CHALMERS
UNIVERSITY OF TECHNOLOGY

Analysis of Spectrum Defragmentation in QoT-Constrained Elastic Optical Networks

Citation for the original published paper (version of record):
Magalhaes, T., Natalino Da Silva, C., Lobato, F. et al (2025). Analysis of Spectrum Defragmentation in QoT-Constrained Elastic Optical Networks. 2025 SBMO/IEEE MTT-S International Microwave and Optoelectronics Conference (IMOC). http://dx.doi.org/10.1109/IMOC65414.2025.11365801

Analysis of Spectrum Defragmentation in QoT-Constrained Elastic Optical Networks

Talles S. Magalhães*, Carlos Natalino†, Fabricio R. L. Lobato*, João C. W. A. Costa*, and Paolo Monti†
*Department of Electrical Engineering, Federal University of Pará, Belém, Brazil
{talles.magalhaes, frl, jweyl}@ufpa.br

† Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden
{carlos.natalino, mpaolo}@chalmers.se

***Abstract*—Spectrum fragmentation and physical-layer impairments pose significant challenges to the efficient operation of elastic optical networks (EONs). While defragmentation techniques can consolidate idle spectrum and reduce blocking, most existing studies neglect the dynamic interplay between amplified spontaneous emission (ASE) noise and nonlinear interference (NLI). In this work, we implement a defragmentation algorithm and conduct a comprehensive performance evaluation of the routing, modulation format, and spectrum assignment (RMSA) problem on the European topology under dynamic network traffic. Three quality-of-transmission (QoT) constraints to the modulation format selection are compared: distance-based reach limits, ASE-only impairment, and ASE+NLI impairment. Our results show that, although defragmentation reduces blocking rates in all cases, relative gains under realistic ASE+NLI conditions are smaller, up to 40% lower than those reported using the simplified distance-based model. At higher loads, the proposed approach triggers fewer defragmentation cycles, although each cycle defragments more lightpaths, which tend to increase control plane signaling overhead. These findings underscore the need for integrating realistic physical-layer impairment models into resource allocation and defragmentation strategies for next-generation EON deployments.**

***Index Terms*—Elastic optical networks, spectrum defragmentation, quality of transmission, physical layer impairments, resource allocation.**

I. Introduction

Elastic optical networks (EONs) have emerged as a promising solution to address the increasing bandwidth demands of modern communication systems [1]. Unlike traditional wavelength division multiplexing (WDM) networks with fixed-grid channel spacing, EONs utilize flexible grid technology, which enables fine-grained spectrum allocation based on actual traffic demands. This flexibility allows for more efficient spectrum utilization through adaptive modulation formats and variable bandwidth allocation, making EONs particularly suitable for handling heterogeneous traffic patterns in core networks [2].

Spectrum fragmentation occurs when dynamic traffic arrivals and departures create multiple small, noncontiguous free spectrum fragments. Individually, these fragments cannot accommodate new connection requests; yet, collectively, they represent a significant amount of unused spectrum [3]. This fragmentation naturally occurs in dynamic environments where connections are frequently established and torn down, resulting in isolated spectrum segments that are unable to meet the continuity and contiguity constraints required for new lightpaths [4]. To mitigate this issue, various defragmentation techniques have been proposed, aiming to consolidate the fragmented spectrum and create larger contiguous blocks capable of accommodating incoming requests [5].

The quality of transmission (QoT) has become an increasingly critical consideration in modern optical networks [6]. In QoT-aware EONs, physical-layer impairments are a primary concern that cannot be ignored in resource allocation decisions. Signal attenuation over long-haul fibers requires optical amplifiers, which introduce amplified spontaneous emission (ASE). Concurrently, the effort to improve spectral efficiency by compacting connections intensifies nonlinear interference (NLI), leading to signal distortion and crosstalk between adjacent channels. Traditional defragmentation approaches often overlook these complex physical-layer interactions, focusing primarily on spectrum consolidation without adequately accounting for the potential QoT degradation that may result from spectrum aspects [7].

