Les Cahiers du GERAD ISSN: 0711–2440

Cross-layer consolidation and defragmentation in OTN-over-EON networks

E. Rezagholizadeh, E. Amiri, S. Kohlert, R. Kashyap, B. Sansò

G–2026–28

June 2026

GERAD HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7

Tél. : 514 340-6053
Téléc. : 514 340-5665
info@gerad.ca
www.gerad.ca

Cross-layer consolidation and defragmentation in OTN-over-EON networks

Ehsan Rezagholizadeh [a, b]

Elahe Amiri [b, c]

Scott Kohlert [d]

Raman Kashyap [a, e]

Brunilde Sansò [a, b]

[a] *Department of Electrical and Computer Engineering, Polytechnique Montréal, Montréal (Qc), Canada, H3T 1J4*

[b] *GERAD, Montréal (Qc), Canada, H3T 1J4*

[c] *Department of Mathematics and Industrial Engineering, Polytechnique Montréal, Montréal (Qc), Canada, H3T 1J4*

[d] *Ciena Corporation, Ottawa (On), Canada, K2K 0L1*

[e] *Department of Engineering Physics and Department of Electrical Engineering, Polytechnique Montréal, Montréal (Qc), Canada, H3T 1J4*

ehsan.rezagholizadeh@polymtl.ca

June 2026
Les Cahiers du GERAD
G–2026–28

Abstract : Optical transport networks (OTNs) carry Ethernet traffic over lightpaths that can be provisioned using elastic optical networks (EONs) with flexible spectrum allocation. Under dynamic traffic, frequent client arrivals and departures create scattered free capacity across parallel lightpaths in the transport layer and spectrum fragmentation in the optical layer, leading to underutilized lightpath capacity and spectrum gaps. To quantify and mitigate the transport-layer inefficiency, we introduce a metric called *lightpath utilization entropy* and propose a corresponding *OTN consolidation* strategy that rearranges existing clients across lightpaths serving the same O/D pair to consolidate free capacity. Because this operation is confined to the electrical domain, it can be executed quickly and with minimal service disruption. To jointly address both impairments, we further develop a cross-layer framework that simultaneously consolidates client allocation in the transport layer and reconfigures lightpath placement in the optical layer. We formulate the resulting cross-layer consolidation and defragmentation (CLCD) problem as a cutting-stock model and solve it efficiently by column generation. Because optical-layer reconfiguration is slower, CLCD is executed less frequently than OTN consolidation and is complemented by the latter at a higher frequency. Simulation results demonstrate that the proposed framework, combining CLCD with OTN consolidation, reduces the blocking ratio by up to 51% under medium traffic loads and outperforms single-layer (EON-only) defragmentation in both performance and resource utilization.

Keywords : Bandwidth fragmentation; defragmentation algorithms; elastic optical networks; optical transport networks; cross-layer consolidation and defragmentation

Acknowledgements: This work was supported by Mitacs and Ciena Corporation.

1 Introduction

Transport networks have evolved significantly in recent years due to the growing demand for high-speed data transmission and flexible resource allocation. Elastic Optical Networks (EONs), with their ability to adjust channel bandwidth dynamically, have emerged as a critical solution for optimizing spectral resources. Unlike traditional Wavelength Division Multiplexing (WDM) networks that rely on fixed wavelength grids, EONs enhance spectral efficiency by enabling Bandwidth-Variable Transponders (BVTs) to provision just enough bandwidth for each connection request based on traffic demands.

Parallel to this, Optical Transport Network (OTN) technology is widely used to encapsulate Ethernet services and groom multiple client flows into high-capacity optical channels, thereby enabling an OTN-over-EON architecture to jointly leverage OTN grooming and EON spectrum flexibility to accommodate heterogeneous Ethernet demands in a scalable manner. Under this architecture, when a new Ethernet demand arrives, the network reuses established lightpaths whenever possible and sets up new optical lightpaths only when additional capacity is required.

However, the dynamic provisioning that makes OTN-over-EON attractive under fluctuating traffic also introduces new challenges over time. As traffic continuously arrives and departs, the accumulation of provisioning and teardown connections produces two distinct forms of inefficiency in the OTN and EON layers. In the OTN layer, lightpaths often become partially filled, so that free capacity is scattered across many parallel lightpaths of the same O/D pair as small residual chunks, with no single lightpath holding enough to accept a large incoming demand. As a result, a new client can be blocked even when the aggregate free capacity is sufficient. In the EON layer, by contrast, when all clients on a lightpath depart, tearing down the lightpath leaves spectrum gaps. This effect is exacerbated in spectrum-flexible operation, where lightpaths may be set up, released, expanded, or contracted using BVTs. Due to spectrum continuity and contiguity constraints, these gaps make it difficult to find usable spectrum blocks for establishing new lightpaths, especially for large-bandwidth demands. Taken together, OTN-layer capacity dispersion and EON-layer spectrum fragmentation degrade overall network efficiency, increase the blocking probability, and complicate resource management.

Several approaches have been proposed to mitigate fragmentation in the EON layer. Preventive techniques carefully place new connections to limit fragmentation. For instance, spectrum partitioning has been investigated to steer connection requests with different bandwidth requirements to dedicated spectral regions [2, 10, 18]. Fragmentation-aware Routing and Spectrum Assignment (FA-RSA) algorithms proactively select routes and frequency slots during lightpath establishment to reduce future fragmentation [6, 13, 20, 21]. These methods rely on quantitative fragmentation indicators, such as a fragmentation index [18] or Shannon-entropy-based measures [20, 21], to guide resource allocation.

Reactive defragmentation techniques reorganize existing connections in response to network conditions, employing strategies like spectrum shifting, make-before-break approaches, or iterative adjustment of established lightpaths [4, 14–16, 19, 22]. These methods are triggered either by lightpath blockage caused by spectrum fragmentation or when fragmentation exceeds a threshold. More recently, learning-based approaches have been investigated to make EON defragmentation proactive and self-adaptive. For example, Etezadi et al. [7] proposed *DeepDefrag*, a deep reinforcement learning framework that jointly decides *when* to trigger defragmentation, *which* connections to reconfigure, and *where* to reallocate them in the spectrum, by considering service attributes, multiple fragmentation metrics, and the reconfiguration cost. Although these methods substantially improve EON-layer defragmentation, they all operate strictly at the optical layer and do not address the payload within lightpaths or the associated transport-layer packing inefficiencies.

Concerning multilayer modeling, earlier studies (e.g., [12]) considered provisioning for IP over GMPLS networks, highlighting the benefits of coordinated resource management across layers. Although fragmentation mitigation is a key motivation for multilayer coordination, not all multilayer reconfiguration frameworks explicitly optimize fragmentation. In the context of multilayer optimization, Tanaka et al. [17] proposed a multiperiod IP-over-EON reconfiguration framework that iteratively updates the

IP virtual topology and then performs RSA in the optical layer. Their heuristic further incorporates adaptive bandwidth resizing together with distance- and bitrate-adaptive modulation to improve energy efficiency under periodically changing traffic matrices. However, the framework primarily targets energy-aware reconfiguration rather than fragmentation-driven optimization.

In [23], a cross-layer spectrum defragmentation approach for IP over EON networks is proposed, in which both IP traffic flows and established lightpaths are reconfigured to alleviate fragmentation, with defragmentation triggered on blocking events. More recently, Liu et al. [11] proposed an adaptive IP/optical cross-layer bandwidth defragmentation algorithm for multi-band optical networks, combining deep reinforcement learning with service clustering and aggregation. Although these studies capture interactions between IP and optical layers, none of the cross-layer or multilayer works discussed above addresses OTN grooming or the transport-layer packing inefficiencies specific to OTN-over-EON architectures.

In this paper, we first introduce a metric called *lightpath utilization entropy* to quantify the transport-layer inefficiency caused by free capacity being scattered across many partially-filled parallel lightpaths. To mitigate this inefficiency, we propose an *OTN consolidation* strategy that repacks Ethernet demands across lightpaths between the same O/D pair, operating entirely in the electrical domain so that it can be executed quickly and with minimal service disruption. We formulate the OTN consolidation problem as a mixed-integer program and, due to its computational complexity at scale, also develop an efficient heuristic to enable frequent online application.

Because OTN consolidation cannot address spectrum fragmentation in the optical layer, we extend our approach to a *cross-layer consolidation and defragmentation* (CLCD) model that simultaneously repacks Ethernet demands at the OTN layer and reconfigures lightpath placement in the EON layer. We first formulate CLCD as a mixed-integer program and then develop a scalable column-generation solution. Since modifying optical lightpaths is costly, CLCD is intended to run less frequently than OTN consolidation, and the combination of the two provides an effective balance between performance and operational overhead.

Our approach distinguishes itself from existing literature through the following key contributions:

Introduction of lightpath utilization entropy: We introduce a novel metric that penalizes half-filled lightpaths and rewards configurations in which each lightpath is either full or empty.

OTN consolidation strategy: We propose both a mixed-integer programming formulation and a heuristic algorithm to reduce lightpath utilization entropy.

Cross-layer consolidation and defragmentation model: We jointly consolidate client allocation at the OTN layer and defragment the optical spectrum at the EON layer within a single optimization model.

Column generation approach: We develop a column-generation-based approach for efficiently solving the CLCD problem at large scale.

