Journal of Internet Services and Applications, 2026, 17:1, doi: 10.5753/jisa.2026.7201
This work is licensed under a Creative Commons Attribution 4.0 International License.

APACHE: Fragmentation and Spectrum-Aware Provisioning for Elastic Optical Networks with Service-Class Differentiation and Revenue Analysis

Carlos Salvador Gomide Richa *[Universidade Federal Fluminense | salvador.richa@hotmail.com]*
Eonassis Oliveira Santos *[Universidade Federal da Bahia | eonassissantos@ufba.br]*
Alex Silva Santos *[Instituto Federal de Educação, Ciência e Tecnologia da Bahia | alex.silva@ifba.edu.br]*
Gustavo Bittencourt Figueiredo *[Universidade Federal da Bahia | gustavobf@ufba.br]*
Antonio Augusto de Aragão Rocha *[Universidade Federal Fluminense | arocha@ic.uff.br]*

Institute of Computing, Fluminense Federal University (UFF), Av. Gal. Milton Tavares de Souza, s/n, Campus da Praia Vermelha, Boa Viagem, Niterói, RJ, 24210-346, Brazil.

Received: 16 December 2025 • **Accepted:** 14 March 2026 • **Published:** 23 May 2026

Abstract

Elastic Optical Networks (EONs) enable fine-grained spectrum allocation that adapts to heterogeneous Internet services. Yet, fragmentation and the lack of contiguous free slots remain persistent causes of connection blocking and SLA violations. This article extends our SBRC paper from 2024 by providing a comprehensive journal version of **APACHE**—an analytical provisioning algorithm that (i) distinguishes blocking due to *fragmentation* from blocking due to *true scarcity*, (ii) incorporates *Class of Service* (CoS) differentiation into the RSA, and (iii) evaluates the economic impact of blocking via *revenue loss*. We detail the algorithmic design, simulation methodology, and an expanded performance assessment on the USA and RNP topologies. Results show that APACHE reduces circuit blocking up to 80% (USA) and 56% (RNP) compared to CFCoSP and TSSCF baselines, while consistently minimizing revenue loss under mixed-demand traffic profiles. The findings highlight how fragmentation-awareness and service-class prioritization jointly improve the efficiency and business viability of EON provisioning for next-generation Internet services.

Keywords: Elastic Optical Networks, RSA, Fragmentation, Class of Service, Revenue Loss, Internet Services

1 Introduction

The growing demand for data services has imposed significant challenges on optical networks, especially with regard to providing efficient, scalable, and reliable connectivity. In *Dense Wavelength Division Multiplexing* (DWDM) networks, the use of fixed spectral grids may lead to bandwidth waste, since the allocated capacity does not always match the heterogeneous demand profile generated by upper-layer applications. This mismatch compromises spectrum efficiency and limits service expansion.

In this context, *Elastic Optical Networks* (EONs) have emerged as a promising alternative to conventional DWDM architectures Jinno *et al*. [2009]; Gerstel *et al*. [2012]. By enabling flexible spectrum allocation, EONs allow optical resources to be assigned according to the bandwidth requirements of each connection request, improving spectral utilization and supporting heterogeneous Internet services. However, this flexibility also introduces important provisioning challenges. In particular, EONs must satisfy routing, modulation, spectrum continuity, and spectrum contiguity constraints, while also dealing with dynamic traffic arrivals, variable holding times, and different service requirements.

Among these challenges, spectrum fragmentation and scarcity of contiguous free slots are major causes of connection blocking. Fragmentation occurs when free spectrum is available but distributed in non-contiguous portions, preventing the establishment of new lightpaths that require adjacent slots. Conversely, spectrum scarcity occurs when the network lacks enough free slots to accommodate a request, regardless of contiguity. Although both situations lead to blocking, they have different operational meanings and require different management actions. Fragmentation may indicate the need for defragmentation or more efficient spectrum assignment, whereas true spectrum scarcity may indicate the need for capacity expansion or traffic engineering adjustments.

Most existing approaches for Routing and Spectrum Assignment (RSA) or Routing, Modulation Level, and Spectrum Assignment (RMLSA) seek to reduce blocking probability, improve resource utilization, or support differentiated services. However, many of them treat blocking as a single aggregated phenomenon, without explicitly identifying whether a request was blocked due to fragmentation or due to the absence of free spectrum. Moreover, economic aspects, such as revenue loss caused by blocked demands, are often analyzed separately from technical provisioning decisions. This separation limits the ability of operators to understand not only how many requests were blocked, but also why they were blocked and what financial impact such blocking produced.

To address this gap, this paper presents APACHE — *Analytical Provisioning Algorithm of Critical Hop Edge*, a fragmentation- and spectrum-aware provisioning algorithm for dynamic traffic in EONs. APACHE evaluates candidate paths according to the criticality of their links, considering the size of the largest contiguous block of free slots along each path. By assigning higher cost to links with reduced contiguous availability, the algorithm favors routes that are less likely to intensify fragmentation and future bottlenecks. In addition, APACHE incorporates Class of Service (CoS) differentiation and holding-time information into the provisioning decision, allowing higher-priority and higher-value demands to be treated according to their service requirements.

The main contribution of this work is a unified provisioning framework that combines: (i) fragmentation-aware path selection, (ii) explicit differentiation between blocking caused by fragmentation and blocking caused by spectrum scarcity, (iii) service-class differentiation in the allocation process, and (iv) revenue-loss analysis associated with blocked requests. This combination provides both technical and economic insight into EON operation, supporting more informed decisions regarding resource allocation, network planning, and service prioritization.

Simulation results on the USA and RNP topologies show that APACHE reduces blocking probability when compared with reference approaches, including CFCoSP and TSSCF. In particular, the results indicate reductions of up to 56% in the RNP topology and up to 80% in the USA topology, while also improving the diagnosis of blocking causes and reducing revenue loss under mixed-demand traffic scenarios. These results suggest that fragmentation awareness and service-class prioritization can jointly improve both provisioning efficiency and the business viability of EON operation.

The remainder of this paper is organized as follows. Section 2 reviews theoretical foundations and related work. Section 3 introduces the problem of link criticality and resource availability in EONs. Section 4 presents the proposed APACHE algorithm. Section 5 describes the simulation setup and discusses the numerical results. Finally, Section 6 concludes the paper and outlines directions for future research.

2 Related Work

Several studies focus on the development of algorithms to address challenges in EONs, as summarized in Table 1. APACHE algorithm appears in Row 18 of the table.[1]

The table highlights the main aspects covered in related work, including the performance metrics considered, the routing, modulation, and spectrum allocation (RMLSA) Chatterjee *et al*. [2015] technique, multi-domain (*MDON*) support Batham and Yadav [2020a], economic treatment, load balancing, protection, disaster recovery, service degradation, fragmentation, application-awareness, multilayer strategies, and negotiation mechanisms.