Despite extensive research into defragmentation techniques for EONs, a substantial gap remains in understanding the trade-offs between fragmentation reduction and its impact on QoT in realistic physical-layer models. Many current studies use simplified impairment models, often fixed distance thresholds [8]. These simplified models can provide an inaccurate estimate of signal quality by neglecting critical NLI effects. Furthermore, recent proactive defragmentation schemes often exacerbate this issue. These strategies usually prioritize optimizing metrics such as fragmentation and blocking rates without considering the nature of NLI. This approach overlooks a fundamental physical penalty: compacting connections to reduce fragmentation intensifies NLI among adjacent channels, potentially compromising the reliability of the outcomes [9].

This paper's main contribution is demonstrating that when ASE and NLI are considered, the estimated performance gains achieved by defragmentation decrease substantially. Ad-

This work was supported by the National Council for Scientific and Technological Development (CNPq) and the Amazon Sustainability Institute for Science and Innovation (iSACI).

ditionally, we assess the underestimation in blocking rate estimations when simplified models are compared against a more robust ASE+NLI model. These estimations are prone to considerable inaccuracies that vary according to the traffic scenario: significant performance gains occur at lower traffic intensities, while the advantages diminish noticeably at higher loads. These findings highlight the importance of incorporating realistic physical-layer impairment models, particularly those accounting for NLI, into defragmentation and spectrum allocation strategies for EONs to ensure accurate and realistic performance assessments.

II. Related Work

Spectrum fragmentation management in EONs has been addressed at two different stages of the network operation: lightpath provisioning and spectrum defragmentation [5]

During lightpath provisioning, heuristics such as shortest-available-path best-modulation first-fit (SAP-BM-FF) allocate spectrum based on predefined routes and modulation formats, limiting fragmentation with low complexity but reduced effectiveness under dynamic traffic and diverse impairments [8].

Defragmentation methods range from reactive push–pull retuning, which gradually shifts active connections to consolidate gaps with minimal disruption, to hop-retuning, enabling more extensive reorganization at the cost of brief outages [9].

Recent Machine Learning (ML) and Deep Reinforcement Learning (DRL) frameworks couple rapid QoT estimation—predicting ASE and NLI impairments via neural predictors [10]—with adaptive defragmentation control. DRL-based DeepDefrag learns when and how to retune the spectrum, consistently outperforming heuristic baselines in both blocking rate and retuning overhead [4].

Despite these developments, recent literature highlights a crucial gap: fragmentation control and QoT management are predominantly treated separately, overlooking their interplay in realistic impairment scenarios [8]. Motivated by this limitation, our study explicitly examines how realistic physical-layer impairments, specifically ASE and NLI, influence defragmentation outcomes. We demonstrate through simulations that QoT-aware modeling reduces the performance gains of defragmentation compared to the results when simplistic assumptions are made, underscoring the need for integrated impairment-aware resource management strategies.

III. Spectral Defragmentation Algorithm

To evaluate the impact of considering impairments on the defragmentation of EONs, we implement the lightpath defragmentation strategy described in Algorithm 1, which aims to reduce fragmentation while respecting impairment constraints.

The defragmentation process is triggered whenever a previously provisioned lightpath leaves the network, freeing spectrum resources. Then, all active lightpaths S_{active} are evaluated for possible defragmentation, following an oldest-first policy [11]. Candidate spectrum allocations are examined in ascending frequency order, prioritizing connections that can be shifted closer to the lower end of the spectrum while still meeting QoT requirements.

The QoT validation is based on verifying that transmission conditions meet the minimum quality requirements of the selected modulation format. In practice, this is expressed as the difference between the available GSNR and the GSNR threshold, i.e., $\text{GSNR} - \Gamma_m > 0$, or, in distance-constrained scenarios, by ensuring that the optical path length does not exceed the maximum reach. Depending on the chosen constraint model, this condition is evaluated using one of the three approaches detailed below.

1) Distance-based model: This is the simplest approach, completely abstracting the physical layer. It only checks if the path length is within a pre-calculated maximum reach for the modulation format: $l(p_s) \leq R_m$ where $l(p_s)$ is the physical length of path p_s and R_m is the maximum allowed reach (in km) for modulation m.

