Traffic-aware network reconfiguration for energy-efficient IPoWDM networks with ZR/ZR+ pluggable optics

Qiaolun Zhang[1], Nisanur Camuzcu[2], Xiaoyang Guo[1], Jiaheng Xiong[1], Ruikun Wang[3], Tianqu Luo[1], Mëmëdhe Ibrahimi[1], Massimo Tornatore[1]

[1] Department of Electronics, Information and Bioengineering, Politecnico di Milano, 20133 Milano, Italy, [2] Rochester Institute of Technology, Rochester, NY, USA, [3] College of Information Science and Technology, Qingdao University of Science and Technology, Qingdao 266061, China

Corresponding author: Ruikun Wang, ruikun.wang@qust.edu.cn

Hyperscale AI workloads, cloud services, and data center interconnect are increasing optical network traffic in future 6G systems. This traffic growth will raise network power consumption. ZR/ZR+ pluggable optics consume less power than traditional long-haul muxponders and are a promising option for IP-over-Wavelength Division Multiplexing (IPoWDM) networks. Yet real network traffic changes during the day, and the energy behavior of ZR/ZR+-based networks under dynamic reconfiguration has not been investigated. This paper studies traffic-aware reconfiguration strategies, including modulation-format adjustment and traffic rerouting at the IP and optical layers, for IPoWDM networks equipped with ZR/ZR+ pluggables or traditional long-haul muxponders. We formulate the 24-hour energy minimization problem and develop auxiliary graph-based heuristic algorithms to solve it. Using the Japan and German (G50) topologies and real AMS-IX traffic statistics, the results show that ZR/ZR+ uses 25.3–34.5% less power at peak load. However, reconfiguration reduces ZR/ZR+'s daily energy by only about 0.6–5.3%, compared to about 8.2–13.6% for long-haul muxponders, because ZR/ZR+ router power is released only when an entire 400G port is turned off, whereas long-haul muxponder power scales continuously with traffic. These results indicate that, although rerouting strategies yield more limited energy gains for ZR/ZR+ than for long-haul muxponders, they remain necessary for energy-aware 6G backhaul and data center interconnect networks, as freeing entire 400G router ports is one of the most effective ways to dynamically reduce energy with ZR/ZR+ pluggables.

Keywords: IPoWDM networks, ZR/ZR+ pluggable optics, network reconfiguration, energy efficiency, optical network

1. INTRODUCTION

Optical transport networks are under increasing pressure to reduce energy consumption. The ICT sector already accounts for approximately 2–3% of global energy consumption [1, 2], and continued traffic growth is expected to further increase the power required by communication networks. This growth is driven by machine-learning training workloads, cloud services, and inter-data-center traffic carried by Data Center Interconnect (DCI) [3, 4, 5], with additional demand expected from edge computing, immersive applications, massive IoT deployments, and future sixth-generation (6G) mobile systems [6, 7]. As these services increase fronthaul, backhaul, metro, and inter-data-center traffic, energy-efficient, intelligent, and dynamically reconfigurable optical transport becomes increasingly important for future networks [8, 9].

ZR and ZR+ pluggable coherent transceivers are increasingly used in optical transport networks. ZR is standardized by the Optical Internetworking Forum (OIF) at 400 Gb/s for short-reach DCI links [10]. OpenZR+ extends this idea to regional links with three rate-reach operating points from 200 to 400 Gb/s and reaches up to 3 000 km [11]. These compact modules place coherent optics directly into an IP router line card, and a ZR module consumes much less power than a long-haul muxponder [12, 13]. Under static, peak-traffic conditions, this power advantage has been well investigated across different network architectures [14, 15, 16, 17, 18]. Moreover, some recent studies further validate ZR/ZR+ feasibility and use cases in metro aggregation and regional IPoWDM segments [19, 20].

The behavior of ZR/ZR+-based IPoWDM networks under dynamic traffic is still not well understood. Real network traffic does not remain constant at peak demand; measurements from the Amsterdam Internet Exchange (AMS-IX) show that aggregate traffic regularly falls to 50–60% of its daytime peak during the early morning hours [21]. For conventional long-haul IPoWDM networks, this daily variation has motivated many studies on traffic-aware reconfiguration, including modulation-format adaptation, demand rerouting, and sleep-mode operation, with reported power or energy reductions on the order of 10–25% or higher: conservative grooming and transponder-configuration changes yield roughly 9–13% power reduction [22], while direct daily IP-over-WDM reconfiguration achieves 16% or more daily energy savings for simple line-card shutoff and substantially larger savings when IP-layer rerouting is included [23]; multi-period studies further show that these savings must be balanced against reconfiguration cost [24]. It is unclear whether the same strategies can deliver similar savings in ZR/ZR+-based networks.

ZR/ZR+ differs from long-haul muxponders in two important ways. First, because ZR/ZR+ modules are hosted on 400G physical-layer router ports, each active 400G port requires a fixed amount of router modular power, regardless of the actual carried traffic. In contrast, long-haul muxponders use 100G client-side cards whose power scales proportionally with carried traffic in 100 Gb/s increments. A reduction in traffic therefore cannot reduce router power for ZR/ZR+ until the load on a given port drops entirely to zero and the port can be turned off. Second, for the Nokia equipment configuration used to parameterize this study [12], ZR offers only a single modulation format (16-QAM at 400 Gb/s), whereas ZR+ supports three rate-reach operating points spanning 200–400 Gb/s; the corresponding long-haul muxponder supports seven operating points from 200 to 800 Gb/s, with reaches up to 5 700 km. Because ZR/ZR+ offers only this limited set of rate–reach operating points, modulation format adaptation can achieve only limited power proportionality. These two differences lead to the main question of this paper: does ZR/ZR+ retain its power advantage over long-haul muxponders when both technologies use the same traffic-aware reconfiguration strategies, and how large is the reconfiguration benefit?

This paper extends our previous conference paper [25]. The main contributions of this paper, as well as the key differences from the conference version, are summarized as follows.

1. We formally define and characterize five traffic-aware reconfiguration strategies for IPoWDM networks, ranging from modulation format-only adaptation to joint IP and optical layer rerouting for both opaque and transparent architectures. In addition, we analyze when and how effectively these strategies apply to ZR/ZR+ versus long-haul muxponder networks using an illustrative example, which was not included in the conference version [25].
2. We develop a family of auxiliary graph-based heuristic algorithms for all five strategies, explicitly modeling the 400G router port constraint that governs ZR/ZR+ interfaces. Compared with the conference version [25], this paper provides the complete description of the proposed algorithm and its variants.
3. We provide a more comprehensive quantitative evaluation across two metro-regional network topologies, namely the 14-node Japan topology and the 50-node German (G50) topology, and two architectures, opaque and transparent. The evaluation uses traffic demands generated from scaled real AMS-IX traffic traces and substantially extends the evaluation in the conference version [25].

The remainder of this article is organized as follows. Section 2 reviews related work. Section 3 presents the system model, power consumption framework, and an illustrative example. Section 4 describes the proposed algorithms. Section 5 reports simulation results, and Section 6 concludes the work.

2. RELATED WORK

2.1 Energy efficiency in optical networks

Energy consumption in optical networks has received considerable attention because IP routers, transponders, and optical amplifiers are among the major contributors to total network power [26, 27, 28]. A first line of work studies power-aware network planning and transport architecture design, showing that energy consumption can be reduced by optimizing the placement, dimensioning, and utilization of network resources [29, 30]. A complementary line of work exploits the load dependence of network operation: lightly-loaded network elements can be placed in low-power or sleep states, and traffic can be routed over fewer active resources during periods of low demand [8, 31, 32, 33]. These ideas also extend to flexible-grid and elastic optical networks, where routing, modulation format, and spectrum assignment can be jointly optimized under time-varying traffic to improve energy efficiency [34, 35].

