

Enabling extended access aggregation with light-trees and point-to-multipoint coherent transceivers

P. SOUMPLIS,^{1,*}  K. CHRISTODOULOPOULOS,^{1,2} P. KOKKINOS,^{1,3} A. NAPOLI,⁴ 
M. HOSSEINI,⁴  K. YIANNOPOULOS,^{1,5}  AND E. VARVARIGOS¹

¹Institute of Communication and Computer Systems, Zografou, Greece

²Department of Informatics and Telecommunications, National and Kapodistrian University of Athens, Athens, Greece

³Department of Digital Systems, University of the Peloponnese, Sparta, Greece

⁴Infinera, Munich, Germany

⁵Department of Informatics and Telecommunications, University of the Peloponnese, Tripoli, Greece

*soumplis@mail.ntua.gr

Received 16 September 2024; revised 19 January 2025; accepted 6 February 2025; published 6 March 2025

Traditionally, traffic from the access nodes is aggregated in metro-aggregation hubs over rings or horseshoes in a fixed configuration using multiple optical point-to-point (P2P) transceivers. However, this static setup limits dynamic scalability and often leads to inefficient resource usage and high costs, especially under increasing and varying traffic conditions. Coherent optical point-to-multipoint (P2MP) transceivers offer a promising solution for aggregation at this network level. These transceivers allow a single (aggregation) node to communicate with multiple (access) nodes simultaneously via digital subcarrier multiplexing (DSCM) technology. Additionally, their long-distance transmission capabilities enable the placement of the P2MP transceiver root deeper in the hierarchy. In this paper, we propose an extended access aggregation architecture that incorporates modified reconfigurable optical add-drop multiplexers (ROADMs) to support both traditional P2P and P2MP connections. This architecture allows the creation of light-trees to enable P2MP communication and hence to boost transmission flexibility and multiplexing gains. We propose a mixed-integer linear programming (MILP) model to determine the optimal placement of P2MP transceivers and the establishment of light-trees, considering physical layer impairments (PLIs) and the required quality of transmission (QoT) for the connections. To tackle the high algorithm complexity, we also introduce a best-fit decreasing heuristic to efficiently exploit the trade-off between execution time and performance. Our simulation experiments using real network topologies showcase that our proposed architecture can greatly enhance multiplexing gains while considering the operational costs tied to network expansions and upgrades. © 2025 Optica Publishing Group. All rights, including for text and data mining (TDM), Artificial Intelligence (AI) training, and similar technologies, are reserved.

<https://doi.org/10.1364/JOCN.542492>

1. INTRODUCTION

The aggregation of traffic from access nodes in metro networks has traditionally been implemented using static configurations, such as ring or horseshoe topologies, which rely on multiple optical point-to-point (P2P) transceivers. While these architectures have proven to be effective, they lack the scalability and flexibility required to manage the increasing and variable traffic demands of modern optical and mobile access networks. Coherent point-to-multipoint (P2MP) transceivers have emerged as a promising solution for cost-effective metro-aggregation, overcoming the shortcomings of conventional ring or horseshoe architectures [1,2]. These transceivers leverage advanced digital signal processing (DSP) and digital subcarrier multiplexing (DSCM) technologies

to provide enhanced spectral efficiency and flexible capacity allocation. This allows the network to adjust to varying traffic patterns and distances. Typically, metro-aggregation networks follow a “hub-and-spoke” paradigm, where several access nodes (spokes) connect on a horseshoe (or ring) and communicate with an aggregation node (hub) at the end of the horseshoe (or ring) which connects to the higher aggregation layer or the metro-core [3]. Within this hub-and-spoke model, coherent P2MP transceivers have emerged as an efficient solution by utilizing low-capacity coherent transceivers at multiple spokes to communicate with a single high-capacity root transceiver at the hub. Each spoke is assigned a distinct set of digital subcarriers (SCs) according to its traffic requirements, and the hub concurrently receives and processes the different sets of SCs

from the spokes. In the other direction, the spokes receive all SCs and isolate their own traffic via coherent demultiplexing. Compared to P2P communications, the P2MP approach reduces the number of transceivers by approximately a factor of 2, since a single high-capacity transceiver is required at the hub; although it might result in a higher cost (but comparable to other same data rate transceivers), this is still better than duplicating the number of devices. At the same time, the use of a single transceiver introduces the potential for multiplexing gains as the hub can accommodate spokes until the available SCs are exhausted. This enables flexible resource management, allowing the allocated SCs to be utilized based on the evolving traffic demands of the spokes.

However, the association of the hub with a single horseshoe restricts the potential multiplexing gains, particularly when some spokes produce traffic loads comparable to the hub transceiver capacity. For example, spokes in urban areas typically generate high traffic loads, making it difficult to locate appropriate low-capacity spokes for multiplexing in the same region. The advent of next-generation transceivers, expected to increase the available capacity by a factor of 2, will temporarily restore the P2MP multiplexing gain. However, the same effect will recur as the spokes' traffic continues to grow. Alternatively, the other solution could be to continuously increase the data rate of the transceivers to better match rising traffic demands. A solution that mitigates this effect is to disassociate the hubs from the horseshoes and position them further in the metro network. In this approach, the hubs utilize the existing metro infrastructure to connect to multiple horseshoes. A spoke can connect to any hub in the vicinity of its horseshoe, and the multiplexing gain is maximized by properly assigning spokes to hubs. Moreover, the spokes can be assigned to hubs based on additional criteria, for instance, specific hubs that accommodate regional data centers or are connected to the backbone network. This omits intermediate aggregation points, thus reducing the total number of P2P transceivers [4,5] and electronic processing equipment (i.e., IP routers) [6].

A second benefit of this approach is the reduced and deterministic latency, as aggregating at the horseshoe-end is avoided. By employing optical aggregation within the P2MP architecture, the need for optical/electrical/optical (O/E/O) conversions, which add delay and typically require additional low-latency switching to prevent excessive queuing, is considerably reduced and, in certain configurations, can be completely eliminated. Low latency is crucial for supporting emerging applications, such as 5G, industrial automation, and AR/VR, to name a few. Additional benefits might include efficient transport of anycast, multicast, and broadcast traffic. Finally, energy savings can be realized by the reduction of the O/E/O conversions and the number of transceivers used. Note that although coherent P2MP transceivers, which have a similar cost to P2P ones, are initially more expensive than lower-capacity P2P transceivers, the potential for multiplexing gains and hardware consolidation can still lower overall costs in the long run, particularly in environments with expanding traffic loads [3].

The proposed solution requires the incorporation of P2MP connections within the metro network, which relies on reconfigurable optical add-drop multiplexers (ROADMs) and is

predominantly designed for P2P connections (light-paths), typically offering multiple paths between its nodes. Our solution establishes light-trees [7–9], with P2MP hubs positioned at metro-nodes as roots of the light-trees, while the spokes reside at intermediate points of the horseshoes forming the leaves. A major challenge is the incompatibility of light-trees with the operation of the current ROADMs in metro, which are designed to support light-paths. To address this issue, we introduce a novel ROADM architecture that, with minimal additions, supports both light-trees and light-paths, making it compatible with the P2MP transceivers and light-trees; this is a major contribution of this work. Notably, the only similar architecture that can support light-trees has been reported in Ref. [10], but it is oriented toward data center applications.