2) ASE-only model: This model considers only the ASE noise. The defragmentation is permitted if:

$$\text{GSNR}_{\text{ASE}} = \frac{p_{ch}}{p_{\text{ase}}} \geq \Gamma_m \tag{1}$$

where the p_{ase} is calculated using the expression:

$$p_{\text{ase}} = B \sum_{i=1}^{N_{\text{sec}}} f_{n_i}\, h\nu(g_i - 1) \tag{2}$$

where B is the channel bandwidth, N_{sec} is the number of fiber spans, f_{n_i} is the noise figure of the i-th amplifier, h is Planck's constant, ν is the carrier frequency, and g_i is the linear gain of the i-th amplifier. This model is faster but less accurate in densely occupied networks.

3) ASE+NLI model: This is the most comprehensive model, evaluating both linear and nonlinear noise effects. The GSNR is calculated and must satisfy:

$$\text{GSNR} = \frac{p_{ch}}{p_{\text{ase}} + p_{\text{nli}}} \geq \Gamma_m \tag{3}$$

where p_{ch} is the channel's signal power, Γ_m is the required GSNR for the modulation format, and p_{ase} and p_{nli} are the noise powers from ASE and NLI, respectively. The NLI power (p_{nli}) is estimated using the enhanced Gaussian noise (EGN) model [7]. Specifically, the NLI power spectral density for a single span, g_{nli}, is given by:

$$g_{\text{nli}} = \mu \left(\ln(\rho b_i^2) + \sum_{j \neq i} \ln \left| \frac{\Delta f_{ij} + b_j/2}{\Delta f_{ij} - b_j/2} \right| \right) \tag{4}$$

where μ and ρ are coefficients dependent on fiber parameters (nonlinearity γ, dispersion β_2, and attenuation α), b_i and b_j are the bandwidths of the channel under test and an interfering channel, respectively, and Δf_{ij} is the frequency spacing between them. The total NLI power is then

$$p_{\text{nli}} = B \sum_{i=1}^{N_{\text{sec}}} g_{\text{nli},i} \tag{5}$$

Once the QoT condition is satisfied, the lightpath is re-established in a new spectral allocation that frees its previous

Algorithm 1: Spectral Defragmentation Algorithm

Data: $\mathcal{S}_{\text{active}}$ // Set of active lightpaths
Result: Updated network state with defragmented lightpaths

for *each lightpath* $s \in \mathcal{S}_{\text{active}}$ **do**
  /* Get free slots */
  $\mathcal{C} \leftarrow$ findFreeSlots(s.path)
  for *each slot* $c \in \mathcal{C}$ **do**
    /* Check position and QoT */
    if $c < s.startSlot \wedge$ *checkQoTValidity*(s, c) **then**
      /* Free current slot */
      deallocateOldSlots(s)
      /* Assign new slot */
      allocateNewSlot(s, c)
      /* Next lightpath */
      break

allocation. This controlled migration compacts spectrum utilization toward lower frequencies, resulting in improved spectral efficiency and more effective use of available resources while maintaining the required transmission performance.

IV. Results

This section presents the simulation setup and discusses the performance evaluation of spectrum defragmentation under various impairment constraints. Across all scenarios, the network topology and traffic pattern are fixed, and the RMSA follows the same heuristic (SAP-BM-FF).

A. Simulation Setup

This study considers the dynamic lightpath provisioning and defragmentation over EONs. The simulations are executed using the European network topology, identified as *nobel-eu* in SNDlib [12]. Each fiber link supports 320 frequency slots, each with 12.5 GHz, in the C-band, with a fixed guardband requirement between neighboring lightpaths to manage interference.

TABLE I
Simulation Parameters

Parameter	Symbol	Value
Fiber type	–	Standard Single Mode Fiber (SMF)
Bandwidth	B	4 THz
Slot bandwidth	Δf	12.5 GHz
Guard-band	G_b	1 slot
Amplifier noise figure	f_n	4.5 dB
Amplifier gain	g	16 dB
Fiber attenuation (α)	α	0.2 dB/km
Span length	L_{span}	max. 80 km
Wavelength	λ	$\approx$ 1534–1566 nm
Fiber dispersion (β_2)	β_2	-21.3 ps^2/km
Nonlinear coefficient	γ	1.3 W^{-1}km^{-1}