More recent work has investigated reconfiguration to improve the energy efficiency in the context of daily traffic patterns [36, 37, 38, 22, 39]. The first direction is to re-route the traffic to reduce the power consumption. For instance, stateful grooming was proposed to exploit daily traffic patterns for energy-efficient IP-over-optical networks [36]. Energy savings in dynamic and resilient optical networks using traffic-aware strategies were studied in [37]. Energy efficiency in next-generation optical networks was investigated in [38], while energy-efficient configuration methods for flexible transponders and grooming switches were proposed in [22]. In addition to re-routing the traffic, another direction for network reconfiguration is rate adaptation. Specifically, the impacts of rate adaptation at daily timescales for reducing network energy consumption were investigated in [39].

While the above studies mainly focus on algorithmic mechanisms for improving energy efficiency through routing, grooming, and rate adaptation, the network architecture itself is another important factor affecting network power consumption. Previous studies of opaque and transparent optical architectures have consistently shown that transparent architectures offer lower power consumption for traditional transponders by eliminating intermediate Optical-Electrical-Optical (OEO) conversion at intermediate nodes [40, 27]. A systematic benchmarking of these architectures at scale further confirms that the energy advantage of transparent architectures increases with node degree and average link length [41].

2.2 Network reconfiguration for reducing power in optical networks

Network reconfiguration has been widely studied as an effective approach to reducing power consumption in optical networks equipped with conventional long-haul transponders [29, 30, 35, 42, 34]. By exploiting temporal traffic variation, early studies optimized network planning and transport architectures to serve traffic with fewer active network elements under changing load conditions [29, 30]. Subsequent work formulated this idea as energy-efficient routing, modulation, and spectrum assignment for elastic optical networks with time-varying traffic [35], while dynamic traffic-grooming heuristics were developed to reduce computational complexity while maintaining near-optimal performance [42, 34]. These studies generally treat traffic-aware reconfiguration as a time-dependent resource-assignment problem, where routing, spectrum allocation, grooming, and modulation choices are periodically adapted to satisfy traffic demands with lower power consumption.

Several algorithmic frameworks have further improved the practicality of reconfiguration in IPoWDM networks. The auxiliary graph framework has been shown to be particularly effective for energy-aware reconfiguration heuristics [43], as it captures routing, grooming opportunities, and incremental power costs within a unified graph representation. More recently, a detailed power model for IPoWDM networks was formalized in [12]. However, these studies mainly focus on networks based on conventional long-haul transponders or muxponders, whose power characteristics and resource constraints differ substantially from those of emerging pluggable optics.

Recently, ZR/ZR+ pluggable transceivers have emerged as a practical alternative to conventional long-haul muxponders, attracting increasing attention due to their compact form factor and low power consumption [19, 44, 12, 13, 45]. Related analyses of multi-band optical networking have also explored the energy and capacity trade-offs of advanced transceiver technologies [46]. Network-scale evaluations have also confirmed the potential benefits of ZR/ZR+ pluggables in practical deployment scenarios. Multiple IP-optical configurations across metro, regional, and long-haul network scales were compared in [14], identifying the conditions under which pluggable optics are preferable to conventional transponder-based solutions. The comparative efficiency of next-generation coherent interfaces for survivable network design was evaluated in [15], while transceiver upgrade strategies and their impact on network traffic and energy consumption were assessed in [47]. Cost-oriented studies in [17, 18] showed that ZR/ZR+ can reduce both capital expenditure and total cost of ownership, especially at data center interconnect and metro scales.

Despite these advances, existing studies on ZR/ZR+ pluggables have mainly focused on static deployment studies, architectural comparisons, or dynamic service provisioning under online RMSA decisions [17, 18, 48]. In particular, recent work has investigated energy-aware RMSA for IPoWDM networks with ZR+ transceivers, including deep-reinforcement-learning-based approaches that adapt routing, modulation, spectrum, and transceiver operational-mode selection to dynamic connection requests [48]. However, these studies do not address traffic-aware reconfiguration of already established services over time-varying traffic profiles. The impact of such reconfiguration on ZR/ZR+-based IPoWDM networks therefore remains largely unexplored. This gap is important because ZR/ZR+ pluggables introduce distinct operational constraints, especially the 400G client-port structure, which changes how traffic grooming, port occupancy, and decisions about when resources can be turned off should be modeled during reconfiguration. Moreover, no prior work has directly compared ZR/ZR+ pluggables and conventional long-haul muxponders under identical dynamic traffic profiles and reconfiguration strategies.

This paper fills this gap by studying energy-saving reconfiguration in IPoWDM networks equipped with ZR/ZR+ pluggables. We build on existing reconfiguration frameworks and adapt the algorithm to handle ZR/ZR+-specific constraints: 400G port occupancy and port turn-off events. We further provide a side-by-side comparison between ZR/ZR+ pluggables and long-haul muxponders under the same traffic dynamics, network architectures, and reconfiguration strategies, and quantify how the energy advantage of ZR/ZR+ changes under dynamic operation.

3. SYSTEM MODEL

3.1 IPoWDM network architectures

We consider two IPoWDM network architectures. In an *opaque* architecture, every intermediate node performs Optical-Electrical-Optical (OEO) conversion: the incoming signal is terminated, electronically processed at the IP layer, and regenerated onto the outgoing fiber. Wavelength multiplexing and demultiplexing are handled by passive Arrayed Waveguide Gratings (AWGs), and there is no optical bypass of intermediate nodes. While the opaque model is simple to implement and manage, the per-hop OEO overhead results in substantial power consumption that scales with both traffic volume and the number of hops [40].

In a *transparent* architecture, optical signals traverse intermediate nodes without OEO conversion, using Reconfigurable Optical Add-Drop Multiplexers (ROADMs) and Add-Drop Blocks (ADBs) to route individual wavelengths [49]. Electronic processing occurs only at the endpoints of each lightpath, reducing transponder count and eliminating intermediate regeneration power [27]. The flexibility to optically bypass nodes also enables a richer set of reconfiguration options: as traffic decreases, lightpaths can be rerouted to longer transparent spans, activating fewer active transponders. Fig. 1 illustrates the node structures for both architectures. We evaluate both architectures to assess the energy efficiency of ZR/ZR+.

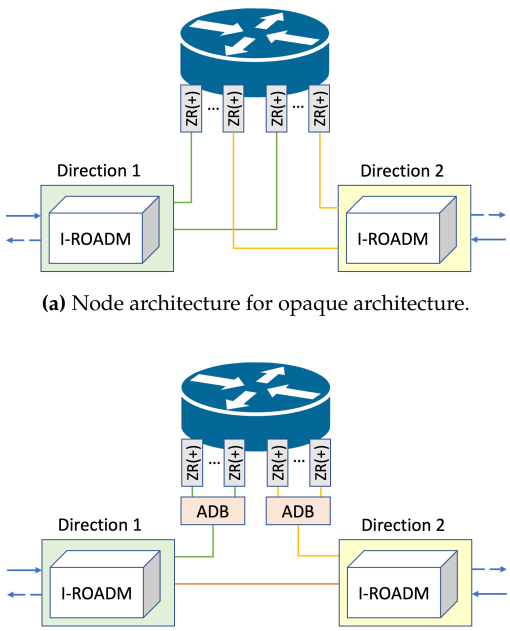

(a) Node architecture for opaque architecture.

