SEASON

Self-Managed Sustainable High-Capacity Optical Networks

This project is supported by the SNS Joint Undertaken through the European Union's Horizon RIA research and innovation programme under grant agreement No. 101096120

Deliverable D2.3

Techno-economic Analysis

Editor: M. Quagliotti (TIM)

Contributors: INF-P, TIM, ADVA, TID, UC3M, UPC, INF-G, ACC, CTTC, ERI, WEST, SSSA, and CNIT

Version: 2.0

Date: November 12th, 2025

Distribution: PUBLIC (PU)

DISCLAIMER

This document contains information which is proprietary to the SEASON consortium members that is subject to the rights and obligations and to the terms and conditions applicable to the Grant Agreement number 101096120. The action of the SEASON consortium members is funded by the European Commission.

Neither this document nor the information contained herein shall be used, copied, duplicated, reproduced, modified, or communicated by any means to any third party, in whole or in parts, except with prior written consent of the SEASON consortium members. In such case, an acknowledgement of the authors of the document and all applicable portions of the copyright notice must be clearly referenced. In the event of infringement, the consortium members reserve the right to take any legal action it deems appropriate.

This document reflects only the authors' view and does not necessarily reflect the view of the European Commission. Neither the SEASON consortium members as a whole, nor a certain SEASON consortium member warrant that the information contained in this document is suitable for use, nor that the use of the information is accurate or free from risk, and accepts no liability for loss or damage suffered by any person using this information.

The information in this document is provided as is and no guarantee or warranty is given that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability.

REVISION HISTORY

Revision	Date	Responsible	Comment
0.1	May 2nd, 2025	M. Quagliotti (TIM)	Table of Contents
1.0	October 7th, 2025	M. Quagliotti (TIM) and A. Souza (INF-P)	Complete draft
1.1	October 15th, 2025	M. Quagliotti (TIM) and A. Souza (INF-P)	Document ready for internal review
1.2	November 4th, 2025	K. Vlachos (CNIT) and L. Rapp (ADVA)	Document internally reviewed
2.0	November 12th, 2025	C. Pinho (INF-P)	Final version quality checked

LIST OF AUTHORS

Partner	Name Surname
CNIT	Filippo Cugini
CTTC	Laia Nadal, Ramon Casellas
ADVA	Arantxa Villavicencio Paz, Achim Autenrieth
TID	José Manuel Rivas-Moscoso, Antonio Melgar, Óscar González de Dios
UC3M	Farhad Arpanaei, José Alberto Hernández, Fernando Díaz de María
TIM	Marco Quagliotti, Mauro Agus
UPC	Marc Ruiz, Luis Velasco, Jaume Comellas, Salvatore Spadaro
INF-P	André Souza, Cátia Pinho
INF-G	Carlos Castro, Mohammad Hosseini, Antonio Napoli
WEST	Andrea Marotta, Carlo Centofanti, Stefano Tennina
SSSA	Nicola Sambo
ACC	Simon Pryor, Stephen Parker
ERI	Gianluca Gambari, Roberto Magri

EXECUTIVE SUMMARY

The SEASON project aims to redefine sustainable, high-capacity optical network architectures for the cloud era, driven by Europe's joint horizon RIA initiative. Deliverable D2.3 details validated innovations, technical and economic studies, and the achievement level of strategic KPIs related to next-generation optical network scalability, flexibility, economic viability, energy efficiency, and self-management.

The validated reference data plane architecture defined in SEASON project is organized into two macro-domains: Access-Metro and Backbone.

The Access-Metro domain features an open, flat structure that seamlessly connects access points to Edge and Far Edge Central Offices (COs) over distances up to 400 km. The proposed technological solutions investigated and whose effectiveness has been proven in the Access-Metro segment are advanced passive optical network (PON), the use of point-to-multipoint (P2MP) networking that makes use of coherent transmission with Digital Subcarrier Multiplexing (DSCM) and the adoption of cloud-based Radio Access Networks (e.g., complaint with O-RAN forum) with Data Processing Units (DPUs) on the servers. DPUs are cards that offload the server's central processor, realizing and accelerating networking, storage and security and enabling, through pluggable transceiver modules, the optical, packet and compute integration and pursuing optical-electro-optical (O/E/O) conversions reductions.

The Backbone Domain supports multi-terabit data flows via sophisticated optical multiplexing, multi-core fiber bundles, and dynamic network function placement. Over medium (by 2028) and long-term (by 2032), the SEASON network architecture supports a transition from conservative, limited-fiber expansion using C and L bands up to radical new deployments that exploit all available bands (C, L, E, S, O, U) and multicore fibers to meet exponential traffic growth and cloud service demands. Key elements for achieving these objectives are the proposed and analyzed node architectures, the network studies on the backbone segment, where the demand for high capacity is highest, and the migration strategies from C band only to multi-band (MB) and space division multiplexing networks.

Key Studies Performed and Main Results

Scalability and Capacity. Modular prototypes employing MB transmission and space division multiplexing (SDM) demonstrated capacity increases from 6x up to 120x versus conventional C-band systems. Scalable node and transceiver designs for MB over SDM (MBoSDM) were realized, offering up to 3.6 Pbs switching capacity and lower blocking probability.

Energy Efficiency. Dynamic spatial lane aggregation and P2MP innovations in PONs delivered 18%–38% energy savings in lab and field tests. Real-time monitoring further refined consumption, confirming significant sustainable operational benefits.

Techno-economic Impact. Case studies revealed credible 40% CAPEX reduction through DPU-based Open-RAN architectures in mature open hardware ecosystems—guiding strategies for

edge-cloud convergence and mobile upgrades. Cost modeling and scenario analyses matched advanced systems against legacy deployments to ensure optimal, future-proof planning.

Artificial Intelligence & Orchestration. Multi-agent AI systems reduced active subcarriers by up to 50% and projected 25% OPEX savings. Telemetry compression achieved three orders of magnitude data reduction for large-scale operations. Context-aware orchestration models supported dynamic capacity allocations, self-management, and predictive maintenance.

Cost Model. The cost model developed in SEASON provides a structured, transparent methodology. Cost model provides a dataset of cost values for evaluating capital and operational (energy) expenditures of advanced optical networks, covering key equipment categories such as transceivers, amplifiers, switches, and servers across multiple transmission bands and technologies. Notably, SEASON delivers a publicly accessible cost model dataset, enabling community benchmarking and reproducibility for forthcoming techno-economic studies and network planning activities.

Open Datasets and Validation. Publicly released datasets—including IQ data, optical power traces, quality-of-transmission metrics, topology templates and cost model as mentioned just above—established benchmarks for empirical validation and reproducibility, ensuring lasting impact for the community.

Achievement of Key Performance Indicators (KPIs)

The WP2 techno-economic analysis assessed twelve KPIs from the project. A summary of their achievement status is presented below. It should be noted that for some KPIs the work was carried out in cooperation with WP3 and WP4. The complete report on KPI achievement is included in Section 1.2 of this document.

Network Scalability and Sustainability (KPI 1). Achieved through validated MB and SDM architectures, with scalability and power efficiency priority across access and backbone domains.

Bandwidth Expansion (KPI 2.1). Achieved with backbone capacity growth up to 120x, MB/SDM deployment, and robust impairment modeling.

CAPEX Reduction (KPI 2.2, 6.1). Achieved; cost model analyses confirmed 40%–50% reductions are feasible using advanced transceivers, fewer aggregation layers, and Open-RAN architectures.

Modular MBoSDM Nodes (KPI 3.1, 3.2). Achieved; prototypes demonstrated switching/modular transceiver capacities of up to 3.6 Pbs across several bands and fiber cores.

Energy Saving (KPI 4.2). Partially achieved; dynamic aggregation approaches validated energy savings of 18%–38%, with further optimization ongoing.

400G RAN Fronthaul Capacity (KPI 4.3). Achieved; scalable 400G network configurations and amplification schemes validated robustness and throughput.

OEO Conversion Reduction (KPI 6.2). Achieved; architecture reduced required transceivers and OEO conversions by up to 50% in field scenarios bringing technical and economic benefits, including CAPX reductions.

Data compression achieved by intelligent data aggregation (KPI 7.1). Achieved; it has been proved that techniques based on supervised feature extraction, data compression, and data summarization provide data compression ratio >90% without significant information loss.

Replace manual/static operation with autonomous operation based on Multi-agent systems to achieve opex reduction (KPI 8.1). Achieved; a reduction of energy consumption of up to 50% has been demonstrated by applying a dynamic adaptation of the active capacity to traffic needs when point-to-multipoint DSCM-based devices are used in metro-aggregation networks.

Soft-failure detection and localization with digital twin to improve speed and accuracy in soft failure identification (KPI 8.3). Achieved: several methods for soft-failure management have been introduced and its performance was presented testing the achievability of KPI 8.3 requirements.

In conclusion, all KPIs taken into consideration in WP2, some of them in cooperation with WP3 or WP4, can be considered fully satisfied.

Overall Conclusion and Community Relevance

Deliverable D2.3 marks a leap forward in the design and validation of next-generation optical networks. SEASON's reference architectures and modular solutions establish new technical and economic paradigms for sustainable, scalable cloud-era networking. Major breakthroughs in capacity scaling (120x), energy savings (up to 38%), and CAPEX reduction (up to 50%) demonstrate both feasibility and practical value. Integrated AI management, open datasets, and cost models further empower the community to benchmark, reproduce, and advance the state of the art. As European digital infrastructure enters an era of exponentially increasing cloud- and AI-driven demand, SEASON's innovations set the trajectory for agile, efficient, and resilient transport networks—offering industry and academia powerful tools and insights for future development.

TABLE OF CONTENTS

1 INTRODUCTION 10
1.1 SEASON NETWORK ARCHITECTURE FROM ACCESS TO CLOUD 12
1.2 Report on KPIs accomplishment 17
1.2.1 KPI 1 17
1.2.2 KPI 2.1 19
1.2.3 KPI 2.2 20
1.2.4 KPI 3.1 21
1.2.5 KPI 3.2 23
1.2.6 KPI 4.2 24
1.2.7 KPI 4.3 24
1.2.8 KPI 6.1 25
1.2.9 KPI 6.2 26
1.2.10 KPI 7.1 27
1.2.11 KPI 8.1 27
1.2.12 KPI 8.3 28
2 Scalabilty and cost effectiveness of MB and SDM systems and networks 30
2.1 Capacity increase due to MB 30
2.2 Realistic Capacity Scaling of Single-Band and MB Optical Transport Systems 31
2.2.1 Simulation Setup and System Optimization 32
2.2.2 Results and Discussion 33
2.2.3 Conclusions 36
2.3 Capex Savings of Node Architectures for High-Capacity MBoSDM Optical Networks 37
2.3.1 Cost Model 38
2.3.2 Results 39
2.3.3 Conclusion 40
2.4 Techno-Economic Study on the benefits of open coherent pluggables in MB over SDM architectures 40
2.4.1 Methodology 41
2.4.2 Results and discussion 43
2.4.3 Conclusions 45
2.5 Nonlinear-Aware Design of Filterless Horseshoe-and-Spur Optical Networks 45
2.6 Empirical latency model for IPoWDM end-to-end connections 48
2.7 Cost effective strategy for migration to the MB systems 51
2.7.1 MB Optical Systems and Migration strategies 52

2.7.2 Physical Layer Impairment Modeling and Optimization 53
2.7.3 Digital Twin-Enabled PaYG Migration Framework 54
2.7.4 Simulations, results and techno-economic evaluation 56
2.7.5 Summary and conclusions 59
2.8 Reliability based on Hierarchical Aggregation 60
3 Power-efficient and Cost-effective Solutions for Access-Metro segment 63
3.1 Dynamic Spatial Aggregation for Energy-Efficient Passive Optical Networks 63
3.2 Power consumption of DSCM transceivers in the metro-aggregation domain 64
3.3 Comparison of different transport solutions for FH in the RAN in different geotype scenarios and different time horizons 69
3.4 Comparison between WDM P2P and P2MP transceivers within a Full Mesh for FH network 74
3.5 Validation of 400 Gb/s within FH network 75
4 Smart Edge Nodes for Packet/Optical Integration with Computing Resources 80
4.1 Railway mobile service coverage scenario definition for DPU application Use Case .. 80
4.1.1 Scenario Definition and Transport Requirements 81
4.1.2 Evaluated network architectures 82
4.1.3 Simulations and Lab Experiments 84
4.1.4 Techno Economic analysis 84
4.1.5 Counting of transceivers and O/E/O conversions 88
5 Artificial Intelligence (AI)/Machine Learning (ML) Service Orchestration and Self-Management and Secure AI 89
5.1 Coordination of Radio Access and Optical Transport Operation to Reduce Optical Capacity Overprovisioning 89
5.1.1 Context-Aware Autonomous Network Operation 89
5.1.2 Results 95
5.1.3 Conclusions 100
5.2 Dynamic Subcarrier Allocation in P2MP Connections To Reduce Energy Consumption 101
5.2.1 Dynamic P2MP connection management and energy consumption analysis 102
5.2.2 Energy consumption performance evaluation 104
5.2.3 Conclusions 104
5.3 Optical Line system automatic setup (amplifier configuration) 105
6 Cost Model 112
6.1 Commercially Available Devices 114
6.1.1 Transceivers 115
6.1.2 C-band Amplifiers 118
6.1.3 Layer 3 Switches and Routers 118

6.1.4 Compute and Storage 119
6.2 Modeling the cost in function of equipment capacity 120
6.3 An approach to estimate the cost of optical systems for O, E, S and L bands for MB systems 123
7 Datasets for Network Studies and AI/ML training 125
7.1 Datasets 125
7.2 IQ Optical Constellation Data 126
7.2.1 Description 126
7.2.2 Illustrative Figures 127
7.3 Optical Power Data 128
7.3.1 Description 129
7.3.2 Illustrative Figures 130
7.4 Optical Span QoT Data 133
7.4.1 Description 133
7.5 Reference Network topologies 134
7.5.1 Geotypes with example of fiber duct tree topologies 135
7.5.2 Metro-Regional topology 136
7.5.3 Combining Access with metro-regional topology 137
7.5.4 Traffic generation for SEASON Refence Network 138
7.5.5 Generating realistic optical topologies for techno-economic studies 140
7.6 Packet Trace Data 141
7.6.1 Description 141
7.6.2 Illustrative Figures 142
7.7 Control Plane Data 142
7.7.1 Architecture and topology 142
7.7.2 Requests and Notifications 143
7.8 Cost Model Dataset 144
7.9 Generative adversarial networks for synthetic traffic generation 145
7.9.1 Methodology 145
7.9.2 Results 146
8 CONCLUSIONS 148
GLOSSARY 149
REFERENCES 152
APPENDIX - Relevant publications 156

Dissemination Level	PUB (Public)

1 INTRODUCTION

The SEASON project aims to design and validate sustainable, ultra-high-capacity optical networks, combining multi-band (MB) and space division multiplexing (SDM) innovations to support services beyond 5G and 6G services. Its main objectives include scaling up capacity, improving energy efficiency, enabling self-managed and autonomous network operations, and supporting new user services and verticals, all while fostering open standards, technological innovation, and European leadership in the field.

Within SEASON, WP2 handles the use cases, network architecture, and techno-economic dimension of the project, acting as a foundation for technical and experimental work. Specifically, WP2 is tasked to:

- Identify future use cases and translate them into system-level requirements and KPIs for the project.
- Define reference SEASON network architectures and topologies for technical development and evaluation.
- Establish detailed high-level technical and functional specifications spanning the Metro-Access, and Metro/Long-Haul (Backbone) domains.
- Provide guidance for the design, planning, and assessment activities in downstream work packages (WP3, WP4, WP5).
- Deliver techno-economic analyses of solutions and subsystems to validate sustainability and business impact.
- Compile and specify required data sets for technical validation and benchmarking of targeted use cases.

WP2 is critical in setting the overall direction, architectural framework, and evaluative baseline for the SEASON project, ensuring that proposed innovations are both technically sound and economically viable in future telecom environments.

The first three points listed above are addressed in deliverable D2.1 [SeaD2.1] delivered in month M14. The fourth bullet is achieved from constant contact, during the plenaries and in additional focused meetings, with the other work packages, especially with WP3 (systems for data plane) and WP4 (architecture and software for management and control planes).

The last two points are addressed in the two deliverables dedicated to results of techno-economic studies, namely D2.2 [SeaD2.2], delivered in month M25, reporting preliminary results, and this one, D2.3, that reports the final results.

This section includes two introductory subsections and a summary of the other chapters of the document. The first subsection (subsection 1.1) describes the final status of the work on the

SEASON data plane network architecture that has been defined in D2.1 and the one on control plane developed in WP4, highlighting the studies and activities which have developed enabling technologies and prototypes, evaluating their technical and economic viability and prove their practical operation through lab experiments and demonstrations. The second subsection of this chapter (subsection 1.2) summarizes the level of achievement of the KPIs considered in WP2, namely those related to architectural and techno-economic aspects, continuing and completing the work reported in deliverable D2.2.

Before presenting the overview of the results that demonstrated the feasibility and benefits of the architecture proposed by SEASON and the related supporting report on the status of KPIs achievement, the content of the following chapters is briefly presented. They include the studies and activities conducted in the third year of the project by WP2 which continue and complete the work reported in the previous deliverables D2.1 and D2.2.

In short, the subsequent chapters of the document are organized as follows.

Chapter 2 develops the studies associated with the exploitation of MB and SDM, with the possible combination of the two, MB over SDM (MBoSDM). Studies involve scalability aspects as well as technical viability of system and networks and migration strategies from single band to multi-band networks.

Chapter 3 focuses specifically on solutions for enabling the development of a high-capacity, efficient, and economically sustainable Access-Metro segment.

Chapter 4 addresses a specific solution designed to improve the performance and costs of edge nodes, namely the use of data processing unit (DPU) cards to integrate compute and networking functions, compact and virtualize the network equipment eliminating unnecessary transport and switching equipment.

Chapter 5 focuses on artificial intelligence solutions for orchestration and management, assessing their potential to increase efficiency and reduce costs, particularly operational costs.

Chapter 6 reports the SEASON project cost model, an attempt to provide reference cost values for network techno-economic studies, not only within the project, but also by defining and releasing a public dataset with the costs of equipment and systems to be applied for technical and economic network assessments.

Finally, Chapter 7 collects all the datasets used by the project in various simulation and experimental studies, which were made public through uploading to unrestricted, accessible storage spaces.

The document concludes with Chapter 8 which summarizes the work developed on this final deliverable of WP2.

1.1 SEASON NETWORK ARCHITECTURE FROM ACCESS TO CLOUD

The SEASON project envisions two distinct time periods for network evolution: medium-term (targeting around 2028) and long-term (targeting around 2032 and beyond), each with fundamentally different requirements and technological constraints. While both periods address traffic growth and network modernization, they differ in their architectural complexity and technology adoption. The medium-term architecture (Fig. 2-3 of [SeaD2.1]) is characterized by conservative constraints, including limited new optical fiber deployment, use of only C+L bands (possibly one additional band), and traffic growth by a maximum factor of 5x compared to current levels. In contrast, the long-term architecture (Fig. 2-4 of [SeaD2.1], also reported in D2.3 Figure 1.1-1) presents a more aggressive evolution with traffic demands increased by a factor of 10-30, no restrictions on new fiber deployment including multicore fibers, and exploitation of all optical bands (C, L, E, S, O, and U).

The SEASON architecture defines *three types of Central Offices* (COs) that play crucial roles in network function distribution. The **Far Edge CO** based on small structures (containers or reinforced cabinets) with space and power constraints, hosting limited computing capacity for Telco and Service functions. The **Edge CO,** which consists of medium-sized structures hosting extensive Telco applications and IT functions, typically providing upgrades of existing metro-regional Central Office. The **Cloud CO,** which implements large backbone structures hosting comprehensive Telco and IT applications with hundreds to thousands of servers.

About Network Segmentation and Domain Architecture, both short-term and long-term architectures establish a two-segment network structure composed of an **Access-Metro Domain** and a **Backbone Domain**.

The Access-Metro Domain is a flat, hierarchically unbounded segment extending from Access Points to Edge COs, spanning 100-400 km. This domain deliberately blurs traditional access/metro boundaries to optimize end-to-end connectivity. The Access-Metro domain is partitioned into four subsegments.

Fixed Access: Complete FTTH deployment with emerging high-capacity passive optical network (PON) technologies (200G+ coherent PON, SDM PON)

Mobile Access: Dense 6G deployment with macro and small cells, supporting advanced geographical areas

Metro Far Edge: Distributed processing points with limited computing capacity

Metro Edge: Full-service nodes with extensive computing and networking capabilities

The Backbone Domain is the traditional long-haul segment for inter-regional transport, characterized by multi-Terabit flows and advanced optical technologies.

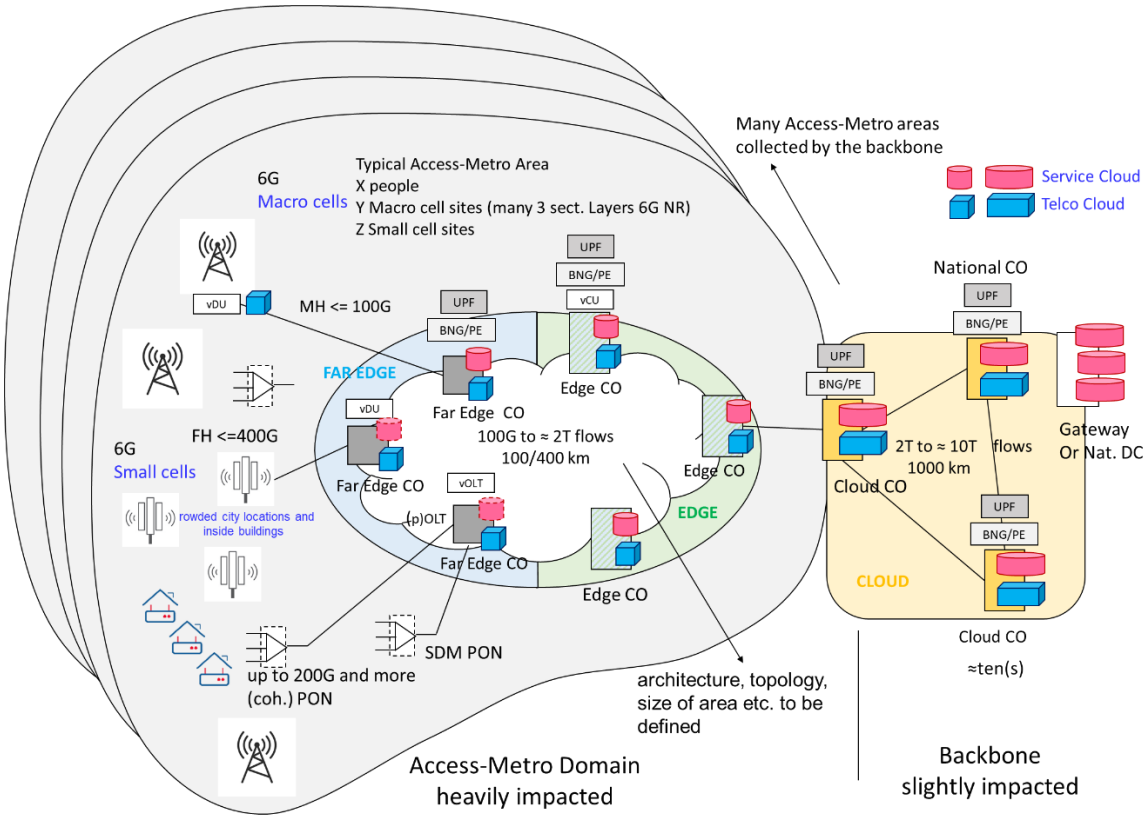

Figure 1.1-1 SEASON data plane architecture for long-term as defined in WP2 [SeaD2.1].

The network sectioning follows a clear hierarchy where access, metro Far Edge, and metro Edge components collectively form the Access-Metro segment, while Cloud COs and backbone infrastructure constitute the backbone segment. This creates a streamlined two-domain architecture that eliminates traditional hierarchical boundaries within the Access-Metro domain while maintaining clear separation from backbone operations. The long-term vision anticipates extended decentralization of Telco and Service functions, positioning network capabilities where needed rather than maintaining centralized cloud-only approaches. Traffic flows in the Access-Metro domain range from 100 Gbps to 1.6-3.2 Tbps depending on the specific implementation timeframe, while backbone segments support multi-Terabit interconnections up to 10 Tbps.

This architectural evolution represents a fundamental shift toward flatter, more flexible network designs that can accommodate the dynamic and unpredictable traffic patterns expected in the 6G era while leveraging advanced optical technologies including full SDM implementation and comprehensive MB operations.

With reference to the reference architecture for the long term presented in D2.1, Figure 1.1-2 shows the studies performed in WP2 and their connection to the data plane architecture, which, as mentioned above, is structured across two macro segments: Access-Metro and Backbone. The studies were conceived to demonstrate whether individual or combined aspects of architecture are feasible and are economically viable and energy sustainable. They have been

grouped into four categories using different colors for their representation. The works that introduce the concept of MB and SDM and address various aspects such as scalability and the economic viability of MB systems (nodes and transceivers) are highlighted in yellow (top in the Figure 1.1-2). The several studies conducted on digital subcarrier multiplexing (DSCM) devices enabling the point-to-multipoint (P2MP) networking used in the access-metro segments are highlighted in light green (under the representation of architecture). The numerous studies specifically dedicated to solutions for both fixed and wireless access segments are highlighted in blue (on the left). Such studies involve both the data plane and resource management in charge of the control plane. Finally, the cross-sectional, end-to-end study on latency estimation is highlighted in white. This is a very important aspect for ensuring service delivery in compliance with assigned quality of service levels.

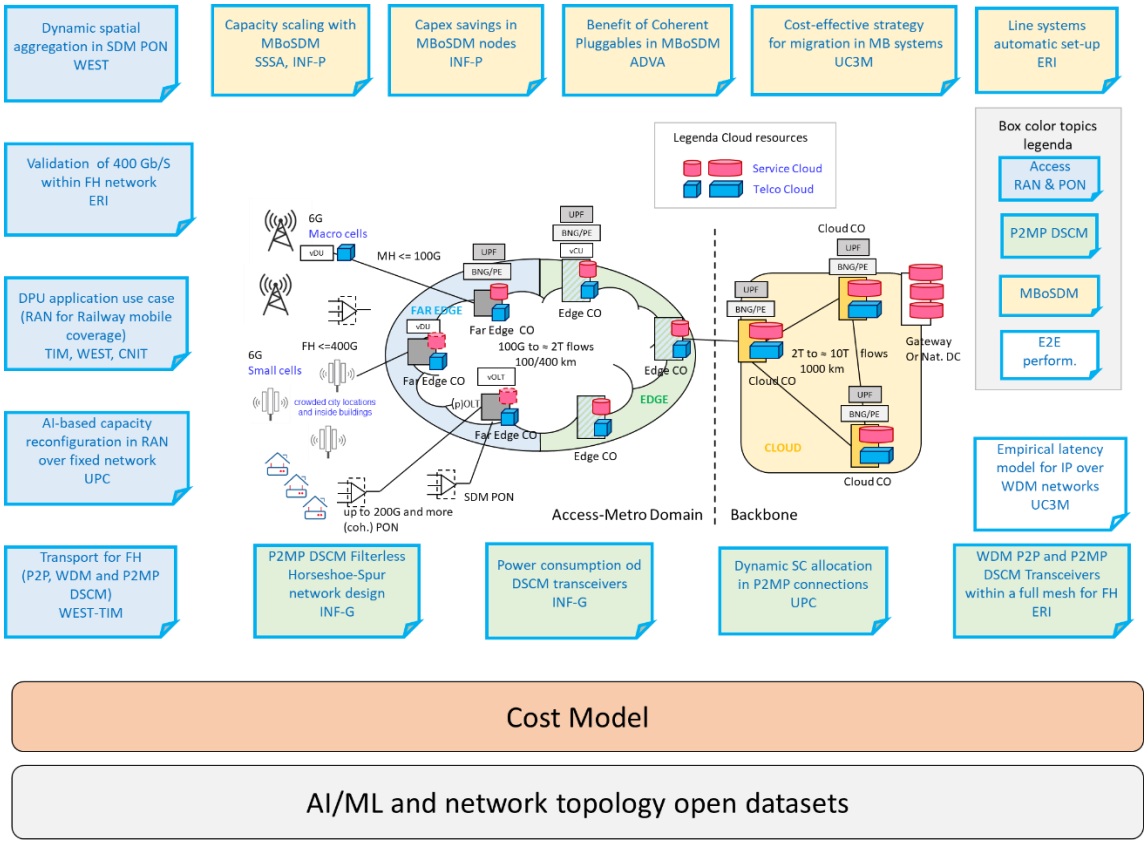

Figure 1.1-2 Techno-economic studies performed and main general outputs of WP2 (Cost Model and Datasets) related to SEASON data plane architecture.

In addition to the various monographic studies mentioned above and presented in dedicated subsections in D2.2 and in the following chapters of this document, two important general outputs of WP2 is included in the bottom part of Figure 1.1-2: the Cost model (in light red) and the Datasets (in grey). Chapters 6 and 7 are dedicated to these two key results respectively, and we refer to them for details.

To complete the overview of the work conducted in parallel with the other SEASON project WPs, we present a summary of their activities and results. These complementary results reinforce WP2's outcomes and collectively demonstrate the feasibility and sustainability of the SEASON network architecture.

Figure 1.1-3 represents a map of the data plane prototypes (systems and subsystems, as well as algorithms for their control, monitoring, and predictive maintenance) developed in WP3, illustrating their relationship to the data plane's reference architecture. The report on the final version of WP3 prototypes can be found in [SeaD3.3].

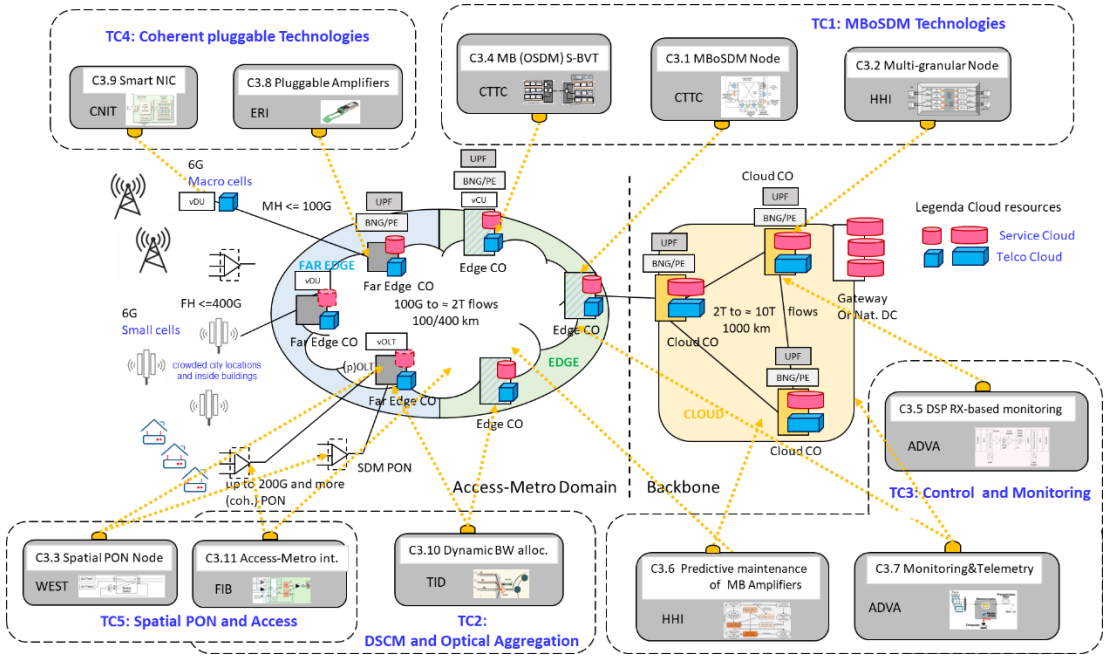

Figure 1.1-3 Data plane architecture updated with data plane prototypes.

Figure 1.1-4 (Fig 2-2 in [SeaD4.3]), on the other hand, represents a schematic representation of the software (SW) blocks developed in WP4 in relation to the main demos with reference to the data plane architecture. As can be seen, the management plane developments assure full orchestration from access (mobile and fixed) to backbone relying on specific components for telemetry and control in the different segments of the network. Details and explanation about this block diagram and its components can be found in [SeaD4.3].

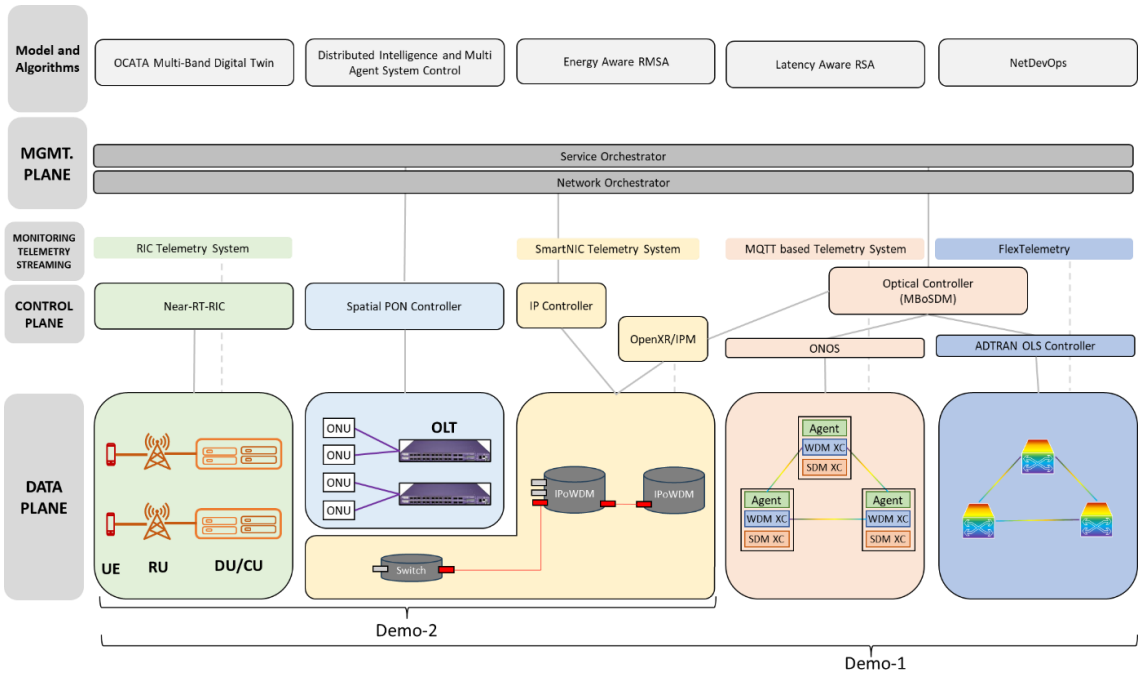

Figure 1.1-4 Final Control plane architecture and components developed in WP4 and related to the two main demos of the project: Demo-1, operator perspective (D1-Berlin demo) and Demo-2, user perspective (D2-L'Aquila demo) [SeaD4.3].

Finally, Figure 1.1-5 schematically represents the two main demos (D1 Berlin - Operator perspective, D2 L'Aquila - User perspective) and the two additional demos (D3 Telefonica – Madrid, and D4 Fibercop - Turin) in relation to the data plane architecture. The arrows indicate the parts of the network (network segments, access systems, Central Offices/node, links) to which the experiments are connected. All these demos have the ambition to prove that the SEASON solutions implemented on part of the network architecture are experimentally demonstrable (D1-Berlin, representing the Operator perspective, offers the widest coverage of network segments, excluding only the access terminals). Demos planning has been described in [SeaD5.1] while the final reports on demos will be reported in D5.2 (D1-Berlin) and D5.3 (D2-L'Aquila) due in M36 of the project and not available at the date of D2.3 delivery.

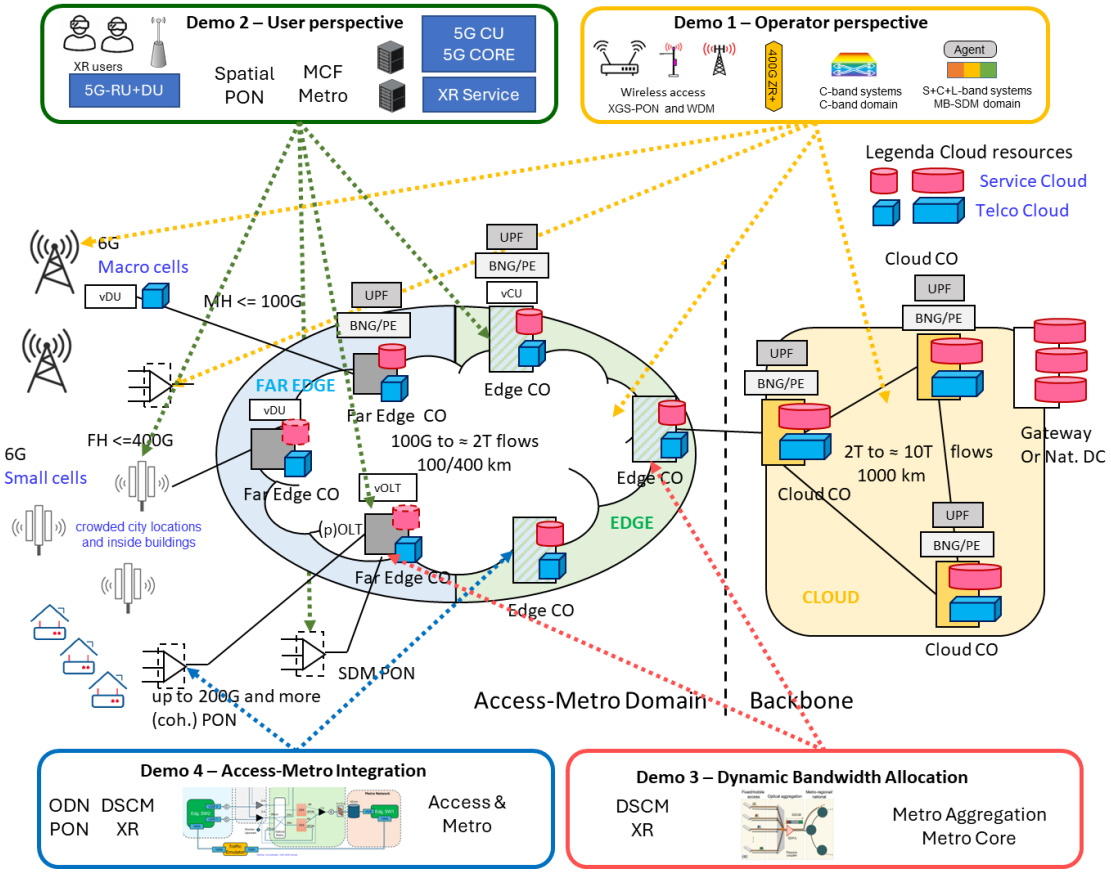

Figure 1.1-5 SEASON data plane architecture updated with project demos.

1.2 REPORT ON KPIS ACCOMPLISHMENT

This subsection reports on KPIs considered in WP2 and their degree of achievement. For each KPI, the declaration (as defined in the project proposal), the classification given in relation to the technical-economic aspects as specified in D2.2 (*mandatory* or *additional*), the degree of achievement at the end of the WP2 work period (M32 of the project) and a brief description of the activities carried out in relation to the KPI are provided.

The statements in the following subsections update what has been reported in chapter 1 of D2.2 where the status of achievement of the KPIs and the related references to the results obtained at the end of the second year of the project (M24 of the project) were documented.

1.2.1 KPI 1

KPI 1: "Design a sustainable network architecture able to scale up network capacity and cope with challenging user needs through improved power efficiency, reliability, and self-management capabilities."

This KPI is classified as ***mandatory***.

Status: Achieved.

Dissemination Level	PUB (Public)

Significant advancements toward a sustainable, scalable, power-efficient, reliable, and self-managed network architecture have been developed and validated within the SEASON project's and documented in final project deliverables D2.3, D3.3, D4.3, D5.2 and D5.3 (last two available in M36, at the end of the project). The main features and results (excluding only final achievements of WP5 which are still ongoing) on specific aspects of KPI 1 statement and proving its fulfillment are summarized below.

Network architecture. An architecture for the SEASON data plane was defined in D2.1 and then used as a reference throughout the project by all WPs. A detailed summary of this result is provided in Section 1.1 of this document. The summary highlights various aspects covered by the individual WPs, across the various network segments and network planes (data plane, control/telemetry, management/orchestration). Some reference networks (provided with topologies and traffic inspired by real operator networks) have been made available by the operators. A reference network covering the Access-Metro segment has also been provided as a public data set (Section 7.5).

Scalability and Capacity. The architecture adopts MB transmission (MBT) and SDM, proven in simulation and prototyping to increase network capacity by up to 6–120 times over conventional C-band systems, with further scalability possible by parallelizing multiple fibers or multi-core bundles. Prototypes (e.g., MBoSDM switching nodes, sliceable bandwidth transceivers) have been validated for their ability to manage multiple bands and spatial channels, demonstrating modular, scalable architectures applicable from access to backbone, aligning with medium- and long-term evolutions outlined in the project. Scalability, migration and techno-economic studies are performed in WP2 and reported in sections 2.1, 2.2, 2.7 of D2.3 and data plane studies and prototypes can be found in [SeaD3.3].

Power Efficiency. Dynamic aggregation and optical P2MP transmission reduce interface and router requirements, resulting in clear operational energy savings: experimental and simulation data report average savings from 18% up to 38% in access-metro PONs, with even higher peaks under low load (section 3.1). Coherent pluggable transceivers and SDM technology further shrink power consumption per bit, especially in edge and aggregation nodes, as reported in sections 3.2 and 5.2.

Reliability. Robust optimization models and advanced monitoring (DSP Rx, predictive maintenance, machine learning (ML)-based fault detection) have been developed and tested in several labs, increasing the reliability of filterless and MB optical segments. The architecture includes provisions for automated amplifier gain control, traffic-aware reconfiguration, and real-time telemetry, elevating the system's operational resilience and fault localization capabilities. The results are reported in deliverable D3.3. Reliability in optical networks planning and design is also extensively discussed in section 2.8 of this deliverable where a traffic engineering modelling approach for multi-homed, multi-layer hierarchical metropolitan area networks (MANs) are presented.

Self-Management. Integration with programmable control planes and self-management frameworks (AI/ML-based orchestration, closed-loop monitoring) is well advanced, with numerous prototypes and algorithms explored, through full demonstration in final deliverables. WP4 and WP5 interfaces and demo settings are in place, enabling future automation and streamlining network operations, in line with KPI requirements. In this deliverable, significant results are reported in the three subsections of section 7.

Conclusion. The architectural solutions, designs, and prototype validations together demonstrate a strong trajectory toward fulfilling KPI 1. Substantial evidence is provided for scalability, power efficiency, reliability, and progress toward self-management in the network. Complete achievement and integration, especially for self-management and full reliability, are projected for the final versions of the demos already done or in progress in WP5.

1.2.2 KPI 2.1

KPI 2.1: "Increase the available bandwidth of the fiber from actual C-band (~35 nm) to O, E, S, L, U bands (~415 nm) that, together with the usage of SDM, e.g., with >10 fibers / cores, will make the available bandwidth to grow by a factor x120. [SRIA, mid-term evo ~2028]."

This KPI is classified as ***additional***.

Status: Achieved.

Investigations carried out in the framework of the SEASON project and presented in Section 2.1 found that a x120 of capacity increase in a sample backbone (a Spanish backbone) with respect to the actual C-band can be achieved with a combination of MB and SDM as follows: 20 fibers (or cores) exploiting L, C, S, and E band (24 THz). Such a result accounts for wideband transmission impairments like stimulated Raman Scattering.

The analysis presented in Section 2.2 shows that the number of parallel fibers required to carry a certain traffic in a point-to-point link can be reduced by more than 50% by using super bands and deploying L, S and O bands in comparison to a C-band-only system, especially for shorter links. For example, for a 75 km link, the number of required fibers drops by 20%, 60%, 73%, and 80% for the Super-C, Super-C+L, Super-S+C+L, and Super-S+C+L+O systems, respectively. At 1425 km, the reductions are more modest for the MB systems: 20% for Super-C and 45% for the other MB systems. The choice of deploying the O instead of the E band is intended to reduce the stimulated Raman scattering effect between transmission channels and increase the robustness of the system against failures.

The choice of potential transmission bands used in the previous analyses are based on performance estimations presented by WP3 in deliverables D3.2 and D3.3. The techno-economics of using of the U band was not evaluated due to difficulties in optical amplification and the pronounced hindering effects of bending losses and inter-core cross talk as shown in D3.2.

1.2.3 KPI 2.2

KPI 2.2: *"50% CAPEX reduction by (1) designing an architecture that jointly leverages on parallel fibers (where fiber resources are abundant), multiple bands (where fiber resources are scarce), and multi-core fibers (where fibers are not present, e.g., for cell densification); (2) limiting intermediate aggregation in routers thanks to the ultra-high capacity of MBoSDM and by exploiting smart coherent pluggable to remove aggregation layers and unnecessary O/E/O conversions."*

This KPI is classified as **mandatory**.

Status: Achieved.

This deliverable presents techno-economic evaluations of several technologies developed to reduce the capital expenditure (CAPEX) of future MBoSDM transmission systems and are therefore related to KPI 2.2.

Several preliminary results were presented in deliverable D2.2 [SeaD2.2] regarding the complexity and number of components of optical transmission systems and serve as basis for a more complete CAPEX analysis in this deliverable, together with the definition of a cost model. Section 2.3 extends the preliminary results showing the benefits of fiber core switching by including cost to the analysis of three node architectures presented in deliverables 2.2 [SeaD2.2] and 3.2 [SeaD3.2]. Using the SEASON cost model, supplemented with additional component cost estimates derived through a similar methodology, the study analyzed CAPEX trends for various proposed node architectures and a reference node as a function of the number of directly connected fiber/core pairs (N_c). The results indicate that CAPEX can be reduced by more than 50%, depending on the value of N_c.

The analysis presented in Section 2.2 shows that the number of parallel fibers required to carry a certain traffic between nodes can be reduced by more than 50% by using super bands and deploying the L, S and O bands in comparison to a C-band only system, especially for shorter links. For example, for a 75 km link, the number of required fibers drops by 20%, 60%, 73%, and 80% for the Super-C, Super-C+L, Super-S+C+L, and Super-S+C+L+O systems, respectively.

In the context of mobile access networks, the study included in section 4.1 demonstrates how a solution based on a virtualized radio access network (RAN) architecture (O-RAN) using servers equipped with DPU cards to support compute and networking functions eliminates the need for switching equipment (routers) and reduces the number of required transceivers, effectively halving optical-electro-optical (O/E/O) conversions. The study demonstrated that, under the assumption of a mature O-RAN ecosystem, overall Capex reductions of 40% can be achieved, with the specification that all RAN costs (including software) are considered in the above-mentioned analysis. The level of cost reduction is not exactly the 50% required by the KPI 2.2, but is very close to that target.

With respect to optical core networks, section 2.4 presents a study that evaluates different evolution strategies including MB and multi-core technologies under coherent pluggables and transponder clients. Techno-economic results across three core topologies highlight the impact

of fiber leasing and MBoSDM architectures cost advantage. The use of MBoSDM with core switching provides the best long-term scalability, while solutions employing single-mode fiber bundles are attractive short-term but penalized by fiber leasing. When considering fiber bundles with transponders or coherent pluggables as baselines, MBoSDM achieves maximum cost reductions of up to 15% and 47%, respectively.

1.2.4 KPI 3.1

KPI 3.1:*"Design and implement flexible and modular MBoSDM node prototypes able to switch/add/drop channels in at least 3 different bands (e.g., S, C, L) in an SDM/MCF fiber infrastructure featuring up to 10 fibers/cores, able to cope with switching capacities scalable up to between 2.4-3.6 Pb/s (considering a 4-degree node with 50% local add/drop and depending on the number of used bands and SDM cores/fibers) [SRIA, mid-term evo ~2028], by approaching (fractional) space-wavelength flexible architectures."*

This KPI is classified as **additional**.

Status: Achieved.

The design of several architectures for a flexible and modular MBoSDM node was presented in section 4.5.3 of [SeaD2.1] (qualitative description of three node architectures and comparison with the conventional node architecture), section 4.3.1 of [SeaD3.1] (numerical analysis of the number of components and port count for the architectures presented in [SeaD2.1]), and section 3.2 of [SeaD2.2] and section 4.2 of [SeaD3.2], where a more polished design was presented of three architectures, including formulas for the number of components and port count, and a study of the number of WSS and optical amplifiers for an increasing number of directly switched core pairs (N_C) and different numbers of in/out core pairs (also published in [Sou24-2]). An implementation of this MBoSDM node was described in section 7.1 of D3.2, which is being used in the WP5 activities.

In order to evaluate the switching capacities of the proposed MBoSDM node designs, let us consider the reference backbone network topology proposed in section 5.2.2 of [SeaD2.1], and calculate how the traffic distributes across the network for different capacities at the HL4 nodes, assuming that the traffic is not distributed uniformly across all links —which would not be realistic—, but that 90% of the traffic propagates along 40% of the links. We also consider lightpaths that, on average, pass through 1-8 links. With those assumptions, we observe in Table 1.2-1 that, in order to achieve switching capacities of 2.4-3.6 Pb/s, a plausible option would be to consider that the average HL4 node capacity should reach 20-30 Tb/s and the lightpaths should traverse 6 links on average.

A switching capacity of 2.4-3.6 Pb/s translates, for an MBoSDM node with 4 degrees and 10 fibers/cores, into a capacity per fiber/core of 60-90 Tb/s. Assuming 800G channels with a spectral occupancy of 150 GHz, the number of channels per fiber/core would be 75-113, which can be accommodated in the super-C and super-L bands (lower range) or with the addition of part of the S band (upper range). An add-drop (A/D) module capable of adding/dropping 50% of the node capacity would need to process between 1500-2250 800G channels. Of course, if such

amounts of traffic need to be added/dropped per node, new MB/SDM transceivers must be developed in order to reduce the required number of ports, but those developments fall outside the scope of the project. An alternative to reduce the number of A/D ports would be to consider that just one of the fibers/cores per degree is processed at the channel level (see Table 4-4 in [SeaD3.1]. In that case, the maximum number of channels that could be added/dropped would be 300-450 (i.e., 240-360 Tbps, or 10% of the switching capacity), assuming colorless-directionless-contentionless (CDC) functionality. For colorless-directionless (CD) architectures, it would decrease to 2.5% (but it could increase up to tenfold by allowing processing of the contents of more fibers/cores, as needed). Note that in current network deployments, 4-degree ROADM nodes with CD functionality are usually equipped to add/drop up to 25% of the node switching capacity.

Table 1.2-1 National network link and node traffic assuming 90% of HL4 traffic is distributed among 40% links. For HL4 nodes with an average 20-30Tb/s capacity, switching capacities in the order of 2-3 Pb/s could be achieved for 6 links/lightpaths.