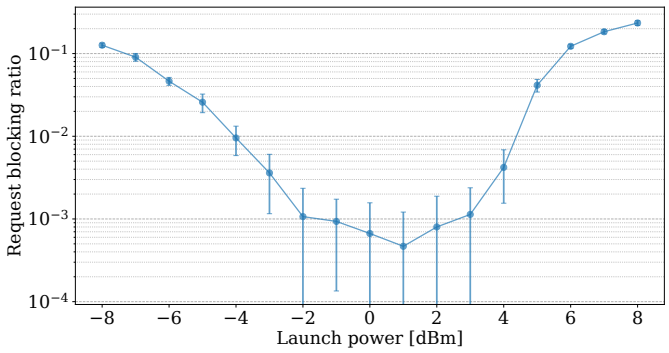

Fig. 1. Request blocking rate vs. launch power for the European topology considering a load of 200 Erlang.

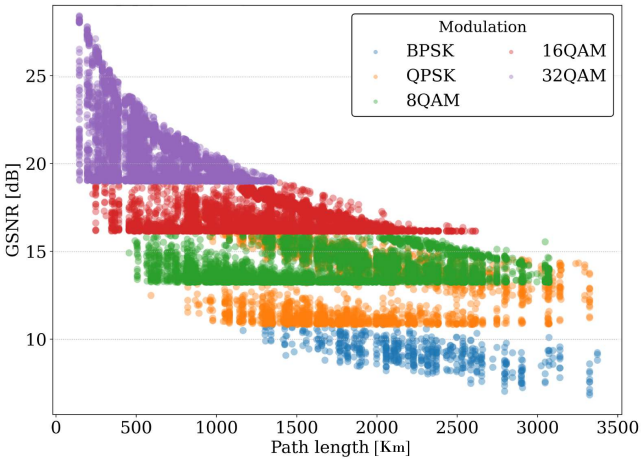

Fig. 2. Modulation vs. path length for the European topology considering a load of 200 Erlang and with 1 dBm as the launch power.

The dynamic traffic model features 50,000 requests (i.e., 50 episodes of 1,000 arrivals each) with uniformly distributed source/destination nodes and uniformly distributed bit rates of {10, 40, 100, 400} Gbps. Table I provides further physical layer details.

The launch power for all lightpaths was optimized based on an extensive search conducted over a range from -8 dBm to 8 dBm, using a traffic load of 200 Erlang and the EGN model [7], with the same network and traffic settings used in the simulations. Fig. 1 shows that the launch power yielding the lowest blocking rate is 1 dBm, providing a balanced trade-off between ASE and NLI [13].

For transmission, we consider the following modulation formats: BPSK, QPSK, 8-QAM, 16-QAM, and 32-QAM. Under GSNR-based constraints, the ASE+NLI model is calculated according to Eq. 3 and the ASE-only model according to Eq. 1. The corresponding thresholds are {3.71, 6.72, 10.84, 13.24, and 16.16} dB, respectively [13].

The distance-based thresholds were derived using an optimistic distance-adaptive procedure with an optimized launch power of 1 dBm (see Fig. 2). This procedure produced modulation-dependent reach values of {3420, 3370, 3130,

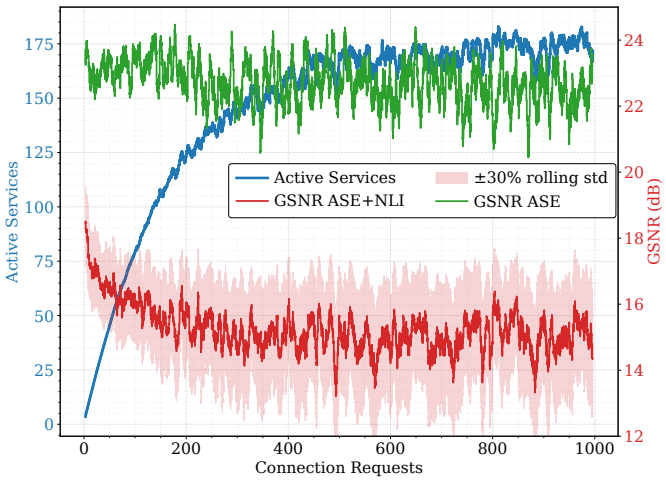

Fig. 3. GSNR vs. request arrivals and active lightpaths for load of 400 Erlang.