(b) Node architecture for transparent architecture.

Figure 1 – Node architectures for transparent and opaque IPoWDM networks [12].

3.2 Power consumption model

We adopt the normalized power model from [12, 13], in which all power values are expressed relative to the draw of a single ZR 400G module. Table 1 lists the normalized power consumption of each network element.

The IP router model is based on a 36 Tb/s chassis with a fixed component of 75 units and a modular component that depends on the number and type of active line cards. The key difference between ZR/ZR+ and long-haul muxponders lies here: ZR/ZR+ pluggables connect via 400G physical-layer interfaces (as operators typically want ZR/ZR+ has the capability to operate at the largest data rate), which consume four units of modular router power regardless of the actual traffic carried on that interface. Long-haul muxponders use 100G client-side cards, each consuming one unit, so the router's modular power scales proportionally with carried traffic in 100 Gb/s increments. Table 1 therefore separates the 400G ZR/ZR+ port entry from the 100G long-haul client-card entry. This difference matters for reconfiguration: reducing the load on a ZR/ZR+ channel does not reduce router power until the load drops to zero and the port can be turned off, whereas any reduction in long-haul traffic releases a proportional amount of router power.

Transponder power is modeled separately and decreases with a lower modulation format, as Table 1 quantifies. For ZR+, the three rate-reach operating points span a power range of 1.04–1.30 units [46]. Similarly, the long-haul muxponder transponder power decreases as a lower modulation format, and hence a lower data rate, is used: the range is 7.22–8.00 units across seven operating points from 200G to 800G. Long-haul muxponders thus have substantially higher absolute power consumption but offer finer reconfiguration granularity, giving them more flexibility to achieve off-peak savings through modulation format reconfiguration.

Table 1 – Normalized power consumption of network elements [11, 10, 13, 12].

Device	Type	Power
Transponder	ZR 400G	1.00
	ZR+ 400G	1.30
	ZR+ 300G	1.17
	ZR+ 200G	1.04
	Long-Haul 800G	8.00
	Long-Haul 700G	7.87
	Long-Haul 600G	7.74
	Long-Haul 500G	7.61
	Long-Haul 400G	7.48
	Long-Haul 300G	7.35
	Long-Haul 200G	7.22
IP Router	Fixed chassis	75
	Modular 400G port (ZR/ZR+)	4
	Modular 100G client card (long-haul)	1
Optical side	Add-Drop Block (ADB)	4.0
	I-ROADM (bidirectional)	4.1
	Optical Amplifier (unidirectional)	1.7
	AWG	0.3
	Shelf	20
Monitoring	Opaque (bidirectional)	0.3
	Transparent (bidirectional)	1.5

3.3 Transceiver reach table

Tables 2 and 3 list the rate-reach operating points for ZR/ZR+ pluggables and long-haul muxponders, respectively. Long-haul muxponders support data rates from 200 to 800 Gb/s with reaches up to 5 700 km across seven operating points, giving them substantially greater routing and format-adaptation flexibility. ZR operates at a fixed 400 Gb/s over up to 120 km, while ZR+ provides three rate-reach operating points (200–400 Gb/s, reaches up to 3 000 km). In the Japan and German (G50) topologies evaluated here, link lengths fall within the maximum reach supported by ZR+, while fixed-rate ZR is applicable only to links no longer than 120 km.

Table 2 – Reach table for ZR/ZR+ pluggable modules [12].

Module	Rate (Gb/s)	Modulation	Spacing (GHz)	Reach (km)
ZR	400	16-QAM	100	120
ZR+	400	16-QAM	75	600
ZR+	300	8-QAM	75	1 800
ZR+	200	QPSK	75	3 000

Table 3 – Reach table for long-haul muxponders [12].

Module	Rate (Gb/s)	Modulation	Spacing (GHz)	Reach (km)
Long-Haul	800	PCS 64-QAM	100	150
Long-Haul	700	PCS 64-QAM	100	400
Long-Haul	600	16-QAM	100	700
Long-Haul	500	PCS 16-QAM	100	1 300
Long-Haul	400	PCS 16-QAM	100	2 500
Long-Haul	300	PCS 16-QAM	100	4 700
Long-Haul	200	PCS 16-QAM	100	5 700

3.4 Reconfiguration strategies

In this work, we consider five reconfiguration strategies, from modulation-format-only adaptation to combined strategies that use both rerouting and format changes.

Modulation Format (MF) Reconfiguration. This strategy only changes the modulation format of active lightpaths. When the traffic carried by a lightpath decreases, the corresponding transponder is switched to the lowest-data-rate modulation format that can still accommodate the demand while satisfying the reach constraint of the link. Fixed-rate ZR has a single 400G operating point. For ZR+, this corresponds to reconfiguration from 400G 16-QAM to 300G 8-QAM or 200G QPSK, depending on the link length. The ZR/ZR+ router interface, however, remains a 400G port that consumes four normalized units regardless of the downshifted data rate, so MF saves only transponder power for ZR/ZR+ unless rerouting makes an entire port unnecessary. By contrast, long-haul muxponders use 100G client-side cards, allowing router modular power to decrease stepwise as carried traffic is reduced.

IP-layer Rerouting (RE(ip)). This strategy performs traffic grooming only at the IP layer. Demands are rerouted so that flows whose aggregate traffic fits within the residual capacity of an existing channel can be moved onto that channel. As a result, previously active transponders can be released and powered down when a channel no longer carries traffic. The underlying optical paths remain unchanged.

IP and optical-layer Rerouting (RE(ip+op)). This strategy extends IP-layer grooming by also allowing optical-layer rerouting. Demands may be shifted to different wavelength routes, enabling transponders associated with empty lightpaths to be powered down. This strategy is particularly beneficial in transparent architectures, where optical rerouting may also remove intermediate regeneration sites by extending the reach of the remaining lightpaths.

MF reconfiguration combined with IP-layer Rerouting (MFRE(ip)). This combined strategy first applies IP-layer rerouting to move traffic onto fewer active channels. Modulation format reconfiguration is then applied to the resulting channel assignment, allowing the remaining transponders to operate at the lowest feasible data rate.

MF reconfiguration combined with IP and optical-layer Rerouting (MFRE(ip+op)). This strategy combines the most flexible rerouting option with modulation-format adaptation. It first performs joint IP and optical-layer rerouting to move traffic onto fewer lightpaths and release unnecessary lightpaths, and then reconfigures the modulation formats of the remaining active lightpaths.

All five strategies defined above are applicable to IPoWDM networks with either ZR/ZR+ pluggables or long-haul muxponders. Section 3 establishes that ZR/ZR+ and long-haul muxponders differ in router port granularity and modulation format range, which affects the magnitude of savings achievable under each strategy; the strategy definitions themselves are transceiver-agnostic, and all five are evaluated for both transceiver types in Section 5.

3.5 Problem statement

The network is modeled as a physical graph $G = (V, E)$ with node set V and bidirectional link set E. Each link $e \in E$ has spectral capacity S_e, expressed as the maximum number of WDM channels available on the link. The 24-hour period is divided into $T = 12$ reconfiguration intervals of duration $\Delta t = 2$ hours ($T = 24\,\text{h}/\Delta t = 12$). A set of traffic demands $\mathcal{D}$ is given, with each demand $d \in \mathcal{D}$ characterized by a source node s_d, a destination node t_d, and a time-varying data rate $r_d(t)$ obtained by scaling the peak demand according to the AMS-IX daily profile.

