



Network Design with Overlapping Spectrum Allocation in DSCM-Enabled Optical Transport

Downloaded from: <https://research.chalmers.se>, 2026-07-12 18:11 UTC

Citation for the original published paper (version of record):

Pendem, S., Natalino Da Silva, C., Monti, P. et al (2026). Network Design with Overlapping Spectrum Allocation in DSCM-Enabled Optical Transport. Optical Fiber Communication Conference (OFC) 2026. <http://dx.doi.org/10.1364/OFC.2026.W4A.3>

N.B. When citing this work, cite the original published paper.

Network Design with Overlapping Spectrum Allocation in DSCM-Enabled Optical Transport

Sai Vikranth Pendem ^{1,*}, Carlos Natalino ¹, Giuseppe Parisi ²,
Pablo Torres-Ferrera ², Antonio Napoli ² and Paolo Monti ¹

¹Chalmers University of Technology, Gothenburg, Sweden; ²Nokia, Munich, Germany

*vikranth@chalmers.se

Abstract: This paper investigates a design strategy for bi-directional transmission with overlapping spectrum allocation under back-scattering noise penalties. Results show up to 18% fewer transceivers and 116% higher carried traffic compared to non-overlapping allocations.

1. Introduction

The next generation of services, including both 6G use cases and data-hungry Generative Artificial Intelligence applications, will further stress the capacity of optical transport networks [1]. As deploying new fiber remains costly, operators seek solutions to scale capacity while minimizing investments, prioritizing reuse of existing infrastructure and minimal new equipment.

A promising approach to improve resource efficiency is the use of Digital Subcarrier Multiplexing (DSCM) transceivers, which digitally synthesize a set of independent subcarriers (SCs) within a shared frequency resource, instead of a single monolithic channel as with conventional coherent transceivers [2]. Each SC can be independently modulated, assigned, and later reallocated, enabling both point-to-point and point-to-multipoint operation, where different SC groups can serve different services or destinations. This allows capacity to be dynamically matched to instantaneous demand, improving spectral utilization and reducing cost and energy consumption [3].

An additional advantage of DSCM transceivers, beyond fine-granular spectrum allocation, is their ability to support coherent bi-directional (bi-di) transmission using a single laser over the same fiber. In this configuration, counter-propagating SCs can operate at identical central frequencies, enabling overlapping spectrum allocation in opposite directions. This increases the traffic per deployed fiber without having to install parallel ones or additional transceivers. While [4] considered DSCM-based bi-di links with independent spectrum in each direction, it did not explore spectral reuse within the same fiber. However, overlapping operation introduces penalties, mainly due to Rayleigh back-scattering, which degrade SCs sharing the same spectral region [5]. Assessing these penalties and trade-offs is essential before bi-di transmission over shared spectrum resources can be widely adopted.

To this end, this paper for the first time formulates and solves a static Routing, Modulation and Spectrum Assignment (RMSA) problem with bi-di transmission and overlapping spectrum allocation, explicitly considering Rayleigh-induced impairments. Results across metro-scale and 6G ultra-dense topologies show up to 18% fewer transceivers and up to 116% higher carried traffic compared to non-overlapping schemes, demonstrating that coherent bi-di transmission with DSCM transceivers offers a scalable and cost-efficient path toward high-capacity, 6G-oriented optical transport systems.

2. System architecture and proposed Overlap RMSA algorithm

The study models the optical transport as a directed graph $G(V, E)$, where V is the set of nodes and E the set of directed edges. Each physical link has up to F fibers, each modeled as a pair of directed edges representing counter-propagating directions. Nodes are optically transparent and support SC-granular switching and filtering. Each node hosts one or more DSCM transceivers, each supporting a fixed number of SCs that can be independently modulated with any format $m \in M$, each with an associated reach. Overlapping spectrum allocation is assumed, where a downlink SC may share spectral slots with uplink traffic on the same fiber, and vice versa. To account for the performance degradation introduced by overlapping spectrum allocation, the optical-reach values used for the overlapping case are adjusted to reflect the impact of Rayleigh back-scattering and other impairments from counter-propagating transmissions within the same spectral region [5]. Traffic is defined by a set of static demands D , where each demand $d_i = \langle s_i, t_i, b_i \rangle$ specifies source, destination, and bit rate.