Performance evaluation: We demonstrate the effectiveness of our approaches through extensive simulations, showing clear improvements relative to a well-established EON-only baseline.

The reader should be aware that we deliberately omit survivability considerations. Incorporating working-backup path constraints would require enforcing disjointness or other operator-specific policies whose interaction with defragmentation is beyond the scope of this work. To isolate the efficiency gains of OTN consolidation and CLCD, all demands are therefore treated as unprotected working connections; extending the framework to support protected services is left to future work.

The rest of the paper is organized as follows. Section 2 introduces the OTN and EON structures and the multilayer network architecture. Section 3 defines the lightpath utilization entropy metric and presents a heuristic for OTN consolidation. Section 4 presents the CLCD MIP formulation and the column-generation technique used to solve it. Section 5 reports simulation results, and Section 6 concludes.

Table 1: Acronyms and Full Forms

Acronym	Full Form
OTN	Optical Transport Network
EON	Elastic Optical Network
WDM	Wavelength Division Multiplexing
MIP	Mixed-Integer Programming
CAG	Collapsed Auxiliary Graph
CLCD	Cross-Layer Consolidation and Defragmentation
BVT	Bandwidth-Variable Transponder
BV-WSS	Bandwidth-Variable Wavelength Selective Switches
FA-RSA	Fragmentation-Aware Routing and Spectrum Assignment
ODU	Optical Data Unit
WSSs	Wavelength Selective Switches
FS	Frequency Slot
O/D	Origin-Destination
OPU	Optical Payload Unit
RMP	Restricted Master Problem

Table 2: Notations and Definitions

Sets	
$\mathcal{D}$	Set of demands between an O/D pair
$\mathcal{H}$	Set of demand types
$\mathcal{P}$	Set of EON-layer paths connecting an O/D pair
$\mathcal{S}_p$	Set of spectrum segments along path p
$\mathcal{L}$	Set of lightpaths connecting an O/D pair
$\mathcal{L}_p$	Set of lightpaths that use path p
$\mathcal{T}$	Set of all operational transponder modes
$\mathcal{T}_p$	Set of transponder modes whose optical reach is sufficient for path p
$\mathcal{A}$	Set of feasible patterns
Parameters / Constants	
W_d	Data-rate requirement of demand d
Q_l	Capacity of lightpath l
M_h	Number of demands of type h
W_h	Bandwidth of demand type h
F_t	Number of FSs required by transponder mode t
O_t	Optical reach of transponder mode t
R_t	Data rate supported by transponder mode t
B_{ps}	Number of FSs available in segment s on path p
G	Threshold to trigger OTN consolidation
$\mathcal{F}_a$	Number of FSs used by pattern a
η	Number of lightpaths between an O/D pair
η_p	Number of lightpaths on path p
φ_{ap}	1 if pattern a uses path p
Decision Variables	
θ_{dl}	1 if demand d uses lightpath l
r_l	Remaining capacity of lightpath l
x_{plts}	1 if lightpath l uses transponder mode t and segment s on path p
y_{dpl}	1 if demand d is on lightpath l of path p
z_h	Number of times demand type h is included in the new pattern
μ_{ha}	Number of times demand type h is included in pattern a
λ_{as}	Number of times pattern a is used in spectrum segment s

2 OTN and EON structure in a multilayer network architecture

This section outlines the structures of EON and OTN, describes the multilayer network architecture they form, and explains how Ethernet clients are mapped into this multilayer framework. EONs divide the optical spectrum into smaller units called Frequency Slots (FSs), typically of 6.25 GHz width. BVTs are integral components of EONs, capable of adjusting their spectral width and modulation formats.

By varying the number of FSs occupied, BVTs provide just enough bandwidth for each lightpath (the end-to-end optical channel) based on the required data rate and optical reach.

OTN technology operates above the EON layer and is responsible for encapsulating client data, such as Ethernet packets, into optical channels. The core of the OTN architecture is the Optical Data Unit (ODU) sublayer, which serves as a signal wrapper. The ODU encapsulates client signals into the Optical Payload Unit (OPU) through a process known as OTN multiplexing or electrical grooming. ODUflex enhances this flexibility by enabling variable-rate ODU containers whose payload is allocated in discrete *tributary slot* granularity (i.e., an integer number of tributary slots, each of 1.25 Gb/s) [9]. Ethernet clients are mapped into ODU tributary slots, with the number of tributary slots allocated corresponding to the client's data rate.

Moreover, OTN switching capabilities allow for the extraction of client signals from an ODU and their insertion back into another ODU [8]. OTN switching is often constrained by the availability of switching equipment at network nodes. The combination of EON and OTN technologies results in a three-layer network architecture:

- **Ethernet Client Layer**: This is the highest layer, representing the Ethernet client demands between Origin-Destination (O/D) pairs. The data rates of these demands follow standard values, such as 10GE, 100GE, and 400GE, which we refer to as Ethernet demands of specific types.
- **OTN Layer**: The middle layer consists of OTN switches and lightpaths. Ethernet clients are encapsulated into ODU tributary slots and transmitted over lightpaths.
- **EON (Physical) Layer**: The lowest layer consists of the physical fiber infrastructure managed by EON technology. Lightpaths in the OTN layer are mapped onto the EON layer and routed over fiber links, taking into account spectrum continuity and contiguity constraints.

In the OTN layer, as the capacity of a single lightpath may be insufficient to accommodate the total traffic volume between an O/D pair, multiple lightpaths may connect the same pair of nodes. However, in OTN technology, Ethernet clients are typically not split across parallel (i.e., same-O/D) lightpaths. Each Ethernet demand is transmitted over a single lightpath or a sequence of interconnected lightpaths via OTN switches. Such splitting is avoided due to challenges such as propagation delays and difficulties in synchronizing data arriving over different paths [9].

In some cases, an Ethernet demand is carried over aggregated lightpaths. This aggregation requires the lightpaths to traverse the same route through the network and operate on a contiguous segment of the spectrum. As a result, a bonded lightpath can be treated as a single 'lightpath,' preserving the validity of the initial assumptions.

Fig. 1a shows the three-layer architecture used in this paper. The bottom layer is the EON, whose fiber links provide FSs. The middle layer is the OTN layer, where lightpaths are established over the EON. The top layer represents the Ethernet demand layer. Fig. 1b zooms into the OTN layer and illustrates how three Ethernet demands of 100, 200, and 400 Gb/s from node A to node F are carried over OTN lightpaths. Each small rectangle on a lightpath represents an abstract capacity unit, which may consist of multiple OTN tributary slots, and is used for ease of illustration. Different fill patterns indicate capacity units occupied by different Ethernet demands. For example, the 100 Gb/s demand uses the fourth capacity unit on the lightpaths AB, BE, and one of the parallel EF lightpaths. Two parallel lightpaths are installed between nodes E and F to carry the higher traffic volume for this O/D pair. In addition, lightpath BE has seven available capacity units, whereas lightpath AC has only four, so BE can multiplex more subwavelength traffic.

Fig. 1c moves down one layer and shows how a single OTN lightpath, AB, is implemented in the EON. The three vertical bars represent the FS spectrum on the physical fiber links AC, CE, and EB that form the route of lightpath AB. The first three FSs on each fiber, drawn with the same hatch pattern, are reserved for lightpath AB. The remaining FSs, shown with different fill patterns, are used

by other lightpaths such as AC, CE, and BE, while the unshaded segments labelled Free correspond to unused FSs.

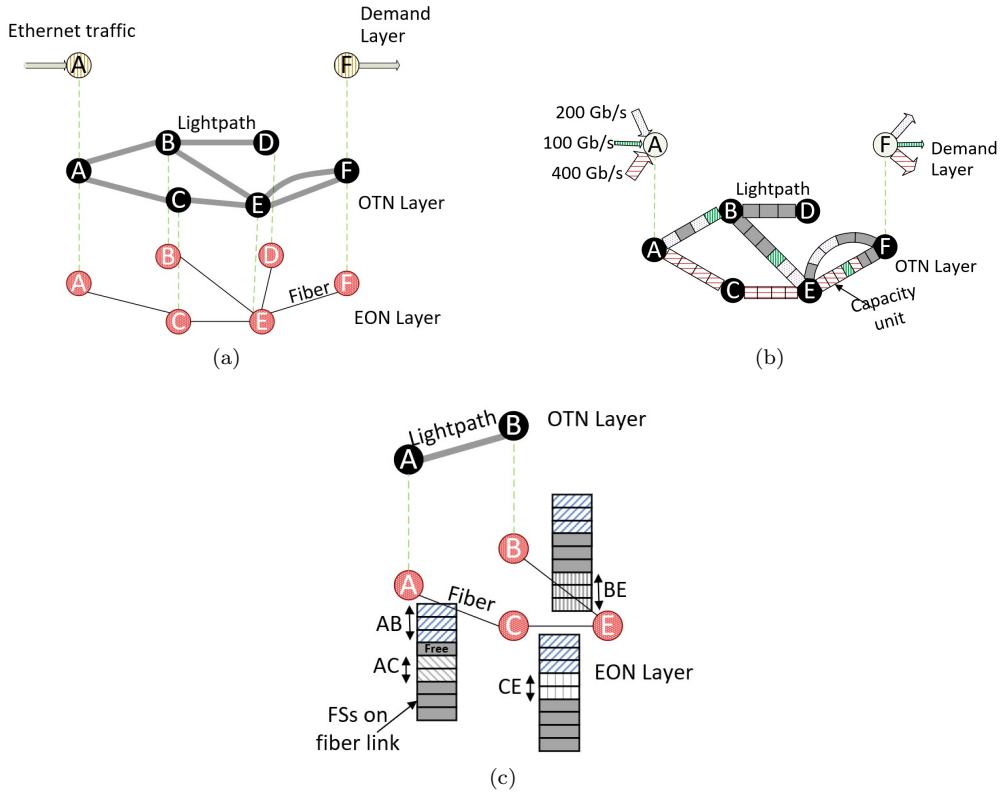

Figure 1: Multilayer network architecture. (a) Three-layer view with Ethernet, OTN, and EON layers. (b) Encapsulation of Ethernet demands into capacity units across multiple lightpaths. (c) Mapping of OTN lightpaths onto EON-layer FSs along the underlying fiber links.