Given: Network topology G and link spectrum S_e for all $e \in E$; chosen architecture (opaque or transparent) and reconfiguration strategy from Section 3.4; traffic demands $\mathcal{D}$ at each time period; power consumption model (Table 1); transceiver reach table (tables 2–3).

Decide: For each time period, the routing path for each demand, the modulation format of each active transponder, and the rerouting decisions at the IP layer or jointly with the optical layer.

Subject to: All demands are served within the available spectrum on each link; modulation format reach constraints are satisfied; ZR/ZR+ ports occupy the full 400G of router interface bandwidth regardless of carried traffic.

Objective: Minimize the total energy consumed by the network over the 24-hour period:

$$\min \sum_{t=1}^{T} \Delta t \cdot P_{\text{total}}(t), \quad (1)$$

where $P_{\text{total}}(t)$ is the total instantaneous network power at time period t, summing contributions from all active routers, transponders, optical infrastructure, and monitoring components, and $T = 12$ time periods cover the full 24-hour day.

3.6 Illustrative example

To illustrate how network reconfiguration can impact power consumption of IPoWDM networks with ZR/ZR+, we present a 3-node example under an opaque architecture. Fig. 2 shows the peak-traffic state. Link distances are 500 km on links (1, 2) and (2,3). Two demands are active:

1. *Req1* (1 → 3): 400 Gb/s, served by 400G ZR+ channels on the path 1−2−3.
2. *Req2* (2 → 3): 100 Gb/s, served by a 200G ZR+ transponder on link 2−3.

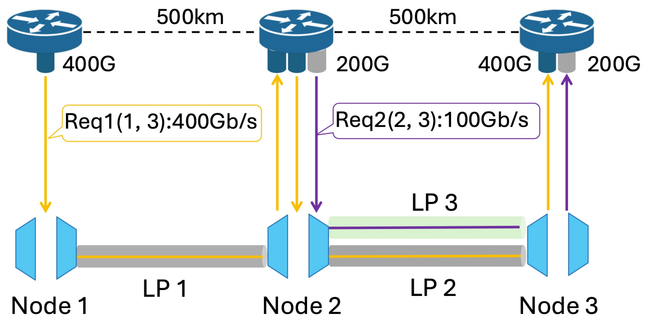

Figure 2 – Three-node opaque network under peak traffic.

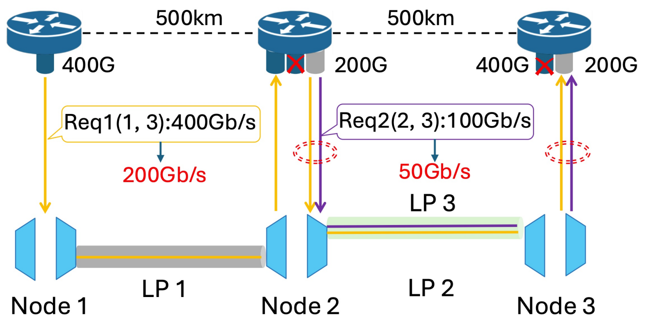

Figure 3 – Three-node opaque network at 50% of peak traffic after IP-layer rerouting. Port power is unchanged.

During the day, when traffic is halved, applying only modulation format (MF) reconfiguration yields marginal savings: transponders are downshifted to the minimum format that supports the reduced demand, but the router port count remains unchanged.

In contrast, the RE(ip) strategy achieves a much larger power reduction. As shown in Fig. 3, the reduced Req1 traffic of 200 Gb/s and Req2 traffic of 50 Gb/s are carried over the existing ZR+ channels on links (1,2) and (2,3). In this case, two ZR+ modules, namely those associated with the request between node pair (2,3), can be turned off. This reduces both the ZR/ZR+ power consumption and the corresponding router traffic load, thereby lowering the associated router power consumption.

The most significant savings are obtained with the combined MFRE strategy. After IP-layer grooming, the 400G ZR+ channel on link (1,2) can be downshifted to 200G ZR+, while the channel on link (2,3) can be downshifted to 300G ZR+, since the groomed traffic demands are 200 Gb/s and 250 Gb/s, respectively. Consequently, the total power consumption can be further reduced by operating the ZR/ZR+ modules at lower modulation formats.

This example reveals two main insights into the impact of network reconfiguration in IPoWDM networks with ZR/ZR+. First, under the 400G ZR/ZR+ port constraint, simple traffic reduction through modulation format reconfiguration does not reduce router power, because the number of active router ports remains unchanged; router power is reduced only when ZR/ZR+ modules are turned off. Second, MFRE can further reduce power consumption by jointly grooming traffic onto fewer lightpaths and lowering the modulation formats of the remaining active ZR/ZR+ modules.

4. HEURISTIC ALGORITHMS

Since the daily energy minimization problem is NP-hard in the general case (it generalizes the offline traffic grooming problem [43]), we develop efficient heuristic algorithms for each of the reconfiguration strategies in Sec. 3.4.

4.1 Auxiliary graph construction

The aim of this work is to evaluate energy-consumption behavior under traffic-aware reconfiguration rather than to provide an exact solution to the network design problem. Auxiliary graph-based heuristics are a widely adopted approach for traffic grooming and resource-allocation problems in optical networks [43, 42, 50, 51, 52], with published results demonstrating close to optimal performance, e.g., within 4–8% of the MILP optimum on tractable instances [34]; ILP formulations become computationally intractable for large-scale topologies, which motivates the heuristic approach adopted here. All algorithms in this section are therefore built on a shared auxiliary graph-based routing procedure.

The auxiliary graph G_{aux} provides a snapshot of all feasible routing options at a given moment, with edge weights reflecting the marginal power cost of each option [43, 42]. Each node in G_{aux} corresponds to a physical node; each directed edge (n_1, n_2) represents a feasible way to carry traffic between those nodes, either by grooming onto an existing channel with residual capacity or by opening a new lightpath. The edge weight is the incremental power cost of that choice.

Algorithm 1 builds G_{aux} for one request. Its inputs are the request bit rate R_b and a grooming flag $G_f \in \{0, 1\}$: G_f is set to 1 for the RE and MFRE strategies, which try to groom traffic onto existing channels before opening new ones, and to 0 for the MF-only strategy, which never reroutes. For a candidate path p, $\ell(p)$ is its length and C_{res} is the residual capacity of the channels already active along p; m denotes a modulation format, with reach reach(m) and data rate rate(m), and w_m is the incremental power cost of serving the request with a new lightpath that uses m. The algorithm starts from an empty graph (line 1) and examines every ordered node pair (n_1, n_2) (line 2). For each pair, it takes the shortest candidate path p (line 3), computes its length $\ell(p)$ and the residual capacity C_{res} along it (line 4), and initializes the best cost $w^\star$ and the best option found so far for this pair (line 5). When grooming is enabled and an active channel still has enough spare capacity ($G_f = 1$ and $C_{\text{res}} \geq R_b$, line 6), the request is groomed onto that channel without activating a new transponder (line 7), and this grooming option is given a cost slightly below that of a new lightpath (line 8), so that grooming is preferred whenever it is feasible. Otherwise, a new lightpath must be opened: the algorithm scans the modulation formats m whose reach covers $\ell(p)$ and whose data rate can carry R_b (line 9), keeps those for which enough spectrum slots are still free along p (line 10), and retains the modulation format with the lowest incremental power cost w_m (line 11). Finally, if a feasible option was found (line 12), the edge (n_1, n_2) is added to G_{aux} with the selected weight $w^\star$ and its serving information (line 13). The grooming check at line 6 is the key step, because reusing an already-active channel is what avoids the power cost of a new transponder. In transparent architectures, paths that avoid intermediate OEO regeneration are additionally given a small weight advantage (10^{-3}), so that optical bypass is preferred where it is allowed.

