A Power-Consumption Analysis for Different IPoWDM Network Architectures with ZR/ZR+ and Long-Haul Muxponders

QIAOLUN ZHANG[1,3,*], ANNALISA MOREA[2], PATRICIA LAYEC[3], MEMEDHE IBRAHIMI[1], FRANCESCO MUSUMECI[1], AND MASSIMO TORNATORE[1]

[1] *Department of Electronics, Information and Bioengineering, Politecnico di Milano, 20133 Milano, Italy*
[2] *Nokia, Vimercate, Italy*
[3] *Nokia Bell Labs France, 91300 Massy, France*
[*] *Corresponding author: qiaolun.zhang@mail.polimi.it*

Compiled March 19, 2025

Operators are constantly faced with the need to increase optical-network capacity to accommodate rapid traffic growth while minimizing the cost-per-bit and power-per-bit. The drastic reduction of power consumption of IP routers and ZR/ZR+ pluggable transponders seen in the last years has renewed the interest in "opaque" optical-network architectures, where no optical bypassing is allowed. In this work, we aim to quantify and compare the power consumption of four "IP over Wavelength Division Multiplexing" (IPoWDM) transport network architectures employing ZR/ZR+ modules vs. long-haul muxponders, considering different grooming, regeneration, and optical bypassing capabilities. We first propose a power consumption model for different IPoWDM node architectures with ZR/ZR+ modules and long-haul muxponders. Then, to obtain the power consumption of different architectures, we propose a compact auxiliary-graph-based network-design algorithm extensible to different network architectures. Moreover, we investigate how the continuous decrease in the power consumption of ZR/ZR+ and IP routers can impact the power consumption of different architectures through a sensitivity analysis. Illustrative numerical results on networks of different sizes show that, despite drastic reductions of power consumption at IP layer, optical bypassing is still the most power-efficient solution, reducing consumption by up to 48%.

http://dx.doi.org/10.1364/ao.XX.XXXXXX

1. INTRODUCTION

Power efficiency of the optical-network infrastructure is a top priority for network operators, that need to accommodate ever-growing traffic without increasing associated power-per-bit [1]. The latest generation of 400G ZR/ZR+ coherent transmission modules enables router client ports to support up to 400 Gb/s using pluggable optics, which have lower equipment cost and lower power consumption compared to traditional long-haul muxponders [2–4]. Meanwhile, power consumption of state-of-art IP routers is also decreasing rapidly [5, 6], being reduced from 14.5 W to 0.022 W per Gb/s in the last decade [5, 6]. "IP over Wavelength Division Multiplexing" (IPoWDM) network architectures can be classified into *opaque* and *transparent*, according to optical switching capabilities [5]. Specifically, in transparent network architectures, lightpaths can optically bypass intermediate nodes between source and destination nodes, while in opaque network architectures, switching and regeneration of traffic are performed directly in IP routers at each traversed intermediate node. The emerging low-cost/low-power silicon technologies have brought back in the foreground "opaque" network architectures [5, 7], also considering that opaque networks can increase their capacity thanks to the shorter transmission distance, ubiquitous electrical grooming and the elimination of the spectrum contiguity constraint [8–10]. Although transparent architectures reduce power consumption of IP layer, they also introduce more complex and costly switching components, such as Reconfigurable Optical Add/Drop Multiplexer (ROADMs), leading to higher power consumption at the optical layer. Recent advances in reducing power consumption at the IP layer lead to the need to investigate if the transparent architectures are still favorable in terms of power consumption.

So far, most research work has been devoted to evaluating the cost reduction enabled by ZR/ZR+ modules compared to long-haul muxponders [2, 17–19], but only a few works investigated the power consumption of network architectures adopting plug-

Table 1. Comparison between this work and other existing works about the energy efficiency of ZR/ZR+ at the network scale.

Related work	Year	Architecture				Analysis of impact of grooming	Analysis of impact of regeneration	Router power model	400G router port constraint	Sensitivity analysis	Approach
		OP-IP	Tr-IP	Tr-O	Tr-IP&O						
[11]	2021	Not specified				✗	✗	✗	✗	✗	Heuristic
[12]	2022	✓	✓	✗	✗	✓	✓	✓	✗	✗	ILP
[13]	2022	✓	✓	✗	✗	✗	✓	✗	✗	✗	Heuristic
[14]	2022	✓	✓	✗	✗	✗	✓	✓	✗	✗	ILP&Heuristic
[15]	2023	✓	✓	✗	✗	✗	✓	✗	✗	✗	Heuristic
[16]	2024	✓	✓	✗	✗	✗	✗	✗	✗	✗	Heuristic
This work	2024	✓	✓	✓	✓	✓	✓	✓	✓	✓	Heuristic

gable transponders, using either heuristic algorithms or Integer Linear Programming (ILP) model [11–16]. Table 1 compares this work to other recent works investigating the energy efficiency of ZR/ZR+ at network scale. Existing works do not investigate the power consumption of ZR/ZR+ using network architectures with optical-layer regeneration, specifically *Tr-O* and *Tr-IP-O*, as outlined in Section 2.A. Besides, existing research studies have not conducted a sensitivity analysis on the power consumption of ZR/ZR+ and routers. This analysis is critical for generalizing the study's findings to scenarios where power consumption for ZR/ZR+ and routers varies, e.g., due to adoption of devices from different vendors or due to advancements in technology. Moreover, unlike other works, this work comprehensively addresses the impact of grooming, regeneration, and router power model, and also performs sensitivity analysis for ZR/ZR+ and router, as shown in Table 1.

This study aims to quantify and compare the power consumption of different IPoWDM network architectures employing ZR/ZR+ modules vs. long-haul muxponders. This is an extended version of our previous work in [3]. The main contributions and differences compared to Ref. [3] are as follows.

- We propose a pragmatic power consumption model for different optical node architectures, considering the state-of-the-art power consumption of both IP-layer and optical-layer devices. Compared to Ref. [3], we apply this model to long-haul muxponders for comparison with ZR/ZR+.
- We develop a Compact Auxiliary-Graph-based Network-Design (CAG-ND) algorithm for the Routing, Modulation format, and Spectrum Assignment problem. The proposed algorithm is extensible to all the different IPoWDM network architectures considered in this study.
- We conduct a sensitivity analysis to evaluate the impact of reducing power consumption in IP routers and ZR/ZR+ on various IPoWDM architectures, which is not investigated in Ref. [3].
- We provide extensive illustrative numerical evaluations to compare the cost and power consumption of different network architectures for both ZR/ZR+ and long-haul muxponders. These results provide a more thorough analysis considering the role of grooming and regeneration in different IPoWDM architectures compared to Ref. [3].

The rest of the paper is organized as follows. Sec. 2 classifies the different IPoWDM network architectures, proposes a detailed power consumption model, and provides a formal problem statement. Sec. 3 describes our proposed auxiliary-graph-based heuristic algorithm, applicable to all the proposed network architectures. Sec. 4 presents and discusses the obtained numerical results. Finally, Sec. 5 concludes the paper and discusses future work.

2. NETWORK ARCHITECTURES AND POWER CONSUMPTION MODEL

A. IPoWDM Network Architectures

We consider four different IPoWDM network architectures where switching and regeneration can be performed at the IP and/or the optical layer (note that each of these four architectures can choose to adopt either ZR/ZR+ modules or long-haul muxponders, for a total of 8 compared scenarios). For the sake of brevity, Fig. 1 only illustrates the four architectures with ZR/ZR+ modules. As shown in the figure, ZR/ZR+ modules can be directly connected to the routers, while long-haul muxponders need to be connected to routers through grey optics [14].

Fig. 1 (a) depicts the **Opaque-IPoWDM (Op-IP)** case, where traffic is groomed and regenerated at each intermediate router. At the optical layer, we assume that lightpaths are MUXed/DeMUXed using arrayed waveguide grating (AWG). Fig. 1 (b) shows the **Transparent with regeneration at IP layer (Tr-IP)** case, where traffic can be groomed and regenerated at intermediate routers (as in *Op-IP*), but nodes are also equipped with ROADMs that enable transparent lightpath bypass. Fig. 1 (c) shows the **Transparent with regeneration at optical layer (Tr-O)** case, where regeneration is not performed in router, but is realized by optical-layer regenerators built using transponders in back-to-back configurations; in this configuration, no intermediate but only end-to-end grooming is allowed. Finally, Fig. 1 (d) shows the **Transparent with regeneration at both IP and optical layer (Tr-IP&O)** case, which combines the *Tr-IP* and *Tr-O* features, allowing both IP-layer and optical-layer regeneration. Note that to accommodate long distances, we consider lower modulation formats, without considering regenerators at non-node sites, i.e., along the links.

B. Power Consumption Modeling for IPoWDM Nodes

The power consumption of an IPoWDM node consists of the power consumption of transceivers, IP routers, and optical devices. Table 2 reports the normalized power consumption of the network elements considered in our power model, which is obtained from Nokia products or verified internally in Nokia. The reported values are normalized with respect to the power consumption of a ZR transmission module, where the power consumption of the ZR module is considered as 1. The power consumption of ZR/ZR+ and long-haul muxponders is obtained internally in Nokia with optical specifications defined in Ref. [20–22]. Note that we focus only on the *node* power consumption

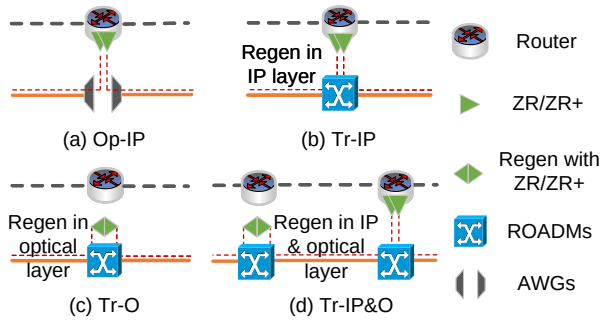

Fig. 1. IPoWDM network architectures with ZR/ZR+.

as power consumption of devices installed in optical fiber links (e.g., amplifiers) does not depend on the IPoWDM architecture being considered.

Table 2. Normalized power consumption of devices.