Moreover, establishing light-trees requires efficient resource allocation and root placement algorithms, particularly as the transmission distances between hub (P2MP root) and spoke (P2MP leaf) nodes increase, and thus, physical layer impairments (PLIs) and the quality of transmission (QoT) play a more important role. Specifically, PLIs could make infeasible certain direct optical communications between a spoke and a desired hub node, e.g., data center or backbone. In such cases, an intermediate metro-node should be selected to terminate the P2MP traffic, obtaining statistical multiplexing gains. Then, a new transceiver, such as P2P, should be employed to restore the QoT and steer the traffic toward the target node. In this work, we propose two algorithms—a mixed-integer linear programming (MILP) and a heuristic—that optimize the placement of P2MP transceivers. These approaches take into account the PLIs and the cost of transceivers and leverage the strengths of both P2P and P2MP technologies in various network and traffic configurations. In Ref. [11], we introduced the architecture that supports light-trees to facilitate the placement of P2MP transceivers deeper within the metro-core network, and in Ref. [12], we investigated how P2MP can complement existing P2P architectures in reducing the number of required transceivers. In the current study, we extend the aforementioned works by providing a formal definition of the problem, which was not explicitly detailed previously. We explore both fixed and flexible P2MP transceiver configurations to further increase the aggregation and reduce the cost. We introduce a MILP formulation aimed at optimizing the placement of these transceivers, considering not only the connectivity and capacity but also cost factors and PLIs that impact the QoT. Note that while the concept of P2MP communication through sliceable-bandwidth variable transceivers (S-BVTs) has been explored for core and metro applications [13,14], our work focuses on P2MP communication as part of a light-tree-based architecture for metro-aggregation. Unlike S-BVTs, which are typically used in limited P2MP scenarios and make use of multiple lasers and modulators, our approach leverages DSCM and flexible P2MP transceiver configurations to maximize statistical multiplexing gains when allocating resources. However, the developed mechanisms are generic, designed to optimize the P2MP transceivers' placement, ensuring scalability and cost-efficiency in different network deployments, and thus maintain compatibility also with S-BVT architectures.

The paper is organized as follows: Section 2 presents the light-tree-based P2MP metro network and the proposed

ROADM architecture that implements the functionalities required to support that. Section 3 discusses resource allocation algorithms that take into account the PLIs and the required QoT to efficiently assign spokes to hub–P2MP root transceivers in the metro network. Initially, we introduce the mathematical formulation of the problem through a MILP, which can be used to calculate the optimal solution. As the execution time can be prohibitive for large problem instances, typically encountered in real-scale scenarios, we also present a suboptimal greedy heuristic to efficiently trade off execution time with performance. Section 4 presents the simulation outcomes for the MILP and heuristic algorithms. Our results show that the proposed approach leads to a significant reduction in the transceivers' cost and achieves the maximum multiplexing gain when the hubs can be placed within 170 km from the horseshoes in realistic networks. Section 4 also details the applicability of flexible P2MP transceivers, supporting multiple modulation formats, with the goal of further reducing the intermediate aggregation points. The results show that the transceivers' cost reduces further when flexible P2MP transceivers are introduced, while there is also an indirect advantage due to the minimization of O/E/O conversions and energy consumption. Finally, Section 5 summarizes the main findings and concludes the article.

2. LIGHT-TREE-BASED P2MP ARCHITECTURE

The proposed architecture is illustrated in Fig. 1 and utilizes the existing horseshoes and light-trees to connect spokes (leaf nodes) to hubs (roots). We consider that the leaf node traffic requires several SCs, generated by one or more PON OLTs, 5G cell sites, or, in general, a traffic aggregation node of the access network. Each leaf node is equipped with a transceiver that fully serves its aggregated traffic and connects to a P2MP transceiver at the root of the tree. In the “upstream” (leaf $\rightarrow$ root) direction, the SCs from all leaf nodes are passively combined in the horseshoes [3]. The modified ROADM, discussed in the next paragraph, treats a terminating horseshoe as one additional degree and forwards the SCs, which belong to the same P2MP root transceiver of a light-tree, toward its root

node. In the “downstream” (root $\rightarrow$ leaf) direction, the root transceivers broadcast their SCs to all connected horseshoes via the modified ROADMs, and the SCs enter the horseshoes of the corresponding leaf nodes. The leaf nodes receive the traffic from multiple light-trees and use coherent demultiplexing to isolate their own SCs [3].

Coherent P2MP transceivers (Fig. 2), placed at both hub and leaf nodes, leverage the DSCM technique [15]. DSCM is a digital communication technique that generates multiple independent data streams simultaneously from a single laser using closely spaced subcarriers, thus enabling near-Nyquist signaling. DSP is employed for efficient digital multiplexing and demultiplexing, acting like filters at the transmitter and receiver, respectively. This method maintains similar spectral occupancy and hardware requirements as a single-carrier transceiver with an equivalent total symbol rate and spectral occupancy. This approach leverages coherent detection with a single transmit and receive laser and enables the proposed architecture to not only handle complex traffic aggregation but also address the challenges of longer distances in metro-core networks compared to access networks. Flexible modulation formats combined with the consideration of the actual OSNR budget left by horseshoes for transmission to the core enable communication over longer distances, making this architecture suited for metro-core scenarios.

The functionalities that are required to form light-trees, “splitting” and “merging” SCs, are not supported by current route-and-select ROADM architectures that are designed to manage P2P connections (light-paths) [16]. Moreover, P2MP transceivers introduce strict constraints on the operation of the ROADMs, since they may only transmit using successive Nyquist-shaped SCs with very tight guard bands (<1 GHz) between them [1]. As a result, it is not possible to isolate and separately manage individual or subsets of SCs inside the ROADMs using the currently available WSS technologies. For this reason, we propose a route-and-select ROADM architecture that enables the “split” and “merge” functionalities and is compatible with existing WSS technologies (see Fig. 3). The “split” functionality applied in the downstream, shown in Fig. 3(a), is implemented using the broadcast-and-select




Fig. 1. (a) Reference P2MP architecture where hubs are associated with horseshoes. (b) Proposed extended access aggregation P2MP architecture based on light-trees.




Fig. 2. Upstream transmission in a P2MP architecture using DSCM: multiple leaf nodes modulate 4–32 digital subcarriers whose signals are merged at a single-hub transceiver.




Fig. 3. Three-degree ROADMs with (a) split and (b) merge functionalities. Light-trees are shown in (c).

concept. All the SCs of the P2MP root are broadcasted toward the leaf nodes, since it is not possible to isolate the SCs of each respective leaf. To be more specific, the downstream traffic is directed to a common port of the input WSS that is connected to a power splitter-combiner module (PSCM) that broadcasts all the SCs (originated at the root) to all output WSSs. An output WSS allows the root SCs to pass only if a leaf node is located toward its corresponding direction. Following Fig. 3(a), the SCs from the west incoming root (green) are broadcasted to and allowed to pass through all outputs, while the SCs from the south root (orange) are only allowed by the western output WSS.