2650, and 1390} km for BPSK, QPSK, 8-QAM, 16-QAM, and 32-QAM, respectively.

Additionally, the fragmentation state of the network is evaluated through three metrics: route cuts, which measure the continuity of available spectrum along a path; Residual spectral spread (RSS), quantifies the compactness and distribution of unused spectrum; and Shannon entropy (SE), which reflects the overall randomness of spectral occupancy [4].

The RMSA problem is solved using a QoT-aware approach under three alternative modulation format constraint methods: *i)* distance-based, which considers the maximum transmission reach of each modulation format; *ii)* ASE-only, which accounts solely for ASE impairment; and *iii)* ASE+NLI, which considers both ASE and NLI impairments. The GSNR in scenarios *ii* and *iii* is calculated using the EGN model [7].

B. Performance Evaluation

Figure 3 traces GSNR evolution from the initial ramp-up to a stable operational regime for all accepted requests. During the ramp-up phase (0–400 requests), the GSNR under ASE+NLI decreases rapidly from approximately 18 dB to around 15 dB, whereas the ASE-only scenario maintains a GSNR consistently above 22 dB. From approximately 400 connection requests onward, the network reaches a steady operational state, with ASE+NLI GSNR stabilizing around an average of 15.24 dB (standard deviation: 0.67 dB), compared to an average of 22.73 dB (standard deviation: 0.63 dB) for the ASE-only scenario, resulting in a mean difference of roughly 7.49 dB. The shaded area around the ASE+NLI curve represents ±30% of the rolling standard deviation, highlighting the variability arising from NLI effects, which consider co-propagating channel interactions and path diversity. Additionally, a positive correlation (r = 0.314) between the number of connection requests and the ΔGSNR (GSNR_{ASE} - $\text{GSNR}_{\text{ASE+NLI}}$) suggests a slight growth in nonlinear penalties with increasing traffic load, emphasizing the importance of incorporating ASE+NLI effects into spectrum defragmentation decisions.

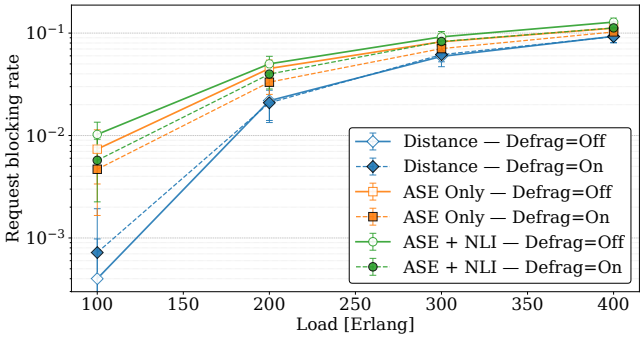

Fig. 4. Blocking rate per QoT and defragmentation scenario.

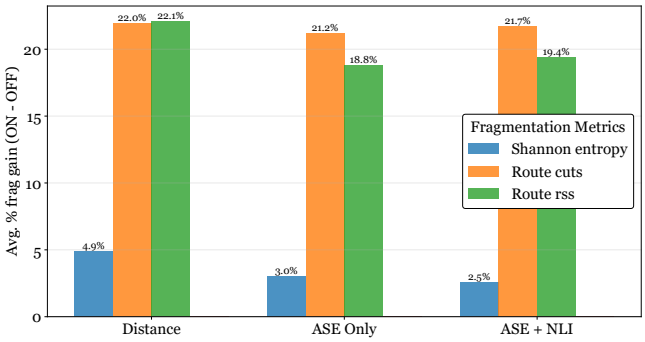

Fig. 5. Average improvement in fragmentation metrics (ON vs. OFF) under different QoT models.

As illustrated in Figure 4, the blocking rate of requests is evaluated under scenarios with defragmentation states enabled and disabled, for varying traffic loads. Across all the tested scenarios, defragmentation consistently reduces blocking rates, although the extent of the improvement varies significantly based on the impairment model considered.