4.2 Rerouting algorithm (RE variants)

The rerouting algorithm, whose overall flow is illustrated in Fig. 4, operates as follows. The peak traffic is first provisioned fully, with demands served using the minimum-cost assignment that achieves at most 1% blocking. At each subsequent reconfiguration interval, the demand rates are updated according to the traffic profile and the current request set is processed one request

Algorithm 1: Auxiliary graph construction.

Input : request bit rate R_b; grooming flag $G_f \in \{0,1\}$
Output: auxiliary graph G_{aux}

Initialize G_{aux} as an empty graph;
foreach *ordered node pair* (n_1, n_2) **do**
  $p \leftarrow$ shortest candidate path between n_1 and n_2;
  Compute length $\ell(p)$ and residual capacity C_{res} of the channels already active along p;
  Initialize best cost $w^\star \leftarrow \infty$ and best option $\leftarrow$ none;
  if $G_f = 1$ ***and*** $C_{\text{res}} \geq R_b$ **then**
    best option $\leftarrow$ groom R_b onto an existing channel on p (no new transponder);
    $w^\star \leftarrow$ grooming cost, set below the cost of opening a new lightpath;
  else
    foreach *modulation format m with* reach$(m) \geq \ell(p)$ ***and*** rate$(m) \geq R_b$ **do**
      if *enough free spectrum slots for m along p* **then**
        Compute new-lightpath cost w_m; **if** $w_m < w^\star$ **then** set $w^\star \leftarrow w_m$ and best option $\leftarrow$ new lightpath with m;
      end
    end
  end
  if *best option* $\neq$ *none* **then**
    Add edge (n_1, n_2) to G_{aux} with weight $w^\star$ and the serving info of the best option;
  end
end

at a time. For each demand, the algorithm tentatively releases the resources currently assigned to it, rebuilds G_{aux} with those resources tentatively released, and routes the demand on the minimum-cost path. If the new total power is lower than the current assignment, the rerouting is accepted; otherwise, the original assignment is restored. This per-request accept-or-restore pass continues until every request in the interval has been evaluated.

The distinction between RE(ip) and RE(ip+op) lies in the scope of G_{aux}: in RE(ip), optical paths are fixed and the algorithm grooms traffic over existing IP-layer edges or wavelength routes with sufficient residual capacity; in RE(ip+op), the auxiliary graph may also modify optical path usage, allowing transponders that IP-layer rerouting alone cannot free to be powered down.

4.3 Modulation Format (MF) reconfiguration

The MF algorithm is applied after peak provisioning. For each active channel, it identifies the total carried traffic and finds the minimum-capacity modulation format (from tables 2–3) whose data rate covers the load

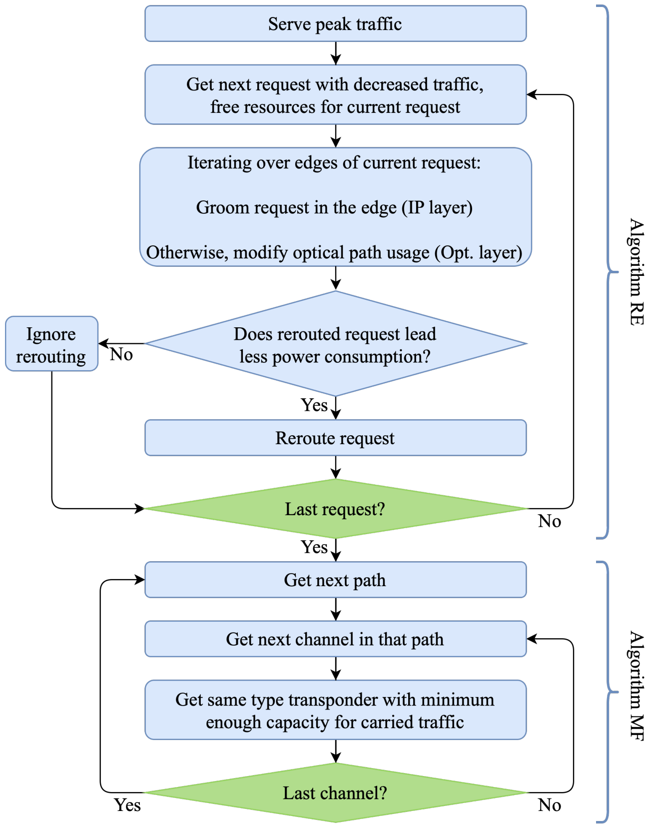

Figure 4 – General flow of the reconfiguration algorithm. The request loop evaluates all requests in the current interval and accepts a tentative reroute only when it reduces total power.

and whose reach is compatible with the lightpath length. Fixed-rate ZR remains at 400G; ZR+ and long-haul channels can be downshifted when a lower-rate operating point is feasible. If a lower-order format is feasible, the transponder is switched, reducing its power draw as per Table 1. For ZR/ZR+, this change does not reduce the 400G router-port contribution unless rerouting makes the port unnecessary; for long-haul muxponders, the associated 100G client-card allocation can decrease with the carried traffic. Since modulation downshifting also reduces the channel's maximum capacity, the algorithm verifies that all groomed demands on the channel still fit under the new format before committing the change.

4.4 Combined MFRE strategies

For MFRE(ip) and MFRE(ip+op), rerouting runs first, producing a traffic assignment with fewer active paths; MF is then applied to the resulting channel state. The two steps are complementary: rerouting concentrates traffic onto fewer active ports and channels, after which MF captures additional savings through format downshifting on the remaining, lighter-loaded channels. Reversing the order would risk locking in format configurations that block subsequent traffic grooming.

5. NUMERICAL RESULTS

5.1 Simulation settings

Simulations are implemented in Python 3.8. We evaluate two topologies:

14-node Japan topology [53]: link lengths range from 40 to 320 km (average 165 km), characteristic of a dense national backbone. All ZR/ZR+ rate-reach operating points are applicable on most links.

50-node German (G50) topology [54]: link lengths range from 27 to 219 km (average 100 km). The high connectivity and short links make this topology particularly favorable for ZR/ZR+ deployment.