The implementation of the “merge” functionality, which is applied in the upstream direction, requires that the ROADMs combines the SCs from multiple leaf nodes that are located at different inputs, as shown in Fig. 3(b). To this end, the upstream (leaf-to-root) traffic is treated as a conventional light-path at the input WSS, and the input WSS transfers the leaf SCs to its output that connects to the merging port of the output WSS in the direction of the associated root. The output WSS has a dedicated merging port that is equipped with another PSCM, where the SCs from all the leaf nodes are passively merged. As an illustrative example, the leaves of the

“blue” color light-tree are located west and south in Fig. 3(b), and the two-input WSS routes their SCs toward the eastern PSCM. There, the traffic from the “blue” and “yellow” leaf nodes is combined and directed toward the corresponding roots.

Compared with the conventional route-and-select ROADMs architecture, both functionalities require an additional port at each WSS for the PSCMs; the input WSS PSCM is used as a splitter and the output WSS PSCM as a combiner. So, in total, assuming that the ROADMs supports N degrees, the input WSS requires $N - 1$ ports for the regular light-paths, $N - 1$ ports for the light-tree leaves, and 1 port for the light-tree roots. In a similar manner, the output WSS requires $N - 1$ ports for the regular light-paths, 1 port for the light-tree leaves, and $N - 1$ ports for the light-tree roots. The introduction of PSCMs incurs negligible cost, especially in the case that the WSSs have the extra ports, and the additional losses that are introduced can be compensated by the existing ROADMs amplifiers.

It should be noted that both the network and ROADMs architectures are tailored to the restrictions of the WSSs and the transceivers, namely that the SCs of each P2MP transceiver are transmitted contiguously and cannot be separated by the

WSSs due to their tight guard bands (sub-GHz). Given these restrictions, 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. 3(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, which can also handle the whole bandwidth (SC wavelength range) of a light-tree, 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 PSCMs.

To complement the P2MP connectivity, P2P connections can also be utilized at intermediate nodes. Their primary role is to extend P2MP communication in the cases where direct optical connectivity from the leaf to the backbone is limited due to the presence of PLIs or insufficient OSNR budget. In these scenarios, a P2P transceiver provides the connectivity between an intermediate metro-core node and a backbone node, ensuring acceptable end-to-end quality of transmission. Each intermediate node is equipped with a router that aggregates traffic from multiple rings/hubs depending on the direction and interfaces with coherent P2MP transceivers. This architecture enables the relocation of IP routers deeper into the metro-core or backbone network, where larger, higher-capacity routers are utilized. Hence, fewer aggregation nodes manage the increased traffic volumes originating from multiple light-trees, with metro IP routers deployed at these locations to support higher port densities. The consolidations of routers at metro-core nodes can result in cost and power consumption reduction [6].

3. PROBLEM FORMULATION

We assume a network that spans from the access to the backbone, consisting of metro-aggregation (horseshoes) and metro-core (mesh) segments, as shown in Fig. 1. The network comprises nodes of two types: (i) access aggregation or spoke nodes S , which aggregate traffic from lower-hierarchy/access networks and are connected onto horseshoes, and (ii) metro-core nodes D , which participate in the metro-core network, and can be the end of one or more horseshoes, and can connect to the higher-hierarchy/backbone network. In turn, the metro-core nodes are divided into two distinct subsets: (ii.a) metro-backbone nodes F , which are metro-core nodes that connect to the higher-hierarchy/backbone network, and (ii.b) pure metro-core nodes A that do not connect to higher hierarchy, so that $D = F \cup A$ and $F \cap A = \emptyset$.

The traffic is generated/terminated at/to spoke nodes S and forwarded/received to/from backbone nodes F . This scenario can be adjusted to consider one or more regional data centers, allowing for traffic termination at these points. If the metro-core node at the end of the horseshoe acts as a backbone, the spokes' traffic can terminate there without traversing further into the metro-core. In the reference P2MP architecture, the

P2MP root transceivers are placed at the metro-core nodes at the end of each horseshoe, comprising the hubs, and then P2P transceivers are used to connect to the backbone nodes. Our proposed extended metro-aggregation architecture allows P2MP transceivers to be placed deeper within the metro network and, depending on sufficient QoT, e.g., OSNR, enables a single hop from the spoke to the backbone node. Thus, in our proposed architecture, the P2MP transceivers can be placed at any metro-core node in $D_s = D$ (greenfield scenario). Placing a P2MP at a metro-core node that does not act as a backbone (thus belonging to A , not F), due to inadequate QoT to reach the backbone, is also possible since it still provides some multiplexing benefits. However, in this case, like in the reference P2MP architecture, a P2P transceiver must be added to connect to a backbone node.

Each metro-core or backbone node in D_s is equipped with $I(d)$, $d \in D_s$ ports, where P2MP transceivers of different capabilities can be placed. We assume flexible P2MP transceivers where we can select among different modulation formats. The case of fixed P2MP transceivers is a subset and can be immediately accounted for. Flexible P2MP transceivers are described by a tuple $T_b = \{(C_t, W_t, B_t)\}$, where C_t represents the transceiver's monetary cost, W_t indicates the number of subcarriers that the transceiver can handle, and B_t is the symbol rate of the transceiver. We assume that these P2MP transceivers may support a diverse set of modulation formats, represented by M_t , where each $m \in M_t$ defines the bits per symbol (e.g., 2 for QPSK and 3 for 8-QAM) alongside the corresponding acceptable OSNR threshold O_m . Thus, the transmission rate R_t of a transceiver of type t is $R_t = W_t \cdot B_t \cdot m$. Here we also allow different spoke nodes of a light-tree to receive at different modulation formats, so different subcarriers can have different modulation formats in the same P2MP transceiver. On the other end, matching P2MP transceivers with sufficient capacity are placed at the leaf-spoke nodes. Actually, a P2MP transceiver at a leaf node connects to one P2MP root transceiver (corresponding to one branch of the light-tree communication with the single-root P2MP node). However, we refer to leaf transceivers as P2MP since they use DSCM technology to match the P2MP root transceivers. We denote the set of available leaf (spoke) P2MP transceivers by T_s .

As outlined above, due to PLI, P2MP transceivers can also be deployed at various intermediate metro-core nodes in D_s along the route to a backbone node. At these intermediate locations, P2P transceivers are placed, matched with corresponding P2P transceivers at backbone nodes, to enable the backbone-access communication. The P2P transceivers are described by tuples T_a . The demand at a spoke node $s \in S$ is defined by the demand size L^s (in Gb/s) and is assumed to be identical for both (uplink and downlink) directions. Each spoke node s is attached to a horseshoe that ends at a metro-core node $d(s) \in D$, and we assume that all spokes connected to the horseshoe reach the end with the OSNR level $O_{d(s)}$ left by horseshoes for transmission through the core, for simplicity in the physical layer calculations; however, the proposed formulation and algorithms can be adapted to account for varying OSNR levels per spoke. This parameter helps in identifying the reach of the transmission with respect to the PLIs that affect the transmission.