In the ASE+NLI scenario, defragmentation reduces the blocking rate from 16.7% at higher traffic loads (400 Erlang) to 63.98% at lower traffic loads (100 Erlang). The ASE-only scenario shows similarly consistent performance gains, with reductions ranging from 16.16% at 400 Erlang to 64.52% at 100 Erlang. By contrast, the distance-based scenario skews conclusions, with gains remaining modest at higher traffic loads (19.37% at 400 Erlang) but appearing very large at lower traffic loads (up to 89% at 100 Erlang). Defragmentation brings real gains, particularly at low loads, yet these gains are often overestimated when using simplified assumptions. For instance, at 100 Erlang the gain is 63.98% under ASE+NLI and 89.53% under a distance-based model, which overstates the gain by 25.55 p.p. ($\approx$ 40% relative), a significant error.

Accordingly, a comparative analysis against the ASE+NLI baseline reveals a consistent and significant underestimate of blocking probabilities. For the ASE-only scenario, errors range from 5.2% at higher traffic loads (400 Erlang) to 13.8% at lower traffic loads (100 Erlang) when defragmentation is disabled. With defragmentation enabled, these errors decrease

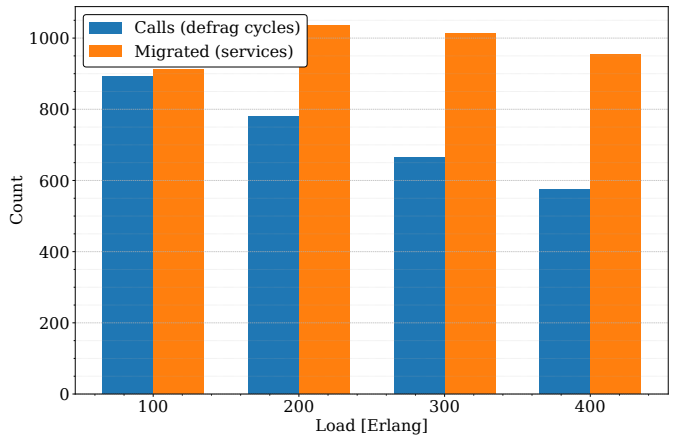

Fig. 6. Number of defragmentation cycles and lightpath defragmentations.

to between 3.3% (400 Erlang) and 12.5% (100 Erlang). The distance-based model shows even greater discrepancies, with errors ranging from 21.1% at 400 Erlang to 91.4% at 100 Erlang when defragmentation is disabled. With defragmentation enabled, the results remain large or even increase, ranging from 23.6% at higher loads to 97.5% at lower loads.

Figure 5 reports improvement of all fragmentation metrics under the different QoT-constraint scenarios. The results show a substantial improvement for both RSS and Route cuts, with gains consistently ranging from 18.8% to 22.1%. This indicates that defragmentation effectively creates a more compact spectrum allocation and a more uniform spectral distribution. In contrast, the SE shows only a modest improvement, ranging from 2.5% to 4.9% across the scenarios. These findings highlight that simplified models not only substantially underestimate blocking rates but also overstate the relative performance gains from defragmentation, particularly at lower traffic intensities. This underlines the importance of using precise models, such as ASE+NLI, to ensure accurate QoT assessments.

Fig. 6 complements this analysis, illustrating the average number of defragmentation cycles ("calls") and lightpath defragmentations ("migrated") per episode. At lower loads, approximately 890 defragmentation cycles result in around 910 lightpaths defragmented per episode. With increased load, cycles decrease to around 580, while lightpaths defragmented remain consistently high ($\approx$ 950). As the load increases, the defragmentation algorithm is triggered less frequently, yet each trigger defragments more lightpaths. This leads to fewer yet more disruptive reconfigurations, resulting in a tendency toward increased control-plane overhead. Consequently, it is necessary to balance the decrease in blocking rate and the higher operational cost.

V. Conclusion

This study assessed the effectiveness of spectral defragmentation in EONs, considering realistic QoT conditions. Defragmentation significantly reduces blocking rates, but the magnitude of these benefits is highly dependent on the impairment model used. Scenarios that neglect nonlinear interference or use simplified constraints substantially underestimate blocking, thereby overstating the benefits of defragmentation. The ASE+NLI scenario, realistically modeling ASE and NLI, showed meaningful but reduced defragmentation gains, underscoring the need for accurate QoT modeling. Additionally, lightpath defragmentation analysis reveals a trade-off between defragmentation frequency and operational overhead, particularly under higher loads.

