. Baseline architecture (a) and proposed P2MP architecture that is based on light-trees (b).

III. EVALUATION OF THE PROPOSED ARCHITECTURE

A. Light-tree Formation Algorithm

The light-trees that connect the aggregation (leaf) nodes with the metro-core (root) nodes (Regional or National) are subject to physical layer and cost constraints. For example, the available OSNR at the output of a horseshoe may not be adequate and a root node must be therefore placed nearby or even at the metro-core node that terminates the horseshoe. On the other hand, the roots should be placed on National nodes whenever the available OSNR allows for it, so as to minimize additional P2P connections between the metro-core and National nodes S-POP. In addition, the root placement also affects the utilization of the corresponding P2MP transceivers, since a combination of nearby leaf nodes that fully occupy the available DSCs must be found. Finally, each leaf node may be equipped with a single transceiver and belong to one light-

tree, since multiple low-capacity transceivers are expected to have a higher cost than a high capacity one.

The light-tree formation problem becomes a variant of bin-packing since each “bin” (root transceiver) can only accommodate “items” (leaf transceivers) within its reach. Our proposed heuristic algorithm for this problem relies on best-fit decreasing [8], which provides close to optimal solutions for bin packing, with the additional physical layer constraints. The leaf nodes from all horseshoes are first sorted according to their requests and higher requests are served first. The leaf nodes are also categorized depending on their available OSNR. The assignment process then takes place in two phases: the light-trees of the first phase are formed from leaf nodes that do not have an adequate OSNR to reach the nearest National node. These nodes are served by roots that are placed on Regional metro-core nodes, and additional P2P connections are required to transport traffic to the nearest National node. When each new light-tree is first formed in this phase, multiple metro-core nodes can possibly host the new root depending on their distance from the first leaf node. Then leaf nodes are added to the tree and the possible root positions are gradually restricted to less (or one) metro-core nodes. In the second phase, the remaining nodes are similarly connected to an existing light-tree or a new light-tree is established.




Fig. 4. The distribution of (a) leaf nodes and (b) traffic requests. Each DSC carries 25 Gb/s.

B. Simulation Parameters

The topology considered for the simulations is shown on Fig. 2. The traffic requests were generated by 342 leaf nodes (aggregation nodes) that were located inside 87 horseshoes. The number of leaf nodes inside a horseshoe ranged between 1 and 8, while most horseshoes were formed from 3 to 5 leaf nodes. The exact distribution is shown in Fig. 4(a). The horseshoes were served by 71 metro-core nodes (66 Regional and 5 National), while all traffic was transported to 5 National nodes and one Datacenter as S-POP. The traffic requests were also based on real data, and their distribution is shown in Fig. 4(b). Most of the leaf nodes requested 4 DSCs, while the maximum requests amounted to 15 DSCs.

We consider a fixed 16-QAM dual polarization constellation for transceivers with capacity of $\{4, 8, 16, 32\}$ DSCs. Each DSC carries 25 Gb/s and a maximum rate of 800 Gb/s can be attained. The transceiver cost is assumed to double when its capacity quadruples [2], amounting to $\{1, 1.5, 2, 3\}$ arbitrary units (a.u.) for the four candidate P2MP transceivers, respectively. The cost of equal-rate P2P transceivers ranges from 30% to 90% of the P2MP device, which is close to the lower and upper limits of the values used in [9]. Unless stated otherwise, the relative P2P cost is assumed equal to 90% of P2MP. Regarding the physical layer considerations, a full study is beyond the scope of this work, thus we considered OSNR degradations that range between 1.5 and 4.5 dB per 80 km span in the metro-core network. These values originate from initial estimations using the GN model and a worst-case approximation of the OSNR degradation. Note that, even with close spacing, there can be an optimized baud rate that minimizes the nonlinearities compared to single channel transmission [10]. The reach of each light-tree will be limited by the available OSNR budgets and the OSNR degradation that is introduced by transmission of optical signals in the metro-core network. Results presented hereafter assume different value of OSNR budget from 0 dB (baseline architecture, DSCs coming from leaf are constrained to be hubbed within horseshoe) to 10.5 dB (with a degradation of 1.5 dB/span this budgeted corresponds to a maximum distance of 560 km within metro core network).




Fig. 5. Total transceiver cost for varying levels of OSNR degradation in the metro-core.




Fig. 6. Number of light-trees for varying levels of OSNR degradation in the metro-core.




Fig. 7. Distribution of P2MP transceivers for an OSNR degradation equal to 3 dB per span.

C. All traffic is directed to National nodes

In our initial simulation setup, we considered that the whole traffic is transported from the Aggregation nodes to the National nodes. The corresponding results for total transceiver cost are shown in Fig. 5, which plots cost as a function of the OSNR budget available at the output of the horseshoes. The OSNR budget is defined as the additional OSNR available above the threshold for error-free operation for 16-QAM, and recent studies have shown that significant OSNR budgets can be achieved by properly designing the horseshoes [11]. This budget can be utilized to propagate optical signals in the metro-core network, thus forming light-trees with extended reach, and the reach of each leaf node can be calculated from the OSNR budget and OSNR degradation in the metro-core.