3 OTN consolidation

3.1 Lightpath utilization entropy

In OTNs, multiple lightpaths may be established between the same pair of nodes to support high traffic volumes. Although this increases aggregate capacity, inefficiency arises when the free capacity is scattered across many partially-filled lightpaths, so that no single lightpath has enough room to accept a large incoming demand. In such cases, a new demand may be blocked even though the total free capacity is sufficient. To illustrate this effect, consider the two lightpaths between nodes E and F shown in Fig. 1b. Although several capacity units remain unused across the two lightpaths, these free capacity units are split between them. If a new 400 Gb/s Ethernet demand arrives, requiring four capacity units on a *single* lightpath, neither lightpath has four free capacity units available. As a result, the demand is blocked despite the fact that the two lightpaths together have more than four free capacity units in total.

This scenario highlights the importance of efficient demand packing across the lightpaths serving the same node pair. Ideally, each lightpath should be either fully utilized (i.e., carrying as much traffic as possible) or completely empty (i.e., available for release or for a new large demand), since intermediate fill levels both waste capacity and block large demands. To characterize the inefficiency caused by many simultaneously half-filled lightpaths, we introduce the *lightpath utilization entropy*, which measures how far, on average, the lightpaths of an O/D pair are from these two preferred states.

To formally define lightpath utilization entropy, consider a pair of nodes (u, v) connected by a set of lightpaths $\mathcal{L}$. Let η be the number of these lightpaths. For each lightpath $l \in \mathcal{L}$, let Q_l denote its total capacity and r_l its remaining (unused) capacity. The lightpath utilization entropy f is defined as the average binary entropy of the utilization ratios of these lightpaths:

$$f = \frac{1}{\eta} \sum_{l \in \mathcal{L}} H\left(\frac{r_l}{Q_l}\right) \tag{1}$$

Here, $H(\cdot)$ is the binary entropy function given by:

$$H(\chi) = -\chi \log_2(\chi) - (1-\chi)\log_2(1-\chi) \tag{2}$$

The binary entropy function is maximized at $\chi = 0.5$ (a half-filled lightpath, which is the worst case for both accepting a new large demand and for being released), and drops to zero as $\chi \to 0$ (an empty, releasable lightpath) or $\chi \to 1$ (a fully packed one); both extremes are desirable. The metric f is therefore *minimized* when each lightpath of the O/D pair is individually pushed toward one of the two extremes.

3.2 Optimization model

To reduce lightpath utilization entropy and improve network efficiency, we propose a process called *OTN consolidation*: existing Ethernet demands are reassigned across the lightpaths serving the same node pair so that free capacity becomes concentrated on as few lightpaths as possible. The operation is confined to the OTN layer; it does not modify lightpath configurations, physical routes, allocated spectrum, or modulation formats.

Fig. 2 illustrates OTN consolidation using a bins-and-items analogy. Each bin corresponds to a lightpath between the same node pair, with fixed total capacity Q_l and remaining free capacity r_l shown below each bin. The colored items represent already established Ethernet demands (red: 100 Gb/s; blue: 10 Gb/s; green: 400 Gb/s). In Fig. 2a, all four bins are partially filled: the free capacity is spread across them, so no single bin has 400 Gb/s available even though the aggregate residual capacity is sufficient. This configuration yields a high utilization entropy, $f = 0.78$. Fig. 2b shows the configuration after OTN consolidation: the same demands are repacked so that three lightpaths become fully utilized and one becomes empty, giving $(r_1, r_2, r_3, r_4) = (0, 0, 70, 600)$ Gb/s and reducing the utilization entropy to $f = 0.13$. The empty lightpath ($r_4 = 600$ Gb/s) can now accommodate the new 400 Gb/s demand. We model OTN consolidation as the following optimization problem that minimizes f:

$$f^* = \min \frac{1}{\eta} \sum_{l \in \mathcal{L}} H\left(\frac{r_l}{Q_l}\right) \tag{3a}$$

$$s.t. \quad \sum_{l \in \mathcal{L}} \theta_{dl} = 1 \qquad \forall\, d \in \mathcal{D} \tag{3b}$$

$$r_l = Q_l - \sum_{d \in \mathcal{D}} \theta_{dl} W_d \qquad \forall\, l \in \mathcal{L} \tag{3c}$$

$$\theta_{dl} \in \{0, 1\} \qquad \forall\, d \in \mathcal{D}, \forall\, l \in \mathcal{L} \tag{3d}$$

$$r_l \geq 0 \qquad \forall\, l \in \mathcal{L} \tag{3e}$$

where:

$$\mathcal{D} : \text{Set of Ethernet demands between the node pair.} \tag{4}$$

$$W_d : \text{Data rate required by demand } d \in \mathcal{D}. \tag{5}$$

$$\theta_{dl} : \text{Binary variable indicating if demand } d \text{ is assigned to lightpath } l. \tag{6}$$

Constraint (3b) ensures each demand is assigned to exactly one lightpath. Constraint (3c) calculates the remaining capacity on each lightpath after assignment. The objective function (3a) is nonlinear due to the entropy term but can be linearized using piecewise approximations, allowing the use of MIP solvers. However, the exact solution of (3) may become computationally intractable for large network instances, motivating the use of more scalable solution techniques.

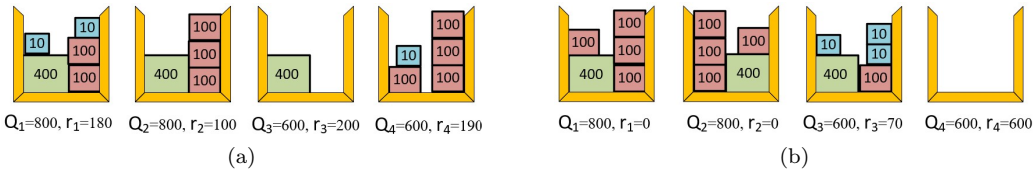

Figure 2: Illustration of OTN consolidation across four parallel lightpaths (bins). (a) Before consolidation: free capacity is spread over several partially-filled lightpaths, giving a high utilization entropy. (b) After consolidation: demands are repacked so that most lightpaths become full and one becomes empty, driving the utilization entropy down.

3.3 Heuristic approach

To address scalability in practical deployments, we developed a heuristic algorithm that approximates the optimal solution of the OTN consolidation problem (3). As summarized in Algorithm 1, the procedure iterates through each O/D pair (u, v) in the network. For a given pair, the algorithm first constructs the set $\mathcal{L}$ of lightpaths serving the node pair and the set $\mathcal{D}$ of Ethernet demands routed over those lightpaths, and then computes the initial lightpath utilization entropy f^0. Once these sets are established, the lightpaths are sorted in descending order of capacity to prioritize the largest ones for consolidation. For each lightpath $l \in \mathcal{L}$, an initially empty assignment set A_l is created to store the demands that will ultimately be mapped to it.

Algorithm 1 Heuristic OTN Consolidation Algorithm

Input: Current network status.
Output: Updated Ethernet demand assignment
for each O/D pair (u, v) **do**
  Initialize $\mathcal{D}$ and $\mathcal{L}$
  $f^0 \leftarrow$ Calculate initial lightpath utilization entropy
  Sort $\mathcal{L}$ in descending order of capacity (tie-breaker: increasing start FS)
  $A_l \leftarrow \emptyset, \forall\, l \in \mathcal{L}$: Create initial mapping
  while $\mathcal{L}$ is not empty **do**
    $l \leftarrow$ Select the first lightpath of $\mathcal{L}$
    $A_l \leftarrow$ Solve a knapsack problem to assign demands from $\mathcal{D}$
    $\mathcal{L} \leftarrow \mathcal{L} \setminus \{l\}$
    $\mathcal{D} \leftarrow \mathcal{D} \setminus A_l$
  $f^* \leftarrow$ Calculate new lightpath utilization entropy using mapping A
  if $f^0 - f^* \geq G$ **then**
    Apply the new assignment A
  else
    Discard changes
Return Updated Ethernet demand assignment

The algorithm then processes the lightpaths in the sorted order. For each l, a dynamic-programming-based knapsack procedure selects a subset of the remaining demands in $\mathcal{D}$ that maximizes the total assigned bandwidth without exceeding the capacity of l. The selected demands are added to A_l and removed from $\mathcal{D}$, and the algorithm moves on to the next lightpath in the sorted list. This process continues until all lightpaths in $\mathcal{L}$ have been examined, yielding a complete reassignment of demands across the available lightpaths.