Device		Power	
		Notation	Value
Transceiver	ZR	p_{zr}	1
	ZR+	p_{zr+}	1.3
	Long-haul muxponders	p_{lm}	8
IP router	IP Fixed part	p_{fix}	75
	Router port	p_{port}	4
	I/O card	$p_{i/o}$	1
Optical side	OA (undirectional)	p_{oa}	1.7
	AWG	p_{awg}	0.3
	I-ROADM (bidirectional)	$p_{I-ROADM}$	4.1
	ADB	p_{ADB}	4
	Shelf	p_{shelf}	20
	Monitoring for opaque (bidirectional)	p_{mo}	0.3
	Monitoring for transparent (bidirectional)	p_{mp}	1.5

The IP-router power consumption model is developed based on an actual Nokia router with a maximum capacity of 36 Tb/s [23]. Note that we provide an approximated model verified by Nokia engineers, as the exact power consumption model is proprietary. The power consumption of the IP router comprises a fixed part (the chassis with fans, power pack, control board, and switching fabric) plus a modular part that depends on the total traffic handled by the router. Table 2 lists the values of the fixed part and modular part of an IP router. When traffic exceeds the capacity of a single router, multiple IP routers with backplane interactions [24, 25] are deployed. Since we are addressing a static problem, we evaluate router traffic first, and then dimension the router counts and modular components accordingly. Regarding the modular part, 100 Gb/s of traffic in the router requires 1 unit of power consumption. Yet, it is important to note that the calculation of the carried traffic differs for ZR/ZR+ and long-haul muxponders. Specifically, since ZR/ZR+ modules must be connected to 400G ports, each ZR/ZR+ module consumes 400 Gb/s traffic (4 units of power consumption) in the router even if the data rate of its operated modulation format is smaller than 400 Gb/s. Instead, long-haul muxponders are connected to routers through 100 G I/O cards, each of which consumes 100 Gb/s traffic in the router, and consequently 1 unit of power consumption. Our assumption of using only 100 G grey interfaces has no impact on the power consumption analysis for two reasons. First, the 400 G grey interface is only used for 400 Gb/s traffic, resulting in equivalent carried traffic in routers. Second, the power consumption of grey optics (100G or 400G) is neglected since it is relatively negligible compared to long-haul muxponders. Without loss of generality, the effects of reduction of power consumption per unit of traffic due to the usage of larger routers and technology advances are analyzed in our numerical evaluation by assuming different percentages of reduction of power consumption. Note that we neglect the inter-connection between multiple routers because the current inter-connection between routers typically relies on optical pluggables, resulting in negligible power consumption [24, 25]. Assume that the total number of carried traffic in node i is t_i Gb/s, and the number of IP fixed parts in node i can be approximated with $\lceil t_i/(5*1000)\rceil$, as one IP router has a maximum capacity of 5 Tb/s. Since the carried traffic is a multiple of the carried traffic of one router interface (400 Gb/s and 100 Gb/s for one router port and one I/O card, respectively), the number of router ports and the number of I/O cards are $t_i/400$ and $t_i/100$ for ZR/ZR+ and long-haul muxponders, respectively. Assume that the power consumption of one IP router fixed part, one router port, and one I/O card are denoted with p_{fix}, p_{port}, and $p_{i/o}$, respectively. In addition, the set of nodes in the network is denoted with N_p. The IP-router power consumption of all the nodes can be calculated with Eqn. (1) and Eqn. (2) for ZR/ZR+ and long-haul muxponders, respectively. Note that the carried traffic with ZR/ZR+ in Eqn. (1) is larger than that with long-haul muxponders in Eqn. (2), even when both are subjected to identical routing for requests because I/O card has a smaller granularity than the router port. For instance, if the carried traffic on the lightpath is 200 Gb/s, the handled traffic in the router is 400 Gb/s for ZR/ZR+ (one router port). Instead, the handled traffic in the router can be reduced to 200 Gb/s for long-haul muxponders (two I/O cards).

$$p_{router} = \sum_{i \in N_p} (p_{fix} * \lceil t_i/(5*1000)\rceil + p_{port} * t_i/400) \quad \textbf{(1)}$$

$$p_{router} = \sum_{i \in N_p} (p_{fix} * \lceil t_i/(5*1000)\rceil + p_{i/o} * t_i/100) \quad \textbf{(2)}$$

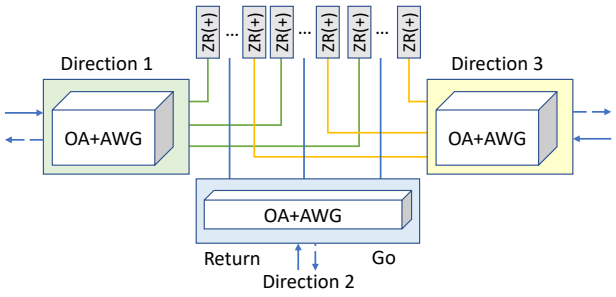

Fig. 2. Optical node architecture for the opaque network.

As for the optical side, the node architectures for opaque and transparent architectures are illustrated in Fig. 2 and Fig. 3, respectively, using ZR/ZR+ modules. In the case of *Op-IP*, as shown in Fig. 2, the optical node is composed of a terminal per direction realized by AWGs (Mux/Demux), each of which is associated with an optical amplifier (OA) to cope with line losses (2 OAs per node direction are required). Demultiplexed optical

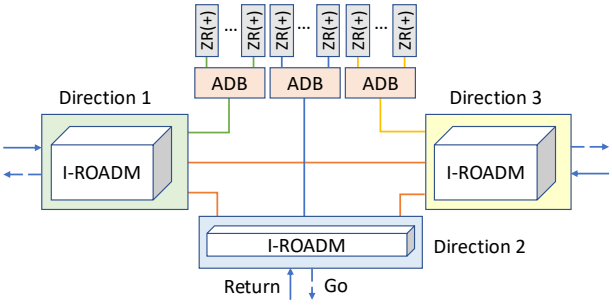

Fig. 3. Optical node architecture for the transparent network.

channels are then directly connected to ZR/ZR+ pluggable or long-haul muxponders through grey optics. Instead, for the transparent cases, to enable optical bypass, WSSs are used as Mux/Demux to switch the optical channels either toward one of the node's output directions or toward transceivers through an add/drop block (ADB). To improve overall power efficiency, a network element called I-ROADM [26], which integrates WSSs and OAs for both directions (Go and Return), is used. Note that, this work considers implementing ADB with AWG instead of WSS because AWG has a lower cost and power consumption. An optical node also requires a number of shelves that must be dimensioned based on the number of I-ROADMs, ADBs, and long-haul muxponders required in the node. Specifically, an optical shelf has 16 slots. Both one long-haul muxponder and one OA occupy 1 slot, and an I-ROADM and an ADB both occupy 2 slots. Hence, each direction for *Op-IP* has 2 OAs and occupies 2 slots. Each direction for the transparent case has 1 I-ROADM, which is connected to one dedicated ADB. Thus, one direction of the transparent case occupies 4 slots. Note that ZR/ZR+ does not require slots on the shelf, since it is directly connected to the router. Instead, long-haul muxponders are connected to the router through grey optics and still require one slot per muxponder in the optical shelf. The AWG is outside the shelf and hosted in the same rack. Our model accounts for the power consumption of the shelves, not of the rack. Lastly, our model also accounts for the monitoring of the optical node. Specifically, each direction of the optical node for opaque and transparent nodes requires 0.3 and 1.5 units of power consumption, respectively, which is verified internally in Nokia. Assume that the degree and the number of long-haul muxponders used in node i are denoted with x_i and m_i, respectively, then the power consumption in the optical layer can be calculated with Eqn. (3) and Eqn. (4) for opaque and transparent architectures, respectively. Note that in addition to the power consumption of devices on the optical side, such as I-ROADM and shelf, Eqn. (3) and Eqn. (4) also consider the power consumption of monitoring of nodes. Specifically, p_{mo} and p_{mt} denote the power consumption of monitoring the node per node degree for opaque architecture and transparent architecture, respectively.

$$p_{optical} = \sum_{i \in N_p} (2x_i(p_{OA} + p_{AWG}) + p_{shelf}\lceil \frac{2x_i + m_i}{16} \rceil + x_i p_{mo}) \quad (3)$$

$$p_{optical} = \sum_{i \in N_p} (x_i p_{I-ROADM} + p_{shelf}\lceil \frac{4x_i + m_i}{16} \rceil + x_i p_{mt}) \quad (4)$$

C. Reach Table and Problem Statement

We obtained the reach table of ZR and ZR+ as in Table 3 based on Ref. [4] and the implementation agreement of 400ZR by Optical Internetworking Forum (OIF) [20] and the optical specification of the standard OpenZR+ [21]. This reach table is also aligned with the one in the whitepaper of OpenZR+ [27] and on the field [28]. The output power for ZR is defined as -20 dBm as defined in Ref. [20] and the output power of ZR+ is 0 dBm as defined in Ref. [21]. The required OSNR values at the FEC limit in a back-to-back configuration for ZR/ZR+ are detailed in Table 3. The reach table of the muxponder is shown in Table 4, which is obtained from the same long-haul muxponder operating under different modulation levels. The maximum transmission distance is based on Ref. [22] and slightly modified under internal verification. The required OSNR values for the long-haul muxponder cannot be provided due to confidentiality concerns. However, for further information on this matter, please refer to Reference in Ref. [22]. Note that, in this work, we do not allow inverse-multiplexing (one traffic request cannot be split and served by more than one transceiver) because it requires precise synchronization between multiple multiplexed channels together with reduced hardware and software complexity.

Table 3. Reach table for ZR/ZR+ pluggable optics.

Modules	Data rate (Gb/s)	Modulation format	Spacing (GHz)	Reach (km)	OSNR (dB/0.1nm)	Baud rate (GBd)
ZR	400	16 QAM	100	120	34	30
ZR+	400	16 QAM	75	600	34	60
ZR+	300	8 QAM	75	1800	34	60
ZR+	200	QPSK	75	3000	34	60
ZR+	100	QPSK	50	3000	34	30

Table 4. Reach table for the considered long-haul muxponder.