The goal is to allow the spokes to communicate directly with the backbone nodes with P2MP transceivers and optimize the multiplexing gains leading to a predominantly P2MP architecture with minimal P2P reliance, which exhibits potential benefits that include reduced (i) total transceiver costs, (ii) O/E/O conversions, and (iii) energy consumption. From an economic perspective, the optimal hub–leaf ratio arises naturally from balancing cost, capacity, and physical constraints in the network. The proposed optimization models implicitly consider how many leaf nodes can be served by each hub without incurring excessive transceiver or infrastructure costs. Factors such as OSNR budgets, traffic demands, and P2MP transceiver capacities directly affect this balance, ensuring that the selected hub–leaf assignment minimizes total expenditures while still meeting quality of transmission requirements. As a result, the most appropriate ratio will vary depending on each network's unique topology, traffic, and equipment cost profile.

The problem is defined as follows: Given the network topology (S , D , A , and F), the traffic demands (L^s from spokes S to any backbone node F), and the descriptions of the P2MP and P2P transceivers (T_h , T_a , and T_s), our goal is to place P2MP at metro-nodes (D) and P2P transceivers if required at intermediate pure metro-nodes (A) and assign P2MP transceivers to spokes to serve all the traffic while minimizing the overall cost. While we primarily examine greenfield deployment scenarios where P2MP transceivers can be placed freely at any metro-core node in D , the proposed algorithms are also adaptable to brownfield deployment scenarios. In this scenario, two distinct sets of nodes can be considered: $\bar{D} \subseteq D$, representing metro-core nodes with existing, nonmodifiable transceiver placements, and $\hat{D} \subseteq D$, representing nodes where additional P2MP transceivers can be deployed. In that case, the optimization mechanisms consider the union $\bar{D} \cup \hat{D}$, accounting for the specific characteristics and constraints of the underlying infrastructure.

A. MILP Formulation

In the following, we present a MILP formulation to optimally solve the problem for the greenfield scenario.

1. Preprocessing Phase

Initially, in a preprocessing phase, we determine the shortest paths among all metro-core nodes, $d, d' \in D$ and assess the OSNR degradation $O_{d,d'}$. For each spoke node $s \in S$ that is attached to the horseshoe that ends at the metro-node $d(s) \in D$, with $O_{d(s)}$ the OSNR degradation on the horseshoe from s to $d(s)$, we identify the highest modulation format denoted as $m_{s,d}$ that can maintain connectivity to metro-node d such that

$$(O_{m_{s,d}})^{-1} \geq (O_{d(s)})^{-1} + (O_{d(s),d})^{-1}. \quad (1)$$

This step is critical as it determines the number of subcarriers required for the feasible communication between the root and spoke nodes.

2. Variables

$x_{d,i,t}$: Binary variable is equal to 1 if a P2MP root transceiver of type $t \in T_h$ is used at node $d \in D$ at port $i \in I(d)$; otherwise, it is equal to 0.

$z_{a,t}$: Integer variable indicates the number of type t of P2P transceivers used at an intermediate pure metro-core node $a \in A$.

$y_{s,d,i}$: Binary variable is equal to 1 if the demand of the spoke node $s \in S$ is fulfilled by a P2MP transceiver placed at the i th port of node $d \in D_s$; otherwise, it is equal to 0. This communication uses the highest feasible modulation format $m_{s,d}$ as found in the preprocessing phase [see Eq. (1)].

$u_{s,t}$: Binary variable is equal to 1 if a P2MP transceiver of type $t \in T_s$ is placed at the spoke node $s \in S$; otherwise, it is equal to 0.

3. Objective Function

Minimize the total monetary cost of the deployed transceivers, including the (i) P2MP transceivers deployed at the metro-core and backbone nodes ($d \in D$), (ii) the cost of P2P transceivers placed for the single-point communication between the intermediate pure metro-core A and backbone F nodes, and (iii) the cost of the spoke node $s \in S$ transceivers:

$$\min \sum_D \sum_{i \in I} \sum_{t \in T_h} x_{d,i,t} \cdot C_t + 2 \cdot \sum_{a \in A} \sum_{t \in T_a} z_{a,t} \cdot C_t + \sum_{s \in S} \sum_{t \in T_s} u_{s,t} \cdot C_t. \quad (2)$$

4. Constraints

C.1. Each spoke node demand belongs to a single light-tree and must be connected to exactly one port i of a metro-core node D (pure or backbone) to serve its demand as follows:

$$\forall s \in S, \quad \sum_{d \in D} \sum_{i \in I(d)} y_{s,d,i} = 1. \quad (3)$$

C.2. Every port i on a node $d \in D$ can host at most one activated P2MP root transceiver as follows:

$$\forall d \in D, \quad \forall i \in I(d), \quad \sum_{t \in T_h} x_{d,t,i} \leq 1. \quad (4)$$

C.3. If a traffic demand is routed through a specific port on a node, a P2MP root transceiver at that port must be activated to provide the necessary connectivity as follows:

$$\forall d \in D, \quad \forall i \in I(d), \quad s \in S, \quad \sum_{t \in T_h} x_{d,t,i} \geq y_{s,d,i}. \quad (5)$$

C.4. The capacity of an activated P2MP root transceiver at node d and port i must be sufficient to handle the aggregated demands of all spoke nodes routed to it as follows:

$$\forall d \in D, \quad \forall i \in I(d), \quad \sum_{t \in T_h} x_{d,t,i} \cdot W_t \cdot B_t \geq \sum_{s \in S} y_{s,d,i} \cdot \lceil L^s / m_{sd} \rceil. \quad (6)$$

C.5. The P2P transceivers placed at the intermediate pure metro-nodes $a \in A$ must have enough capacity to forward all incoming P2MP demands to the backbone nodes as follows:

$$\forall a \in A, \sum_{t \in T_h} z_{a,t} \cdot R_t \cdot B_t \geq \sum_{s \in S} \sum_{i \in I(a)} y_{s,a,i} \cdot \lceil L^s / m_{sa} \rceil. \quad (7)$$

C.6. The transceiver placed at spoke nodes s must have sufficient capacity to meet the node demands as follows:

$$\forall s \in S, \forall a \in A, \sum_{t \in T_h} u_{s,t} \cdot W_t \cdot B_t \geq \lceil L^s / m_{sa} \rceil. \quad (8)$$

In the aforementioned MILP formulation, let $|I|$ denote the maximum number of ports available at any metro-core node, and the number of variables are $|D| \cdot |T| \cdot |I|$ binary variables for the $x_{d,t,i}$, $|S| \cdot |D| \cdot |I|$ binary variables for the $y_{s,d,i}$, $|S| \cdot |T|$ binary variables for the $u_{s,t}$, and $|A| \cdot |T_A|$ integer variables for the $z_{a,t}$ that take values between 0 and $|I|$. The MILP formulation requires a total of $|S| + |D| \cdot |I| \cdot (2 + |S|) + |A| \cdot (|S| + 1)$ constraints broken down as follows: $|S|$ constraints for ensuring each spoke node's demand are routed to exactly one P2MP root in a metro-core node in D (C.1), $|D| \cdot |I|$ constraints for the placement of transceivers at the different ports of the node (C.2 and C.4), $|D| \cdot |I| \cdot |S|$ constraints for activating transceivers based on the routed demands (C.3), $|A|$ constraints for the capacity of the P2P transceivers at the intermediate pure metro-core nodes (C.5), and $|S| \cdot |A|$ constraints for the transceiver capacity at spoke nodes (C.6).