For example, Horota *et al*. [2015] (Row 10, Table 1) proposed the RSA-MFPF (*RSA Algorithm with Most Fragmented Path First*), which prioritizes selecting the most fragmented path in order to reduce fragmentation. Its performance is compared with contemporary algorithms such as DF (*Degree of Fragmentation*), AP (*Acceptance Prone*), and P-KSP (*Precomputed K-Shortest Paths*) Wang *et al*. [2011] (Row 11, Table 1). The latter selects the least-cost path among k precomputed path candidates with sufficient spectrum, without explicitly considering spectrum state. As one of the earliest algorithms in this field, P-KSP remains widely used as a baseline, including this paper.

Santos *et al*. [2018] (Row 15, Table 1) evaluated an algorithm that reduces spectrum fragmentation in EONs by choosing lightpaths that minimize fragmentation after allocation. This work also introduced the EON simulator later adapted and extended for the present study.

Other examples include Batham and Yadav [2020b,a] (Rows 5 and 3, Table 1), who developed cost-function-based algorithms to reduce blocking in different scenarios, and Batham *et al*. [2020] (Row 2, Table 1), who proposed the TSSCF (*Traffic Scheduling Strategy based on Cost Function*). TSSCF incorporates metrics such as path length, holding time, and CoS into a cost function and extends to multi-domain provisioning. Its evaluation considered blocking probability, revenue, and resource utilization. TSSCF is also used here as a comparison baseline.

Similarly, Dixit *et al*. [2020] (Row 1, Table 1) presented the CFCoSP (*Cost Function-based CoS Provisioning*), which classifies traffic into CoS1, CoS2, and CoS3 and uses a cumulative cost function based on CoS, path length, and bandwidth demand. Simulation results showed that CFCoSP outperforms simple CoS provisioning (CoSP) in terms of established connections, bandwidth blocking, and network utilization. The model also accounts for guard bands and spectral efficiency, assumptions adopted in our APACHE algorithm.

Recent research has also addressed service degradation in EONs. Lourenço *et al*. [2018] (Row 14, Table 1) studied provisioning under resource crunch conditions, where connections may tolerate reduced bandwidth to minimize revenue loss while accommodating new requests. Santos *et al*. [2022] (Row 16, Table 1) proposed an application-aware EON degradation strategy that leverages QoS models and cross-layer information to mitigate the impact of resource unavailability on delay- and bandwidth-sensitive applications. Unlike these approaches, the present work does not focus on degradation.

3 The Problem of Link Criticality in Resource Availability in Elastic Optical Networks

Efficient resource provisioning in EONs requires more than simply tracking how much spectrum is free across the network; it requires understanding why spectrum becomes unavailable and where bottlenecks emerge. Although the literature often emphasizes reducing blocking probability or mitigating physical-layer impairments, most algorithms treat blocking as a single, unified phenomenon Júnior *et al*. [2025]. In practice, however, the causes of blocking differ substantially, and distinguishing between fragmentation-driven failures and genuine spectrum scarcity is fundamental for operational decisions, maintenance strategies, and long-term capacity planning.

Central to this discussion is the notion of link criticality. The criticality of a link is not defined solely by its overall occupancy but by its spectral structure, particularly the size of its largest contiguous block of free slots. Because EON con-

[1] Although several efforts, to add the authors and their references directly into the table, this was not possible, due to space limitations, these details are instead cited in the text (e.g., Row 1 of Table 1).

Table 1. Related Work

Ref.	Metrics	RMLSA	MDON	Econ.	Bal.	Prot.	Dis.	Degr.	Frag.	Aware	Multi.	Negot.
1	BP, RUR	RSA	✓	✗	✗	✗	✗	✗	✗	CoS	✗	✗
2	BP, RUR, Rev.	RSA	✓	✓	✗	✗	✗	✗	✗	CoS	✗	✗
3	BP, RUR	RSA	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗
4	BP, RUR	RSA	✓	✗	✗	✗	✗	✗	✗	CoS	✗	✗
5	BP, RUR	RSA	✓	✗	✗	✗	✗	✗	✗	CoS	✗	✗
6	BP, SUR, FI	RSA	✓	✗	✓	✗	✗	✗	✓	CoS	✗	✗
7	BP, RUR	RSA	✓	✗	✓	✓	✗	✗	✗	✗	✗	✗
8	RO, LL, BP	RSA	✓	✗	✗	✓	✗	✗	✗	CoS	✗	✗
9	BP, RUR	RSA	✓	✗	✓	✗	✗	✗	✗	✗	✗	✗
10	BP, FI	RSA	✗	✗	✗	✗	✗	✗	✓	✗	✗	✗
11	BP, BBP	RSA	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
12	Rev., RUR	RSA	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗
13	BP, Rev.	RSA	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗
14	BP, Rev.	RSA	✗	✓	✗	✗	Crunch	✓	✗	CoS	✗	✗
15	BP, BBP	RMLSA	✗	✗	✗	✗	✗	✓	✓	✗	✗	✗
16	BP, BBP	RMLSA	✓	✓	✗	✗	Crunch	✓	✗	App	✓	✗
17	BP, BBP, Rev.	RMLSA	✗	✓	✗	✗	✗	✗	✗	CoS	✗	✗
18	APACHE	RMLSA	✓	✓	✓	✓	✓	✓	✓	CoS	✓	✓

nections must satisfy continuity, contiguity, and modulation constraints, even links with considerable free spectrum may become severe bottlenecks if heavily fragmented. Thus, accurately characterizing contiguity at the link level is essential for effective provisioning.

Importantly, link criticality is not a static or purely topological attribute. It emerges dynamically from routing decisions Rodrigues and Oliveira [2025]. As requests are served over time, some links naturally receive more traffic due to their position in the k-shortest-paths set, their proximity to central nodes, or the routing strategy adopted. These frequently used links tend to accumulate fragmentation faster, exhausting their ability to accommodate contiguous blocks. As traffic imbalances grow, they become the first points at which new requests fail, not because the network lacks spectrum globally, but because routing choices have concentrated fragmentation and slot exhaustion on specific edges. In this sense, link criticality is a direct product of routing policies and traffic evolution, rather than an inherent characteristic of the physical topology.

This dynamic behavior exposes an important limitation in most RMLSA algorithms. While many solutions improve blocking probability, spectral efficiency, or service differentiation, they rarely distinguish between blocking caused by fragmentation and blocking caused by true scarcity. When an allocation fails, operators are often unable to determine whether: (1) the network had enough free slots in aggregate but could not form a contiguous block on a critical hop, or (2) the network had genuinely exhausted its spectrum across all feasible routes.

The absence of this diagnostic insight hinders network management. If fragmentation is the dominant cause, defragmentation routines Shi *et al*. [2014]or spectrum optimization techniques may restore performance. If scarcity dominates, capacity upgrades must target specific high-stress links rather than the network as a whole. When these causes vary by Class of Service (CoS) Luo and Zhu [2021]; Santos *et al*. [2024a], adjustments to SLA priorities, pricing models, or traffic engineering policies may be necessary. Without this differentiation, operators risk reacting to symptoms rather than underlying structural problems.