Traffic demands are generated by scaling hourly AMS-IX traffic statistics to each topology's capacity, with a target blocking rate of 1% for ZR/ZR+ at peak load, used as a peak-load reference point: the AMS-IX traffic profile is scaled so that the network operates close to full utilization at peak load while still serving almost all peak traffic. This provides a balanced operating point for comparing ZR/ZR+ and long-haul muxponder reconfiguration. The blocking target specifies the maximum fraction of requests that may be rejected at peak load; the AMS-IX traffic profile is scaled upward until peak blocking reaches this value. At 0.1% blocking (99.9% peak acceptance rate), the network must maintain more residual capacity, so scaling stops at a lower offered traffic level. At 5% blocking (95% peak acceptance rate), the offered traffic can be scaled higher before the criterion is triggered, resulting in less residual capacity for grooming. The exact numerical savings may therefore change with the blocking-rate target. Nevertheless, the qualitative gap between ZR/ZR+ and long-haul muxponder reconfiguration benefit is expected to remain, because ZR/ZR+ router power can only decrease when an entire 400G port is shut down, regardless of the specific blocking-rate target used to set the peak-load level. For the Japan topology, this corresponds to approximately 147 Tb/s peak load in the opaque case and 141 Tb/s in the transparent case. For the German (G50) topology, the corresponding peak loads are approximately 301.6 Tb/s in the opaque case and 325.5 Tb/s in the transparent case. The reconfiguration interval is set to 2 hours ($T = 12$ cycles per day). This choice follows prior work in which a 24-hour daily traffic variation profile is divided into 12 two-hour periods [36], and it is consistent with energy-aware backbone studies showing that most sleep-mode savings can be captured with a limited number of daily configurations [55]. At the same time, multi-period reconfiguration studies emphasize that frequent reconfiguration increases rerouted traffic, signaling, QoS impact, and service-disruption risk [24, 56]. We therefore use 2 hours as a cost-performance tradeoff: it captures the broad AMS-IX daily trough and the associated energy-saving opportunity, while limiting reconfiguration cost in terms of control plane signaling, rerouted traffic, and disruption exposure. Daily energy is computed as the time-weighted integral of $P_{\text{total}}(t)$ over the 24-hour period.

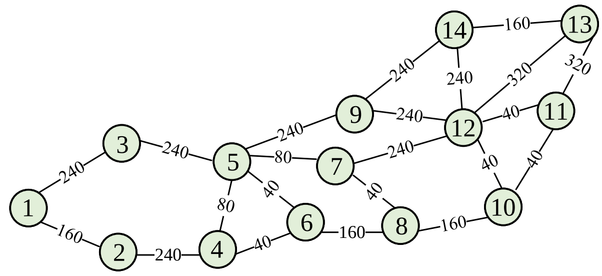

(a) 14-node Japan topology.

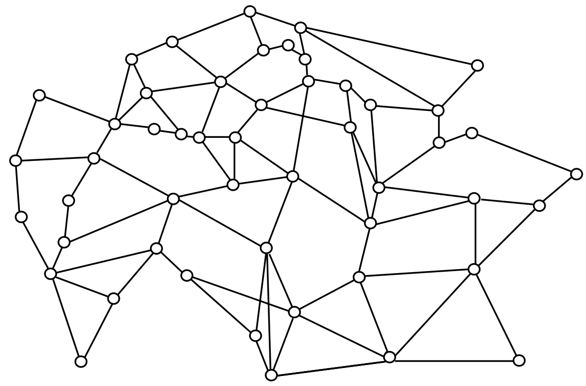

(b) 50-node German (G50) topology.

Figure 5 – Network topologies used in the simulation.

For the energy reduction percentage plots, we decompose the total daily savings into contributions from the three main component categories: transponders, IP router, and Optical Cross-Connect (OXC, encompassing I-ROADMs, ADBs, OAs, AWGs, and monitoring).

5.2 Impacts on energy consumption in Japan topology

5.2.1 Daily power profiles

Fig. 6(a) and Fig. 6(b) show the 24-hour power profiles for ZR/ZR+ and long-haul muxponders in the Japan opaque architecture. For ZR/ZR+, peak power is 13 178 units; the MFRE(ip+op) algorithm reduces this by 10.5% to 11 797 units at minimum-traffic hours. For long-haul muxponders, starting from 19 704 units at peak, MFRE achieves a 25.2% reduction to 14 739 units. ZR/ZR+ absolute power remains about 20–33% below long-haul over the Japan opaque profile.

IP-only rerouting RE(ip) and joint IP-optical rerouting

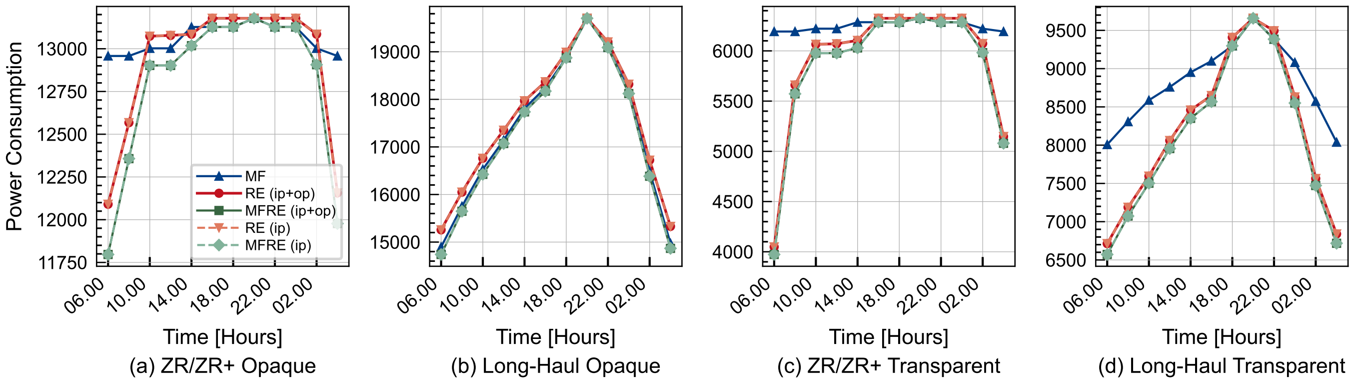

Figure 6 – Daily power consumption in the Japan topology for all four architecture-transceiver combinations.

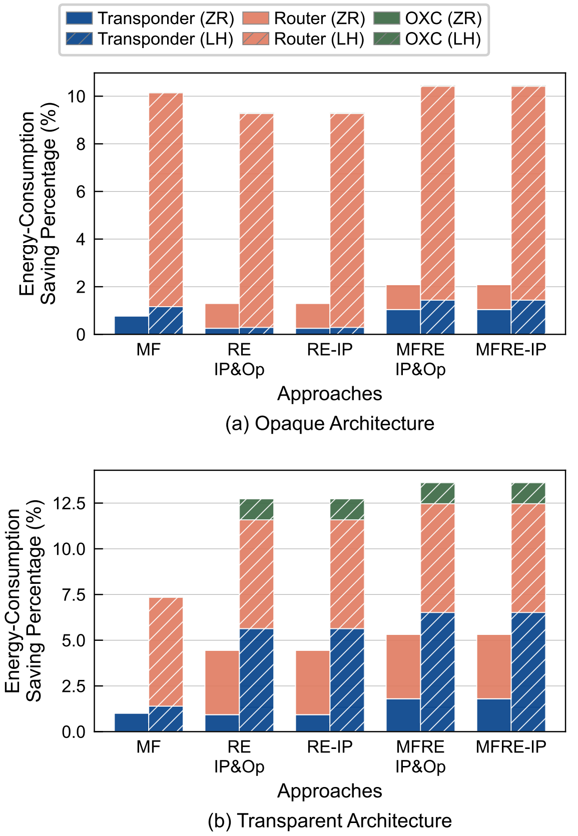

Figure 7 – Daily energy-consumption reduction percentages for the Japan topology, decomposed by component type.

RE(ip+op) yield nearly identical results in this topology. The Japan network is moderately dense (average node degree of roughly 3.1) and the IP layer already identifies near-optimal traffic-grooming paths, so optical-layer flexibility adds negligible savings. In the metro-regional topologies evaluated here, IP-layer rerouting alone captures essentially all available energy savings.