References

[1] R.-J. Essiambre, G. Kramer, P. J. Winzer, G. J. Foschini, and B. Goebel, "Capacity limits of optical fiber networks," *Journal of Lightwave Technology*, vol. 28, no. 4, pp. 662–701, Feb. 2010. DOI: 10.1109/JLT.2009.2039464.

[2] Y. Zhang, J. Xin, X. Li, and S. Huang, "Overview on routing and resource allocation based machine learning in optical networks," *Optical Fiber Technology*, vol. 60, p. 102355, 2020. DOI: 10.1016/j.yofte.2020.102355.

[3] E. Rezaghoilzadeh, S. Kohlert, R. Kashyap, and B. Sansò, "Defragmentation-aware multilayer route and spectrum assignment in elastic optical network," GERAD, Tech. Rep. G-2025-03, January 2025.

[4] E. Etezadi, C. Natalino, R. Diaz, A. Lindgren, S. Melin, L. Wosinska, P. Monti, and M. Furdek, "Deep reinforcement learning for proactive spectrum defragmentation in elastic optical networks," *Journal of Optical Communications and Networking*, vol. 15, no. 10, pp. E86–E96, 2023. DOI: 10.1364/JOCN.489577.

[5] B. C. Chatterjee, S. Ba, and E. Oki, "Fragmentation problems and management approaches in elastic optical networks: A survey," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 1, pp. 183–210, 2018. DOI: 10.1109/COMST.2017.2769102.

[6] L. Zhang, X. Li, Y. Tang, J. Xin, and S. Huang, "A survey on QoT prediction using machine learning in optical networks," *Optical Fiber Technology*, vol. 68, p. 102804, 2022. DOI: 10.1016/j.yofte.2021.102804.

[7] M. R. Zefreh, F. Forghieri, S. Piciaccia, and P. Poggiolini, "Accurate closed-form real-time EGN model formula leveraging machine-learning over 8500 thoroughly randomized full C-band systems," *Journal of Lightwave Technology*, vol. 38, pp. 4987–4999, 2020. DOI: 10.1109/JLT.2020.2981194.

[8] D. Yan, N. Feng, J. Lv, D. Ren, J. Hu, and J. Zhao, "DRL-based fragmentation- and impairment-aware resource allocation algorithm in C + L band elastic optical networks," *Optical Fiber Technology*, vol. 90, p. 104133, 2025. DOI: 10.1016/j.yofte.2025.104133.

[9] R. K. Jana, B. C. Chatterjee, A. P. Singh, A. Srivastava, B. Mukherjee, A. Lord, and A. Mitra, "Machine learning-assisted nonlinear-impairment-aware proactive defragmentation for C + L band elastic optical networks," *Journal of Optical Communications and Networking*, vol. 14, no. 3, pp. 56–68, 2022. DOI: 10.1364/JOCN.440214.

[10] G. Bergk, B. Shariati, P. Safari, and J. K. Fischer, "ML-assisted QoT estimation: a dataset collection and data visualization for dataset quality evaluation," *Journal of Optical Communications and Networking*, vol. 14, pp. 43–55, 2022. DOI: 10.1364/JOCN.442733.

[11] J. Comellas, L. Vicario, and G. Junyent, "Proactive defragmentation in elastic optical networks under dynamic load conditions," *Photonic Network Communications*, vol. 36, no. 1, pp. 26–34, 2018. DOI: 10.1007/s11107-018-0767-7.

[12] S. Orlowski, R. Wëssäly, M. Pióro, and A. Tomaszewski, "Sndlib 1.0–survivable network design library," *Networks*, vol. 55, pp. 276–286, 2010. DOI: 10.1002/net.20371.

[13] C. Natalino, T. Magalhães, F. Arpanaei, F. R. L. Lobato, J. C. W. A. Costa, J. A. Hernández, and P. Monti, "Optical Networking Gym: An open-source toolkit for benchmarking resource assignment problems in optical networks," *Journal of Optical Communications and Networking*, 2024. DOI: 10.1364/JOCN.532850.