The complexity increases when traffic heterogeneity is introduced. Services may request different bandwidths, have diverse holding times, or belong to distinct CoS categories. High-priority CoS traffic often requires large contiguous bandwidth blocks and stricter guarantees, while lower-priority traffic can tolerate delays or blocking. These differences shape fragmentation patterns: for example, large-bandwidth requests intensify fragmentation, and premium CoS demands require differentiated handling. Consequently, routing decisions that ignore the interplay between spectral state, CoS requirements, and economic impact can lead to cascading inefficiencies.

Addressing these challenges requires an integrated strategy that combines spectrum awareness, fragmentation sensitivity, CoS prioritization, and economic considerations. Effective provisioning demands evaluating each candidate path not only by distance or hop count but by the spectral structure of its links and the dynamic criticality produced by previous routing decisions Santos *et al*. [2024a]. APACHE addresses this need by explicitly modeling the critical hop edge and assigning higher cost to links with small contiguous free blocks. This cost is further coupled with the number of required slots, holding time, and CoS weight. As a result, APACHE shifts the paradigm from simply selecting a feasible k-shortest path to choosing the route that minimizes future fragmentation propagation, avoids evolving bottlenecks, and preserves high-value traffic classes. Equally important, APACHE separates the cause of each blocking event, identifying whether a failure stemmed from fragmentation or true lack of spectrum—an essential capability for precise network planning.

Viewed holistically, the interaction between routing decisions and link criticality is central to understanding EON performance. Traffic routed early in the network influences spectral availability for all subsequent requests, and the cumulative effect of these decisions determines which links become bottlenecks later. By incorporating this interdependence directly into its cost function and decision-making logic Oliveira and Oliveira [2025], APACHE achieves significant reductions in blocking probability and revenue loss, while also providing deeper operational insight into network behavior. This integrated perspective allows providers to act proactively—managing fragmentation, planning upgrades, and adjusting CoS policies—rather than reacting to unexpected congestion and failures.

4 The APACHE Algorithm

This work proposes a heuristic for circuit allocation in response to EON network requests, called APACHE (*Analytical Provisioning Algorithm of Critical Hop Edge*). This section presents the algorithm developed to guide the choice of the best path alternatives through spectral analysis, while also differentiating the causes of blocking.

4.1 Network and Lightpath request Models

The network is modeled as a Graph $G(N, E, DST)$, where N corresponds to nodes, E are links, and DST represents the physical distance between nodes. The physical layer is not part of this modeling. Lightpath requests arrive dynamically and must be provisioned according to spectrum availability, modulation constraints, service-class priorities, and link criticality. Each request is represented by the tuple

$$r_i = (s_i, d_i, b_i, h_i, c_i),$$

where s_i and d_i denote the source and destination nodes, b_i is the requested bandwidth (Gb/s), h_i is the holding-time, and c_i is the Class of Service (CoS). This set of parameters defines both the spectral requirements of the request and its relative priority in the allocation process.

Unserved bandwidth is associated with a lost revenue U$. Depending on the physical distance between source and destination (sum of the DST values of the links in the path), different modulations can be used to increase the transmission rate, as shown in Table 2.

Table 2. Relationship between modulation, reach, and transmission rate

Modulation	Reach (km)	Transmission rate (Gb/s)
BPSK	4000	12.5
QPSK	2000	25.0
8-QAM	1000	37.5
16-QAM	500	50.0

The requested bandwidth b_i, combined with the physical distance between s_i and d_i, determines the modulation format used for transmission. APACHE selects the highest-order modulation whose reach is compatible with the total path length, choosing among BPSK, QPSK, 8-QAM, and 16-QAM, as summarized in Table 2. The selected modulation defines the number of contiguous optical slots required to establish the lightpath. Slot assignment follows the First-Fit policy, ensuring continuity and contiguity across all links of the chosen route.

Requests are grouped into three service classes. CoS 1 corresponds to high-priority or real-time services, CoS 2 represents moderately sensitive applications, and CoS 3 includes delay-tolerant or background traffic. The CoS value c_i directly influences APACHE's cost function through a normalization factor that adjusts the penalty associated with each candidate path, thus favoring higher-priority classes during contention for scarce spectral resources.

For each incoming request r_i, APACHE evaluates the k-shortest candidate paths between s_i and d_i (computed via Yen's algorithm) and, for every link in these paths, inspects the size of the largest contiguous block of free slots. This information is used to compute the *criticality* of each edge according to the formulation in Section 4.2 (Equations(1)–(3)), where links with smaller contiguous blocks become more critical. The resulting per-edge criticality values are then combined with the number of required slots, the holding time h_i, and the CoS c_i to obtain the total cost of each candidate path.

A request is successfully provisioned only if a contiguous block of slots is available along all links of at least one candidate path; otherwise, it is blocked. The model explicitly records whether blocking occurred due to fragmentation (insufficient contiguous slots) or true scarcity (no free slots available), enabling a detailed analysis of how the request parameters and the evolving criticality of links jointly affect provisioning outcomes in elastic optical networks.

4.2 Link criticality computation

The Equation 1 below calculates the criticality of each edge in a candidate path for serving a request in an EON. The concept of criticality involves assigning an importance or priority measure to each link based on the maximum size of contiguous slot blocks that can be accommodated on that link. Essentially, the smaller contiguous slot block size on a link, the higher the criticality assigned to that edge. This means that links with smaller contiguous capacity are considered more critical to the proper functioning of the network. This equation is important because it helps identify and prioritize paths without links that may become bottlenecks. Assigning higher criticality to such links enables more efficient resource allocation, helping to avoid congestion and maximize the use of available capacity.

The criticality of each edge (P_{edge}) is assigned according to Equation 1, where b_l is the maximum contiguous block of slots on link l ($l \in L_c$), continuous along the entire candidate path L_c:

$$P_{edge} = \frac{1}{max(b_l)} \tag{1}$$

The cost function of the edge (F_A) is represented in Eq. 2. In this equation, FS_i represents the number of slots required to serve the i-th request, and HT_i represents the holding time.

$$F_A = FS_i \times HT_i \times P_{edge} \tag{2}$$

The final cost function (F_c) represents the sum of edge cost functions across the analyzed path. It also includes the fraction of the request's CoS divided by the total number of CoS (CoS_{total}) existing in the network, to handle Classes of Service. F_c is given by Eq. 3, representing the cost that penalizes or favors the acceptance of a request:

$$F_c = \frac{CoS}{CoS_{total}} + \sum_{l \in L_c} F_A \tag{3}$$

4.3 The proposed algorithm

The APACHE algorithm looks for the pre-computed k-shortest paths between source and destination for each new request. From this list, k path alternatives are tested, with their intermediate nodes between source and destination, excluding paths that traverse edges with no free slots. For each combination, the objective function cost is computed. After sorting, the path with the lowest cost and available resources is selected. Resources are then allocated and recorded.