	Number of links/lightpaths	Number of Links	Number of Nodes	Number of Paths	Capacity (Tb/s)			
HL4 nodes			694		10	20	30	40
Links HL3-HL2-HL1	**1**	190		190.0	82.18	164.37	246.55	328.74
Links HL3-HL2-HL1	**4**	190		47.5	328.74	657.47	986.21	1314.95
Links HL3-HL2-HL1	**5**	190		38.0	410.92	821.84	1232.76	1643.68
Links HL3-HL2-HL1	**6**	190		31.7	493.11	986.21	1479.32	1972.42
Links HL3-HL2-HL1	**7**	190		27.1	575.29	1150.58	1725.87	2301.16
Links HL3-HL2-HL1	**8**	190		23.8	657.47	1314.95	1972.42	2629.89
Nodes HL3-HL2-HL1	**1**		98		318.67	637.35	956.02	1274.69
Nodes HL3-HL2-HL1	**4**		98		796.68	1593.37	2390.05	3186.73
Nodes HL3-HL2-HL1	**5**		98		956.02	1912.04	2868.06	3824.08
Nodes HL3-HL2-HL1	**6**		98		1115.36	**2230.71**	**3346.07**	4461.43
Nodes HL3-HL2-HL1	**7**		98		1274.69	2549.39	3824.08	5098.78
Nodes HL3-HL2-HL1	**8**		98		1434.03	2868.06	4302.09	5736.12

Finally, let us have a look at the number of components (and their port count) of the MBoSDM node for 4 degrees, 10 fibers/cores per degree and 3 bands per fiber/core. A conventional CD route-and-select reconfigurable optical add-drop multiplexer (ROADM) architecture, as depicted in Figure 4-14 of [SeaD3.1], would require **120** dual wavelength selective switches (WSSs) for the line connections and a minimum of **12** 20-port dual WSS for the A/D (or **6** 64-port dual WSS). Of course, in order to add/drop 1500-2250 800G channels (50% of the switching capacity), many more WSS would need to be added (112 20-port dual WSS or 35 64-port dual WSS). For SEASON MBoSDM node, let us consider design #3 in section 4.3.1 of [SeaD3.1] and architecture #1 in section 3.2 of [SeaD2.2]. They would require **10** and **3** spatial-optical cross connects (S-OXCs) with **28** and **80** (considering N_C= 0) ports, respectively; and **12** and **3** NxN WSS, respectively. By adding **120** 1x20 WSS, design #3 of [SeaD3.1] can add/drop over 50% of the switching capacity of the node (2400 800G channels), but the architecture would become

colourless-directionless (not contentionless). The number of NxN WSS for design #3 of [SeaD3.1] can be reduced to 3 if, as mentioned above, we only process the contents of 1 fiber/core per degree.

1.2.5 KPI 3.2

KPI 3.2 *"MBoSDM transceivers able to increase the capacity of SoA transceivers up to 2× - 4× by exploiting enhanced wavelength/space dimensions while enabling appropriate slice/band/core/fiber selection according to the network path."*

This KPI is classified as **additional**.

Status: Achieved.

The required capacity scaling up to 2×–4× can be achieved by expanding transmission beyond the C-band and introducing SDM technology to exploit additional spatial paths. In this regard, the programmable MBoSDM sliceable bandwidth/bitrate-variable transceiver prototype (S-BVT), proposed in SEASON and implemented in WP3, enables dynamic bandwidth allocation for flexible, high-capacity, and resilient operation within the S+C+L-bands. This technology has been experimentally validated across the S, C, and L bands and in both single-core and multi-core fiber (MCF) scenarios (see deliverables [SeaD3.2] and [SeaD3.3]). This approach provides scalability, improved network utilization, and optimal slice/band/core selection according to network demands. The transceiver can be reconfigured via software-defined networking (SDN) agents, which manage and adapt programmable elements and reconfigurable devices to meet dynamic traffic requirements. This design aligns with SEASON's KPIs by enabling adaptable use of spectral and spatial resources, directly supporting targeted capacity improvements. Section 4.1 of SEASON deliverable [SeaD3.3] indicate that utilizing the full S+C+L band spectra over one core of a 19-core MCF, an aggregated capacity of 28.2 Tb/s can be achieved, considering 350 single-sideband (SSB) channels in the S-band, 175 SSB channels in the C-band, and 150 double-sideband (DSB) channels in the L-band. By comparison, a state-of-the-art C-band S-BVT (175 SSB channels) can achieve about 7 Tb/s. Therefore, adopting MB transmission (S+C+L) offers a theoretical 4× capacity increase, which further scale with the number of fiber cores in an MCF system. However, operating across multiple bands and cores introduces additional impairments such as stimulated Raman scattering (SRS) and crosstalk (XT), which can impact the overall achievable capacity/scaling factor.

On the other hand, for operators transitioning from C-band to MB (C+L+S) and MBoSDM optical systems, the most cost-effective approach is a phased, "Pay-as-You-Grow" (PaYG) strategy. This approach involves demand-driven, incremental upgrades to add new bands (L, S) and spatial channels (SDM) on selected links, which aligns with actual traffic growth. This minimizes disruption and allows for flexible network scaling by matching investments with revenues. Subsection 2.7 shows that PaYG with aggressive reconfiguration (PaYG-WAR) maximizes throughput and energy efficiency while minimizing capital CAPEX by up to 70% compared to a

"Day One" (DO) deployment. These benefits are achieved with only moderate reconfiguration overhead due to SDN/NFV automation.

Key metrics show significant progress: a fully achieved 120x backbone capacity increase (using four bands and 20 fibers/cores), notable preliminary progress toward a 50% CAPEX reduction by optimizing network layers and node design, and measured energy savings in the 20–38% range for select scenarios, with ongoing research toward meeting a 50% savings target.

1.2.6 KPI 4.2

KPI 4.2: *">50% contribution in energy saving via dynamic spatial channels aggregation and deactivation of unused transceivers at the OLT side basing on traffic conditions over total 70% energy saving targeted by [SRIA]."*

This KPI is classified as ***mandatory***.

Status: Achieved.

The dynamic spatial aggregation architecture developed in SEASON effectively achieves the targeted energy efficiency improvements by dynamically activating and deactivating spatial lanes in response to real-time traffic demands. As discussed in D2.2, simulation results demonstrate an average potential energy saving of 38%, with peak savings up to 87.5%, depending on the ratio of fixed to variable power consumption and the configuration of spatial lanes. Experimental validation on a commercial PON system, characterized by a relatively high baseline power, confirmed measurable energy savings of around 4%. The lower experimental figure is primarily due to the fact that the commercial hardware used was not originally designed to support the activation/deactivation mechanisms explored within SEASON, thus limiting the achievable power variation at the transceiver. When considering the joint activation and deactivation of radio units (RUs) and distributed units (DUs) in the mobile network—together with the optical access domain as described in Section 3.1—the overall saving increases to approximately 8%. Further optimization of radio equipment operation could raise the total system-level energy saving up to 19%. These results confirm the effectiveness and scalability of the proposed architecture for next-generation energy-aware optical and radio access networks.

1.2.7 KPI 4.3

KPI 4.3: "400Gb/s RAN fronthaul ports capacity."

This KPI is classified as ***additional***.

Status: Achieved.

The described fronthaul network architecture builds upon the foundations outlined in Section 3.3 and Section 3.4, interconnecting vDUs and remote radio units (RRUs) through packet switches and an optical line system (OLS) equipped with 400G QSFP-DD transceivers. This setup forms a bidirectional fiber ring and logical mesh topology designed to balance high throughput

with flexible wavelength management. The OLS design adopts a filterless approach, relying on optical add-drop multiplexers (OADMs) with simple splitting and pluggable erbium-doped fiber amplifiers (EDFAs). Two configurations are considered: one employing a single 2:2 coupler for add/drop, and another with separate splitters for the add and drop paths. The latter configuration improves signal quality by mitigating attenuation penalties. Amplification strategies include automatic gain and power control, with both single- and dual-amplifier modules available to deliver different output powers. The coherent 400G transceivers used in the system come in low- and high-power variants, with the latter significantly enhancing optical signal-to-noise ratio (OSNR) and link budgets. The drop-and-continue ring topology requires N(N–1)/2 wavelengths for N nodes, meaning that for practical 400G operation the network supports up to six nodes, while 200G transceivers extend this scalability to about ten nodes. The deployment of high-power transceivers and stronger amplifiers further improves performance and reach. Future enhancements under evaluation include variable gain amplifiers and gain flattening filters, which would enable dynamic control and improved stability across optical paths. Overall, the architecture provides a flexible and scalable solution that ensures robust signal quality, while supporting multi-node optical rings for next-generation fronthaul transport.

1.2.8 KPI 6.1

KPI 6.1: *"40% CAPEX reduction by collapsing computing, IP networking, and usage of high-speed coherent optical transmission in a single element (i.e., DPU) not designed for the Telecom market but for much wider computing markets and verticals (e.g., automotive)."*

This KPI is classified as ***mandatory***.

Status: Achieved.

This KPI is fully achieved under specific conditions related to the evolution of the Open-RAN ecosystem. The techno-economic analysis, reported in subsection 4.1 and focusing on a 5G/6G railway coverage use case, compares three alternatives: two conventional RAN architectures (Baseline 1 using routers at mobile sites and Baseline 2 using distributed baseband units (BBUs) with enhanced networking) and an innovative Open-RAN solution based on ruggedized servers with DPUs at the mobile site and virtualized centralized units (vCUs) at the central office.

In summary, the techno-economic "what if" study reported in subsection 4.1 robustly supports the KPI 6.1 target: a 40% CAPEX reduction is feasible under optimal industry and ecosystem conditions with DPU-based Open-RAN, and this is substantiated by detailed architecture and cost modeling against relevant legacy alternatives. Final results will require ongoing updates to reflect real-world technology and market developments.

Economic Evaluation

Three Alternatives has been investigated in the study:

- Baseline 1 implements conventional routers at each site alongside dedicated radio and transport equipment.
- Baseline 2 employs distributed BBUs with improved networking, avoiding routers at intermediate sites but still relying on specialized telecom hardware.
- Open-RAN with DPU consolidates compute, networking, and optical transport using server-based platforms with Smart Network Interface Cards (NICs) and DPUs, facilitating virtualization and flexible deployment.

The "what if" cost model considers variable scenarios for hardware and software pricing, openness of the RAN ecosystem (i.e., the availability and maturity of open-source radio software and open RUs), and technology trends. The analysis demonstrates that under favorable conditions—such as a mature O-RAN ecosystem and competitive open hardware—up to a 40% reduction in CAPEX is credible compared to legacy solutions. The magnitude of savings depends on the relative costs of the ruggedized server/DPU platform versus legacy BBUs, potential cost advantages on RUs in a O-RAN ecosystem, and reductions in software licensing fees. In best case scenarios (with full O-RAN ecosystem maturity and hardware cost advantages), savings even above 40% are predicted; in less favorable conditions, the reduction may be lower or neutral.

Practical and Methodological Considerations

The economic projections in subsection 4.1 must be interpreted cautiously. The analysis is inherently speculative, employing sensitivity analysis to account for uncertainty in price evolution, vendor strategies, and market uptake of virtualized solutions. In addition, results are scenario-dependent, drawing attention to the importance of robust supply chains, open standards, and continuous monitoring of market conditions. Despite these limitations, the methodology offers valuable strategic insight into transformative potential and risk factors, supporting evidence-based decision making for network evolution.

1.2.9 KPI 6.2

KPI 6.2: *">40% reduction of O/E/O conversions in edge-edge and edge-cloud communications by developing smart edges with high-speed coherent intelligent pluggables and by moving 5G functions closer to the cell sites. [SRIA, short/mid-term evo]."*

This KPI is classified as ***mandatory***.

Status: Achieved.

The study proposed to validate this KPI is the same with the one set for KPI 6.1 and is reported in subsection 4.1. The study on 5G/6G railway coverage use case demonstrated that the use of an O-RAN architecture that virtualizes the RAN DU function (vDU) by placing it at the site and using a server equipped with DPU board for processing and networking, reduces the number of total transceivers required by 33% (50% if only the Extended Reach transceivers, the most expensive ones, are considered) and the number of O/E/O conversions by 50%. The reduction is evaluated taking as a reference the legacy RAN architecture (Baseline 1) that uses a BBU and an

aggregator router at the site. In the study carried out in subsection 4.1, non-coherent grey transceivers were used given the distances (from a few to fifty km) and single wavelength transmission involved, however similar advantages can be obtained for longer distances with coherent transceivers, using wavelength-division multiplexing (WDM) or DSCM networking if multiple optical signals would be needed, and exploiting the potential of DPU cards that integrate compute and networking function.

1.2.10 KPI 7.1

KPI 7.1: *"Intelligent data aggregation to provide data compression ratio >90% without significant information loss."*

This KPI is classified as ***additional***.

Status: Achieved.

This KPI was already achieved and documented in Section 7.1 of D2.2 where telemetry data compression results that are achieved by the different intelligent data aggregation techniques developed within SEASON project are summarized. Those techniques have been included as part of the distributed telemetry system developed and assessed within WP4/WP5 activities. The techniques are specifically designed for measurements of larger size, i.e., optical spectrum and IQ constellations (see dataset in Section 7.2), and consist of supervised feature extraction, data compression, and data summarization. Results using openly available datasets show that Intelligent data aggregation provides data compression ratio >90% without significant information loss. In fact, by combining compression and summarization, up to 3 orders of magnitude of data reduction can be achieved compared with baseline telemetry.

1.2.11 KPI 8.1

KPI 8.1: "*Autonomous operation based on multi-agent systems to reduce >25% OPEX w.r.t. manual/static operation."*

This KPI is classified as ***mandatory***.

Status: Achieved.

In D2.2 Sec. 5.1, several methods based on multi-agent systems to control point-to-multipoint DSCM-based connections were presented, i.e., to activate/deactivate subcarriers (SCs) in order to dynamically adapt the capacity to traffic needs. Moreover, this dynamicity allows deploying systems with oversubscription, which allows supporting more Txs on a same Rx than a simple static capacity assignment. The different methods were evaluated in terms of effective oversubscription and number of active subcarriers for different traffic scenarios. The latter can be reduced up to 50%, which is a promising value towards the target KPI.

In this deliverable, we extend previous results in D2.2 with numerical power consumption evaluation based on the energy consumption model and parameters presented in Section 3.2.

Detailed values in Section 5.1 show that energy savings larger than 25% can be clearly reached by dynamic management of P2MP connections with respect to static point-to-point (P2P) deployments. Moreover, the oversubscription supported by the proposed methods allows reducing the number of active hub nodes (receivers) to support the same number of leaf (transceiver) nodes, which leads to a reduction of circa 50% of power consumption in the hub with respect to static P2MP deployments.

Optimizing EDFA power in dense WDM (DWDM) networks improves energy efficiency and network performance by dynamically adjusting amplifier gains based on real-time data. Two main operation modes exist: Constant-Power mode, fitting for networks with real-time channel monitoring, and Constant-Gain mode, suitable for networks without extensive monitoring but requiring accurate span loss estimates. Using OLCs, the system collects traffic matrix data, network topology, and received power measurements to compute optimal gains. This optimization is framed as a Linearly Constrained Quadratic Programming problem, minimizing the difference between actual and target received powers while ensuring channels stay within receiver power margins. The method allows automated, network-wide gain adjustments that enhance signal quality, OSNR and power levels, reducing manual calibration need and adapting dynamically to network changes. Simulations show significant improvements in maintaining signal quality and compliance with design specs after applying this automated gain optimization, as discussed in Section 5.2.

Preferred EDFA operating modes depend on network monitoring capabilities: In heavily instrumented DWDM systems (for example, long-haul cores) where channel counts can be tracked in real time via optical channel monitors (OCMs) or similar devices, Constant-Power mode is ideal because it delivers precise per-wavelength output control.

In other environments, where you can't reliably count channels at each amplifier, Constant-Gain mode is more practical. However, it still requires an accurate, real-time estimate of the preceding span loss, so that the gain can be set correctly. If gain is misconfigured, some channels may fall outside the receiver's power or OSNR requirements, and these errors can accumulate over multiple hops. Simulation results (see Figure 5.3-2) show that improper gain settings greatly increase the risk of channels dropping below spec, further underlining the need for the optimization framework introduced in Section 5.2.

1.2.12 KPI 8.3

KPI 8.3: *"Optical layer digital twin for gradual soft-failure detection and localization with at least 1min before major impact on the service. >90% accuracy in soft-failure identification."*

This KPI is classified as ***additional***.

Status: Achieved.

This KPI was already achieved and documented in Section 8.3 of D2.2 where several methods for soft-failure management are introduced and its performance was presented. The methods

are deployed within the OCATA optical digital twin, mainly designed and assessed in WP4/WP5 activities. The results presented in this document illustrate the performance of models and algorithms that exploit IQ constellation features analysis for soft-failure detection, identification, and severity estimation. Considering challenging scenarios where an incipient soft-failure becomes hard failure in the short period of 1 day, different failures affecting Tx, EDFAs, and WSS are promptly detected (>40 min before hard failure) and localized/identified (>18 min before hard failure). With these results, the target objectives of the KPI (*Optical layer digital twin for gradual soft-failure detection and localization with at least 1min before major impact on the service/ >90% accuracy in soft-failure identification*) are clearly achieved.

2 SCALABILITY AND COST EFFECTIVENESS OF MB AND SDM SYSTEMS AND NETWORKS

The following chapters present studies related to specific areas such as the access segment (Section 3), the use of specific technologies (DPU cards) at the EDGE (Section 4), and applications using AI (Chapter 5). As mentioned previously, Chapters 6 and 7 are dedicated to cross-cutting topics such as the cost model and public data sets.

2.1 CAPACITY INCREASE DUE TO MB

MB optical networking can increase the lifetime of fibers by using unexploited portions of the spectrum. The C-band and, more recently a combination of the C-band and the L-band have been typically exploited. The adoption of MB poses several aspects to be taken into account: i) EDFA won't be used in S- or E-band, thus amplifiers used in bands different than the C-band might have different noise figures; ii) transmission parameters (e.g., attenuation, chromatic dispersion) experience a dependency from wavelength/frequency and cannot be approximately considered as constant over a so wide transmission spectrum; iii) Stimulated Raman Scattering (implying a power transfer from longer to shorter wavelength channels) should be considered [Sam20] since – as an example – can increase the non-linearities in the C-band and the L-band.

In the framework of the SEASON project, we evaluated the capacity increase of MB in a backbone network, in particular by referring to the Spanish backbone topology shown in [Sam22]. The network performance has been evaluated with simulations and with the routing and spectrum assignment proposed in [Sam20] giving preference to the bands according to the following rank to allocate channels: first C-band, then L-band, S-band, or E-band, and accounting for Stimulated Raman Scattering.

A custom-built event-driven C++ simulator has been used and simulations have been carried out on an Intel(R) Core(TM) i5 @2.67 GHz, with 4 RAMs of 4 GB. The reference Spanish transport network is of 30 nodes and 55 bi-directional links, all composed of standard standard single-mode fibers (SSMFs). Each span is 80 km long. The inter-arrival process of 200 Gb/s connection requests is assumed to be Poisson. Inter-arrival and holding times are exponentially distributed with an average of $1/\lambda$ and $1/\mu$ = 500 s, respectively, with the connection requests uniformly distributed among all node pairs. Traffic load is expressed as λ/μ. The following modulation formats and symbol rates are considered: DP-16QAM and DP-QPSK, both at 32 Gbaud. The same forward error correction (FEC) overhead of 21% is assumed in all cases. Single carrier transmission is adopted for PM- 16QAM, while dual carrier is assumed for PM-QPSK, thus always guaranteeing a 200 Gb/s net rate. We assume 200 Gb/s connections using 37.5 GHz and 75 GHz if PM-16QAM or PM-QPSK is adopted, respectively. For simplicity, we set a per-band power with flat spectral load [Cur15]. The generalized signal-to-noise ratio (GSNR) of the worst channel (comprising the effect of cross-phase modulation) is assumed for each band. This way, when

new channels are set up, any existing lightpath will still show acceptable performance thus avoiding possible reconfigurations. Assuming the pre-FEC bit error rate (BER) threshold 4×10^{-3}, according to the theoretical BER-vs-GSNR curve, the following threshold values are considered for the GSNR: 15.1 dB for DP-16QAM and 8.5 dB for DP-QPSK.

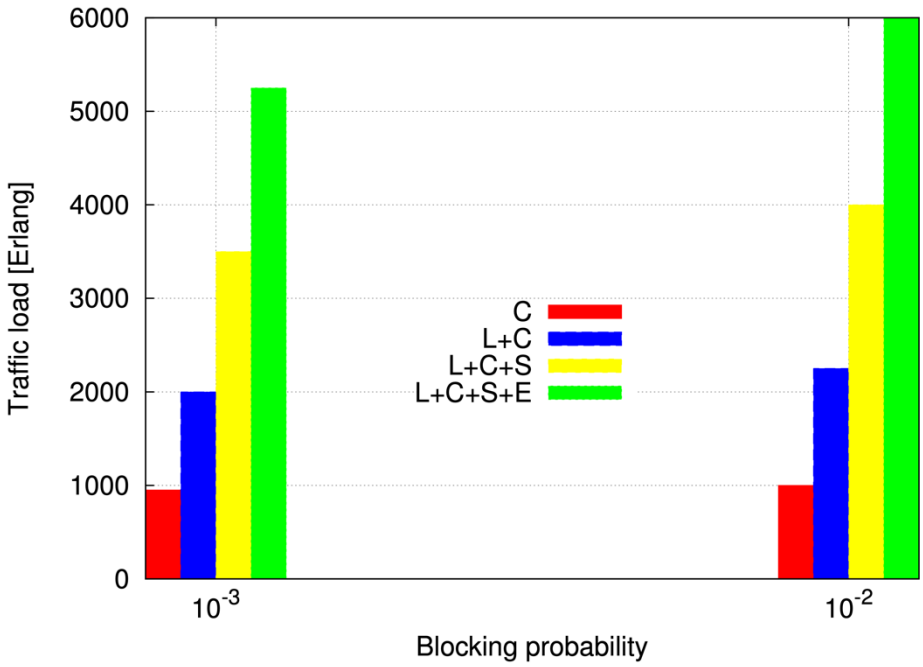

Figure 2.1-1 Load for a given connection blocking probability.

Figure 2.1-1 shows the supported traffic load at a given blocking probability (i.e., 10^{-3} or 10^{-2}) in the case of C, L+C, L+C+S, L+C+S+E systems. The figure clearly shows that the activation of other bands increases the supported network capacity. In particular, for a blocking probability of 10^{-2}, the capacity of a L+C+S+E-band system increases by six times (6x) the load of a C-band network in the considered conditions. Note that the activation of the O-band in a backbone network, with typically very high distances, may not bring any appreciable benefits, as observed in [Sam20].

Thus, if larger capacity increase is required by an operator, also the adoption of SDM is suggested. In particular, as mentioned in Sec. 1.2.2, in order to achieve a 120x of capacity increase with respect to C-band and considering the 6x of increase activating L-, S-, and E-band, at least 20 SMF fibers would be needed.

2.2 Realistic Capacity Scaling of Single-Band and MB Optical Transport Systems

As global data traffic continues to grow exponentially (driven by cloud computing, 5G, internet of things (IoT), and AI applications) traditional single-band optical transmission systems, typically

operating in the C-band, are approaching their capacity limits. To meet future bandwidth demands and ensure scalable, cost-effective network evolution, MBT has emerged as a compelling solution.

Focused on validating KPI 2.1, this section discusses the challenges and opportunities to increase the available bandwidth of deployed fibers with MBT, that together with the usage of spatial division multiplexing techniques such as multi-fibers or multi-core fibers, will make the available bandwidth grow by a factor of 120. The content of this section was published in [Sou25]. It analyzes the optical performance and capacity of single-band (SB) and MB transmission systems, ranging from a 6 THz single super-band to a 24 THz four-super-band configuration. Configurations include Super-C, Super-C+L, Super-S+C+L, and Super-S+C+L+O systems. A 4.8 THz C-band-only system is also evaluated for benchmarking. This framework enables comparison of the spectral efficiency, complexity, capacity, and fiber requirements of mid-term future optical transmission systems.

2.2.1 Simulation Setup and System Optimization

The simulation setup, system modeling and optimization was already discussed in Deliverable 3.3 [SeaD3.3]. Therefore, only a summary of the simulation setup and optimization framework will be presented in the following paragraphs.

Table 2.2-1 presents the transmission bands that compose the four configurations of the MBT system analyzed in this work. All systems include a minimum 500 GHz guardband between adjacent bands. As a reference, we first consider a C-band system spanning 191.3–196.1 THz. The Super-C system extends this to 190.8–196.8 THz. The Super-C+L system covers 184.3–196.8 THz (using the super C and super L bands), while the Super-S+C+L system spans 184.3–203.3 THz (using the super S, super C and super L bands). Finally, the Super-S+C+L+O configuration adds the Super O band (220.5–226.5 THz) to the three-band system, maintaining a guard band of over 17 THz from the S-band, as recommended in [Sam22, Sou25-2]. In this case, nonlinear interference (NLI) and the SRS-induced power transfer between the O-band and other bands are considered negligible. Although some transmission bands are formally defined with bandwidths exceeding 6 THz, we assume that the optical amplifiers used across these bands will have the same bandwidth as those in the super C-band, due to the challenges in extending amplifier bandwidth.

Table 2.2-1 Minimum and maximum frequencies of each amplification band.

Band	Start Frequency [THz]	Stop Frequency [THz]
Super L	184.3	190.3
C	191.3	196.1
Super C	190.8	196.8
Super S	197.3	203.3
Super O	220.5	226.5

The transmission system operates with 40 channels per super-band and 32 channels within the C-band, all spaced on a 150 GHz channel grid. Each channel carries a 120 GBd signal shaped with a roll-off factor of 0.15.

The fiber characteristics are adopted from the model described in [Sou24]. Optical amplifiers are characterized by constant noise figures: 5 dB for the super L and super C/C bands, and 7 dB for the super S and super O bands. The total insertion loss introduced by the multiplexing and demultiplexing components is assumed to be 1.5 dB.

The objective of launch power optimization is to maximize the overall system GSNR while minimizing GSNR fluctuations across each spectral band. After optimization, analysis is performed on the capacity and spectral efficiency of the different transmission systems in a point-to-point link with a variable number of 75 km spans. For each band, the GSNR of channel i at the end of a lightpath with N spans of length L_n is given by Eq. 1, where $\mathrm{GSNR}_{\mathrm{OPT}}^{\mathrm{L_N}}(i)$ is the optimized single-span GSNR of channel i and M is the system margin defined by Eq. 2. The system margin comprises a fixed 1 dB margin and a variable contribution that depends on the number of optical amplifiers traversed (N_{OLAs}).

$$\mathrm{GSNR}_{\mathrm{N}}(i) = \mathrm{GSNR}_{\mathrm{OPT}}^{\mathrm{L_N}}(i) - 10\log_{10}(N) - M \quad \text{Eq. 1}$$

$$M = 1 + 0.05N_{OLAs} \quad \text{Eq. 2}$$

A lightpath is feasible for a given signal configuration and transmission band if the required signal-to-noise ratio (SNR) is less than $GSNR_N$. In this work, five different bit rates are considered and their required OSNR and SNR are presented in Table 2.2-2. These values were taken from the OpenRoadm multi-source agreement (MSA) specifications version 7.0 [O-RDM] for 400 Gb/s and 800 Gb/s formats at 118 GBd (the values for the remaining bit rates were linearly interpolated). The SNR_{req} values are calculated directly from the $OSNR_{req}$ values using the relation given in [Ess12], that is, using Eq. 3 with B_{ref} equal to 12.5 GHz.

Table 2.2-2 Required OSNR (defined in 0.1 nm) and SNR for each considered bit rate.

Bit rate [Gb/s]	**400**	**500**	**600**	**700**	**800**
$\mathrm{OSNR}_{\mathrm{req}}$ [dB]	20.0	21.8	23.5	25.8	27
$\mathrm{SNR}_{\mathrm{req}}$ [dB]	10.2	12.0	13.7	16.0	17.2

$$SNR_{req} = OSNR_{req} - 10\log_{10}\left(R_S/B_{ref}\right) \quad \text{Eq. 3}$$

2.2.2 Results and Discussion

The optimized launch power and the corresponding GSNR values for all considered transmission systems of a 75-km span are presented in Deliverable 3.3 [SeaD3.3]. A summary of the results is that single-band systems exhibit the highest average GSNR (27.0 dB and 26.9 dB for the C and Super-C systems, respectively). The lower GSNR of MB systems is due to the additional

MUX/DEMUX losses, higher fiber attenuation, nonlinear effects, and amplifier noise figures. Additionally, the super S and super O bands show significantly lower average GSNR than the super C and super L bands, with the super O band performing the worst.

These GSNR variations across bands result in diverse capacity-reach profiles for each system. To illustrate this, consider a link composed of 75-km spans and a transceiver operating in the modes listed in Table 2.2-2.

Figure 2.2-1 and Figure 2.2-2 show the capacity vs. reach and spectral efficiency vs. reach profiles, respectively, with reach measured in number of 75-km spans. Single-band systems consistently achieve the highest spectral efficiency (SE) due to their superior average GSNR on a single span ($GSNR_{OPT}^{L_n}$). The maximum SE of 5.33 bit/s/Hz (meaning that all channels transmit 800 Gb/s) is maintained at a maximum of 6 spans for the C and Super-C systems, 5 spans for the Super-C+L, 3 for the Super-S+C+L and 1 span for the Super-S+C+L+O system.

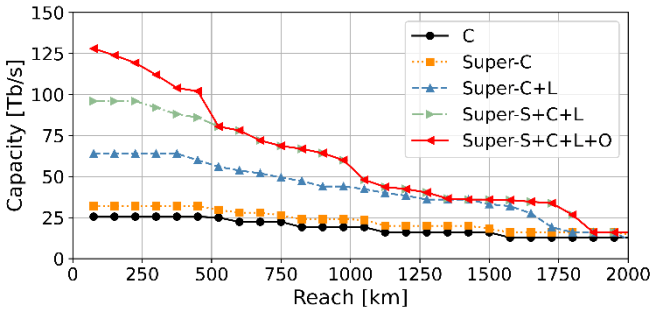

Figure 2.2-1 Capacity of the different transmission systems for different link lengths.

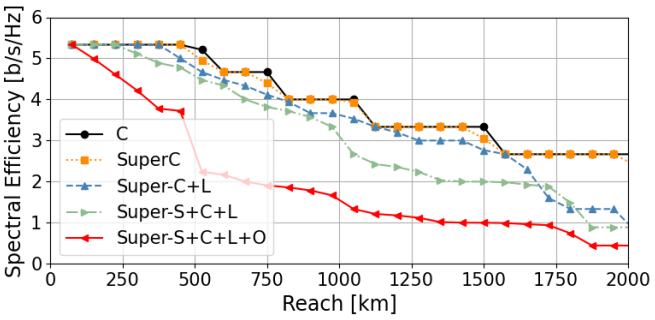

Figure 2.2-2 Spectral efficiency of the different transmission systems for different link lengths.

Despite the superior average GSNR and SE of single-band systems, MB systems generally offer higher total capacity across most distances, thanks to their greater channel count. The capacity gain from adding an amplification band is most pronounced at shorter distances, where the optical performance degradation is less severe. For example, adding the super L band to a Super-C system doubles the capacity (from 32 to 64 Tb/s) for links up to 375 km. However, this gain diminishes with distance and disappears beyond 1800 km.

Adding the super S band to a Super-C+L system yields a smaller benefit: a 50% capacity increase for links up to 225 km, tapering off to no gain beyond 1800 km. Compared to a system with Super-C only, the Super-S+C+L configuration triples the link capacity up to 225 km.

The super O band supports transmission only over distances up to 450 km. When added to a Super-S+C+L system, it increases capacity by 33% for a single 75-km span. However, this benefit gradually decreases with distance, becoming negligible at 450 km. Beyond this point, the capacities of the Super-S+C+L and Super-S+C+L+O systems are the same (because of the guard band between the super O and the other bands). As a result, the SE of the Super-S+C+L+O system drops sharply from 450 to 525 km, since the super O-band channels, which initially carried 400 Gb/s, cease to transmit any data.

Lastly, Figure 2.2-3 illustrates the number of parallel fibers required by various transmission systems to achieve a capacity 120 times greater than that of a single-fiber C-band-only system, across different link lengths. Table 2.2-3 complements this figure by providing specific values for 75 km, 300 km, 675 km, and 1425 km.

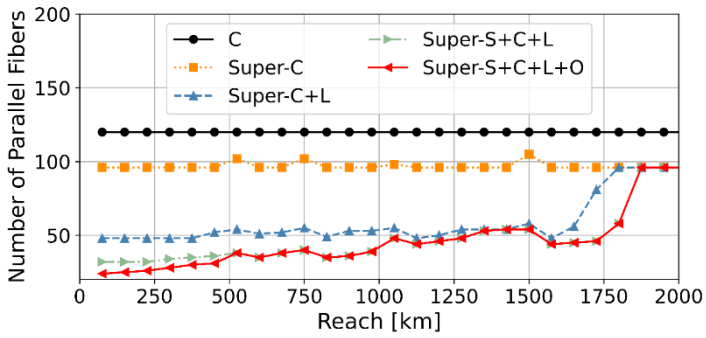

Figure 2.2-3 Number of required parallel fibers to achieve a 120 times capacity increase of the C-band-only system for different point-to-point link lengths.

Table 2.2-3 Required parallel fibers to achieve a 120 times capacity increase of the C-band-only system for selected reaches.

System	Reach [km]			
	75	300	675	1425
C	120	120	120	120
Super-C	96	96	96	96
Super-C+L	48	48	52	54
Super-S+C+L	32	34	38	54
Super-S+C+L+O	24	28	38	54

As expected, the C-band-only system consistently requires 120 parallel fibers at all distances. The Super-C system typically requires 96 parallel fibers across most distances, due to its similar average GSNR and 20% higher channel count compared to the C-band only system. Minor deviations at 525, 750, 1050, and 1500 km arise from slight GSNR differences and the limited set

of transceiver operating modes. For instance, while the C-band-only system may transmit at 800 Gb/s at a certain distance and frequency slot, the Super-C system might drop to 700 Gb/s due to a marginal GSNR reduction (0.1 dB), lacking an intermediate mode closer to 800 Gb/s.

The trends for the other systems mirror their respective capacity curves. The inclusion of the super O band allows us to reduce the number of required fibers only for links up to 450 km. For instance, for a 300-km link, only 23.3% of the 120 fibers are required. Beyond 450 km, the Super-S+C+L and Super-S+C+L+O systems require the same number of fibers. For a target 675-km link, the former provides the same capacity as 120 parallel C-band systems with only 31.7% of the fibers.

Since only the super L band in these MB systems maintains GSNR levels comparable to single-band systems, while the super C, super S, and super O bands have lower GSNR, the number of required fibers increases rapidly beyond 1575 km (for the Super-C+L system) and 1750 km (for the Super-S+C+L and Super-S+C+L+O systems), eventually converging to 96, matching the Super-C system. Still, for a target 1425 km link, a Super-C+L-band system can offer the same capacity as the reference (120 fiber) C-band system at the expense of only 45% of the number of fibers.

2.2.3 Conclusions

MB transmission offers a cost-effective route to significantly increase the capacity of existing optical systems, leveraging the wide low-loss window (~54 THz) of ITU-T G.652D fiber. However, adding more amplification bands tends to reduce spectral efficiency, particularly over long distances. As a result, full exploitation of the entire 54 THz window remains impractical for data transmission. This section analyzed the optical performance and capacity of four MBT system configurations: Super-C, Super-C+L, Super-S+C+L and Super-S+C+L+O and compared them to the benchmark C-band-only system.

Results indicate that a practical balance between spectral efficiency and bandwidth expansion appears to be the use of two or three super bands, each occupying 6 THz (such as Super-C+L or Super-S+C+L). For localized capacity upgrades in smaller network segments, the super O band can be selectively deployed.

This capacity enhancement translates into a reduced need for parallel fibers to increase the capacity of a single-fiber single-band system with 4.8 THz. For a 75 km link, the number of required fibers drops by 20%, 60%, 73%, and 80% for the Super-C, Super-C+L, Super-S+C+L, and Super-S+C+L+O systems, respectively. At 1425 km, the reductions are more modest for the MB systems: 20% for Super-C and 45% for the other MB systems.

All the analysis presented in this section validates KPI 2.1 by indicating different solutions on how to increase the available bandwidth by a factor x120 by combining MBT and SDM.

2.3 CAPEX SAVINGS OF NODE ARCHITECTURES FOR HIGH-CAPACITY MBOSDM OPTICAL NETWORKS

This study explores the cost savings enabled by multi-granular, modular, and adaptable switching network node architectures with advanced capabilities, leveraging MB and SDM technologies. The nodes are designed to meet the escalating traffic demands and the rigorous requirements of forthcoming optical networks. Deliverable 3.2 [SeaD3.2] introduced three architectural solutions grounded on current technologies, which present viable approaches for efficient MBoSDM systems for the next generation of optical networks (see Figure 2.3-1).

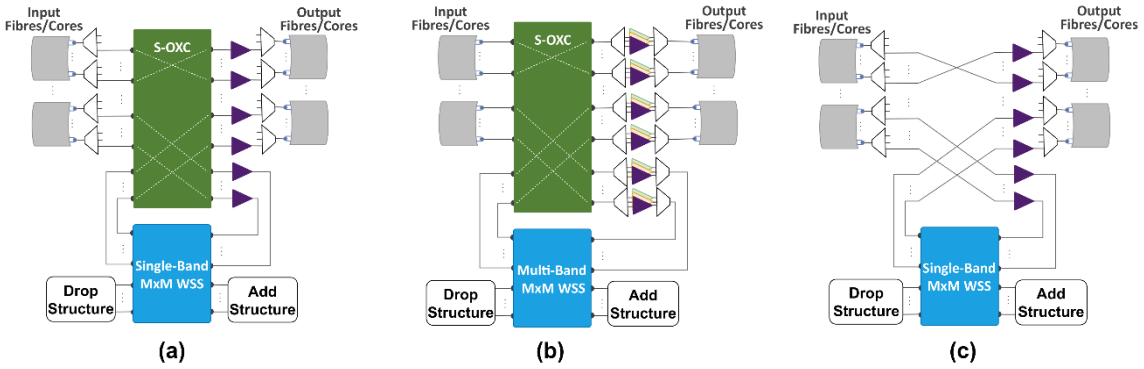

Figure 2.3-1 (a) Node architecture #1 (SBS-OXC): single-band matrix-switch-based MBoSDM node with a S-OXC. (b) Node architecture #2 (MBS-OXC): MB matrix-switch-based MBoSDM node with a S-OXC. (c) Node architecture #3 (SB): single-band matrix-switch-based MBoSDM node without a S-OXC.

The main simplifying feature of the proposed architectures is the possibility of directly interconnecting fibers/cores (bypassing the WSS and the add-drop structure) via physical connections or a S-OXC. This direct connection holds the potential to simplify the node by reducing the WSS and optical amplifiers count in relation to a reference architecture that does not allow direct fiber/core switching [SeaD3.2]. The capital expenditure of the node architectures depends on the number and cost of S-OXCs, MxM WSSs and optical amplifiers, assuming that the cost of band filters is negligible.

For simplicity, an equal number of bands per core (B) and fiber cores per degree (C), equations can be found for the number of ports and number of devices required by each architecture [SeaD3.2]. Table 2.3-1 shows these equations. Here, D represents the nodal degree, N_C denotes the number of core pairs directly switched, and $N_{A/D}$ indicates the number of ports of the WSSs connected to the A/D structure. It is worth noting that these figures exclude the A/D structure since the same one can be employed independently of the node architecture.

Dissemination Level	PUB (Public)

Table 2.3-1 Number of components and port count for the different node architectures.

	Reference Arch.	SBS-OXC	MBS-OXC	SB
S-OXC Count	0	B	1	0
S-OXC Port Count	--	$4CD - 2N_C$	$4CD - 2N_C$	--
MxM WSS Count	B	B	1	B
MxM WSS Port Count	$2(CD + N_{A/D})$	$2(CD - N_C + N_{A/D})$	$2(CD - N_C + N_{A/D})$	$2(CD - N_C + N_{A/D})$
Optical Amplifier Count	$2BCD$	$2BCD - BN_C$	$2BCD - BN_C$	$2BCD - BN_C$

The analysis of Table 2.3-1 shows that each directly switched core pair saves two ports on the MxM WSS and B optical amplifiers. It is important to note that it is assumed that all bands of a core/fiber pair are directly switched and it is not possible to directly switch only some of the transmission bands from a core/fiber.

2.3.1 Cost Model

The cost of the optical amplifiers is taken from the cost model for dual stage amplifiers (see Section 6.1.2 and 6.3). However, the cost of WSSs and S-OXCs are not available and should therefore be estimated using a similar strategy as the one used in the cost model. The cost premium is assumed to be 10% for the L-band amplifier and 25% for the S-band.

To estimate the cost of the MxM WSS, it is assumed that it is implemented by a concatenation of smaller 1xN_w WSSs in a route and select structure [Sah17]. With this assumption, the number of required 1xN_w WSSs to construct an MxM WSS is given by $2\left(M - N_{A/D}\right)$ if $M \leq N_w$. (additional WSSs are not required for the add and drop ports because they are taken directly from the N_w ports of the WSSs as shown in Figure 2.3-2). This work does not consider the case of $M > N_w$.

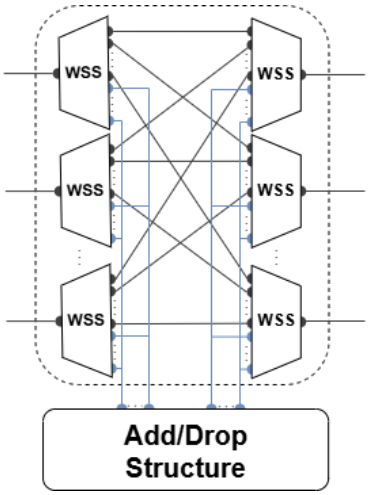

Figure 2.3-2 Schematics of an implementation of an MxM WSS by a concatenation of smaller 1xN_w WSSs in a route and select structure.

The costs and power consumption values of a 1x12 and a 1x32 WSS and of a 16x16 and 32x32 S-OXCs are estimated using the following framework. The costs are [2.5, 4.0, 5.1, 12.1] c.u. (cost units as defined in chapter 6) for the 1x12 and the 1x32 WSSs, the 16x16 and the 32x32 S-OXCs, respectively. The S-OXC cost values are taken from the same public website as the cost model (described in chapter 6) and using the same strategy (European cost value excluding the value-added tax). Even though there is only a single exemplary of each device, the values are consistent with expert partners consultations. The WSS values are hard to find in public databases and are therefore based solely on expert knowledge from the partners. For the architecture using MB devices (MBS-OXC), we assume that a 3-band MB WSS costs 95% of three single-band WSSs and that the integrated amplifiers and band filters device for amplification of three bands cost 95% of three single-band amplifiers.

2.3.2 Results

The evolution of the CAPEX of a network node for an increasing number of directly switched core pairs is evaluated. The evaluation considers a transmission system with three bands B (S, C and L) and a node with four degrees D, three parallel fibers per degree C (leading to 12 input/output fiber pairs) and 4 WSS add/drop ports $N_{A/D}$.

Figure 2.3-3 shows the CAPEX evolution for the different proposed node architectures and the reference node as a function of N_C. The cost of the reference node remains at 225 c.u. because it does not offer the possibility of core/fiber switching. When $N_C = 0$, the cost of architecture SB is the same as the reference node because they have the same number of devices. The other two proposed node architectures have the additional cost of the S-OXCs, but the reduced cost of the integrated MB devices make the total cost of the MBS-OXC architecture to be similar to the reference and SB node architectures. On the other hand, when all fiber pairs are directly connected ($N_C = CD = 12$), the cost of the proposed architectures is reduced to half of the initial number of amplifiers plus the S-OXCs, given that no WSS is needed. The faster reduction of the cost of the proposed architectures from $N_C = 3$ to $N_C = 4$ and from $N_C = 7$ to $N_C = 8$ are caused by the reduction of the number of ports of the WSS (from 1x32 to 1x12) and the S-OXC (from 32x32 to 16x16), respectively.

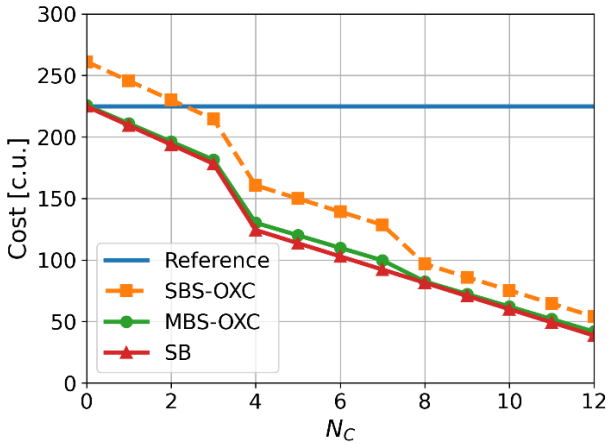

Figure 2.3-3 Cost evolution of the different network node architectures for an increasing number of directly switched core pairs (N_C).

2.3.3 Conclusion

This section evaluated the potential cost savings enabled by MB and SDM node architectures. A key simplification in the proposed designs is the ability to directly interconnect fibers or cores—either through physical links or a S-OXC—bypassing the WSS and the add-drop structure. This direct interconnection can significantly streamline node architecture by reducing the number of WSS units and optical amplifiers compared to a reference design that does not support direct fiber/core switching.

Using the SEASON cost model, supplemented with additional component cost estimates derived through a similar methodology, the study analyzed CAPEX trends for various proposed node architectures and a reference node as a function of the number of directly connected fiber/core pairs (N_c). The results indicate that CAPEX can be reduced by more than 50%, thereby validating KPI 2.2. In practical network scenarios, N_c at a given node depends on factors such as traffic patterns, the number of parallel fibers between nodes, routing and wavelength assignment algorithms, and the overall network configuration. Further analysis on real network topologies is identified as future work.

2.4 Techno-Economic Study on the Benefits of Open Coherent Pluggables in MB over SDM Architectures

The continuous growth of traffic led by video streaming, cloud computing and AI applications is pushing optical network infrastructure to the adoption of new technologies. MB and MCF technologies are promising solutions which can increase spectral and spatial capacity. For near-term capacity growth, MB technology has been adopted over the past few years. Although in the longer term, the combination of both MB and MCF technologies, referred to as MBoSDM, will be needed [Arp24]. Within this context, there is an ongoing debate between the use of

Internet Protocol (IP)+WDM vs IP-over-WDM (IPoWDM) in optical metro/core networks [Mar24].

2.4.1 Methodology

2.4.1.1 Multi-period network planner

This work performs a multi-period network planning study that considers the installation of new resources as the traffic growth requires. This approach allows for a more realistic techno-economic assessment, as it simulates the operator's investment behavior and avoids over-investment in the early years. The planner receives as input a traffic demand matrix per period, a table of maximum optical reach (MOR), and a cost and power table for all network devices. The traffic demand matrix follows the model presented in [Pat20], where the traffic is focused on datacenters. The traffic capacities of this model are scaled according to the topology, in order to reach upgrades that lead to the use of MB, as well as multicore technologies. The MOR table is used to encode device, capacity, modulation, band and fiber characteristics.

At each period, the planner tries to allocate the traffic for each demand. If the demand has assigned client equipment with the possibility of upgrade, the planner performs an upgrade to fulfill the requested traffic. In the case that the requested traffic is still not fulfilled after upgrades of existing client equipment or the demand has no previous ones assigned, the planner checks the k-shortest paths against the MOR table, to obtain a suitable candidate (transponder, coherent pluggable). With this candidate, the planner tries to allocate spectrum resources. When the spectrum assignment is not successful, the planner triggers an upgrade in order to fulfill the demand. The following upgrade strategies are considered:

- Single-mode fiber (SMF) architecture: Considers a SMF C-band solution and the addition of extra fibers in C-band when capacity is exhausted.
- Multicore fiber (MCF) architecture: Considers a sequential evolution from SMF C-band, to SMF C+L-band, and finally to MCF C+L-band solutions with/without core switching.

As mentioned, upgrades are performed in the least-cost manner, adding only the equipment and fiber strictly needed to meet the current traffic requirements. Devices such as transponders, pluggables, amplifiers, and ROADMs are installed incrementally. The unused capacity in already deployed elements is reused in subsequent periods. The information of the network evolution, which indicates the amount of equipment deployed at each period forms the basis for the techno-economic evaluation.

2.4.1.2 Optical reach model

The feasibility of a lightpath is determined by the MOR, obtained from a Gaussian Noise (GN) model considering the following noise sources: amplified spontaneous emission (ASE), NLI and inter-core crosstalk (ICXT). The model in this work is approximate but fit for purpose. We employ pre-computed optical reach tables due to their computational efficiency. These tables provide

maximum distance values according to the device, data rate, modulation format, band and fiber type. For clarity, we summarize the results into Table 2.4-1, that reports representative MOR values for the different coherent pluggables and transponders. The complete MOR tables are available on request. For the calculations, fiber parameters were derived from the uncoupled 7-core MCF design reported in [Hay12]. This fiber exhibits ≈2.4 dB better SNR than SSMF due to its lower attenuation in the C-band, despite higher dispersion. Consequently, MCF shows longer reach than SSMF in the C-band, while L-band performance remains comparable, consistent with [Sha18]. This table highlights the main trends. Mainly, the longer reach of MCF relative to SMF and the differences in reach between the different client devices, being the pluggable ZR+ with 0dBm output power the closest in MOR values to the transponder.

Table 2.4-1 Representative MOR values (km).

	SMF(C)	SMF(C+L)		MCF(C+L)	
	C	C	L	C	L
400ZR	80	80	80	240	160
800ZR	160	160	160	320	160
400ZR+	1600	1600	1520	2080	1520
800ZR+	880	880	880	1040	720
TR 400G	2240	2080	2240	3600	2960
TR 800G	2640	2560	2720	3040	2640
TR 1.2T	960	880	960	1040	960

2.4.1.3 Cost model

The total cost of ownership (TCO) is based on both CAPEX and operational expenditure (OPEX). The CAPEX includes the purchase of all equipment. The model accounts for the cost of transponders, coherent pluggables, optical amplifiers, and switching elements. All prices are expressed in cost units (c.u.), and are normalized to allow direct comparison across architectures. The OPEX considers energy and fiber leasing costs. Energy costs are calculated as the product of device power consumption, annual operation hours, and the price of electricity for the respective period. Fiber leasing costs are considered per km per year and are differentiated between SMF and MCF. Table 2.4-2 presents a summary of the equipment for cost and power consumption. Information from Section 6 (Cost model) was considered for the case of C-band/L-band amplifiers, as well as for coherent pluggables. The rest of the values are estimated from vendor price lists and datasheets. For equipment not available on the market, estimations and interpolations were performed. The costs are normalized to the price of a C-band EDFA amplifier (1 c.u. = 2200 euros). For WSSs, coherent pluggables and transponders, a 20% increase in price was considered for L-band respect to C-band.

Table 2.4-2 Cost model summary.

Category	Example Components	Cost (c.u.)	Power (W)
Amplifiers	C-band / L-band, C+L, / MCF C+L	1 / 1,2 2,4 / 4,4	16 / 23 40 / 60
Couplers/splitters	C+L coupler/splitter	0,2	0,25
SDM / MCF devices	SDM fan-in/fan-out	0,6	–
WSS (C-band)	1×4 / 1×12 / 1×32 1×80 / 1×160	2,1 / 3,8 / 6 10 / 15,8	10 / 22 / 50 119 / 233
OXCs (MEMS-based)	32×32/ 64×64 128×128 / 196×196	3,2 / 6,4 12,8 / 19,6	12 / 28 58 / 91
Coherent pluggables (C-band)	400ZR / 800ZR 400ZR+/800ZR+	2,7 / 4,5 3,2 / 5,3	17,5 / 29 23,5 / 29
Transponders (C-band)	400G / 800G / 1.2T	7,3 / 12,7 / 16,4	330
Fiber lease (OpEx)	SMF / MCF 7-cores (per core per km/year)	0,3 / 0,4	–

2.4.2 Results and discussion

The network planner is applied to 9 different scenarios, which comprehend the cross-product of three upgrade strategies and three client technologies. As mentioned in the methodology, the three upgrade strategies are using (1) bundles of SMF, (2) MCF without core switching and (3) MCF with core switching. The three client technologies refer to the use of coherent pluggables only, transponders only or a mixed case, where both types of devices can be used. Each scenario is tested on three reference topologies: Coronet30, which represents a medium-scale US backbone with 30 nodes and 36 links; Cost266, corresponding to a continental European network with 28 nodes and 41 links; and Germany50, which models a dense national backbone with shorter average spans, having 50 nodes and 88 links. The evaluation metrics include cumulative CAPEX, OPEX and TCO over the planning horizon.