To consider brownfield deployment scenarios, the MILP formulation needs to be modified to prioritize the utilization of existing P2MP transceivers at specific metro-core nodes. Specifically, nodes with existing transceivers are given priority because they are already purchased. Consequently, the objective function [Eq. (2)] is modified to account solely for the cost of deploying new transceivers at nodes in $\hat{D}$. Furthermore, to ensure the utilization of existing transceivers, Eq. (4) is restricted to select exclusively from the preinstalled transceivers at nodes in $\bar{D}$, while nodes in $\hat{D}$ retain the flexibility to deploy additional transceivers as needed.

B. Heuristic Algorithm

Since solving MILP formulations can be time-consuming, particularly for large metro networks with numerous horseshoes and spoke nodes $|S|$, we also developed a heuristic algorithm to solve the light-tree formation and transceiver placement problem more effectively. Given the cost implications, with multiple smaller capacity transceivers being costlier than a single higher-capacity one, an efficient approach is to serve each leaf node in a single light-tree and also use large root transceivers. The assignment of leaf nodes to roots is a variant of the bin packing problem, where the bin size is equivalent to the transceiver capacity, denoted as W_t .

To solve this problem efficiently, we used a novel heuristic based on the best-fit decreasing algorithm [17]. This heuristic begins by ordering all leaf nodes $s \in S$ in descending order according to their traffic capacity requirements L^s , ensuring that nodes with the largest demands are prioritized. This ordering is crucial because handling larger demands first allows for more flexibility and ensures that these high-capacity nodes can be accommodated without difficulty. Leaf nodes are then

assigned to roots based on two criteria: (a) they are within a distance that meets the QoT requirements, namely a sufficient OSNR threshold O_m , and (b) have the necessary remaining capacity to serve the demand $\lceil L^s / m_{s,d} \rceil$, where $m_{s,d}$ is the highest feasible modulation format for node s when connected to node d .

The process of assigning leaf node demands to light-trees is executed through a multiphased approach as illustrated in Algorithm 1. Initially, the algorithm sorts the demands of spoke nodes S in decreasing order based on their traffic capacity requirements and then examines the demands sequentially (line 5). This is done to ensure that more options are available for the bigger demands, which are harder to accommodate and avoid the fragmentation of the utilized subcarriers. In the first phase, the focus is on the demands at leaf nodes that are outside the reach of metro-backbone nodes (F). These demands necessitate the use of P2MP and P2P transceivers placed at intermediate metro-core nodes $a \in A$ (lines 6–13). For such demands, the highest modulation format (i.e., 16-QAM in our simulations) is selected (line 7) to maximize the data throughput and efficiency. It then identifies potential light-trees with sufficient free subcarrier capacity exceeding the required capacity for the demand. Among these candidate trees, the one offering the minimal excess capacity is selected to serve the demand (lines 8–10). If no such light-tree exists (line 11), the algorithm considers all reachable metro-core nodes as candidates for the placement of the root transceiver of the newly created light-tree. This involves identifying the possible roots of the new light-tree (lines 12–14). Demands allocated to light-trees in this phase are then removed from those pending for assignment (line 13).

The subsequent phase addresses the remaining demands by sequentially processing each remaining leaf node across all modulation formats (line 14), from the highest (16-QAM) to the lowest (BPSK). For each spoke node s that has not been assigned to a light-tree yet (line 15), the algorithm calculates the necessary subcarriers $\lceil L^s / m \rceil$ required to meet the demand under the current examined modulation format m (line 17). All possible light-trees with enough SCs that can serve the demand are identified, and the one with the minimum number of free SCs is selected (lines 18–21). Subsequently, the demand is removed from the pending demands (line 21). If no suitable light-tree is found, the algorithm creates a new light-tree and identifies potential new root locations among the backbone nodes (F) (lines 22–24). The end of this phase marks the end of the assignment of leaf node demands to light-trees. The next step involves optimizing the transceivers used within the light-trees. This is achieved by selecting smaller transceivers that match the remaining free SC capacity in the light-trees (line 25). Finally, we tackle the challenge of determining optimal root nodes for light-trees based on their possible root sets (line 26). This involves selecting a minimal set of root nodes that can efficiently serve all spokes. Given the combinatorial nature of this challenge, we employ a greedy heuristic to assign root nodes to light-trees, iteratively covering all spokes.

Algorithm 1. Pseudocode of the Best-Fit Decreasing Heuristic Algorithm

```

1 Input: Network topology ( $S, D, A, F$ ), traffic demands ( $L^s$ ), transceiver tuples ( $T_b, T_a$ ), OSNR budget, and modulation formats  $M$ 
2 Output: Placement and configurations of P2MP and P2P transceivers
3 Initialize placement, configurations, and tree roots to an empty set
5 Sort  $S$  in descending order based on traffic demand  $L^s$ 
5 For each spoke node  $s$  in  $S$  do:
6   If spoke  $s$  cannot reach any backbone node in  $F$  with  $\min(M)$  then
7   Identify reachable light-trees from  $s$  with free SC capacity exceeding  $\lceil L^s / \max(M) \rceil$ 
8     If at least one light-tree exists, then
9     Select the light-tree with the minimum free SCs
10    Update possible tree roots, free SCs, and configuration
11    else (no light-tree exists)
12      Create light-tree with  $\max_{t \in T_a} (W_t) - \lceil L^s / \max(M) \rceil$  free SCs and roots from  $A$  reachable by  $s$  with  $\max(M)$ 
13      Remove  $s$  from  $S$ 
14 For each modulation format  $m$  from highest to lowest do:
15   While  $S$  is not empty do:
16   Calculate the required number of SCs  $\lceil L^s / m \rceil$ 
17   Identify reachable light-trees from  $s$  with free SC capacity exceeding  $\lceil L^s / m \rceil$ 
18     If at least one light-tree exists, then
19     Select the light-tree with the minimum free SCs
20     Update light-tree's possible roots, configurations, and available SCs
21     Remove  $s$  from  $S$ 
22     else if no light-tree exists and  $s$  can reach a metro-backbone node in  $F$  do:
23       Create a light-tree with  $\max_{t \in T_b} (W_t) - \lceil L^s / m \rceil$  free SCs and roots from  $F$  reachable by  $s$  with  $m$ 
24       Remove  $s$  from  $S$ 
25 Optimize light-tree transceivers by selecting appropriately sized ones based on free SCs
26 Assign root nodes to light-trees using a best-fit greedy approach based on root sets
27 return placements and configurations

```

4. NUMERICAL EVALUATION OF THE PROPOSED ARCHITECTURE

A. Simulation Parameters

We evaluated the proposed architecture that moves P2MP transceivers deeper in the metro and supports light-trees in the metro through simulations. We assumed a metro-core mesh network (the respective nodes are denoted with D) where each node is connected to a horseshoe that comprises several leaf nodes (S). Certain metro-core nodes are the desired destinations F ($F \subset D$) having attached a data center or connected to the backbone.