Routing in this work uses the Yen Algorithm (YEN, 1971), which is widely used to find multiple shorter (or nearly shorter) paths between a pair of nodes in a graph. Specifically, it is known to find the best single disjoint paths (paths without cycles) between two nodes. As modulation formats, Binary Phase-Shift Keying (BPSK), Quadrature Phase-Shift Keying (QPSK), 8, and 16 Quadrature Amplitude Modulation (QAM) were used. The frequency spectrum allocation policy is implemented using First Fit.(ZHANG; ANDREWS; JR., 2005)

The main input parameters of APACHE are the vectors: $topology[]$ containing network state data; $caminhos[]$ containing candidate paths; and $request[origin, destination]$ with source, destination, bandwidth, holding-time, and CoS. The output is the vector $path[]$ containing the ordered paths.

Algorithm 1 APACHE

Input: $topology[], paths[], request[o, d][]$
Output: $paths[]$
1: $cost \leftarrow [][]$
2: $slots\_req \leftarrow []$
3: $slots\_free \leftarrow []$
4: $holding\_time \leftarrow request[o, d][holding\_time]$
5: **for** $i \leftarrow 0$ **to** $|paths| - 1$ **do**
6: $\quad slots\_req \leftarrow modulation(paths[i][distance],$
7: $request[o, d][bandwidth])$
8: $\quad slots\_free \leftarrow free\_slots(paths[i])$
9: $\quad$ **if** $slots\_free \leq 0$ **then**
10: $\qquad Continue$
11: $\quad$ **end if**
12: $\quad cost[i][pos] \leftarrow i$
13: $\quad cost[i][value] \leftarrow slots\_req \times holding\_time \times (1/slots\_free)$
14: **end for**
15: $cost \leftarrow ascending\_sort(cost[][value], cost[][])$
16: $paths[] \leftarrow paths(cost[][pos])$

The algorithm begins with the declaration of the $[cost]$ matrix and the solution and free-slot vectors as empty (Lines 2–4). In Line 5, the variable $holding\_time$ receives the holding time of the request.

In Line 6, the algorithm begins iterating over the number of alternative paths to be considered (in this study, five paths, corresponding to the size of $caminhos[]$). In Line 7, the variable $slots\_sol$ receives the result of the modulation function, which, given the requested bandwidth and the source-destination distance (sum of link lengths in the path), computes the modulation and returns the required number of slots.

In Line 8, the variable $slots\_livres$ receives the result of the s_livres function, which, based on the path under test, returns the contiguous free slots along the path. If $slots\_livres = 0$ (Line 9), there is no availability, and the loop continues to the next iteration (Line 10).

Next, the $[custo]$ vector stores the path position $(Line12)$ and the cost calculated (Line 13) as the product of required slots and holding-time divided by free slots. The loop ends in Line 14.

This part of the algorithm computes the equations introduced earlier. The results are then sorted: in Line 15, the $[custo]$ matrix is sorted in ascending order, and in Line 16, the caminhos[] vector receives the ordered paths. The output is thus the set of alternative paths mapped with their respective costs, prioritizing the lowest-cost one.

4.3.1 Complexity Analysis

The computational complexity of the APACHE algorithm is as follows. Let N be the number of nodes, E the number of links in the topology, and S the number of slots.

In the offline phase (first stage), the k-shortest paths are found between source-destination pairs (here $k = 5$) using Yen's algorithm, with complexity $O(KN(E + N \log N))$.

In the online phase, the search for available frequency slots between source-destination pairs for each new request has complexity $O(SE)$.

Thus, the final complexity of the APACHE algorithm is given by $O(KN(E + N \log N) + SE)$

5 Numerical Examples

An ad hoc Python simulator was developed using the NetworkX and SimPy libraries Santos *et al.* [2024b] to enable the performance evaluation of the proposed APACHE Algorithm. Simulations were conducted using the independent replications method to achieve a 95% confidence level (note that although plotted, the confidence intervals are not visible in the graphs because they are too small relative to the figure size). A predefined number of 100,000 randomly generated requests was used in each of the sixteen simulation runs for every point in the plotted curves. Dynamic traffic was employed in this work. Records were ordered over time, allowing the observation of traffic modeling. Request generation among source-destination pairs followed a uniform probability distribution. These requests were uniformly distributed into three service classes: CoS1 (real-time traffic), CoS2 (non-real-time traffic), and CoS3 (delay-tolerant traffic), with CoS1 being the highest priority and CoS3 the lowest. Requests arrived at the network following a Poisson process, with holding times exponentially distributed with a mean of 2 time units (TU). Bandwidth demands were evenly divided among 10 Gbps, 20 Gbps, 40 Gbps, 80 Gbps, 160 Gbps, 200 Gbps, and 400 Gbps. At the start of the simulation, the network had no occupied slots. The offered load in Erlangs was defined by the product of the arrival rate of requests and their average holding time. As new requests arrived, they were allocated accordingly Santos *et al.* [2025].

The simulations used the USA network topology shown in Figure 1a, and the RNP topology (Brazilian National Research Network), shown in Figure 1b.

The USA network has 24 nodes and 43 edges, while the RNP Ipê network has 28 nodes and 50 edges. The distances

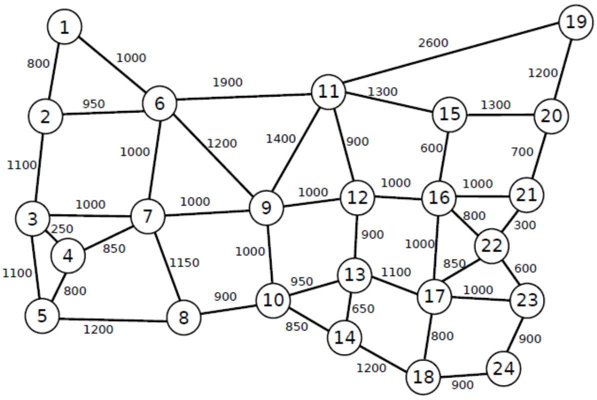

(a) USA topology.

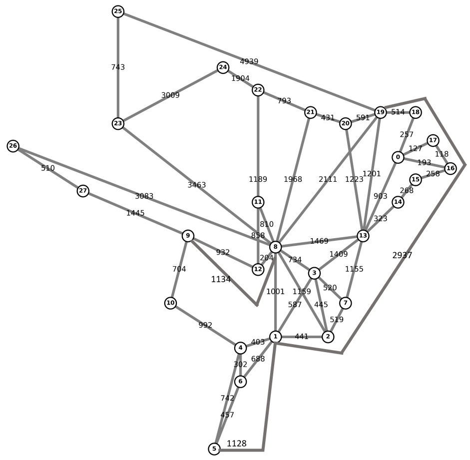

(b) RNP topology.

Figure 1. Network topologies used in the simulations: (a) USA backbone with 24 nodes and 43 links; (b) RNP Ipê backbone with 28 nodes and 50 links. Link lengths approximate the physical distances between cities (in km).

of the edges (in kilometers) approximate the actual physical distances between cities (nodes). Each link in the network topologies has 300 frequency slots, each with 12.5 GHz.