The impact of the available OSNR budget on the required transceiver costs is also shown in Fig. 5. In the “baseline architecture”, a budget of 0 dB is available, and optical transmissions must terminate at horseshoe endpoint metro-core nodes. P2MP root transponders are placed at these nodes, and numerous P2P connections are required to reach the backbone. Since the baseline relies heavily on P2P, associated




Fig. 8. Distribution of P2P transceivers for an OSNR degradation of (a) 1.5 dB/span, (b) 3.0 dB/span, and (c) 4.5 dB/span.

cost is high. However, if the OSNR budget increases, placing P2MP root transceivers further from horseshoes and deeper in the metro-core network [3] becomes possible. This approach reduces the number of light-trees constructed, as it becomes easier to find the correct combination of leaf node capacities that fully utilize P2MP roots when multiplexing multiple horseshoes.

The number of light-trees constructed is shown in Fig. 6, and it can be verified that the baseline requires 68 light-trees. However, the multiplexing gain increases with the OSNR budget, and the number of light-trees reduces to approximately 50. It can also be observed that the OSNR budget required to attain the maximum multiplexing gain increases with metro-core OSNR degradation, as expected. The reduced number of light-trees leads to an equal reduction in the number of P2MP root transceivers that are required.

The P2MP transceiver distribution is shown in Fig. 7 for an OSNR degradation of 3 dB per span. It can be verified that the total number of P2MP transceivers initially reduces and then stabilizes to its final value for an OSNR budget of 4.5 dB. After this point, which translates to a metro-core distance of approximately 120 km, the P2MP cost has been minimized. Similar results hold true for both OSNR degradations of 1.5 and 4.5 dB, which are omitted for clarity. Still, the main

contribution to the cost reduction of Fig. 6 comes from the elimination of P2P transceivers at higher OSNR budgets.

The distribution of P2P transceivers is shown in Fig. 8, which demonstrates that a significant number of P2P transceivers is required in the baseline architecture, but they are gradually phased out as the leaf nodes acquire enough OSNR budget to reach the National nodes. According to the figure, the total minimum cost is achieved for the budget that is required is 4.5, 7.5 and 10.5 dB for an OSNR degradation of 1.5, 3.0 and 4.5 dB per span, respectively. The three budgets correspond to propagation distances between 190 and 240 km in metro-core network.

Finally, Fig. 9 illustrates the total transceiver cost comparison, showcasing a range of lower transceiver costs from 70% down to 30%. For brevity, we only consider the 3 dB per span OSNR degradation scenario. It is evident from the data that the overall transceiver cost tends to decrease as the cost of P2P transceivers is reduced. As anticipated, the minimum cost threshold is effectively reached when the OSNR budget is set to 7.5 dB, a level at which the cost benefits are realized irrespective of the varying P2P costs. This occurs because, at this specific budget point, the reliance on P2P connections significantly diminishes and eventually vanishes.

D. Traffic is directed to the National nodes and the regional Datacenter

In the second simulation setup, we assumed that 50% of the traffic is directed towards the National nodes and the other 50% is designated to reach the Data Center. For this scenario, we consider that the traffic is efficiently split either at the high-capacity root node or at the individual leaf nodes. In the former case, which we denote as “P2P-split”, the light-trees are optimally formed using the total traffic requests. If the root of the tree is placed on a Regional metro-core node, two lower capacity P2P connections are utilized to reach the closest National node and the Data Center, otherwise, if the root is situated in the National node, a single P2P connection is used to forward from the National node to the Data Center the portion of traffic (50%) that is destined for the Data Center. In the latter case, which we denote as “P2MP-split”, separate light-trees are individually formed for the National nodes and Data Center traffic. Consequently, each Aggregation node logically corresponds to two leaf nodes, each with half the original capacity, and since each leaf node may only belong to a single light-tree, two P2MP transceivers are required per




Fig. 9. Total transceiver cost for different relative P2P costs. The OSNR degradation is equal to 3 dB per span.




Fig. 10. Comparison between P2MP and P2P split for the transport of back-bone and Datacenter traffic.

Aggregation node. Following the aforementioned logic, the P2P split requires approximately twice as many P2P transceivers, whereas the P2MP split necessitates approximately twice as many P2MP leaf node transceivers.

The simulation results are shown in Fig.10 plotting total transceiver cost vs OSNR budget. Both P2P- and P2MP-split approaches demonstrate trends observed in the previous subsection, with cost reducing as horseshoe OSNR budget increases and metro-core OSNR degradation decreases. A minimum steady cost is attained for high OSNR budgets, but for different reasons. In the P2MP-split, cost is lowered because P2P connections vanish at high OSNR budgets, where light-tree roots are placed on National nodes and the Data Center. In the P2P-split, cost is lowered since each light-tree initially requires two P2P connections to the National and Data Center nodes at low OSNR budgets, but a single P2P connection is needed when the light-tree root is placed on either node. In both approaches, P2MP transceivers reach their minimum number before the P2P ones.