After the reassignment is completed, the algorithm recomputes the lightpath utilization entropy, denoted f^*, for the updated demand-lightpath assignment. The new mapping is accepted only if it

yields a sufficiently large reduction; that is, if $f^0 - f^* \geq G$, where $G \in [0, 1]$ specifies the minimum required improvement. Otherwise, the original assignment is retained.

By introducing the threshold G, we ensure that consolidation is applied only when it yields a significant reduction in lightpath utilization entropy, avoiding unnecessary disruption to active connections. This heuristic provides a balance between optimization quality and computational efficiency, making it suitable for large-scale networks where exact solutions are impractical.

OTN consolidation reorganizes Ethernet demands across existing lightpaths but does not modify the underlying EON spectrum allocation. The next section addresses that limitation.

4 Cross-layer consolidation and defragmentation

CLCD acts on both layers at once by merging or releasing lightpaths serving the same O/D pair, adjusting their spectral widths, and repositioning their spectrum allocations, while concurrently rearranging Ethernet demands on top. Two constraints are considered on these operations. First, the physical routing paths of established lightpaths are assumed to remain fixed, thereby avoiding rerouting operations that would introduce additional latency and cause service disruption through Bandwidth-Variable Wavelength Selective Switch (BV-WSS) adjustments. Second, any modification to a lightpath configuration must comply with the capabilities of the underlying BVTs, including maximum data rate, supported modulation formats, and optical-reach limitations.

To perform these operations, we first identify the spectrum regions that each O/D pair can use. For an O/D pair on a fixed route, an *available spectrum segment* is a continuous group of FSs that is either free or already occupied by a lightpath of the same O/D pair. In this way, the available spectrum segments define the spectrum range within which the O/D pair's lightpaths can extend, merge, or shift.

In Fig. 3a, the spectrum allocation on the links along the shortest path between a node pair is depicted, including four established lightpaths l_1 to l_4. The solid black portions indicate spectrum occupied by lightpaths belonging to other O/D pairs and are therefore unavailable to the current optimization. The gray portions correspond to the spectrum used by the lightpaths of the considered node pair. The available spectrum segments for this pair are denoted by S_1 to S_4.

Due to spectrum fragmentation and partial utilization, these four lightpaths do not use the spectrum efficiently. By strategically expanding lightpaths l_1 and l_2 toward the lower end of the spectrum, the demands currently served by l_3 and l_4 can be migrated onto l_1 and l_2, effectively packing the traffic into fewer lightpaths. Fig. 3b shows the resulting allocation. The extended lightpaths l_1^{ex} and l_2^{ex} now carry the entire demand of the O/D pair; the hashed portions indicate the additional FSs assigned to these lightpaths within segments S_1 and S_2. Lightpaths l_3 and l_4 become redundant and can be released, which frees spectrum in S_3 and S_4 and reduces fragmentation. In this example, the consolidation is enabled by selecting a higher-baud-rate transponder mode to support the increased load on fewer lightpaths and by shifting their spectrum allocation toward lower-index FSs (see Table 3, which lists the supported modes, adapted from the Ciena WaveLogic 5 data sheet [1]).

More generally, merging or adjusting the lightpaths serving an O/D pair reduces the number of guard bands and can replace longer routes with shorter merged ones that consume fewer FSs at the same data rate. As a result, CLCD removes spectral gaps in the EON layer and reduces the number of active transponders. To formalize the procedure, we next describe how the available spectrum segments are computed for each O/D pair.

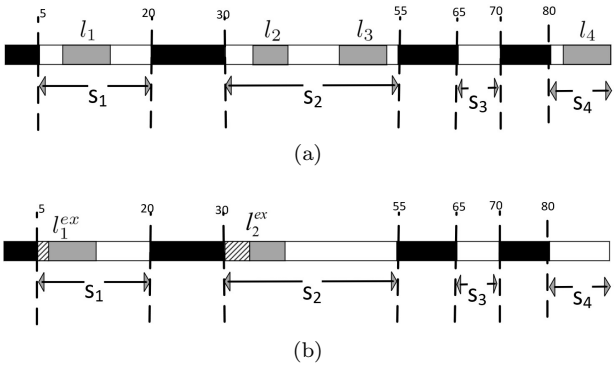

Figure 3: Spectrum optimization through CLCD. (a) Initial configuration showing lightpaths l_1–l_4 and available spectrum segments S_1–S_4. (b) After CLCD: the demands on l_3 and l_4 are migrated onto the expanded lightpaths l_1^{ex} and l_2^{ex}, freeing segments S_3 and S_4.

4.1 Segment calculation

When an O/D pair is routed over a single path, the available spectrum segments are obtained directly from the spectrum occupancy along that path. In the multi-path case, any FSs on shared links are allocated to the shorter path, since such FSs cannot be used by more than one path simultaneously.

Fig. 4 illustrates this for two candidate paths between nodes A and F in the EON layer. The green path is the first shortest path (1-SP), and the red path is the second shortest path (2-SP). In Fig. 4a, the two paths share fiber link AB. On the 1-SP, the free FSs form two contiguous segments, S_1 and S_2. To avoid conflicts on the shared link, any free FS assigned to a segment on the 1-SP is blocked (red crosses) on the 2-SP; the remaining free FSs on the 2-SP are grouped into segments S_3 and S_4.

Fig. 4b shows the contrasting case where the 1-SP and 2-SP are link-disjoint. Since the two paths do not share any fiber link, all free FSs may be used independently on each path. The 1-SP therefore contains segments S_1 and S_2, while the 2-SP contains S_3 and S_4; the same FS indices may appear on both paths without causing spectrum conflicts.

With the available spectrum segments characterized for every path of each O/D pair, we are now ready to formulate the CLCD optimization problem.

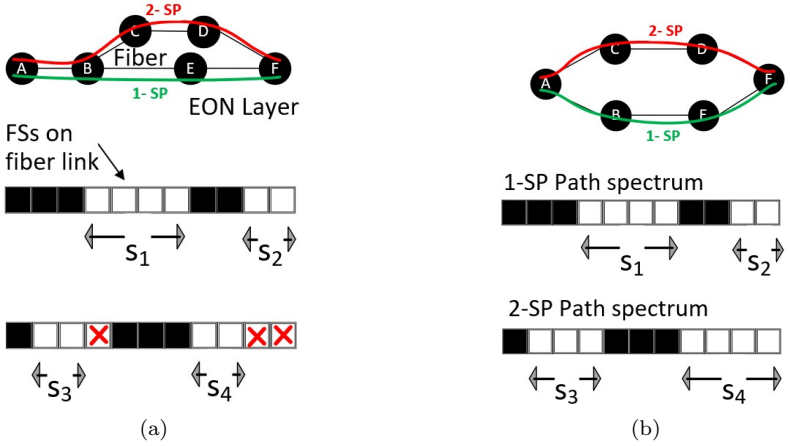

Figure 4: Segment calculation for different path configurations. (a) Paths sharing a fiber link: FSs already assigned to the 1-SP are blocked on the 2-SP (red crosses). (b) Link-disjoint paths: all free FSs can be grouped independently on each path.

4.2 Optimization model

To model the CLCD process, we propose an MIP formulation with the objective of minimizing the total spectrum usage while satisfying all Ethernet demands between the O/D pair. We introduce the following constants and variables to formulate the problem.

$\mathcal{T}$: Set of all operational modes of transponders, where each mode $t \in \mathcal{T}$ is defined by a tuple (F_t, O_t, R_t). (7)

F_t : Number of FSs required for transponder mode t. (8)

O_t : Optical reach of transponder mode t. (9)

R_t : Data rate supported by transponder mode t. (10)

$\mathcal{P}$: Set of EON-layer paths connecting the O/D pair. (11)

$\mathcal{T}_p$: Set of transponder modes whose optical reach O_t is sufficient for path $p \in \mathcal{P}$. (12)

$\mathcal{S}_p$: Set of spectrum segments along path p. (13)

B_{ps} : Number of available FSs in the s^{th} spectrum segment along the p^{th} shortest path. (14)

$\mathcal{L}_p$: Set of all lightpaths that use the path p. (15)

x_{plts} : Binary variable indicating whether lightpath $l \in \mathcal{L}_p$, uses transponder mode $t \in \mathcal{T}_p$ and spectrum segment $s \in \mathcal{S}_p$. (16)

y_{dpl} : Binary variable indicating whether Ethernet demand d is assigned to lightpath $l \in \mathcal{L}_p$ on path $p \in \mathcal{P}$. (17)

The optimization problem is formulated as follows:

$$Z_{MIP} = \min \sum_{p\in\mathcal{P}, l\in\mathcal{L}_p, t\in\mathcal{T}_p, s\in\mathcal{S}_p} x_{plts} F_t \tag{18a}$$

$$\text{s.t.} \quad \sum_{d\in\mathcal{D}} y_{dpl} W_d \le \sum_{t\in\mathcal{T}_p, s\in\mathcal{S}_p} x_{plts} R_t \qquad \forall p \in \mathcal{P}, l \in \mathcal{L}_p \tag{18b}$$

$$\sum_{p\in\mathcal{P}, l\in\mathcal{L}_p} y_{dpl} = 1 \qquad \forall d \in \mathcal{D} \tag{18c}$$

$$\sum_{l\in\mathcal{L}_p, t\in\mathcal{T}_p} x_{plts} F_t \le B_{ps} \qquad \forall p \in \mathcal{P}, s \in \mathcal{S}_p \tag{18d}$$

$$\sum_{t\in\mathcal{T}_p, s\in\mathcal{S}_p} x_{plts} \le 1 \qquad \forall p \in \mathcal{P}, l \in \mathcal{L}_p \tag{18e}$$

$$x_{plts}, y_{dpl} \in \{0, 1\} \tag{18f}$$

The objective function (18a) minimizes the total number of FSs used across all lightpaths. Constraint (18b) ensures that the total data rate of all Ethernet demands assigned to a lightpath does not exceed the capacity of the selected transponder mode. Constraint (18c) guarantees that each Ethernet demand $d \in \mathcal{D}$ is assigned to exactly one lightpath. Constraint (18d) limits the number of FSs used in any spectrum segment to stay within its capacity. Lastly, Constraint (18e) restricts each lightpath to at most one transponder mode. If no transponder mode is assigned (i.e., $x_{plts} = 0$ for all t and s), the lightpath becomes redundant and can be released.