5.1 Analysis of the Blocking Probability

The performance of the algorithm was evaluated using the Blocking Probability (BP – blocked requests relative to total requests) and the Bandwidth Blocking Probability (BBP – blocked bandwidth relative to total requested bandwidth). Additional metrics included blocking probability per service class, analyzed in detail by bandwidth in Gbps, and revenue loss, which reflects the financial value of blocked requests. Performance was first compared against baseline algorithms, P-KSP, CFCoSP, and TSSCF, in terms of blocking probability, without considering service classes or bandwidth sizes.

Analyzing both circuit blocking probability (BP) and bandwidth blocking probability(BBP) is essential in EON simulations because they capture different aspects of network performance. Circuit blocking measures how often connection requests are rejected, reflecting control-plane efficiency. In contrast, bandwidth blocking quantifies the total amount of requested capacity that goes unserved, offering a more accurate view of resource utilization and spectral efficiency.

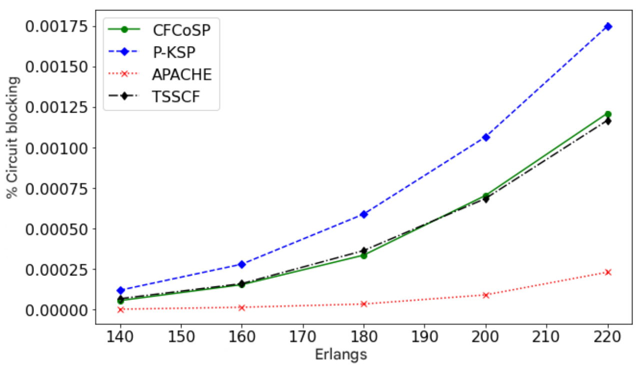

Figure 2. USA Topology.

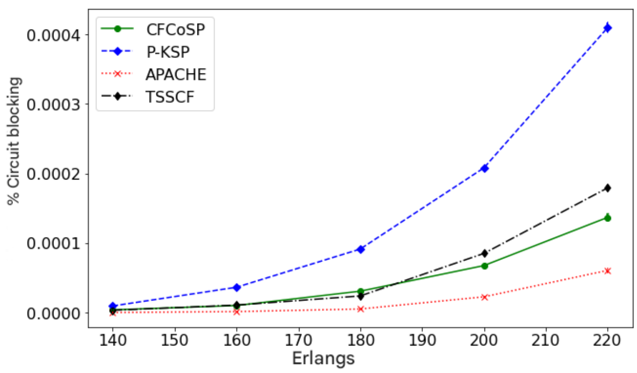

Figure 3. RNP Topology.

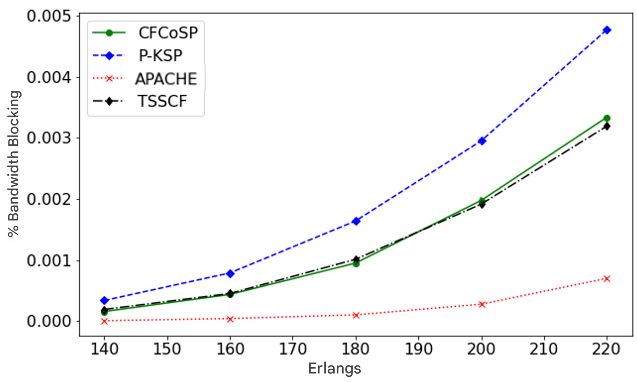

Figure 4. USA Topology.

This distinction is critical for performance evaluation shown in BP Figures 2 and 3, also BBP Figures 4 and 5.

APACHE algorithm clearly outperforms the benchmarks. In the USA topology, APACHE reduces the blocking probability (BP) by up to 45% compared to the best baseline, and in the RNP topology the reduction reaches 38%. This highlights the advantage of considering link criticality and adaptive prioritization in the path selection and spectrum allocation process.

APACHE excels not only in reducing the number of blocked circuits, but also — more importantly — in minimizing bandwidth blocking by intelligently managing spectrum

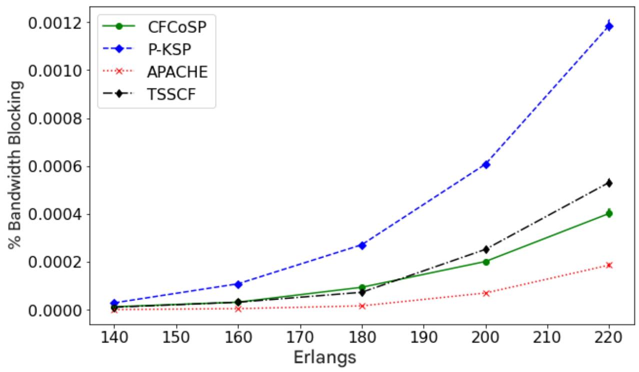

Figure 5. RNP Topology.

fragmentation and allocation. Since large high-bandwidth requests have a greater impact on network performance, an algorithm that prioritizes bandwidth efficiency like APACHE demonstrates superior scalability and robustness under tested load. BBP measures are shown in as shown in Figure 4 regarding USA topology, and Figure 5 over RNP topology.

Thus, reporting both metrics provides a complete, scientifically rigorous assessment of an algorithm's real-world effectiveness in EONs.

5.2 Load Levels

Three load levels were defined for the analyses. Below 140 Erlangs is the Low Load range. Between 140 and 220 Erlangs is the Medium Load range, which is used in most of the results. Above 220 Erlangs is considered the High Load range, where blockages appear more frequently and the different effects of Fragmentation and the Absence of Free Spectrum can be observed. The analysis is extended by increasing the network load, this time up to 500 Erlangs, to verify the response behavior of the algorithms in this high load situation.

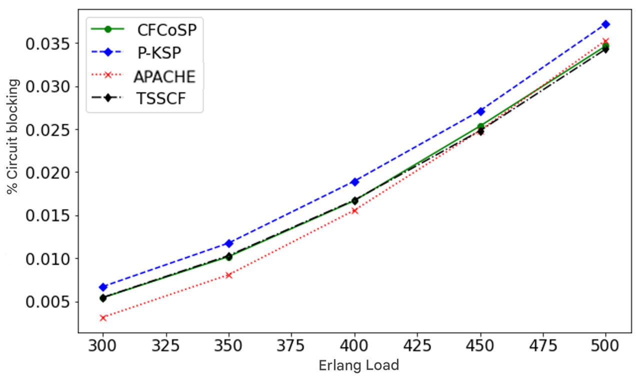

Figure 6. USA high load Topology.

When higher load volumes are applied to RNP, APACHE algorithm underperforms CFCoSP and TSCF starting at 400 Erlangs, as can be seen in Figure 7.

Since APACHE algorithm is Aware of Service Classes, we begin the analysis regarding these service modalities.