Figure 2.4-1 displays the TCO breakdown for each scenario of the three topologies. OPEX related to fiber leasing (FL) dominates in all cases, being the proportion of contribution to the TCO higher in topologies with larger average spans, like Coronet30 and Cost266. For both of these topologies, the total TCO ranking regarding client technologies follows the order of Coherent pluggables > Mixed > Transponders solutions. In the case of Germany50, the proportion of fiber leasing is close to the CAPEX values. Therefore, the TCO behavior changes respect to the client technologies, obtaining an order of Mixed > Transponders > Coherent plugabbles. This proves that for backbone topologies with shorter link lengths, coherent pluggables present a cost-effective solution.

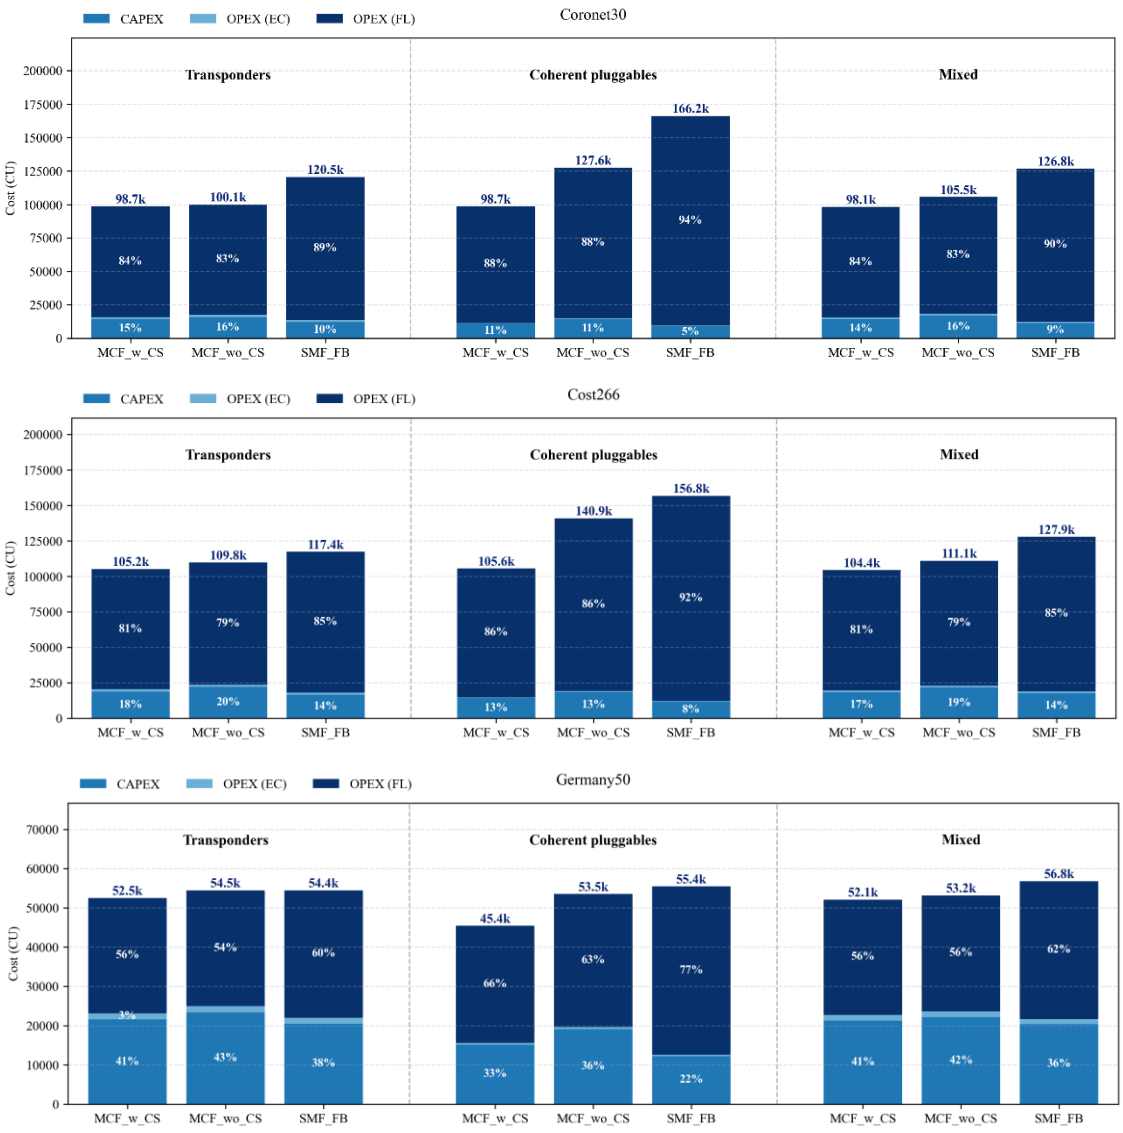

Figure 2.4-1 TCO breakdown.

In order to have a better view of the TCO differences, total TCO values are normalized with the minimum value in each solution (less expensive scenario). In Figure 2.4-2, a heatmap is presented with these values for each topology. For Coronet30, the least expensive scenarios correspond to MCF with core switching upgrade strategy. Similar values are obtained for coherent pluggables, transponders and mixed cases. The most expensive scenario is the one using SMF with fiber bundles and coherent pluggables technology, having a TCO 58% higher than the lowest one. For Cost266, the MCF with core switching upgrade strategy remains as the best option in terms of cost. The difference between client technologies is also minimal in this upgrade strategy. Nevertheless, Germany50 shows results that highlight the cost efficiency of coherent pluggables compared to transponders and the mixed case. Germany50 exhibits 13–15% higher savings with coherent pluggables. As mentioned before, this can be attributed to the average path lengths being shorter for its demands.

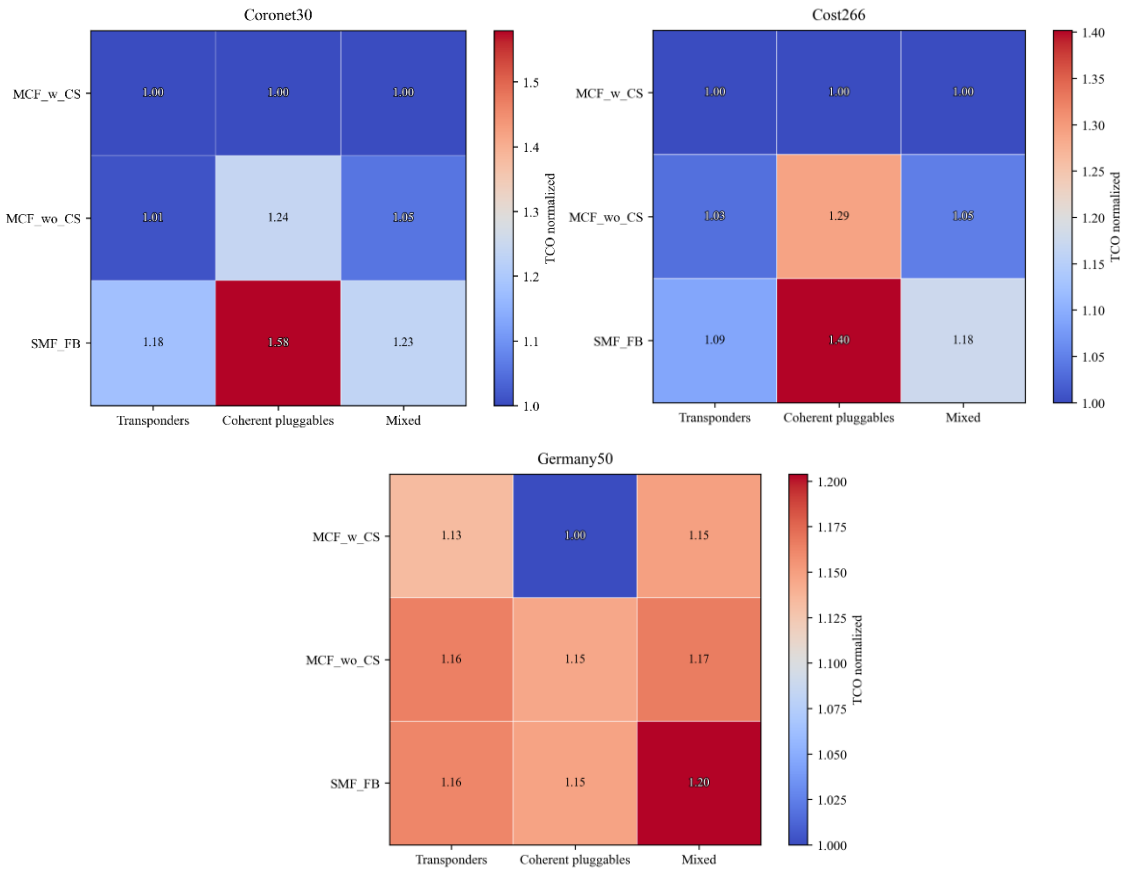

Figure 2.4-2 TCO normalized heatmap.

2.4.3 Conclusions

The multi-period network planning study was applied to compare 9 scenarios, defined by different node architectures and client technologies across three realistic topologies. Results show that MCF with core switching provides the best long-term scalability. SMF bundles are an attractive short-term solution but are penalized by fiber leasing. Considering the case of SMF bundles with transponders or coherent pluggables as baselines, a maximum cost reduction of 15% and 47% can be achieved via MCF with core switching, respectively. Coherent pluggables offer a cost-effective solution in medium-sized topologies, while transponders display comparable results in large-sized topologies.

2.5 Nonlinear-Aware Design of Filterless Horseshoe-and-Spur Optical Networks

Horseshoe topologies are well-suited for resilient hub-and-spoke traffic patterns. Filterless architectures, coupled with coherent-based P2MP transceivers (using DSCM), offer cost-effective, efficient all-optical aggregation of traffic flows destined for the same hub node. While

these technologies hold promise for long-term scalability and resource utilization, network reconfiguration remains a challenge when accommodating dynamic changes like adding a new leaf node. Recent research [Hos24] proposed a scalable horseshoe-and-spur optical network architecture to minimize reconfiguration and operational costs while enabling long-term traffic growth as shown in Figure 2.5-1(a). By decoupling the design of the static horseshoe from the dynamic spurs (trees), the network can adapt to changing traffic demands without disrupting the core infrastructure.

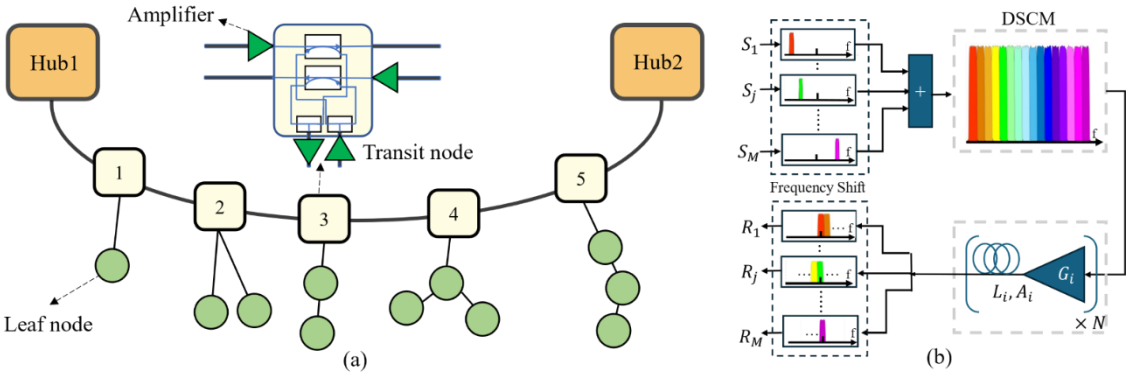

Figure 2.5-1 Illustration of a typical horseshoe-and-spur topology and (b) the simulation schematic.

To enhance network performance and reduce operational costs, this work relaxes the stringent nonlinearity constraints that typically limit amplifier gain. By allowing the use of higher-gain amplifiers, we significantly increase the available power budget, thereby decreasing the required number of amplifiers. To ensure the feasibility of operating with lower margins, we developed a specialized nonlinear noise modeling tool based on the split-step Fourier method (SSFM) considering unequal span/amplifier gain and DSCM (see Figure 2.5-1 (b)). This tool is tailored to the unique challenges of this problem, enabling accurate prediction of system performance under nonlinear conditions as a post-analysis.

As shown in

Figure 2.5-2, we calculated the non-linear noise efficiency (η) spectrally located over the central SC, normalized by P_{SC}^{-3}. This was done for a DP-16QAM single-carrier system with a consistent total symbol rate, transmitted over 1 to 6 spans of 30 km using the SSFM. The near linearity of this curve suggests that most of the nonlinear interference, primarily arising from self-phase modulation (SPM), accumulates within the initial fiber spans.

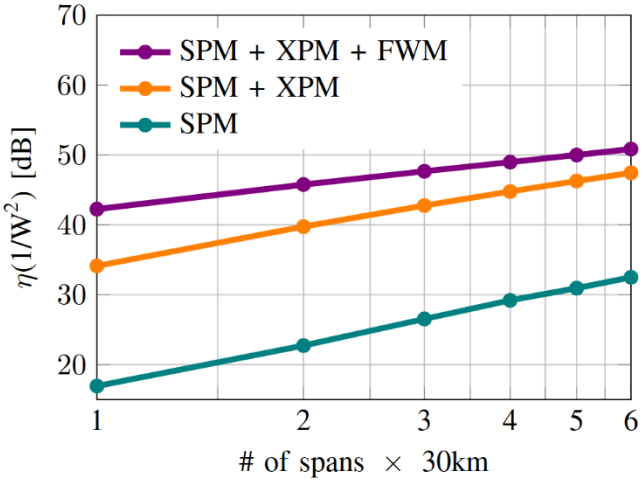

Figure 2.5-2 η in dB for 16 SCs at 4 GBaud with DP-16QAM when considering all nonlinearity components (violet), neglecting FWM (orange), and single SC (teal), respectively, transmitted over 1 to 6 spans of 30 km.

Figure 2.5-3 presents the average number of amplifiers required for a 5-transit node network under various power budgets, coupler configurations, and nonlinearity level scenarios. The figure includes 90% confidence intervals, which are shown as highlight. For baseline horseshoes, results are shown in Figure 2.5-3(a). The high-nonlinearity approach (circle) is shown to have a clear advantage in terms of amplifier reduction (e.g., for the same fixed power budget) and a significantly increased power budget when compared to the low-nonlinearity design. Note that, even with a higher level of nonlinearity, the quality of transmission (QoT) is meeting the requirement. It is also important to point out that with designing in low nonlinearity regime, power budgets exceeding 13 dB have become infeasible for the balanced coupler scenario. This is because there is less flexibility in controlling power levels to satisfy the constraints compared to when using a set of different power splitting ratios. Similarly, Figure 2.5-3(b) presents the results for extended horseshoes. As shown, an amplifier reduction of up to 20% can be achieved for a given power budget of up to 13 dB. Beyond 13 dB, only the High-nonlinearity case can consistently provide feasible solutions.

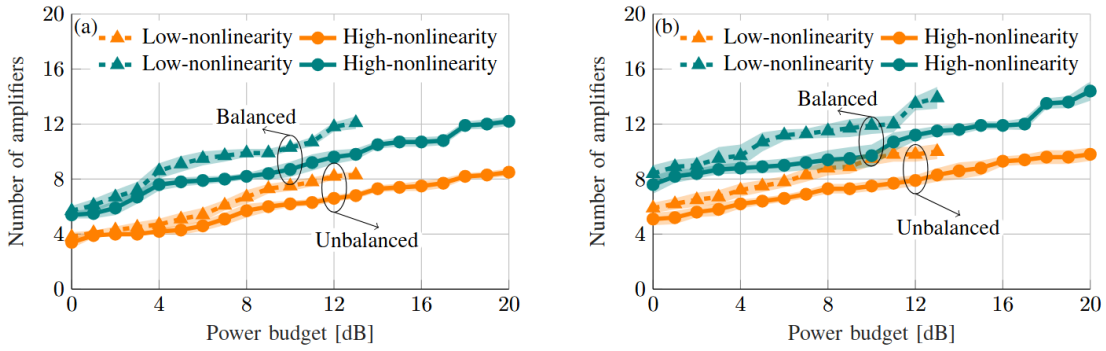

Figure 2.5-3 Average number of amplifiers versus spurs' power budget for (a) baseline and (b) extended horseshoes when using balanced and unbalanced coupler ratios while allowing low and high nonlinearities.

Figure 2.5-4 depicts an example of an extended network configuration optimized for a 20 dB power budget while considering balanced couplers for the transit node architecture (worst-case conditions in terms of both power budget flexibility and OSNR). As shown in the figure, on top of the simulated constellations, the worst-case OSNR instances (Hubs to the leaf nodes located in the furthest spur) are 14.8 dB and 17.1 dB for central SCs exceeding the 12 dB OSNR requirement. As can be seen, some optical amplifiers have taken their maximum gain value, helping the 20 dB power budget constraint to be met, which was only possible by tolerating a certain impact of nonlinear interference.

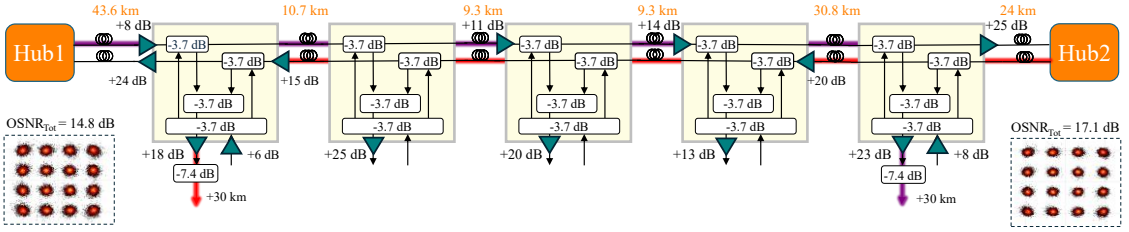

Figure 2.5-4 An illustrative optimized design of horseshoe-and-spur network and the corresponding OSNR validation for the longest Hub-to-leaf links.

2.6 EMPIRICAL LATENCY MODEL FOR IPOWDM END-TO-END CONNECTIONS

Deterministic Networking (DetNet) addresses the need for predictable network performance—including bounded latency and reliability—essential for emerging applications like AR/VR, robotic industry, and connected cars. It ensures strict Quality-of-Service (QoS) through resource allocation and deterministic data paths, yet faces challenges in modeling queuing delays accurately, particularly with bursty traffic or limited independent sources. This complexity complicates network design, especially for IPoWDM infrastructures requiring precise latency guarantees.

Time-Sensitive Networking (TSN) complements DetNet by enabling microsecond-level timing precision, critical for applications such as autonomous vehicles and robotic systems. For instance, 5G networks must maintain end-to-end latency below 50–80 ms for cloud gaming and under 120 ms for responsiveness, while industrial IoT and telemedicine demand near-lossless packet forwarding with bounded delays (e.g., under 6.25 µs per node). These requirements drive the adoption of slotted, scheduled architectures to meet stringent performance thresholds.

Existing traffic models (e.g., Gaussian/log-normal distributions) and M/G/1 queuing theory struggle to compute delay percentiles due to the lack of closed-form solutions for cumulative density functions. To overcome this, work presented in this section proposes a simple, practical queuing delay model as an upper bound for latency estimation, facilitating reliable planning for IPoWDM networks with strict deterministic constraints [Kon25]. This model is based on the well-known M/M/1 queuing model where packets arrive at the switches following a Poisson process with λ packet/s, and service times are exponentially distributed with a mean of 1/µ seconds per packet. Since this model is not realistic for Internet traffic behaviour, but M/G/1 models can be

used in those links with highly aggregated and uncorrelated traffic sources, we can try to find an M/M/1 upper bound model (let us call it the "envelope" M/M/1 model) that acts as an upper bound for latency analysis in links above 100 Gb/s.

It is worth remarking that the M/M/1 model states that the latency experienced by a packet traversing a packet switch (including transmission and queuing) can be modelled exactly with an exponential distribution. In other words, the Cummulative Density Function of the latency *D* experienced by a packet follows Eq. 4. As shown, the M/M/1 model is fully characterized by the average service time μ^{-1} and the link load $\rho = \lambda/\mu$.

$$F_D(t) = 1 - e^{-\mu(1-\rho)t}, t \geq 0 \qquad \text{Eq. 4}$$

The goal is then to find a new value $\rho_{env} \geq \rho = \lambda/\mu$ which acts as an upper bound of real traffic for latency percentiles above the median (i.e., 50%). Figure 2.6-1 shows and example of real latency histogram and such an envelope M/M/1 model whose load ρ_{env} is higher than the real one.

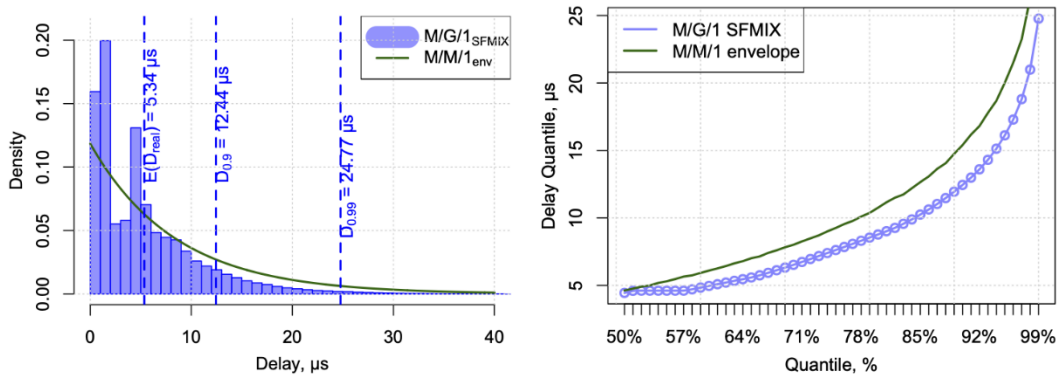

Figure 2.6-1 In green, the Probability Density Function (left) and Cumulative Density Functions of the envelope M/M/1 model. In blue, the real latency histogram (left) and Quantile functions (right).

To find such an envelope M/M/1 model upper bound, the following Algorithm maps real M/G/1 queues (with empirically observed traffic) to an M/M/1 model serving as a conservative upper bound. The algorithm is summarized in the next table:

1. **Inputs**: Simulated M/G/1 delays, average service time $E(X)$, and real load ρ_{real}.
2. **Process**:
   - Computes delay percentiles (50–99%) for M/G/1 traffic.
   - Iterates candidate M/M/1 loads (ρ_{env}) to find the smallest ρ_{env} where all M/M/1 percentiles exceed M/G/1 values [Figure 2.6-1].
3. **Regression Fitting**: Real-world traffic from AMS-IX (Amsterdam) and SFM-IX (San Francisco) exchanges informed polynomial models:
   - **SFM-IX**: $\rho_{\text{env}} = 0.50 + 0.16\rho_{\text{real}} + 0.34\rho_{\text{real}}^2$

- **AMS-IX**: $\rho_{\text{env}} = 0.43 + 0.13\rho_{\text{real}} + 0.47\rho_{\text{real}}^2$
- These equations enable network planners to convert real loads into envelope loads for latency calculations.

Figure 2.6-2 shows the mapping between real traffic load and the envelope M/M/1 model, where the envelope load has been obtained using non-linear regression models from empirical data (degree-2 polynomial fit). Simulations using the simmer discrete-event simulator confirmed the model's accuracy across traffic profiles, showing differences between simulated delays and envelope bounds below 0.5 µs (average) and 1 µs (99th percentile) at bitrates above 100 Gb/s.

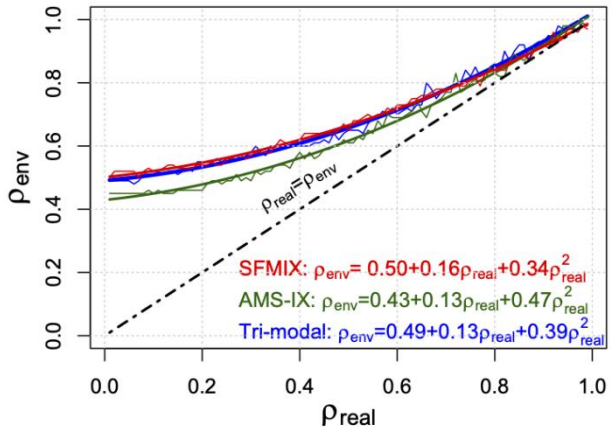

Figure 2.6-2 Mapping between real load and envelope M/M/1 load for real traffic observed in San Francisco Internet Exchange (SFMIX in red), Amsterdam Internet Exchange (AMS-IX in green) and the classical.

As an example of application, consider a 400 Gb/s link operating at 40% real load (that is carrying 160 Gb/s, 40% of the total capacity). Let us consider an average service time of 20 ns (1000 Bytes at 400 Gb/s), which translates to a service rate µ = 1 packet / 20 ns = 50 G packet/s.

According to the envelope M/M/1 model, using the SFM-IX regression model, 40% real load becomes an envelope load of 62% as Eq. 5.

$$\rho_{\text{env}} = 0.50 + 0.16 \times 0.40 + 0.34 \times 0.40^2 = 0.62 \quad \text{Eq. 5}$$

This envelope load allows us to use the M/M/1 equations for finding the median delay (percentile 50) and other percentiles like 90th and 99th percentile. To find any latency percentile, we just need to apply the CDF equation as Eq. 6 which becomes Eq. 7. For the 99-th percentile, the upper latency bound is given by Eq. 8.

$$0.5 = 1 - e^{-(50Gp/s)(1-0.62)t_{.50}}, \qquad t \geq 0 \quad \text{Eq. 6}$$

$$t_{.50} = \ln\left(\frac{1}{1-0.5}\right)\frac{1}{1-0.62}20\ ns\ = 15.8\ ns \quad \text{Eq. 7}$$

$$t_{.99} = \ln\left(\frac{1}{1-0.99}\right)\frac{1}{1-0.62}20\ ns\ = 105\ ns \quad \text{Eq. 8}$$

As shown, at 400 Gb/s and 40% load, 99% of the packets experience a latency of 105 ns or less. Thus, at very high bitrates, not only average latencies are small but also very high latency percentiles. Thus, the majority of latency is due to propagation of signals (which occurs at 200,000 Km/s over silica fibers), hence the variable component of latency (queuing) becomes negligible compared with the 5 µs per Km of fiber traversed.

This methodology has been used to implement an AI Agent (based on Google's large language model Gemini 2.5 Flash although other large language model (LLM) APIs can be used, i.e., OpenAI, Claude, etc) for calculating latency percentiles. The code has been written in open-source programming R, uses the ellmer library and is publicly available in Github[1]. As shown in the repo, the user can interface the AI Agent using natural language and the Agent finds that it has to use internal tools for calculating such latency percentile values. An example follows:

chat4$chat("Can you calculate the 0.8 latency percentile for two links in series, the first one operates at 100 Gbps link loaded with 8 Gbps, while the second one operates at 400G and is loaded with 210 Gbps? Calculate the two latency values separately")

The AI Agent (based on Gemini-2.5 Flash) then returns:

Agent:> The 0.8 latency percentile for the first link is 0.3082 microsecs, and for the second link, it is 0.116 microsecs.

2.7 COST EFFECTIVE STRATEGY FOR MIGRATION TO THE MB SYSTEMS

Traditional single-band (C-band) WDM systems, with limited spectral bandwidth (~4.8 THz), are reaching their capacity limits. To meet future scalability and capacity requirements, network operators must make the transition to MB optical systems that incorporate extended spectral ranges such as C+L (12 THz) and further to C+L+S band (up to ~20 THz) operation, leveraging additional fiber bands for higher data throughput without proliferating fiber counts. Beyond spectral expansion, MCF-based SDM introduces parallel spatial channels within fibers, multiplying capacity. The convergence of these two paradigms—MBoSDM—represents a foundational approach to future-proof optical networks.

However, migrating from legacy single-band C-band infrastructure to MB systems introduces critical challenges including increased deployment costs, complex cross-band impairments like ICXT and Inter-Channel Stimulated Raman Scattering (ISRS), and potential service disruptions during upgrades. Hence, economically viable and operationally smooth migration strategies are essential for network operators to scale capacity efficiently while minimizing CAPEX and service downtime. This section aims to synthesize insights from recent key studies and frameworks, focusing on cost-effective and performance-aware migration strategies to MB systems,

[1] J A Hernández, AI Agent for estimating upper bounds on Latency percentiles, https://github.com/josetilos/AI_agent_LatencyEstimation/, Oct 2025

emphasizing the role of advanced modeling, digital twin platforms, and incremental "Pay-as-You-Grow" (PaYG) approaches.

2.7.1 MB Optical Systems and Migration strategies

Indeed, MB systems extend the usable optical bandwidth beyond the C-band by incorporating adjacent spectral bands. In particular, the C+L band (Extended Band) approximately spans 12 THz spectrum, combining C-band (around 4.8 THz) with L-band (approximately 6 THz). In addition, the C+L+S band (Ultra-Wideband) further expands to encompass the S-band (~8 THz), reaching an aggregate bandwidth close to 20 THz. This spectral expansion, when combined with SDM through multi-core fiber technologies, enables ultra-high capacity (>100 Tb/s per fiber) transmission capabilities essential for data center interconnects and core network backbones.

However, migrating to MB systems involves many difficulties. First, there is a need for compatible equipment, including advanced amplifiers (e.g., broadband EDFAs), ROADMs capable of handling multiple bands, and transceivers adapted for wider spectral coverage. Second, inter-band impairments that degrade signal quality may appear, as for example ICXT in MCFs, which increases with core count and proximity, and ISRS causing power tilt and cross-band nonlinear penalties. Third, maintaining QoT while managing increasing nonlinear and noise effects is not an easy task. In addition, service continuity risks during migration due to unstable cross-band interactions and necessary network reconfigurations. Finally, significant CAPEX are required to deploy expanded spectrum-capable equipment and doped fiber amplifiers, especially over large network footprints. Thus, constructing migration methodologies that balance incremental investment and progressive capacity scaling, while safeguarding ongoing network operations, is not an easy task. In this sense, two different migration strategies are possible: Day-One vs. PaYG.

The **Day-One** (DO) deployment strategy entails a high initial CAPEX as it requires the deployment of MB components across the entire network irrespective of the immediate traffic demand. This approach carries the risk of underutilization during the early stages of network operation, as the fully provisioned capacity may exceed actual usage. Moreover, DO deployment offers limited flexibility in responding to uncertain and evolving traffic growth patterns.

The **Pay-as-You-Grow** (PaYG) approach represents an incremental migration strategy that introduces MB capabilities progressively based on actual traffic demand and strategic priorities. It begins with the existing C-band infrastructure and gradually upgrades selected links or nodes to extend spectrum into the L-band or S-band, dynamically adapting routing and spectrum assignments to efficiently manage resources while minimizing upfront capital expenditures. This method aligns CAPEX with revenue-generating traffic growth, offers flexibility by upgrading only critical network segments, and helps reduce redundant capacity and operational complexity. However, it faces challenges such as the complexity of service and network reconfiguration during transitions, increased cross-band impairments in mixed-band operations, and the need for advanced impairment-aware planning to maintain QoT. Key contributions in this domain

include enhanced modeling of cross-band nonlinearities and ISRS effects, advanced algorithms for resource optimization and reconfiguration in mixed-band environments, and digital twin frameworks for accurate simulation and upgrade validation. Performance evaluations demonstrate that PaYG strategies incorporating comprehensive reconfiguration, particularly the PaYG-With All Reconfiguration (WAR) method, outperform traditional Day-One (DO) deployments in both throughput and CAPEX efficiency.

2.7.2 Physical Layer Impairment Modeling and Optimization

Enhanced Generalized Gaussian Noise (EGGN) Model: Accurate impairment estimation is essential for effective MB system design and migration planning. The Enhanced Generalized Gaussian Noise (EGGN) model integrates several critical factors affecting signal quality. These include ASE noise, NLI components such as self-phase modulation, cross-phase modulation, and four-wave mixing, as well as ICXT effects, which are particularly frequency-dependent in MCFs. Additionally, the model accounts for ISRS, a nonlinear phenomenon that causes wavelength-dependent power shifts, significantly impacting the accumulation of noise across the optical spectrum. The EGGN model enables dynamic calculation of the GSNR for each channel in the network. This calculation is vital for determining which modulation formats are feasible—ranging from low-order schemes like BPSK up to high-order formats such as 64QAM. It also facilitates the estimation of achievable bit rates per channel and provides insights into the effects of partial network upgrades and mixed-band operations, helping to maintain service quality and network performance during migration phases.

Power Optimization Strategies: To manage signal power distribution effectively and reduce nonlinear penalties across different bands, power optimization strategies, like Flat Launch Power (FLP) and Flat Receive Power (FRP) are employed. These approaches aim to maintain balanced power levels, thereby minimizing impairments that degrade signal quality. Furthermore, advanced techniques such as hyper-accelerated power optimization leverage system-wide constraints to determine optimal launch powers. This method enhances overall network efficiency by adjusting power levels dynamically to suit current network conditions and anticipated traffic demands.

Impact of ISRS and ICXT on Migration: The nonlinear effect of ISRS causes power tilt across the spectrum bands, which leads to a reduction in effective GSNR and diminishes the modulated capacity of channels. This impact is especially significant during the transition from a C-band-only configuration to a MB operation, where the spectrum is expanded. Meanwhile, ICXT effects become increasingly prominent as more cores and bands are activated, and spatial and spectral layouts become denser. These impairments complicate capacity planning and require incremental network reconfiguration strategies to mitigate capacity shortfalls. Careful management of these factors is crucial to prevent service interruptions and maintain QoT during the migration to MB optical networks.

2.7.3 Digital Twin-Enabled PaYG Migration Framework

The Digital Twin (DT) concept in optical networks entails creating a real-time, synchronized virtual replica of the physical optical infrastructure. This virtual environment allows network operators to validate migration scenarios safely without any disruption to live services. By simulating different configurations such as amplifier settings, route changes, and spectrum allocations, operators can experiment within a realistic impairment model framework. This enables the quantification and understanding of cross-band nonlinear effects and the assessment of service-level impacts prior to actual physical deployment. Additionally, the digital twin framework supports automated network orchestration by integrating with SDN control planes and telemetry systems, providing seamless interaction between the virtual and physical network layers.

The framework is structured into several key layers, each serving a critical role in managing and optimizing the optical network. The Data Collection Layer interfaces directly with essential network elements such as ROADMs, EDFAs, and transceivers using standardized protocols like gNMI and SNMP, ensuring timely access to essential KPIs. The Core Digital Twin Layer embeds a detailed representation of network topology and models fundamental physical impairments using advanced models like the EGGN model, alongside path and spectrum resource modeling. Built upon this is the Analytics and AI Layer, which harnesses data analytics and machine learning techniques to advise on network optimization, resource scheduling, and upgrade planning decisions. For user and operator interaction, the Visualization and Control Layer provides interactive dashboards and orchestration APIs, facilitating intuitive network monitoring and management. Lastly, the Integration Layer ensures connectivity and interoperability with Operations Support Systems (OSS), Business Support Systems (BSS), and cloud platforms, enabling wide operational integration across different IT and network domains.

This digital twin-enabled framework is particularly powerful for supporting incremental migration strategies like PaYG, which dynamically add spectral bands in alignment with actual demand growth. By simulating such strategies in a controlled virtual environment, operators can anticipate and mitigate risks, notably reducing the likelihood of service disruptions by pre-validating complex reconfiguration and traffic grooming plans before they are applied to the live network. Furthermore, the framework enables precise techno-economic analyses by quantifying both CAPEX and OPEX trade-offs associated with various upgrade paths. This capability equips decision-makers with reliable information to make informed investment and deployment strategies, ensuring cost-effective, scalable, and resilient MB optical network evolution.

The PaYG migration framework introduces five novel algorithms, each distinguished by their level of reconfiguration aggressiveness and the underlying computational assumptions used to manage network upgrades. These algorithms are designed to incrementally expand optical network capacity by upgrading spectral bands in alignment with actual demand growth while balancing operational complexity and cost efficiency.

The **PaYG-MinG** algorithm adopts a conservative approach by assigning the minimal GSNR per band and avoids any reconfiguration of existing lightpaths. This simplicity makes it safe in terms of avoiding service disruptions during migration. However, the trade-off is that it under-utilizes network capacity and often requires earlier-than-necessary upgrades, which can lead to inefficient resource use and higher costs in the long run.

In contrast, **PaYG-WoR** assumes a fully loaded MB network upfront to estimate impairments, yet it also avoids any reconfiguration activities. While this reduces the complexity and operational challenges associated with reconfiguration, its assumption tends to overestimate impairments, again resulting in premature network upgrades and underutilization of available capacity, which impacts cost-effectiveness.

PaYG-WPR introduces partial reconfiguration by targeting only the unserved residual traffic (URT), migrating just this portion to the new band while leaving the rest intact. This method manages network capacity more efficiently and limits reconfiguration overhead. However, over time, as incremental changes accumulate, the approach may lead to spectrum fragmentation, potentially complicating future network expansions and spectrum management.

The **PaYG-WFR** algorithm implements full reconfiguration, where entire affected lightpaths are migrated to the new band. This strategy significantly frees up capacity in the legacy C-band and improves the network's future-proofing skills by optimizing spectral usage. The trade-off is a higher operational effort and reconfiguration overhead, which requires careful orchestration and management to minimize service impact.

The most aggressive and comprehensive method, **PaYG-WAR**, goes further by reconfiguring both affected and unaffected lightpaths to new bands. This maximizes spectral efficiency and delivers the best throughput and CAPEX savings among all strategies. However, it entails the highest volume of reconfiguration events and demands sophisticated orchestration capabilities to handle the operational complexity effectively.

All these algorithms depend on advanced impairment-aware GSNR calculations based on the EGGN model, explicitly considering ISRS effects. Channel capacity and modulation formats are dynamically reassigned as network links are upgraded to wider spectral bands. The migration planning optimizes the timing and selection of link upgrades based on current traffic load and spectrum occupancy. This process is executed and evaluated within a digital twin environment, enabling risk-free simulation and validation of upgrade scenarios before real-world deployment.

Figure 2.7-1 shows a snapshot of a Legacy-Agnostic Band Integration (LABI) migration flowchart concept. LABI enables seamless coexistence between existing legacy C-band services and newly activated spectral bands such as L-band or S-band. Crucially, LABI ensures that ongoing communications in the legacy bands remain uninterrupted during network upgrades, avoiding service disruptions that could otherwise impact users and operational stability.

LABI achieves this by integrating advanced multi-layer service migration and reconfiguration algorithms that maintain both QoT and bitrate continuity throughout the incremental upgrading

process. These algorithms carefully manage how traffic and service routes adapt to changing spectral conditions, ensuring that established connections sustain required performance metrics despite evolving network conditions.

A vital aspect of LABI is its inclusion of band-agnostic power optimization and modulation format assignment strategies. These techniques consider the complex nonlinear interactions that arise between different spectral bands—especially cross-band impairments like ISRS—enabling optimized signal power and modulation choices regardless of whether a flow is carried on legacy or new bands. This capability maximizes resource utilization across the entire spectrum.

Practically, LABI serves as a critical enabler for migration, as it allows operators to upgrade network segments progressively without necessitating full shutdowns or wholesale hardware replacements. This flexibility is vital for real-world deployments, where maintaining continuous service is paramount and extensive infrastructure overhaul is often cost-prohibitive or operationally unfeasible.

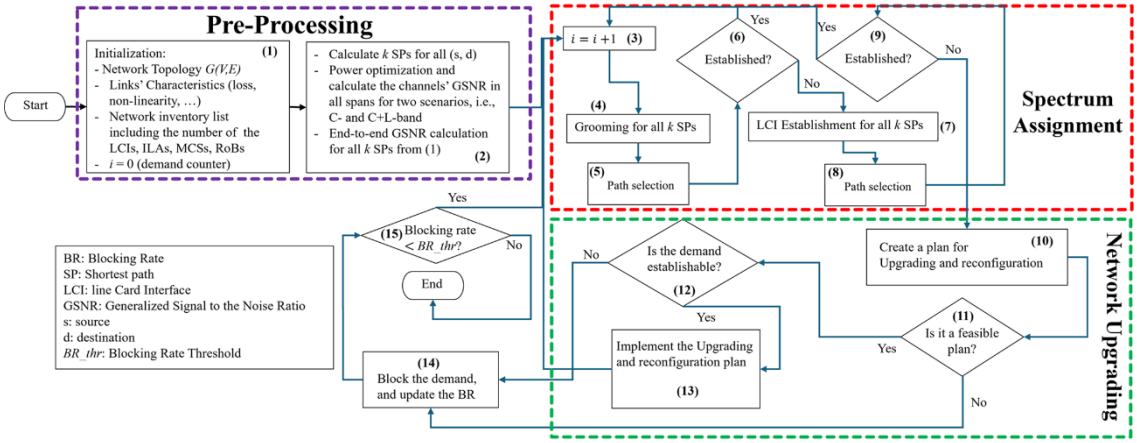

Figure 2.7-1 The flowchart of the proposed PaYG methods: LABI.

2.7.4 Simulations, results and techno-economic evaluation

The performance and economic evaluation of the proposed PaYG migration framework is based on extensive simulations conducted on three representative backbone network topologies: the USB14 (United States), SPNB14 (Spain), and JPNB12 (Japan). These topologies vary in scale and link characteristics, providing a robust testbed for analyzing the impact of migration strategies. Simulations employ advanced, flexible 64 Gbaud transceivers that support variable data rates between 100 and 600 Gbps by utilizing mixed modulation formats adapted dynamically to channel conditions. The optical spectrum is partitioned into 6 THz wide C-band and L-band segments, complemented by an 8 THz S-band with appropriate guard bands to prevent inter-band interference. The simulations accurately model full system impairments, including ASE noise, NLI, ICXT, and ISRS, along with hardware-related penalties such as filtering and aging impacts. Traffic requests follow exponential arrival and holding time distributions, realistically portraying fluctuating network demand.

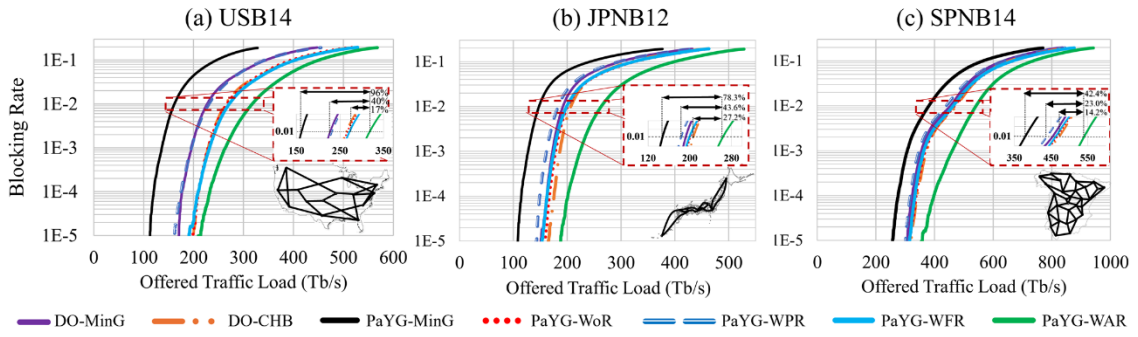

Figure 2.7-2 Established traffic of different approaches at various blocking rates in USB14, JPNB12, and SPNB14 networks.

The traffic simulation results shown in Figure 2.7-2 demonstrate that the PaYG-WAR (With All Reconfigurations) algorithm consistently outperforms all other PAYG variants and Day-One (DO) approaches in terms of throughput while keeping blocking probabilities below the critical 1% threshold. This performance advantage ranges between 15% to 40%, with the greatest gains observed in larger, long-haul backbone networks like USB14, where nonlinear effects such as ISRS tend to accumulate significantly along extended fiber spans.

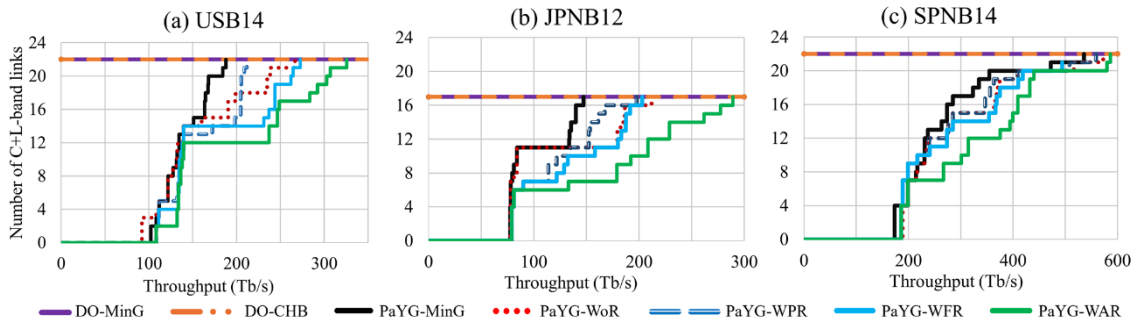

Figure 2.7-3 Number of links upgraded to C+L-band during network growth.

As shown in Figure 2.7-3, the PaYG-WAR algorithm's aggressive reconfiguration strategy efficiently frees up spectral capacity, enabling better accommodation of incoming traffic. Conversely, the PaYG-WoR (Without Reconfiguration) approach delivers moderate throughput performance but does so without incurring reconfiguration overhead, making it suitable for operational scenarios prioritizing simplicity and stability over maximum capacity utilization. PaYG-MinG (Minimum GSNR) represents a conservative strategy that significantly underutilizes network capacity due to its pessimistic assumptions about impairments, which results in premature network upgrades and consequently higher capital expenditures.

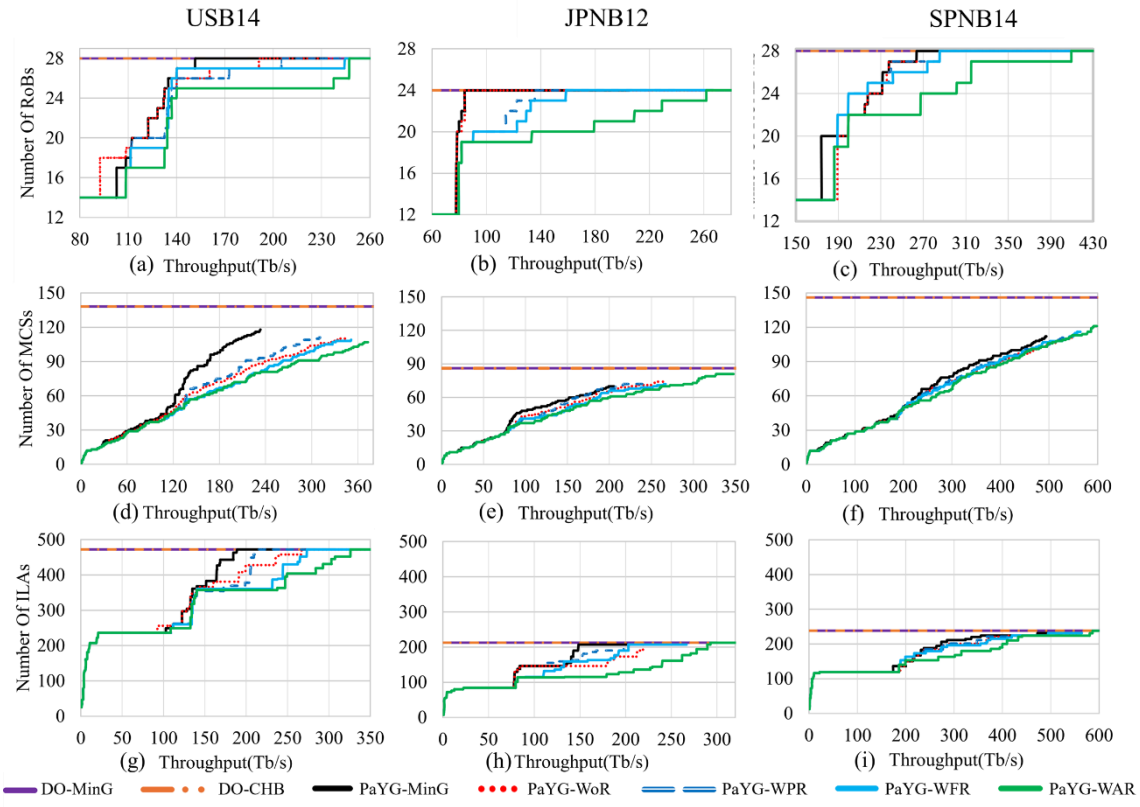

Figure 2.7-4 Number of C-band and L-band components (RoBs, MCSs, and ILAs) deployed at varying traffic rates.

CAPEX analyses shown in Figure 2.7-4 incorporate the costs of all critical network components, including ROADM-on-blade cards, inline amplifiers (ILAs), multicast switches (MCSs), and line cards (LCIs) across the C, L, and S bands. The model assumes that L-band components are approximately 20% more expensive than their C-band counterparts, reflecting current market trends and manufacturing costs.

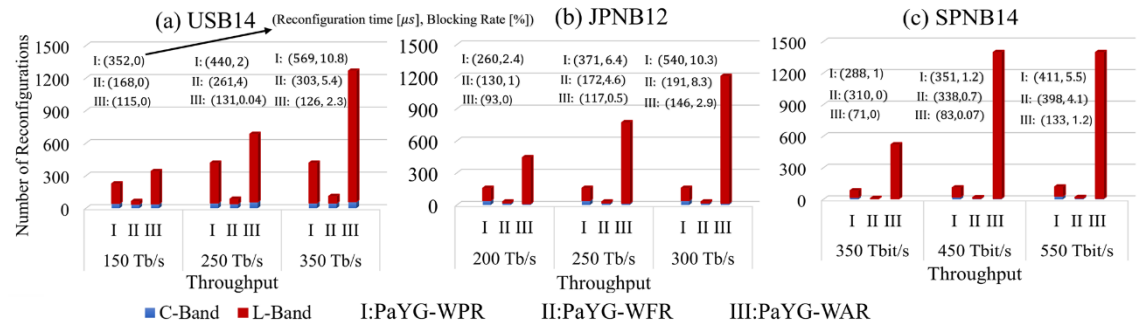

Figure 2.7-5 Number of reconfigured channels to either C-band or L-band channels for three different established traffic values.

Operational efficiency, shown in Figure 2.7-5, is further enhanced under PaYG-WAR by minimizing the total number of upgraded links involved in transitioning to C+L-band use,

reducing the physical footprint of new hardware deployments. The intelligent use of grooming techniques coupled with advanced spectrum assignment ensures that the addition of network elements such as line cards and switches remains optimized, preventing costly overprovisioning. While PaYG-WAR involves a higher frequency of reconfiguration events relative to other algorithms, this is mitigated by leveraging SDN and digital twin platforms. These technologies enable rapid, automated reprogramming of network elements, significantly reducing the risk of service disruption. Furthermore, the incremental nature of PaYG migration facilitates energy-efficient network scaling by activating additional spectral bands only when traffic demands necessitate expansion, aligning resource use with actual consumption patterns and supporting sustainable operation.