This MIP model has a polynomial number of variables and constraints. However, for large-scale networks, it becomes computationally intractable even with state-of-the-art MIP solvers. In the next subsection, we present a reformulation that can efficiently handle large-scale instances using the column generation algorithm.

4.3 Problem reformulation

To address the computational challenges of large-scale instances, we reformulate model (18) around the notion of a *pattern*: a feasible grouping of Ethernet demands that can be carried on a single lightpath with a given transponder mode. For instance, a 600 Gb/s lightpath can serve either six 100GE demands or one 400GE and two 100GE demands, and each such combination is a distinct pattern. The following notations are introduced for the reformulation:

$$\mathcal{A} : \text{Set of feasible patterns.} \tag{19}$$
$$\mathcal{H} : \text{Set of all demand types.} \tag{20}$$
$$M_h : \text{Number of Ethernet demands of type } h. \tag{21}$$
$$\mu_{ha} : \text{Number of times Ethernet demand type } h \text{ is included in pattern } a \in \mathcal{A}. \tag{22}$$
$$\varphi_{ap} : \text{Binary parameter indicating whether pattern } a \text{ uses path } p. \tag{23}$$
$$\mathcal{F}_a : \text{Number of FSs used by pattern } a. \tag{24}$$
$$\eta_p : \text{Number of lightpaths on path } p. \tag{25}$$
$$\lambda_{as} : \text{Integer variable denoting the number of times pattern } a \tag{26}$$
$$\text{is used in spectrum segment } s \in \mathcal{S}_p. \tag{27}$$

The reformulated problem is:

$$Z^{1*}_{MP} = \min \sum_{a\in\mathcal{A}, s\in\mathcal{S}_p, p\in\mathcal{P}} \mathcal{F}_a \lambda_{as} \tag{28a}$$
$$\text{s.t.} \sum_{a\in\mathcal{A}, s\in\mathcal{S}_p, p\in\mathcal{P}} \mu_{ha}\lambda_{as} \geq M_h, \qquad \forall h \in \mathcal{H} \tag{28b}$$
$$\sum_{a\in\mathcal{A}} \mathcal{F}_a \varphi_{ap} \lambda_{as} \leq B_{ps}, \qquad \forall p \in \mathcal{P}, s \in \mathcal{S}_p \tag{28c}$$
$$\sum_{a\in\mathcal{A}, s\in\mathcal{S}_p} \varphi_{ap}\lambda_{as} \leq \eta_p, \qquad \forall p \in \mathcal{P} \tag{28d}$$
$$\lambda_{as} \in \mathbb{Z}_+ \tag{28e}$$

The objective (28a) minimizes the total FSs used by the selected patterns. Constraint (28b) ensures that all demands of each type are served. Constraint (28c) enforces the per-segment FS capacity. Constraint (28d) caps the number of patterns assigned to path p at the number of lightpaths already established on that path (η_p), so that no new transponders are deployed. Finally, (28e) defines the domain of the decision variables. By shifting from per-demand assignments to pattern selection, this reformulation exposes the problem to column-generation techniques, as detailed next.

4.4 Column generation framework

The pattern-based reformulation significantly increases the number of variables, as each pattern introduces a new variable. For large-scale networks, this results in exponential growth in the variable space, making direct optimization computationally infeasible. To address this challenge, we adopt a column generation framework [5].

Column generation is an iterative approach that decomposes the problem into a master problem and a set of pricing subproblems. In this context, the reformulated model (28) serves as the master problem, which focuses on determining the optimal mix of patterns. A linear relaxation of the Restricted Master Problem (RMP), which optimizes over a limited subset of possible patterns, is solved at each iteration. The pricing subproblems, guided by dual variables obtained from the RMP, identify new patterns (columns) with negative reduced cost that can improve the objective. These new patterns are added to the RMP, and the process iterates until convergence.

Pricing subproblems

For each path $p \in \mathcal{P}$ and transponder mode $t \in \mathcal{T}_p$, the pricing subproblem solves a knapsack problem (via dynamic programming) to generate patterns with negative reduced cost. Let π_h be the dual value associated with Constraint (28b). Since the duals of the segment-capacity and max-lightpath constraints enter the reduced cost only as terms independent of the pattern composition, they do not affect pattern generation and are omitted from the subproblem.

$$W_h : \text{Bandwidth requirement of Ethernet demand type } h. \tag{29}$$

$$z_h : \text{Number of times demand type } h \text{ is included in the new pattern.} \tag{30}$$

The subproblem is formulated as follows:

$$Z_{SP}^{1*} = \min \; F_t - \sum_{h \in \mathcal{H}} \pi_h z_h \tag{31a}$$

$$\text{s.t.} \quad \sum_{h \in \mathcal{H}} W_h z_h \leq R_t \tag{31b}$$

$$z_h \in \mathbb{Z}_+, \tag{31c}$$

The objective function (31a) minimizes the reduced cost of the master problem. Constraint (31b) ensures that the total capacity used by the selected demands does not exceed the transponder's available capacity.

4.5 Algorithm implementation

The column-generation framework for CLCD is presented in Algorithm 2. It iterates over each O/D pair (u, v). For a given pair, the algorithm first extracts the set of paths $\mathcal{P}$ currently used by (u, v) and identifies the corresponding available spectrum segments $\mathcal{S}_p$ along these paths. An initial set of feasible patterns $\mathcal{A}$ is constructed from the current network state, and a RMP is initialized from these patterns. The initial objective value Z_{MP}^1 corresponds to the FS usage of all lightpaths currently established for the O/D pair.

Next, the column generation loop is executed. At each iteration, the linear relaxation of the RMP is solved to obtain the dual variables associated with the demand-satisfaction constraint. These dual values are passed to the pricing subproblems, which, for every path $p \in \mathcal{P}$ and transponder mode $t \in \mathcal{T}_p$, solve a knapsack problem to search for new patterns with negative reduced cost. Any improving patterns found are added to $\mathcal{A}$, and the process repeats until no further patterns with negative reduced cost can be generated.

Once the column generation phase converges, the final RMP is solved to integrality using an MIP solver, yielding the optimal values λ_{as}^* and the corresponding objective value Z_{MP}^{1*}. The selected patterns λ_{as}^* determine which lightpaths are established on which segments, and a first-fit spectrum allocation policy arranges them in the spectrum. If the resulting configuration reduces the total FS usage compared to the initial state, it is then applied to the network for the considered O/D pair. To limit service disruption during this reconfiguration, a make-before-break strategy is employed, whereby new resources are activated before tearing down the old ones whenever possible.

5 Performance evaluation

In this section, we assess heuristic-based OTN consolidation and column-generation-based CLCD under dynamic traffic. We first describe the simulation setup, then evaluate OTN consolidation alone, CLCD alone against an EON-only baseline, and finally the combination of the two.

Algorithm 2 Column-Generation-Based Cross-Layer Consolidation and Defragmentation Algorithm

Input: Current network status, transponder types $\mathcal{T}$
Output: Updated lightpath configuration and Ethernet demand assignment
for each O/D pair (u, v) **do**
  Initialize $\mathcal{P}$ based on paths used for O/D pair (u, v)
  $\mathcal{S}_p \leftarrow$ Identify available spectrum segments
  $\mathcal{A} \leftarrow$ Construct initial patterns from current demands to lightpaths assignment
  $Z^1_{MP} \leftarrow \sum_{a \in \mathcal{A}, s \in \mathcal{S}_p} \mathcal{F}_a \lambda_{as}$
  repeat
    $\pi_h \leftarrow$ Solve linear relaxation of RMP
    newPatternFound $\leftarrow$ *false*
    for each path $p \in \mathcal{P}$ **do**
      for each transponder type $t \in \mathcal{T}_p$ **do**
        $\mathcal{A}' \leftarrow$ Solve the subproblem to find patterns with negative reduced cost
        if $\mathcal{A}' \neq \emptyset$ **then**
          $\mathcal{A} \leftarrow \mathcal{A} \cup \mathcal{A}'$
          newPatternFound $\leftarrow$ *true*
  until newPatternFound = *false*
  $\lambda^*_{as} \leftarrow$ Solve RMP with an MIP solver
  $Z^{1*}_{MP} \leftarrow \sum_{a \in \mathcal{A}, s \in \mathcal{S}_p} \mathcal{F}_a \lambda^*_{as}$
  if $Z^{1*}_{MP} < Z^1_{MP}$ **then**
    Update the lightpaths and Ethernet demands based on λ^*_{as} using first-fit spectrum allocation
Return Updated lightpath configuration and Ethernet demand assignment

5.1 Simulation setup

All experiments are conducted under a dynamic traffic model in which, for each O/D pair, Ethernet demands arrive according to a Poisson process with exponentially distributed holding times. The offered load per O/D pair is expressed in Erlangs; across the different traffic granularities, the aggregate load is varied between $\rho = 6$ and $\rho = 20$. Time is measured in Erlang time units, and defragmentation operations are executed once per time unit. We consider demand data rates of 10GE, 100GE, and 400GE, representing heterogeneous Ethernet traffic. The network is initially empty and gradually populated with demands until steady-state performance is observed. Each fiber link supports a spectrum of 768 FSs with a slot granularity of 6.25 GHz. The transponder operational modes, including their data rates and required number of FSs per lightpath, are summarized in Table 3; the reported FS values already include the spectral guard requirements for each mode. Simulations are performed on the 56-node NA and the 28-node USB topologies, shown in Fig. 5a and Fig. 5b, respectively.