Data rate (Gb/s)	Modulation format	Spacing (GHz)	Reach (km)	Baud rate (GBd)
800	PCS 64 QAM	100	150	90
700	PCS 64 QAM	100	400	90
600	16 QAM	100	700	90
500	PCS 16 QAM	100	1300	90
400	PCS 16 QAM	100	2500	90
300	PCS 16 QAM	100	4700	90
200	PCS 16 QAM	100	5700	90

The problem solved in this paper can be formally stated as follows: **Given** a network topology, selected network architectures, a set of traffic requests, and a reach table for ZR/ZR+ modules and long-haul muxponders, **decide** the Routing, Grooming, Modulation format, and Spectrum Assignment (RGMSA), together with transceiver (i.e., either ZR or ZR+, in case an IPoWDM architecture with ZR/ZR+ modules is adopted) to be used for each traffic request, **constrained** to reach and data rate of each modulation formation, the link capacity, and no inverse-multiplexing, with the **objective** to minimize the cost of transceivers as the primary objective function and the number of router interfaces as the secondary objective function, as an approximation for network cost[1]. For the obtained solution, we calculate the overall power consumption a-posteriori and compare the various IPoWDM architectures.

[1]Note that we do not consider the cost of AWG and I-ROADM to approximate the network cost, as the number of AWG and I-ROADM is fixed and does not affect the optimization.

3. AUXILIARY-GRAPH BASED HEURISTIC ALGORITHM

A. Auxiliary Graph Model

To evaluate the performance of different network architectures, we propose a novel Compact Auxiliary-Graph-based Network-Design (CAG-ND) algorithm that can be employed for all the network settings by just assigning different weights for the auxiliary graph. Note that our aim is to provide a principled performance comparison of different network architectures rather than an optimal solution. Hence, we select the auxiliary-graph-based approach, because it has been proven to be a good heuristic for the routing problem for optical networks with grooming, regeneration, and different line rates [29–33]. The difference between the proposed CAG-ND algorithm and the previous auxiliary-graph-based algorithms (as, e.g., [31, 32]) is twofold. First, the proposed CAG-ND algorithm applies to all the architectures covered in this work, where regeneration can be done in the optical layer and/or IP layer, and traffic grooming together with multiple modulation formats is allowed. Second, the proposed CAG-ND algorithm adopts a more compressed representation of the auxiliary graph compared to the layered auxiliary graph model as in Ref. [32], which reduces the computational complexity by simplifying the auxiliary graph.

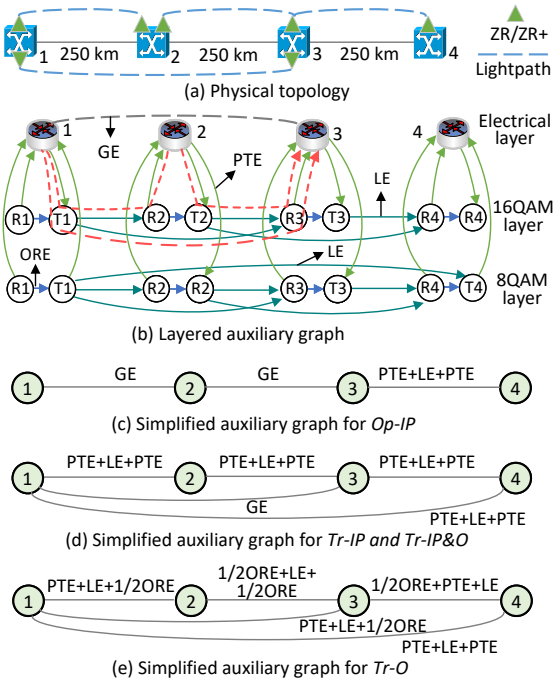

Fig. 4. Illustration of the CAG model.

Fig. 4 demonstrates an example of the proposed compact auxiliary graph (CAG) model in our work. Fig. 4(a) shows the physical topology, and Fig. 4(b) shows a layered auxiliary graph that contains all the possible edges in *Tr-IP&O* because *Tr-IP&O* has the highest flexibility of switching and regeneration. Hence, the edges of all other architectures are just a subset of *Tr-IP&O*. The layered auxiliary graph model in Fig. 4(b) contains two types of layers, namely, *electrical layer* and *modulation layer*. The electrical layer models traffic grooming, and the modulation layer models the possibility of establishing lightpath between node pairs with a given modulation level. Note that each row in Table 3 and Table 4 corresponds to one modulation layer to take into account different spacing and types of transceivers (ZR or ZR+). Each node in the physical topology corresponds to two nodes (defined as *transmitter* node and *receiver* node) in each modulation layer. Specifically, the transmitter node and the receiver node of the node i in the physical topology are denoted with Ri and Ti respectively, which models the usages of the transceiver and receiver in each node. For simplicity, Fig. 4 only lists two modulation layers, namely, 16-QAM layer and 8-QAM layer for ZR+ operating at 16-QAM and 8-QAM, respectively. The auxiliary graph in Fig. 4 (b) is utilized to construct different CAGs for different architectures, as shown in Fig. 4 (c)-(e), which will be described later in Sec. 3.B. The layered auxiliary graph is constructed when serving a new request, and the details of different types of edges are listed below.

Potential Transceiver Edge - (PTE): Exists between the node in the electrical layer and the *transmitter/receiver nodes* in each modulation layer.

Optical Regeneration Edge - (ORE): Exists between the transmitter node and the receiver node of a physical node in each modulation format layer.

Lightpath Edge - (LE): Exists between the *transmitter node* and *receiver node* of different physical nodes at the same modulation layer if there exists one path (note that we only consider k-shortest path between each node pair) satisfying the reach and spectrum constraints without regeneration (all the traversed edges have the same available spectrum slots).

Grooming Edge - (GE): Exists between the nodes in the electrical layer. *GE* exists if at least one lightpath with enough residual capacity to serve the request with traffic grooming.

Table 5. Weights of different types of edges.

Edge	PTE	ORE	LE	GE
Weight	c	2 c	$10^{-6} \cdot h \cdot s \cdot d/\overline{d}$	$2c - 10^{-3} + 10^{-9} \cdot h \cdot s \cdot d/\overline{d}$

B. Weight Assignment Policies

The weight of different types of edges is listed in Table 5, which is set to minimize the cost of serving requests as the main objective and minimize the spectrum occupation as the second objective. The weight of edges is determined by two factors: the first part is associated with the cost of transceivers, while the second part is associated with spectrum occupation. To prioritize cost reduction for transceivers, the weight assigned to the first part is significantly higher than that of the second part to minimize the number of transceivers and regenerators used. Specifically, the weight of *PTE* equals the cost of one transceiver, which models the cost of transmitters and receivers. In addition, the weight of *ORE* equals the cost of two transceivers to model the two transceivers used for optical regeneration. Unlike *PTE* and *ORE*, the weight *LE* and *GE* should incorporate the part for spectrum occupation. For an edge between node pair (i, j), an *LE* exists in one modulation layer if we can find a path p with available spectrum to connect the node pair (i, j) without regeneration. Assume that h denotes the number of hops of path p and s denotes the number of spectrum slots required of the corresponding level. The weight of *LE* is set to a small number (compared to the weight of *PTE* and *ORE* related to the cost of transceivers) related to the spectrum slots occupied to save spectrum as the second objective. In addition, assume the requested data rate is d and the data rate of the selected modulation format is $\overline{d}$, the number of slots in all the traversed edges $h \cdot s$ in the weight of *LE* is multiplied by $d/\overline{d}$ to give priority

to the modulation format with larger residual capacity for traffic grooming. Thus, establishing a new lightpath between two nodes requires two *PTE* and one *LE*, which has a total weight of $2c + 10^{-9} \cdot h \cdot s \cdot d/\overline{d}$. Instead of establishing a new lightpath, traffic grooming can also be used to serve a request. Since traffic grooming does not require transceivers, we set the weight of it to be smaller than the weight when establishing a new lightpath. In addition, the weight of *GE* is set to be a large number (slightly smaller than the cost of creating a new lightpath (i.e., $2c$) to make sure *GE* is only used when it can avoid establishing a new lightpath, and hence reduce the number of transceivers.

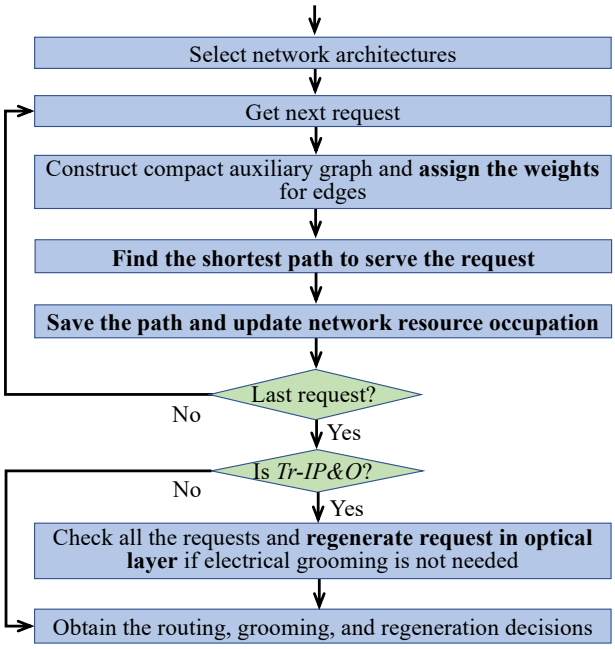

Fig. 5. Flowchart of the proposed CAG-ND algorithm.

C. Compact Auxiliary-Graph-Based Algorithm

The proposed auxiliary-graph-based heuristic algorithm is illustrated with the flowchart in Fig. 5. For all the network architectures, to serve each request, the algorithm first constructs a CAG and serves the request with the shortest path. Since optical and IP layer regeneration have the same cost of transceivers, *Tr-IP&O* adopts the same auxiliary graph for *Tr-IP* to prioritize IP layer regeneration, and hence residual capacity can be reserved for future traffic grooming. After serving all the requests, *Tr-IP&O* checks all the lightpath and converts IP layer regeneration to optical-layer regeneration if electrical grooming is not needed.

The construction of CAGs and the procedure for serving requests are illustrated as follows. All the CAGs only contain one layer, and each node pair is connected with at most one edge, representing the most cost-effective way to connect the node pair. Specifically, regarding *Op-IP*, *Tr-IP*, and *Tr-IP&O*, for each node pair, the algorithm will check the weight of grooming (weight of *GE*) and the weights of establishing new lightpath in all *modulation layers*, and select the one with minimum weight as the weight for the edge. Instead, for *Tr-O*, we cannot use different modulation formats in a single CAG because the data rates of different modulation formats are different, which makes it infeasible to change the modulation format while performing optical-layer regeneration. Thus, we construct a CAG for each

Algorithm 1. Construction of compact auxiliary graph.

