Enabling Extended Access with Light-Trees and Point-to-Multipoint Coherent Transceivers

K. Yiannopoulos[1, *], P. Soumplis[1], K. Christodoulopoulos[1,2], P. Kokkinos[1], A. Napoli[3], M. Hosseini[3], and E. Varvarigos[1]
[1]*Institute of Communication and Computer Systems, Zografou, Greece*
[2]*Department of Informatics and Telecommunications, National and Kapodistrian University of Athens, Athens, Greece*
[3]*Infinera, Munich, Germany*
**kyianno@uop.gr*

***Abstract*—We present an architecture that supports light-trees so that point-to-multipoint (P2MP) transceivers can be placed deeper in the metro-core network. It enables flexible aggregation from multiple access nodes in horseshoes or rings, maintaining the P2MP multiplexing gain under heavy traffic. The architecture simulation shows that the advantages are more pronounced in settings where the light-trees span a metro-core distance equal to 200 km irrespective of the traffic conditions.**

Keywords—Coherent transceivers, coherent point-to-multipoint, node architecture, access optical network, metro optical network

I. INTRODUCTION

The deployment of P2MP transceivers based on digital subcarrier multiplexing (DSM) in the optical access and metro segments of the network has been extensively studied. It was shown that it can lower the equipment cost over point-to-point (P2P) solutions, and also provide flexibility in the allocation and management of resources [1], [2]. An important aspect to be investigated is the use of coherent P2MP transceivers in an “extended access” network, where numerous access nodes (spokes) are interconnected using horseshoe (or ring) topologies, and their traffic is directed and aggregated in P2MP hubs at the horseshoe ends [2]. This omits, thanks to the use of DSM, intermediate aggregation points, thus reducing the total number of transceivers [3]. A second benefit of this approach is the reduced latency, since each aggregation step involves O/E/O conversion that adds delay on its own and should be combined with appropriate low-latency switching techniques to avoid additional queueing delay. Note that low latency is of key importance in emerging applications such as 5G, industrial automation, and AR/VR, to name a few [4]. In addition, energy savings can be achieved from the reduction of O/E/O conversions and of the utilized transceivers. Finally, the transport of anycast, multicast and broadcast traffic is inherently supported via P2MP connections.

The current approach is to install at the two ends of a horseshoe P2MP hubs that multiplex traffic from spokes belonging to that horseshoe [2]. Since the number of spokes per horseshoe is relatively small and some spokes can have high loads close to the total capacity of the P2MP transceivers, a limited number of efficient aggregation combinations is available, which reduces the multiplexing gains. The advent of next-generation P2MP transceivers, expected to double the available capacity, will temporarily restore the multiplexing gain, however, the same issue will re-appear as the spoke traffic grows. A solution that mitigates this effect is to disassociate the hubs from the horseshoes. In this approach, the hubs reside deeper in the metro-core network and connect to multiple horseshoes over the existing metro-core infrastructure. A spoke may now freely connect to any hub in the vicinity of its horseshoe end and the multiplexing gain is maximized by properly assigning spokes to hubs. Moreover, the spoke assignment can be performed based on additional criteria that, for instance, favor the assignment to specific hubs accommodating the regional datacenters or backbone nodes. This further extends the access, reducing the P2P metro connections from hubs to the datacenters or backbone nodes.

In the proposed solution, the P2MP communication over the metro-core is implemented using light-trees as described in [5]: the P2MP hubs reside at the root of the light-tree and the spokes constitute the leaves. However, light-trees are not compatible with ROADMs in metro networks designed for light-paths [5]. To address this, we present a novel ROADM architecture that is suitable for P2MP in Section II. In addition, moving the P2MP transceivers deeper in the network requires a suitable allocation algorithm. We propose in Section III an algorithm that takes into account the required quality-of-transmission (QoT) and can be extended to include additional constraints, such as latency, to efficiently assign spokes to hubs. Our results show that the proposed approach imparts a significant reduction in the cost of transceivers and achieves the maximum multiplexing gain when the hubs are placed within 200 km from the horseshoes in realistic networks.