5.3 Analysis of the QoS Classes

Working with these Service Classes serves to prioritize certain types of applications preferred by the client, which would yield higher revenue values, considering, for example, the

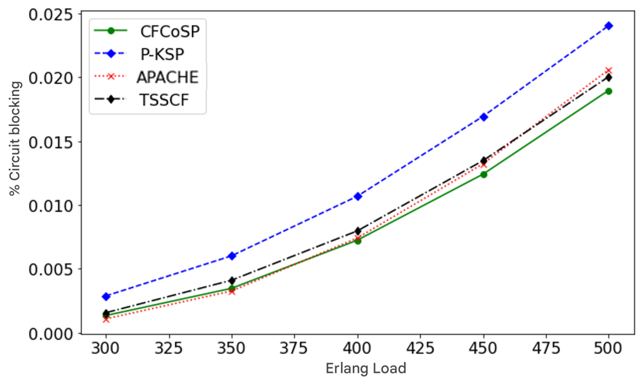

Figure 7. RNP high load Topology.

need for real-time transmission in CoS1, non-critical time transmission in CoS2, and delay-tolerant traffic in CoS3. The simulator allows varying the percentage of each *CoS*, as the provider can charge different values to increase their profit. In this graph, we consider a uniform class ratio *CoS*1:*CoS*2:*CoS*3 (1:1:1) to compare the results in the same way as they were handled in the other algorithms.

Figures 8 and 9 show the Blocking Probabilities of Bandwidths in Gbps, differentiating their Classes among all algorithms in both topologies.

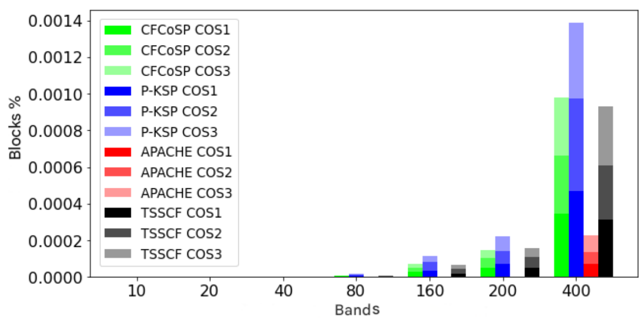

Figure 8. BP 220 Erlangs USA.

APACHE algorithm presents the lowest blocking probability in all three Service Classes, across all request bandwidths, in both topologies, with significant visually observable differences compared to the other algorithms, as shown in Figures 8 and 9 for a 220 Erlang load.

Figures 8 and 9 show the Bandwidth Blocking Probability in Gbps, differentiating their Classes among all algorithms in both topologies.

APACHE algorithm exhibits the lowest blocking probability across all three Service Classes, across all request bandwidths, in both topologies, with significant visually observable differences compared to the other algorithms, as shown in Figures 8 and 9 for a load of 220 Erlangs.

Next, the analysis considers the effects of different request bandwidths using high load.

Over high load, APACHE algorithm presents the lowest probability of circuit blocking in the USA topology up to a load of 450 Erlangs, at which point its performance ceases to have an advantage over the CFCoSP algorithm. This happens because the maximum free block strategy, under very high load conditions, ceases to be a differentiating factor when compared to other algorithms.

Once again, APACHE algorithm shows the lowest probability of blocking in Gbps bandwidths in the USA and RNP

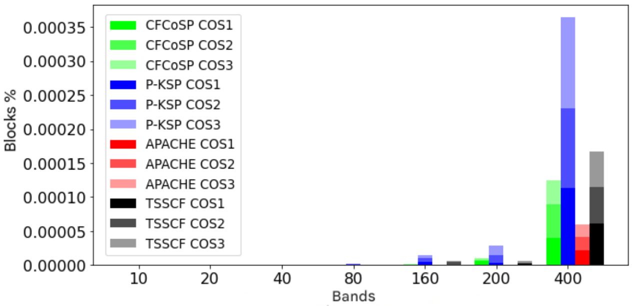

Figure 9. BP 220 Erlangs RNP.

topologies, with a smaller but still significant advantage compared to the others in Figures 10 and 11 for a load of 350 Erlangs. It can be seen that at 400 Gbps requests, the performance of APACHE algorithm is almost equivalent to CFCoSP.

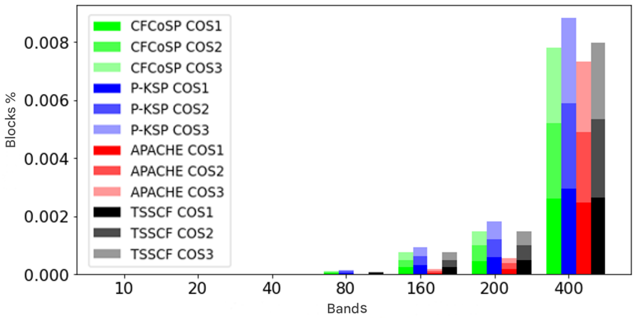

Figure 10. BP 350 Erlangs USA.

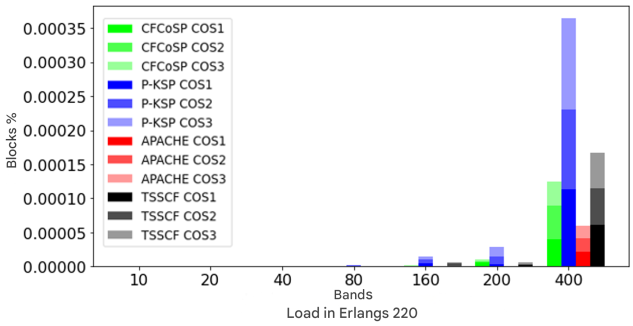

Figure 11. BP 350 Erlangs RNP.

5.4 Analysis of the Network Revenue

Then the economic analysis begins by applying prices to the circuits. Financial values equal to the bandwidths were considered, i.e., a linear relationship, where 10 Gbps bandwidth costs U$10 up to 400 Gbps bandwidth which costs U$400.

APACHE algorithm exhibits the lowest revenue loss across all bandwidths, in both topologies, with a significant advantage compared to other algorithms, as shown in Figures 12 and 13 for all requests with any bandwidth in Gbps.

Considering various fixed loads, we will now analyze the lost revenue generated by Class, comparing all algorithms across all Classes in Figures 14 and 15. APACHE algorithm presents the lowest lost revenue in both topologies and all classes.

The lines representing the percentage of bandwidth blocked are plotted on the same graph.

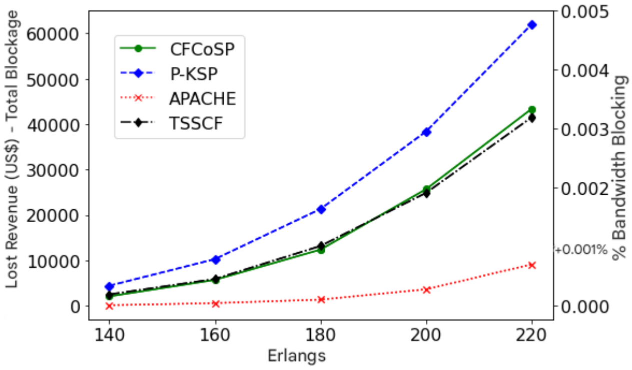

Figure 12. Lost revenue due to USA total band blocking.