Input: requested data rate d_i for request i
Output: compact auxiliary graph for request i

Initialize a topology without edges
for each node pair p **do**
  Set weight $w_p = Inf$
  if Tr-O and end-to-end grooming available **then**
    if End nodes of i is same as p **then**
      Add edge between p and set weight according to *GE*
      Continue
  if not *Tr-O* and exist LP l between p with $d_l > d_i$ **then**
    Select the LP l_s with the lowest weight
    Add edge between p and log the LP l_s
    Set edge weight as weight of using LP l_s as *GE*
  else
    for each shortest path j between p **do**
      if *Op-IP* and path j has more than one edge **then**
        Continue
      for each layer with modulation m **do**
        if Distance of path j > reach of m **then**
          Continue
        if *Tr-O* **then**
          if Two nodes of p are ends nodes of i **then**
            Set w_c as weight of *PTE+LE+PTE*
          else if One node of p is end nodes of i **then**
            Set w_c as weight of *PTE+LE+1/2+ORE*
          else
            Set w_c as weight of *1/2ORE+LE+1/2ORE*
        else
          Set w_c as weight of *PTE+LE+PTE*
        Set $w_p = w_c$ if $w_p > w_c$ and log j and m
    Add edge between p and set the weight as w_p

modulation format layer, and the shortest path among these CAGs will be selected.

Now let us illustrate the construction of CAGs using ZR/ZR+ with an example in Fig. 4 (b)-(e). Assume that the network already served a request requiring $100Gb/s$ of traffic between node pair $(1,3)$, and we need to serve another request between node pair $(1,4)$ requiring $200Gb/s$ of traffic. The lightpaths used to serve these two requests are denoted with two blue dotted lines in Fig. 4 (a). These lightpaths are mapped to the red dotted line in the layered auxiliary graph in Fig. 4 (b). For *Op-IP*, as shown in Fig. 4 (b), the request between node pair $(1,3)$ is served with the red dotted path $(1, T1, R2, 2, T2, R3, 3)$, in which the lightpaths between node pair $(1,2)$ and $(2,3)$ both utilize 16QAM and have $300Gb/s$ of residual capacity. Fig. 4 (c) shows the CAG for *Op-IP*. The edges between node pair $(1,2)$ and $(2,3)$ are constructed from *GE* (traffic grooming), and the edges between node pair $(3,4)$ are constructed from *PTE+LE+PTE* (establishing a new lightpath). Unlike *Op-IP*, all the transparent architectures serve the request between node pair $(1,3)$ with the red dotted path $(1, T1, R3, 3)$ by bypassing node 2. The auxiliary graphs for *Tr-IP* and *Tr-IP&O* do not consider optical-layer regeneration at this stage since IP layer regeneration and optical-layer regeneration have the same cost of transceivers. As shown in Fig. 4 (d), only the node pair $(1,3)$ has an edge constructed from *GE* with traffic grooming. In this case, path $(1,4)$ is selected to serve request $(1,4)$ instead of using traffic grooming in path $(1,3,4)$ since two paths have the same cost of transceivers and the residual capacity between node pair $(1,3)$ can be used to serve future requests (e.g., request between $(1,3)$). Lastly, the auxiliary graph for *Tr-O* is shown in Fig. 4 (e). In this case, we only consider the *PTE* related to the source node and destination node since IP layer regeneration is not allowed. In addition, the weight of *ORE* is modeled by adding

the weight of $1/2ORE$ in edges connected to intermediate nodes. The procedure of creating the compact auxiliary graph is also illustrated with Algorithm 1, which creates an edge between each node pair based on grooming and different modulation layers in Fig. 4 (b). Since the weight of *GE* is smaller than the weight of traversing in other layers, Algorithm 1 first creates an edge for node pair p if grooming is available for *Tr-O* (line 4-6) and other architectures (line 8-11). Otherwise, Algorithm 1 loops over each shortest path and each layer and calculates the weight of connecting p for *Tr-O* (line 19-25) and other architectures (line 27). Then, the most cost-effective path is selected to connect node pair p and set as the weight of the edge between node pair p (line 27-28).

The procedures for serving requests are as follows. For *Op-IP*, *Tr-IP*, and *Tr-IP&O*, the CAG-ND algorithm first constructs one CAG using the electrical layer and all the modulation layers for each request and then finds the shortest path between the end nodes of the request. Instead, for *Tr-O*, the CAG-ND algorithm first constructs several CAGs, where each auxiliary graph is constructed with one modulation layer and the electrical layer. Then, the CAG-ND algorithm finds the shortest paths among all the CAGs and selects the path with the smallest weight. The obtained shortest path is saved and utilized to update the resource occupation (i.e., spectrum occupation, the residual capacity of channels, and traffic in routers). For *TR-IP&O*, after obtaining all the regeneration performed in the IP layer by *Tr-IP*, the CAG-ND algorithm identifies the necessary IP layer regenerations, specifically those associated with traffic grooming to reduce the number of transceivers in the optical layer. The algorithm then attempts to replace these IP-layer regenerations with optical-layer regenerations. Specifically, for each request, if the optical signal is only regenerated in the IP router without electrical grooming, the CAG-ND algorithm chooses to regenerate the signal in the optical layer instead of the IP layer. Finally, the CAG-ND algorithm makes the routing, grooming, and regeneration decisions.

Complexity of the algorithm: The complexity of the CAG-ND algorithm for *Op-IP*, *Tr-IP*, and *Tr-O* is mainly determined by constructing the CAG and finding the shortest path in the auxiliary graph, while *Tr-IP&O* also needs to consider the complexity of replacing IP layer regenerations with optical-layer regenerations. The complexity of replacing IP layer regenerations with optical-layer regenerations can be neglected since it only needs to check the found path for each request, which is less complex than finding the shortest path. Assume that the number of physical nodes and edges in the network is $|N_p|$ and $|E_p|$, respectively, and the number of spectrum slots in each edge is $|S|$. In addition, the number of modulation layers and the number of predefined shortest paths between each node pair are denoted with M and K, respectively. For *Op-IP*, *Tr-IP*, and *Tr-O*, for each request, we need to construct only one CAG and the complexity of constructing nodes is $O(|N_p|)$. Since the maximum number of lightpaths is $O(|E_p||S|/s_{min})$, where s_{min} is the minimum number of required slots among all modulation formats. To construct links, the complexity of building links with *GE* and the modulation layer is $O(|E_p||S|/s_{min})$ and $O(KM|S||N_p|^2)$, respectively. For *Tr-IP&O*, for each request, we need to construct M CAG and the complexity of constructing nodes and links in all the auxiliary graphs is $O(M|N_p|)$ and $O(K|E_p||S|/s_{min} + KM|S||N_p|^2)$, respectively. Hence, the complexity of constructing the CAG is $O(KM|S||N_p|^2)$ for all the network architectures. Since the number of nodes and links in the CAG is $O(|N_p|)$ and $O(|N_p|^2)$, respectively, for *Op-IP*, *Tr-IP*, and *Tr-O*, the complexity of finding the shortest path in the constructed auxiliary graph is $O(|N_p|^2 + |N_p|log(|N_p|))$. Instead, for *Tr-O*, the complexity of finding the shortest path in all the constructed auxiliary graphs is $O(M|N_p|^2 + M|N_p|log(|N_p|))$. Compared to previous auxiliary-graph-based algorithms that incorporate multiple layers for different modulation formats (as, e.g., [32]), our version of the auxiliary-graph-based algorithms leads to a reduction of computational complexity for finding shortest path. Specifically, prior works employed networks with a larger number of nodes in the auxiliary graph, specifically $2M|N_p| + |N_p|$, which results in a higher complexity of $O(|N_p|^2(2M+1)^2 + (2M+1)|N_p|log((2M+1)|N_p|)$.

4. NUMERICAL RESULTS

A. Simulation Settings

We perform numerical evaluations with a custom-built Python simulator, working for all the network architectures introduced in Sec. 2.A. We first evaluate performance on two distinct network topologies, namely, Japan topology (J14) [34] and NSF topology (N14) [35]. These two topologies have same number of nodes and are characterized by short and long link lengths, with average values of 166 km and 1081 km, respectively. Then we evaluate performance on topology with a larger number of nodes, namely, 50-node German topology (G50) [36] with 42 edge nodes and 8 core nodes and an average link length of 100 km. Each fiber has a capacity of 6-THz (C-band transmission). The normalized cost of ZR, ZR+, and long-haul muxponders is 1, 2, and 4, respectively. For J14, we assume that the requested data rate of requests is randomly selected from 100, 200, 300, 400 Gb/s, while for N14, we assume that it is randomly selected from 100, 200, 300 Gb/s as the transmission distance with data rate of 400 Gb/s for ZR/ZR+ is shorter than the maximum link length in N14, and hence requests with data rate greater than 300 Gb/s may be rejected even with enough spectrum resources. To evaluate the generalizability of the results under different traffic distributions, for G50, we evaluate the performance of cost and power consumption under a hub-and-spoke traffic pattern as in Ref. [36] among core nodes and edge nodes. Specifically, requests are generated with a mix of 67% edge-core (between the edge node and the two closest core nodes) and 33% core-core. The requested data rate is selected from 100, 200, 300, 400 Gb/s with a probability of 20%, 25%, 30%, 25%. We first evaluate the rejection probability of different network architectures using ZR/ZR+ pluggable optics and long-haul muxponders under different network loads. Specifically, the rejection probability is determined by computing the average across all evaluated instances. This evaluation offers insights into the tradeoff between energy efficiency and network capacity of different network architectures. Then we compare the cost and power consumption of ZR/ZR+ pluggable optics and long-haul muxponders under 1% rejection rate. Note that since different network architectures have different carried traffic, we compare the cost and power consumption per Tb/s (normalized by the total carried traffic) to have a fair comparison. Note that we perform the comparison between the various architectures at 1% rejection rate to have a comparison in terms of energy required per transported bit. We remark that performing a comparison under the same traffic load would not reflect the benefits of long-haul muxponders, as they have a much larger data rate. Comparing under the same traffic load would lead to resource underutilization of long-haul muxponders, resulting in residual capacity waste and increased