The PaYG-WAR algorithm excels in minimizing capital expenditure by deferring widespread network upgrades until absolutely necessary, thus preventing premature investments (see Figure 2.7-6). By aggressively reconfiguring existing lightpaths to the L-band and effectively releasing capacity within the C-band, this approach reduces the number of links and amplifiers requiring upgrades, culminating in CAPEX reductions of up to 70% compared to DO deployment scenarios. The PaYG-WoR and PaYG-WFR algorithms also provide significant CAPEX savings, balancing between minimized operational complexity and improved resource utilization, albeit at varying levels of reconfiguration effort.

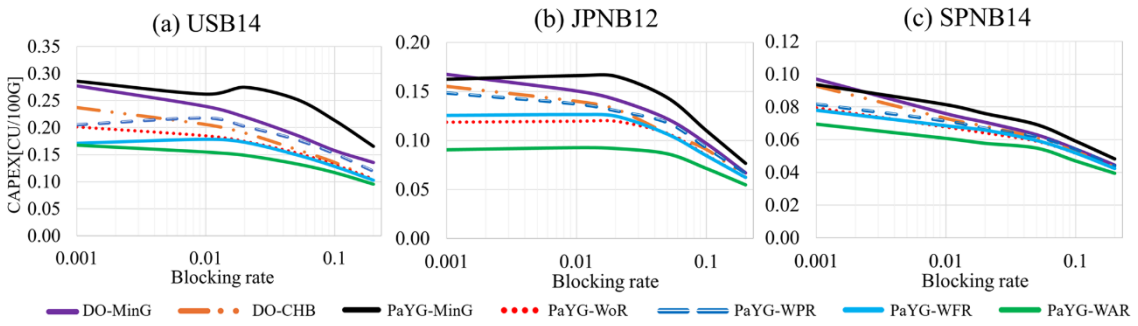

Figure 2.7-6 CAPEX [CU/100G] for different approaches.

2.7.5 Summary and conclusions

This section addressed the escalating demand for enhanced data transmission rates and increased network capacity in elastic optical networks. We propose a novel legacy-agnostic band integration approach for cost-effective and seamless C+L- band optical network migration, focusing on a “pay-as-you-grow” (PaYG) strategy to optimize CAPEX. This approach leverages advanced modeling techniques, including the EGGN model, to account for ISRS and other nonlinear effects, ensuring optimal QoT.

Through extensive simulations on real-world network topologies, we have compared the performance of various PaYG algorithms, against a traditional Day-One (DO) approach, demonstrating significant gains in throughput and CAPEX efficiency. The proposed PaYG-WAR (PaYG-With All Reconfigurations) outperforms other strategies, maximizing throughput while

minimizing the deployment of costly infrastructure components. Our results provide valuable insights for telecom operators seeking a financially viable and practical solution for migrating optical networks to C+L-band operation. Specifically, PaYG-WAR achieves 20% higher throughput compared to DO and PaYG-WoR (PaYG-Without Reconfigurations) and 17% more than PaYG-WFR (PaYG-With Full Reconfigurations), while requiring 25% less CAPEX than DO-CHB (DO-Channel Based). These findings underscore the potential of PaYG-WAR in facilitating efficient, cost-effective network upgrades to meet next-generation communication demands.

2.8 RELIABILITY BASED ON HIERARCHICAL AGGREGATION

This section provides a detailed explanation of the network and traffic engineering modelling approach for multi-homed, multi-layer hierarchical MANs, considering both the primary routers and their co-located lower hierarchical-level (HL) routers [San25, Arp24-2]. The hierarchical optical MAN under study consists of seven distinct subsets of IP nodes: N_4, N_3, N_2, N_1, N_4c, N_3c, and N_2c. These correspond respectively to the HL4, HL3, HL2, and HL1 nodes, and the co-located HL4, HL3, and HL2 nodes. The network topology can be represented as a graph G(N, L), where N denotes the set of nodes (vertices) and L represents the set of links (edges). The node set N can be further divided into hierarchical subsets Nh, where h ∈ {1, 2, 3, 4}. The overall set of nodes is the union of these subsets, excluding the co-located nodes.

Since each node is assumed to have co-located routers corresponding to all lower hierarchical levels, the co-location relationships can be expressed as:

- $N_4c = N_3 \cup N_2 \cup N_1$
- $N_3c = N_2 \cup N_1$
- $N_2c = N_1$

These relationships are illustrated in Figure 2.8-1.

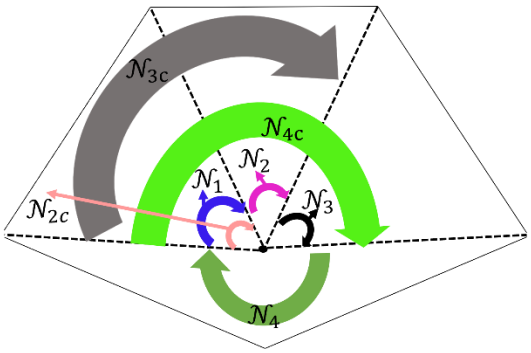

Figure 2.8-1 Venn diagram of the nodal sets.

The real-world MANs based on the geographical distribution of nodes in urban areas exhibit more complex interconnections. Specifically, HLh nodes are not always hubbed solely through the immediately lower HL(h–1) nodes. Instead, they may be physically or virtually connected to

different HL(h–m) nodes, where 1 ≤ m < h, due to the co-location of lower HL routers within higher HL levels. To illustrate this, consider the example shown in Figure 2.8-2, which depicts a hierarchical optical MAN with eight nodes—two from each hierarchical level. The corresponding optical-layer topology (layer zero) is shown in the top-left corner of the figure.

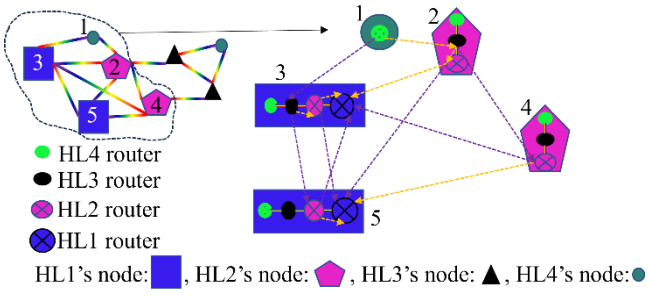

No	Node		Router		OCh bit rate	IP Flow bit rate	P/S
	Source	Destination	Source	Destination			
1	1	2	HL4	HL3	100G	50G	P
2	1	3	HL4	HL3	100G	50G	S
3	2	2	HL3	HL2	50G	25G	P
4	2	4	HL3	HL2	50G	25G	S
5	3	3	HL3	HL2	50G	25G	P
6	3	5	HL3	HL2	50G	25G	S
7	2	3	HL2	HL1	33G	12.5G	P
8	2	5	HL2	HL1	50G	12.5G	S
9	3	3	HL2	HL1	33G	12.5G	P
10	3	5	HL2	HL1	25G	12.5G	S
11	4	5	HL2	HL1	50G	12.5G	P
12	4	3	HL2	HL1	33G	12.5G	S
13	5	5	HL2	HL1	25G	12.5G	P
14	5	3	HL2	HL1	25G	12.5G	S

Figure 2.8-2 An illustrative example for hierarchical multi-layer traffic engineering in an optical MAN.

In this example, the HL4 node (node 1) has no direct optical connection to any HL3 nodes but is directly connected to a HL2 (node 2) and a HL1 (node 3). Traffic survivability is achieved through a multi-homed approach. Each HLh node's traffic demand is received by at least two HL(h–m) nodes, ensuring robustness, where 1 ≤ m < h. Following shortest-path routing, nodes 2 and 3 serve as the primary and secondary hubs, respectively, as indicated by flows 1 and 2 in the left-hand table of Figure 2.8-2. In this case, a 100 Gbps optical channel (OCh) is allocated for both the primary and secondary connections to accommodate a 100 Gbps demand. Consequently, the IP flows carried by these OChs are limited to 50 Gbps each. The co-located HL3 routers at nodes 2 and 3 serve as aggregation hubs for these IP flows.

Subsequent intra-datacenter optical connections—flow 3 from the HL3 to HL2 router at node 2, and flow 5 from the HL3 to HL2 router at node 3—each use 50 Gbps OChs to transport two 25 Gbps IP flows. This hierarchical traffic distribution continues down to the HL1 routers. Ultimately, the original 100 Gbps traffic is successfully delivered through two HL1 nodes using link- and node-disjoint paths. Notably, if a conventional 1+1 protection scheme were used, the required resources would be quadrupled. The proposed multi-homed approach therefore offers a resource-efficient and resilient design, maintaining full traffic survivability against single fiber cuts. To further ensure robustness against node failures, two important design constraints are defined:

Remark 1: The hub of a co-located router in an HLh node cannot itself be a co-located router in a lower HL(h–m) node, where h ≤ 3 and 1 ≤ m < 3.

Remark 2: The optical channel (OCh) capacity in HLh nodes (for h < 4) is calculated as: OCh capacity = R / (Np – Np–c), where:

- R is the aggregated traffic at the HL2 node,

- Np is the total number of primary and secondary paths hubbed to the HL1 node under study, and
- Np–c is the number of paths originating from HL2 or HL3 nodes co-located with other HL1 nodes and hubbed to the HL1 site under study.

These two remarks are clarified with the following examples. For **Remark 1**, consider flow 4 in Figure 3. Although node 3 is geographically close to node 2, the hub of the co-located HL3 router at node 2 must be the HL2 router at node 4—not the co-located HL2 at node 3. Otherwise, a node failure at the top HL1 would reroute only 50% of the original 100 Gbps traffic. For **Remark 2**, consider the flows entering the top HL1 router. If each OCh were configured with a 25 Gbps capacity, an HL1 failure would cause complete loss of the original traffic. For instance, the bottom HL1's failure would eliminate the 25 Gbps flow provided by flow 14. If the remaining OChs were also 25 Gbps, only two-thirds of the traffic would survive. Therefore, to guarantee full traffic recovery, the OChs corresponding to flows 7, 9, and 12 should each have a capacity of 33 Gbps.

The LAND (Link and Node Disjoint) algorithm presented in [Arp24-2] is a routing method designed to enhance the reliability of optical metropolitan networks by ensuring complete separation between the primary and backup paths used to carry traffic. It works by calculating pairs of routes between network nodes so that the two paths do not share any common links or intermediate nodes. This guarantees that if one path is disrupted -by a fiber cut, node malfunction, or other failure - the alternative path remains completely unaffected and can immediately take over the traffic. The algorithm operates on the network's topology by first determining several shortest path options between nodes using Dijkstra's algorithm. It then filters and pairs these paths to select combinations that are fully link- and node-disjoint, prioritizing those with the fewest hops and shortest total distance. In this way, LAND ensures complete physical diversity between the main and backup routes, forming the foundation of the network's survivability strategy.

Summarizing, in combination with the traffic engineering mechanisms of the 6D-MAN planner [Arp24-2], the LAND algorithm enables a multi-homed network structure where each lower-level node connects to multiple higher-level nodes through disjoint paths. This architecture allows traffic to be dynamically rerouted in case of failures while maintaining service continuity. As a result, the metropolitan network achieves high resilience and robustness, ensuring uninterrupted operation even under adverse conditions.

3 POWER-EFFICIENT AND COST-EFFECTIVE SOLUTIONS FOR ACCESS-METRO SEGMENT

3.1 DYNAMIC SPATIAL AGGREGATION FOR ENERGY-EFFICIENT PASSIVE OPTICAL NETWORKS

The spatial PON architecture proposed in SEASON leverages dynamic spatial aggregation to enhance the energy efficiency of Passive Optical Networks (PONs). The evaluation proceeded beyond isolated PON performance to assess energy saving levels achievable when integrating spatial PON with Radio Unit (RU) and Distributed Unit (DU) scaling in a real 5G O-RAN deployment.

A field trial was conducted in L'Aquila, Italy, demonstrating the first deployed multi-core fiber (MCF) infrastructure that combines spatial PONs with O-RAN-based DU scaling. The system dynamically activates and deactivates both optical line terminal (OLT) ports and spatial lanes in coordination with 5G radio resources based on actual traffic conditions.

The architecture operates through coordinated control across multiple domains: when a RAN Intelligent Controller (RIC) detects increased traffic demand, it triggers not only DU activation but also coordinates with the PON controller to enable corresponding spatial lanes, ensuring adequate bandwidth for the additional capacity. This traffic-driven orchestration integrates access and radio domains to optimize both performance and energy consumption.

Figure 3.1-1 presents the measured power consumption dynamics during experimental validation, showing three operational scenarios:

- **Baseline (blue line)**: Both RUs and PON ports continuously active, consuming 185W
- **Measured operation (yellow line)**: Dynamic system transitioning from ~170W (single RU/PON active) to ~185W (dual operation) and back
- **Achievable scenario (red dashed line)**: Theoretical operation with complete RU shutdown capability

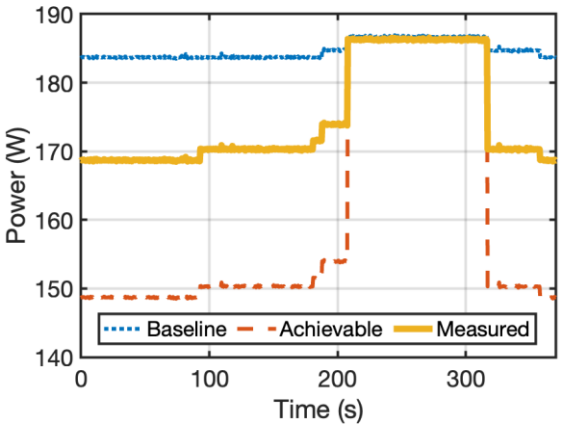

Figure 3.1-1 Power consumption dynamics during RU/DU scaling experiments.

The experimental results demonstrate concrete energy efficiency gains:

- **15W (8%) measured reduction**: Achieved through traffic-driven orchestration that deactivates one RU and its corresponding PON port during low-traffic conditions
- **35W (19%) achievable saving**: Theoretical maximum if complete RU shutdown were available in the commercial O-RAN equipment used

The system successfully scales from single-lane operation (one active RU/DU with one spatial PON core) to dual-lane configuration (two active RU/DUs with two spatial PON cores) in approximately 50 seconds, including all control-plane and data-plane reconfigurations. This automated orchestration is coordinated through a Network Service Orchestrator (NSO) that manages the O-RAN RIC, PON Controller, and Metro Controller to achieve end-to-end service delivery.

The 8% measured energy reduction represents savings per small cell under low traffic conditions. These results validate that coordinated control of spatial PON resources and radio access elements provides measurable energy savings while maintaining service quality for demanding applications such as augmented/virtual reality. The integrated architecture demonstrates that spatial PON energy efficiency gains are amplified when coordinated with dynamic radio resource management in converged access-metro deployments.

3.2 POWER CONSUMPTION OF DSCM TRANSCEIVERS IN THE METRO-AGGREGATION DOMAIN

To investigate power consumption in the metro-aggregation domain, we generalized the network topology, producing a reference construction that would allow us to study not only how power consumption behaves but also how performance is affected as network characteristics change (e.g., by increasing traffic volume). In D2.2 [SeaD2.2], we described the methodology

behind the creation of the network scenarios: we randomly generate horseshoe topologies based on a given number of horseshoes, leaf nodes, fiber spans, and traffic demands, which are all drawn from probability distributions based on realistic network examples provided by Telecom Italia [Cas24, Cas24-2].

Since DSCM slices the spectrum of a broadband signal into narrower spectral pieces (e.g., the 64-GHz-wide spectrum of the 400 Gb/s signal is divided among 16 digital subcarriers (DSCs)), we can think of and operate the individual DSCs as independent "channels", allowing us to activate or deactivate these bandwidth resources as needed. This flexibility permits us to operate the transceiver pluggables either at full capacity (i.e., as traditional single-carrier transceivers) or at an adjusted operation (i.e., with the minimum required number of subcarriers to address traffic demands). To carry out an analysis of the operating conditions of the transceiver modules, it is necessary to create a model for their power consumption. Here, the overall power consumption of a transceiver is determined by that of the application-specific integrated circuit (ASIC) and that of the optical components (e.g., laser, modulator, coherent receiver, etc.). Figure 3.2-1 shows the estimated overall power consumption of a digital subcarrier multiplexing (DSCM)-capable transceiver for different operating conditions (i.e., capacity values) as a function of the number of active DSCs. In this graph, we can observe and infer how it can be possible to lower power consumption across the network by configuring the operating conditions of the pluggables in accordance with the traffic conditions. These savings are achieved by turning off elements–or operating them at limited capacity—of the transceiver such as FEC, equalizers, clock recovery, etc. It must be noted, however, that there are common elements that are shared among all of the DSCs (e.g., analog-to-digital and digital-to-analog converters).

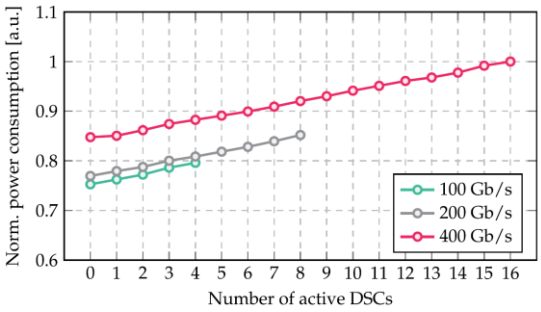

Figure 3.2-1 Normalized overall power consumption of a DSCM-capable transceiver for different operating conditions as a function of the number of active DSCs.

To reduce power consumption across the network, an adjusted operation of the transceivers is not the only measure we can adopt; a point-to-multipoint (P2MP) communication paradigm offers us the possibility of consolidating elements to establish multiple connections simultaneously, therefore reducing the number of transceivers across the network. With regards to DSCM, P2MP refers to a realization that delivers high bandwidth and simplifies the network architecture by substituting multiple low-speed P2P pluggables at the hub node(s) for a single higher-speed unit capable of communicating with multiple endpoints simultaneously. Following

this, Figure 3.2-2 exemplifies four possible alternatives to fully connect a four-sublink network topology and operate as P2MP.

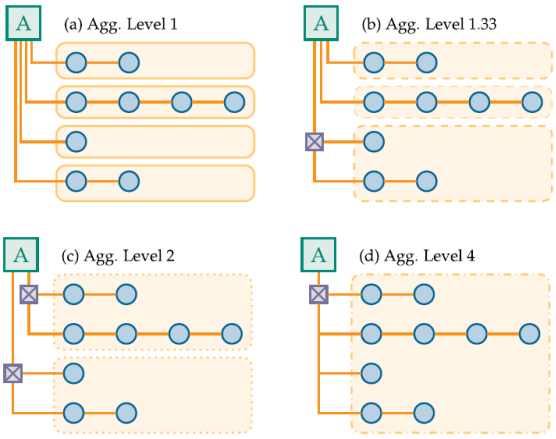

Figure 3.2-2 Example of a network scenario with four sublinks, illustrating different levels of optical aggregation: (a) optical aggregation within a single horseshoe, (b) and (c) show pairings of different subgroups of horseshoes, and (d) optical aggregation of all possible horseshoes.

These examples represent different levels of optical aggregation: from no aggregation at a horseshoe level (i.e., only optically aggregating traffic from leaf nodes within a single horseshoe) to a complete aggregation of all possible horseshoes and, therefore, leaf nodes to a single transceiver source. To quantify the level of optical aggregation, the degree to which transceiver resources are (re)utilized at the hub node, we defined a metric called level of aggregation as Eq. 9 where *Agg* represents how many sublinks are simultaneously connected to a given (set) of transceivers at the hub node, and C_s denotes how many sublink combinations are necessary to guarantee full connectivity across the network. Optical aggregation, by being a generalization of the P2MP operation of a single horseshoe, has the implication of requiring fewer transceiver pluggables at the hub node as the aggregation level increases. This, however, is not a trivial decision for the deployment of a network, as there is a delicate balance between operating conditions of the network, traffic requirements, and optical performance.

$$Agg.Level = \frac{1}{C_s} \cdot \sum_{i=1}^{C_s} Agg_s \qquad \text{Eq. 9}$$

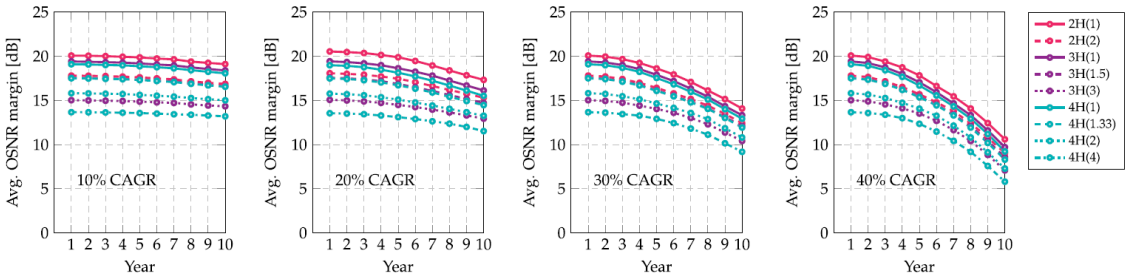

Figure 3.2-3 Average OSNR margin value as a function of the type of network and optical aggregation level over the years as a function of the simulated traffic growth. In the legend, the numbers that precede and follow the letter "H" indicate the number of horseshoes present in the network topology and the optical level of aggregation, respectively.

Figure 3.2-3 illustrates the change in optical performance, represented by the optical signal-to-noise ratio (OSNR) margin as a function of traffic growth across time and level of aggregation. These graphs evidence not only the performance degradation of the network over time, as the number and operating conditions of the transceivers change, but also how different growth rates accelerate this process. For instance, at a 30% or 40% year-over-year growth in terms of traffic volume, major network upgrades (e.g., multiple new transceiver units, and the corresponding passive optical components) will have to be considered in a 'pay-as-you-grow' model much more quickly than under more relaxed conditions (e.g., 10% or even 20% growth). Moreover, based on this analysis, it is possible to define conditions for a desired margin of operation to account for unexpected degradations, avoiding interruption of the service, and plan when and how to carry out enhancements to the network.

After ensuring the desired optical performance of the network and addressing the capacity requirements of the nodes, and by knowing the operating conditions of the transceiver modules, we can determine the relative power consumption of the transceivers according to the data depicted in Figure 3.2-1. More specifically, the calculation contemplates the number of units at the hub and leaf nodes, their maximum capacity (e.g., 100 Gbit/s or 400 Gbit/s), and the minimum required number of DSCs to cover the data demands. Then, from these conditions and Figure 3.2-1, we calculate a normalized power consumption for a point-to-point (P2P) deployment and all the possible P2MP implementations based exclusively on DSCM-capable transceivers.

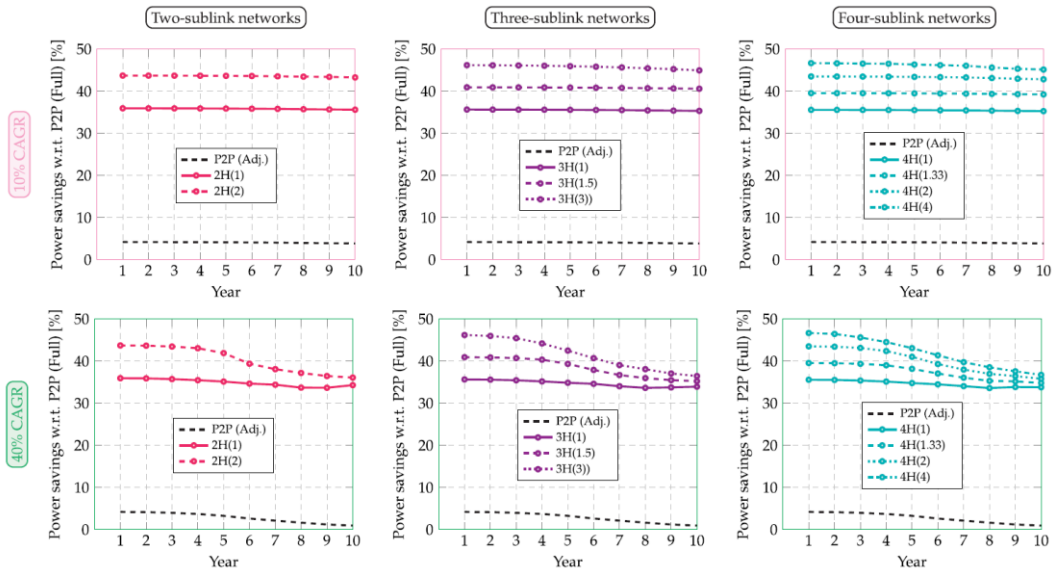

Figure 3.2-4 Relative power savings with respect to P2P (full capacity) as a function of time. Data are segregated according to type of network and year-over-year traffic growth.

To determine potential savings due to the flexible adjustment of SCs and the implementation of a P2MP communication strategy, our baseline/reference is represented by the P2P scenario, where the DSCM-capable transceivers operate at full capacity regardless of the traffic demands of the leaf nodes in the network. Figure 3.2-4 presents a summary of power savings according to size of the network and traffic growth as a function of time. The results reveal that, while flexibly controlling bandwidth resources (i.e., DSCs), the savings originating from this single measure alone are limited. On the other hand, reducing the number of pluggables by means of optical aggregation appears to significantly reduce the power footprint of the network. At this point, before going any further, there is one important disclaimer regarding Figure 3.2-4 (and by extension, Figure 3.2-1): power savings are, in this case, determined by the consumption ratio between the types of transceivers considered in the study. Here, the data for a 100 Gbit/s pluggable does not correspond to a native, produced transceiver, but rather represents a 400-Gbit/s module that's operating at a lower rate. Consequently, the relation between the higher-speed and lower-speed modules defines the limits and expectations regarding power savings [Cas24-2].

Now, to fairly evaluate P2P and P2MP—in the sense of single-carrier traditional systems and DSCM, respectively—, we first need to consider that transceivers specifically designed for 100 Gbit/s operation should be able to lower power consumption compared to Figure 3.2-1 by implementing changes in the DSP architecture and their design. Then, the question changes to: what should be the reference power consumption of a native, DSCM-capable 100 Gbit/s unit for a P2MP deployment to consume as much power as a P2P implementation utilizing single-carrier P2P modules [Adt22].

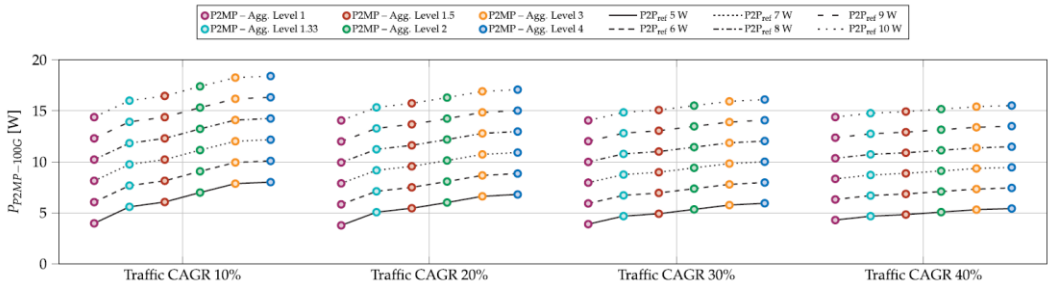

Figure 3.2-5 Maximum possible power consumption for native DSCM-capable 100 Gbit/s transceivers to ensure neutral consumption compared to P2P single-carrier 100 Gbit/s units.

Figure 3.2-5 summarizes our results, creating guidelines for the design of effective transceivers to deploy energy-efficient networks. These data points represent a comparison at a transceiver level (i.e., ignoring the power consumption of all other components in the horseshoe networks) between P2P with single-carrier units and P2MP with DSCM. By knowing the power consumption of the DSCM-capable 400 Gbit/s module and of the single-carrier pluggables, we can determine the maximum amount of power that a 100 Gbit/s transceiver can consume and still deploy DSCM networks that are as energy-demanding as traditional P2P ones. In this diagram, the multicolor markers indicate the levels of aggregation, while the line styles denote the power consumption of the reference 100 Gbit/s P2P transceivers. This investigation demonstrates how both optical aggregation and the reduction of transceivers can relax the requirements for the lower-speed DSCM-based modules, while also illustrating how high traffic volume reduces the effectiveness of P2MP, as the network practically operates as a P2P system.

3.3 COMPARISON OF DIFFERENT TRANSPORT SOLUTIONS FOR FH IN THE RAN IN DIFFERENT GEOTYPE SCENARIOS AND DIFFERENT TIME HORIZONS

With the evolution towards 6G, the fronthaul segment is expected to become even more bandwidth demanding. As such, coherent technologies are promising candidates to deliver the required levels of bandwidth. However, adopting traditional Point-to-Point (P2P) and Wavelength Division Multiplexing (WDM) approaches can prove costly and inflexible when scaling to meet the increasing demands of 6G networks.

In this subsection, an update of the study on the comparison of different transport solutions to carry the front-haul (FH) flows in the RAN between the RU located at mobile site and DU function located in the Central Office (CO) is reported. Preliminary results of the study were already included in subsection 3.4 of D2.2. New results have been obtained with updated cost parameter values achieved in the final version of the SEASON Cost model (see Section 6). Additional views and a new analysis of the results are also provided. A description of the scenario and model is provided here for clarity, although all the details on scenario setting can be found in the previous deliverables D2.1 [SeaD2.1] and D2.2 [SeaD2.2]

The analysis has been made in different conditions of macro and small cell deployments (four geotypes configurations reflecting areas with different population densities and socio-economic context) and in two different time frames, medium- and long-term (MT and LT, respectively), each of them with specific configurations and requirements of radio layers at mobile sites. The assumption is that four operators share the same mobile site and the transport service is one for all of them. In Table 3.3-1 the radio carrier parameters applied in the study are specified. The table shows for both category of cells (macro and small) each Bands range identifier, the corresponding Bands range, the number of bands within the Bands range, the carrier width in MHz, the MIMO layers and the resulting front haul datarate in Gb/s (FH Gbps). The values of FH data rates given Table 3.3-1 are for Split 7.2 option [ORA21].

Table 3.3-1 Assumptions on radio carrier parameters in each bands range for macro and small cells.

Cell type – BR id.	Bands Range	N. of Bands	Carr. Width [MHz]	MIMO layers	FH data rate [Gb/s]
Macro - BR1	Sub GHz	4	10	16x16	4.32
Macro - BR2	1-3 Ghz	4	20	16x16	8.64
Macro - BR3	3-7 Ghz	2	100	8x8	21.16
Small – BR1	3-7 Ghz	3	200	4x4	10.58
Small – BR2	7-15 GHz	1	2000	4x4	211.6
Small – BR3	24-26 GHz	2	1000	4x4	105.8

In Table 3.3-2, the number of carriers for each geotyope and for the two timeframes (MT and LT) are provided. For each carrier, three cells are considered in macro cell sites (i.e., three-sectorial sites), while small cell sites are single-cell sites. The data rate transportation need of a macro cell site or a small cell site is obtained combining of FH data rates in Table 3.3-1, number of carriers in each radio site as in Table 3.3-2 and number of cells per carrier. Please note that FH streams originate from each individual RU, and as such they can be transported separately or aggregated, depending on the solution chosen for transport between the site and the central office.

Table 3.3-2 Number of site carriers for different geotypes and the two timeframes.

Cell type Band Range Identifier	Dense Urban		Urban		Suburban		Rural	
	MT	LT	MT	LT	MT	LT	MT	LT
Macro - BR1	8	8	8	12	8	16	4	12
Macro - BR2	12	16	8	16	8	12	4	8
Macro - BR3	8	8	4	8	4	8	4	4
Small – BR1	8	12	4	8	0	4	0	0
Small – BR2	0	4	0	4	0	0	0	0
Small – BR3	4	8	4	4	0	4	0	0

The FH signal transport problem consists in carrying the streams from each RU to the Distributed Unit (virtual, vDU, or physical, like a conventional base-band unit) located at the CO. The size of

the stream depends on the channel width and the number of MIMO layers, depending in turn on the type of RU. As already said, the values of FH datrates are given in Table 3.3-1 for Split 7.2 option [ORA21]. These flows can be transported separately or aggregated, in the last case all or subset of them can be aggregated.

The three options taken into account for the transport of FH flows are illustrated in Figure 3.3-1 which refers to a macro cell site. On the left we have the radio mobile site with its multiple radio carriers active, each of them made of three cells (if the site was small cell, it would be one), one RU per cell. The central office is where the RAN Distributed Unit (DU) function is located, which is typically run on servers in virtualized mode (vDU) [ORA21]. The DU processes all FH flows and forwards them to the Centralized Unit (CU) of the RAN. Typically, a router performs the aggregation function of multiple FH flows. Between the mobile radio site and the central office, a transport system based on fiber connectivity is required.

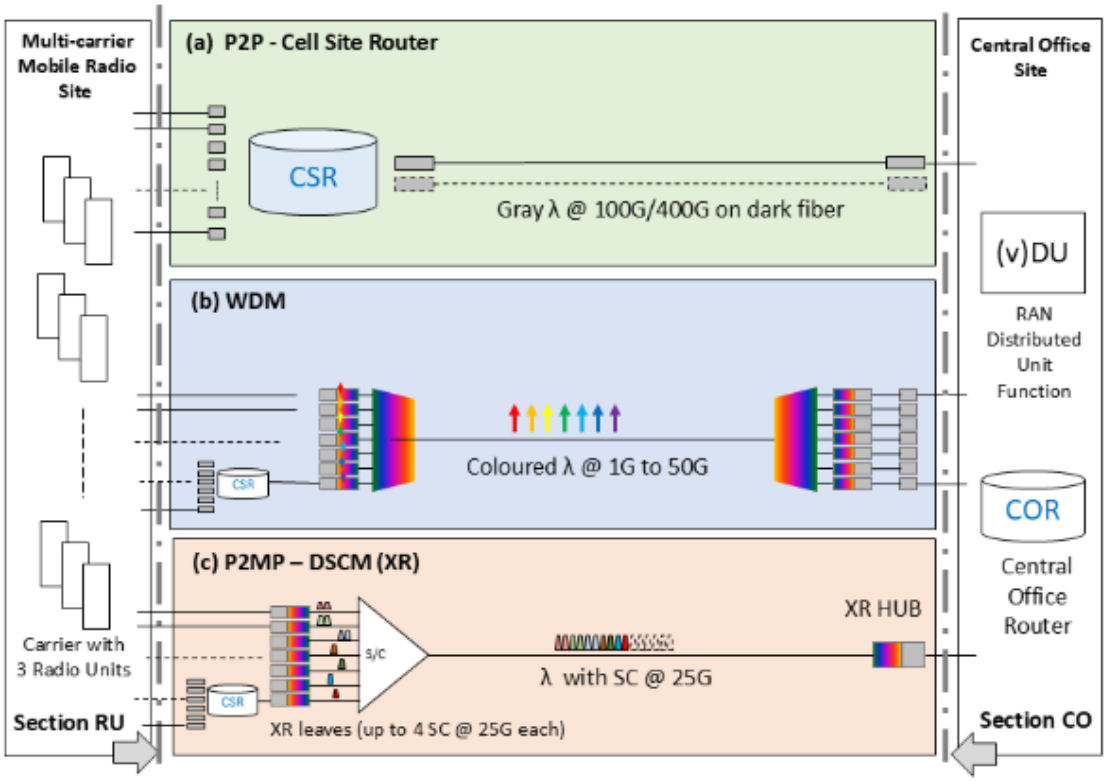

Figure 3.3-1 Alternatives solutions for FH traffic transportation examined.

The first option (inset (a) of Figure 3.3-1, P2P - Cell Site Router) uses a layer 3 aggregator switch for all FH flows coming from the RUs and exchange the traffic with the CO trough circuits made with grey transponders and fibers. A dedicated couple of fiber (one for each direction of monodirectional transmission) is used for each circuit. Depending on the traffic collected and the data rate of the grey transceivers used, one or more circuits could be required.

The second option considered, uses a WDM system (typically CWDM, inset (b) of Figure 3.3-1, WDM) to optically multiplexing the signals from the RUs onto a single line system and carry the

traffic to the CO. This system proves adequate when the data rate of the FH flows makes good use of the WDM system; this happens when the single flows are in the order of one Gb/s or higher, but it is inefficient for the transport of lower bit rate flows. To overcome this drawback, it is possible to insert a low-capacity aggregator router that aggregates the low bit rate flows into a single circuit at 1 Gbps or higher to be transported over the WDM system.

Finally, the third option considered (inset (c) of Figure 3.3-1, DSCM P2MP DSCM) uses coherent transceivers that implement digital subcarrier multiplexing and which allow for point-to-multipoint networking with a hub at the CO and leaves at the radio base station. Since the granularity of DSCM flows is 25 Gb/s, packet multiplexing of low-bit-rate FH signals is necessary, while high-bit-rate FH flows (>10 Gb/s) can be efficiently transported by a single subcarrier or by two subcarriers if data rate is greater than 25 Gb/s (according to hypothesis there are no FH flows from a single RU exceeding 50 Gb/s). This is the reason for the presence of the multiplexing router at the radio base station.

We can summarize that in the first case we have a solution based on full-packet multiplexing and gray optics, in the second we have analog optical multiplexing (WDM), and in the third we are dealing with optical digital subcarrier multiplexing. In the second and third options, however, partial-packet multiplexing of low data rate FH flows becomes necessary to optimize the utilization of optical signals. In accordance with the hypotheses presented, the results shown in Figure 3.3-2 were obtained. These results update those reported in [SeaD2.2] and [Mar25] as they apply the cost values presented in this document, which were not previously available.

The costs used for the economic evaluation are those reported in Table 6.1-2 for the grey transceivers, in Table 6.1-3 for the WDM transponder, Table 6.1-4 for DSCM coherent transceivers, and Table 6.1-6 for the Router respectively. A factor of x4 is applied in the last case as router is expected to cost four times a layer 3 switch for the same capacity. The WDM multiplexer in case (b) was assigned the cost of 0.24 CU. For the DSCM modules, the best and worst case scenarios were evaluated using the minimum and maximum cost values from the ranges provided in Table 6.1-4. The results presented in Figure 3.3-2 demonstrate the techno-economic performance of the three transport solutions across different deployment scenarios. The analysis reveals several key insights regarding both CAPEX and energy consumption patterns across the four geotypes considered.

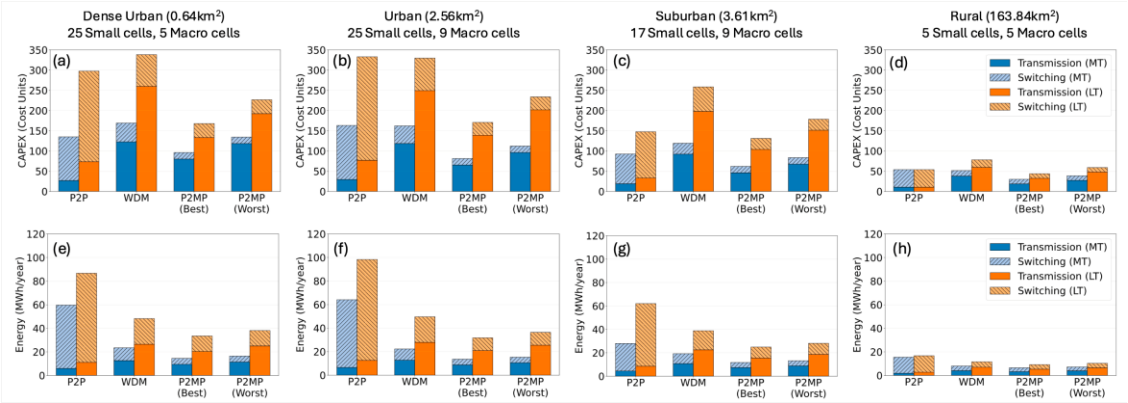

Figure 3.3-2 Capex and annual energy consumption comparison.

The P2P solution consistently exhibits the highest CAPEX across all geotypes, particularly in dense deployments, where it reaches approximately 330 c.u. in the dense urban scenario. This increased cost stems primarily from the requirement for dedicated fiber pairs and multiple grey transceivers for each fronthaul link. The WDM solution offers substantial CAPEX savings compared to P2P, with reductions ranging from 40-50% across different geotypes. This efficiency results from the optical multiplexing capability that reduces the number of required fibers and leverages economies of scale in the optical transport layer.

The P2MP DSCM solution demonstrates the most attractive CAPEX profile, especially when considering the best case scenario with lower cost assumptions for DSCM transceivers. In this configuration, CAPEX reductions of up to 60% compared to P2P are achieved in dense urban scenarios. Even in the worst case with higher DSCM transceiver costs, the P2MP solution remains competitive with WDM across most deployment scenarios. The rural geotype shows the smallest absolute differences between solutions, reflecting the lower traffic aggregation requirements and reduced complexity in sparse deployments.

Energy consumption patterns largely follow the CAPEX trends, with P2P solutions requiring the highest power consumption due to the larger number of active transceivers and Layer 3 switching equipment. Dense urban deployments show P2P energy consumption reaching approximately 85 MWh/year. WDM solutions achieve energy savings of 30-40% compared to P2P across all geotypes, primarily through the elimination of multiple grey transceivers and the efficient optical multiplexing approach. The P2MP DSCM solution demonstrates the most favorable energy profile, with consumption reductions of up to 50% in the best case scenario. This efficiency results from the point-to-multipoint architecture that reduces the overall number of active optical components and eliminates the need for multiple point-to-point connections.

The comparison between Medium Term (MT) and Long Term (LT) scenarios reveals the impact of evolving radio requirements. The LT scenarios generally show higher absolute costs and energy consumption due to increased capacity requirements from additional radio carriers and higher MIMO configurations. However, the relative performance rankings between transport solutions remain consistent across time frames. Dense urban environments show the largest

absolute differences between solutions, making solution selection most critical from both CAPEX and OPEX perspectives, while rural deployments show minimal differences between solutions, suggesting that solution selection may be driven by factors other than pure cost optimization.

These results indicate that P2MP DSCM solutions offer the most promising techno-economic performance, particularly in high-density deployments where the benefits of point-to-multipoint architecture and digital subcarrier multiplexing are most pronounced.

3.4 Comparison between WDM P2P and P2MP transceivers within a Full Mesh for FH network

The RAN fronthaul physical layer, essential for scalable 5G deployments, interconnects virtualized distributed units (vDUs) with remote radio units (RRUs) over robust optical transport. As detailed in Section 5.1.1 of D3.3, this layer integrates transceivers, optical fibers, splitters, amplifiers, and optical line system components to enable high-speed bidirectional transmission through WDM and 400G modules.

Building on this foundation, we designed, implemented, and validated an upgraded optical network (Figure 3.4-1) featuring a bidirectional ring topology across two hubs and two leaf sites. By employing P2MP-capable XR pluggables and 2×2 splitters, we reduced the number of transceivers by one-third, minimized insertion losses, and extended reach without EDFAs. The architecture is modular and scalable, supporting additional nodes and transceivers with minimal changes.

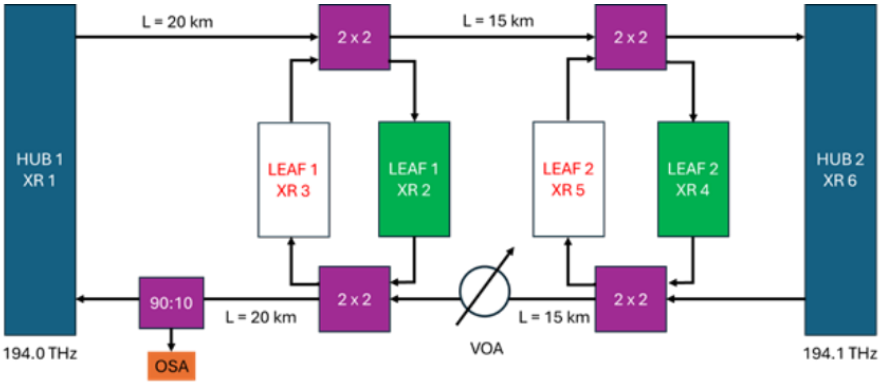

Figure 3.4-1 Fronthaul architecture using XR transceivers.

A key result is the 32.8% reduction in energy consumption, achieved by cutting power use from 274 W to 184 W, thanks to XR transceivers' higher optical output and fewer active components. This validates the potential of XR-based fronthaul transport to deliver sustainable, cost-efficient, and scalable solutions for next-generation mobile networks.

Table 3.41 summarizes the cost comparison between the two fronthaul architectures. By adopting XR transceivers and simplifying the optical design, the total equipment cost is reduced from **80.7 kUSD to 48 kUSD**, primarily due to the elimination of EDFAs and the lower number of standard transceivers required. This confirms that the proposed XR-based solution not only enhances performance and energy efficiency but also significantly decreases capital expenditure.

Table 3.3.41 Cost Comparison between the two architectures.

Components	Present previous setup	Present new setup	Single cost	Total Cost without XR	Total cost with XR
Transceivers	12	0	6000$ ~	72.000$	0$
Transceivers XR	0	8	6000$ ~	0$	48.000$
EDFA Amplifier	3	0	2500$ ~	7.500$	0$
Splitter	12	4	100$ ~	1.200$	400$
Total cost				80.700$	48.000$

3.5 VALIDATION OF 400 GB/S WITHIN FH NETWORK

The network architecture we developed for the Fronthaul segment features an arrangement where vDUs and RRUs are linked through packet switches and Optical Line Systems (OLS), employing 400G QSFP-DD transceivers to enable high-speed data transfer. This setup consists of a bidirectional fiber ring combined with a logical mesh topology.

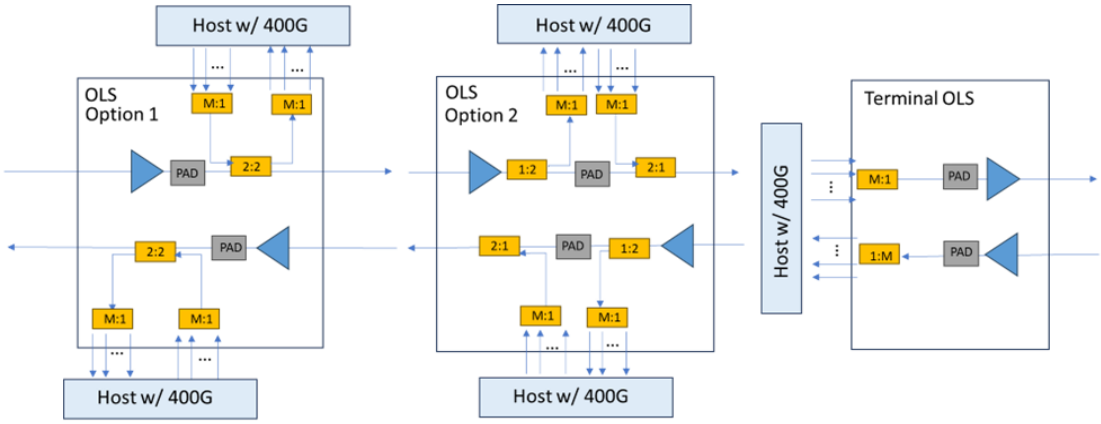

Figure 3.5-1 Filterless OADM and OLT nodes (OLS).

Figure 3.5-1 illustrates the internal layout of the OLS. Its reliance on straightforward, non-selective splitting and the versatility of pluggable EDFAs—combined with wavelength control via tunable transceiver lasers—enables channels to be added or dropped on the fly. This makes the design both adaptable and scalable to diverse network requirements. In essence, the architecture represents a sophisticated optical network optimized for high-speed throughput, consistent signal levels, and flexible wavelength assignment. Physically, it forms an open-ring (or

horseshoe) topology, while logically all nodes are fully meshed through dedicated wavelength paths.

Because of the drop-and-continue functionality, an N-node ring uses N(N–1)/2 wavelengths. The "Terminal OLS" sits at the end of the open ring; its amplifiers and pad attenuators are optional. Each OADM OLS employs a 2:2 splitter for in-line add/drop operations on the trunk fiber and M:1 splitters to merge or separate drop/add channels for the transceivers. A PAD attenuator balances the power between added channels and the through channels. Two configurations are considered:

- Option 1: a single 2:2 coupler handles both in-line, colorless add and drop functions
- Option 2: separate 1:2 splitters for the add and drop portions of the in-line path

Both single- and dual-amplifier pluggable modules include automatic gain control (AGC) and automatic power control (APC) to stabilize output under varying inputs. The single-amplifier version additionally provides gain flatness control and a tunable output power range of 15–17 dBm, whereas the dual-amplifier model offers a fixed 15 dBm output. The comparison between both amplifiers is presented in Table 3.5-1.

Table 3.5-1 Comparison between Single and Dual EDFA amplifier.

Property	Dual EDFA	Single EDFA
Channels	8	15-48
AGC	Yes	Yes
APC	Yes	Yes
Flatness Control	No	Possible
Output Power	15 dBm	15-17 dBm
NF Typical	6.5 dB	6.5 dB

The reference performance of coherent pluggable 400G transceivers, based on the latest OpenZR+ specifications, includes the high-level parameters essential for end-to-end feasibility assessment reported in Table 3.53.5-2.

Table 3.53.5-2 Comparison between transceivers 400G Coherent and 200G Coherent.

	400G Coherent	200G Coherent
TX o/p	-10 dBm / 3 dBm	-9 dBm / 3 dBm
TX OSNR	34 dB	34 dB
RX OSNR	24 dB	16 dB
RX i/p	-12 dBm	-18 dBm
CD	20000 ps/nm	50000 ps/nm

The low-power variant tops out at –10 dBm launch power, whereas the high-power model reaches +3 dBm by incorporating an optical amplifier and, if desired, a tunable noise-reduction filter. As expected, the fixed-gain nature of the amplifier and the modest output of the entry-

level transceiver impose the greatest constraints. Achieving adequate OSNR and receiver sensitivity at these high data rates also presents a significant challenge.

In the filterless OADM layout described above, a fixed attenuator (node pad) is needed to match the amplified pass-through channels to the locally inserted low-power channels. That extra insertion loss cuts into the overall power budget and degrades the receive margin. While link-budget calculations with low-power optics show the concept can work, separating the drop and add paths—as in OLS option 2—prevents the pad from reducing the dropped signal level. By decoupling the equalization pad from the drop function, we avoid that penalty, although this requires an extra splitter in the main fiber line. Under this scheme, the achievable span budget is (the calculation details are presented in Table 3.5-3 and Table 3.5-4):

- A four-node ring using six wavelengths
- 3 dB of fiber loss allowance per span
- 24 dB of EDFA gain

Table 3.5-3 Calculation details for the 4 nodes, low-power case.

Field	Value	Unit
Power EDFA	17.0	dBm
TX power	-10.0	dBm
Span Loss	24.0	dBm
Nodes	4.0	
N wavelengths	6.0	
Max Pch	9.2	dBm
Pch	3.5	dBm
Gross OSNR	26.2	dB

Table 3.5-4 Additional calculation details for the 4 nodes, low-power case.

	OLS opt 1	OLS opt 2	Unit
Net OSNR	24,7		dB
Node PAD	17.5	14.0	dB
Net Budget	3.0	3.0	db
RX Power	-23.5	-5.9	dBm

By opting for the high-power (+3 dBm) transceiver module, we can shrink the node-pad attenuation needed for channel equalization. This in turn lowers the required EDFA gain, boosts OSNR, and extends span lengths. Under this scenario, the link budget supports (the calculation details are presented in Table 3.5-5 and Table 3.5-6):

- Six nodes with 15 wavelengths in use
- 10 dB of fiber loss allowance per span
- 21 dB of EDFA gain

Key calculation points:

- A 17 dBm output lets each channel sit at roughly +3.5 dBm after amplification.
- Gross OSNR is reduced only by the small contributions of net OSNR degradation, gain ripple, and PDL.
- Both the received power and OSNR margin meet 400G-coherent receiver requirements over six hops.