We initially conducted a limited set of simulations using the optimal MILP algorithm and compared it against the heuristic one, both outlined in Section 3. Subsequently, we extended the simulations using the heuristic to validate the performance of the proposed architecture for various network settings. Both the heuristic and the optimal MILP algorithms were implemented in MATLAB, and the MILP was solved using the Gurobi optimizer. Due to the vast search space inherent in real-size network problems, a time limit of 120 min was used when solving the MILP. The best solution obtained within this timeframe is reported, although it might not necessarily be the global optimum. The simulations were performed on a Windows PC with 16 GB RAM and a Ryzen 5 2600 CPU with six cores.

For the simulations, we considered the regional topologies (domains A–E) of Telefonica's (TID) Spanish network that each contains 30 nodes. Detailed topology diagrams can be

found in Fig. 16 of Ref. [18]. In each domain, six backbone nodes are connected to a national/ultra-high-speed network and play the role of the desired roots (backbone nodes F). Each metro-core node connects to a horseshoe with 10 leaf nodes and can be equipped with up to $I = 10$ root P2MP transceivers. The leaf nodes request SCs that are distributed uniformly between 1 and the maximum number of SCs set equal to 10, 20, or 30 to assess the effect of varying traffic loads [2]. We also considered fixed modulation format and flexible modulation format P2MP transceiver capacities of $W_t = \{4, 8, 16, 32\}$ SCs. Each SC was carrying 25 Gb/s using 16-QAM in the fixed transceiver setting, while the available modulation formats were 16-QAM, 8-QAM, Q-PSK, and B-PSK in the flexible transceiver setting, and the SC capacity was modified accordingly. The corresponding OSNR thresholds for both types of transceivers were set equal to $O_m = \{5.5, 8.5, 12.5, 15.1\}$ dB [19], starting with B-PSK. The transceiver cost was assumed to double when its capacity quadruples [2], amounting to $C_t = \{1, 1.5, 2, 3\}$ arbitrary units (a.u.) for the P2MP transceivers, respectively, where 1 a.u. corresponds to the cost of a 4SC P2MP transceiver. The cost of the equal-rate P2P transceivers is set to 90% of the corresponding P2MP, a value that is close to the upper limit of values in Ref. [20], and a more conservative 70%, which is the average of the values in Ref. [20].

With respect to the physical layer, we performed simulations of the architecture using the GN model [21], assuming a worst-case scenario with 32 fully occupied 32-SC light-trees (1024 consecutive SCs) and SSMF fibers for the metro-core

links. Each SC occupied a 4 GHz bandwidth, totaling over 4 THz, with -9 dBm transmission per subcarrier. The 80 km fiber links were considered to be transparent, with an optical amplifier ($NF = 5.5$ dB) compensating for losses. The modified ROADMs added an OSNR degradation equal to one fiber span, accounting for extra amplification that is introduced by the power splitter-combiner. The OSNR degradation between the horseshoe-end node, $d(s)$, and any other metro node, d , exhibited an almost logarithmic behavior. This degradation was approximated by the following equation:

$$O_{d(s),d}(\text{dB}) = 26.0 - 10\log_{10}(2N_s), \quad (9)$$

where N_s denotes the number of 80 km spans that are traversed. The OSNR at the horseshoe output was also set equal to the 16-QAM threshold (15.1 dB) [19] plus a budget of a few dB that can be exploited to transmit the SCs in the metro network [3]:

$$O_{d(s)}(\text{dB}) = 15.1 + \text{budget}, \quad (10)$$

and the end-to-end OSNR from s to d was calculated in accordance with Eq. (1) as follows:

$$(OSNR_{s,d})^{-1} = (O_{d(s)})^{-1} + (O_{d(s),d})^{-1}. \quad (11)$$

In the considered topologies, the fiber links were not always 80 km, so N_s took noninteger values, calculated by dividing the distance between metro-core nodes by the 80 km span. Note that an OSNR budget of $O_{d(s)} = 0$ dB corresponds to the deployment of P2MP transceivers at the ends of horseshoe (ring) topologies (baseline scenario), leaving no additional OSNR margin for core transmission. In contrast, positive OSNR budgets provide extra margin for longer-distance metro-core links, enabling more flexible hub placements and reducing reliance on intermediate P2P connections.

B. Performance Comparison of the MILP and Heuristic for Flexible P2MP Transceiver Cost Optimization under Varying OSNR Budgets

Initially, in Fig. 4, we compared the performance of the heuristic with the MILP, assuming flexible modulation format P2MP transceivers as a function of the OSNR budget available at the horseshoe output for transmission to the core and for three different SC traffic loads. For each network setting, we conducted 10 simulations and averaged the results over those. As discussed above, a time limit of 120 min was used when solving the MILP, and we report the best solution obtained within that, although it might not necessarily be the global optimum. The MILP approach consistently achieved superior performance, optimizing the traffic multiplexing and effectively utilizing the flexible modulation formats to minimize the transceiver's usage.

Specifically, for the reference P2MP scenario where we reach the horseshoe-end with a 0 dB OSNR margin and have to place the P2MP transceivers there, the heuristic mechanism exhibited a performance deviation of 4.7%, 7.5%, and 6.2% for average traffic demands of 10, 20, and 30 SCs at the spoke nodes, respectively. However, as the OSNR budget is higher, especially from 2 dB and above, the heuristic mechanism




Fig. 4. Comparison of total flexible modulation format transceiver cost for the MILP and the heuristic as a function of the OSNR budget available at the horseshoe output for transmission to the core, for three different maximum traffic loads (10, 20, and 30 SCs).

approximates the optimal solution, with the difference narrowing to less than 0.5%. This is due to the fact that the P2MP roots can be placed at the backbone nodes for this budget, and the complexity of the problem is reduced. The heuristic then exhibits an increased efficacy, especially under higher OSNR conditions, managing to achieve close to optimal solutions with a significant difference in execution time. On average, the heuristic completed in 1.61 s, while the MILP finished earlier than the 120-min limit for lower traffic loads but reached this limit for traffic loads of 30 SCs.

C. Simulation Results for Fixed Modulation Format P2MP Transceivers

In the following results, we only used the heuristic algorithm, which was shown to be quite effective in the previous subsection. We start with simulations assuming fixed modulation format P2MP transceivers and then continue with flexible ones in the next subsection. The simulation results using the heuristic algorithm for fixed modulation format transceivers are shown in Fig. 5, where the total transceiver cost is plotted against the OSNR budget that is available at the horseshoe output. For each network setting, we conducted 1000 simulations and averaged the results over those. As expected, a higher traffic load results in the placement of a higher number of transceivers and a higher cost. Regional domain A exhibits the highest cost since the nodes are located at a higher average distance than the other two domains.