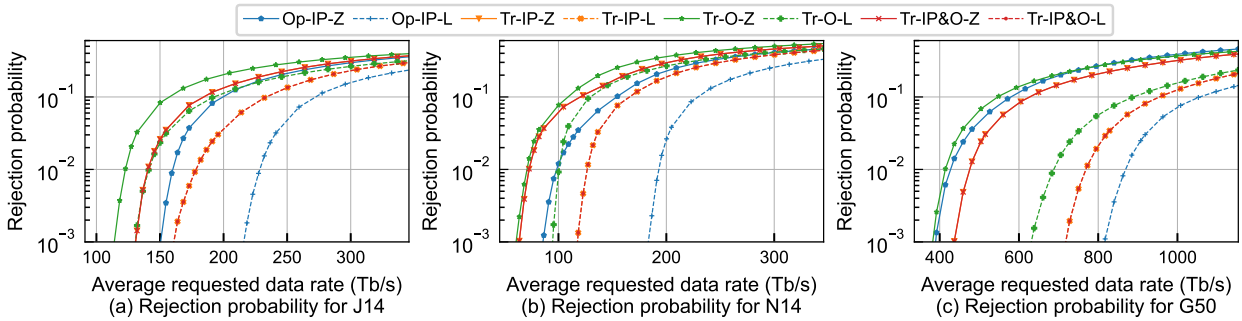

Fig. 6. Rejection probability vs. loads for J14, N14, and G50.

power consumption per bit. To clarify this, we also compare the power consumption per Tb/s under the same traffic load for J14, specifically when the traffic for *Tr-O-Z* results in less than 1% rejection, as *Tr-O-Z* can accommodate the least amount of traffic among all the considered architectures for J14. All the results are averaged over 100 traffic instances, achieving confidence intervals for rejection probability, with a 95% of confidence level, within an absolute difference of ±0.38% for all the evaluated requested data rates.

B. Evaluation of Rejection Probability

We first evaluate the rejection probability of different architectures on J14 in Fig. 6(a). The architectures with ZR/ZR+ and long-haul muxponders are denoted with suffix *-Z* and *-L*, respectively. Results show that *Op-IP* achieves the lowest rejection probability. This is because each intermediate node regenerates signals, resulting in a shorter transmission distance than transparent architectures and allowing the LP to operate at a higher modulation format. In addition, *Op-IP* regenerates signals at the IP router of each intermediate node, and hence has the largest capability of traffic grooming and does not have spectrum contiguity constraint, leading to lower spectrum occupation and hence lower rejection probability. Instead, *Tr-O* has the highest rejection probability for both ZR/ZR+ and long-haul muxponders since it can only perform end-to-end grooming. For instance, under 1% of rejection probability, with ZR/ZR+, *Op-IP* carries on average 158 Tb/s total traffic in the network, while *Tr-O* carries 23% less traffic than *Op-IP*, with a total traffic of 121 Tb/s. The rejection rate of *Tr-IP* and *Tr-IP&O* is the same since to obtain the routing decisions of *Tr-IP&O*, the CAG-ND algorithm only replaces unnecessary IP regeneration of *Tr-IP* with optical regeneration, which does not affect spectrum occupation and the carried traffic as described in Sec. 3. For long-haul muxponders, the performance of different network architectures keeps the same trend as that for ZR/ZR+. Regarding the performance difference between ZR/ZR+ and long-haul muxponders, for the same network architecture, ZR/ZR+ always has a higher rejection probability than long-haul muxponders under the same average requested data rate since long-haul muxponders have higher spectral efficiency. Specifically, under 1% rejection probability, by using long-haul muxponders, *Op-IP* can carry 43% additional traffic compared to *Op-IP* with ZR/ZR+, achieving maximum network capacity (total carried traffic under 1% of rejection rate) of 225 Tb/s. Instead, transparent architectures can only carry up to 26% additional traffic with long-haul muxponders since transparent architectures tend to use lightpaths with a larger number of hops, resulting in a smaller data rate, and hence smaller improvement in carried traffic.

We now compare the rejection probability of different architectures on N14 as shown in Fig. 6(b). Compared to J14, all network architectures have a similar trend of rejection probability but start to have rejected requests earlier. We then discuss the performance differences between ZR/ZR+ and long-haul muxponders. Compared to a small topology, the additional carried traffic of long-haul muxponders compared to ZR/ZR+ increases from up to 43% to 92% for *Op-IP* and from 26% to 75% for transparent architectures because the advantage of longer reach and higher data rate of long-haul muxponders is more evident in a larger topology.

The rejection rate for G50 is shown in Fig. 6(c). The overall trend for all transparent architectures remains the same as in J14 and N14. The difference is that with ZR/ZR+, opaque architecture has a higher rejection rate than *Tr-IP* and *Tr-IP&O*. This is because, for G50, most of the links are within the maximum reach of ZR, making it preferable compared to ZR+ due to its lower cost. Consequently, a larger number of ZR is used in G50, thereby resulting in a larger spectrum occupation (note that one ZR occupies at least 25 GHz more spectrum than ZR+ as in Table. 3). Moreover, the maximum network capacity increases compared to J14 and N14 due to shorter link lengths in G50, achieving 465 and 862 Tb/s for ZR/ZR+ and long-haul muxponders, respectively.

C. Evaluation of Cost and Power Consumption for J14

We first compare the number of regenerators and the percentage of groomed lightpaths under 1% rejection probability to interpret how the requests are served. Then, we compare the cost of the network in terms of the number of transceivers and the number of interfaces (i.e., the number of router ports and I/O cards for ZR/ZR+ and long-haul muxponders, respectively). At last, we analyze the power consumption of the network.

Fig. 7 (a) shows the averaged number of regenerators per request in IP layer (the cases with ZR/ZR+ and the case with long-haul muxponders are named as *IP-Z* and *IP-L*, respectively) and optical layer (the cases with ZR/ZR+ and the case with long-haul muxponders are named as *O-Z* and *O-L*, respectively) for J14. First, regarding the number of regenerators, *Op-IP* with ZR/ZR+ requires around one regenerator per request, and transparent architectures with ZR/ZR+ require only around 0.06 regenerator per request due to optical bypassing. Note that the maximum number of regenerators per request is only around one because several requests may be groomed and share the same regenerator. By enabling both optical regeneration and IP layer regeneration, *Tr-IP&O* achieves a significant reduction of 60% in IP layer regenerators compared to *Tr-IP*. This reduction not only saves the number of router interfaces but also

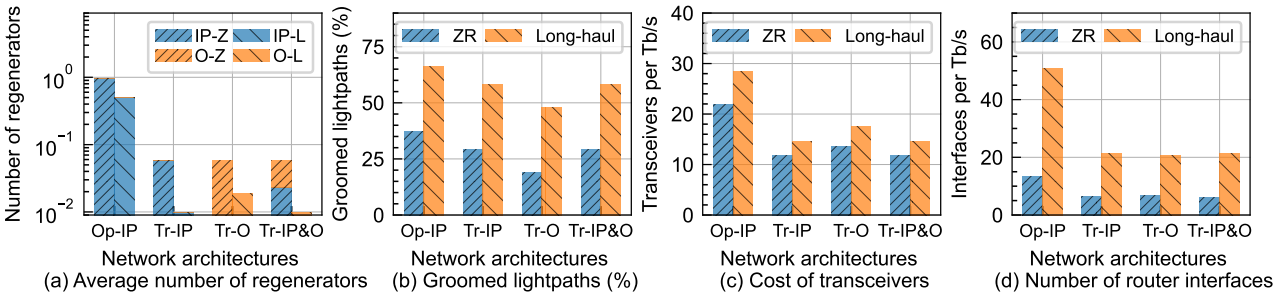

Fig. 7. Regenerators, groomed lightpaths, cost of transceivers and router interfaces under 1% rejection probability for J14.

decreases the traffic to be processed in routers. For long-haul muxponders, the number of regenerators of *Op-IP* reduced to around 0.5 per request, and the regenerators reduced to only around 0.01 per request for transparent architectures due to the longer reach of long-haul muxponders and more considerable grooming capabilities (one regenerator can be shared by more requests). Fig. 7 (b) shows the percentage of groomed lightpaths, defined as the number of lightpaths served with grooming divided by the total number of established lightpaths. First, the percentage of groomed lightpaths of transparent architectures is up to 49% lower than that of opaque architecture for ZR/ZR+. Moreover, among the transparent architectures, since *Tr-O* only allows end-to-end grooming, *Tr-O* with ZR/ZR+ has the lowest percentage of groomed lightpaths, which is around 35% lower than other transparent architectures with ZR/ZR+. When using long-haul muxponders, the percentage of groomed lightpath of *Op-IP* increases 77% compared to that with ZR/ZR+, and the increment for transparent architectures can be up to 153% since transparent architectures tend to have paths with a longer path due to bypassing and hence the advantage in grooming with long-haul muxponders becomes more evident.

Fig. 7 (c) shows the cost of transceivers for J14. For ZR/ZR+, transparent architectures reduce at most 46% of transceivers of *Op-IP* due to optical bypassing as in Fig. 7 (a). Moreover, with ZR/ZR+, the cost of transceivers of *Tr-IP* and *Tr-IP&O* with ZR/ZR+ is 13% lower than *Tr-O* since *Tr-IP* and *Tr-IP&O* have a higher grooming capability (i.e., electrical grooming) as shown in Fig. 7 (b), which reduces the number of required transceivers. In addition, compared to ZR/ZR+, the normalized cost of transceivers per Tb/s with long-haul muxponders is around 30% higher in opaque and 29% higher in transparent architectures, due to the higher cost of a single long-haul muxponder relative to a ZR/ZR+. We then compare the number of router interfaces used for different network architectures in J14 as shown in Fig. 7 (d). With ZR/ZR+, transparent architectures reduce at most around 55% of router ports of *Op-IP* with optical bypassing and optical-layer regeneration. Among all the transparent network architectures with ZR/ZR+, *Tr-IP&O* requires the least router ports. Specifically, *Tr-IP&O* reduces 5% of router ports of *Tr-IP* by optical-layer regeneration. In addition, *Tr-IP&O* reduces 9% of router ports of *Tr-O* due to higher grooming capabilities. The abovementioned reduction of transceivers and interfaces in transparent architectures is achieved with 88 I-ROADMs in J14. Unlike ZR/ZR+, long-haul muxponders use I/O cards as router interfaces, and each router port requires the capacity of four I/O cards. For *Op-IP*, the number of I/O cards with long-haul muxponders is 3.8 times that of router ports with ZR/ZR+, smaller than 4 times, indicating that long-haul muxponder saves the carried traffic in routers compared to the cases with ZR/ZR+. Instead, for transparent architectures, the number of I/O cards is only up to 3.5 times that of router ports with ZR/ZR+ because long-haul muxponders have a longer transmission distance and can reduce more interfaces through bypassing compared to ZR/ZR+. Lastly, with long-haul muxponders, all the transparent architectures consume almost the same number of I/O cards since almost all the network architectures have a very small number of regenerators, as in Fig. 7 (a), resulting in interfaces are mainly used in the source and destination nodes (identical for different architectures).