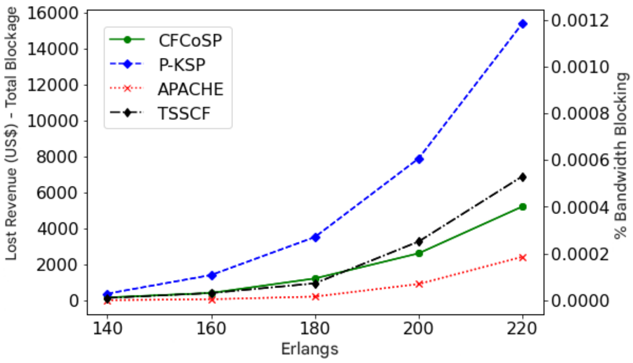

Figure 13. Lost revenue due to RNP full bandwidth blocking.

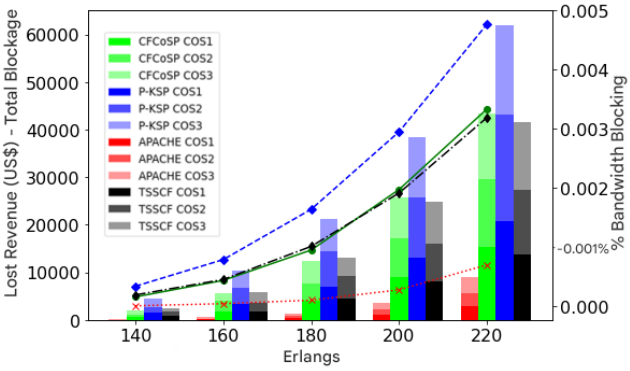

Figure 14. Revenue loss of classes in USA topology.

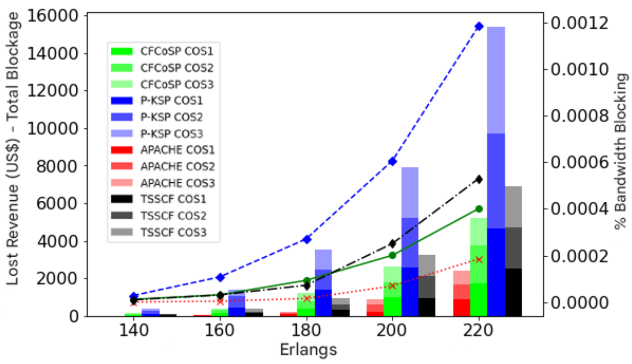

Figure 15. Revenue loss of classes in the RNP topology.

Expanding on the revenue analysis, the potential of U$200,000 as a financial value to be received for the total requests made, reflects the financial impact of unallocated bandwidth across service classes 1, 2, and 3. This loss is directly tied to bandwidth blocking, which reveals not just how many requests were rejected, but how much revenue-generating capacity was wasted—particularly for high-bandwidth services. Figures 16 and 17 show that even with moderate network load, significant revenue is lost due to inefficient.

Once again, we evaluated metrics under increased demand to assess the impact on financial outcomes at higher loads. As traffic intensifies, blocking rises—especially for high-bandwidth requests—leading to greater revenue loss.

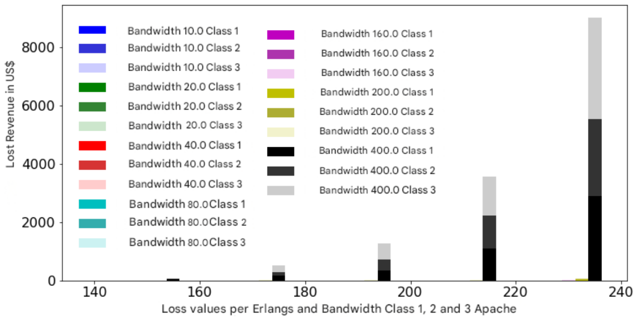

Figure 16. Revenue loss opened by bands and classes — USA.

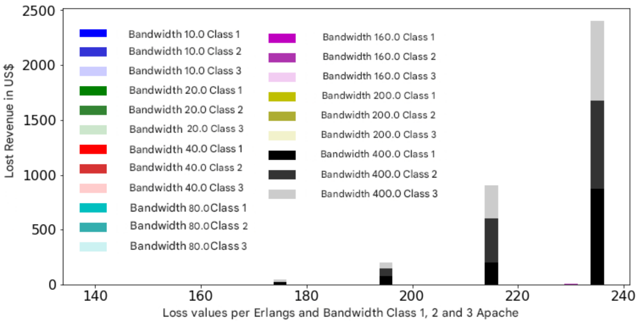

Figure 17. Revenue lost due to a breakdown by band and classes — RNP.

When the load is increased to 500 Erlangs in Figures 18 and 19, , the focus shifts to the APACHE algorithm, which demonstrates superior performance under stress. Despite its efficiency, a clear concentration of blocking—and thus revenue loss—occurs at 400 Gbps requests, with minor impacts on 200 and 160 Gbps. This highlights a critical insight: high-capacity services, while fewer in number, contribute disproportionately to revenue loss when blocked.

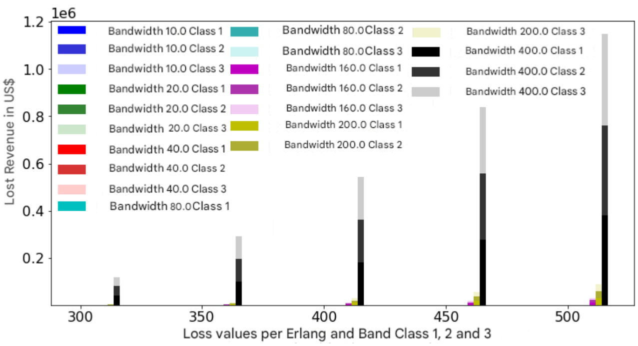

Figure 18. Revenue loss opened by band and by classes — USA high load.

Therefore, evaluating both blocking probability and financial impact provides a more comprehensive view of network performance. APACHE's ability to minimize overall blocking translates into substantial cost savings, as shown by its strong performance in reducing normalized revenue loss. This

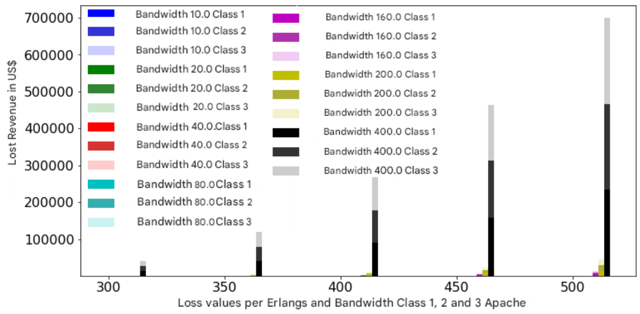

Figure 19. Revenue loss opened by band and by classes — RNP high load.

makes APACHE not only technically effective but also economically advantageous in real-world EON deployments.

6 Conclusions and Future Work

This paper presented APACHE algorithm, a novel approach for resource allocation in Elastic Optical Networks that incorporates link criticality in the path selection and spectrum assignment process. Simulation results demonstrated that APACHE achieves superior performance compared to state-of-the-art algorithms, significantly reducing circuit and bandwidth blocking probability, minimizing revenue loss, and improving fairness among different Classes of Service.