Because the drop-and-continue architecture still falls below the receiver's sensitivity threshold when using the filterless single-coupler design (option 1), the dual-splitter layout of option 2 remains necessary at 400 G.

Table 3.5-5 Calculation details for the 6-nodes case.

	Value	Unit
Power EDFA	17.0	dBm
TX power	3.0	dBm
Span Loss	21.0	dBm
Nodes	6.0	
N wavelengths	15.0	
Max Pch	5.2	dBm
Pch	3.5	dBm
Gross OSNR	27.2	dB

Table 3.5-6 Additional calculation details for the 6-nodes case.

	OLS opt 1	OLS opt 2	Unit
Net OSNR	24,7		dB
Node PAD	7.5	4.0	dB
Net Budget	10.0	10.0	db
RX Power	-17.6	-10.1	dBm

Six nodes seems to be the limit for 400G operation. However, up to 10 nodes can be feasible at 200Gb/s with the following characteristics (the calculation details are presented in Table 3.5-7 and Table 3.5-8):

- 10 nodes (15 wavelength used)
- 14 dB available for the fiber on each span
- 25 dB gain EDFA

Table 3.5-7 Calculation details for the 10 nodes case.

	Value	Unit
Power EDFA	17.0	dBm
TX power	3.0	dBm
Span Loss	25.0	dBm

Nodes	10.0	
N wavelengths	45.0	
Max Pch	0.5	dBm
Pch	0.5	dBm
Gross OSNR	25.5	dB
GV	10.5	dB
PDL_P	0.9 - 1.3	dB
GV_P	2.6	dB

Table 3.5-8 Additional calculation details for the 10 nodes case.

	OLS opt 1	OLS opt 2	Unit
Net OSNR	16,9		dB
Node PAD	7.5	4.0	dB
Net Budget	14.0	14.0	db
RX Power	-25.8	-18.3	dBm

The received optical power is somewhat marginal, so using a more powerful amplifier — such as one with an 18 dBm output — would be preferable. The estimations made for the 16/48 channel scenarios assumed that a passive gain flattening filter (GFF) is present, which is not yet available in current pluggable EDFA modules. Performance enhancements could also be achieved by integrating variable gain amplifiers. These amplifiers offer dynamic gain control, allowing them to adjust to fluctuations in input power and other changing network conditions. This feature adds flexibility to the system, removing the need to manually adjust the optical path length or insert attenuators. Additionally, our findings suggest that deploying high-power transceivers contributes to better scalability in the network design.

4 SMART EDGE NODES FOR PACKET/OPTICAL INTEGRATION WITH COMPUTING RESOURCES

4.1 RAILWAY MOBILE SERVICE COVERAGE SCENARIO DEFINITION FOR DPU APPLICATION USE CASE

The scenario analyzed was presented [SeaD2.2] and concerns the coverage of a railway line or a motorway section where the high-bandwidth, low-latency and high-reliability mobile radio connectivity service must be provided. Due to the inherent topography features of this service, it must be provided within a very narrow but continuous surface along the railway or motorway line. In terms of scenarios defined in the SEASON context, this means considering an additional geotype that is added to the four already identified in D2.1 (i.e., rural, suburban, urban and dense urban, the ones considered in the study reported in subsection 3.3) [SeaD2.1].

The study investigates the application of smart NIC or Data Processing Units (DPU) cards in servers (in an O-RAN and could RAN perspective) for a Railway 5G/6G scenario coverage use case as an alternative to traditional Radio Access Network (RAN) architectures that use legacy equipment for RAN (e.g., aggregated BBU) and transport. The overall objective is to analyze and evaluate the technical feasibility and potential economic benefits of using DPUs in this specific context.

The study is centered on a specific scenario defined "Railway geotype". This scenario, shown in Figure 4.1-1, involves a chain of mobile stations covering a segment of railway line, starting at a "tail TIP" site and ending at a "tail END" site. The tail END site is connected to the operator's Central Office. Each station hosts two radio cells (Remote Units or RUs) primarily propagating in opposite directions along the railway pathline. The scenario considers trains moving at speeds up to 300 Km/h, with the possibility of more than one train being present on the same segment simultaneously.

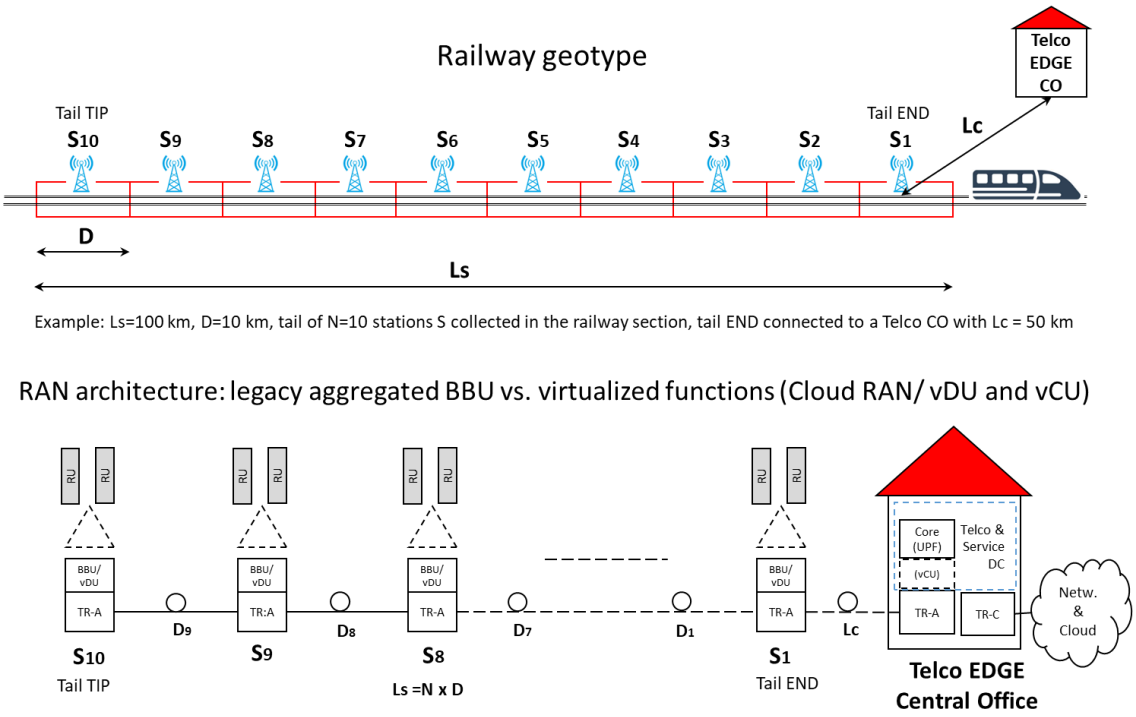

Figure 4.1-1 Railway scenario representing a segment covered by a chain of mobile stations whose generated traffic is directed to Central Office where Telco functions are available. On top geometric model, in the part below an high level view of network topology (mobile sites and Central office) and architecture (RAN, mobile core and transmission capabilities. i.e., TR-X equipment).

The study encompasses several key areas: scenario definition, simulations, lab experimentation, and techno-economics.

4.1.1 Scenario Definition and Transport Requirements

The scenario definition phase involved modeling the Railway geotype with parameters such as the number of sites and distances. Radio parameters were also defined, as were the assumptions for RAN and transport systems and architecture. For the reference scenario analyzed below, the geo and plant parameters are the following ones. The length of railway segment covered (Ls) is 50 km, the number of sites (Ns) placed uniformly along the segment is 10 and this means an inter-site distance (D) of 5 km. Distance between Tail TIP (S10) and Tail END (S1) turns out to be 45 km. Distance between tail END and the Telco Central Office is assumed to be 50 km. Obviously, a range of analyses can be carried out, in which these parameters can be changed to analyse the sensitivity of the solution to variations in the geometric configuration of the geotype.

High-level transport requirements were estimated based on radio characteristics. The radio setup assumes 3.7 GHz frequency, 100 MHz bandwidth, and 4 MIMO layers per cell. Massive MIMO is not considered to bring significant benefit in this specific railway context.

Based on this radio configuration, estimated data rates for different segments of the network were derived:

- Fronthaul (FH), connecting the RU to the Distributed Unit (DU), is estimated to require 15-50 Gb/s, depending on the chosen functional split option between the RU and the DU.
- Midhaul (MH) (connecting DU to Centralized Unit - CU) and Backhaul (BH) (connecting CU to the core network) are estimated at 1-4 Gb/s per cell at full load. However, it is noted that these values are likely significantly lower in real conditions, potentially only a few hundred Mb/s, based on simulations. The total MH/BH traffic collected by all radio stations in the entire segment is estimated to be less than 10 Gb/s, assuming one to a few trains in transit at a given time and that up to 10 mobile stations take part of the railway segment.

It should be noted that the assumptions on transport requirements reported above (50 Gb/s maximum for FH, 4 Gb/s maximum for BH/MH, for each site) should not vary in the case of 6G since the available frequencies and bandwidths for this type of service are constrained by the spectrum assignments to the operators (see Section 2.4 of deliverable D2.1 [SeaD2.1], where a detailed analysis is made for Italy and Spain) and will not vary significantly in the next ten years. The use of frequencies in the 26 Gb/s band that could lead to significantly higher x-haul data rates (FH > 100 Gb/s and MH/BH > 10 Gb/s) is not suitable for the type of service under analysis, or would be at most in the limited parts of the railway located in urban contexts with distances between antennas in the order of hundreds of meters, not a few km as realistic for a railway line that for most of its development crosses long stretches of suburban and rural areas.

4.1.2 Evaluated Network Architectures

The study compared three distinct RAN network architectures, two taken from the legacy and one proposed as alternative which includes innovations proposed by SEASON.

The first one, named Baseline 1 and shown on top part of Figure 4.1-2, assumes a RAN architecture employing standard BBU and Cell Site Router (CSR) for aggregating and forwarding the traffic hop by hop from mobile sites to the Central Office. The second one is Baseline 2 depicted in middle part of Figure 4.1-2 and relies on enhanced BBU which, thanks to networking capability added in the BBU (grey card), avoids the use of routers for traffic aggregation and transport. A router is, however, provided at the tail END site to aggregate the traffic of the entire railway segment and exchange it with the central office. This router could also be avoided if the BBU networking card had all the necessary capabilities and capacity to handle the entire aggregate traffic of the segment.

The third and innovative RAN architecture is shown in the middle part of Figure 4.1-2 and leverages virtualization for RAN functions and use of systems and solutions proposed by SEASON project. On the mobile site a ruggedized server hosts the virtualized Distributed Unit (vDU). Ruggedized servers are particularly robust and suitable to be installed in unprotected environments such as mobile sites in the open air along the railway. The server placed at mobile

site includes a NIC with 2x25 Gb/s ports for Fronthaul (FH) connectivity and a DPUs (Data Processing Units) with 2x25 Gb/s card for Midhaul (MH) flows.

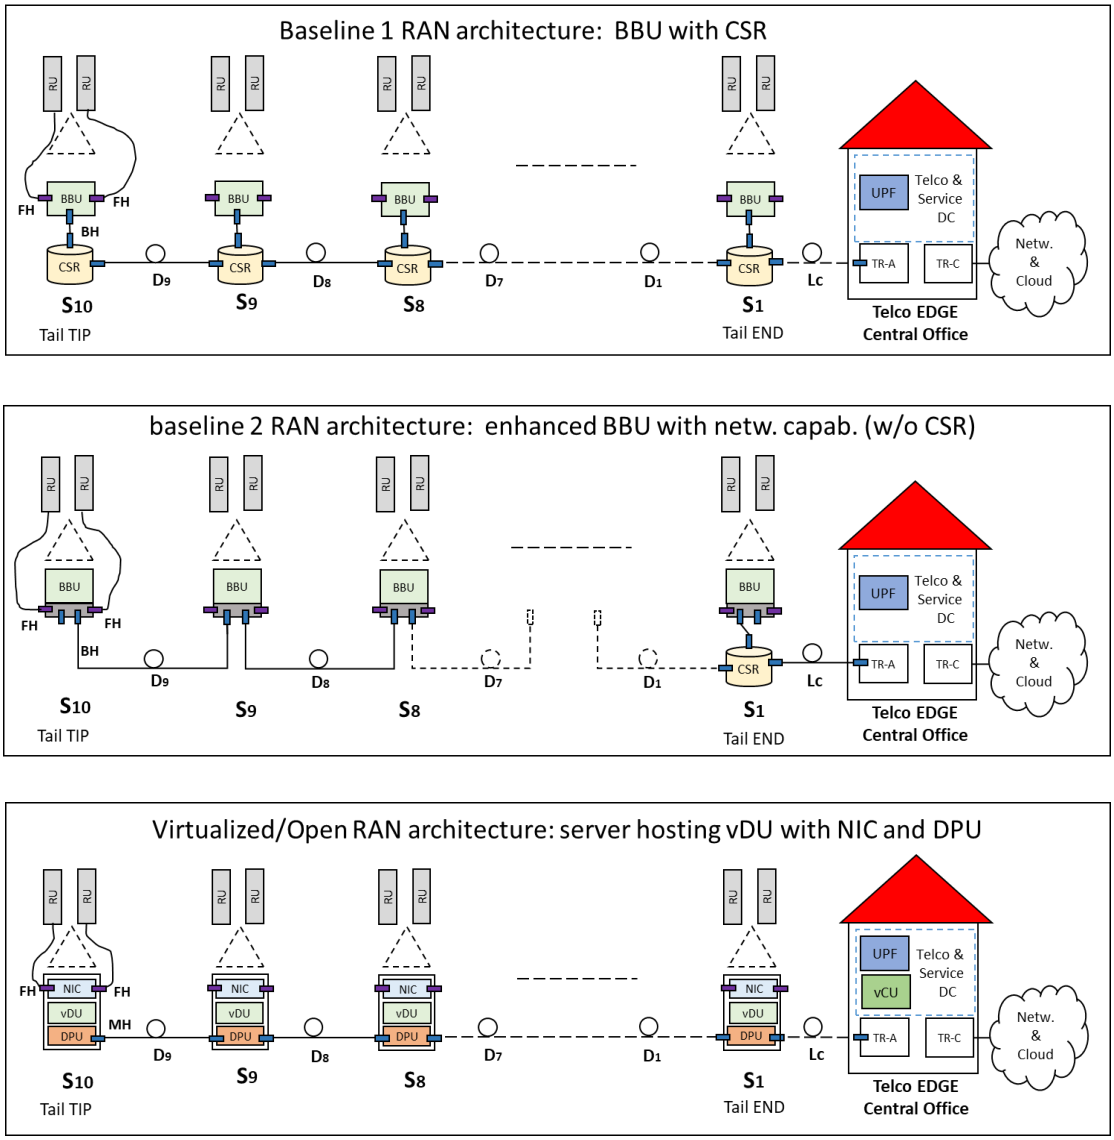

Figure 4.1-2 RAN architectures considered in the analysis. On top Baseline 1 architecture relying on Cell Site Routers for aggregate and forward BH traffic hop by hop up to the Central Office. In the middle, Baseline 2 architecture which uses BBUs with enhanced networking capabilities. This option avoids the use of a router in mobile sites, except the tail END (S1) where a router is placed to handle the BH traffic of the whole railway segment and exchange it with the Central Office. At bottom side the innovative architecture adopting Open-RAN and virtualized functions (cDU and vCU) on servers and NIC and DPU cards for x-haul flows handling.

These data rates matches with the requirements of x-haul interfaces reported above (subsection 4.1.1), taking also into account that the traffic collection is hop by hop along the chain and the last DPU, the one at S1 (end TAIL), must process and route the MH traffic for all the mobile sites. If the data rate requirements for x-hauling were higher than the one reported in subsection 4.1.1, as it could be in the future for 6G with a reorganization of the frequency bands (for example 100/400 Gb/s for FH and 100 Gb/s or more for MH aggregated at the tail END site may

be required) an upgrade on the cards would be necessary. Cards with 200 or 400 Gb/s per port are already available: so, in fact, implementations with these features are already feasible today. In the Virtualized open RAN architecture, at the Telco Central Office (CO) a standard server hosts the virtualized Centralized Unit (vCU) and other functions and physical equipment (i.e., UPF and IP and optical transport equipment).

The scenario described above has been subjected to three types of analysis that complement each other. The first two types of analysis are experiments and concern a simulation of the radio network to verify the reliability of the data rate values achievable in a system like the one modeled and a laboratory test which demonstrates the RAN and networking functions relating to network architecture depicted in the bottom part of Figure 4.1-2. They are reported in detail in D3.3 [SeaD3.3] and briefly summarized below in subsection 4.1.3. The third part of the study is the techno-economic evaluation of the solution, and it is reported here in subsection 4.1.4.

4.1.3 Simulations and Lab Experiments

The simulations of the radio system include as scenario setting parameters the channel bandwidth (20 MHz, 50 MHz, and 100 MHz), the MIMO configuration (SISO, 2×2, 4×4, 8×8), and the RU inter-site distance. The objective was to assess how these factors influence the achievable downlink throughput for a single train in motion, under realistic channel conditions that account for Doppler effects, path loss, and fading. The results demonstrate how increased spectral bandwidth and MIMO parallelism can compensate for the adverse propagation conditions typical of high-speed scenarios, albeit with diminishing returns at large inter-site distances. In summary, simulations reveal that with 100 MHz and 4×4 MIMO, the system can deliver downlink rates exceeding 500 Mbps per RU under optimal positioning, even in the presence of high-speed mobility

The experiment in the lab testbed employs two commercial servers with PCIe 5.0 bus and DPUs was used to evaluate a 5G infrastructure. Tests confirmed that a DPU with DOCA Flow could operate at maximum throughput for VLAN encapsulation/decapsulation actions, entirely in hardware, without relying on ARM cores. This demonstrated the DPU's capability for efficient packet processing.

The detailed descriptions of both experimental parts of the study are included in Section 5.2 of Deliverable D3.3 [SeaD3.3] and not replicated here to avoid content redundancies among project deliverables.

The third part of the study is the techno-economic evaluation of the solution, and it is reported here in the following subsection.

4.1.4 Techno Economic Analysis

A cost evaluation was performed comparing the three architectures, applying a specific cost model for transceivers, routers, and servers, along with estimated costs for legacy RAN

components like RUs, BBUs, and RAN software. The cost parameters used for the evaluation are included in Table 4.1-1. Values in table are given in Cost Unit and are taken in part, from what concern transceivers, servers and routers, from cost model tables included in Section 6. Routers are assumed to cost the same as Layer 3 switches, since these devices are assumed to have non-carrier-grade functionality, as they are intended for access use. However, even significantly increasing the cost of these devices (for example, by a factor of 2 or even 4), the results do not significantly change the relative cost differences between the three compared solutions. The receivers used are of only two types, 10G Extended Reach (L<10km) used between mobile sites (hypothetically spaced 5km apart) and 10G ZR (L< 80km) between the final site (END TAIL) and the central office (hypothetically distanced 50km). Indeed, radio simulations show that the MH/BH throughput of each RU does not exceed values in the order of 500 Mbit/s. Even in the most extreme case where all 20 RUs in the considered railway coverage scenario (2 RUs per site, 10 sites) generate 500 Mb/s of downlink traffic each, the total aggregate traffic collected on the last site is at most 10 Gb/s, but it will likely be much lower. For this reason, 10G transceivers for connections between mobile sites and even the one from the last site to the central office are sufficient to guarantee traffic transport in the most traffic demanding condition.

To carry out a complete cost assessment of the architectural solutions under examination, in addition to the costs of the compute and HW for the transport, the costs of the radio part (RU) and the costs of the RAN SW were also considered. Considerations and hypotheses were then made on how these costs could vary depending on the ecosystem that would be created if the O-RAN framework became relevant. Moderate savings on the radio component and significant savings on the RAN SW component are in fact expected in the case in which the O-RAN ecosystem was to reach a penetration such as to lead to significant market opening and volume.

Table 4.1-1 Cost parameters used in railway mobile service coverage scenario evaluation.

Description	Cost (c.u.)
RU - Active Antenna Unit (AAU) MIMO 4x4, 3.7GHz , 100 MHz	0.800
Standard Base Band Unit (BBU)	1.000
Enhanced BBU with networking capabilities	1.300
Ruggeddized 16 cores CPU server for vDU with 2x25G-NIC and 2x10/25G DPU	1.000
2 x 24 core CPU standard server	2.800
Grey transceiver 10 Gb/s LR	0.005
Grey transceiver 10 Gb/s ZR	0.080
Router with total of 400G interface bidirectional capacity	0.600
Router with total of 800G interface bidirectional capacity	0.900
RAN SW legacy (BBU)	1.000
RAN SW O-RAN (RAN part, vDU+vCU)	1.000

The bill of material for the three RAN architecture shown in Figure 4.1-2 are created taking into account capacity needs (as required by RAN x-haul interfaces) and distances to be covered by transceivers (according to scenario geotype parameters).

A what if analyses to understand the impact of price variations is performed, focusing on three key variational parameters, namely α (alpha) β (beta) and γ (gamma), with the following meaning and values.

- α (alpha) represents the relative cost of Legacy RAN Hardware (BBU) compared to O-RAN Hardware (servers for vDU). Values considered for α are 0.5, 1, 1.5. Reference case I for α = 1. With α = 0.5, the Legacy BBU cost is half the cost of a vDU server (-50%) while for α = 1.5 the Legacy BBU cost one and a half times the cost of a vDU server (+50%).
- β (beta) characterizes the relative cost of O-RAN Software (O-RAN SW) compared to Legacy RAN Software. RAN SW cost for the O-RAN case is expected to be equal to or lower than RAN SW cost for the legacy, this because O-RAN SW is based on the open-source model. Values considered for β are 0.25, 0.5, 1. For β = 0.25 O-RAN SW cost is 25% of legacy SW cost (or -75%, corresponding to a significant lower cost for SW in the O-RAN architecture compared with the legacy). With β = 1 the SW cost the O-RAN is the same as legacy RAN SW (corresponding to achieving any saving with open-source model, which can be considered rather unlikely). β = 0.5 is considered the more likely and reference value (O-RAN SW cost is 50% of legacy).
- γ (gamma) signifies the relative cost of O-RAN Radio Units (RU) compared to Legacy RUs. O-RAN RUs are expected to be equal to or lower than the cost of the legacy RUs, due to a likely tougher price competition typical of an open environment. Values considered for γ are 0.7 (-30% reduction, optimistic), 0.85 (-15% reduction, affordable and assumed as reference value), 1 (same RU cost for RU for both legacy RAN and O-RAN, rather pessimistic and unlikely value).

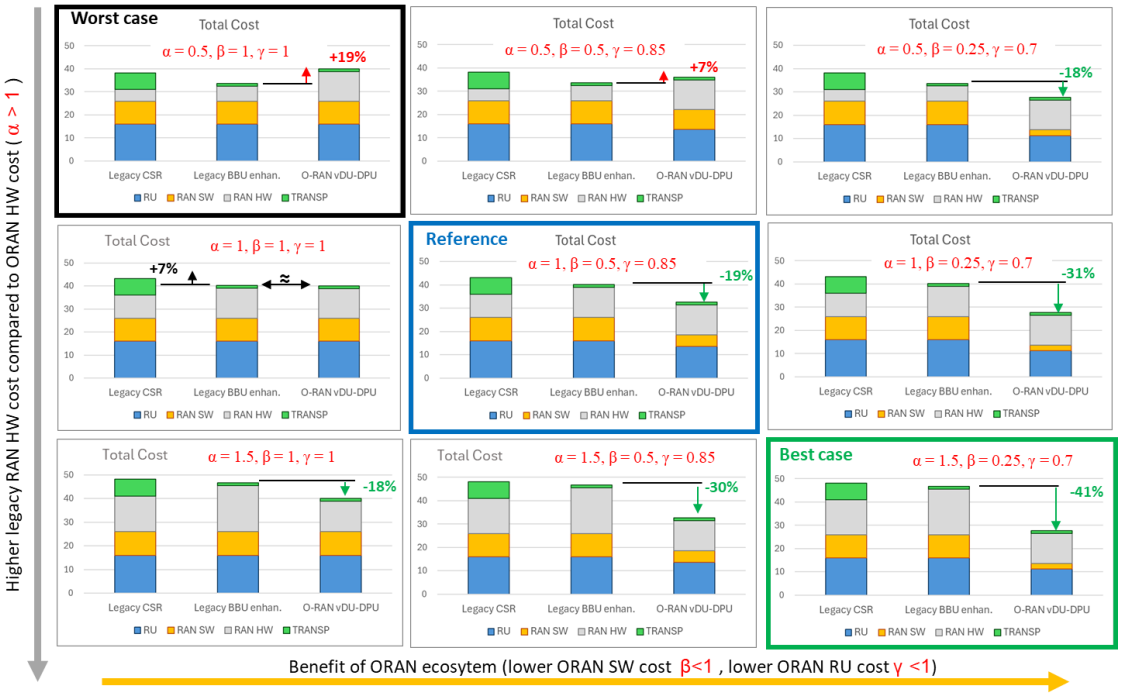

Figure 4.1-3 Summary of the cost evaluation. In each diagram the cost of the three RAN architecture is shown highlighting the α , β and γ parameter values used for cost calculation.

The evaluation results are given in Figure 4.1-3 where nine combinations of α, β and γ parameters were applied.

The representation of the figure shows how the benefits of the O-RAN ecosystem increase (from left to right) and as the presumed costs of the legacy HW (aggregated and mono-vendor HW and SW) increase compared to the O-RAN HW, the O-RAN vDU-DPU architecture achieve significant cost saving in comparison with both legacy baseline architectures.

The reference case is the one in the center of Figure 4.1-3 with the legacy BBU HW at the same cost with the Oran HW that hosts only the vDU (α=1) and with moderate advantages on RU (-15%) and on SW (-50%). In this case, classified as Reference, the O-RAN architecture based on ruggedized servers equipped with NIC and DPU cards at the site, allows a cost reduction of 19% compared to the best of the two baseline solutions, baseline 2 that uses the enhanced BBU and avoids the use of on-site routers. The worst case for the O-RAN architecture is when the BBU costs half the rugeddized server at site and there are no cost advantages for RU and SW with O-RAN (a very unlikely situation). In this case the penalty on the cost of O-RAN compared to Baseline 2 is symmetrical to the Reference case, i.e., +19% instead of -19%.

When all the conditions are favorable to O-RAN, i.e., there are full cost benefits from the O-RAN ecosystem for both RU (-30%) and SW (-75%) and even the legacy site HW is more expensive than the ORAN one (+50%), the economic advantages of the O-RAN solution reaches 41% compared to baseline 2, as it can be seen in the lower-right part of Figure 4.1-3. In all cases the baseline 1 solution that uses routers at the site is, even if only slightly, the worst of the 3 solutions examined.

4.1.5 Counting of Transceivers and O/E/O Conversions

In addition to the technical-economic analysis, some simple considerations can be made regarding the use of transceivers in the cases examined.

Very simply, the upper part of Figure 4.1-2 shows that in the Baseline 1 case, the number of transceivers used at each site is 2 SR transceivers for the FH (10G or 25G, depending on the radio parameters) and 4 10G ER transceivers for the BH. For Baseline 2 case (middle Figure 4.1-2) and the one with virtualized vDU O-RAN functions at the site with DPU cards for processing and networking (bottom Figure 4.1-2), the number of 10G ER transceivers is reduced from 4 to 2.

In the second and third cases, there is therefore a 33% reduction in the number of total transceivers and a 50% reduction in the number of ER transceivers (the more expensive ones). If we extend the analysis to the total number of optical-to-electro-optical (O/E/O) conversions, we can observe that in the Baseline 1 case, there are two conversions at each site: one in the BBU between the optical FH signal coming from the RU and directed to the connection to the router (CSR), and one between the two router ports, the one receiving the BH traffic from the BBU, and the one directed to the next router in the chain.

In the case of the Baseline 2 solution or the virtualized vDU O-RAN function with DPU cards, there is only one O/E/O conversion at the site: the one between the FH signal coming from the RU and the one that exchange the MH signal on the link to the next site in the chain. It can therefore be stated that in this scenario, the number of O/E/O conversions of the Baseline 2 and O-RAN-vDU-DPU solution is 50% lower than in a Baseline 1 scenario, corresponding to the legacy reference scenario.

5 ARTIFICIAL INTELLIGENCE (AI)/MACHINE LEARNING (ML) SERVICE ORCHESTRATION AND SELF-MANAGEMENT AND SECURE AI

5.1 COORDINATION OF RADIO ACCESS AND OPTICAL TRANSPORT OPERATION TO REDUCE OPTICAL CAPACITY OVERPROVISIONING

The advent of B5G/6G will revolutionize the way RAN will operate. Expected massive small cell deployments and features, such as adaptive functional splitting, are expected to change not only the volume but also the requirements of the traffic to be supported by the fixed transport network. In D2.2, an insight into smart RAN operation, focusing on how such operation will impact the autonomous operation of the fixed network, was presented. Based on tailored traffic models, a preliminary numerical analysis was done to highlight that smart and dynamic RAN operation lead to sharp fixed network traffic changes that require from coordination between RAN and fixed transport network to achieve efficient use of optical capacity resources while guaranteeing end-to-end (e2e) QoS.

In this work, we extend those contributions and results with the following novel contents:

- A context-aware autonomous network operation procedure is presented. This procedure is based on using contextual information sharing between slice management and fixed transport domains for an autonomous e2e network operation that coordinates and encompasses both the RAN and fixed network operation. The main algorithms and models involved are presented, which include asynchronous context updates and context-aware AI-based capacity reconfiguration modules.
- Numerical results to illustrate the benefits of the proposed context-aware autonomous network operation compared to benchmarking approaches, where slice management and the fixed network operate independently

5.1.1 Context-Aware Autonomous Network Operation

Figure 5.1-1 sketches the architecture and workflows of the proposed context-aware autonomous network operation. By means of this procedure, the RIC shares with the fixed network orchestrator the relevant information about the slice reconfigurations that come before slice changes are actually performed. With this relevant input, the transport network orchestrator can generate a context for the different agents running autonomous network control loops. In particular, those agents are in charge of dynamically allocating capacity to the optical capacity according to the current and predicted traffic [Vel21]. Although the procedure is focused on the reconfiguration of existing slices, note that it can be easily extended to the provisioning of new e2e slices, i.e., a first reconfiguration notification can be the provisioning of the slice. Table 5.1-1 introduces the notation consistently used in this section.

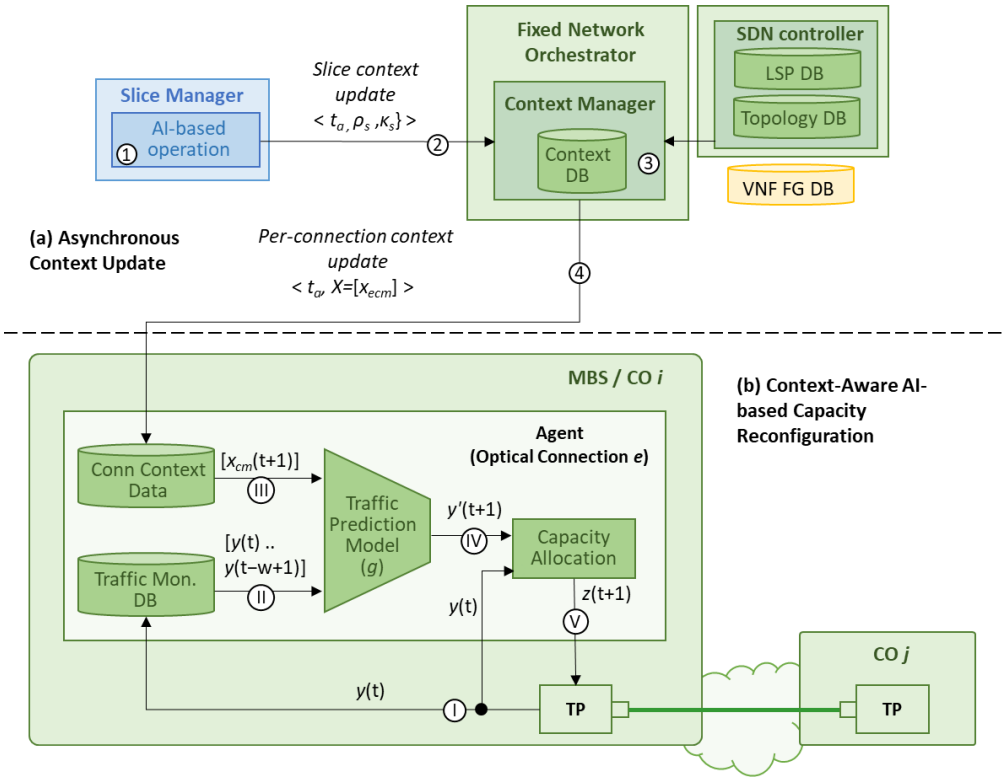

Figure 5.1-1 Context-aware autonomous network operation scheme containing: (a) asynchronous context update and (b) context-aware AI-based capacity reconfiguration

Dissemination Level	PUB (Public)

Table 5.1-1 Notation.

Symbol	Description
N	*Set of MBS sites and COs*
E	*Set of optical connections*
V	*Set of computing nodes*
S	*Set of e2e slices*
$F(s)$	*Set of virtual functions, i.e., DU, CU, UPF, of slice s*
$L(s)$	*Set of slice links of slice s*
C	*Set of service classes, i.e., eMBB, URLLC, mIoT*
M	*Set of RAN segments, i.e., radio, F-H, M-H, B-H*
δ_{ne}	*1, if MBS site/CO n is adjacent with optical connection e*
δ_{vn}	*1, if computing node v is located in MBS site/CO n*
δ_{sc}	*1, if slice s belongs to service class c*
δ_{lm}	*1, if slice link l belongs to segment m*
δ_{lf}	*1, if slice link l is adjacent to function f*
δ_{le}	*1, if slice link l is supported by optical connection e*
κ_s	*Capacity (in normalized PRBs) assigned to slice s*
$\rho_s = [< f, v >]$	*Node-based graph of slice s, consisting of an ordered vector of tuples. Each tuple contains the function f and the computing node v where the function is placed.*
$\Pi_s = [< l, e >]$	*Link-based graph of slice s, consisting of an ordered vector of tuples. Each tuple contains the slice link l and connection e that supports the link.*
$Y_e(t)$	*Traffic monitored in optical connection e at time t*
$x_{ecm}(t+1)$	*Capacity (in normalized PRBs) supported by optical connection e and belonging to segment m of class c expected at time $t+1$*
$z_e(t+1)$	*Capacity (in optical capacity units) to be allocated to connection e at time t+1*

The proposed procedure can be divided into two main processes. Figure 5.1-1a illustrates the asynchronous context update workflow that is executed when the slice manager decides to perform a reconfiguration, which is assumed to happen at time t_a (labeled as 1 in Figure 5.1-1a). Let us assume that reconfiguration decision is computed by an AI-based operation module (e.g., one that is similar to that of [Zhu21]) that analyzes the slice monitoring data and predicts a new (better) configuration for one or several end-e2e slices. Without loss of generality, we assume that such a reconfiguration is performed at a given time $t_b > t_a$. For instance, the AI-based module can predict an increase in user equipment (UE) activity and decide, at time t_a, that a new µBS in a given cell is needed 5 min later. Consequently, the capacity of some slices increases accordingly once the µBS is active, i.e., at time $t_b = t_a + 5$ minutes. Note that, in the case of a

reactive RAN reconfiguration due to some unexpected and unpredictable event, it might happen that $t_b = t_a$ limits the performance of the context-aware operation in the short period of time between the performance of the reconfiguration and the convenient updating of the context.

As soon as the decision is taken, a slice reconfiguration notification is sent to the context manager at the fixed network orchestrator for each of the slices that are to suffer some modification (2). This notification for slice s contains the time where the reconfiguration is expected to happen (t_b) and the new characteristics of the slice, namely the capacity κ_s (in capacity units like the number of PRBs) allocated to the radio segment and the graph ρ_s with the placement of each virtual function $f \in F(s)$ in the set of computing nodes V. Although both characteristics are related to RAN configuration, they are also indicative of the actual UE activity and, consequently, correlated with the actual slice traffic. Hence, this notification procedure enables the passing of relevant information from the RAN to the fixed network in an asynchronous and efficient way, without the need of continuous monitoring data exchange between domains.

Once the context manager receives a slice reconfiguration notification, it executes the context computation process (3). This process aims at computing the changes in the context due to the new reconfiguration. In a timely manner, those changes are updated in each of the agents managing optical connections (4), in order to guarantee that they can use up-to-date contextual information when running autonomous operation actions. Recall that the slice manager triggers the needed reconfiguration workflows with sufficient anticipation to ensure that they are active at time t_b.

More formally, we denote X as the context variable set at a given reconfiguration time t_b. This variable contains, for each optical connection $e \in E$, the expected capacity (in normalized capacity units) that such a connection will support for each service class $c \in C$ and segment $m \in M$. More formally, each $x_{ecm} \in X$ can be expressed by Eq. 10.

$$x_{ecm} = \sum_{s \in S} \delta_{sc} \sum_{l \in L(s)} \delta_{lm} \cdot \delta_{le} \cdot \kappa_s \qquad \text{Eq. 10}$$

Note that this context omits particular details of the individual slices since it aggregates all the slice capacity per class and per segment. Thus, the privacy of individual slices (clients) is guaranteed when operating optical connections. Moreover, to compute the context variables, static information about the mapping of slices to classes (δ_{sc}) and links to segments (δ_{lm}) is needed from the slice definition, which can be easily obtained at slice provisioning time. However, the mapping of slice links to optical connections (δ_{le}) is dynamic and depends not only on the slice graph ρ_s but also on the placement of the computing nodes in the COs (δ_{vn}), which can change in time depending on the fixed network reconfiguration actions.

Algorithm in Table 5.1-2 details the context computation process, which is required for every single slice reconfiguration notification. Thus, the algorithm receives all the data of one slice reconfiguration notification at a given time t_b, computes the new context for that time, and updates the context so that those planned reconfigurations happen at time t_b (and after it).

Moreover, the process has access to the internal database (DB) that stores the static slice characteristics, context status, and transport network configuration and is conveniently kept up to date by the fixed network controller.

Note that the DB can store multiple context variable sets and fixed network configurations if several reconfigurations are planned to happen at different times. Then, for convenience, the slice reconfiguration and context data are stored and indexed by time; so, they can be easily retrieved and modified every time a new reconfiguration notification is processed. For those slice static parameters, we assume that they follow a non-temporal indexing by slice id.

Table 5.1-2 Algorithm 1: Context computation (Context Manager).

INPUT: s, t_b, ρ_s, κ_s	
OUTPUT: -	
1:	**if** $DB[t_b] == \emptyset$ **then**
2:	t' ← getPreviousTime(DB, t_b)
3:	$DB[t_b]$ ← $DB[t']$
4:	t'' ← getHighestTime(DB)
5:	T ← getTimeRange(t_b, t'')
6:	c ← $DB[s]$[“c”]
7:	**for each** $t_i \in T$ **do**
8:	X={x_{ecm}} ← $DB[t_i]$[“X”]
9:	κ^0 ← $DB[t_i]$[“κ_s”]
10:	π_s^0 ← $DB[t_i]$[“π_s”]
11:	π_s ← computeMapping(ρ_s) (Algorithm 2)
12:	**for each** <l, e> ∈ π_s^0 **do**
13:	m ← $DB[s][l]$[“m”]
14:	x_{ecm} ← x_{ecm} − κ^0
15:	**for each** <l, e> ∈ π_s **do**
16:	m ← $DB[s][l]$[“m”]
17:	x_{ecm} ← x_{ecm} + κ_s
18:	$DB[t_i]$[“π_s”] ← π_s
19:	$DB[t_i]$[“κ_s”] ← κ_s
20:	$DB[t_i]$[“X”] ← X
21:	**return** -

The algorithm starts by generating a new entry in the DB for time t_b (if it does not exist yet), copying the same data in the entry indexed in the immediately previous time to t_b (lines 1–3 in Table 5.1-2). Then, the vector of times T is obtained; this vector contains the list of available times in the DB between t_b and the highest available time (lines 4–5). In other words, this identifies all the different contexts that need to be updated due to the new reconfiguration. Next, after retrieving the class of the slice, a loop is executed to update all the contexts (lines 6–7). Thus, for each $t_i \in T$, the stored context X, capacity κ^0, and mapping π_s^0 of the slice links to the optical connections are retrieved from the DB (lines 8–10). Then, the new mapping π_s is computed from graph ρ_s (line 11). After this computation, the stored slice capacity is subtracted from all the x_{ecm} variables belonging to the stored mapping (lines 12–14), whereas the new

capacity is added to the x_{ecm} variables of the new mapping (lines 15–17). Finally, the slice capacity, mapping, and context are updated (lines 18–20).

Algorithm 2 in Table 5.1-3 details the process of computing a mapping π_s from a new slice graph ρ_s, i.e., line 11 of Algorithm 1. After some initialization, the process iterates in the graph in order to obtain the pair of tuples <*f*,*v*> and <*f'*,*v'*> that identify each graph adjacency (lines 1–4 of Table 5.1-3). Then, queries to the DB are performed to retrieve the slice link *l* that supports that graph adjacency (line 5), as well as the optical connection *e* that connects the sites where the computing nodes are running (lines 6–8). Afterwards, the tuple <*l*,*e*> is added to π_s, which is eventually returned (lines 9–10).

Table 5.1-3 Algorithm 2 Compute Mapping.

	INPUT: ρ_s
	OUTPUT: Π
1:	π_s = []
2:	**for** i =1..$\lvert\rho_s\rvert$ − 1 **do**
3:	<f, v> ← $\rho_s[i]$
4:	<f', v'> ← $\rho_s[i + 1]$
5:	l ← $DB[s]$.query("l" | δ_{lf} == 1 & $\delta_{lf'}$ == 1)
6:	n ← $DB[s]$.query("n" | δ_{vn} == 1)
7:	n' ← $DB[s]$.query("n" | $\delta_{v'n}$ == 1)
8:	e ← DB.query("e" | δ_{ne} == 1 & $\delta_{n'e}$ == 1)
9:	π_s.add(<l,e>)
10:	**return** Π

As already mentioned, in addition to the asynchronous context update workflow, the *context-aware AI-based capacity reconfiguration* workflow (Figure 5.1-1b) is executed periodically at every time *t*, i.e., when a new traffic monitoring sample $y(t)$ is collected (labeled as I in Figure 5.1-1b). To support this workflow, the following elements are deployed in the agent: (*a*) the traffic monitoring DB storing the last *w* traffic values; (*b*) the context DB with the current and future contexts affecting the optical connection; (*c*) the traffic prediction model that estimates the expected traffic $y(t + 1)$; and (*d*) the capacity allocation module that decides the amount of optical capacity units $z(t + 1)$ (e.g., digital sub-carriers) to be configured in the optical transponder according to traffic prediction $y(t + 1)$ to stay below a desired (target) maximum load.

Once the new measurement is collected, the traffic monitoring DB is updated (II) and the expected context at the prediction time, i.e., $t + 1$, is retrieved from the context DB (III). For convenience, since the connection context DB contains only context variables for the specific connection, we removed the sub-index *e* from the *X* variables. Then, the traffic prediction model *g* can be formally expressed as Eq. 11, which results in $w + |C|\cdot|M|$ inputs and provides one single output (IV). The training model *g* can be autonomously conducted by the agent as soon as it receives context variables and can learn from them. Without loss of generality, we assume that the agent, at its provisioning time, receives an initial model that has been trained offline

with generic traffic data and null context inputs. Then, the model is re-trained online, combining the contextual variables and monitored traffic and improving the accuracy of the initial one.

$$y'^{(t+1)} = g([y(t-i), \forall i = 0..w-1], [x_{cm}(t+1), \forall c \in C, m \in M\ ]) \quad \text{Eq. 11}$$

The traffic prediction $y'(t+1)$ is then processed by the capacity allocation module (V), jointly with the current traffic $y(t)$. Thus, with λ_{max} as the target maximum load, u as the traffic supported by each optical capacity unit, and z_{max} as the maximum number of optical capacity units supported by a transponder, $z(t+1)$ can be formally defined as Eq. 12.

$$z(t+1) = max\left(z_{max}, \left\lceil \frac{\max(y(t), y'^{(t+1)})}{\lambda_{max} \cdot u} \right\rceil\right) \quad \text{Eq. 12}$$

Without loss of generality, we assume that the transponder internally implements an agent that allocates the target number of capacity units detailed in $z(t+1)$ [Vel21]. Note that, while model g is being retrained online and the $y'(t+1)$ predictions have not reached a target accuracy performance, a more conservative policy can be temporarily used for capacity allocation, e.g., $z(t+1) = z_{max}$.

As a final remark, it is worth mentioning that the proposed context-aware approach does not increase execution time with respect to the existing optical capacity management approaches based on optical connection monitoring and local decision making [Vel21]. Note that the connection context is stored and available at the local agent and is simply added as another input to the prediction model. Moreover, this connection context is updated asynchronously, e.g., every time a slice reconfiguration update is processed by the context manager. Thus, context management and AI-based optical capacity reconfiguration are completely independent processes; therefore, the former does not increase the computing time of the latter.

5.1.2 Results

For numerical evaluation purposes, we developed a Python-based simulator that allows reproducing network scenarios that mix RAN and fixed network technologies and segments. In a nutshell, the simulator generates flow traffic by means of statistics-based generators that emulate RAN cells and fixed access points and propagate them through a system of fluid-flow continuous queues that model the different network elements (packet and optical interfaces and connections) and segments (RAN cell, optical access, etc.). A detailed architecture and building blocks of this simulation, as well as the different implemented policies and functionalities and parameter configuration can be found and openly accessible in [Wan24].

As a summary of simulator configuration, following we detail those more relevant aspects:

- Regarding RAN configuration, we considered RAN cells consisting of 1 macro BS (MBS) working at the sub-6 GHz band (2 × 2 MIMO, 20 MHz bandwidth, 30 KHz subcarrier spacing, and 55 PRBs), and a variable number of micro BS (µBSs) configured in the mmWave band (8 × 8 MIMO, 100 MHz bandwidth, 30 KHz subcarrier spacing, and 275 PRBs). Without loss of

generality, we assumed that every active BS has one slice per service class. Moreover, the F-H and B-H traffic were computed according to the formulation and references in [Wan23] for a mix of eMBB, URLLC and mIoT services.

- Two different RAN cell scenarios were configured according to the distribution of UE within the cell areas. For the sake of a fair comparative analysis, the total demand and the demand per service class are the same in both scenarios. Then, in the *clustered scenario*, the UE types of the same class are grouped in a given region within the area and separated from the other classes. This means that a given µBSs will typically provide service to (mainly) the UE of one single class. Therefore, the clustered scenario tends towards µBSs with a predominant slice and a more variable load (depending on the predominant class) during the day. On the other hand, the *distributed scenario* assumes that the UE types are randomly placed in the cell area and, therefore, that there are no clusters of UE of the same class. Consequently, all the µBSs have a similar load (which fluctuates with the overall traffic) with a mix of slices from all the service classes.
- We implemented a policy to run in the slice manager that consisted of rules oriented towards reconfiguring the RAN capacity resources and slices by means of the computation of traffic loads (defined as UE traffic over RAN capacity), and we compared them with thresholds in order to determine the actions to perform. Whenever an action is detected at time *ta*, slice reconfiguration notification (if needed) is created with time $t_b = t_a + 1$, thus giving time for a context update and the proper use of contextual information for AI-based optical connection operation. Note that the slice reconfiguration actions detected and notified at time t_a are actually performed at time t_b. We conducted an exhaustive set of simulations in order to find the configuration of rules and thresholds that produces a remarkable energy consumption reduction (measured in terms of active µBS) that guarantees robust QoS (no traffic loss):
    - When one µBS exceeds a total load of 75%, the closest inactive µBS (if any) is switched on. Otherwise, if the load drops below 15% and the neighbor µBSs are active, then the µBS is switched off and its current demand is then served by its active neighbor µBS.
    - The PRBs assigned to each slice are assigned and released in order to keep (or approach) a slice load between 60% and 70%.
    - Depending on the number and utilization of the µBSs, we define both low-load and high-load regimes, in which different splits and virtual function placements are considered, e.g., when the load exceeds 30% in URLLC slices, the functions approach UEs to reduce the risk of a potential delay degradation due to high load (DU at MBS and CU at access CO).

Assuming the aforementioned configuration, we first focus on the traffic prediction carried out in the agent of an access optical connection connecting the MBS of a RAN cell with 1 MBS and 16 µBSs with its associated access CO site. In particular, we aim to evaluate the accuracy of predicting y(t + 1) using the model in Equation (2) (named as *context*) and to compare it with

the model used in [Rui20] (named as the *benchmark*) that showed high accuracy in predicting traffic in fixed networks supporting 4G and residential services. Note that the benchmark model predicts y(t + 1) only as a function of the last *w* traffic values stored in the monitoring DB. Both models have been implemented as deep neural networks with 2 hidden layers with 24 and 12 neurons, respectively, and hyperbolic tangent activation functions. Thus, the comparison between both models enables a fair way to evaluate the use of contextual data in B5G scenarios.

Figure 5.1-2 shows the normalized prediction error (normalized to zero-day error) as a function of the elapsed operation time (in days) from the provisioning time instant. As can be seen, for both the distributed and clustered scenarios, convergence to a stable error (one order of magnitude lower than that of the initial model) is reached after few days of operation (3 to 5). As the connectivity between RAN cells and automatic cell optimizers (ACOs) typically lasts for very long time periods (weeks to months), the duration of this initial tuning phase is negligible with respect to optical connectivity lifetime. Thus, context-based prediction models can be practically deployed since they are easily trained and remarkably improved with the data available in the optical connection agent.

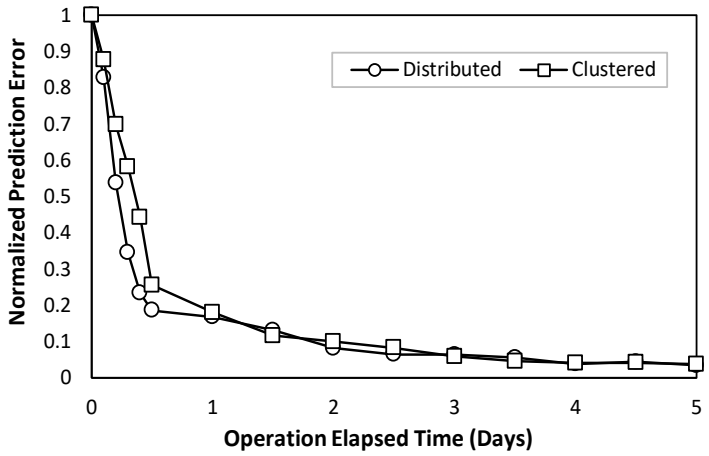

Figure 5.1-2 Online training performance of context traffic prediction model.

Next, regarding the error convergence performance obtained in Figure 5.1-2, where a stable prediction error for both the benchmark and context models is reached, Figure 5.1-3 and Figure 5.1-4 show a detail of the performance of the traffic prediction for one of the sudden traffic changes introduced by smart slice operation in clustered and distributed scenarios, respectively. The figures show real traffic and predicted traffic using both benchmark and context methods, as well as the relative error. Note that benchmark prediction presents large inaccuracies (errors around 50%) in the presence of those sudden changes. Conversely, the context model provides accurate and smooth prediction since it anticipates sudden changes thanks to context information.