ZR/ZR+ can not only reduce the cost of transceivers per Tb/s as in Fig. 7 (c) but also reduce the power consumption per Tb/s as in Fig. 8 compared to long-haul muxponders for J14. Fig. 8 shows the power consumption (PC) of different network elements with different percentages of reduction of power consumption of IP routers. Let us first discuss the power consumption for different components (i.e., transceiver, IP router, and optical layer) for different network architectures, as shown in Fig. 8 (a). The power consumption of transparent architectures is up to 48% lower than opaque architectures for ZR/ZR+, as transparent architectures reduce the average number of regenerators from 1 to 0.06 with optical bypassing, and hence reduce the power consumption for IP routers and transceivers. Specifically, compared to the opaque architecture scenario, transparent architecture results in power savings of up to approximately 53% for IP routers and 50% for transceivers in ZR/ZR+ configurations, due to the reduction of IP and optical layer regenerators, as shown in Fig. 7 (a). Among all the network architectures, *Tr-IP&O* achieves the lowest normalized power consumption, which is 3% lower than *Tr-IP* due to power consumption savings in routers with optical-layer regeneration. Since more than half of the IP-layer regenerations of *Tr-IP* have been reduced by *Tr-IP&O*, even optimal solution (obtained, e.g., with ILP) with fewer IP-layer regenerations could not yield significant additional power savings. In addition, *Tr-IP&O* can reduce around 11% of power consumption of *Tr-O* due to the power consumption savings in both transceivers and routers with larger grooming capability. Instead, if we use long-haul muxponders, the total power consumption increases up to 53% and 47% for opaque and transparent architectures, respectively, due to the increment of power consumption with long-haul muxponders (note that the normalized power consumption of one long-haul muxponder is 8, while the normalized power consumption of one ZR+ is only 1.3). The power consumption reduction of *Tr-IP&Op* compared to *Tr-IP* decreases from 3% for ZR/ZR+ to nearly 0 for long-haul muxponders. This occurs because *Tr-IP&Op* cannot achieve further reductions in power consumption through optical-layer regeneration, as the number of regenerators of *Tr-IP* is almost 0. Moreover, the power consumption reduction of *Tr-IP&Op*

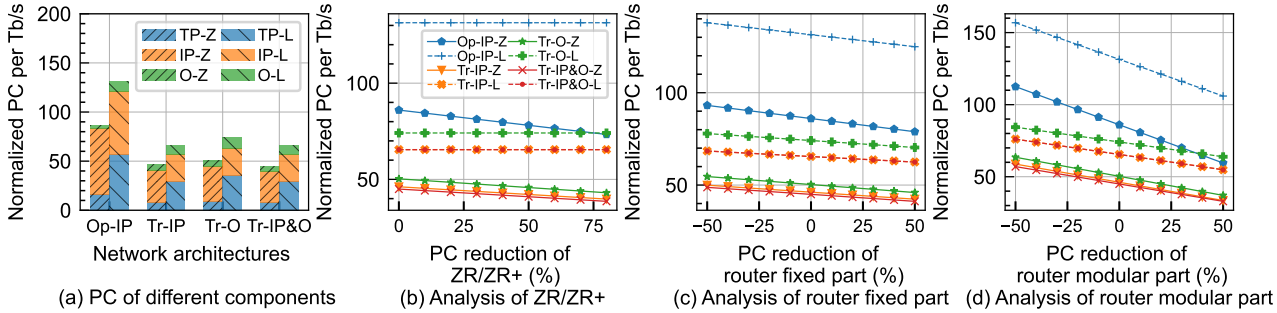

Fig. 8. Power consumption under 1% rejection rate when varying power consumption of routers and ZR/ZR+ for J14.

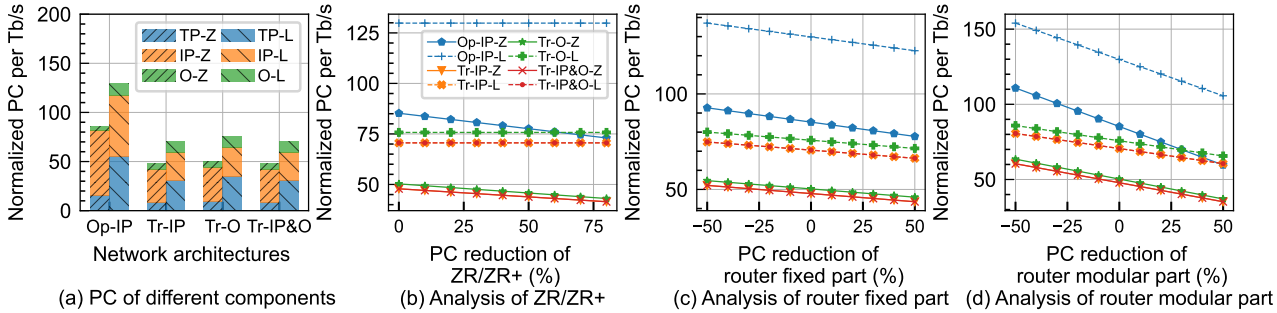

Fig. 9. Power consumption under same traffic load when varying power consumption of routers and ZR/ZR+ for J14.

compared to *Tr-O* increases from 11% for ZR/ZR+ to 12% for long-haul muxponders. This is because long-haul muxponders have a larger data rate, leading to more traffic being groomed to reduce power consumption.

We now perform sensitivity analysis on power consumption. Since the power consumption of the optical layer only consumes less than 12% of the total power consumption, the power consumption of transceivers and IP routers are the dominating factors for the total power consumption. Thus, we only perform sensitivity analysis on transceivers and routers. Since ZR/ZR+ technology is still evolving, the power consumption of ZR/ZR+ is expected to decrease with technological advances. Thus, we perform the sensitivity analysis of ZR/ZR+ by varying (reducing) the power consumption of ZR/ZR+ by up to 80% to evaluate the changes in total power consumption of different architectures. As shown in Fig. 8 (b), when using ZR/ZR+, varying the power consumption of ZR/ZR+ does not alter the relative performance among different architectures significantly. For ZR/ZR+, when the power consumption of ZR/ZR+ reduces by 80%, the total power consumption of both opaque and transparent architectures reduces by around 12% and 13%, respectively. Moreover, the advantage of power consumption with ZR/ZR+ compared to that with long-haul muxponders increases with a higher reduction of power consumption of ZR/ZR+. Specifically, when the power consumption reduction of ZR/ZR+ increases to 80%, the advantage of total power consumption of ZR/ZR+ compared to long-haul muxponders increases from up to around 26% to 35% for opaque architecture and 32% to 41% for transparent architectures, respectively.

We now conduct a sensitivity analysis on router power consumption, which is crucial due to potential variations from adopting routers of different capacities or from various vendors, as well as technological advances. We vary the power consumption as input from -50% to 50% of the current power consumption for both the fixed and modular parts of the router. This analysis offers insights into the impact of different parts of the router for different network architectures. As shown in Fig. 8 (c) and (d), when varying the power consumption of router fixed part and router modular part by up to 50%, the relative power consumption among different architectures remains the same. With a higher reduction of power consumption in IP routers, the architectures with ZR/ZR+ modules become increasingly power-efficient than the architectures with long-haul muxponders. For instance, with a 50% reduction in power consumption of router fixed part, the additional total power consumption with long-haul muxponders compared to ZR/ZR+ increases from around 26% to 30% and 32% to 35% for opaque and transparent architectures, respectively. This is because, with a higher reduction in the cost of routers, the power consumption savings in transceivers with ZR/ZR+ become more evident. Moreover, as shown in Fig. 8, the variation in total power consumption is more significant when reducing the router modular part. For instance, with ZR/ZR+, reducing the router's fixed part by 50% results in only a 9% reduction in the total power consumption of *Tr-IP*, as shown in Fig. 8(c). However, when the router's modular part is reduced by 50%, the total power consumption of *Tr-IP* decreases by 31%.

In the following, we evaluate the power consumption under the same traffic load for J14. As shown in Fig. 9, the power consumption trends of different architectures remain consistent with those observed under a 1% rejection rate. However, a key difference is that long-haul muxponders become more power-hungry than ZR/ZR+ under the same traffic load. Notably, for Tr-IP, the additional power consumption with long-haul muxponders compared to ZR/ZR+ increases from 42% under the 1% rejection rate to 47% under the same traffic load. This is because long-haul muxponders are more underutilized under the same traffic that ZR/ZR+ can accommodate.

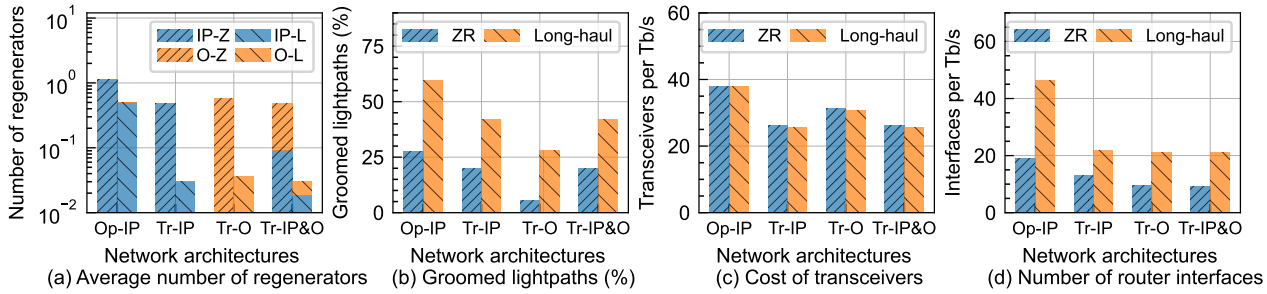

Fig. 10. Regenerators, groomed lightpaths, cost of transceivers and router interfaces under 1% rejection probability for N14.