Given the P2MP and P2P transceiver costs we have considered, the P2P-split approach exhibits a considerably lower cost. The best strategy for the reference network with the given traffic requests is therefore to (a) attain a high OSNR budget at the horseshoes so that long reach light-trees are formed, and (b) interconnect the National nodes and the Data Center via P2P connections. If P2MP-split is to become competitive, then the cost of the P2MP transceivers needs to be lower than the equal-rate P2P counterparts. In our simulations, this was achieved when the P2MP transceiver cost was set equal to 86% of the reference values {1, 1.5, 2, 3}, while the P2P transceiver cost remained fixed to 90%.

IV. CONCLUSION

In this study, we have utilized data, inspired by a real network configuration, to investigate the introduction of P2MP transceivers within the metro-core domain. P2MP communications require a fundamentally different approach from the currently utilized light-paths, and for this purpose, we have utilized light-trees as the underlying structure to solve this complex problem. Our research demonstrated that the OSNR budget available at the horseshoe endpoints plays a

critical key role in determining the number and type of transceivers that are required. Low OSNR budgets necessitated the extensive use of P2P transceivers, with a subsequent increase in associated costs. Conversely, high OSNR budgets significantly favour the deployment of P2MP technologies, reducing both the overall number of light-trees that must be constructed and the number of P2P connections required to effectively connect the light-trees with the National nodes and a Data Center. For the considered network topology, the minimum transceiver cost can be achieved if the OSNR budget is sufficient for the optical signals to effectively go across for approximately 200 km within the metro-core network. In this case, the maximum multiplexing gain is achieved (with the minimum number of light-trees) and only a minimal number of P2P connections remain. Finally, a full P2MP architecture is both technically feasible and economically competitive when the traffic is directed exclusively to the backbone. When a Data Center is introduced into the network, the best approach is to use P2MP connections to link the National nodes or Data Center and then interconnect them with P2P connections. Based on these findings, future research will focus on enhancing the dynamic adaptability of the proposed P2MP architecture by integrating protection mechanisms, machine learning algorithms for real-time network management and impairment estimation along with the joint consideration of the availability of computing resources at nodes. These enhancements are expected to further improve the scalability and resiliency of next-generation network deployments.

REFERENCES

- [1] D. Welch, A. Napoli, J. Bäck, *et al.*, "Digital Subcarrier Multiplexing: Enabling Software-Configurable Optical Networks," *J. Lightwave Technol.*, vol. 41, no. 4, pp. 1175-1191, 2023.
- [2] J. Bäck, A. Napoli, E. Riccardi, *et al.*, "A Filterless Design with Point-to-multipoint Transponders for Cost-Effective and Challenging Metro/Regional Aggregation Topologies," *ONDM 2022*, pp. 1-6.
- [3] A. Napoli, Z. Stevkovski, J. Jiménez, *et al.*, "Enabling Router Bypass and Saving Cost using Point-to-Multipoint Transceivers for Traffic Aggregation," *OFC 2022*, paper W3F.5.
- [4] 5G-ACIA, "Industrial 5G Edge Computing – Use Cases, Architecture and Deployment", white paper.
- [5] ITU-T Technical Report, "FG-NET2030 – Focus Group on Technologies for Network 2030; FG-NET2030-Sub-G1 - Representative use cases and key network requirements for Network 2030", online.
- [6] P. Pavon-Marino, N. Skin-Pop and A. Napoli, "Tree-determination, routing, and spectrum assignment using point-to-multipoint coherent optics," *J. Opt. Commun. Netw.* 15, C29-C40, 2023.
- [7] K. Yiannopoulos, P. Soumplis, K. Christodouloupoulos, *et al.*, "Enabling Extended Access with Light-Trees and Point-To-Multipoint Coherent Transceivers," *ONDM 2024 (submitted)*.
- [8] D. S. Johnson, A. Demers, J. D. Ullman, *et al.*, "Worst-Case Performance Bounds for Simple One-Dimensional Packing Algorithms," *SIAM J. Sci. Comput.*, vol. 3, no. 4, pp. 299-325, 1974.
- [9] N. Skorin-Kapov, F. J. Moreno Muro, *et al.*, "Point-to-Multipoint Coherent Optics for Re-thinking the Optical Transport: Case Study in 5G Optical Metro Networks," *ONDM 2021*, pp. 1-4.
- [10] P. Poggiolini *et al.*, "Analytical and Experimental Results on System Maximum Reach Increase Through Symbol Rate Optimization," in *Journal of Lightwave Technology*, vol. 34, no. 8, pp. 1872-1885, 15 April 2016.
- [11] C. Castro, A. Napoli, M. Porrega, Johan *et al.*, "Scalable filterless coherent point-to-multipoint metro network architecture," *J. Opt. Commun. Netw.* 15, B53-B66, 2023.