The cost is also higher at lower OSNR budgets when the relative P2P transceiver cost amounts to 90% since P2MP connections cannot reach the desired backbone nodes. By considering the reference P2MP architecture, where the available OSNR budget is 0 dB, the P2MP roots are placed at the metro-core nodes that constitute the endpoints of each horseshoe, and each P2MP transceiver only serves a single horseshoe. In this reference P2MP architecture, 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




Fig. 5. Cost comparison of fixed modulation format transceivers for TID regional domains A, C, and E when the relative P2P transceiver cost equals (a) 70% and (b) 90% of the P2MP transceiver cost.

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 toward the backbone nodes, thus reducing the number of P2P connections [4]. It is also interesting to note that all domains converge to a minimum cost at a relatively limited distance irrespective of the traffic load. This distance is attained for a 2.0 dB OSNR budget and corresponds to approximately 170 km, which is the maximum distance between any metro-core node and its closest backbone node. Therefore, the light-trees do not need to span the whole domain, which could be challenging to achieve, depending on the available OSNR budgets at the horseshoes.

The distribution of transceivers that are required to serve the traffic in domain A is illustrated in Fig. 6. Domain A was selected since it involves the highest average distance between nodes, while the maximum traffic request per leaf was set equal to 30. It can be verified from the figure that numerous P2P connections are required for the reference P2MP architecture (OSNR budget = 0 dB), but they gradually reduce as the OSNR budget increases, and they totally vanish at the distance where the minimum overall cost is attained. Simultaneously, the number of P2MP transceivers also reduces with the OSNR

budget, but in a smoother manner. The number of P2MP transceivers stabilizes at a lower budget of 1.0 dB (approximately 100 km), at which point the maximum multiplexing gain has also been achieved for this specific traffic scenario and network topology. This result highlights the primary benefit of P2MP transceivers: their ability to enable flexible, high-efficiency resource aggregation that reduces the need for additional P2P links, thereby lowering overall network costs. Similar observations can be made for the two other domains and traffic request scenarios.

D. Simulation Results for Flexible Modulation Format P2MP Transceivers

The results of the previous subsection show that the number of P2P connections that are required depends on the OSNR degradations in the metro-core domain. The conversion from P2MP to P2P requires O/E/O conversions that increase latency, require extra mechanisms to avoid varying queuing delay in packet switching, and contribute to the overall cost and energy consumption. It is therefore desirable to further reduce the use of P2P connections and rely predominantly on




Fig. 6. Number of (a) P2P and (b) P2MP transceivers for domain A and maximum traffic request equal to 30 SCs.

P2MP technologies. This goal can be achieved by introducing flexible modulation format transceivers that lower their modulation order when the OSNR degrades below predefined thresholds [19]. In this approach, a leaf node that is located at a longer distance will be assigned an increased number of lower-rate SCs to serve its traffic. The associated P2MP cost is expected to increase, but the P2P cost will reduce, and it is interesting to investigate whether the flexible modulation format transceiver approach is worth pursuing given that P2MP is slightly costlier.

To study this, we employed the heuristic algorithm described in Section 3 to form light-trees that support SCs with varying capacities, determined by the available modulation formats described in previous sections. The primary goal was to assess how flexible P2MP transceivers, when integrated into the network, could impact overall cost and performance. As shown in Fig. 7, the simulation results reveal a significant reduction in total transceiver costs when flexible P2MP transceivers are utilized, despite the expected increase in both SC count and associated costs. The cost reduction is mainly due to the fact that the P2P transceivers are also reduced by more than 50% compared to the proposed architecture with fixed modulation format transceivers. This is detailed in Fig. 8(a),

where we consider the regional domain A with leaf requests up to 30 SCs for a direct comparison with Fig. 6(a). It is also interesting to note that P2P connections completely vanish for all OSNR budgets when the requests are limited to 10 or 20 SCs. The corresponding plots in Figs. 7(a) and 7(b) are identical despite the different P2P cost, which is a clear indication that a full P2MP implementation can be achieved with flexible transceivers when the transceiver capacities significantly exceed the leaf node requests.

Regarding the P2MP transceivers, their number and SC count initially increase due to the use of lower-capacity modulation formats (predominantly 8-QAM), which are utilized to avoid having the leaf traffic aggregated at intermediate metro-core nodes. As illustrated in Fig. 8(b), 32 SC transceivers prevail at low OSNR budgets, reflecting a preference for higher-capacity configurations to optimize network efficiency. When higher OSNR budgets become available, the overall cost of deploying P2MP transceivers begins to decrease, and it stabilizes to its final value for a budget of 2.0 dB, where the minimum number of P2MP transceivers is used. Beyond this threshold, all leaf nodes can reach the backbone nodes using higher-order modulation formats, such as 16-QAM, which maximizes spectral efficiency and minimizes the number of




Fig. 7. Cost comparison of flexible transceivers for TID regional domains A, C, and E when the relative P2P transceiver cost equals (a) 70% and (b) 90% of the P2MP transceiver cost.




Fig. 8. Number of (a) P2P and (b) flexible P2MP transceivers for domain A and maximum traffic requests equal to 30 SCs.




Fig. 9. Traffic percentage that requires P2P connections for maximum traffic requests equal to (a) 10, (b) 20, and (c) 30 SCs.

transceivers needed. At this point, the total cost converges with that of the fixed modulation format transceiver architecture, highlighting the diminishing returns of further increasing the OSNR budget. Thus, the flexible P2MP architecture not only adapts to varying network conditions but also achieves cost parity with fixed modulation format transceiver setups.

The other possible benefits (reduced latency, O/E/O conversions, and energy consumption) are indirectly evaluated by recording the percentage of traffic transported to the backbone nodes using P2P transceivers. So, these have an extra hop as opposed to the proposed architecture using the P2MP transceivers to directly go from the leaves to backbone nodes. This percentage of traffic requiring P2P transceivers is summarized in Fig. 9 for domains A, E, and C, with three different maximum traffic request levels. The figure highlights that in the reference P2MP architecture, where fixed modulation format transceivers are utilized and the OSNR budget is set to 0 dB, around 80% of the leaf nodes are not directly connected to the backbone nodes. Instead, their traffic is aggregated at local horseshoe-end nodes and then forwarded to the backbone using P2P transceivers, introducing additional delays and energy consumption due to the extra hop.

As light-trees are introduced and root transceivers are placed deeper within the regional network, this reliance on P2P connections significantly decreases. When the leaf nodes request up to 20 SCs, almost 100% of the traffic is transported directly over P2MP connections, as shown in Figs. 9(a) and 9(b). This eliminates the need for P2P transceivers, optimizing network efficiency by reducing latency and energy consumption. However, as traffic requests increase to 30 SCs, illustrated in Fig. 9(c), the percentage of P2P traffic drops from 80% to 30% for the 0 dB OSNR budget across all domains. This is because the leaf nodes may not be able to lower their modulation order sufficiently without exceeding the transceiver's capacity, thereby necessitating some use of P2P connections to accommodate the higher demand.