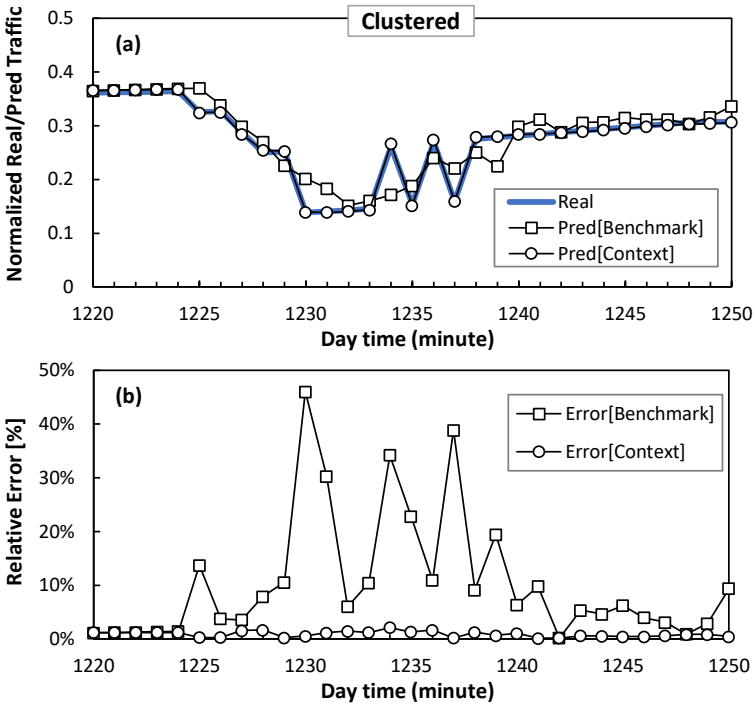

Figure 5.1-3 Traffic prediction detail in access optical connection for clustered scenario in terms of absolute (a) and relative (b) error.

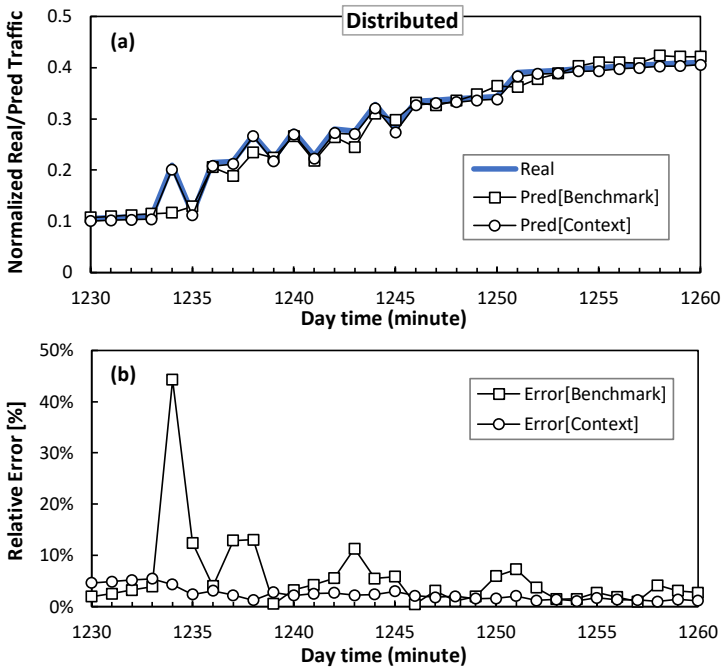

Figure 5.1-4 Traffic prediction detail in access optical connection for distributed scenario in terms of absolute (a) and relative (b) error.

Dissemination Level	PUB (Public)

A comparable performance can be seen in metro network scenarios. Thus, in metro connections, where the RAN and fixed access traffic is mixed, the benchmarking model provides a much larger prediction error (~30%) than the context one (~2%), even when the fixed access traffic doubles that of the RAN traffic. Now that the accuracy of the context model has been validated and its improvement with respect to benchmark one has been shown in terms of the prediction error, let focus on evaluating the impact of such predictions to dynamically allocate optical connection capacity. Figure 5.1-5 shows the obtained results for two operational objectives. On the one hand, the capacity minimization objective (Figure 5.1-5a-c) aims at adjusting the allocated optical capacity so that it is as close as possible to the actual traffic. Thus, the load of optical connections can grow up to 95%, in order to exploit the optical capacity resources. On the other hand, the delay reduction objective (Figure 5.1-5d-f) allocates an excess of capacity to avoid congestion in the fixed segment, which could introduce additional delay to the RAN slices. Therefore, the maximum load is kept below 80%.

The results include the average and maximum values of the capacity allocated, the load, and the traffic loss as functions of the RAN cell density. One initial observation is that both models allocate similar overall optical capacity (with small differences in the averages that are hard to observe). Moreover, the (large) relative errors in model prediction shown in the previous section for the low-density RAN cells (16 µBS) do not produce an effect on the decisions made by the connection capacity allocation module. This is because the absolute fluctuations stay between the size of the optical capacity units (recall that we consider 25 Gb/s sub-carriers); consequently, the benchmark model is still a valid traffic predictor.

However, as soon as the density of the cells increases, the real magnitude of the prediction inaccuracies exceeds the sub-carrier capacity. Thus, the benchmark model (sometimes) underestimates the required capacity to support the target load, which leads to unacceptable performance (traffic loss) for the RAN cells that are denser than 32 µBS. Conversely, the optical connection capacity allocation based on the context model prediction provides accurate operation, which is always below the desired target load and eliminates traffic loss, for a wide range of RAN cell densities. In light of these results, we can finally validate the proposed context-aware autonomous network operation procedure, which guarantees smooth and smart fixed transport network operation in the presence and during the activity of the smart B5G slice operation.

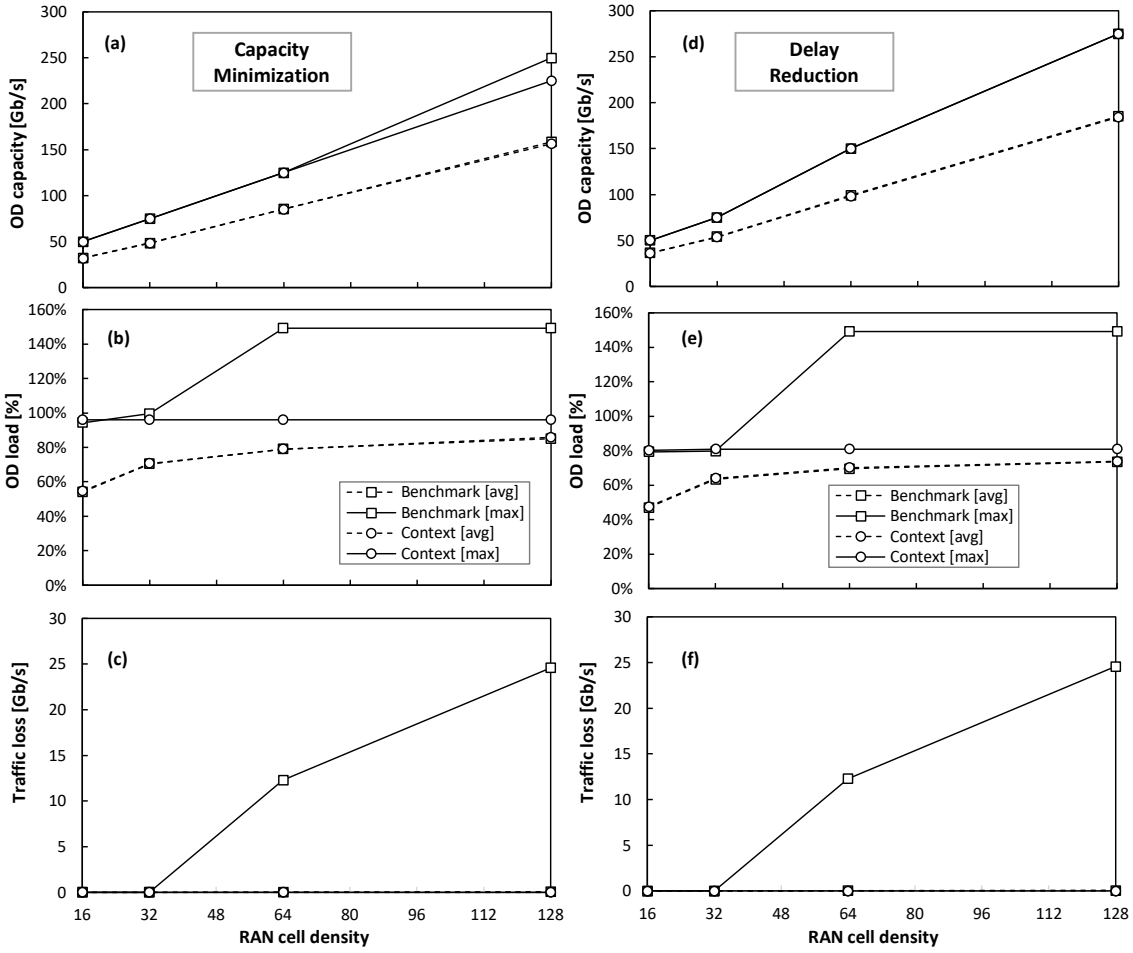

Figure 5.1-5 Performance of optical connection capacity reconfiguration for capacity minimization (a–c) and delay reduction (d–f) objectives.

5.1.3 Conclusions

In this Section, we focused on the impact of B5G slice operation on the underlying transport network infrastructure in charge of providing computing and connectivity resources. In particular, we analyzed how smart slice operation based on pillars, such as dynamic capacity reconfiguration, adaptive µBS management, and flexible functional split configuration, injects highly variable and hard-to-predict traffic into access and metro fixed networks. Thus, effective autonomous fixed network operation cannot be achieved without coordination between the slice manager and the fixed networks. In order to maximize the autonomous fixed network performance while hiding the sensitive information of the individual slices, we explored the concept of context sharing, which allows the slice manager to asynchronously provide information about the slice reconfigurations to be performed in an aggregated and private way with enough anticipation to synchronize with an autonomous fixed network operation. This slice context is then incorporated into the AI-based prediction models that are used for dynamic capacity allocation in optical access and metro networks.

The proposed architecture (presented in D2.2) and algorithms for context update and usage for prediction purposes in optical connections (presented in this chapter) were validated by means of exhaustive simulations that considered a realistic network operator infrastructure. The results showed that adding the slice context to traffic prediction greatly enhanced prediction accuracy. In particular, the maximum prediction error was remarkably reduced compared with the benchmarking approaches for traffic prediction. This prediction accuracy gain is observed under different RAN demand scenarios with heterogeneous services, as well as for optical access networks carrying only RAN traffic and for metro networks transporting a mix of RAN and fixed (residential and business) traffic.

The value of this prediction accuracy improvement was eventually evaluated for the dynamic allocation of capacity to the optical connections supporting slice links, under two different network operational objectives: capacity minimization and delay reduction. In both cases, the benefits of the proposed context-based operation increased with the size and density of the RAN cells. Therefore, we can conclude that, with the increase in RAN traffic and the density of RAN cells that is foreseen for B5G scenarios, there is a need for context-based coordination to achieve actual e2e smart operation.

Extended details and results about this work can be found in [Wan24].

5.2 Dynamic Subcarrier Allocation in P2MP Connections To Reduce Energy Consumption

Optical point-to-multipoint (P2MP) connectivity [Rui14] based on digital subcarrier multiplexing (DSCM) has been shown as a solution for the metro-access segment that is able to reduce capital and operational costs and support the capacity and high dynamicity needs of future 6G services [Wel21]. To achieve maximum performance, activation and deactivation of subcarriers must be done near-real-time to provide just the capacity needed to support the input traffic.

In addition, OPEX savings can be obtained by exploiting DSCM capability for operating the different SCs independently, so just enough capacity to support the input traffic is provided, which reduces energy consumption. However, that requires near-real-time operation, as well as some anticipation to give time to timely activate the SCs, so that they become available when it is required.

In D2.2, four approaches supporting near-real-time SC operation in a P2MP connection were presented and their performances compared. Distributed approaches eliminate the need for centralized decision making by moving decision making to the agents of the transponders participating in the specific P2MP connection. The performance of the distributed approaches is, in general, below that of the centralized one, and MARL is the approach with the best performance among the distributed ones. The most significant finding, however, is that dynamic DSCM allocation, regardless of the method, allows reducing the number of active SCs during

extended periods throughout the day. This reduction is directly linked to energy consumption costs.

In this deliverable, we translate these previous performance results into OPEX reduction evaluation, to quantify the energy consumption savings of dynamic DSCM allocation compared to fixed and statically planned operation.

5.2.1 Dynamic P2MP Connection Management and Energy Consumption Analysis

For the sake of a better understanding of the energy consumption analysis, we start re-introducing some concepts already explained in D2.2 about dynamic P2MP connection management.

DSCM facilitates the deployment of P2MP optical connectivity, since SCs sourced from a single hub node can be assigned to different leaf nodes. In the reverse direction (denoted MP2P), SCs generated from different leaves can merge to connect the source nodes to the hub node. For illustrative purposes, Figure 5.2-1(a) presents an example where four-leaf nodes are connected to a hub node in a P2MP connection. In the example, the hub node can support 16 SCs and each leaf node is assigned 4 contiguous SCs, while ensuring that each SC is assigned to a single leaf node, so as to avoid SC overlap (also referred to as oversubscription) since it leads to data loss. Comparing the P2MP optical connection in Figure 5.2-1(a) to a regular P2P one, we observe the reduction in the total number of transponders that are required, i.e., four Txs for the leaf nodes and one Rx for the hub node in the case of the P2MP connection in contrast to four Txs and four Rxs, respectively, in the case of optical P2P.

The above observation can also be stated in a slightly different way. In the example shown in Figure 5.2-1(a), leaf nodes are serviced since DSCM Txs can generate up to 4 contiguous SCs and the DSCM Rx can process up to 16 SCs. The number of leaf nodes in a P2MP connection can be further increased with dynamic SC management, which can assign SCs dynamically to the leaf nodes, so those not requiring the full capacity of the transponder to support the local traffic can give one or more of their SCs up to other leaf nodes with higher capacity requirements. This can result in cost savings, as more leaf nodes [e.g., five leaves, as in Figure 5.2-1(b)] can be serviced, as well as in power savings, since not every SC might be required to be active for the current traffic. For these gains to be fully realized, a control mechanism is necessary to ensure the proper P2MP connection operation, specifically from leaves to the hub (MP2P), to avoid oversubscriptions and to assure that the capacity needs of every leaf node are met.

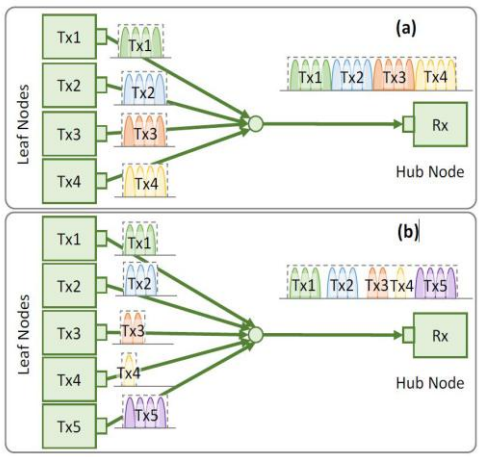

Figure 5.2-1 Example of a P2MP connection with four (a) and five(b) leaf nodes.

Now let us provide some initial numbers about energy consumption, when more than 4 leaf nodes are considered, in particular up to 7 that is the operational limit shown in [Sha24]. Table 5.2-1 presents the maximum energy consumption (all SCs active) assuming the normalized energy consumption values provided in Section 3.2. In particular, we consider leaf nodes (Txs) of 100 Gb/s with 4 SCs each and different configurations for the receiver (Rx) side:

- P2P: One 100 Gb/s Rx for each 100 Gb/s Tx
- P2MP – Static (Figure 5.2-1a): one hub node with 400 Gb/s transponder equipped with 16 SCs for every 4 100 Gb/s leaf nodes. Note that, when the number of leaves is greater than 4, another 400 Gb/s is needed. In that case, only the number of SCs necessary to support the number of associated TXs is considered for energy consumption calculations.
- P2MP – Dynamic (Figure 5.2-1b): a single hub node with 400 Gb/s transponder equipped with 16 SCs.

Table 5.2-1 Normalized Maximum Power Consumption for different configuration and number of leaf nodes.

Number of Leaf Nodes	P2P			P2MP – Static			P2MP – Dynamic		
	Tx	Rx	Total	Tx	Rx	Total	Tx	Rx	Total
4	3.20	3.20	6.40	3.20	1.00	4.20	3.20	1.00	**4.20**
5	4.00	4.00	8.00	4.00	1.88	5.88	4.00	1.00	**5.00**
6	4.80	4.80	9.60	4.80	1.92	6.72	4.80	1.00	**5.80**
7	5.60	5.60	11.20	5.60	1.96	7.56	5.60	1.00	**6.60**

As it can be observed, the maximum energy consumption obtained by provisioning P2MP connections in a static way is clearly reduced when we allow more than 4 leaf nodes per Rx and apply dynamic capacity allocation to avoid oversubscription and traffic loss. Therefore, dynamic P2MP management is expected to provide remarkable energy consumption savings.

5.2.2 Energy Consumption Performance Evaluation

In this section, we go a step beyond the previous analysis and compute the energy consumption along the day time, where SCs are dynamically activated and de-activated according to actual traffic needs. To this end, we start from the results presented in D2.2 about the number of SCs allocated every day for different approaches. In particular, we focus on the Distributed deterministic MAS (DD-MAS), where Tx agents use a deterministic algorithm for spectral allocation based on the capacity required for the local transponder, as well as the current spectral allocations composed of the information shared by the other Tx agents.

Figure 5.2-2a shows the normalized power consumption along the day for a P2MP connection with 6 leaf nodes when it is dynamically operated with the DD-MAS approach. For comparison purposes, the power consumption for both P2P dynamic and P2MP static benchmarking approaches is shown. Moreover, Figure 5.2-2b shows the relative savings of DD-MAS with respect to both benchmarking approaches. As anticipated in Table 5.2-1, DD-MAS operation offers the best performance in terms of energy efficiency, being between 20% and 40% more energy efficient than benchmarking approaches.

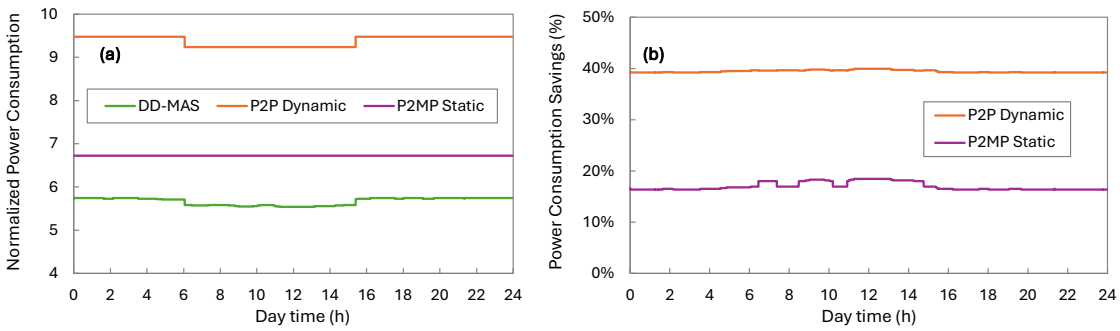

Figure 5.2-2 Detailed power consumption values (a) and relative savings (b) considering 6 leaf nodes.

5.2.3 Conclusions

OPEX savings in terms of energy consumption reduction are obtained by exploiting DSCM capability for operating the different SCs independently, so just enough capacity to support the input traffic is provided. The proposed dynamic near-real-time operation procedure based on MAS allows supporting more than 4 100Gb/s Txs in a single 400Gb/s receiver in a P2MP configuration, which leads to remarkable energy consumption savings, i.e., between 20% and 40% with respect to alternative connectivity management such as dynamic P2P operation and static P2MP operation.

5.3 Optical Line System Automatic Setup (Amplifier Configuration)

Besides dynamically configuring DSCM-based optical connections, optimizing EDFAs power has been identified as a promising way to minimize energy consumption. An OLS automatic setup approach for Access/Aggregation/Metro networks is presented. This method is based on dynamically configuring EDFASs in constant power mode, which is often preferred because it allows direct control over the power levels of each wavelength. The main drawback of this operation mode is that it necessitates precise and real-time knowledge of the number of channels entering the amplifier. A way to efficiently and practically implement automation based on monitoring to achieve constant power mode setup is to use optical line controllers (OLCs). The solution relies on the traffic matrix and the network's topology. It begins with an approximate estimate of the span loss, which serves as the initial basis for determining the optimal gain. Additionally, it requires measurements of the received power for all channels at every point where they are added or dropped, using the power monitors available on transceivers. This approach estimates the optimal gain for each amplifier in the network to minimize the difference between the actual and target receiver power for all channels at any drop node. It also ensures that all channels remain within the receiver's acceptable range, with a specified margin. With this method, network operators can achieve a more accurate and automated calibration of amplifier gains, enhancing overall network performance and reliability. This approach reduces the need for manual intervention and allows the network to dynamically adapt to changing conditions, ensuring optimal performance across various network segments.

Preferred EDFA operating modes depend on network monitoring capabilities:

- In heavily instrumented DWDM systems (for example, long-haul cores) where channel counts can be tracked in real time via OCMs or similar devices, Constant-Power mode is ideal because it delivers precise per-wavelength output control.
- In other environments, where you can’t reliably count channels at each amplifier, Constant-Gain mode is more practical. However, it still requires an accurate, real-time estimate of the preceding span loss so the gain can be set correctly. If gain is misconfigured, some channels may fall outside the receiver’s power or OSNR requirements, and these errors can accumulate over multiple hops. Simulation results (Figure 5.3-1) show that improper gain settings greatly increase the risk of channels falling outside the valid operating range in terms of received power or OSNR, leading to service degradation over multiple hops.

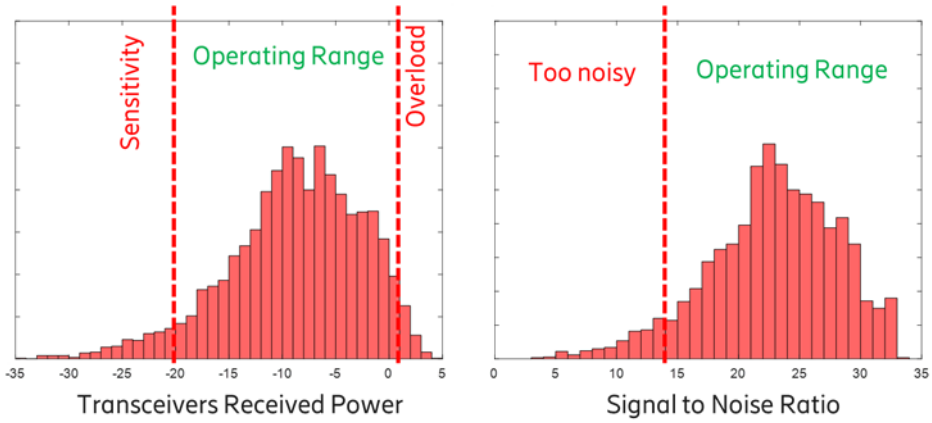

Figure 5.3-1: Simulated distributions of received power and OSNR at the receivers.

In this discussion, we focus on straightforward DWDM systems without built-in monitoring equipment, which therefore run in Constant-Gain mode. Such configurations are common in cost-sensitive IP-over-WDM deployments at the access, aggregation, and metro levels. While amplifiers are needed to offset span losses and those introduced by FOADM filters or splitters, it's impractical and too expensive to install per-node channel monitors throughout these networks. Fortunately, modern OLCs, like SDN controllers for the optical domain, can connect to every node over a management network. This centralized controller can gather optical performance metrics from all amplifiers and other devices and then automatically adjust settings (for example, tuning each EDFA's gain). By applying advanced algorithms or AI-driven routines, an OLC delivers automation, reduces manual intervention, and improves both operational efficiency and service reliability.

In practice, when a DWDM system runs in Constant-Gain mode, each EDFA's gain is set based on the loss measured over the preceding span. There are three common ways to obtain that loss figure:

- Transmitter/Receiver Power Monitors
    - Install photodiodes at the output of the upstream EDFA and at the input of the downstream EDFA.
    - The difference between those two total-power readings equals the span loss.
    - Drawback: any channels added or dropped in that span will skew the loss calculation.
- Use of an Optical Supervisory Channel (OSC)
    - An OSC signal—already present to link nodes to the management network—is launched after one EDFA and tapped before the next, using dedicated filters.
    - The loss is computed from the OSC transmitter's output and receiver's input, with corrections for filter insertion loss and wavelength offsets.

  - Drawback: typical OSC receivers have ±2 dB measurement uncertainty, so the calculated span loss can be off by as much as ±4 dB.

- Manual Span-Loss Calibration
  - At initial deployment, a technician measures each fiber span's loss with external test gear (e.g., a laser source and power meter) and programs the corresponding EDFA gain.
  - Subsequent adjustments after fiber aging or maintenance faults are also done manually.
  - Drawbacks: no automatic compensation for loss changes over time, high operational cost, and reliance on skilled field staff.

Because of these challenges, we adapted a technique that Ericsson patented in 2020 (US20240031024A1) for use in the filterless coherent WDM architecture defined by SEASON in access and long-access networks. It can also be applied in the metro and core environments. Our approach uses two sources of data: the network's traffic matrix and its physical topology. To kick off the gain-setting process, we start with a coarse estimate of each span's loss. We also collect per-channel received power readings at every add and drop point data already available from the transceiver power monitors. Using that information, the algorithm computes the optimal gain for each EDFA to minimize the deviation between each channel's actual received power and its designed target level at every drop node. Crucially, the "target" isn't just any value within the receiver's acceptable range—it's the specific design-average power level expected for that node and transceiver type. By steering the amplifiers toward those precise setpoints, the system automatically aligns not only received power but also OSNR and launched power with their intended design values.

The method is formulated as a "Linearly Constrained Quadratic Programming" problem as following:

- The solution guarantees stable and no traffic-affecting operation since every step goes toward the minimization of the error and each setting is constrained to lay in the required range
- Fast: each step evaluates all the gains in the network at once allowing convergence in very few steps (1 to 3).
- The method converges with the whole system expected operating point by only monitoring the received power of traffic.

The following parameters are used:

- **Topology of the DWDM network** – It is necessary to know the types of nodes present and their exact sequence for each end-to-end service in the network in both directions (that can be supposed identical in case of bidirectional connection, but that can be generalized to the case of non-symmetrical link because the algorithm operates separately on each direction

- **Traffic Matrix [$\overline{TM}$]** – This is a matrix with a column per each Amplifier in the network and a row per each configured service: each element of the matrix can be 0 or 1 depending on the presence (1) or absence (0) of a service in the link preceding the amplifier represented in a specific column of the matrix. For "service" we mean a connection from an add node till a drop node
- **Target Received Power [$\overline{TRx}$]** – This is an array which, for each service, contains the desired receiver power defined as a target by design. It can differ on a service basis depending on the type of service (e.g., Bit rate, modulation format, type of receiver …) and on the type of node (e.g., Terminal or fixed OADM (FOADM)).
- **Margin [$\overline{Mo}$ & $\overline{Mr}$]** – These are two arrays which, for each service, contains the desired Margin on the overload ($\overline{Mo}$) and on the sensitivity ($\overline{Mr}$) that, by design, we want to apply. This must consider the confidence in the receiver power measurement (e.g., if the Rx power is measured with a confidence of ±2 dB we can assume a margin of 2 dB on both the overload and the sensitivity).
- **Span Loss rough estimation [$\overline{SLr}$]** – This is an array which, for each span, contains the value of its loss, roughly estimated.
- **Receiver Power [$\overline{Rx}$]** – This is an array which, for each service/traffic path, contains the measurement of the optical power at relevant receiver.
- **Receiver Range [$\overline{Ro}$ & $\overline{Rs}$]** These are two arrays which, for each service, contain the Receiver overload ($\overline{Ro}$) and the receiver sensitivity ($\overline{Rs}$) of each relevant transceiver.
- **Delta Gain Constraint [$\bar{\Delta}$]** – This is an array which, for each amplifier, contains the maximum value of variation of the Gain that, by design, is acceptable to set at each step of the algorithm.

This procedure can be repeated throughout the network's lifetime, either on a regular schedule or whenever a receiver-power alarm indicates potential drift (for example, from aging or other network changes). The goal is to adjust each amplifier so that the average receive power at every node, $\overline{Rx}$, matches its designated target level, $\overline{TRx}$.

These target values ($\overline{TRx}$.) vary by node type or configuration and are supplied externally (for instance, by a network planner or as part of the system design). A FOADM drop node might use one $\overline{TRx}$., while a terminal node uses another. Likewise, different data rates (10G, 25G, 100G, etc.) can have their own $\overline{TRx}$. because they're assigned distinct launch-power settings. By driving all receivers to their specific design-average power, the method not only ensures correct receive levels but inherently aligns other critical optical metrics, such as OSNR, gain tilt, and launch power, to their intended operating points.

An additional benefit of this strategy is that it lets us calculate and apply gain adjustments network-wide in a single pass rather than tuning each node individually. Because the system runs in Constant-Gain mode, any dB change in one amplifier's gain propagates equally to every downstream amplifier and thus to all channels. As a result, the resulting receive power at any

point is simply the sum of the individual gain corrections (in dB) applied along the path. This additive property enables us to predict and enforce the correct power levels for every channel across the entire network simultaneously.

Let $\overline{\boldsymbol{TM}}$ be the traffic matrix where each row represents a traffic service from one node to another in the optical link and each column represents an amplified link node from one terminal to the other. This matrix will have value '1' where the given service is traveling the amplifier of that node, 0' otherwise. So, if the old received powers are $\overline{\boldsymbol{Rx}}_{\boldsymbol{old}}$, the new received power $\overline{\boldsymbol{Rx}}_{\boldsymbol{new}}$, for all traffic paths, is given by Eq. 13, where $\overline{\boldsymbol{Rx}}$ are array vectors, whose dimension is the number of traffic services. $\overline{\boldsymbol{\Delta G}}$ is the array of the gain correction, that is, the unknown to be found and optimized and its dimension is the number of amplified nodes in the link.

$$\overline{\Delta Rx} = \overline{Rx}_{new} - \overline{Rx}_{old} = \overline{TM} \cdot \overline{\Delta G} \qquad \text{Eq. 13}$$

The objective function to be minimized is the distance between $\overline{\boldsymbol{Rx}}_{\boldsymbol{new}}$ and the target powers $\overline{\boldsymbol{TRx}}$. This can be expressed in many ways. Without losing generality, we suggest the use of the squared Euclidean distance (the *Norm-2*) as in Eq. 14 since it is suitable for the *Quadratic Programming Optimizer*, it guarantees convexity and, therefore, a global optimum.

$$\begin{aligned} \text{Objective function} &= \frac{1}{2} \|\overline{Rx}_{new} - \overline{TR}x\|^2 \\ &= \frac{1}{2} \|\overline{TM} * \overline{\Delta G} + (\overline{Rx}_{old} - \overline{TRx})\|^2 \end{aligned} \qquad \text{Eq. 14}$$

In general, the LCQP minimizes a quadratic objective function given equality and inequality linear constraints ($\bar{x}$ is the target array to be found), that is in the standard formulation given by Eq. 15.

$$\min_x \frac{1}{2} \bar{x}^T \cdot \bar{H} \cdot \bar{x} + \bar{f}^T \cdot \bar{x}, \text{such that} \begin{cases} \bar{A} \cdot \bar{x} \leq \bar{b}, \\ \overline{Aeq} \cdot \bar{x} = \overline{beq} \\ \overline{lb} \leq \bar{x} \leq \overline{ub} \end{cases} \qquad \text{Eq. 15}$$

Using the definition in Eq. 16, we can express our *Objective function* in the formalism of Eq. 15, where $\bar{\boldsymbol{H}}$ and $\bar{\boldsymbol{f}}^{\boldsymbol{T}}$ are given by Eq. 17 and Eq. 18, respectively.

$$min\ (\|\bar{A} \cdot \bar{x} - \bar{c}\|^2) = min\ (\bar{x}^T \cdot \bar{A}^T \cdot \bar{A} \cdot \bar{x} - \bar{c}^T \cdot \bar{A} \cdot \bar{x}) \qquad \text{Eq. 16}$$

$$\bar{H} = \overline{TM}^T \cdot \overline{TM} \qquad \text{Eq. 17}$$

$$\bar{f}^T = (\overline{Rx}_{old} - \overline{TRx})^T \cdot \overline{TM} \qquad \text{Eq. 18}$$

The Quadratic Programming formalism includes also linear inequality and equality constraints.

We use the ***inequality constraints*** ($A \cdot x \leq b$) to impose that the new received powers $\overline{\boldsymbol{Rx}}_{\boldsymbol{new}}$ are within the receiver dynamic range that is above sensitivity limits and below overload limits. These constraints are given by Eq. 19 and Eq. 20. Using the same formalism of Eq. 15, the above constraints translate into Eq. 21 and Eq. 22 for the overload constraint (Eq. 19) and Eq. 23 and Eq. 24 for the sensitivity constraint (Eq. 20).

$$\overline{Rx}_{new} = \overline{TM} \cdot \overline{\Delta G} + \overline{Rx}_{old} \leq \overline{Ro} - \overline{Mo} \quad \text{Eq. 19}$$

$$\overline{Rx}_{new} = \overline{TM} \cdot \overline{\Delta G} + \overline{Rx}_{old} \geq \overline{Rs} + \overline{Ms} \quad \text{Eq. 20}$$

$$A = \overline{TM} \quad \text{Eq. 21}$$

$$b = \overline{Ro} - \overline{Mo} - \overline{Rx}_{old} \quad \text{Eq. 22}$$

$$A = -\overline{TM} \quad \text{Eq. 23}$$

$$b = \overline{Rx}_{old} - \overline{Rs} - \overline{Ms} \quad \text{Eq. 24}$$

These elements can be combined into single vectors through concatenation. Although optional, including these constraints is advisable because they guarantee that the solution stays within design specifications and prompt the optimizer to issue warnings if any limits are violated. Similarly, you can define lower- and upper-bound vectors ($\overline{lb}$ and $\overline{ub}$) to restrict the decision variable $\bar{x}$, in our case, the gain adjustments $\overline{\Delta \boldsymbol{G}}$, to a specified range. While these bounds aren't mandatory, they help prevent excessively large corrections.

If the Quadratic Programming solver fails to find a feasible solution, it indicates that the constraint set is inconsistent, revealing deeper issues in the system that warrant investigation. The solver will flag these conflicts, providing valuable feedback for debugging.

In the SEASON project, we tested this gain-optimization procedure on a simulated, four-node filterless ring using MATLAB. We encoded the network topology and all design parameters, and we set the target receive power at each drop node to –7 dBm. This value corresponds to a 3.5 dBm channel power after amplification, minus 7 dB for the 4:1 drop splitter, and minus 3.5 dB for the 2:2 pass-through splitter.

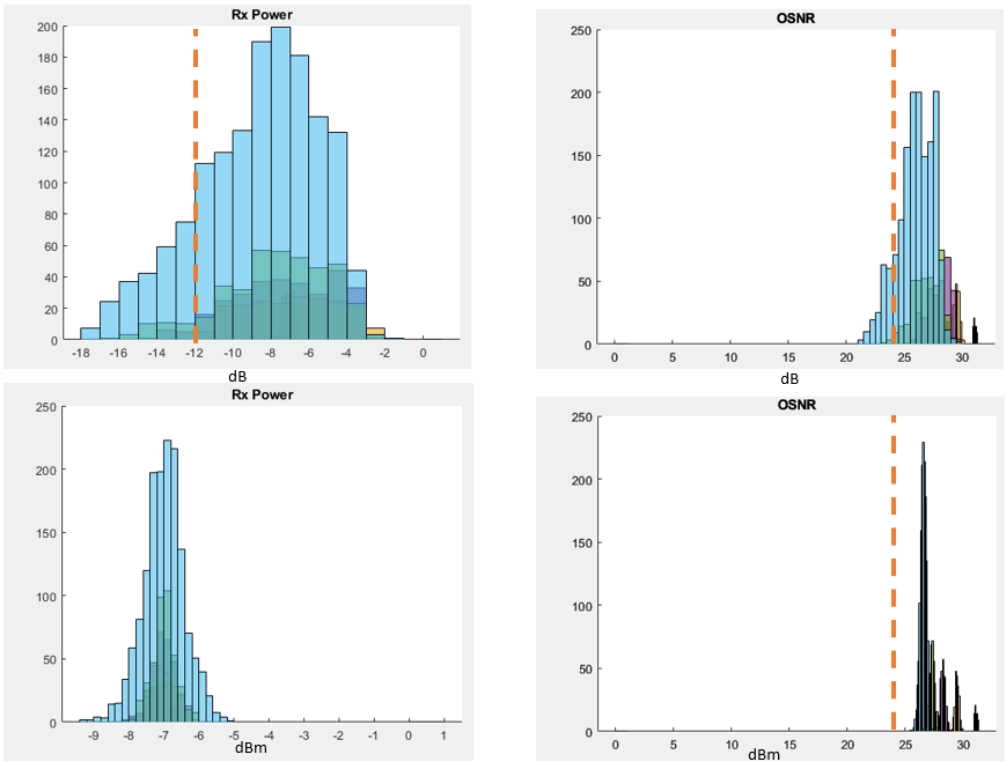

Figure 5.3-2 Gain Optimization – Simulation Results.

We began by modeling the initial gain errors as a uniform random distribution up to ±4 dB, reflecting a ±2 dB uncertainty in power measurements, and by including a statistical model of each EDFA's channel flatness. Figure 5.3-2 illustrates the resulting impact on system performance. The upper histograms show the distribution of received power (left) and OSNR (right) before any correction. Due to the uncorrected gain errors, both parameters exhibit a wide spread—several channels fall outside their target ranges, as indicated by the dashed orange lines representing the design thresholds. This variability demonstrates how amplifier gain inaccuracy can significantly degrade link performance and compromise uniformity across wavelengths.

Once our automatic optimization algorithm is applied, the improvement is evident in the lower histograms. The distributions of both Rx power and OSNR narrow substantially, and all channels are brought back within the acceptable operating window. This confirms that the proposed optimization effectively compensates for amplifier nonuniformities, restoring the system's optical budget and ensuring compliance with design specifications.

6 COST MODEL

In the realm of innovation and technology development, techno-economic analysis (TEA) serves as a critical tool for evaluating the viability and scalability of emerging technologies. At the heart of any meaningful TEA lies a robust and comprehensive capital and operating cost model. This model should not merely be a financial exercise. It is a framework that informs strategic decisions, guides research priorities, and ultimately determines the commercial potential of a technology [Bur17, Cha22]. Moreover, a well-structured cost model enables benchmarking against existing technologies and industry standards. A comprehensive cost model captures all relevant CAPEX and OPEX. By encompassing both fixed and variable costs, the model provides a realistic picture of the financial landscape across different production scales and market scenarios. The robustness of a cost model is equally vital. It must be built on accurate, up-to-date data and validated assumptions.

In summary, a robust and complete cost model transforms TEA from a theoretical exercise into a powerful decision-making tool. It bridges the gap between laboratory innovation and market implementation, ensuring that technological breakthroughs are not only scientifically sound but also economically feasible and strategically aligned with broader societal needs.

Cost modeling is a powerful tool in techno-economic analysis, but it comes with several challenges that can affect the accuracy, reliability, and usefulness of the results.

Table 5.3-1 presents some of the most common challenges, their impact on the TEA and suggestions on how to mitigate them [Fas25]. A critical issue is the availability and quality of data, as reliable cost information for emerging technologies is often limited, outdated, or proprietary, leading to potentially misleading results if not validated. Uncertainty and variability in input parameters further complicate analyses, since fluctuating values can reduce the reliability of projections unless stochastic approaches such as Monte Carlo simulations are applied. The inherent complexity of multi-step systems also presents difficulties, as overly simplified models risk omitting key cost drivers while excessively detailed ones may obscure insights; modular approaches can help balance detail with usability. Scale-up assumptions pose another challenge, as extrapolating from laboratory or pilot data to industrial scale introduces uncertainty that requires benchmarking and sensitivity analysis. Additionally, regional and temporal variability means that cost estimates are not universally transferable, necessitating localized and regularly updated models. Finally, the lack of standardization in methodologies and assumptions across studies hinders comparability, underscoring the importance of following established guidelines and explicitly stating all assumptions.

Table 5.3-1 Challenges of cost modeling, their impact on the techno-economic analysis and mitigation suggestions.

Challenge	Description	Impact	Mitigation
1. Data Availability and Quality	Reliable cost data (especially for emerging technologies) can be scarce, proprietary, or outdated	Inaccurate data can lead to misleading conclusions	Use validated sources, expert input, and sensitivity analysis to test assumptions
2. Uncertainty and Variability	Many input parameters are uncertain or fluctuate over time	Results may not reflect real-world performance or future scenarios	Incorporate stochastic modeling or Monte Carlo simulations to capture uncertainty
3. Complexity of Systems	Complex systems with multiple interdependent steps are difficult to model accurately	Oversimplification can omit key cost drivers; overcomplication can obscure insights	Strike a balance between detail and usability; modularize the model where possible
4. Scale-Up Assumptions	Extrapolating lab- or pilot-scale data to commercial scale introduces uncertainty	Cost estimates may not hold true at industrial scale	Use scale-up factors based on industry benchmarks and include scale sensitivity analysis
5. Regional and Temporal Variability	Costs vary by location and over time (e.g., inflation, market trends)	A model valid in one region or year may not apply elsewhere or later	Localize models and update them regularly with current data
6. Lack of Standardization	Different studies use different methodologies, assumptions, and cost categories	Difficult to compare results across technologies or studies	Follow established guidelines and clearly state document assumptions

Moreover, cost modeling in optical networks has evolved significantly through joint efforts among academia, industry, and network operators. Notable initiatives such as the STRONGEST [STRGST], IDEALIST [IDLST], and Metro-Haul [MTRHL] projects have contributed valuable methodologies and data sets for estimating costs across various segments of the network. In particular, the Metro-Haul project developed a widely referenced cost model for techno-economic analyses in metro-access networks [Her20]

More recently, the B5G-Open [B5GOPN] project introduced a suite of techno-economic models covering long-haul, metropolitan, and metro-aggregation domains. However, those models are standalone and tailored to specific use cases. Additional research published in journals and conferences also presents cost models, though these are typically confined to their respective domains and often rely on distinct methodologies for cost value estimation [Jan23, Sou24, Arp24].

This work proposes a unified and up-to-date cost model designed to span multiple network segments and device types, built upon a consistent and transparent methodology. The work outlines the approach used to develop a comprehensive yet practical cost model for the SEASON

project. Its main goal is to support a wide range of techno-economic analyses within the project, while maintaining enough generality to serve as a reference framework for future studies. By promoting consistency in assumptions and cost parameters, the model contributes to addressing the lack of standardization and enables other researchers to adopt the same methodology for complementing or updating the model's data.

- **Section 6.1** details the approach adopted for sourcing reliable data from publicly available online resources and expert consultations (addressing Challenges 1 and 2). It also discusses the trade-offs between model complexity and accuracy (Challenge 3), and highlights the regional specificity of the data employed (Challenge 5).
- **Section 6.2** outlines a recommended methodology for extrapolating existing data to account for non-commercial or future equipment, such as higher-capacity transceivers.
- **Section 6.3** introduces a modeling framework for estimating the cost and power consumption of doped-fiber optical amplifiers operating in transmission bands beyond the conventional C-band.

The final outputs (including cost estimates and power consumption figures) are compiled in a structured spreadsheet, which is made publicly accessible as described in Section 7.8.

6.1 Commercially Available Devices

One of the most common challenges in cost modeling is the lack of accurate and reliable data (Challenge 1). Since data forms the foundation of any cost model, insufficient or imprecise information can lead to misleading or incorrect results. Therefore, prices for commercially available systems were gathered from web sites that provide them publicly, and cross-validated with data provided by various partners involved in the SEASON project but cannot be publicly referenced. It should be noted right away that prices can be found easily for widely available items subject to open competition (such as transceivers, packet switches and servers). However, for equipment that is less available, more expensive, and that require customization and ad hoc offerings (for example, optical nodes and transponders), it is much more difficult to find prices online.

For transceivers and switches, the primary source used to obtain pricing data was [FS], chosen for its extensive catalog of optical and packet networking devices, competitive pricing, and user-friendly interface for querying products. Additionally, [APPnet] was consulted to gather transceiver pricing data, enabling comparisons and the calculation of average values. For compute and storage components, pricing information was sourced from [DELL] and supplemented through expert consultations. The data collection and analysis efforts resulted in the cost tables presented in the following sections. It is important to note that the cost values provided in this article are based on surveys conducted during the first quarter of 2025. As is commonly known, prices may vary over time due to various factors. Moreover, when defining device or system items, some simplification was necessary. Similar items often exist in multiple variants, which may differ significantly in price. To address this, we defined a ``typical'' item (an

abstraction representing the average market offering for a given device or system). Each typical item includes an estimated average price, expected power consumption, and representative performance parameters. These abstractions do not correspond to specific commercial products but rather serve as generalized models for analysis.

The cost values exclude value-added tax (VAT) and are expressed in euros, assuming a European country context. This approach ensures consistency and applicability across different European regions. To further mitigate temporal, geographical, and source-related variations, costs are normalized and presented as ratios relative to a defined cost unit (c.u.) of five thousand euros, rather than as absolute euro values. These normalized values are then compared and validated against internal data from project partners.

The objective of the cost model is to be as comprehensive as possible while maintaining accuracy and avoiding excessive complexity. Accordingly, the model focuses on components based on mature, commercially available technologies, grouped into four categories: transceivers, amplifiers, Layer 3 switches and routers, and compute/storage devices. Components for which reliable data could not be obtained are excluded from the model. Similarly, elements with negligible cost impact are not considered.

The following subsections describe the methodology for gathering the information for each component category.

6.1.1 Transceivers

The transceivers are divided into three categories based on their application: grey transceivers, colored transceivers for WDM systems, and coherent DSCM transceivers which allow subcarrier-level capacity allocation, enabling coherent P2MP communication. The application categories are further divided into reference reach categories as presented in Table 6.1-1 where the assignments are as follows: SR is Short Reach, LR is Long Reach, ER is Extended Reach, ZR is Extended Range/Zero Dispersion, ZR+ is a variant of Extended Range/Zero Dispersion with longer reach, LH is Long Haul.

Table 6.1-1 Transceiver reach categories.

Reach Class	Maximum distance
SR	300 m
LR	10 km
ER	40 km
ZR	80 km
ZR+	450 km
LH	2000 km

6.1.1.1 Grey Transceivers

Table 6.1-2 presents the cost and power consumption of grey transceivers. Data rates from 10G to 400G and three reach categories are selected, SR, LR, and ER. SR grey transceivers use multimode fiber, while LR and ER use single mode fiber. The types of the form factor range from small form-factor pluggable (SFP) to quad SFP- double density (QSFP-DD), depending on the data rate and power consumption. Gray transceivers transmit in the O-band unless otherwise indicated (only one case, 10G ER C-band).

Table 6.1-2 Grey transceivers cost (in c.u.) and power consumption (in W).

Transceiver	Cost [c.u.]	Power [W]
10G SR	0.004	1
25G SR	0.008	1
50G SR	0.054	1.5
100G SR	0.02	2.5
400G SR	0.04	10
10G LR	0.005	1
25G LR	0.012	1.2
50G LR	0.152	3.5
100G LR	0.08	3.5
400G LR	0.16	9
10G ER	0.008	1.5
10G ER C-band	0.026	1.5
25G ER	0.056	2
50G ER	0.36	4.5
100G ER	0.4	4.5
400G ER	1.4	10

6.1.1.2 WDM Transceivers

Table 6.1-3 reports a selection of WDM colored transceivers for different reaches. All transceivers are based on MSA form factor and are tunable in the extended C-band (~4.8 THz), enabling DWDM multiplexing (48 to 96 channels), except those specified otherwise, which are 10G ER and ZR (CWDM, up to 12 channels, and fixed), and 100G ZR a (WDM, 20 channels, and fixed). The 10G models use non-return-to-zero (NRZ) detection, 100G ZR a coarse WDM (CWDM) fixed uses pulse amplitude modulation with 4 levels (PAM4), while the remaining ones are all coherent. For 100G ZR+ transceiver, power consumption is specified for the QSFP-DD form factor. The same device with a smaller form factor (e.g., QSFP28) consumes much less, around 6 W, with comparable performance and cost. Data for 400G ZR+ is taken from a variable data rate transceiver that can be configured from 100 Gbit per second (Gbps) to 400 Gbps. When configured at 100 Gbps, the reach extends to LH (<2000 km).

Table 6.1-3 WDM colored transceivers cost (in c.u.) and power consumption (in W).

Transceiver	Cost [c.u.]	Power [W]
10G ER (CWDM fixed)	0.044	1.7
10G ZR a (CWDM fixed)	0.08	2
10G ZR b	0.12	2.5
10G ZR+	0.18	3
100G ZR a (WDM fixed)	0.48	7
100G ZR b	0.8	5.5
100G ZR+ (QSFP-DD)	1	17
400G ZR	1.2	18
400G ZR+ (100G LH)	1.5	22.5

6.1.1.3 DSCM Coherent Transceivers

Coherent transceivers based on DSCM represent a new type of device that enables flexible networking by allowing P2MP communication [Wel23]. The target specifications for this type of device are defined in the Open XR forum [XR_forum].

Currently, commercial offerings are limited to 100G modules (typically used as leaf modules) and 400G modules (used as hub modules), both within the ZR+ reach class. As no public price lists for these specific devices were found, pricing estimates were inferred from similar coherent single-carrier modules. For each module type, a range of values is provided for both cost (in c.u.) and power consumption (in W) to account for potential future variability. The minimum price for a DSCM coherent module was assumed to match that of a comparable single-carrier coherent device with the same bit rate and reach class. The maximum price was set to 25% above this baseline. A similar approach was applied to power consumption, with the maximum value set 10% higher than the minimum.

Table 6.1-4 presents 100 and 400 Gbps modules with 4 and 16 subcarriers (SC), respectively. Each subcarrier supports 25 Gbps of traffic and the data rate of a point-to-point (P2P) or P2MP device can be configured with this granularity. Two reach classes, ZR and ZR+, are proposed, corresponding to two different price levels and with slightly different power consumption. Although the coherent technology currently used complies with class ZR+ (with amplification), given the potential applications in the access-metro aggregation segment, the idea was to also introduce the more economical class ZR, which can be achieved with a lower-cost (and lower performance) implementation. It is not certain whether the market will offer such devices in the future, but the combination of a competitive cost with the advantages of P2MP networking could meet the market's favor.

Table 6.1-4 DSCM transceivers cost (in c.u.) and power consumption (in W).

Transceiver	Cost [c.u.]	Power [W]
4 SC 100G ZR	0.80 - 1.00	5.50 - 6.05

4 SC 100G ZR+	1.00 - 1.25	17.00 - 18.70
16 SC 400G ZR	1.20 - 1.50	18.00 - 19.80
16 SC 400G ZR+	1.50 - 1.88	22.5 -24.75

Out of all the transceivers presented in these sections, the coherent DSCM ones are the ones subject to the greatest uncertainty regarding prices and actual market availability. However, we thought it is important to include them because this information could be key to realizing meaningful comparisons of transport solutions for access and metro networks.

6.1.2 C-band Amplifiers

Cost and power consumption estimates for optical amplifiers operating outside the C-band are discussed separately in subsection 6.3 , as these bands exhibit distinct characteristics. While L-band amplifiers show less uncertainty due to the availability of some commercial devices, future cost values will still be influenced by production volumes and ongoing technological development. To streamline the model, we reduced the number of devices modeled. Among the factors influencing amplifier pricing, the most significant are the amplification band and the number of amplification stages. Consequently, the model includes cost and power consumption data for both standard C-band (4.0 THz) and extended C-band (4.8 THz) EDFAs, while excluding parameters with lower impact on pricing—such as noise figure, gain, gain flatness, maximum output power, and number of operating modes since these characteristics tend to be similar across available devices (typical amplifiers used in metro/regional networks).