II. LIGHT-TREE BASED P2MP ARCHITECTURE

The proposed architecture in Fig. 1 utilizes the existing horseshoes and light-trees to connect spokes (leaves) to hubs (roots). We consider that the leaf traffic requires several digital sub-carriers (SCs), generated for example by a PON OLT, a 5G cell site, or a user aggregation node at the access network. Each leaf node is equipped with a transceiver that fully serves its capacity and connects to a P2MP transceiver at the root of the tree. In the “downstream” (root-to-leaf) direction, the root transceivers broadcast their SCs to all connected horseshoes via modified ROADMs, discussed in the next paragraph. The leaf nodes receive the traffic from multiple light-trees and make use of coherent demultiplexing to isolate their own SCs. In the “upstream” (leaf-to-root) direction, the SCs from all the leaf nodes are passively combined inside the horseshoes. The

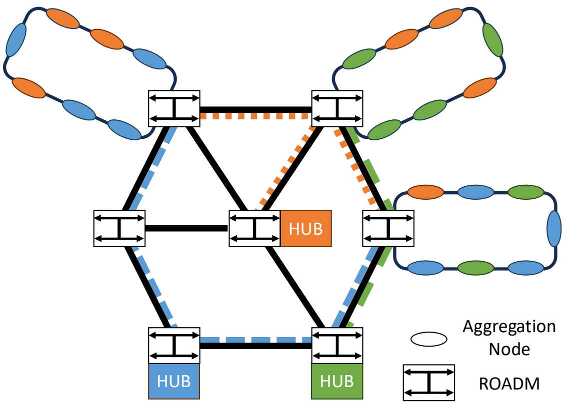

Fig. 1. Proposed extended access P2MP architecture based on light-trees.

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

978-3-903176-62-1 © 2024 IFIP

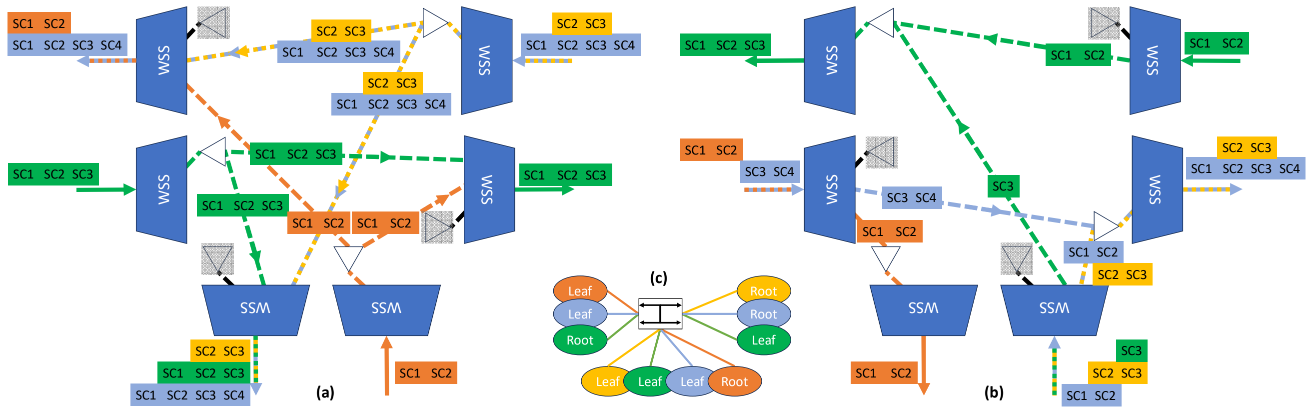

Fig. 2. Three-degree ROADMs with (a) downstream/split and (b) upstream/merge functionalities. (c) Formed light-trees.

modified ROADMs combine the SCs that belong to the same P2MP root transceiver over a light-tree.

The functionalities required to form light-trees, “splitting” and “merging” SCs, are not supported by current ROADMs [5]. An architecture that performs these functionalities and is compatible with WSS technologies is proposed in Fig. 2. Downstream traffic from the roots is directed to a common port of the input WSS and a power splitter-combiner (PSC) broadcasts the SCs (originated at the root) to all output WSSs. The output WSSs allow the root SCs to pass only if a leaf is located towards the corresponding direction. Following Fig. 2(a), the SCs from the West incoming root (green) are broadcasted and allowed to pass to all outputs, while the SCs from the South root (orange) are only allowed by the Western output WSS. In the upstream direction, shown in Fig. 2(b), the input WSS transfers the SCs from leaf nodes to the port that connects to the direction of their root. Since the leaf nodes of a light-tree may connect to different ROADM inputs, another PSC merges the SCs from all inputs before they enter the output WSS (or the hub). Thus, an additional port at each WSS is required for the PSCs; the input WSS PSC is used as a splitter and the output WSS PSC as a combiner. Compared with the conventional ROADM architecture, the introduction of PSCs incurs negligible cost and the additional losses that are introduced are offset by the existing ROADM amplifiers.