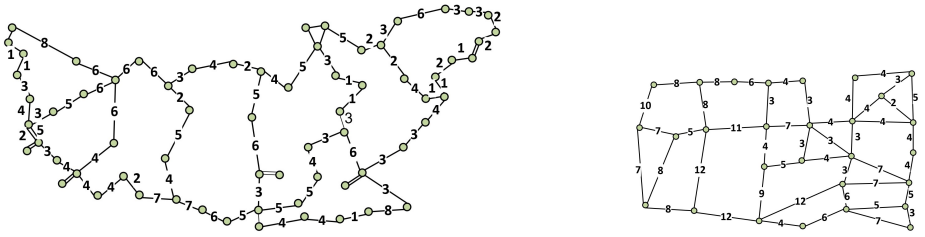

(a) 56-node NA topology (b) 28-node USB topology

Figure 5: Network topologies used in the simulations. Link lengths are expressed in number of fiber spans.

The simulations were executed on the Narval high-performance computing cluster, operated by Calcul Québec under the Digital Research Alliance of Canada, using 100 parallel instances. For CLCD, the restricted master problem was solved to a 0.02% MIP gap with Gurobi; column generation converged in 8–39 iterations (34 on average), and at most 153 distinct negative-reduced-cost patterns were generated across all instances. The overall solution process (pricing subproblems, LP relaxation of the master, and final MIP solution) completed in under 2 seconds per instance.

Table 3: Transponder Operational Modes

Total Capacity (Gb/s)	Modem Bitrate (Gb/s)	Modem Baudrate (Gbaud)	FSs required	#Modems	Maximum Spans
600	600	95	19	1	18
700	700	95	19	1	9
800	800	95	19	1	4
700	350	95	35	2	68
900	450	95	35	2	42
1100	550	95	35	2	23
1300	650	95	35	2	14
300	300	56	12	1	34
400	400	56	12	1	10
300	150	56	22	2	96
500	250	56	22	2	42
100	100	35	8	1	75
200	200	35	8	1	16
300	150	35	14	2	44

The evaluation uses the four performance metrics listed below:

Blocking Ratio: Ratio of total blocked Ethernet bandwidth to total requested bandwidth: $\frac{\sum_i W_i^{\mathrm{b}}}{\sum_i W_i^{\mathrm{r}}}$, where W_i^{b} is the total blocked bandwidth and W_i^{r} is the total requested bandwidth for O/D pair i.

Average Spectrum Usage %: Ratio of used to total FSs across all fiber links: $\frac{\sum_e \mathcal{F}_e^{\mathrm{u}}}{\sum_e \mathcal{F}_e^{\mathrm{t}}} \times 100\%$, where $\mathcal{F}_e^{\mathrm{u}}$ and $\mathcal{F}_e^{\mathrm{t}}$ denote the used and total FSs on link e.

Average Lightpath Capacity Usage: This reflects the utilization efficiency of all active lightpaths: $\frac{\sum_l (Q_l - r_l)}{\sum_l Q_l} \times 100\%$.

Lightpath Reconfiguration Ratio: Ratio of reconfigured to existing lightpaths: $\frac{L^{\mathrm{reconf}}}{L^{\mathrm{exist}}} \times 100\%$, where L^{reconf} and L^{exist} denote the corresponding quantities.

Note that we adopt a bandwidth-based definition of the blocking ratio instead of the traditional count-based metric used in EON-only networks, which measures the number of blocked demands over the total number of demands. Our approach is motivated by the heterogeneous nature of Ethernet traffic demands, which may range from 10GE to 400GE. In such scenarios, a simple request count fails to capture the actual load and resource usage. By measuring the proportion of blocked bandwidth instead of the number of blocked requests, we obtain a more accurate and meaningful representation of the network's performance under diverse traffic demands.

Moreover, for the Lightpath Reconfiguration Ratio, a lightpath is counted as reconfigured only if its optical configuration changes, for example through a transponder-mode modification, spectrum resizing, spectrum shifting, merging, or release. We restrict this metric to optical-layer changes because they are the primary source of service disruption. By contrast, Ethernet-demand reassignment between lightpaths whose optical configurations remain unchanged is performed entirely in the OTN domain and is fast and effectively hitless.

5.2 Effect of the OTN consolidation threshold

Fig. 6 shows the network's blocking ratio when periodic OTN consolidation is applied at each time unit using the heuristic algorithm described in Algorithm 1. As the threshold G (i.e., the minimum required reduction in lightpath utilization entropy to trigger consolidation) decreases from 0.5 to 0, the blocking ratio decreases, indicating that a lower threshold, and hence more frequent consolidation, enables the network to accommodate more demands by reorganizing resources more efficiently.

This improvement is corroborated by Fig. 7: at lower thresholds, OTN consolidation concentrates free capacity on fewer lightpaths, raising the effective per-lightpath utilization. Fig. 8 further shows that reducing G increases the number of consolidation operations while simultaneously lowering the average lightpath utilization entropy in the network, highlighting the trade-off between blocking-ratio gains and consolidation overhead.

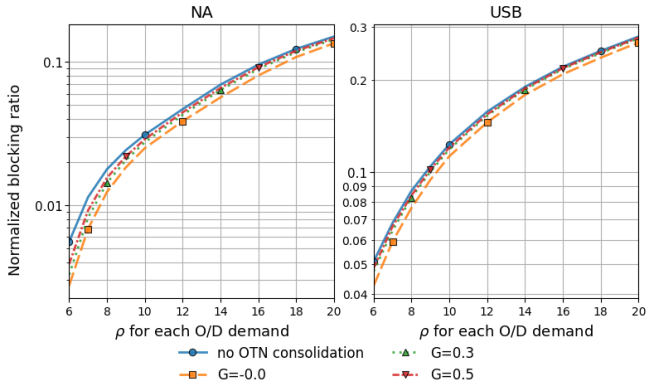

Figure 6: Blocking ratio of the network with periodic OTN consolidation applied, as a function of the threshold G.

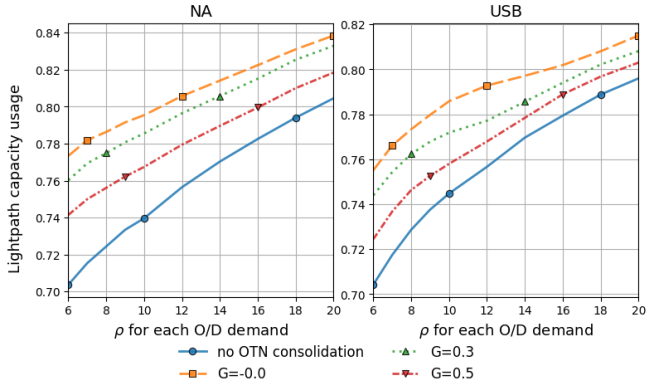

Figure 7: Lightpath capacity utilization as a function of the threshold G.

5.3 CLCD vs. EON-only defragmentation

CLCD is benchmarked against the well-established Oldest-First First-Fit (OF-FF) algorithm as the EON-only defragmentation strategy. In OF-FF, the longest-running lightpaths are reconfigured first and their new spectrum positions are decided by first-fit allocation. OF-FF is widely used as a baseline because of its strong blocking-ratio performance [3]. The algorithm has two parameters: the *period* between consecutive defragmentation cycles and the *number of lightpath reallocations per cycle*. To obtain the strongest baseline, we set the period to one time unit (matching the CLCD frequency) and allow *all* existing lightpaths to be reallocated in each cycle. Fig. 9 reports the blocking ratio for both strategies: CLCD significantly outperforms OF-FF, especially at higher traffic loads, because it acts on both layers, whereas OF-FF is limited to rearranging lightpaths in the spectrum. Figs. 10 and 11 provide insights into resource utilization:

Spectrum Utilization (Fig. 10): OF-FF can lead to higher spectrum utilization because aggressive compaction turns fragmented free FSs into larger *contiguous and usable* blocks, which satisfies continuity/contiguity constraints for more new lightpaths and thus raises the total number of used FSs.

Lightpath Capacity Utilization (Fig. 11): CLCD improves utilization by jointly repacking OTN-layer demands and selectively reconfiguring optical lightpaths, reducing residual unused capacity inside lightpaths and often decreasing the FS requirement per carried bit.