The cost and power consumption values used in the model represent average values derived from publicly available devices. The values are summarized in Table 6.1-5.

Table 6.1-5 Amplifiers cost (in c.u.) and power consumption (in W).

Amplifier	Cost [c.u.]	Power [W]
C-band Single Stage	0.43	15
C-band Dual Stage	0.86	20
Extended-C-band Single Stage	0.48	15
Extended-C-band Dual Stage	0.96	20

6.1.3 Layer 3 Switches and Routers

In view of the benefits of the IP-over-WDM architecture, there is a significant uptick in its adoption in current and future network deployments. In this architecture, coherent pluggable transceivers are hosted directly in switches or routers. This further motivates to jointly consider the L0 up to L3 equipment to be deployed.

The data collected and analyzed for packet switches resulted in Table 6.1-6, which presents devices categorized as Layer 3 switches. It should be clarified that there are different types of

packet switching nodes and the two main categories are the so-called Layer 3 switches and the IP routers. In fact, Layer 3 switches and IP routers both perform network layer (Layer 3) packet routing, but they differ significantly in their roles, hardware design, and application scope within networks. Layer 3 switches are specialized devices optimized for high-speed internal local area network (LAN) routing and switching, while IP routers are designed to connect different networks and provide advanced wide area network (WAN) and edge capabilities.

The price source used in this case is [FS] and the devices are of grey or white box type Layer 3 switch for networking. They are equipped with basic software (SW), i.e., the operating system, but do not include the SW necessary for routing and management (which must be quoted separately and in addition). The overall price of Layer 3 switch ready to operate in a network could be much higher than the cost of the hardware itself (e.g., by a factor of 2 or more). An all-inclusive model is beyond the scope of this work.

Providing pricing for routers is more difficult as public pricing is less available and is more subject to buyer-seller negotiation and volume. A router with the same raw switching/forwarding capacity as a Layer 3 switch can typically cost 2 to 5 times as much (and sometimes significantly more), depending on features, redundancy, and service contracts. A factor of $\times 4$ can be taken as a rule of thumb.

The Layer 3 switches in Table 6.1-6 are supplied in seven models of different size (from the smaller extra extra small, EES, to the bigger extra extra large, XXL) in which the maximum cumulative capacity of interfaces it can host is specified. For example, the S model, which has a total capacity of 1.6 Tb/s, can be equipped with 4 interfaces at 400 Gbps, or with 16 interfaces at 100 Gbps (or with other combinations with a sum of 1.6 Tb/s). At full load, and in theory, the switch is capable of switching double the total capacity of interfaces since they are bidirectional (i.e., they allow the nominal data rate to be transmitted in both directions on two separate ports). It must be said that the switch will always switch a smaller fraction of the maximum traffic it can receive from its interfaces, often a significantly smaller fraction (typically less than 50%).

Table 6.1-6 Layer 3 Switches cost (in c.u.) and power consumption (in W).

Switch Model	Cost [c.u.]	Power [W]
EES - 400 Gbps	0.6	110
ES - 800 Gbps	0.9	130
S - 1.6 Tb/s	1.2	260
M - 3.2 Tb/s	1.5	520
L - 6.4 Tb/s	2.3	850
XL - 12.8 Tb/s	4	1500
XXL - 25.6 Tb/s	6	2400

6.1.4 Compute and Storage

Just as Layer 3 switches are useful for evaluating IP over WDM architectures, compute and storage servers are useful for evaluating network solutions that use virtualized Telco and Service functions (such as cloud-based mobile radio access network (RAN), among others). Indeed, for a comprehensive evaluation of all the components involved, evaluation of optical, IP networking, and compute and storage levels is required. A selection of HW elements used in the cloud as compute, storage, and networking card elements to support Telco and Service functions are included in Table 6.1-7. The first five rows include server configurations classified according to their size (XS with single 24-core CPU, S with two 24-core CPUs, M with two 32-core GPUs) and their memory and GPU cards, if present. Servers with distributed storage have storage units with a total capacity of 10 TBytes. Server configurations are in reality countless and subject to rapid evolution over time (think, for example, of the growing need for GPUs for AI applications), so the selection is made purely for illustrative purposes and to allow for a high-level economic evaluation of simple scenarios.

The first two servers are suitable for RAN applications and specifically for implementing virtual distributed unit (vDU) functions. Their price difference is due not only to their compute capacity (1 or 2 processors), but also to their network cards which have different data rate ports (25G for the first, 100G for the second). Size M servers are used for standard data center applications or for Telco functions like virtual centralized unit (vCU) or user plane function (UPF). Other items in the table are a graphics processing unit (GPU), whose reference model is Nvidia L40, and cards, standard network interface cards (NIC) and more powerful data processing unit (DPU), which are cards with additional processing capabilities that can relieve the central processing unit (CPU) and speed up processing by performing tasks such as processing, routing, and forwarding functions.

Table 6.1-7 Compute and Storage Devices cost (in c.u.) and power consumption (in W).

HW system	Price [c.u.]	Power [W]
SERVER XS with accelerator, DPU and NIC @25G	1.200	1000
SERVER S with DPU and NIC @100G	5.280	1200
SERVER M basic for compute	2.800	1000
SERVER M for distributed storage	4.040	1100
SERVER M with GPU and distributed storage	6.160	1450
GPU	1.760	350
NIC 2x200G	0.400	30
Super NIC 1x400G	0.400	200
DPU 2x25G	0.200	350
DPU 2x200G	0.700	450

6.2 Modeling the Cost in Function of Equipment Capacity

Technology continually becomes cheaper each year due to a confluence of factors, primarily driven by rapid innovation, economies of scale, and fierce market competition. A cornerstone of this trend is Moore's Law, an observation that the number of transistors on an integrated circuit doubles approximately every two years, leading to exponential increases in computing power and efficiency at a proportionally decreasing cost per unit. This continuous miniaturization and improvement in chip manufacturing means that devices can become more powerful while using less material and energy. Furthermore, as new technologies move from niche markets to mass production, manufacturers benefit from significant economies of scale.

Optical communication primary equipment such as transceivers and routers has been improving in terms of capacity, power efficiency, size, and cost. For example, nowadays coherent pluggables can transmit up to several 100G of data while costing less and consuming less power per transmitted bit.

Operators must carefully evaluate their CAPEX deployment to ensure profitable operation and even increase their profit margin. Therefore, accurate planning for the future is crucial, although the future is uncertain. With appropriate modeling frameworks (e.g., Stochastic optimization, Monte Carlo analysis, sensitivity and robustness analysis), the risks and uncertainties can be modeled and then be mitigated with techniques such as margin estimation, flexible networking, softwarization, and adding redundancy.

Regarding technological evolution uncertainty, it is useful to predict the pace and direction of new optical technologies (e.g., new data rate, higher capacity routers) and their impact on network design in multi-period planning scenarios. Trend analysis and extrapolation using e.g., regression analysis can provide insightful information by analyzing historical pricing data for various optical components (e.g., transceivers of different data rates). Regression analysis is a statistical method used to model the relationship between a dependent variable and one or more independent variables. It helps to understand how the value of the dependent variable changes when any one of the independent variables is varied, allowing for prediction and forecasting.

Here we only focus on transceiver prices, but the method and technique can be applied to other equipment as well. The cost of higher data rate transceivers can be predicted by fitting a curve to available data and then extrapolating for those that do not exist yet. The price of a transceiver can be expressed by Eq. 25, where $P(r)$ is the price of transceivers with data rate r, and k_0, k_1, and k_2 are coefficients while s is an exponent smaller than one. k_0 term can model the fixed cost of transceivers while $k_1 r$, and $k_2 r^s$ can model the linear and exponential decay trends, respectively.

$$P(r) = k_0 + k_1 r + k_2 r^s \qquad \text{Eq. 25}$$

Since there is not much data and for the sake of simplicity, we use a simpler curve and only keep the last term (which is the most important one) as shown in Eq. 26, where k, and s are the

unknowns and must be estimated. Nonlinear Least Squares (NLS) fitting is a powerful statistical technique used to find the best-fit parameters for the above-mentioned nonlinear model.

$$P(r) = k.r^{s} \quad \text{Eq. 26}$$

We use a normalized data set with 50 samples of which only 12 of them are original and the rest are fabricated. Figure 6.2-1 shows the data set (blue points), and the optimal curve fitted to the data. The developed mathematical framework can be used to estimate k, s and consequently predict higher data rate transceiver prices while providing average values and uncertainty ranges. Also, the 90% confidence intervals for k, and s are reported. Obviously, the more data, the more accurate and less uncertain the model is.

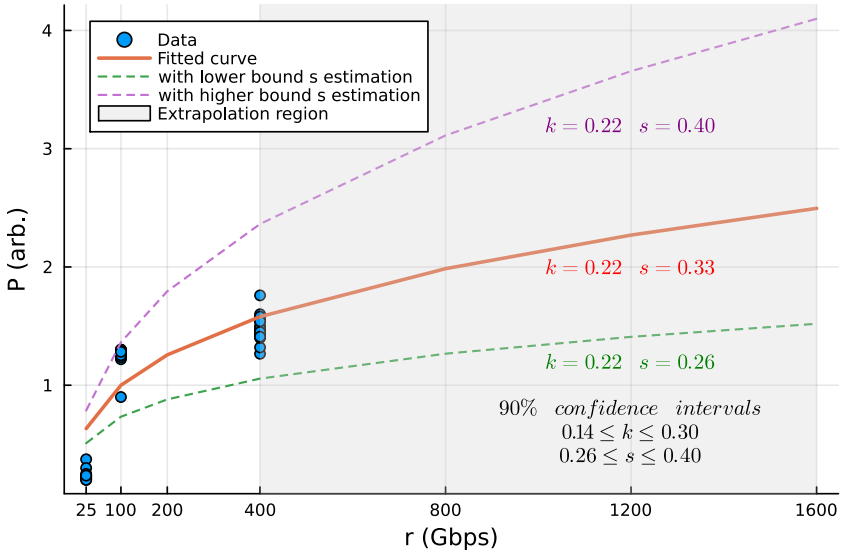

Figure 6.2-1 Curve fitting to a data set of transceivers price and the confidence intervals.

We also used the same modeling approach for fitting a curve to the price of routers by collecting 15 data points of capacity-price from [FS]. Figure 6.2-2 shows the data collected as scatter blue circles and the red line is the fitted curve obtained by NLS. As can be seen, higher capacity routers are more cost efficient as the price/capacity ratio decreases with the increase in capacity. These models can help run more accurate and consistent techno-economic analysis. Recognizing the inherent sensitivity and scarcity of public pricing data, we proposed that this mathematical framework can bridge the gap, offering a robust benchmark even with limited information.

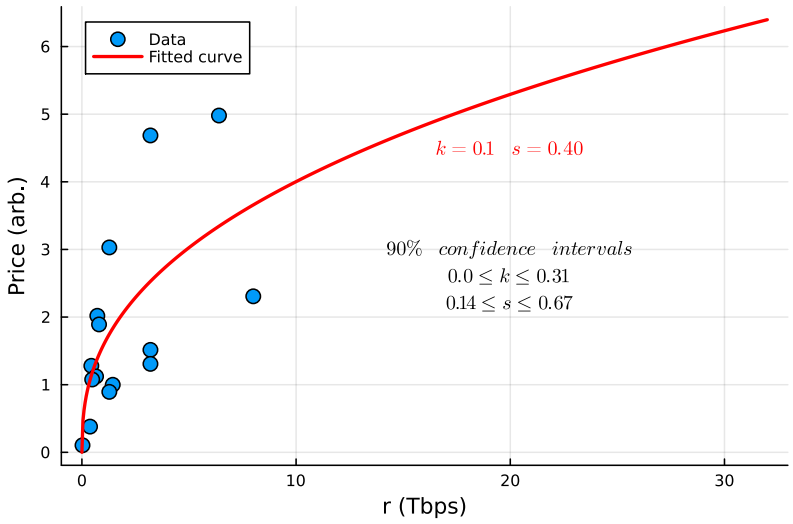

Figure 6.2-2 Curve fitting to a data set of routers price and the confidence intervals.

6.3 An Approach to Estimate the Cost of Optical Systems for O, E, S and L Bands for MB Systems

Optical amplifiers are critical components in long-haul and metro optical networks [Rap22], enabling signal regeneration without electrical conversion. While the C-band remains the most widely deployed due to its favorable gain characteristics and mature component ecosystem, spectrum expansion into adjacent bands is necessary to accommodate future capacity requirements (the widely deployed ITU-T G.652D fiber offers a low-loss transmission spectrum spanning approximately 54 THz from the O- to the L-band [Fer20]). However, realizing this expansion requires addressing additional challenges in terms of cost and power efficiency.

The cost of optical amplifiers outside the C-band is generally higher due to several factors. For instance, the mature, reliable, and cost-effective optical amplifiers widely deployed rely on erbium-doped amplification technology. However, this doping material is only suitable for two of the wavelength bands, namely the C–band and the L–band. Thus, new amplification technologies have to be developed to fully exploit the potential of MB transmission. Building these new amplifiers is complex (alternative doping materials are required) and faces limitations in terms of amplification efficiency and lack of large volume production. For instance, the existing component ecosystem is optimized for C-band operation and, hence, components for other bands, such as pump lasers, isolators, and wavelength filters are either not manufactured or only manufactured in small scale, leading to higher unit costs. Non-C-band amplifiers also suffer from higher cost due to lower amplification efficiency and more complex designs [Rap24]. For instance, L-band EDFAs require longer erbium-doped fiber lengths and dual-stage configurations to achieve comparable gain and noise performance, increasing both material and

integration costs, whereas S-band Thulium-doped amplifiers require more complex pumping schemes to avoid self-termination [Aly05].

To account for future uncertainties, ranges of cost and power consumption values are defined. Table 6.3-1 summarizes the cost and power consumption increment of MB amplifiers compared to their C-band counterparts. The cost increment primarily reflects the anticipated lower production volumes of amplifiers designed for transmission bands other than C-band.

Table 6.3-1 Cost and Power increment for MB amplifiers in relation to the C-band.

Amplifier Band	Cost Increment	Power Increment
L	5% - 15%	5% - 10%
S	25% - 50%	5% - 10%
E	30% - 50%	160% - 175%
O	15% - 30%	60% - 70%

EDFAs can be effectively used in the L-band, resulting in cost and power consumption values that are comparable to those of C-band amplifiers. A cost increment of approximately 5% to 15% is expected due to reduced production volumes, while power consumption is projected to be 5% to 10% higher, attributed to lower power conversion efficiency [Rap22]. For other transmission bands, the model does not incorporate the currently elevated costs associated with emerging technologies, under the assumption that these challenges will be mitigated over time. Instead, the increments reflect differences in production volume and amplification efficiency relative to C-band amplifiers. The model considers multiple production volume scenarios for each band to account for varying adoption rates. Among the bands analyzed, the L-band is expected to see the highest demand for amplifiers due to its similar characteristics and performance to the C-band. This is followed by the O-band, which is anticipated to be more relevant in metro/access networks, the S-band due to its spectral proximity to L and C, and finally the E-band. Power consumption increments are estimated based on differences in power conversion efficiency compared to C-band amplifiers [Rap22].

7 DATASETS FOR NETWORK STUDIES AND AI/ML TRAINING

This section presents the datasets generated by the consortium within the framework of SEASON project. Datasets, that are publicly available, have been specially designed to be used for (but not restricted to) network studies and AI/ML model training.

This Section extends the collection of datasets identified in milestone MS2.3 (June 2024). The chapter starts with a summary of the datasets, including the links to openly access the data. Then, a brief description of each dataset is provided, including relevant details and illustrative numerals and figures, for the sake of a quick understanding of the data. Note that additional exhaustive details about datasets are included in the form of description and readme files in each of the provided links. Finally, a methodology to generate novel synthetic traffic data closes the chapter.

7.1 DATASETS

Table 7.1-1 lists the generated datasets, identifying the type of data, a short abstract and the URL link where each dataset and its related info can be found. The list shows great diversity of data types, including few classes of data plane samples and measurements, network topologies, packet traffic traces, control plane records, and cost model coefficients.

Table 7.1-1 Summary of generated datasets.

Name	Type	Abstract	Link
Measurements of Single Carrier and DSCM Optical Signals	IQ Optical Constellation Data	IQ optical constellation samples obtained by simulation for QPSK, 16QAM and 64QAM optical signals.	https://doi.org/10.34810/data1407
Optical Power and QoT Estimator for MB Optical Networks	Optical Power Data	Synthetic optical power data generated by advanced ML-based GN/EGN models.	https://doi.org/10.5281/zenodo.17192226
Optical Fiber Span Data	QoT Data	Simulated QoT measures in terms of GSNR for optical spans in EONs.	https://github.com/gonz-mart/OpticalFiberSpan_MLmodeling
SEASON Access-Metro Reference Network Topology	Network Topologies & Traffic	Anonymized access-metro topologies and related parameters including reference traffic.	https://zenodo.org/records/16737162
Simulated Flow Traffic Data	Packet Traces	Simulated network flow traffic traces based on given link capacity, load, and packet/flow characteristics.	https://github.com/josetilos/FlowTraceSimulator

TAPI Streaming over MQTT	Control Plane Data	Traces generated by an optical controller following a Linux Foundation TR-548 TAPI Streaming approach.	https://zenodo.org/records/17189453
Cost Model	Model Coefficients	A generalized cost model for techno-economic analysis in optical networks	https://zenodo.org/records/17304012

7.2 IQ Optical Constellation Data

The dataset contains IQ optical constellation samples and QoT metrics obtained by means of simulation for QPSK, 16QAM and 64QAM optical signals based on both single carrier and digital subcarrier multiplexing. Two optical network scenarios considering wavelength selective nodes and filterless nodes have been configured in a MATLAB-based simulator and reproduced to obtain and collect the aforementioned measurements.

7.2.1 Description

Figure 7.2-1 depicts the setup of the numerical simulations. A MATLAB-based simulator of a digital coherent system is employed to reproduce both SC and DSCM signal transmissions. In both cases, we consider a dual polarization Nyquist-shaped single-λ channel with a roll-off factor of 0.06, operating at 64 GBd symbol rate. The DSCM signal is composed of 16 DSCs all operating at 4 GBd spaced with a guard-band of 100 MHz. The pulse propagation along the fiber is modelled by solving the Manakov equation using the split-step Fourier method with a propagation step-size of 1 km including effects such as fiber loss, arbitrary fiber birefringence, group velocity dispersion, polarization mode dispersion (PMD) and self-phase modulation. Each link consisted of 80-km SSMF spans characterized by fiber loss 0.21 dB/km, dispersion 16.8 ps/nm/km, PMD value 0.04 ps/km1/2 and nonlinear coefficient 1.14 W-1km-1.

The EDFAs' noise figure is 5 dB whereas its gain is set to compensate for the span loss and to restore the power to the launch one which is near-optimal (1 dBm for both SC and DSCM signals). Note that each subcarrier has a launch power of -12.2 dBm and that at the transmitter side we assume the same background noise both for SC and DSCM signals leading to a normalized OSNR ratio 49 dB.

In this analysis, we consider three m-quaternary modulation formats with index m equal to 4, 16 and 64. Moreover, we considered two types of network nodes: wavelength selective (WS-nodes), also known as ROADMs and filterless nodes. In particular we consider Route-and-Select (R&S) ROADMs containing two filters: one 1xN route wavelength selective switch (WSS) steering the incoming wavelength to the drop port or to another degree and one Nx1 select WSS that can block unused wavelengths and add new ones. Differently the F-nodes have the so-called Drop & Waste design: an optical splitter broadcast copies of every channel to every network degree and to the drop port whereas an optical combiner inject all the channels in the fiber.

Note that this architecture is not able to erase channels in order to reuse a wavelength. Finally, we assume that both the ROADMs and the filterless (FL) nodes insert similar losses (~ 8 dB). To emulate the WSS, we consider filter shapes that were obtained by characterizing four 1x4 WSSs. These latter have a –3 dB optical bandwidth of approximately 71 GHz and were concatenated with a random pattern in the simulator to mimic an optical lightpath having similar but not coincident filter shapes in each WSSs.

For each configuration, we generate 20 samples and for each sample we store the in-phase (I) and quadrature (Q) optical constellations of the X polarization composed by 218,450 IQ symbols and optical spectra with 150 MHz resolution. Finally, the BER and the Q factor were evaluated over 1,747,600 symbols for 16 QAM and 64QAM (i.e., over 6,990,400 bits and 10,485,600 bits, respectively) and over 3,495,200 symbols for QPSK (6,990,400 bits).

At the receiver side, we considered digital signal processors (DSP) performing ideal chromatic dispersion compensation, ideal carrier phase recovery and equipped with a MIMO equalizer based on least mean squares algorithm with optimized parameters.

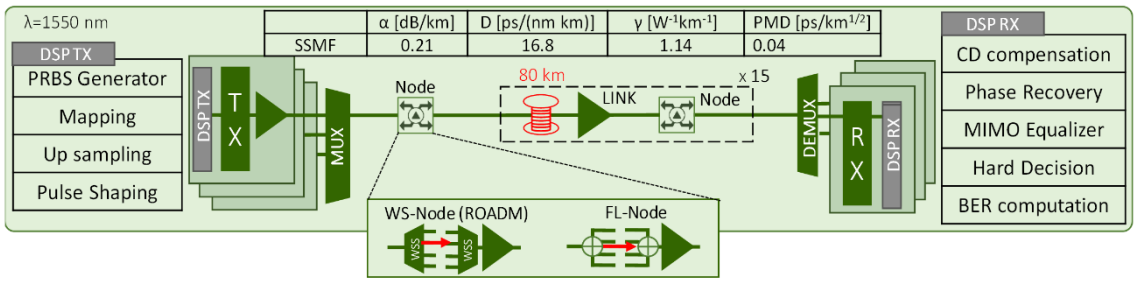

Figure 7.2-1 MATLAB-based simulator architecture.

We provide data having the following five configurations:

a) Considering WS-Nodes for DSCM and SC signals having 16QAM as MF.

b) Considering WS-Nodes for DSCM and SC signals having 64QAM as MF.

c) Considering WS-Nodes for DSCM and SC signals having QPSK as MF.

d) Considering FL-Nodes for DSCM and SC signals having 16QAM as MF.

e) Considering FL-Nodes for DSCM and SC signals having 64QAM as MF.

7.2.2 Illustrative Figures

In particular, Figure 7.2-2 shows the Q factor in dB collected at the node output as function of the number of crossed network nodes for all these configurations. For each of these configurations, we collected 600 IQ Optical Constellations. Namely we collected 20 IQ constellation samples both at the input and at the output of 15 cascaded Network Nodes. Moreover, Figure 7.2-3 shows for illustrative purposes the density diagrams of three IQ

constellation samples collected from the X polarization of signals having different modulation formats (i.e., QPSK, 16QAM and 64QAM).

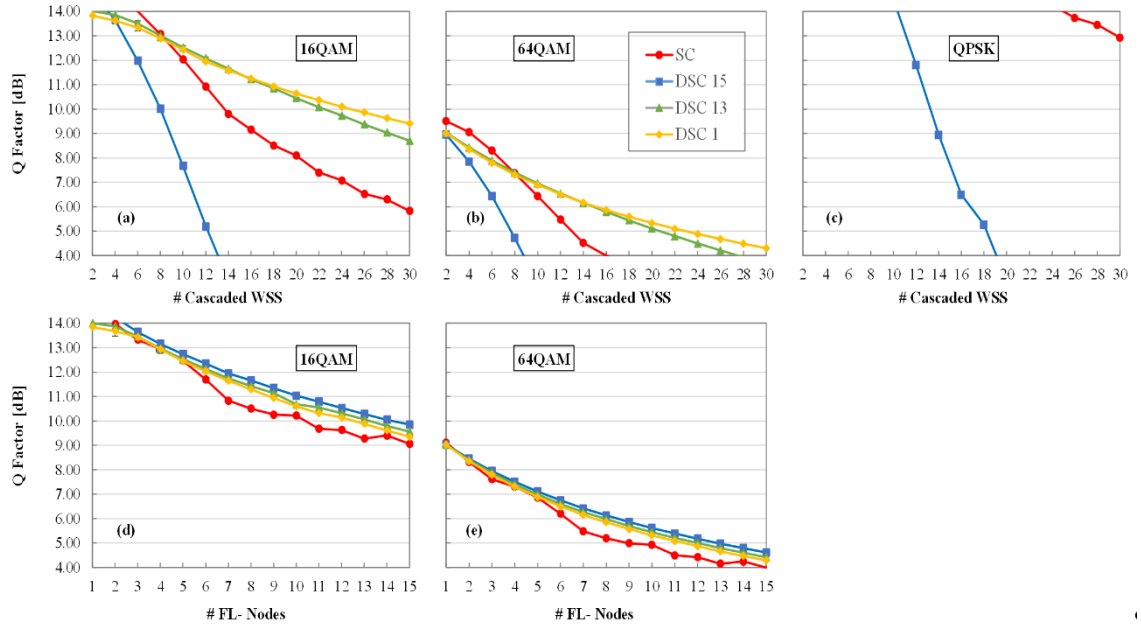

Figure 7.2-2 Q factor as function of the number of network node crossed for different modulation formats.

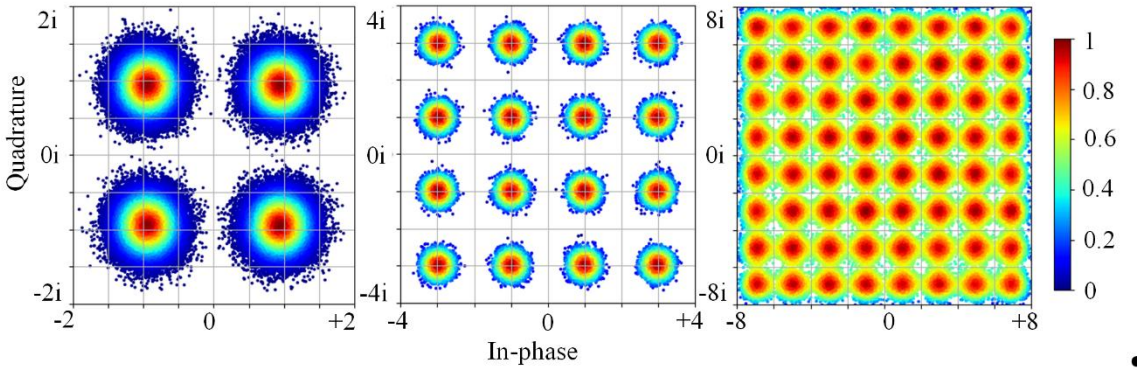

Figure 7.2-3 IQ Constellations Samples of QPSK, 16QAM and 64QAM Signals.

This data is currently being used for AI/ML training purposes. Some models have been already trained and used for optical digital twin modelling and decision-making purposes [Dev24]

7.3 Optical Power Data

This work employs a semi-closed form model (CFM) based on the (enhanced) Gaussian noise (GN/EGN) analytical equations to estimate the GSNR, generating datasets. ML models are then trained using these datasets to accurately estimate optical power and GSNR profiles, achieving errors below 0.04 dB for power and 0.1 dB for GSNR in 99% of cases. The use of ML models is justified by their computational efficiency, aiding in online network management. We adopt the

ML models in a power optimization algorithm that maximizes network total capacity. The power optimization using the ML models take 25-50 times lower time, while resulting in a maximum of 0.1 dBm error versus the analytical semi-CFM.

7.3.1 Description

To generate the synthesized datasets, we utilize the advanced ML-based GN/EGN model from, which has been validated through both the split-step Fourier method and experimental testing. This model incorporates essential features such as dispersion and modulation format correction terms, ensuring a balance between accuracy and speed in estimating GSNR for each span. Since the trained ML model should be able to provide QoT estimations in different scenarios, it needs to be exposed to a wide range of network configurations and trained accordingly. Therefore, a dataset based on diverse conditions and random parameters has been created. Four primary features are considered to generate the scenarios, including span length, loading factor (percentage of activity of each band), modulation format of each channel, and launch power. Each scenario is characterized by these four features. This is in contrast with previous studies, where authors generated biased datasets, the synthesized datasets in this work are topology independent.

Two ML-assisted methods are proposed for power and GSNR estimation. Both methods utilize two powerful and well-known ML models, namely gradient boosting and neural network (NN). Regarding the power estimation, the power at the end of each span can be estimated using the proposed ML model, replacing the time-consuming solution of coupled differential Raman equations, particularly burdensome for MB EONs. For GSNR estimation, we estimate the GSNR of each span using the three proposed ML models instead of relying on the time-consuming semi-CFM. The features for power and GSNR estimation are detailed in Table 7.3-1.

Table 7.3-1 Features for ML models.

No	Parameter
1	Span Length
2	Channel Launch Power
3	Total launch power of all channels
4	Channel position
5	Percentage of activity of all channels
6	Percentage of activity of channels to the left of the channel under test
7	Percentage of activity of channels to the right of the channel under test
8	Idle channels to the left of the channel under test
9	Idle channels to the right of the channel under test
10	Percentage of activity of the first quarter channels
11	Percentage of activity of the second quarter channels
12	Percentage of activity of the third quarter channels
13	Percentage of activity of the fourth quarter channels
14	Modulation format used in the channel
15	Maximum modulation formats used

16	Minimum modulation formats used
17	Average modulation formats used
18	Standard deviation of modulation formats used
19	Channel power at the end of the span

7.3.2 Illustrative Figures

To show the applicability and usefulness of the data for AI/ML training purposes, Figure 7.3-1, Figure 7.3-2 and Figure 7.3-3 show some illustrative results of the gradient boosting and two NNs with different architectures (with 1 and 2 hidden layers respectively) considered as ML models.

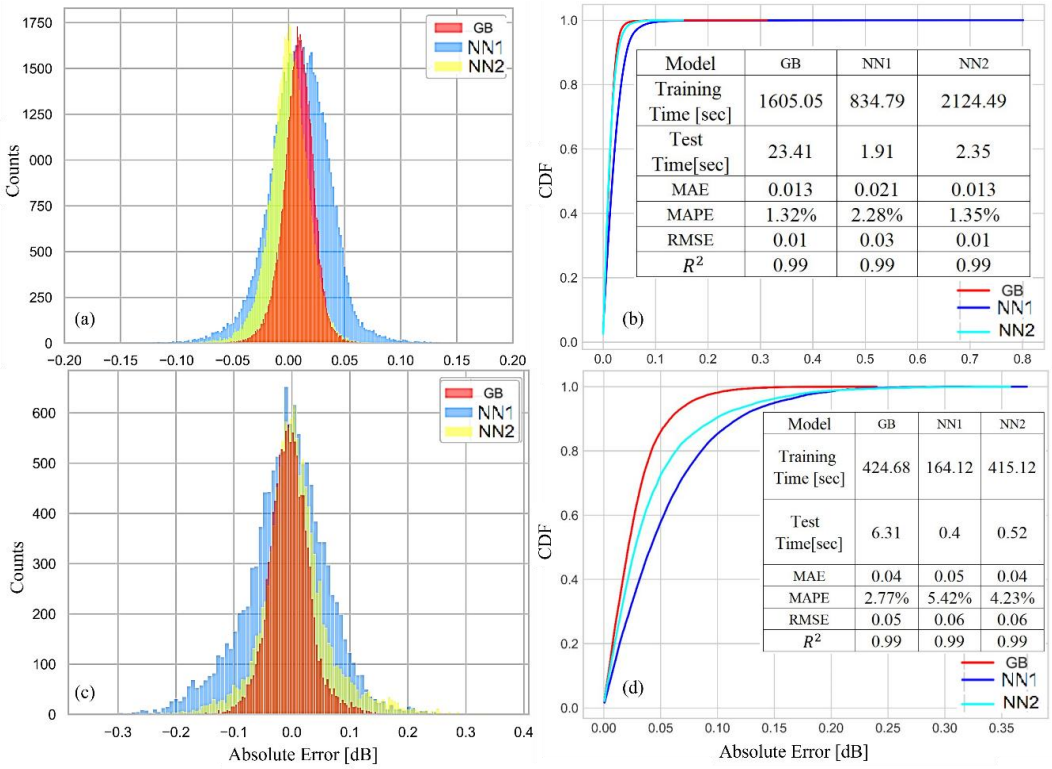

Figure 7.3-1 Histogram of the power (a) and GSNR (c) absolute error for different gradient boosting (GB), neural network 1 (NN1), and neural network 2 (NN2) models. Cumulative distribution function (CDF) of the power (b) and GSNR (d) absolute error for GB, NN1, and NN2 models, along with the corresponding KPIs and training and test times.

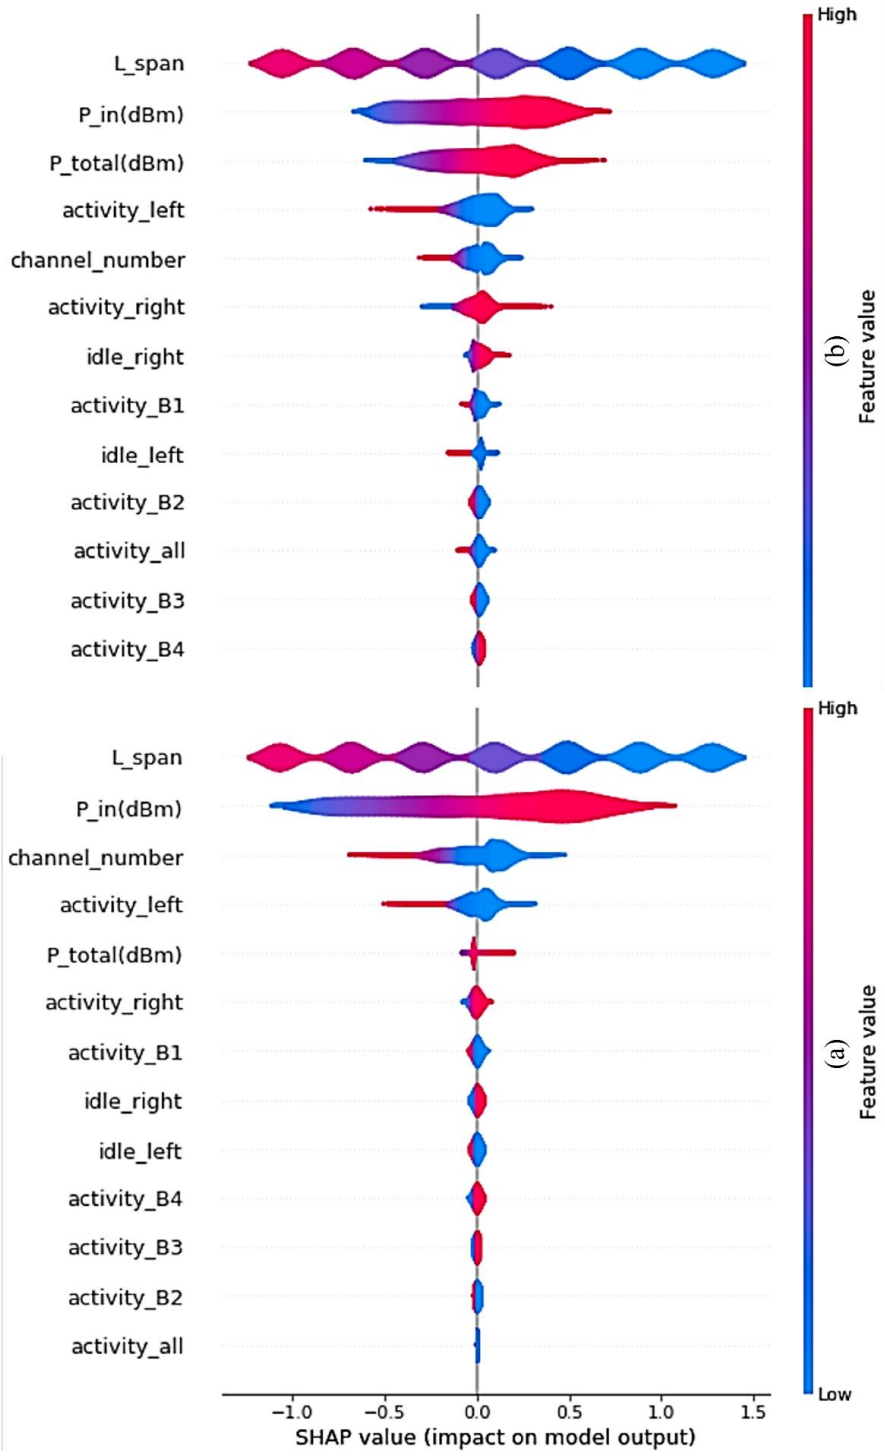

Figure 7.3-2 SHAP values for the power estimation (a) NN1 (b) NN2.

Dissemination Level	PUB (Public)

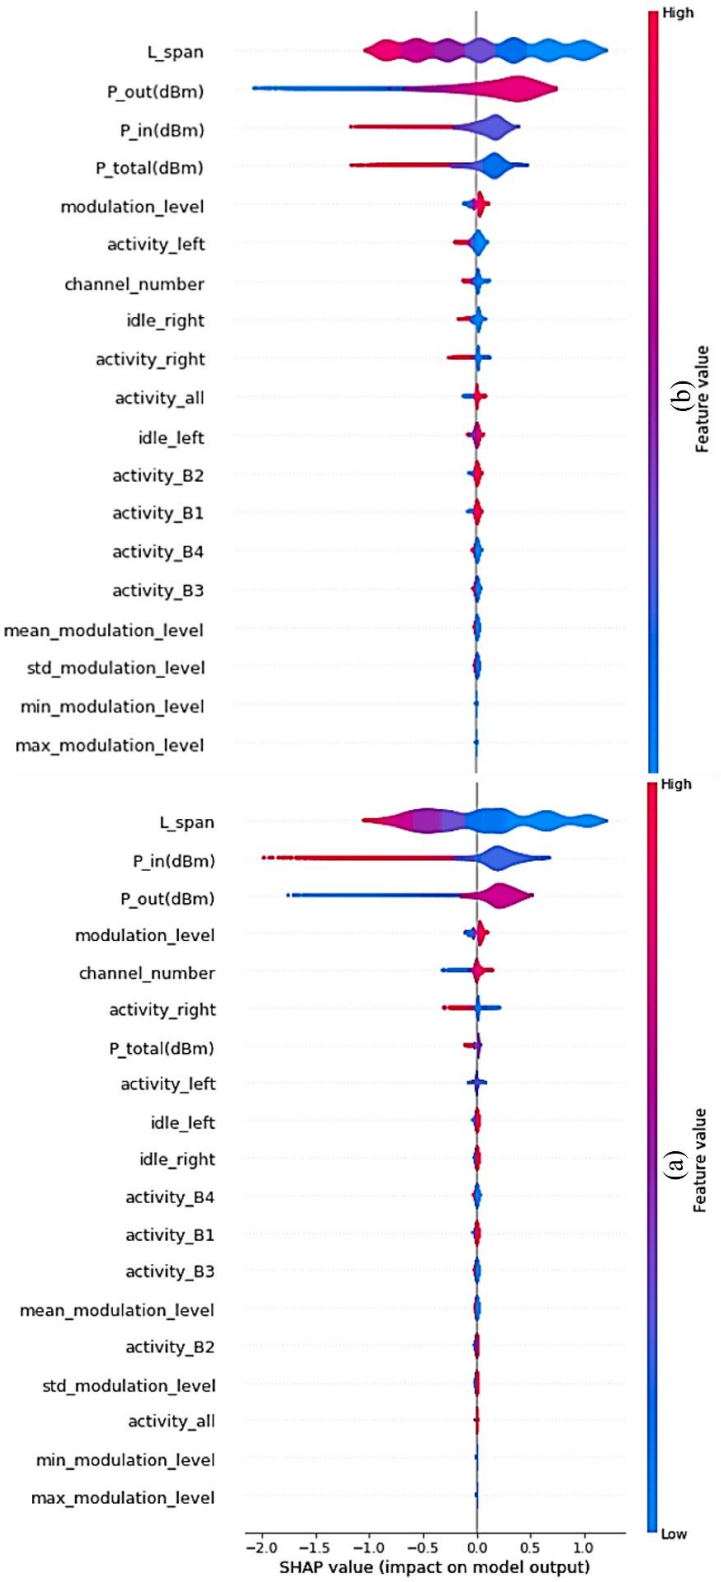

Figure 7.3-3 SHAP values for the GSNR estimation (a) NN1 (b) NN2.

7.4 Optical Span QoT Data

This dataset contains simulated data for estimating QoT in terms of Generalized Signal to Noise Ratio (GSNR) for optical spans in Elastic Optical Networks (EONs). It includes three datasets with different configurations of symbol rates and channel loading factors.

7.4.1 Description

This dataset is generated using an incoherent Gaussian Noise (GN) model to estimate the GSNR of optical spans in EONs. The datasets are created to study the impact of different configurations on QoT estimation accuracy using machine learning techniques. The dataset consists of several CSV files, each containing 10,000 samples with the following configurations:

- DS1: Fixed symbol rate (64 GBaud) with varying channel loading factor (25-100%)
- DS2: Random symbol rates (32, 64, 96 GBaud) with varying channel occupancy state
- DS3: Fixed symbol rate (64 GBaud) and fixed channel loading factor (25%)

Each dataset includes the following input parameters such as: Symbol Rate (GBaud), Span Length (km), Launch Power (dBm), Channel Under Test (CUT) Position, Channel Loading Factor (%). Moreover, additional derived features are: CUT's deviation from the center (left and right), Right and left Channel Loading Factor, Empty adjacent channels (left and right), Non-Linear Interference (NLI) Noise, Amplified Spontaneous Emission (ASE) Noise. The target variable is the GSNR of the span in dB.

The CSV files contain the following columns: "SymbolRate", "SpanLength", "LaunchPower", "CUTPosition", "CLF", "CUTDevLeft", "CUTDevRight", "CLFLeft", "CLFRight", "EmptyAdjLeft", "EmptyAdjRight", "NLINoise", "ASENoise", "GSNR".

This dataset can be used for developing and benchmarking machine learning models for QoT estimation in optical networks, particularly for span-level analysis in EONs.

7.5 REFERENCE NETWORK TOPOLOGIES

In Section 5 of D2.1, a set of reference network topologies for access, metro and backbone levels are presented. They are used in several designs and techno-economic studies made in SEASON project as reported in D2.2 and in this deliverable. The metro and backbone networks described in D2.1 are subject to confidentiality restrictions imposed by the operators who made them available, and this has made it impossible to allow access to these reference networks outside the consortium.

To address this issue, it was decided to create an additional reference network that is appropriately anonymized (i.e., removing any reference to real operator networks) but at the same time highly representative of an operator's network. The topology of this network has been made available for use within SEASON, but it has also been opened to the community as an open dataset exploitable for network studies.

The new defined reference network concerns the access-metro segment, one of the two segments into which the SEASON architecture is partitioned as is shown in Figure 1.1-1. In fact, the access-metro segment is the one where most of the SEASON innovations are focused. The new access-metro reference network is composed of two parts: the access, from each CO to the termination points (TPs) collected, i.e., cabinets or mobile stations, and the metro-regional network, which is the network between COs.

From the access part (from central office to the user termination points) as for confidentiality reasons it is not possible to provide scenario based on operators' actual deployments, a model based on statistical data and adaptable to a wide range of situations. This model, named "geotype" model and introduced in Subsection 5.1.3 of D2.1, defines how the access transmission points (FTTX cabinets and radio mobile stations of two types, macro cell and small cell) are distributed in a given area. For the metro-regional network, a 38-node meshed network including three level of Central offices and two levels of hierarchical segmentation of the topology: an upper-level meshed core and some lower-level aggregation networks (horseshoes, spurs or mix of both) connected to the core.

Excel files describing layouts of of geotypes and a file with metro regional network topology and traffic are available in a folder on ZENODO public repository:

- https://doi.org/10.5281/zenodo.16737162

Public folder includes also a document with figures of layout of geotypes, including an example of a fiber duct tree interconnecting access points with local central office, the metro-regional network topology, tables with summary of topological data and technical parameters, and an example of traffic flows generated for the combined access-metro network.

An example of application of the access geotype is given in [Sta25] where a parallel, decoupled fiber infrastructures versus a single passive P2MP connectivity scheme are examined. The analysis has been made under the framework of fiber scarcity and for residential users with

multi-Gb/s line rates, while small cells exploit mmWave technology. In the article it is show that fiber utilization becomes prohibitively under both alternatives (paralell Vs. convergent) when low-cost, gray interfaces are employed, and the deployments become viable in terms of fiber utilization only when colored interfaces and WDM are introduced.

Hereafter some details are provided on what is included in geotype files and in the metro-regional topology and traffic file.

7.5.1 Geotypes with Example of Fiber Duct Tree Topologies

The four types of geotypes defined in Subsection 5.1.3 of D2.1 are enriched in the excel files with one example of fiber duct topology for each geotype. There are many possibilities and strategies to connect the TPs together when topology is not assigned as in this case where only the TP positions are assigned, for instance following a Manhattan spanning tree [Chi16]. The duct typology for the four geotopes defined in the project was obtained with a heuristic algorithm that, starting from the most peripheral TPs, creates connections to nearby nodes, favoring in the tree building process those closest to the center of the area.

The proposed duct topology can be used in the design of systems and consequent calculation of required fibers in the ducts once network architecture (i.e., type of optical, packet and compute systems used) and traffic requirements have been assigned.

Figure 7.5-1 a) illustrates the Dense urban geotype topological points distribution on the grid which is made up of 36 squares with sides of 133 m for a total square area of 800 m on each side. On the left, the legend is specified for the characterization of the topological points, which can be cabinet (green circle), macro (orange triangle), or small (purple diamond) sites. Small and macro cells can be combined. Figure 7.5-1 (b) shows the duct topology of the Dense urban geotype obtained with the heuristic algorithm mentioned above. The green and purple numbers identify the TP codes (1 to 36 in green for cabinets or joined cabinet and combined Macro and Small cell sites, 37 to 55 in purple for Small cell only sites). For combined mobile sites, they are joined with the nearest cabinet to constitute a single TP in the duct topology, which is the case of 8, 11, 26 and 29 TPs. The black branches and their corresponding numbers of the same color are the main branches of the duct tree, onto which the secondary branches in blue and the third-level branches in red are grafted. The numbers in black, blue, and red encode the branches of the topology. Please note that the coding of branches is not sequential, for instance for main branches, in black, we have 1,2,3, 4, 5, 7, 11, for secondary branches, in blue, we have 6, 9, 10 etc. As can be seen the branches converge to the central topology point (TP 0, in blue) which is the place where the Central Office is located.

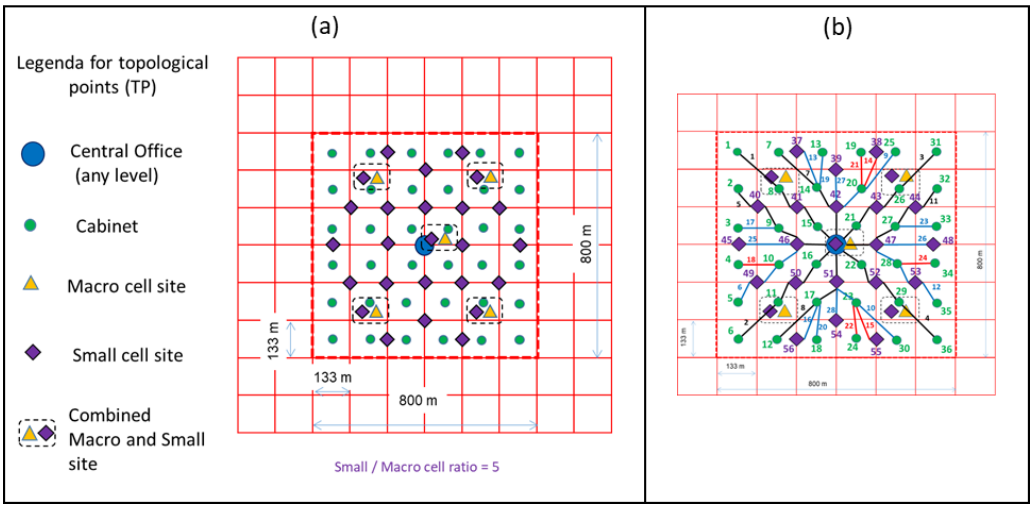

Figure 7.5-1 Dense urban geotype: (a) TPs distribution on the grid with legenda and (b) example of duct topology tree.

In the excel and power point files available on the Zenodo public repository the same information and figures as for the Dense urban geotype can be found for the other three geotypes.

7.5.2 Metro-Regional Topology

Figure 7.5-2 shows the topology of the reference Metro-regional network. It includes 38 nodes connected by 46 fiber links. According to the Telefónica terminology (following TIP MUST IP [Gon20]), the nodes involved can be categorized as HL5 (the lowest level of central office, aggregating packet flows from radio access points and mobile base stations; optical line terminals (OLTs), which collect the residential and business traffic from the fiber-to-the-x (FTTx) access networks; broadband fixed business direct traffic, etc.), HL4 (comprising IP routers that handle residential and business traffic classification, subscriber credentials authentication, validation of user's access policies, routing data to their destination, and aggregation of traffic from different locations of the metro network, as well as from OLTs and radio access points in densely populated areas) and HL3 (which collect IP/MPLS traffic from a metro/regional MAN and perform IP grooming toward the next hierarchical level, on the backbone network segment). According to the SEASON architecture terminology, the nodes correspond to Far Edge CO, and Edge COs (HL5≈Far-Edge, HL4 ≈ Edge-lower, HL3 ≈ Edge-upper). The topology is double tiered: a more peripheral structure made of horseshoes or spurs that interconnect the 26 HL5 nodes with two HL4 nodes (on some occasions, co-located with HL3 nodes), and central horseshoe structures of HL4 nodes (10 nodes) terminating at two HL3 nodes, located in a major city within the region. The length of the links in this representative metro-regional network topology ranges from 2 to 75 km, with an average of 21 km.

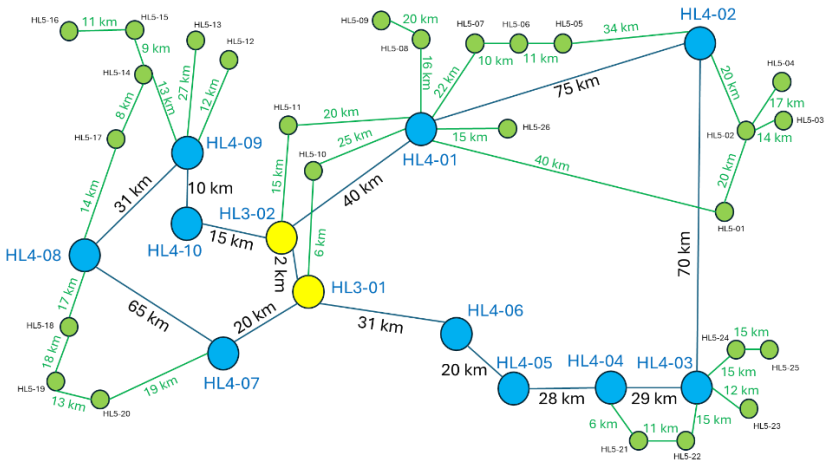

Figure 7.5-2 Metro-regional topology of SEASON Reference network.