It should be noted that this architecture inherits the restrictions of the WSSs and the transceivers, namely the SCs of each P2MP transceiver are transmitted contiguously and cannot be separated by the WSSs due to their tight guardbands (sub-GHz), while the WSSs can handle the whole bandwidth (wavelength) of a light-tree. Given these, each transceiver may belong to a single light-tree. However, there is no restriction on the placement of root/leaf nodes, and multiple root and leaf nodes may connect over the same links, as long as they use different light-tree wavelengths, as shown in Fig. 2(c). The leaf nodes can also be assigned to any light-tree within their reach and are allocated with the requested number of SCs up to the respective P2MP transceiver capacity. Due to the flexible grids of the WSSs, the light-tree capacity is not fixed and can be chosen according to the available P2MP transceiver capacities and the traffic of the leaf nodes that it aggregates. Finally, the capability to switch conventional light-paths is not affected; they are directed to regular WSS ports without crossing the PS-Cs.

III. EVALUATION OF THE PROPOSED ARCHITECTURE

The proposed solution was simulated to evaluate its benefits. We assumed a metro-core mesh network where each node is connected to a horseshoe that comprises several leaf nodes. Certain metro-core nodes are the desired destinations having attached a Datacenter or connecting to the backbone. Given that a leaf node needs a transceiver for each tree it connects to, and assuming multiple smaller capacity transceivers are costlier than a single higher capacity one, an efficient approach is to serve each leaf node in a single light-tree. The assignment of leaf nodes to roots is a variant of the bin packing problem, where the bin size equals the transceiver capacity. To solve this problem, we used a novel heuristic that is based on the best-fit decreasing algorithm [6], where the capacities of all leaf nodes are sorted in descending order. Leaf nodes are then sequentially assigned to roots that (a) are located within the distance dictated by QoT requirements, namely a sufficient OSNR threshold for 16-QAM, and (b) have the requested remaining capacity. This process takes place in two phases: in the first phase, the algorithm considers only leaf nodes that are outside the reach of the desired roots. For those, the P2MP root transceivers are placed at intermediate metro-core nodes and P2P transceivers are used to connect the metro-core node and the closest backbone node. The remaining leaf nodes are assigned during the second phase to either (a) the light-trees that were formed in the first phase, or (b) new light-trees with roots located at the desired destinations only. If multiple candidate roots exist during a phase, the leaf node is assigned to the one with the lowest remaining capacity, while if no root can serve the leaf node, a new P2MP root transceiver is placed.

For the simulations, we consider the regional topologies (domains A-E) of Telefonica’s Spanish network [7]. In each domain, six backbone nodes are connected to a national/ultra-speed network and play the role of the desired roots. Each metro-core node connects to a horseshoe with 10 leaf nodes and can be equipped with root P2MP transceivers. The leaf nodes request SCs that are distributed uniformly and the maximum number of SCs is set equal to 10, 20, or 30 to assess the effect of varying traffic loads [2]. We also consider transceiver capacities of {4, 8, 16, 32} SCs with each SC carrying 25 Gb/s. 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 P2MP transceivers, respectively. The cost of the equal-rate P2P transceivers is set to 90%, close to the upper limit of values in [8], and a more conservative 70%, which is the average of the values in [8]. A full study of the physical layer is beyond the scope of this work. As a simplification, we considered an OSNR degradation of 3 dB per 80 km in the regional network. This is based on a basic estimation using the GN model and a worst-case approximation of the