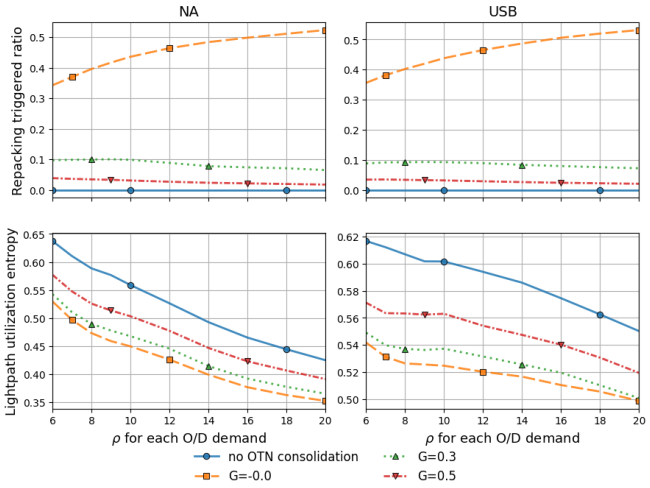

Figure 8: Number of consolidation operations and the average lightpath utilization entropy of the network as a function of the threshold G.

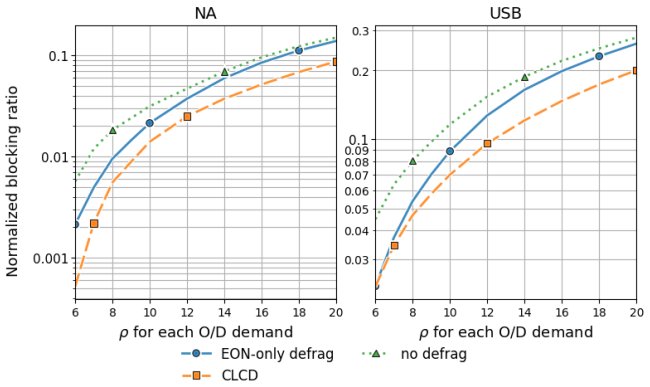

Figure 9: Blocking ratio: CLCD versus EON-only defragmentation (OF-FF).

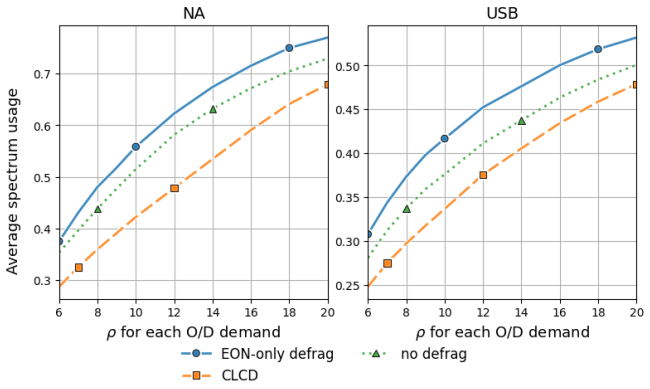

Figure 10: Spectrum utilization: CLCD versus EON-only defragmentation.

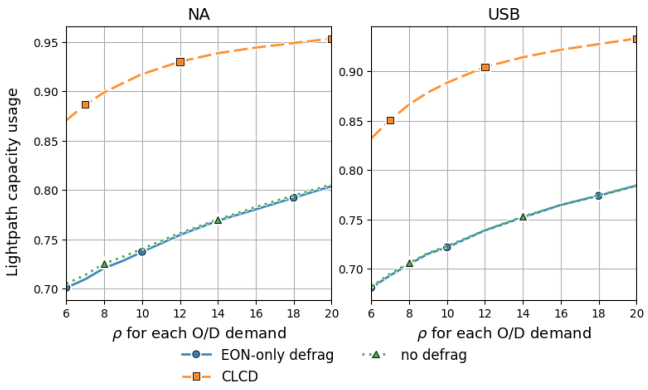

Figure 11: Lightpath capacity utilization: CLCD versus EON-only defragmentation.

5.4 Joint effects of OTN consolidation and CLCD

Fig. 12 compares the blocking ratio under OTN consolidation and CLCD. OTN consolidation alone improves performance, but the larger gains come from CLCD, which addresses both OTN-layer utilization entropy and EON-layer spectrum fragmentation. For the *NA* topology, OTN consolidation alone reduces the blocking ratio by about 9% at 20 Erlangs, 13% at 14 Erlangs, and 39% at 6 Erlangs; CLCD yields about 41% at 20 Erlangs and 43% at 14 Erlangs; the combined scheme reaches about 45% (20 Erlangs), 51% (14 Erlangs), and 90% (6 Erlangs) relative to the baseline. The second topology shows the same qualitative trends. Figs. 13 and 14 further show that CLCD improves both lightpath-capacity utilization and spectrum usage, with OTN consolidation providing an additional (smaller) improvement by consolidating residual OTN capacity between CLCD runs.

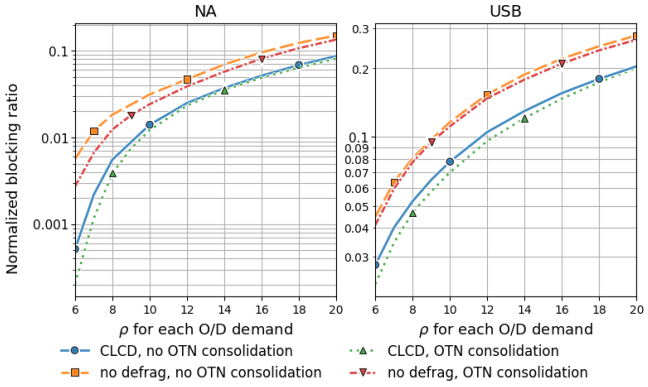

Figure 12: Blocking ratio under OTN consolidation, CLCD, and their combination.

We next evaluate how the CLCD *interval* affects performance. Fig. 15 compares CLCD applied every 1, 5, or 10 time units. When OTN consolidation is performed every time unit (Fig. 15a), increasing the CLCD interval has only a limited impact on blocking, indicating that OTN consolidation partially mitigates OTN-layer packing inefficiency between CLCD runs. Without OTN consolidation (Fig. 15b), performance degrades more noticeably as the interval grows. Hence OTN consolidation is a useful complement, while CLCD remains essential when optical-layer fragmentation is the dominant limiter.

Fig. 16 shows the *lightpath reconfiguration ratio*. As the CLCD interval grows, more lightpaths must be modified per run because utilization entropy and spectrum fragmentation have more time to accumulate; more frequent execution keeps the network better packed and requires fewer changes. OTN

consolidation further reduces the reconfiguration ratio by maintaining efficient transport-layer packing with inexpensive electrical-domain moves. These results suggest that CLCD alone is not lightweight and is better complemented by frequent OTN consolidation than by a shorter CLCD interval.

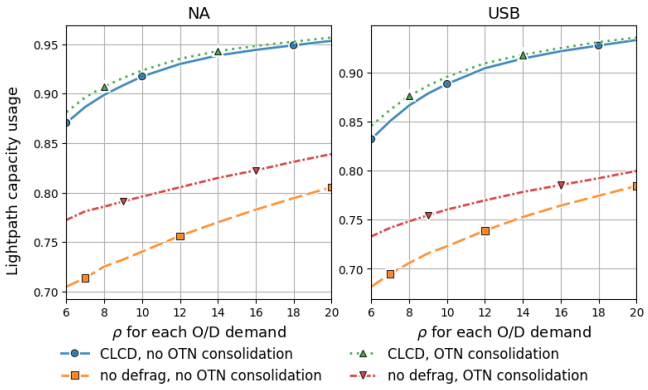

Figure 13: Lightpath capacity utilization under CLCD, with and without OTN consolidation.

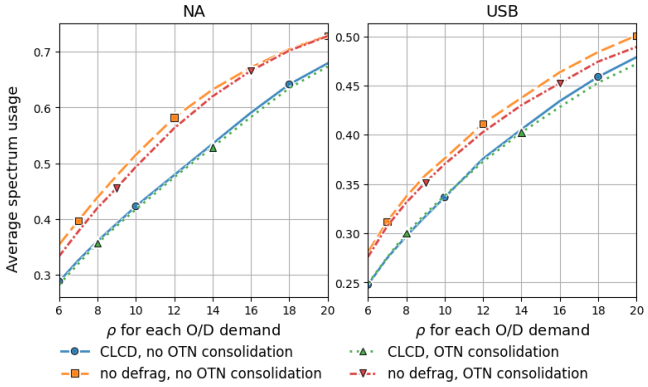

Figure 14: EON-layer spectrum usage under CLCD, with and without OTN consolidation.

6 Conclusion

Dynamic Ethernet traffic in OTN-over-EON networks produces two distinct inefficiencies: lightpath underutilization at the OTN layer and spectrum fragmentation at the EON layer. We introduced the *lightpath utilization entropy* metric and a fast *OTN consolidation* heuristic to address lightpath underutilization. We also proposed a *CLCD* framework, formulated as an FS-minimization problem and solved by column generation, to jointly address lightpath underutilization and spectrum fragmentation. The combination of FS minimization with first-fit deployment packs traffic onto the fewest slots at the lowest available indices, thereby effecting spectrum defragmentation.