7.5.3 Combining Access with Metro-Regional Topology

The regional metro network presented in Figure 7.5-2 can be combined with the geotypes defined in Section 7.5.1. The criteria used for assigning geotypes to COs are based on total population and population density at different points representative of regions organized around a medium-sized city (typically 150,000-300,000 inhabitants by European standards), which serves as the neuralgic center of the region, and the point where all the traffic is ultimately aggregated before entering the backbone network. These demographic data are considered to estimate the user areas collected by each node based on the amount of fixed and mobile access, which are eventually translated into geotypes. Depending on the aforementioned data, more than one geotype can be associated to a metro-regional node, in the case that a single geotype does not cover all the access requirements for that metro-regional node. The results of the association between metro-regional nodes and geotypes is given in Figure 7.5-3. As can be seen in this mapping, the most peripheral HL5 nodes are associated with rural or suburban geotypes, while the HL4 nodes have a prevalence of suburban typology. The HL3 nodes or those in urban areas are instead associated with more urban or dense-urban geotypes.

Dissemination Level	PUB (Public)

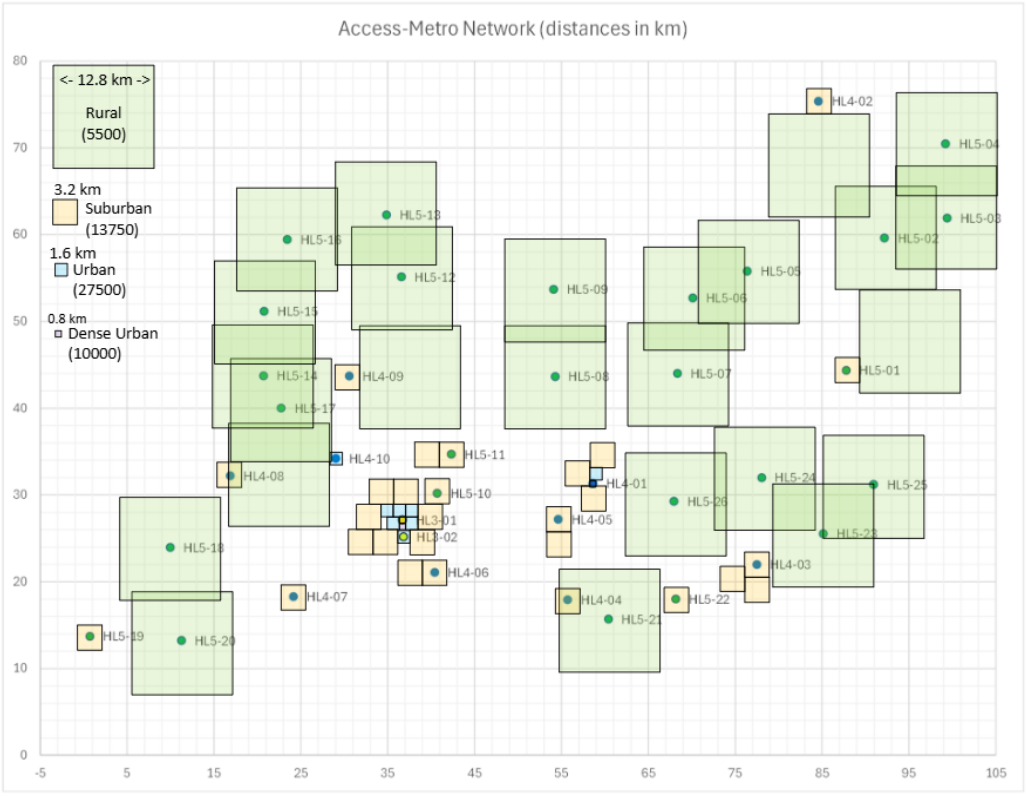

Figure 7.5-3 Combining geotypes with metro-regional topology.

7.5.4 Traffic Generation for SEASON Refence Network

The network traffic modeling process begins with the definition of geotype-based models and assumptions regarding traffic parameters connected to access points. From now on, traffic refers to the flow in terms of data rate (bits/s). Adaptation to the network layers (optical, packet) and related protocols (IP, Ethernet, OTN) will be based on specific architectural assumption. For fixed networks, traffic estimation considers the average household demand and associated peak factor. For mobile networks, traffic is derived from the installed radio layers at mobile sites, their configuration parameters: this allow to calculate the fronthaul constant bit rate (FH) and full load data rates for midhaul (MH) and backhaul (BH). In addition, the aggregated demand generated by mobile users at cell level is assigned as input parameter as it is upper limited by full load BH/MH data rate, but in fact is rather lower and depends on socio-economic factors of users attached to the cell (i.e., density, profile and behavior of mobile users). FH, BH, MH data rates are used to design the RAN part of the network while household average demand and mobile users average demand at cell level are used to design the metro-regional part of the network. A set of peak factors and statistical multiplexing parameters are introduced to take into account that the average values of traffic generated by the users, once aggregated or multiplexed, have different impact on network resources depending on where it is considered. The placement of core telecommunications functions—such as the virtual Distributed Unit (vDU), virtual Centralized Unit (vCU), User Plane Function (UPF), and Broadband Network

Gateway (BNG)—has to be taken into account in traffic flows generation. In fact, the assumed distribution of these functions influences how and where service data center capabilities are integrated across the network and determines the magnitude and direction of traffic flows.

Once assumptions on architecture are established and traffic flow generated, the resulting model, consisting of both the topology and traffic flows, serves as a flexible and technology-agnostic framework. This model enables the evaluation of different transport network solutions, supporting both technical viability assessments and techno-economic comparisons of alternative infrastructure and deployment strategies.

As an example on traffic demand generation for access-metro network dimensioning, an architecture based on centralized function placement at the HL3 level is assumed, meaning that the key Telco functions are deployed in a centralized architecture (i.e., this means that functions like UPF and BNG are placed at HL3 level). The two options considered to derive traffic flows in the metro regional networks are shown in Figure 7.5-4. In Case1 (upper part of Figure 7.5-4) traffic collected by nodes of level HL5 or HL4 is directly sent to HL3. Each flow is divided into two sub-flows (in network design should be routed on disjoint paths, if allowed). Resources for subflows (ideally one half of the total traffic of the flow) can include a margin for protection in case one of the two subflow is blocked by a fault. In Case 2 (bottom part of Figure 7.5-4) The flow generated by HL5 are sent to the reference couple of HL4 and aggregated here with subflows gathered by other HL4 nodes. Aggregated traffic of each HL4, which includes the traffic generated by HL4 itself, are sent to HL3 through two sub-flows to be routed along different paths and with the option of adding capacity for fault protection, if required.

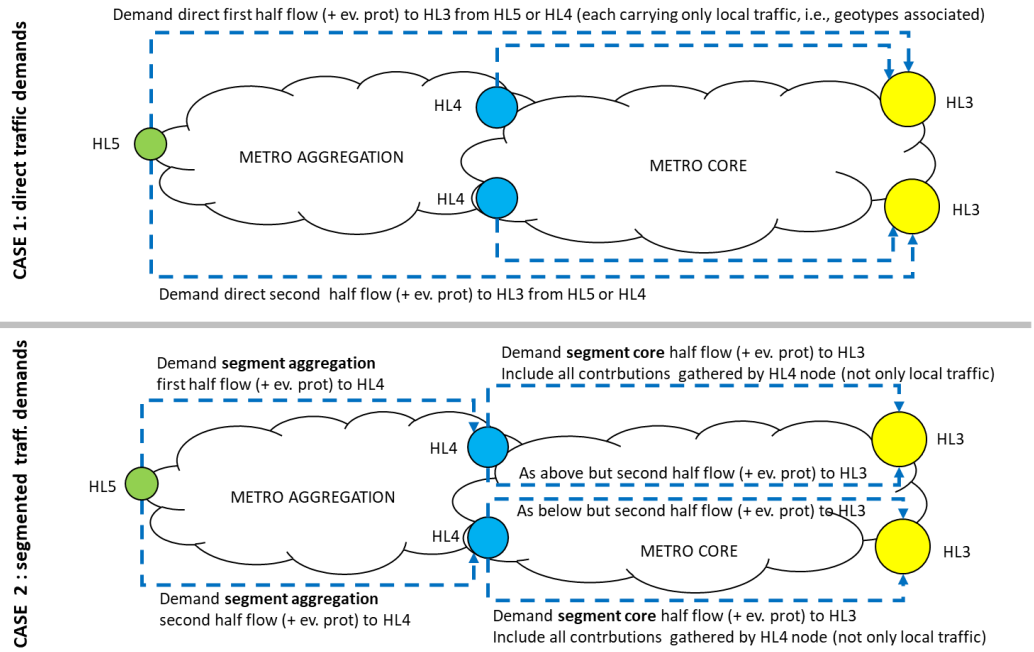

Figure 7.5-4 Traffic flows for centralized architecture in two different segmentation options: CASE 1 direct flows to Network functions placed in in HL3, CASE 2 – segmented flows with HL4 nodes used as intermediate aggregators Between HL5 and HL3.

The outcome of these modeling and calculations activities is captured in the supporting dataset available on Zenodo link provided above and titled *Geotype-Access&Metro-Regional_Traffic_Model_V1.xlsx*, which contains quantified traffic values derived under the specified assumptions. More specifically in the excel following sheets are included:

- "Main parameters": statistical traffic parameters (e.g., average traffic per household and per cell, peak and multiplexing factors) and system parameters (E.g, FH, MH/BH interface data rate for each cell type)
- "Rural", "Suburban", "Urban", "Dense Urban": geotytpe parameters (e.g., number of households, radio layers in a site) and aggregated traffic calculation for each geotype ("geotype" =)
- "Metro Regional": list of Metro regional nodes with hierarchical dependencies. Geotype associations to each metro-regional nodes (one or more than one geotype are associated to one metro regional node) and calculation of traffic generated on each node.
- "Geotype Mapping": list of Geotype with their assigned codes and their central point distances to reference metro-regional node. Visualization of such mapping on a scale map.
- "Demand list direct": traffic demand in the metro-regional network calculated according to Case 1 of Figure 7.5-4.
- "Demand list segmented": traffic demand in the metro-regional network calculated according to Case 2 of Figure 7.5-4.

Traffic data generated with the approach presented above can be used as a baseline for evaluating different transport scenarios, comparing possible solutions, and guiding strategic planning for efficient and scalable network deployment.

This structured approach ensures that network planning is based on realistic demand forecasts, functional architecture considerations, and provides a solid foundation for techno-economic analysis in the context of evolving transport requirements across fixed and mobile access domains.

7.5.5 Generating Realistic Optical Topologies for Techno-Economic Studies

In a collaborative effort with the Allegro project [ALLGRO], a methodology and tool for generating network topologies based on indicators derived from real networks, such as those of operators, was developed [San25]. The tool, called MoleNetwork and coded in Phyton, allows for the generation of a very large number of topologies and therefore for conducting statistical studies, overcoming the limitation of obtaining results from a single or a few topology samples.

The article provides a link to the open-source code and some generated instances, which are then made available to the community for network studies.

7.6 PACKET TRACE DATA

This dataset has been generated with a R script that simulates network flow traces based on given link capacity, load, and packet/flow characteristics. It generates a trace file that can be used for network analysis and research purposes.

7.6.1 Description

This dataset comes from a piece of code for generating a packet trace as it would arrive to a 100G port in a packet-optical router.

The R script simulates network flow traces based on given link capacity, load, and packet/flow characteristics. It generates a trace file that can be used for network analysis and research purposes. The script then outputs the trace data to a CSV file.

As input parameters, the simulator takes into account:

- *Link Capacity*: The capacity of the link in bits per second (bps); by default 100 Gbps.
- *Link Load*: The load on the link as a fraction of the total capacity, by default 0.4 (or 40% load).
- *Packet Size Distribution*: A vector specifying the sizes of packets in bytes. The script uses a trimodal distribution with packet sizes of 64, 596, and 1500 bytes. Each packet size is characterised by the weights, in this case 4/12, 2/12, and 6/12.
- *Flow Characteristics:* including the total number of flows (10,000 flows) and the Heavy-Hitter coefficient alpha parameter for determining the structure of the flow size distribution, which configures the Zipf distribution for the number and size of heavy-hitter flows in the packet trace.
- *Trace duration*: by default, the time window specified is 100 ms.

The script produces a CSV file named `Trace100G.csv`, which contains the simulated trace data. The columns in the CSV file are:

- ArrivalTime: The arrival time of the packet.
- FlowID: The flow ID to which the packet belongs.
- PacketSize: The size of the packet in bytes.

The following is an example of the first rows of the CSV:

"ArrivalTime","FlowID","PacketSize"

0.00000123,5,64

0.00000234,7,1500

0.00000345,3,596

...

7.6.2 Illustrative Figures

Figure 7.6-1 shows an example of the flow size distribution (Zipf-like) obtained using the code for the default parameters.

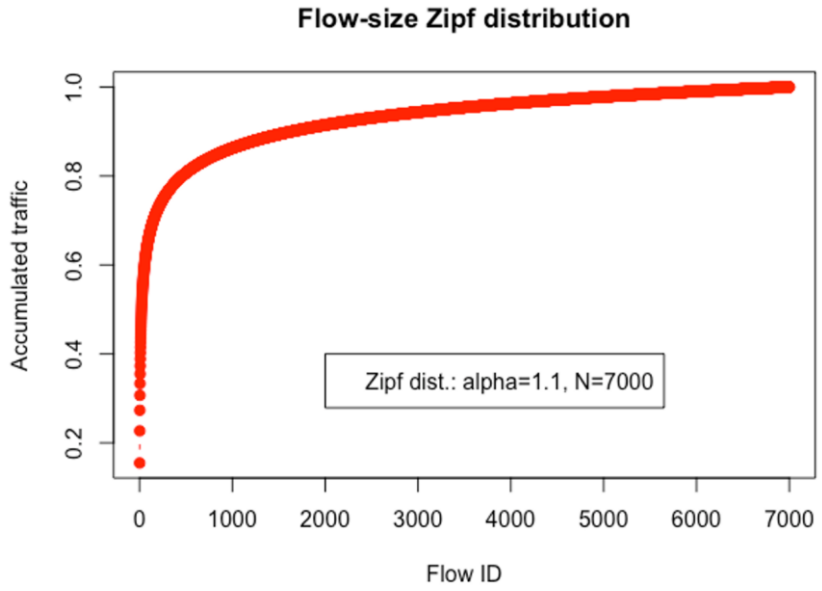

Figure 7.6-1 Accumulated flow-size distribution obtained after running FlowTraceSimulator for 7000 flows with heavy-hitter coefficient alpha=1 on a 100 Gb/s port.

7.7 CONTROL PLANE DATA

This dataset corresponds to the traces generated by an optical controller following a Linux Foundation TR-548 TAPI Streaming approach. In this case the protocol transport used is MQTT. A MQTT Broker (Mosquitto, Open Source) is responsible for collecting traces generated by the optical SDN Controller, and to forward them to subscribed clients depending on specific topics. In this case, the topic is "tapi/streaming".

7.7.1 Architecture and Topology

The SDN Controller is connected to the MQTT broker (localhost, 1883) at startup. A first sync dump is realized upon connection. Subscribers are connected to the same MQTT broker. A TAPI client requests and deletes media channel services using a TAPI North Bound Interface (NBI). Based on specific events on the network, streaming log records are generated. The considered topology corresponds to the Telefonica Network depicted in Figure 7.7-1. There are 65 ROADMs

(Nodes) and 270 links. Links are flexi-grid with C, L and S bands. Each node has 10 transceivers, so there is a total of 650 Service Interface Points. The full TAPI Context is provided in the file “context.json” in the available repository.

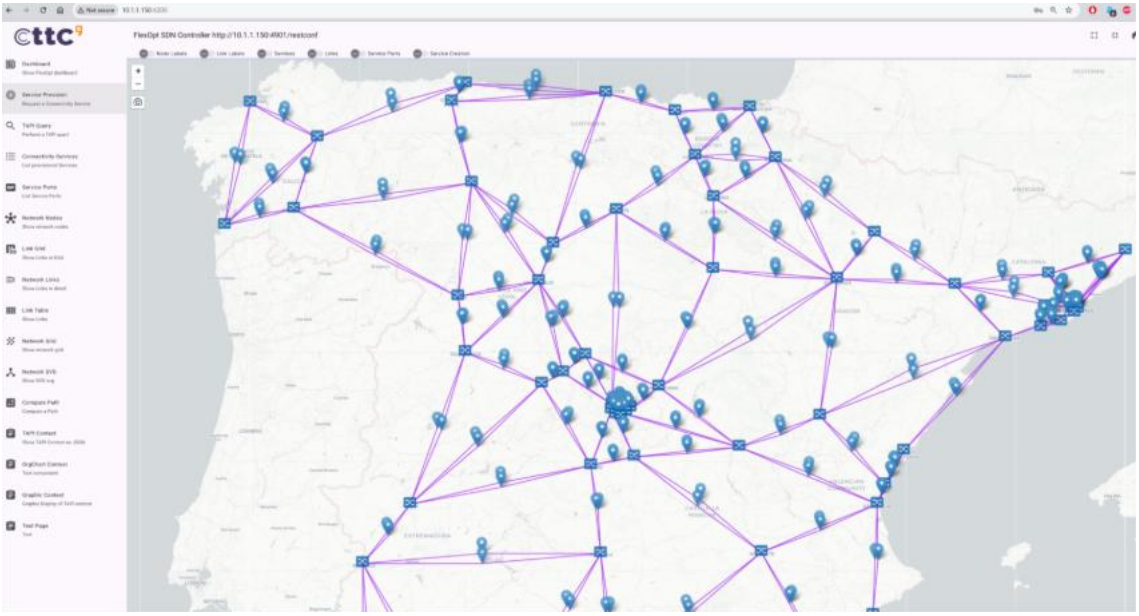

Figure 7.7-1 Reference topology for control plane data.

7.7.2 Requests and Notifications

Requests are generated randomly, with the following parameters.

- Number of connections: 1000 (N)
- Holding Time: 5 “time units”
- Interarrival Time: 1 “time units”

The file “services.log” contains the details of all requests, according to the following format:

<timestamp> LSP: # (<name) - <sip-uuid-src> <sip-uuid-dst> S: <arrival-time> R: departure

Illustrative example of requests is provided below:

```
2025-09-22 18:02:08,204 LSP: 0 (SERVICE-0) - 1e7ad3f6-1b37-5f98-9a31-de5392233a17 to
98c93ee0-e1e1-5d77-8544-5e85c8f4b8f0 S: 0.06901986161910273 s. R: 0.6332621871205555 s.


2025-09-22 18:02:08,204 LSP: 1 (SERVICE-1) - ac1c1264-0989-5197-9f9c-df248cd46e6d to
e8d90252-65cc-5572-96b2-a89de38d6ce3 S: 0.753480363939603 s. R: 1.3990739232324951 s.


2025-09-22 18:02:08,204 LSP: 2 (SERVICE-2) - c6ac0aef-c455-51a3-bba4-73499fa78754 to
12aef0d9-85fd-5c95-a4d1-b7f8b948de75 S: 1.5749482110569066 s. R: 9.016432333058367 s.
```

In addition, events are logged, sequenced as follows:

```
2025-09-22 18:02:08,275 events: 1999

2025-09-22 18:02:08,275 SETUP LPS: SERVICE-0 1e7ad3f6-1b37-5f98-9a31-de5392233a17 to
98c93ee0-e1e1-5d77-8544-5e85c8f4b8f0 (100 GHz) [5dfb444e-20ec-41ad-9918-6f1eafe21397]

2025-09-22 18:02:08,275 Starting new HTTP connection (1): localhost:4901

2025-09-22 18:02:08,289 http://localhost:4901 "POST /restconf/data/tapi-
common:context/tapi-connectivity:connectivity-context HTTP/1.1" 201 0

…

2025-09-22 18:02:13,009 events: 1990

2025-09-22 18:02:13,009 RELSE LSP: SERVICE-5 [54682c87-26cb-412a-8a69-2412ae241a50]

2025-09-22 18:02:13,010 Starting new HTTP connection (1): localhost:4901

2025-09-22 18:02:13,012 http://localhost:4901 "DELETE /restconf/data/tapi-
common:context/tapi-connectivity:connectivity-context/connectivity-service=54682c87-
26cb-412a-8a69-2412ae241a50 HTTP/1.1" 204 0
```

Finally, notifications are logged in the file “mqtt_messages.log”, with the format:

[<timestamp>] Topic: <topic>, Message: JSON message

Illustrative example of a notification is provided below (message is truncated):


```
[2025-09-22 18:02:01.310927104] Topic: tapi/streaming, Message:
{"metadata":{"measurement":"EventTelemetry","index":"cttc-index"},"data":{"tapi-
streaming:log-record":{"log-record-header":{"tapi-context":"02a8bff0-dd1e-55f1-b31f-
80a02491e44b","token":"0","log-append-time-stamp":"2025-09-22 16:02:01.310221693
UTC","entity-key":"f92aef68-2fef-54e9-b6fe-a4b684354474","record-
type":"RECORD_TYPE_CREATE_UPDATE"},"log-record-body":{"event-time-stamp":{"primary-time-
stamp":"2025-09-22 16:02:01.310256269 UTC"},"node":{"uuid":"f92aef68-2fef-54e9-b6fe-
a4b684354474","name":[{"value-name":"clli","value":"A Coruña-1"},{"value-
name":"cttc.node.type","value":"cttc.node.type.ROADM"},{"value-
name":"latitude","value":"43.352505"},{"value-name":"local-
name","value":"node_11_1.0.0.11"},{"value-name":"longitude","value":"-
8.433284"},{"value-name":"ls.pcepls.lsid","value":"97573921074472"}...(message
truncated)}
```


7.8 COST MODEL DATASET

The cost model values for commercially available devices, as outlined in Section 6.1, are provided in an Excel spreadsheet published on Zenodo:

https://zenodo.org/records/17304012.

A detailed explanation of the methodology used to derive these values, along with frameworks for estimating the cost and power consumption of devices not yet commercially available, can be found in Section 6.

7.9 Generative Adversarial Networks for Synthetic Traffic Generation

Generative Adversarial Networks (GANs) have emerged as a transformative technology in artificial intelligence, enabling machines to generate highly realistic synthetic data that closely mimics real-world distributions. They have become very popular AI tools for generating photos of persons that look real (but do not exist).

The core architecture of a GAN lies in their unique adversarial training approach, where two neural networks—the generator and the discriminator—compete in a game-theoretic setup. Starting from random noise, the generator creates synthetic samples, while the discriminator evaluates their authenticity, driving both networks to improve iteratively until the generated data becomes indistinguishable from real samples.

In what concerns this deliverable, one application of GANs is the generation of synthetic network traffic data. As networks grow in complexity, managing and optimizing them requires realistic traffic datasets for simulation, testing, and research. However, real network traffic data is often limited due to privacy, security, or availability constraints. In this sense, GANs can provide a methodology for generating such realistic traffic traces by learning the intricate statistical properties and dependencies present in real network traffic, enabling the creation of synthetic datasets that closely mirror real-world conditions.

7.9.1 Methodology

The methodology for generating synthetic network traffic matrices using GANs and CTGANs involved four key components, leveraging Python's Synthetic Data Vault (SDV) and custom implementations:

1. Data preparation: Since it is very difficult to find real traffic matrices, we have used the publicly available dataset from SNDLib website[2] which offers over 25 real network topologies with various traffic matrices. In particular, the **US Abilene network** topology (12 nodes, 15 links, 132 source-destination demands) includes 167 real traffic matrices recorded at 5-minute intervals over 6 months collected in 2004. The values have been rescaled to 2025 numbers by assuming a 45% Cummulative Annual Growth Rate (CAGR).

2. Preprocessing: With the help of the LabelEncoder and MinMaxScaler functions of Python's sklearn, categorical features have been encoded and normalized.

3a. Custom GAN implementation: The GAN includes a Generator Neural Network that takes a 10-dimensional noise vector as input and has two fully connected layers (128 → 64 nodes) with ReLU activation; the output is a source/destination demand values (numerical) using softplus activation to ensure positivity. The second Neural Network is a discriminator with a 3-

[2] SNDLib website, URL: https://sndlib.put.poznan.pl/home.action, last access June 2025.

dimensional data (source, target, demand) and two fully connected layers (64 nodes each) with Leaky ReLU activation; the output is a Binary classification algorithm (aimed at identifying real vs. synthetic) via sigmoid activation. For training, PyTorch's binary cross-entropy loss has been used along with 1,000 epochs on Adam optimizer.

3b. SDV-CTGAN Approach: This second architecture leverages the SDV library's CTGAN class for tabular data synthesis and modeled conditional dependencies between features (e.g., time-of-day patterns). Again, it has been trained for 1,000 epochs with batch size 128.

4. Evaluation: In order to assess whether or not the resulting synthetic traffic matrices look real, three evaluation metrics have been considered: Statistical Similarity, Link Load Analysis and Correlation Structure. In the first case, statistical similarity aims to check basic statistical properties of both real and synthetic traffic; Network Link Load evaluates difference in terms of link load between synthetic and real traffic matrices, and their impact in the network; finally, the correlation structure pattern identifies whether or not traffic dependencies between source-destination pairs are preserved.

7.9.2 Results

While GANs tend to overestimate traffic load (higher mean and median), CTGAN matches the real data closer in both mean and median but has problems regarding the standard deviation. Regarding Network Link Load, CTGANs generate traffic such that the relative impact on the network link load is below 3% difference with respect to real traffic.

Finally, regarding correlation structure, Figure 7.9-1 shows the correlation heatmap between every source and destination node for the real traffic matrices (left), and the synthetic generated with a GAN (center) and CTGAN (right). As shown, the GAN captures main positive correlations but misses finer negative patterns, while the CTGAN outperforms at diversity and slight negative correlations, closer to real traffic behavior.

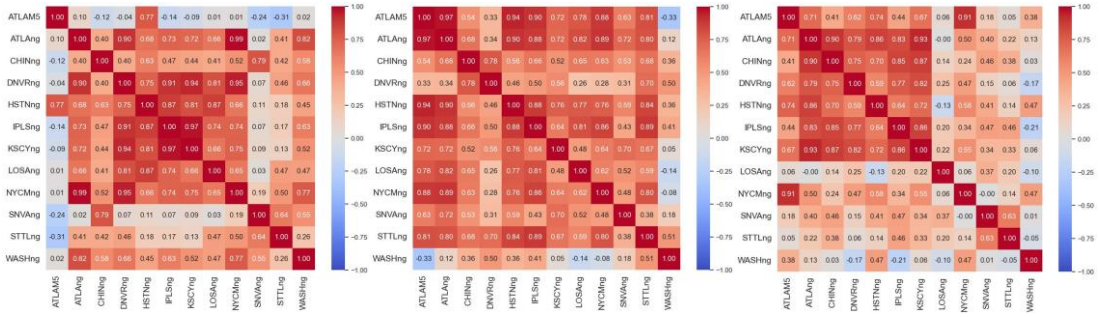

Figure 7.9-1 Correlation heatmap between source-destination pairs for real (left), GAN (center) and CT-GAN (right) traffic matrices.

In conclusion, both GANs and CTGANs can be used as AI-based techniques for generating synthetic traffic matrices in network scenarios, but always starting from real ones taken as baseline. In particular, statistical analysis reveals that CTGANs achieve superior fidelity, with

81.9% overall similarity to real data and a 57% reduction in link load error compared to traditional GANs. The synthetic data preserves key correlations between source-destination pairs while mitigating privacy risks.

8 CONCLUSIONS

The SEASON project set out to redefine sustainable, high-capacity optical network architectures for the cloud era, with objectives focused on developing modular, scalable designs, enabling advanced MB and space-division multiplexing technologies, and enhancing techno-economic viability and energy efficiency.

Deliverable D2.3 demonstrates that up to 120x backbone capacity growth is achievable through MB and SDM innovations, substantiated by prototype validations, simulation results, and comprehensive migration strategies. The implementation of flexible modular nodes and transceivers, strategic adoption of cost-reducing architectures, and deployment of AI-driven orchestration have enabled substantial CAPEX reductions (up to 50%) and energy savings (in the range of 30-50%), meeting all targeted KPIs.

Importantly, the publication of open-access datasets, including cost models and network topologies, empowers the wider research and industrial community with benchmarking and reproducibility tools for ongoing studies. As the project approaches final validation phases, the demonstrated results in D2.3 provide a robust foundation for the progression of self-managed, energy-optimized, and cost-effective optical networks.

Future steps will build upon these advances by further refining AI-enabled automation, MB technology integration, and empirical dataset enrichment to accelerate the deployment of sustainable digital infrastructure across Europe and beyond.

GLOSSARY

ACRONYM	DESCRIPTION
A/D	Add-Drop
ACO	Automatic Cell Optimizer
AI	Artificial Intelligence
ARM	Advanced RISC Machines
ASE	Amplified Spontaneous Emission
ASIC	Application-Specific Integrated Circuit
BBU	Baseband Unit
BER	Bit Error Rate
BH	Backhaul
BNG	Broadband Network Gateway
BSS	Business Support Systems
c.u.	Cost Unit
CAGR	Cumulative Annual Growth Rate
CAPEX	Capital Expenditure
CDC	Colorless-Directionless-Contentionless
CD	Colorless-Directionless
CFM	Closed Form Model
CO	Central Office
CPU	Central Processing Unit
CSR	Cell Site Router
CU	Centralized Unit
CUT	Channel Under Test
CWDM	Coarse WDM
DB	Database
DETNET	Deterministic Networking
DO	Day One
DPU	Data Processing Unit
DSB	Double-Sideband
DSC	Digital Subcarrier
DSCM	Digital Subcarrier Multiplexing
DSP	Digital Signal Processor
DT	Digital Twin
DU	Distributed Unit
DWDM	Dense WDM
E2E	End-to-end
EDFA	Erbium-Doped Fiber Amplifiers
EGGN	Enhanced Generalized Gaussian Noise
EON	Elastic Optical Networks
FEC	Forward Error Correction
FH	Fronthaul
FL	Fiber Leasing

FL	Filterless
FLP	Flat Launch Power
FOADM	Fixed OADM
FRP	Flat Receive Power
FTTX	Fiber-To-The-X
GAN	Generative Adversarial Network
GN	Gaussian Noise
GSNR	Generalized Signal-to-Noise Ratio
ICXT	Inter-Core Crosstalk
ILA	Inline Amplifier
IOT	Internet of things
IP	Internet Protocol
IPOWDM	IP-over-WDM
IQ	In-Phase and Quadrature
ISRS	Inter-Channel Stimulated Raman Scattering
KPI	Key Performance Indicator
LABI	Legacy-Agnostic Band Integration
LLM	Large Language Model
MAN	Metropolitan Area Network
MB	Multi-band
MBOSDM	MB over SDM
MBS	Macro BS
MBT	MB Transmission
MCF	Multi-Core Fiber
MH	Midhaul
ML	Machine Learning
MOR	Maximum Optical Reach
MSA	Multi-Source Agreement
NBI	North Bound Interface
NIC	Network Interface Card
NLI	Nonlinear Interference
NLS	Nonlinear Least Squares
NN	Neural Network
NRZ	Non-Return-To-Zero
O/E/O	Optical-Electro-Optical
OADM	Optical Add-Drop Multiplexer
OCh	Optical Channel
OCM	Optical Channel Monitor
OLC	Optical Line Controllers
OLS	Optical Line System
OLT	Optical Line Termination
OPEX	Operational Expenditures
O-RAN	Open-RAN
OSC	Optical Supervisory Channel
OSNR	Optical Signal-To-Noise Ratio
OSS	Operations Support Systems
P2MP	Point-To-Multipoint

P2P	Point-To-Point
PAYG	Pay-as-You-Grow
PMD	Polarization Mode Dispersion
PON	Passive Optical Network
QOS	Quality-of-Service
QOT	Quality Of Transmission
QSFP-DD	Quad SFP- Double Density
RAN	Radio Access Network
RIC	Radio Intelligent Controller
RISC	Reduced Instruction Set Computer
ROADM	Reconfigurable Optical Add-Drop Multiplexer
RRU	Remote RU
RU	Radio Unit
SB	Single-band
S-BVT	Sliceable Bandwidth/Bitrate-Variable Transceiver
SC	Subcarrier
SDM	Space Division Multiplexing
SDN	Software-Defined Networking
SDV	Synthetic Data Vault
SFP	Small Form-Factor Pluggable
SMF	Single-Mode Fiber
SNR	Signal-To-Noise Ratio
S-OXC	Spatial-Optical Cross Connects
SPM	Self-Phase Modulation
SRS	Stimulated Raman Scattering
SSB	Single-Sideband
SSMF	Standard SMF
SW	Software
TCO	Total Cost of Ownership
TEA	Techno-Economic Analysis
TP	Termination Point
TSN	Time-Sensitive Networking
UE	User Equipment
UPF	User Plane Function
URT	Unserved Residual Traffic
VAT	Value-Added Tax
VCU	Virtual Centralized Unit
VDU	Virtualized DU
VLAN	Virtual Local Area Network
WAN	Wide Area Network
WDM	Wavelength-Division Multiplexing
WSS	Wavelength Selective Switch
XT	Crosstalk

REFERENCES

[Adt22] Adtran, “Optical Engines Coherent 100ZR”, [Online] https://www.adtran.com/-/media/adtran/resources/data-sheets/pdfs/coherent-100zr.pdf?rev=-1&hash=FE0F36401A9A7428276D4F509C57B706, [Accessed on] June 2025.

[Aly05] Aly, M.; Nasr, M.; Seleem, H. Thulium doped fiber amplifier in the S/S/sup +/ band employing different concentration profiles. In Proceedings of the 2005 Fifth Workshop on Photonics and its Applications, 2005, pp. 39–49. https://doi.org/10.1109/PAIA.2005.205055.

[ALLGRO] Agile ultra-low energy secure networks (ALLEGRO) project website, https://www.allegro-he.eu/

[APPnet] Approved Networks website. https://www.approvednetworks.com/

[Arp24] Arpanaei, F. *et al*. (2024). Ultra-high-capacity band and space division multiplexing backbone EONs: Multi-core versus multi-fiber. *Journal of Optical Communications and Networking*, *16*(12), H66–H78. Journal of Optical Communications and Networking. https://doi.org/10.1364/JOCN.533086

[Arp24-2] F. Arpanaei *et al*., "Enabling seamless migration of optical metro-urban networks to the multi-band: unveiling a cutting-edge 6D planning tool for the 6G era," in *Journal of Optical Communications and Networking*, vol. 16, no. 4, pp. 463-480, April 2024, doi: 10.1364/JOCN.505490.

[B5GOPN] Beyond 5G – OPtical nEtwork coNtinuum (B5G-OPEN) project website. https://www.b5g-open.eu/ . Accessed: 2025-09-24.

[Bur17] Burk, C. Techno-economic analysis for new technology development. Advanced Materials 2017, 1, 266–269.

[Cas24] C. Castro et al., “OSNR-Based Hardware Optimization of a Filterless Point-to-Multipoint Network Using Digital Subcarrier Multiplexing,” *2024 International Conference on Optical Network Design and Modeling (ONDM),* Madrid, 2024.

[Cas24-2] C. Castro et al., “Power and Spectral Savings in Metro-Aggregation Networks Exploiting Coherent Point-to-Multipoint Transceivers,” *50th European Conference on Optical Communications (ECOC 2024),* Frankfurt, 2024.

[Cha22] Chai, Slyvester & Phang, Frederick & Yeo, Lip & Ngu, Lock & Shen, How. (2022). Future era of techno-economic analysis: Insights from review. Frontiers in Sustainability. 3. 924047. 10.3389/frsus.2022.924047.

[Chi16] Chiu, W. Y., Yen, G. G., & Juan, T. K. “Minimum manhattan distance approach to multiple criteria decision making in multiobjective optimization problems,” IEEE Transactions on Evolutionary Computation, 20(6), 972-985. (2016).

[Cur15] V. Curri et al., “Design strategies and merit of system parameters for uniform uncompensated links supporting Nyquist-WDM transmission,” J. Lightw. Technol., vol. 33, no. 18, pp. 3921–3932, 2015.

[DELL] Dell Technologies website. https://www.dell.com/

[Dev24] M. Devigili, M. Ruiz, N. Costa, C. Castro, A. Napoli, J. Pedro, and L.Velasco, “Applications of the OCATA time domain digital twin: from QoT estimation to failure management, “ IEEE/OPTICA Journal of Optical Communications and Networking (JOCN), vol. 16, pp. 221-232, 2024.

[Ess12] R.-J. Essiambre and R. W. Tkach, “Capacity trends and limits of optical communication networks,” Proc. IEEE 100, 1035–1055 (2012).

[Fas25] Faster Capital, “Cost Modeling Challenges: How to Overcome the Common Difficulties and Problems of Cost Modeling” Available: https://fastercapital.com/content/Cost-Modeling-Challenges--How-to-Overcome-the-Common-Difficulties-and-Problems-of-Cost-Modeling.html

[Fer20] Ferrari, A.; Napoli, A.; Fischer, J.K.; Costa, N.; D’Amico, A.; Pedro, J.; Forysiak, W.; Pincemin, E.; Lord, A.; Stavdas, A.; et al. Assessment on the Achievable Throughput of MB ITU-T G.652.D Fiber Transmission Systems. J. Lightwave Technol. 2020, 38, 4279–4291.

[FS] FS website. https://www.fs.com/

[Gon20] Ó. González de Dios, L. Á. Muñoz, M. Vázquez, et al., “MUST IP SDN controller NBI technical requirements,” TIP MUST OOPT Technical report (2020).

[Hay12] T. Hayashi, T. Taru, O. Shimakawa, et al., “Uncoupled multi-core fiber enhancing signal-to-noise ratio,” Opt. Express 20, B94–B103 (2012).

[Her20] Hernandez, J. A.; Quagliotti, M.; Serra, L.; Luque, L.; Lopez da Silva, R.; Rafel, A.; Gonzalez de Dios, O.; Lopez, V.; Eira, 604
A.; Casellas, R.; et al. Comprehensive model for technoeconomic studies of next-generation central offices for metro networks. 605
Journal of Optical Communications and Networking 2020, 12, 414–427.
https://doi.org/10.1364/JOCN.402167

[Hos24] Hosseini, Mohammad M., et al. "Optimized design of horseshoe-and-spur filterless networks leveraging point-to-multipoint coherent pluggable transceivers." *Journal of Optical Communications and Networking*, 2024.

[IDLST] Industry-driven elastic and adaptive lambda infrastructure for service and transport networks (IDEALIST) project website. 600
https://cordis.europa.eu/project/id/317999/reporting/ . Accessed: 2025-09-23.

[Jan23] Jana, R.K.; Srivastava, A.; Lord, A.; Mitra, A. Optical cable deployment versus fiber leasing: an operator’s perspective on CAPEX savings for capacity upgrade in an elastic optical core network. Journal of Optical Communications and Networking 2023, 609
15, C179–C191. https://doi.org/10.1364/JOCN.483200

[Kon25]] N. Koneva, A. Sánchez-Macián, J. A. Hernández, F. Arpanaei and O. González de Dios, "On finding empirical upper bound models for latency guarantees in packet-optical networks", 2025 International Conference on Optical Network Design and Modeling (ONDM), Pisa, Italy, May 2025.

[Mar24] Martínez, R. *et al*. (2024). “Autonomous Control Operations for Energy-Efficient Packet Optical Networks.” *2024 24th International Conference on Transparent Optical Networks (ICTON)*, 1–4. https://doi.org/10.1109/ICTON62926.2024.10648263

[Mar25] Marotta, A., Centofanti, C., Graziosi, F., Quagliotti, M., & Agus, M. “Coherent Point-to-Multipoint as a Viable Techno-Economic Solution for Mobile Network Fronthaul”, 2025 International Conference on Optical Network Design and Modeling (ONDM), Pisa, Italy, May 2025

[Moh25] Hosseini, M. M., Parisi, G., Pedro, J., Torres-Ferrera, P., & Napoli, A. Nonlinear-Aware Design of Filterless Horseshoe-and-Spur Optical Networks., ONDM 2025

[O-RDM] OpenROADM, “The open ROADM multi-source agreement (MSA). [online] available: http://www.openroadm.org,”.

[ORA21] O-RAN ALLIANCE Specification O-RAN.WG9.XTRP-REQ-v01.00, O-RAN Open Xhaul Transport Working Group 9, Xhaul Transport Requirements, Feb. 2021. [online] https://orandownloadsweb.azurewebsites.net/specifications

[Pat20] S. K. Patri, A. Autenrieth, J.-P. Elbers y C. M. Machuca, «Planning Optical Networks for Unexpected Traffic Growth,» de European Conference on Optical Communication (ECOC), Brussels, Belgium, 2020.

[Rap22] Rapp, L.; Eiselt, M. Optical Amplifiers for Multi–Band Optical Transmission Systems. Journal of Lightwave Technology 2022, 40, 1579–1589. https://doi.org/10.1109/JLT.2021.3120944.

[Rap24] Rapp, L. Challenges for Multiband Optical Amplification and Solutions Therefore. Journal of Lightwave Technology 2024, 42, 4202–4212. https://doi.org/10.1109/JLT.2024.3356309.

[Rui14] M. Ruiz and L. Velasco, “Performance evaluation of light-tree schemes in flexgrid optical networks,” IEEE Commun. Lett. 18,1731–1734 (2014).

[Rui20] M. Ruiz, F. Tabatabaeimehr, and L. Velasco, "Knowledge Management in Optical Networks: Architecture, Methods and Use Cases [Invited]," IEEE/OSA Journal of Optical Communications and Networking, vol. 12, pp. A70-A81, 2020

[Sah17] Akio Sahara, Hiroki Kawahara, Shuto Yamamoto, Shingo Kawai, Mitsunori Fukutoku, Takayuki Mizuno, Yutaka Miyamoto, Kenya Suzuki, and Keita Yamaguchi, "Proposal and experimental demonstration of SDM node enabling path assignment to arbitrary wavelengths, cores, and directions," Opt. Express 25, 4061-4075 (2017)

[Sam20] N. Sambo *et al.*, "Provisioning in MB Optical Networks," in *Journal of Lightwave Technology*, vol. 38, no. 9, pp. 2598-2605, 1 May, 2020.

[Sam22] N. Sambo *et al.*, "Network upgrade exploiting multi band: S- or E-band?," in *Journal of Optical Communications and Networking*, vol. 14, no. 9, pp. 749-756, September 2022.

[San25] A. Sánchez-Macián, N. Koneva, M. Quagliotti, J. M. Rivas-Moscoso, F. Arpanaei, J. A. Hernández, and J. P. Fernández-Palacios, "Generating realistic optical topologies for techno-economic studies using MoleNetwork," J. Opt. Commun. Netw. 17, 847-862 (2025)

[SeaD2.1] SEASON Project, “Definition of use cases, requirements and reference network architecture,” (2024). Deliverable D2.1.

[SeaD2.2] SEASON-Project, “Techno-economic Analysis: First Results,” (2025). Deliverable D.2.2.

[SeaD3.2] SEASON Project, “Optical systems enabling ultra-high-capacity access/metro networks,” (2025). Deliverable D3.2.

[SeaD3.3] SEASON Project, “The SEASON solution for data plane infrastructure,” (2025). Deliverable D3.3.

[SeaD4.3] SEASON Project, “Final SEASON control plane infrastructure,” (2025). Deliverable D4.3

[SeaD5.1] SEASON Project, “Report on the integration and functional validation of integrated SEASON 1.0 solution,” (2024). Deliverable D5.1

[Sha18] B. Shariati, A. Mastropaolo, N.-P. Diamantopoulos, J. M. Rivas-Moscoso, D. Klonidis y I. Tomkos, “Physical-Layer-Aware Performance Evaluation of SDM Networks Based on SMF Bundles, MCFs, and FMFs,” Journal of Optical Communications and Networking, vol. 10, nº 9, pp. 712-722, 2018

[Sha24] H. Shakespear-Miles, Q. Lin, S. Barzegar, M. Ruiz, X. Chen, and L. Velasco, “Centralized and distributed approaches to control optical point-to-multipoint systems near-real-time,” IEEE/OPTICA Journal of Optical Communications and Networking, 16, 565-575 (2024).

[Sou24] A. Souza, N. Costa, J. Pedro, and J. Pires, “Comparative assessment of S+C+L-band and E+C+L-band systems with hybrid amplification,” in 2024 Optical Fiber Communications Conference and Exhibition (OFC), (2024), pp. 1–3.

[Sou24-2] A. Souza et al., "Node Architectures for High-Capacity MB Over Space Division Multiplexed (MBoSDM) Optical Networks," 2024 24th International Conference on Transparent Optical Networks (ICTON), Bari, Italy, 2024, pp. 1-7, doi: 10.1109/ICTON62926.2024.10647929

[Sou25] A. Souza, A. Napoli and J. Pedro, "Realistic Capacity Scaling of Single-Band and MB Optical Transport Systems," 2025 25th Anniversary International Conference on Transparent Optical Networks (ICTON), Barcelona, Spain, 2025, pp. 1-5, doi: 10.1109/ICTON67126.2025.11125240.

[Sou25-2] A. Souza, N. Costa, and J. Pedro, “Performance and transient-resiliency of S- and E-band upgrades of a C+L-band system,” in Optical Fiber Communication Conference (OFC) 2025, p. M4G.6, Optica Publishing Group, 2025.

[Sta25] A. Stavdas, E. Kosmatos, M. Quagliotti, M. Agus, A. Rafel, C. Matrakidis, and I. Cooper, "Technology convergence is reshaping the 6G access network architecture, but are our infrastructures ready to cope? [Invited]," J. Opt. Commun. Netw. 17, E94-E108 (2025)

[STRGST] Scalable, tunable and resilient optical networks guaranteeing extremely-high speed transport (STRONGEST) project website, https://cordis.europa.eu/project/id/247674/reporting/ . Accessed: 2025-09-23.

[SRIA] Strategic Research and Innovation Agenda 2021-27, European Technology Platform NetWorld2020 "Smart Networks in the context of NGI” https://bscw.5g-ppp.eu/pub/bscw.cgi/d367342/Networld2020%20SRIA%202020%20Final%20Version%202.2%20.pdf.

[Vel21] L. Velasco, S. Barzegar, D. Sequeira, A. Ferrari, N. Costa, V. Curri, J. Pedro, A. Napoli, and M. Ruiz, "Autonomous and Energy Efficient Lightpath Operation based on Digital Subcarrier Multiplexing," IEEE Journal on Selected Areas in Communications (JSAC), vol. 39, pp. 2864-2877, 2021

[Wan23] S. Wang, M. Ruiz, and L. Velasco, "Optical network traffic analysis under B5G/6G RAN operation," in Proc. IEEE International Conference on Transparent Optical Networks (ICTON), 2023.

[Wan24] S. Wang, M. Ruiz, and L. Velasco, "Context-based e2e Autonomous Operation in B5G Networks," MDPI Sensors, vol. 24, pp. 1-25, 2024.

[Wel21] D. Welch, A. Napoli, J. Bäck, et al., “Point-to-multipoint optical networks using coherent digital subcarriers,” J. Lightwave Technol. 39, 5232–5247 (2021).

[Wel23] Welch, D.; Napoli, A.; Bäck, J.; Buggaveeti, S.; Castro, C.; Chase, A.; Chen, X.; Dominic, V.; Duthel, T.; Eriksson, T.A.; et al. Digital Subcarrier Multiplexing: Enabling Software-Configurable Optical Networks. Journal of Lightwave Technology 2023, 41, 1175–1191. https://doi.org/10.1109/JLT.2022.3211466.

[XR_forum] Open XR Optics Forum website. https://www.openxropticsforum.org/

[Zhu21] Y. Zhu and S. Wang, “Joint Traffic Prediction and Base Station Sleeping for Energy Saving in Cellular Networks,” In Proceedings of the IEEE International Conference on Communications 2021, Virtual, 14–23 June 2021

APPENDIX - RELEVANT PUBLICATIONS

This appendix includes the publications linked to the work done in WP2 of project SEASON. Some of the content of this deliverable is based on findings reported in those publications, which explicitly cite the project SEASON as a partial or total source of funding. Other publications, both conference papers and journal articles, with work partly or entirely developed within SEASON WP2, have been submitted but are awaiting acceptance by conference committees or publishers and are not included in the list below. Publications accepted before the end of the project (31 December 2025) will be reported in the deliverable D6.3 of WP6 (Final project report on exploitation, standardization, communication, and dissemination activities), planned for the project M36. The final project review (expected in early 2026) will highlight all articles or papers published or accepted up to that date.

Table A-1 Relevant publications related to SEASON WP2 which include a citation to SEASON as funding project.

A. Souza, A. Napoli and J. Pedro, "Realistic Capacity Scaling of Single-Band and MB Optical Transport Systems," 2025 25th Anniversary International Conference on Transparent Optical Networks (ICTON), Barcelona, Spain, 2025.
Hosseini, M. M., Parisi, G., Pedro, J., Torres-Ferrera, P., & Napoli, A. Nonlinear-Aware Design of Filterless Horseshoe-and-Spur Optical Networks., ONDM 2025.
N. Koneva, A. Sánchez-Macián, J. A. Hernández, F. Arpanaei and O. González de Dios, "On finding empirical upper bound models for latency guarantees in packet-optical networks", 2025 International Conference on Optical Network Design and Modeling (ONDM), Pisa, Italy, May 2025.
C. Castro et al., "OSNR-Based Hardware Optimization of a Filterless Point-to-Multipoint Network Using Digital Subcarrier Multiplexing," *2024 International Conference on Optical Network Design and Modeling (ONDM),* Madrid, 2024.
C. Castro et al., "Power and Spectral Savings in Metro-Aggregation Networks Exploiting Coherent Point-to-Multipoint Transceivers," *50th European Conference on Optical Communications (ECOC 2024),* Frankfurt, 2024.
Marotta, A., Centofanti, C., Graziosi, F., Quagliotti, M., & Agus, M. "Coherent Point-to-Multipoint as a Viable Techno-Economic Solution for Mobile Network Fronthaul", 2025 International Conference on Optical Network Design and Modeling (ONDM), Pisa, Italy, May 2025
S. Wang, M. Ruiz, and L. Velasco, "Optical network traffic analysis under B5G/6G RAN operation," in Proc. IEEE International Conference on Transparent Optical Networks (ICTON), 2023.
S. Wang, M. Ruiz, and L. Velasco, "Context-based e2e Autonomous Operation in B5G Networks," MDPI Sensors, vol. 24, pp. 1-25, 2024.
H. Shakespear-Miles, Q. Lin, S. Barzegar, M. Ruiz, X. Chen, and L. Velasco, "Centralized and distributed approaches to control optical point-to-multipoint systems near-real-time," IEEE/OPTICA Journal of Optical Communications and Networking, 16, 565-575 (2024).
M. Devigili, M. Ruiz, N. Costa, C. Castro, A. Napoli, J. Pedro, and L.Velasco, "Applications of the OCATA time domain digital twin: from QoT estimation to failure management, " IEEE/OPTICA Journal of Optical Communications and Networking (JOCN), vol. 16, pp. 221-232, 2024.

A. Sánchez-Macián, N. Koneva, M. Quagliotti, J. M. Rivas-Moscoso, F. Arpanaei, J. A. Hernández, and J. P. Fernández-Palacios, "Generating realistic optical topologies for techno-economic studies using MoleNetwork," J. Opt. Commun. Netw. 17, 847-862 (2025)

A. Stavdas, E. Kosmatos, M. Quagliotti, M. Agus, A. Rafel, C. Matrakidis, and I. Cooper, "Technology convergence is reshaping the 6G access network architecture, but are our infrastructures ready to cope? [Invited]," J. Opt. Commun. Netw. 17, E94-E108 (2025)

M. Shammakhi, F. Arpanaei, H. Beyranvand and H. Sheikhzadeh, "Legacy-Agnostic Band Integration Approach: A Cost-Effective and Seamless C+l-Band Optical Network Migration," in Journal of Lightwave Technology, vol. 43, no. 19, pp. 9050-9065, 1 Oct.1, 2025