Future work will focus on extending APACHE to multi-layer and multi-domain scenarios, as well as integrating predictive models based on machine learning to anticipate traffic dynamics and further optimize resource allocation. Another research direction involves the inclusion of energy efficiency metrics, aligning with the growing demand for greener and more sustainable optical infrastructures.

Acknowledgements

This work was partially funded by CNPq under grant number 313623/2023-6, 405852/2025-8, FAPESB INCITE PIE0002/2022, and CAPES.

Declarations

Authors' Contributions

CSGR, GBF and AAAR formalized the problem and provided the solution. CSGR, EOS and ASS played key roles in the design of the experiments, implementation of the simulator, and in the performance evaluation. All authors actively contributed to the analysis and interpretation of the results, as well as to the writing and revision of the manuscript. The final version reflects the collective effort and validation of all authors, who read and approved the manuscript

Competing interests

The authors declare that they do not have competing interests.

Availability of data and materials

Not applicable.

References

Batham, D., Pathak, S. K., Yadav, D. S., and Prakash, S. (2020). A traffic scheduling strategy based on cost function for differentiated class of service in multi-domain optical networks. *Optical Fiber Technology*, 60:102337. DOI: 10.1016/j.yofte.2020.102337.

Batham, D. and Yadav, D. S. (2020a). Hpdst: Holding pathlength domain scheduled traffic strategy for multi-domain optical networks. *Optik*, 222:165145. DOI: 10.1016/j.ijleo.2020.165145.

Batham, D. and Yadav, D. S. (2020b). A weight function rsa strategy based on path length and bandwidth demand for static traffic in elastic optical network. In Pandit, M., Srivastava, L., Venkata Rao, R., and Bansal, J. C., editors, *Intelligent Computing Applications for Sustainable Real-World Systems*, pages 88–96, Cham. Springer International Publishing. DOI: 10.1007/978-3-030-44758-8_9.

Chatterjee, B. C., Sarma, N., and Oki, E. (2015). Routing and spectrum allocation in elastic optical networks: A tutorial. *IEEE Communications Surveys & Tutorials*, 17(3):1776–1800. DOI: 10.1109/COMST.2015.2431731.

Dixit, S., Batham, D., and Narwaria, R. P. (2020). Cost function-based class of service provisioning strategy in elastic optical networks. *International Journal of Communication Systems*, 33(18):e4634. e4634 IJCS-20-0735.R1. DOI: 10.1002/dac.4634.

Gerstel, O., Jinno, M., Lord, A., and Yoo, S. J. B. (2012). Elastic optical networking: A new dawn for the optical layer? *IEEE Communications Magazine*, 50(2):s12–s20. DOI: 10.1109/MCOM.2012.6146481.

Horota, A., Reis, L., Figueiredo, G., and da Fonseca, N. L. S. (2015). Routing and spectrum assignment algorithm with most fragmented path first in elastic optical networks. *2015 7th IEEE Latin-American Conference on Communications (LATINCOM)*, pages 1–6. DOI: 10.1109/LATINCOM.2015.7430114.

Jinno, M., Kozicki, B., Takara, H., *et al*. (2009). Spectrum-efficient and scalable elastic optical path network. *IEEE Communications Magazine*, 47(11):66–73. DOI: 10.1109/MCOM.2009.5307471.

Júnior, P. J. d. S., Rodrigues, L., and Marotta, M. (2025). Spectrum defragmentation window in sdm-eon networks. *Journal of Internet Services and Applications*, 16(1):492–507. DOI: 10.5753/jisa.2025.5474.

Lourenço, R. B. R., Tornatore, M., Martel, C. U., and Mukherjee, B. (2018). Running the network harder: Connection provisioning under resource crunch. *IEEE Transactions on Network and Service Management*, 15(4):1615–1629. DOI: 10.1109/TNSM.2018.2875103.

Luo, Y. and Zhu, Z. (2021). Service differentiation and qos-aware provisioning in elastic optical networks. *IEEE/OSA Journal of Optical Communications and Networking*, 13(8):D67–D79. DOI: 10.1364/JOCN.427323.

Oliveira, R. A. and Oliveira, H. M. N. d. S. (2025). Multi-criteria, crosstalk-sensitive flexible topology approach for routing in sdm-eons. *Journal of Internet Services and Applications*, 16(1):54–68. DOI: 10.5753/jisa.2025.5041.

Rodrigues, E. A. F. and Oliveira, H. M. N. d. S. (2025). Request handling in elastic optical data center networks: A routing algorithm approach. *Journal of Internet Services and Applications*, 16(1):131–141. DOI: 10.5753/jisa.2025.5048.

Santos, A. S., de Santi, J., Figueiredo, G. B., and Mukherjee, B. (2022). Application-aware service degradation in elastic optical networks. *IEEE Transactions on Network and Service Management*, 19(2):949–961. DOI: 10.1109/TNSM.2022.3154331.

Santos, A. S., Horota, A. K., Zhong, Z., De Santi, J., Figueiredo, G. B., Tornatore, M., and Mukherjee, B. (2018). An online strategy for service degradation with proportional qos in elastic optical networks. In *2018 IEEE International Conference on Communications (ICC)*, pages 1–6. DOI: 10.1109/ICC.2018.8422781.

Santos, A. S., Santos, E. O., Rahman, S., Wosinska, L., de Santi, J., and Figueiredo, G. B. (2024a). Multi-criteria decision approach for lightpath restoration after resource crunch. *IEEE Transactions on Network and Service Management*, 21(5):5521–5531. DOI: 10.1109/TNSM.2024.3435544.

Santos, E., Richa, C., Santos, A., Santi, J., Rocha, A., and Figueiredo, G. (2024b). f-sim eon: Avaliação de desempenho em redes Ópticas elásticas sob falhas em cascata. In *Anais Estendidos do XLII Simpósio Brasileiro de Redes de Computadores e Sistemas Distribuídos*, pages 65–72, Porto Alegre, RS, Brasil. SBC. DOI: 10.5753/sbrc_estendido.2024.3340.

Santos, E. O., Ferreira, L. F., de Santi, J., Zou, R., Rahman, S., Subramaniam, S., and Figueiredo, G. B. (2025). Post-disaster restoration based on path stability in elastic optical networks under cascading failures. In *Proceedings of the IEEE International Workshop on Future Networks and Wireless Facilities (FNWF)*. DOI: 10.1109/FNWF66845.2025.11317565.

Shi, Y., Zong, Z., and Zhu, Z. (2014). Efficient spectrum defragmentation with holding-time awareness in elastic optical networks. *IEEE/OSA Journal of Optical Communications and Networking*, 6(11):997–1009. DOI: 10.1364/JOCN.6.000997.

Wang, Y., Cao, X., and Pan, Y. (2011). A study of the routing and spectrum allocation in spectrum-sliced elastic optical path networks. *2011 Proceedings IEEE INFOCOM*, pages 1503–1511. DOI: 10.1109/INFCOM.2011.5934939.