Simulations on two network topologies show that while OTN consolidation alone reduces the blocking ratio, the larger gains come from CLCD; combining the two reduces the blocking ratio by up to 51% at medium traffic loads and also improves spectrum usage and lightpath-capacity utilization relative to single-layer (EON-only) defragmentation. Frequent OTN consolidation complemented by periodic CLCD therefore provides a favorable trade-off between resource efficiency and reconfiguration overhead. Extending the framework to support survivability and redundancy constraints is left as future work.

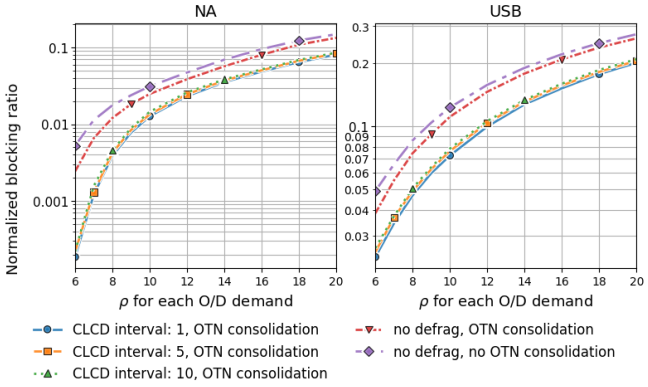

(a) With OTN consolidation

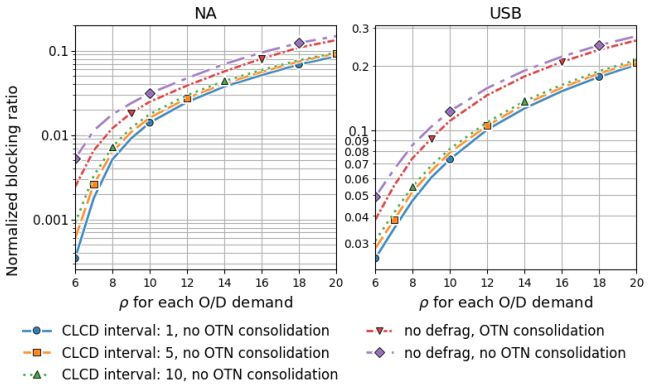

(b) Without OTN consolidation

Figure 15: Impact of CLCD frequency on the blocking ratio. (a) With OTN consolidation at every time unit. (b) Without OTN consolidation.

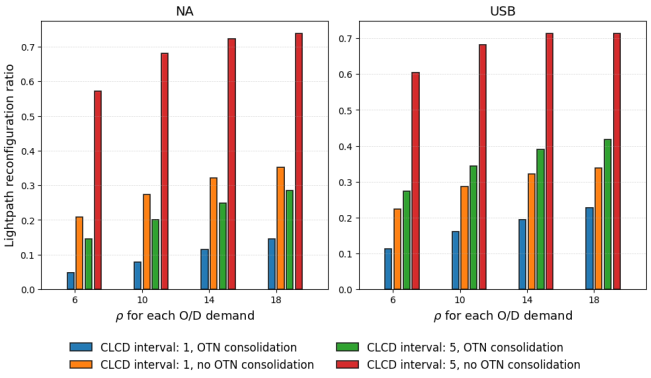

Figure 16: Lightpath reconfiguration ratio under CLCD at different intervals, with and without OTN consolidation.

References

[1] Ciena Corporation. Wavelogic 5 extreme and nano coherent optical solutions. White paper, Ciena, 2020. Data Sheet.

[2] J. Comellas, A. Asensio, M. Ruiz, G. Junyent, and L. Velasco. Using spectrum management techniques to differentiate services in elastic optical networks. In 2015 17th International Conference on Transparent Optical Networks (ICTON), pages 1–4, 2015.

[3] J. Comellas, L. Vicario, and G. Junyent. Proactive defragmentation in elastic optical networks under dynamic load conditions. Photonic Network Communications, 36:26–34, 2018.

[4] F. Cugini, F. Paolucci, G. Meloni, G. Berrettini, M. Secondini, F. Fresi, N. Sambo, L. Poti, and P. Castoldi. Push-pull defragmentation without traffic disruption in flexible grid optical networks. Journal of Lightwave Technology, 31(1):125–133, 2013.

[5] Guy Desaulniers, Jacques Desrosiers, and Marius Solomon. Column Generation. GERAD 25th Anniversary Series. Springer, 2005.

[6] Ruman Dutta, S. Sai Krishna, R. Gowrishankar, G. J. Subhanu, Siva Sankara Sai Sanagapati, Prabhat Praveen Behere, and Sai Kishore Bhyri. Fragmentation-aware routing, wavelength and spectrum assignment (rwsa) scheme in flex-grid optical networks. In 2017 IEEE International Conference on Advanced Networks and Telecommunications Systems (ANTS), pages 1–6, 2017.

[7] Ehsan Etezadi, Carlos Natalino, Renzo Diaz, Anders Lindgren, Stefan Melin, Lena Wosinska, Paolo Monti, and Marija Furdek. Deep reinforcement learning for proactive spectrum defragmentation in elastic optical networks. Journal of Optical Communications and Networking, 15(10):E86–E96, 2023.

[8] International Telecommunication Union . Characteristics of Optical Transport Network Hierarchy Equipment Functional Blocks. ITU-T Recommendation G.798, International Telecommunication Union, September 2023.

[9] International Telecommunication Union - OTN . Interfaces for the optical transport network (otn), g.709 recommendation. ITU-T Recommendation G.709/Y.1331, International Telecommunication Union, June 2020.

[10] Huan-Lin Liu, Lei Lv, Yong Chen, and Chengying Wei. Fragmentation-avoiding spectrum assignment strategy based on spectrum partition for elastic optical networks. IEEE Photonics Journal, 9(5):1–13, 2017.

[11] Jiaxin Liu, Ziyi Xi, and Rentao Gu. Adaptive cross-layer bandwidth defragmentation for multi-band optical network. In 2023 Asia Communications and Photonics Conference/2023 International Photonics and Optoelectronics Meetings (ACP/POEM), pages 1–5, 2023.

[12] E. Oki, K. Shiomoto, D. Shimazaki, N. Yamanaka, W. Imajuku, and Y. Takigawa. Dynamic multilayer routing schemes in gmpls-based ip+optical networks. IEEE Communications Magazine, 43(1):108–114, 2005.

[13] M. Quagliotti, D. Cifuentes Arango, M. Schiano, A. Carena, M. Cantono, and V. Curri. Spectrum fragmentation metrics and their use in optical channel allocation algorithms. In 19th Italian National Conference on Photonic Technologies (Fotonica 2017), pages 1–4, 2017.

[14] Nicola Sambo, Francesco Paolucci, Gianluca Meloni, Francesco Fresi, Luca Potì, and Piero Castoldi. Control of frequency conversion and defragmentation for super-channels [invited]. Journal of Optical Communications and Networking, 7(1):A126–A134, 2015.

[15] Zhi-shu Shen, Hiroshi Hasegawa, and Ken-ichi Sato. Integrity enhancement of flexible/semi-flexible grid networks that minimizes disruption in spectrum defragmentation and bitrate-dependent blocking. Journal of Optical Communications and Networking, 7(4):235–247, 2015.

[16] Domenico Siracusa, Attilio Broglio, Andrea Zanardi, Elio Salvadori, Gabriele Galimberti, and Domenico La Fauci. Hitless network re-optimization to reduce spectrum fragmentation in distributed gmpls flexible optical networks. In 39th European Conference and Exhibition on Optical Communication (ECOC 2013), pages 1–3, 2013.

[17] Takafumi Tanaka, Tetsuro Inui, Akihiro Kadohata, Wataru Imajuku, and Akira Hirano. Multiperiod ip-over-elastic network reconfiguration with adaptive bandwidth resizing and modulation. Journal of Optical Communications and Networking, 8(7):A180–A190, 2016.

[18] Rui Wang and Biswananth Mukherjee. Spectrum management in heterogeneous bandwidth networks. In 2012 IEEE Global Communications Conference (GLOBECOM), pages 2907–2911, 2012.

[19] Rui Wang and Biswanath Mukherjee. Provisioning in elastic optical networks with non-disruptive defragmentation. Journal of Lightwave Technology, 31(15):2491–2500, 2013.

[20] Paul Wright, Michael C. Parker, and Andrew Lord. Simulation results of shannon entropy based flexgrid routing and spectrum assignment on a real network topology. In 39th European Conference and Exhibition on Optical Communication (ECOC 2013), pages 1–3, 2013.

[21] Paul Wright, Michael C. Parker, and Andrew Lord. Minimum- and maximum-entropy routing and spectrum assignment for flexgrid elastic optical networking [invited]. Journal of Optical Communications and Networking, 7(1):A66–A72, 2015.

[22] Mingyang Zhang, Weiran Shi, Long Gong, Wei Lu, and Zuqing Zhu. Bandwidth defragmentation in dynamic elastic optical networks with minimum traffic disruptions. In 2013 IEEE International Conference on Communications (ICC), pages 3894–3898, 2013.

[23] Yunrong Zhang, Ya Zhang, Yongcheng Li, Gangxiang Shen, Yonghu Yan, and Wei Chen. Cross-layer spectrum defragmentation for ip over elastic optical network. In 2018 Asia Communications and Photonics Conference (ACP), pages 1–4, 2018.