The problem under consideration is a static RMSA problem. The inputs are: the network graph $G(V, E)$, the traffic matrix D , and the set of modulation formats M with their respective reach. The objective is to provision the entire traffic matrix D while minimizing the number of required transceivers. For each demand d_i , a path between s_i and t_i must be found with enough SCs to support b_i , a number that is determined by the chosen m .

A demand may be served using non-contiguous SCs aggregated across multiple transceivers. SCs may either be assigned from already-deployed DSCM transceivers or by activating new transceivers at the s_i and t_i . In the former




Fig. 1. Overlap RMSA algorithm.

case, the SCs must belong to transceivers using the same central frequencies and must be available on all links of the selected path. In the latter case, all links of the selected path must have sufficient free spectral resources to host the newly allocated SCs. If no path satisfying both spectrum and reach constraints can be found for at least one demand, the static RMSA problem is deemed infeasible.

The proposed Overlap RMSA algorithm processes demands sequentially (Fig. 1), sorted in decreasing order of hop count. For demand d_i up to K candidate paths (computed using Yen's algorithm) are examined. Starting from the shortest one ($k = 0$), the algorithm evaluates path k and selects the highest-order modulation format m whose maximum supported reach is at least the path length. Based on that choice of m , the number of SCs required to serve b_i is then computed. The SC assignment has two stages. *Existing-SC Allocation* attempts to assign SCs from those available on transceivers deployed at the s_i and t_i . Two conditions must be met: (i) SCs must belong to transceivers operating at the same central frequency, and (ii) the associated spectral slots must be available on every link along path k . If all required SCs can be assigned, the network state is updated, and the algorithm moves to the next demand, until all of them are provisioned. Otherwise, the algorithm triggers *New-Transceiver Allocation* during which new transceiver pairs are activated at the endpoints of path k to allocate the remaining SCs, provided that enough continuous free spectral slots exist on all its links. If neither stage can provision the demand on path k , the next candidate path is evaluated. If no feasible allocation is found after examining all K paths, the problem instance is declared infeasible.

3. Performance evaluation

Overlap RMSA is evaluated using a Python-based simulation platform built on NetworkX. To emulate current *metro-scale* and future 6G *ultra-dense* networks, we consider two topologies, each with $|V| = 50$ nodes and 88 links ($|E| = 176$), representing a single-fiber infrastructure ($F = 1$) with an average nodal degree of 4. Link lengths range from 2–25 km and 1–13 km for the metro and ultra-dense topologies, respectively, following a log-normal distribution. Two traffic scenarios are investigated: *mixed-5G*, representing a typical metro-aggregation data-rate distribution [3], and *mixed-6G*, obtained by scaling the mixed-5G rates by 10× to reflect 6G uplink/downlink demands [6]. Each link offers 4 THz of spectrum within the C-band (1530–1565 nm), corresponding to 1066 SC slots assuming 3.75 GHz per SC, independent of m . Guard bands are neglected. Each SC operates with one of the modulation formats in $M = \{8, 16, 32, 64\}$ -QAM, providing per-SC data rates of $\{18.75, 25, 31.25, 37.5\}$ Gbps, respectively. DSCM transceivers support 16 independently configurable SCs.

To benchmark Overlap RMSA, we define a Non-overlap RMSA algorithm with an identical design logic, except that each spectral resource can be allocated to only one transmission direction. The optical reach values for Non-overlap RMSA are taken from [7], while the corresponding overlap reaches are derived from the same baseline but with modulation-dependent reach-reductions reflecting counter-propagation impairments. Two reach-reduction cases are considered: *conservative* and *optimistic*. The *conservative* assumes reaches of $\{100, 30, 15, 8\}$ km for each index $m \in M$ based on experimental results in [5], where 16-QAM SCs were transmitted over 32 km. The *optimistic* case assumes $\{250, 125, 50, 25\}$ km, representing forward-looking improvements in transceiver design and digital signal processing. For each source–destination pair, up to $K = 5$ candidate paths are computed.