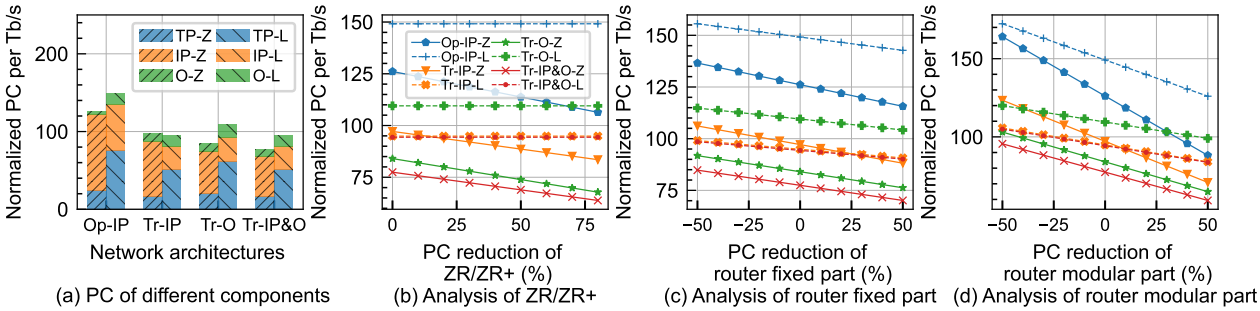

Fig. 11. Power consumption under 1% rejection rate when varying power consumption of routers and ZR/ZR+ for N14.

D. Evaluation of Cost and Power Consumption for N14

Long-haul muxponders in a larger continental-wide topology, such as the N14, show more advantages compared to ZR/ZR+ in terms of both cost and power consumption.

We first compare the average number of regenerators in Fig. 10 (a). Firstly, the reduction of regenerators in transparent architectures compared to opaque architectures decreases because more regenerators need to be placed in transparent architectures due to longer paths in N14, and less requests can share the same regenerator due to a smaller data rate. For instance, for ZR/ZR+, the reduction of regenerators in transparent architectures is reduced from up to 94% in J14 to up to 58% in N14. Moreover, the number of IP-layer regenerators in *Tr-IP* is reduced by 81% with Tr-IP&O for N14, showing the effectiveness in reducing the power consumption with the proposed CAG-NG algorithm. Regarding the percentage of groomed lightpaths, the overall trend between different architecture is still the same as in a small topology, as shown in Fig. 7 (b). The only difference is that the percentage of groomed lightpaths is smaller than that in a small topology.

Unlike the case in J14, long-haul muxponder becomes comparable with ZR/ZR+ in terms of both the cost of transceivers and the number of interfaces. Regarding the cost of transceivers, as shown in Fig. 10 (c), for opaque architectures, the additional cost of long-haul muxponders reduces from 30% to less than 1%. Moreover, for transparent architectures, the cases with long-haul muxponders even have up to 3% reduction of the cost of transceivers. This is because long-haul muxponders have a longer transmission distance than ZR/ZR+, resulting in less transceivers used in a large topology. Regarding the number of router interfaces, as shown in Fig. 10 (d), among three transparent architectures with ZR/ZR+, *Tr-IP&O* achieves the lowest number of interfaces per Tb/s. Notably, the reduction of router ports of *Tr-IP&O* compared to *Tr-IP* increases from around 5% in J14 to 31% in N14 because optical-layer regeneration plays a more important role when the number of regenerations increases in a large topology, as shown in Fig. 10 (a). The reduction of transceivers and interfaces in transparent architectures, as previously mentioned, is accomplished using 84 I-ROADMs in N14. Moreover, the ratio of router ports with ZR/ZR+ to the number of I/O cards with long-haul muxponders increases from 26% in J14 to 41% in N14 for Op-IP, and from 32% in J14 to 61% in N14 for transparent architectures. This is because ZR/ZR+ requires 400Gb/s ports but mainly only operates at a low data rate (less than 300 Gb/s for most paths), which increases the number of router ports required per Tb/s.

Let us now evaluate the power consumption of different architectures for N14. As shown in Fig. 11 (a), similar to the results for J14, transparent architecture is still more power-efficient (up to 39%) than opaque architectures due to a smaller number of IP layer regeneration and hence less power consumption in IP routers and transceivers, which are still dominant factors of power consumption. Among all the transparent architectures, *Tr-IP&O* with ZR/ZR+ has around 8% and 21% lower power consumption than *Tr-O* and *Tr-IP* due to higher grooming capability and optical layer regeneration, respectively, as shown in Fig. 10. Unlike a smaller topology, long-haul muxponders start to show compelling performance in power consumption in N14. For instance, as shown in Fig. 11 (a), the power consumption of *Op-IP* with long-haul muxponders is only 18% higher than that with ZR/ZR+, and the power consumption of *Tr-IP* with long-haul muxponders is even 2% lower than that with ZR/ZR+. This is because, in a large topology, ZR/ZR+ is likely to operate at less than 300 Gb/s data rate but consume 400 Gb/s ports, resulting in increased router ports as in Fig. 10 (d) and hence increased power consumption per Tb/s.

The sensitivity analysis for power consumption of ZR/ZR+ is shown in Fig. 11 (b). Unlike J14, where the power consumption with long-haul muxponders is always lower than that with ZR/ZR+, for *Tr-IP*, the total power consumption with long-haul

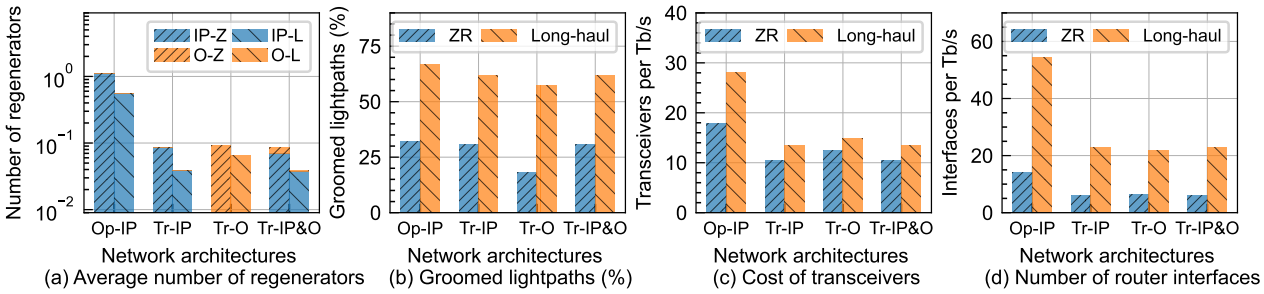

Fig. 12. Regenerators, groomed lightpaths, cost of transceivers and router interfaces under 1% rejection probability for G50.

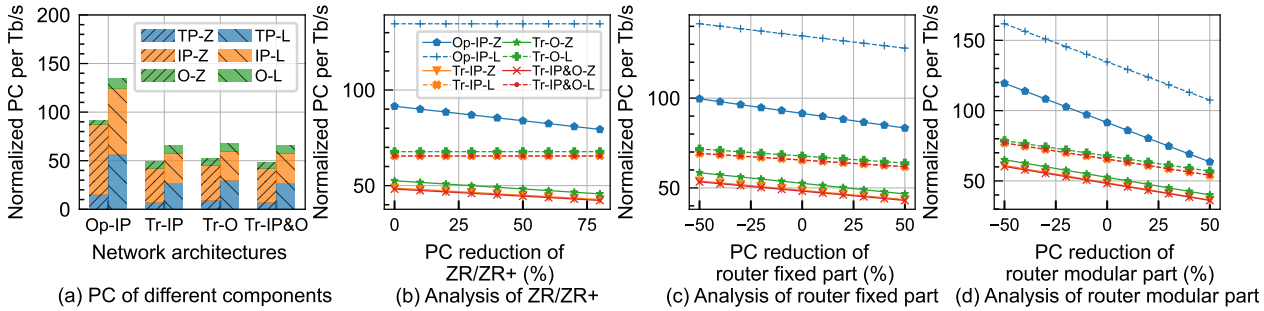

Fig. 13. Power consumption under 1% rejection rate when varying the power consumption of routers and ZR/ZR+ for G50.

muxponders is lower than that with ZR/ZR+ when the power consumption reduction is less than 14%. When the power consumption reduction of ZR/ZR+ increases to 80%, the total power consumption with ZR/ZR+ becomes up to 32% and 38% lower than that with long-haul muxponders for opaque and transparent architectures, respectively. Regarding the sensitivity analysis of routers, as shown in Fig. 11 (c) and Fig. 11 (d), the relative performance among different network architectures remains the same as in J14 when varying the power consumption of router fixed part and router modular part. In addition, the relative performance between ZR/ZR+ and long-haul muxponders also remains the same as in J14 for all network architectures except *Tr-IP*. Specifically, for *Tr-IP*, the total power consumption using long-haul muxponders is lower than that using ZR/ZR+ only when the reduction of power consumption of the router fixed part is less than 24%. When the power consumption reduction of the router fixed part increases to 50%, the total power consumption with ZR/ZR+ is 3% lower than that with long-haul muxponders. Moreover, the impact of variation in power consumption of the router modular part is more evident. Specifically, the power consumption of ZR/ZR+ is less than that of long-haul muxponders as long as the reduction in power consumption for the router modular part remains below 7%. Moreover, a reduction of 50% in the power consumption of router modular part makes ZR/ZR+ 16% more power-efficient than long-haul muxponders. Conversely, if the power consumption of the router modular part increases by 50%, ZR/ZR+ becomes 17% less power-efficient than long-haul muxponders.

E. Evaluation of Cost and Power Consumption for G50

In our evaluation with G50, we observe consistent trends compared to those in the smaller national-wide J14 topology in terms of both cost and power consumption, as shown in Fig. 12 and Fig. 13, respectively. Since G50 has more nodes, the number of I-ROADMs introduced for transparent architectures increases from 88 in J14 to 352 in G50. Notably, *Tr-IP&O* has almost no optical-layer regeneration (See Fig. 12 (a)). This is mainly because most traffic originates between the edge nodes and adjacent core nodes, resulting in more traffic being groomed and hence using electrical regeneration rather than optical regeneration. Consequently, with almost no optical-layer regeneration, *Tr-IP&O* shows similar cost and power consumption as *Tr-IP*. Specifically, with ZR/ZR+, *Tr-IP&O* is 1% more power-efficient than *Tr-IP*, as shown in Fig. 13 (a), due to optical-layer regeneration, and 8% more power-efficient than *Tr-OP* due to higher grooming capability. Regarding the sensitivity analysis of ZR/ZR+ and routers in G50, as shown in Fig. 13 (b)-(d), the relative trend among different architectures, as well as the trend between ZR/ZR+ and long-haul muxponders, does not change, as in the other national-wide topology, namely, J14.