Figs. 6(c) and (d) show the transparent architecture results. Peak power for ZR/ZR+ is 6 324 units, nearly half the opaque value, confirming the benefit of eliminating per-hop OEO regeneration. The MFRE algorithm achieves up to 37.2% instantaneous power reduction at minimum traffic and about 5.3% daily energy savings. For long-haul muxponders, starting from 9 658 units, MFRE achieves a 32.0% instantaneous reduction at minimum traffic hours and about 13.6% daily energy savings (Fig. 7(b)), substantially exceeding ZR/ZR+.

The larger reconfiguration benefit for long-haul muxponders in the transparent architecture stems from their greater reach: as traffic decreases, long-haul lightpaths can extend across more hops without regeneration, allowing a single lightpath to absorb traffic previously served by multiple shorter paths with intermediate OEO. ZR/ZR+ cannot exploit this to the same degree because its shorter reach limits the maximum transparent span and the extent to which intermediate nodes can be bypassed.

5.2.2 *Energy savings decomposition*

Fig. 7(a) decomposes the daily energy savings in the Japan opaque architecture. ZR/ZR+ achieves about 2.1% daily energy reduction with MFRE, while long-haul muxponders achieve about 10.5%. Within the ZR/ZR+ total, the router and transponder contributions are both small, each around 1%. For long-haul muxponders, the router contribution dominates and reaches about 9–10%. The explanation lies directly in the router power model: for ZR/ZR+, a 400G port is either fully active or fully inactive; partial traffic reduction on a port releases no power. For

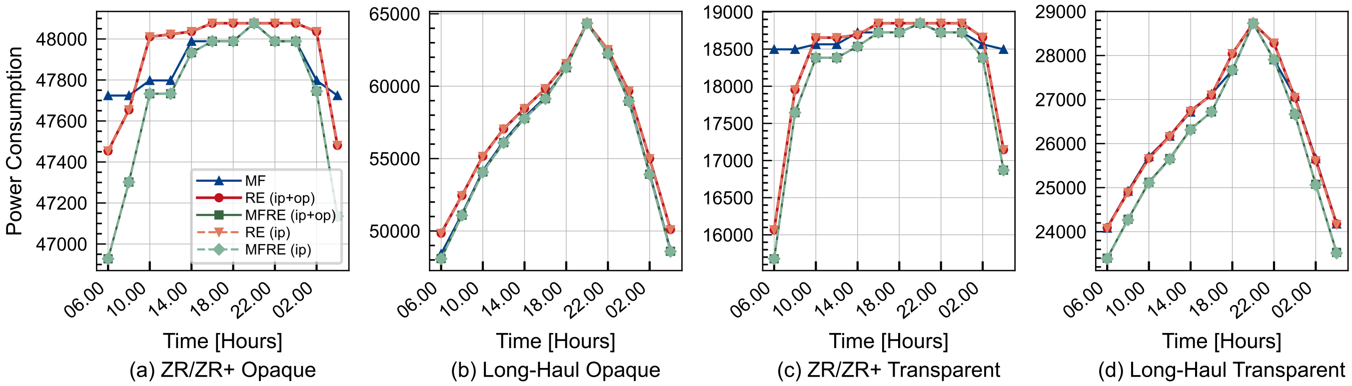

Figure 8 – Daily power consumption in the German (G50) topology for all four architecture-transceiver combinations.

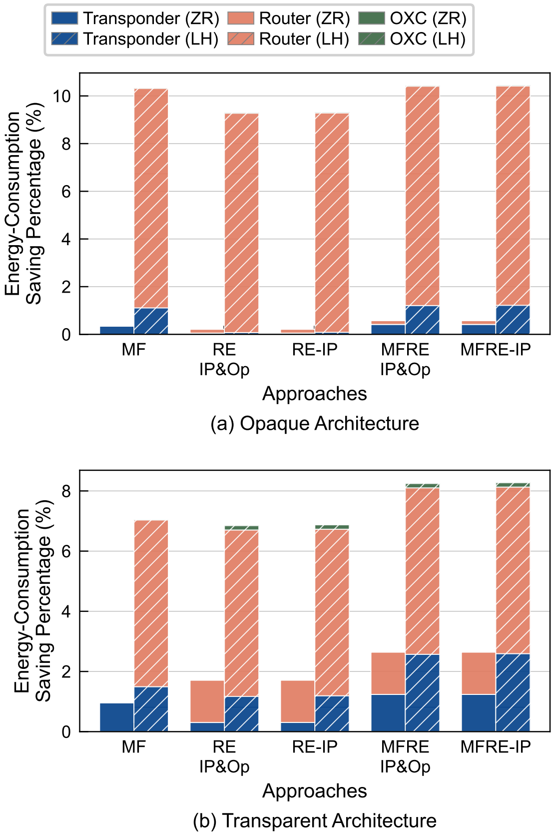

Figure 9 – Daily energy-consumption reduction percentages for the German (G50) topology, decomposed by component type.

long-haul muxponders, every 100 Gb/s of traffic reduction releases 1 unit of modular router power. Consequently, the smooth traffic decline from peak to minimum, which releases router power incrementally for long-haul, produces a step-wise power profile for ZR/ZR+ in which savings appear only when traffic becomes low enough to make an entire port unnecessary.

Fig. 7(b) shows the corresponding decomposition for the Japan transparent architecture. The router again dominates long-haul savings, while transponder and OXC reductions drive ZR/ZR+ savings, consistent with the opaque case.

5.3 Impacts on energy consumption in German (G50) topology

5.3.1 Daily power profiles

Fig. 8(a) and Fig. 8(b) show the 24-hour power profiles for ZR/ZR+ and long-haul muxponders in the German (G50) opaque architecture. ZR/ZR+ peak power is 48 077 units; MFRE reduces this by 2.4% to 46 928 units. For long-haul muxponders, the same strategy achieves a 25.3% reduction from 64 348 to 48 080 units. The markedly smaller relative saving for ZR/ZR+ in the G50 opaque case (0.6% daily, compared to 2.1% in Japan opaque) has two causes. ZR/ZR+ router power can only decrease when an entire 400G port is shut down, whereas long-haul muxponder router power decreases in 100 Gb/s increments whenever client capacity is released; partial traffic reduction on a ZR/ZR+ port yields no router power saving. In G50, clearing a full 400G port is particularly difficult: only about 7% of node pairs are directly connected (176 directed edges out of 2 450 directed pairs), compared with approximately 24% in Japan (44 directed edges out of 182 directed pairs), so most G50 traffic is multi-hop. When a demand spans multiple hops, the per-request algorithm evaluates each request in isolation and cannot coordinate the sequence of moves needed to empty all segments

Table 4 – Summary of power and energy savings across the evaluated topologies and architectures (MFRE strategy).

Topology	Architecture	Transceiver	Peak Power	Max Inst. Reduction	Daily Saving
Japan	Opaque	ZR/ZR+	13 178	10.5%	~2.1%
Japan	Opaque	Long-Haul	19 704	25.2%	~10.5%
Japan	Transparent	ZR/ZR+	6 324	37.2%	~5.3%
Japan	Transparent	Long-Haul	9 658	32.0%	~13.6%
German (G50)	Opaque	ZR/ZR+	48 077	2.4%	~0.6%
German (G50)	Opaque	Long-Haul	64 348	25.3%	~10.4%
German (G50)	Transparent	ZR/ZR+	18 848	16.8%	~2.6%
German (G50)	Transparent	Long-Haul	28 733	18.6%	~8.2%

of a port simultaneously; it therefore rarely clears a full port in G50. An algorithm that selects candidate 400G ports and reroutes all their traffic together would address this limitation; the remaining hardware constraint is that ZR/ZR+ power can only be reduced by turning off the entire 400G port. We leave the design and evaluation of such an algorithm to future work.