The performance of Overlap RMSA is evaluated in terms of average transceiver count, and Maximum Feasible Traffic (MFT), which quantifies the largest traffic volume for which a feasible network design exists. Both metrics are analyzed as a function of the total injected symmetric traffic, each level evaluated over 30 randomly generated demand matrices, enough to guarantee a confidence interval of at most $\pm 5.2\%$ with a 95% confidence level.

Fig. 2 summarizes the results for the metro-scale topology. The non-overlapping assignment consistently requires more transceivers than the overlapping schemes, especially at higher traffic loads (Figs. 2(a) and 2(d)). The trend is more pronounced as the injected traffic increases, reflecting the benefits of spectral reuse between counter-propagating directions. For instance, in the mixed-5G scenario, overlapping allocation reduces the transceiver count by approximately 15–18%, depending on the assumed reach case, while for mixed-6G, the relative savings are smaller. This is because higher-rate demands in the mixed-6G traffic often occupy most or all of a transceiver's SCs, thereby limiting opportunities for reuse as seen in Fig. 2(e), where the values are computed for 190 Tbps of injected traffic. Conversely, mixed-5G demands use a similar average number of subcarriers across all three cases




Fig. 2. Metro-scale topology with mixed-5G (a)–(c) and mixed-6G (d)–(f) traffic.

(Fig. 2(b) computed at 60 Tbps), making spectral reuse more effective. As a result, the achievable network capacity (i.e., MFT) increases significantly under overlapping operation, with gains of roughly 80% in mixed-5G (Fig. 2(c)) and around 10% in mixed-6G (Fig. 2(f)) for the conservative case, and with even higher improvements under the optimistic case. In mixed-5G, the design becomes infeasible at lower injected traffic (i.e., after 60 Tbps) due to higher fragmentation caused by numerous low-rate demands. In contrast, the mixed-6G case tolerates higher injection levels before reaching saturation. Fig. 3 compares the results for the ultra-dense topology. The relative gains




Fig. 3. Reduction in the number of transceivers (a) and Maximum Feasible Traffic (MFT) gain (b) for the ultra-dense topology in mixed-5G and mixed-6G traffic.

from overlapping allocation are more substantial than in the metro-scale network, mainly because shorter links enable higher-order modulations and more efficient spectral reuse. The overlapping scheme achieves transceiver savings up to 18% (Fig. 3(a)) and can almost double the value of the feasible traffic injection (Fig. 3(b)). These improvements highlight that the benefit of overlapping allocation increases with network density, where shorter paths mitigate the penalties from counter-propagating impairments.

4. Conclusion

This work proposes a static RMSA design for bi-di transmission with overlapping spectrum in DSCM-based transport networks. Results confirm the intuition that overlapping spectrum allocation improves spectral reuse, reduces the transceiver count, and increases the feasible traffic, especially in dense, short-reach networks. The approach offers a scalable and cost-efficient path toward higher fiber utilization in future 6G-oriented optical transport systems.

Acknowledgments: This work has been supported by the European Union’s Horizon 2022 – MSCA project EMPOWER-6G (101120332) and by the SNS project PROTEUS-6G (Grant agreement ID:101139134). The topologies used in this study are available at [Github](#).

References

1. C. Natalino, *et al.*, in *ICMLCN*, (2025), pp. 1–6. DOI: [10.1109/ICMLCN64995.2025.11140431](#).
2. A. Napoli, *et al.*, in *Proc. ECOC*, (2025). Invited paper, Tu01.06.1.
3. C. Castro, *et al.*, *J. Opt. Commun. Netw.* **17**, 526–542 (2025). DOI: [10.1364/JOCN.559237](#).
4. P. Pavon-Marino, *et al.*, *J. Opt. Commun. Netw.* **15**, C29–C40 (2023). DOI: [10.1364/JOCN.482457](#).
5. P. Torres-Ferrera, *et al.*, in *Proc. OFC*, (2024), pp. Tu3E–5. DOI: [10.1364/OFC.2024.Tu3E.5](#).
6. A. Larrañaga, *et al.*, *IEEE Commun. Mag.* **63**, 86–93 (2025). DOI: [10.1109/MCOM.002.2400419](#).
7. F. David, *et al.*, *IEEE Netw. Lett.* **6**, 55–59 (2023). DOI: [10.1109/LNET.2023.3337041](#).