Note that while the P2MP architecture reduces the number of transceivers, its broadcast nature may require more bandwidth compared to P2P DWDM, potentially leading to higher fiber utilization. To assess this, we analyzed the number of light-trees per fiber across different scenarios. Instances where a single fiber's capacity was insufficient were rare, occurring in

fewer than 2% of cases under the most demanding traffic scenarios. These occurrences could be further minimized through rerouting, as the available degrees of each node typically provide alternate paths. Hence, in realistic deployments, it is unlikely that the fiber bandwidth will be exceeded, particularly as most nodes possess multiple degrees (commonly four) that can distribute load and mitigate congestion by routing traffic along alternative paths.

5. CONCLUSION

In this study, we presented mechanisms that determine the optimal placement of coherent optical P2MP transceivers to form light-trees in extended access architectures, facilitating flexible and efficient resource allocation. This proposed architecture facilitates the move of hubs within metro and regional networks. The advantage is better multiplexing gains, but it requires careful consideration of the PLIs to satisfy the QoT of the connections. The flexibility is further enhanced through the use of coherent P2MP transceivers, reducing the total cost and the number of required IP routers. We proposed two algorithms, based on a MILP formulation and a heuristic, to solve the corresponding design problem. Our simulation experiments indicated that the proposed architecture can achieve significant multiplexing benefits, efficiently consolidating the spoke node demands and thus lowering the infrastructure expenditures. This is particularly beneficial for managing peak traffic loads and supporting diverse traffic types and volumes. Moreover, the proposed algorithms are generic and applicable to different P2MP architectures, including S-BVTs. This adaptability makes our architecture particularly suitable for future-proofing metro and regional networks to support emerging technologies such as 5G mobile networks and the Internet of Things (IoT), which will both gain from the efficient and scalable nature of the proposed network design.

Funding. This research has been funded by the EU H2022 ALLEGRO project (grant agreement ID: 101092766) and by HFRI project DetermIneD, under the call "Basic Research Financing (Horizontal support of all Sciences)" under the National Recovery and Resilience Plan "Greece 2.0," funded by the EU.

REFERENCES

1. D. Welch, A. Napoli, J. Bäck, *et al.*, "Digital subcarrier multiplexing: enabling software-configurable optical networks," *J. Lightwave Technol.* **41**, 1175–1191 (2023).
2. J. Bäck, A. Napoli, E. Riccardi, *et al.*, "A filterless design with point-to-multipoint transceivers for cost-effective and challenging metro/regional aggregation topologies," *International Conference on Optical Network Design and Modeling (ONDM)*, Warsaw, Poland, 2022, pp. 1–6.
3. C. Castro, A. Napoli, M. Porrega, *et al.*, "Scalable filterless coherent point-to-multipoint metro network architecture," *J. Opt. Commun. Netw.* **15**, B53–B66 (2023).
4. A. Napoli, Z. Stevkovski, J. Jiménez, *et al.*, "Enabling router bypass and saving cost using point-to-multipoint transceivers for traffic aggregation," in *Optical Fiber Communication Conference (OFC)* (2022), paper W3F.5.
5. C. Castro, J. Sime, T. Eriksson, *et al.*, "Power and spectral savings in metro-aggregation networks exploiting coherent point-to-multipoint transceivers," *Proceedings European Conference on Optical Communications (ECOC)*, Frankfurt, Germany, September 2024.
6. P. Soumplis, K. Christodoulopoulos, P. Kokkinos, *et al.*, "Optimization of router cost and consumption in a point-to-multipoint metro-core architecture," in *Optical Fiber Communication Conference (OFC)* (to be published).
7. P. Pavon-Marino, N. Skorin-Kapov, and A. Napoli, "Tree-determination, routing, and spectrum assignment using point-to-multipoint coherent optics," *J. Opt. Commun. Netw.* **15**, C29–C40 (2023).
8. S. Krannig, R. Mögel, M. Gunkel, *et al.*, "How to design an optimized set of fibre-trees for filterless optical networks-the elegance of a multi-goal evolutionary Pareto optimization versus a deterministic approach," *Photonic Networks; 17. ITG-Symposium; Proceedings of*, Leipzig, Germany, 2016, pp. 1–8.
9. M. Ruiz and L. Velasco, "Performance evaluation of light-tree schemes in flexgrid optical networks," *IEEE Commun. Lett.* **18**, 1731–1734 (2014).
10. A. S. P. Khope, R. Helkey, S. Liu, *et al.*, "A scalable multicast hybrid broadband crossbar wavelength selective switch for datacenters," in *IEEE 11th Annual Computing and Communication Workshop and Conference (CCWC)* (2021), pp. 1585–1587.
11. K. Yiannopoulos, P. Soumplis, K. Christodoulopoulos, *et al.*, "Enabling extended access with light-trees and point-to-multipoint coherent transceivers," *International Conference on Optical Network Design and Modeling (ONDM)*, Madrid, Spain, 2024, pp. 1–3.
12. P. Soumplis, K. Christodoulopoulos, P. Kokkinos, *et al.*, "Techno-economic and feasibility study of point-to-multipoint communications in the metro-core," *14th International Symposium on Communication Systems, Networks and Digital Signal Processing (CSNDSP)*, Rome, Italy, 2024, pp. 575–580.
13. N. Sambo, P. Castoldi, A. D'Errico, *et al.*, "Next generation sliceable bandwidth variable transponders," *IEEE Commun. Mag.* **53**(2), 163–171 (2015).
14. M. Rapisarda, J. A. Hernández, A. Gatto, *et al.*, "All-optical aggregation and distribution of traffic in large metropolitan area networks using multi-Tb/s S-BVTs," *J. Opt. Commun. Netw.* **14**, 316–326 (2022).
15. M. Qiu, Q. Zhuge, M. Chagnon, *et al.*, "Digital subcarrier multiplexing for fiber nonlinearity mitigation in coherent optical communication systems," *Opt. Express* **22**, 18770–18777 (2014).
16. T. A. Strasser and J. L. Wagener, "Wavelength-selective switches for ROADM applications," *IEEE J. Sel. Top. Quantum Electron.* **16**, 1150–1157 (2010).
17. D. S. Johnson, A. Demers, J. D. Ullman, *et al.*, "Worst-case performance bounds for simple one-dimensional packing algorithms," *SIAM J. Sci. Comput.* **3**, 299–325 (1974).
18. IDEALIST Project, "Elastic optical network architecture: reference scenario, cost and planning," Deliverable D1.1 (2013). <http://cordis.europa.eu/docs/projects/cnect/9/317999/080/deliverables/001-D11ElasticOpticalNetworkArchitecture.doc>.
19. D. J. Ives, P. Bayvel, and S. J. Savory, "Physical layer transmitter and routing optimization to maximize the traffic throughput of a nonlinear optical mesh network," *International Conference on Optical Network Design and Modeling (ONDM)*, Stockholm, Sweden, 2014, pp. 168–173.
20. N. Skorin-Kapov, F. J. Moreno Muro, M.-V. Bueno Delgado, *et al.*, "Point-to-multipoint coherent optics for re-thinking the optical transport: case study in 5G optical metro networks," *International Conference on Optical Network Design and Modeling (ONDM)*, Gothenburg, Sweden, 2021, pp. 1–4.
21. P. Poggiolini, "The GN model of non-linear propagation in uncompensated coherent optical systems," *J. Lightwave Technol.* **30**, 3857–3879 (2012).