5. CONCLUSION

Even in the presence of the drastic reduction of power consumption of the routers and ZR/ZR+ seen in last years, the opaque network architecture is still more power-hungry than the transparent architectures employing optical bypass. We demonstrated that using optical bypassing saves up to 48% and 39% in terms of power consumption compared to opaque IPoWDM in a national-wide topology as J14 and a continental-wide topology as N14. Moreover, in a national-wide topology, ZR/ZR+ can reduce up to 35% of power consumption together with around 23% reduction in the cost of transceivers compared to long-haul muxponders. Instead, in a continental-wide topology, long-haul muxponders can be more energy-efficient than ZR/ZR+ unless considerable power savings are achievable for routers. In the future, we plan to investigate the energy efficiency achievable in optical networks equipped with ZR/ZR+ devices when re-routing based on traffic demand intensity is allowed.

FUNDING

Ministero dell'Università e della Ricerca; Horizon 2020 Framework Programme (101016); Italian National Recovery and Resilience Plan (PE00000001).

ACKNOWLEDGMENTS

This work was supported by Italian Ministry of University and Research (MUR) and the European Union (EU) under the PON/REACT. This work was also supported by EU Horizon 2020 B5G-OPEN Project (101016). Moreover, this work was supported by the European Union -Next Generation EU under the Italian National Recovery and Resilience Plan (NRRP), Mission 4, Component 2, Investment 1.3, CUP D43C22003080001, partnership on "Telecommunications of the Future" (PE00000001 - program "RESTART").

REFERENCES

1. M. Radovic, A. Sgambelluri, F. Cugini, and N. Sambo, "Power-aware high-capacity elastic optical networks," IEEE/Optica J. Opt. Commun. Netw. **16**, B16–B25 (2024).
2. T. Zami and B. Lavigne, "Optimal deployments of 400 gb/s multihaul cfp2-dco transponders in transparent ipowdm core networks," in *Optical Fiber Communications Conference and Exhibition (OFC),* (2022).
3. Q. Zhang, A. Morea, and M. Tornatore, "A pragmatic power-consumption analysis for ipowdm networks with zr/zr+ modules," in *European Conference on Optical Communication (ECOC),* (2022).
4. B. Zhu, R. Lingle, and D. J. DiGiovanni, "High Capacity 400Gb/s Real-time Transmissions over AllWave ULL Fibres by 400ZR/ZR+ Pluggable Modules," in *European Conference on Optical Communication (ECOC),* (2021).
5. F. Musumeci, M. Tornatore, and A. Pattavina, "A power consumption analysis for ip-over-wdm core network architectures," IEEE/Optica J. Opt. Commun. Netw. **4**, 108–117 (2012).
6. Cisco, "Cisco 8000 Series Routers Data Sheet," https://www.cisco.com/c/en/us/products/collateral/routers/8000-series-routers/datasheet-c78-742571.html, accessed on 13/07/2023.
7. P. Tomlinson, G. Hill, and A. Tzanakaki, "Comparison of transparent and opaque optical transport network designs," in *European Conference on Optical Communication (ECOC),* (2002).
8. M. Gagnaire, "From opacity to transparency via translucent optical networks," in *Networks 2008-The 13th International Telecommunications Network Strategy and Planning Symposium,* (IEEE, 2008), pp. 1–77.
9. W. Fawaz, "Improved edf-based management of the setup of connections in opaque and transparent optical networks," Photonic Netw. Commun. **27**, 8–15 (2014).
10. T. Zami, B. Lavigne, and O. B. Pardo, "Benchmarking of opaque versus transparent core wdm networks featuring 400zr+ qsfp-dd or cfp2 interfaces," in *European Conference on Optical Communications (ECOC),* (2020).
11. R. Sadeghi, B. Correia, E. Virgillito, A. Napoli, N. Costa, J. Pedro, and V. Curri, "Comparison of transceiver and c+ l band upgrades: Network traffic and energy assessment," in *International Conference on Electrical, Communication, and Computer Engineering (ICECCE),* (IEEE, 2021).
12. F. Christou, T. Enderle, and A. Witt, "Towards a hybrid architecture by introducing coherent pluggable transceivers in ip-optical core networks with optical cross-connects," in *Photonic Networks; 23th ITG-Symposium,* (VDE, 2022), pp. 1–8.
13. R. Sadeghi, B. Correia, A. Souza, N. Costa, J. Pedro, A. Napoli, and V. Curri, "Transparent vs translucent multi-band optical networking: Capacity and energy analyses," IEEE/Optica J. Light. Technol. **40**, 3486–3498 (2022).
14. A. Gumaste, M. Sosa, H. Bock, and P. Kandappan, "Optimized ip-over-wdm core networks using zr+ and flexible muxponders for 400 gb/s and beyond," IEEE/Optica J. Opt. Commun. Netw. **14**, 127–139 (2022).
15. F. Masoud, A. Napoli, C. Castro, J. Pedro, B. Spinnler, A. Chase, D. Hillerkuss, and D. Welch, "Environmental impact of coherent point-to-multipoint pluggables in metro aggregation optical networks," in *2023 International Conference on Optical Network Design and Modeling (ONDM),* (IEEE, 2023).
16. R. Davey, M. Iqbal, and P. Wright, "Zr 400 gbit/s and 800 gbit/s use cases, trials, deployments, and future prospects," IEEE/Optica J. Opt. Commun. Netw. **16**, A33–A39 (2024).
17. P. Wright, R. Davey, and A. Lord, "Cost model comparison of zr/zr+ modules against traditional wdm transponders for 400g ip/wdm core networks," in *European Conference on Optical Communications (ECOC),* (2020).
18. A. Eira and J. Pedro, "On the comparative efficiency of next-generation coherent interfaces for survivable network design," in *International Conference on the Design of Reliable Communication Networks (DRCN),* (2021).
19. S. Melle, T. Zami, O. Bertran-Pardo, and B. Lavigne, "Comparing IP-Optical Architectures amp; WDM Transport Technologies in Metro, Regional and Long-Haul Networks," in *Optical Fiber Communications Conference and Exhibition (OFC),* (2021).
20. OIF, "400ZR," https://www.oiforum.com/technical-work/hot-topics/400zr-2/, accessed on 18/02/2024.
21. OpenZR+, "OpenZR+," https://www.openzrplus.org/, accessed on 18/02/2024.
22. T. Zami, B. Lavigne, and M. Lefrançois, "Added value of 90 gbaud transponders for wdm networks," in *Optical Fiber Communications Conference and Exhibition (OFC),* (IEEE, 2020).
23. Nokia, "7750 service router," https://www.nokia.com/networks/ip-networks/7750-service-router/, accessed on 18/02/2024.
24. Nokia, "Nokia 7250 IXR-10/IXR-6/IXR-s Interconnect Routers," https://onestore.nokia.com/asset/201561?_ga=2.165446397.352188883.1698411507-1835375626.1690936125, accessed on 18/02/2024.
25. Cisco, "Cisco ACI Multi-tier Architecture White Paper ," https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-742214.html, accessed on 18/02/2024.
26. J. Kundrát, O. Havliš, J. Jedlinskỳ, and J. Vojtěch, "Opening up roadms: let us build a disaggregated open optical line system," IEEE/Optica J. Light. Technol. **37**, 4041–4051 (2019).
27. OpenZR+, "OpenZR+ 400G Digital Coherent Optics for Multi-Haul," https://openzrplus.org/site/assets/files/1074/openzrplus_whitepaper_-_sept_29_2020_final.pdf, accessed on 04/09/2024.
28. Cisco, "Why Coherent Pluggables Revolutionize Optical Infrastructure," https://www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2024/pdf/PSOSP-2726.pdf, accessed on 05/09/2024.
29. M. Batayneh, D. A. Schupke, M. Hoffmann, A. Kirstaedter, and B. Mukherjee, "On routing and transmission-range determination of multi-bit-rate signals over mixed-line-rate wdm optical networks for carrier ethernet," IEEE/ACM Transactions on Netw. **19**, 1304–1316 (2011).
30. M. Cavalcante, H. Pereira, D. Chaves, and R. Almeida Jr, "An auxiliary-graph-based methodology for regenerator assignment problem optimization in translucent elastic optical networks," Opt. Fiber Technol. **53**, 102008 (2019).
31. Q. Zhu, X. Yu, Y. Zhao, A. Nag, and J. Zhang, "Auxiliary-graph-based energy-efficient traffic grooming in ip-over-fixed/flex-grid optical networks," IEEE/Optica J. Light. Technol. **39**, 3011–3024 (2021).
32. L. R. Costa, Í. B. Brasileiro, and A. C. Drummond, "Energy efficiency in sliceable-transponder enabled elastic optical networks," IEEE Transactions on Green Commun. Netw. **5**, 789–802 (2021).
33. B. Chen, W. Zheng, J. Yu, D. Zheng, H. Chen, M. Gao, W. Ju, P.-H. Ho, J. P. Jue, and G. Shen, "Crosstalk-aware virtual network mapping in space-division-multiplexing optical data center networks," IEEE Transactions on Commun. **72**, 3526–3542 (2024).
34. M. Ibrahimi, H. Abdollahi, C. Rottondi, A. Giusti, A. Ferrari, V. Curri, and M. Tornatore, "Machine learning regression for qot estimation of unestablished lightpaths," IEEE/Optica J. Opt. Commun. Netw. **13**,

B92–B101 (2021).
35. B. G. Bathula and J. M. Elmirghani, “Constraint-based anycasting over optical burst switched networks,” IEEE/Optica J. Opt. Commun. Netw. **1**, A35–A43 (2009).
36. S. Melle, T. Zami, N. Rossi, and B. Lavigne, “Optimal transponder technology for transporting 800 gbe services in ip-over-wdm backbone networks,” in *Optical Fiber Communications Conference and Exhibition (OFC),* (2023).