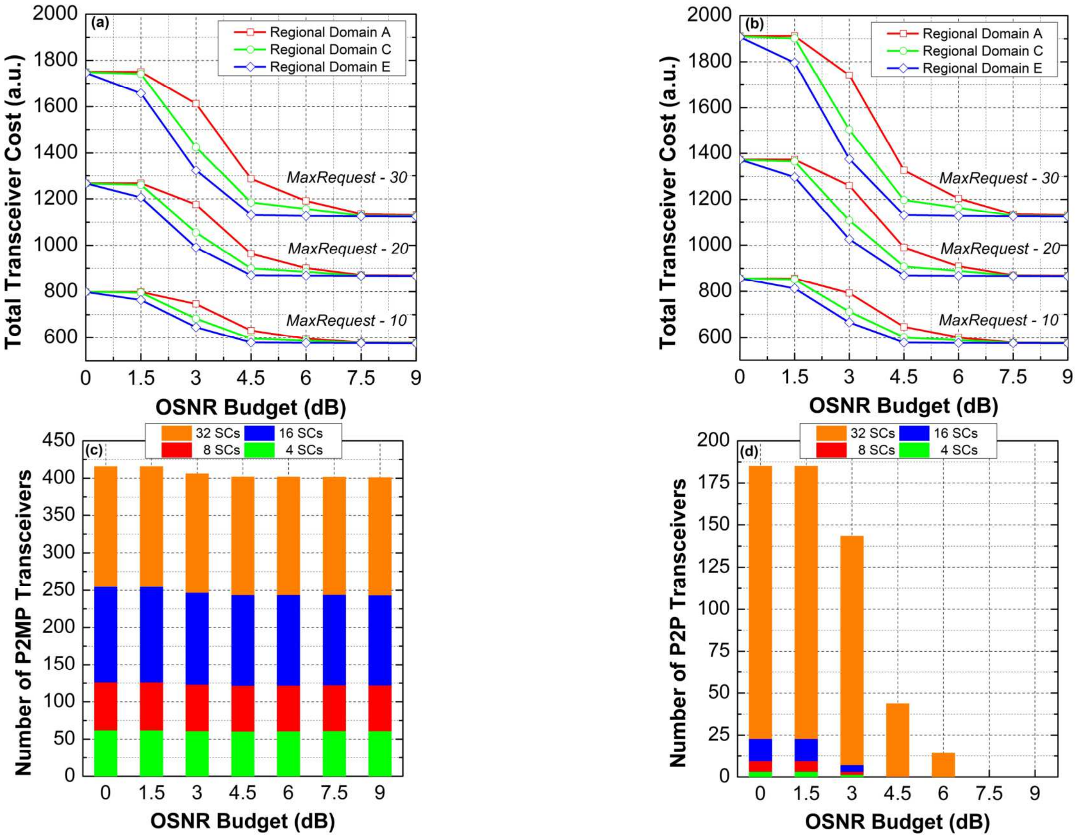

Fig. 3. Transceiver cost for TID regional domains A, C, E and a relative P2P cost equal to 70% in (a) and 90% in (b). Number of (c) P2MP and (d) P2P transceivers for domain A and maximum traffic requests equal to 20 SCs.

OSNR degradation, given that nonlinearities seem to play an important role due to the close spacing of the SCs (~ 4 GHz).

Fig. 3(a) and (b) plot the total transceiver cost against the available OSNR budget at the horseshoes' ends, when entering the metro-core network. As expected, a higher traffic load results in the placement of a higher number of transceivers and a higher cost. The cost is also higher at lower OSNR budgets when the relative P2P transceiver cost amounts to 90% since P2MP cannot reach the desired backbone nodes. Considering a baseline architecture, where the available OSNR budget is 0 dB, P2MP roots need to be placed at the metro-core nodes that constitute the endpoints of the horseshoes and each P2MP transceiver only serves a single horseshoe. Thus, numerous P2P connections are required from the metro-core nodes to the backbone nodes, while the multiplexing gain is reduced as it is challenging to find proper combinations of leaf nodes to efficiently utilize the root SCs on a single horseshoe. When the OSNR budget is increased, longer distances can be achieved and the P2MP transceivers migrate towards the backbone nodes thus reducing the number of P2P connections [3], while the multiplexing gain is increased since the light-trees include multiple horseshoes. This reduces the cost and, following Fig. 3(c), the minimum P2MP cost (maximum multiplexing gain) is achieved at an OSNR budget equal to 4.5 dB, which translates to a horseshoe-to-root distance of 120 km in our simulations. After this point, the cost further reduces due to the elimination of P2P connections and the total cost converges to the same value irrespective of the relative P2P cost for an OSNR budget of 7.5 dB (or 200 km). Fig. 3(d) shows that the P2P connections vanish at this distance, and traffic is served exclusively by P2MP transceivers. This provides an extra benefit to the operators in terms of efficient storage and maintenance.

IV. CONCLUSION

We proposed an extended access architecture that employs coherent P2MP transceivers and light-trees for flexible and efficient resource allocation. The architecture enables the placement of P2MP roots deeper inside the metro-core network so as to connect to leaf nodes located at multiple horseshoes, thus yielding higher multiplexing gains. Simulations show that our approach achieves considerable improvements without requiring excessive reach inside the metro network and maintains them even for high traffic loads.

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, 15 Feb.15, 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] P. Pavon-Marino, N. Skorin-Kapov and A. Napoli, "Tree-determination, routing, and spectrum assignment using point-to-multipoint coherent optics," J. Opt. Comm. Netw. 15, C29-C40, 2023.

[6] 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.

[7] IDEALIST Project, "Elastic Optical Network Architecture: Reference scenario, cost and planning," Deliverable D1.1, 2013.

[8] 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.