Fig. 8(c) and Fig. 8(d) show the transparent results for the German (G50) topology. Under ZR/ZR+, the initial peak power is 18 848 units; MFRE reduces this by 16.8% to 15 678 units. Under long-haul muxponders, the same strategy achieves 18.6% reduction from 28 733 to 23 384 units.

The step-wise energy saving behavior is particularly visible in the German (G50) topology: ZR/ZR+ power remains essentially constant from peak until the late-night low-traffic period, at which point it drops sharply. This contrasts with long-haul muxponders, which begin saving power as soon as traffic starts to decline. Again, the 400G port constraint is the mechanism: no power is released until an entire port can be turned off, which requires traffic to fall below the threshold that makes a complete port shutdown feasible.

5.3.2 *Energy savings decomposition*

Fig. 9 decomposes the daily energy savings for the German (G50) transparent architecture. ZR/ZR+ achieves 2.6% daily energy savings compared to 8.2% for long-haul muxponders. Of all network components, the router provides the largest savings, with reductions of up to 5.5% for long-haul muxponders and 1.5% for ZR/ZR+, reflecting the same constraint that ZR/ZR+ router power is only reduced when an entire 400G port is turned off.

5.4 Summary and discussion

Key findings. Table 4 summarizes the key metrics across all evaluated configurations. ZR/ZR+ has lower absolute power than long-haul muxponders in all evaluated cases; at peak load, the advantage ranges from 25.3% to 34.5%, while after MFRE at minimum traffic it ranges from 2.4% in the German (G50) opaque case to 39.5% in the Japan transparent case. The reconfiguration benefit is systematically smaller for ZR/ZR+: daily savings of about 0.6–5.3% versus about 8.2–13.6% for long-haul, driven by the constraint that ZR/ZR+ power is only reduced when an entire 400G port is turned off. For both evaluated topologies, IP-layer-only rerouting essentially captures all available reconfiguration savings, with optical-layer rerouting adding negligible benefit. Transparent architectures offer roughly 50% lower absolute power than opaque in the evaluated cases, and the combined MFRE strategy is the most effective reconfiguration approach across both topologies and architectures.

Sensitivity to traffic profile shape. The AMS-IX profile used in our evaluation has a pronounced daily variation, with off-peak traffic falling to approximately 50–60% of peak [21]. DCI and enterprise networks may show less pronounced day-night traffic variation due to workloads that run more continuously, such as storage replication, data synchronization, and AI training jobs. When off-peak traffic stays closer to its daytime peak, fewer 400G ports become eligible for shutdown and ZR/ZR+ reconfiguration savings are expected to decrease accordingly. Long-haul muxponders are less sensitive to how far traffic drops during off-peak hours, since their modular power scales continuously with traffic and even a modest traffic reduction yields proportional router power savings. This implies that the relative advantage of long-haul muxponders in reconfiguration benefit may be *more pronounced* in networks where traffic variation is limited, a finding that operators should factor into architecture selection.

Operational cost of reconfiguration. This paper focuses on the energy savings obtained from traffic-aware reconfiguration and does not explicitly model the operational cost of applying such reconfiguration. In practice, rerouting events may introduce transient traffic disruption, packet reordering, and additional control plane signaling overhead. These effects are particularly important in 6G backhaul deployments, where operators must balance energy savings against service-continuity requirements. Practical implementations can mitigate disruption risk through make-before-break or hitless/low-disruption reconfiguration techniques [57], where new paths are established before existing paths are released. Incorporating explicit reconfiguration-cost metrics, such as rerouted traffic volume, number of signaling actions, and disruption duration, is left as future work.

6. CONCLUSION

We studied traffic-aware network reconfiguration for energy minimization in IPoWDM networks with ZR/ZR+ pluggable transceivers and conventional long-haul muxponders. Five reconfiguration strategies, from modulation-format adaptation to joint IP and optical layer rerouting, were formalized and evaluated using auxiliary graph-based heuristic algorithms across the Japan and German (G50) topologies in both opaque and transparent configurations with real AMS-IX daily traffic profiles. ZR/ZR+ maintains lower absolute power consumption than long-haul muxponders in the evaluated metro-regional networks, with a 25.3–34.5% advantage at peak load. However, ZR/ZR+ achieves only about 0.6–5.3% daily energy savings through reconfiguration, compared to about 8.2–13.6% for long-haul muxponders, because ZR/ZR+ router power is released only when an entire 400G port can be turned off, whereas long-haul port power scales continuously with traffic. The joint MFRE strategy, which combines traffic grooming with modulation format downshifting, is the most effective approach that reliably turns off complete ZR/ZR+ ports, so reconfiguration controllers for ZR/ZR+-based networks should prioritize identifying when a full 400G port can be shut down. Future work will pursue optimization benchmarks to quantify the heuristic optimality gap [35] and AI-driven controllers that learn from historical traffic when ports can be turned off.

ACKNOWLEDGEMENT

This work was supported by the Italian Ministry of University and Research (MUR) and the European Union (EU) under the PON/REACT program. This work was also supported by the PRIN PNRR project GASTON, funded by the Italian Ministry of University and Research under grant number P2022XLRFH, Natural Science Foundation of China under Grant 62501345, and the Open Fund of IPOC (BUPT) under Grant IPOC2025A05.

REFERENCES

[1] Guidehouse Insights. *Market Data: Distributed Generation and Energy Storage in Telecom Networks*. https://www.businesswire.com/news/home/20210209005153/en/Guidehouse-Insights-Report-Finds-Telecom-Networks-Are-Expected-to-Install-122-GW-of-New-Distributed-Generation-and-Distributed-Energy-Storage-Capacity-from-2021-2030. Accessed: 2026-05-10. 2021.

[2] Jens Malmodin, Nina Lövehagen, Pernilla Bergmark, and Dag Lundén. "ICT sector electricity consumption and greenhouse gas emissions – 2020 outcome". In: *Telecommunications Policy* 48.3 (2024), p. 102701. DOI: 10.1016/j.telpol.2023.102701.

[3] Xin Wang, Fansong Kong, Ruikun Wang, Qiaolun Zhang, Jingjing Wang, and Massimo Tornatore. "All-Optical Network Coding-Enabled Efficient Collective Communication Over OTN-Interconnected Inter-DC Networks". In: *IEEE/Optica Journal of Lightwave Technology* 44.13 (2026), pp. 5464–5478. DOI: 10.1109/JLT.2026.3688671.

[4] Aryanaz Attarpour, Ehsan Etezadi, Memedhe Ibrahimi, Francesco Musumeci, Andrea Castoldi, Marija Furdek, and Massimo Tornatore. "On the Benefits of Programmable FON for Low-Cost OTN-Over-WDM Metro Networks". In: *2025 International Conference on Optical Network Design and Modeling (ONDM)*. IEEE. 2025, pp. 1–5.

[5] Ruikun Wang, Jiawei Zhang, Zhiqun Gu, Memedhe Ibrahimi, Bojun Zhang, Francesco Musumeci, Massimo Tornatore, and Yuefeng Ji. "Digital-twin-assisted meta learning for soft-failure localization in ROADM-based optical networks". In: *Journal of Optical Communications and Networking* 16.7 (2024), pp. C11–C19. DOI: 10.1364/JOCN.517551.

[6] Hafiz Faheem Shahid, Bilgehan Akdemir, Johirul Islam, Ijaz Ahmad, Ijaz Ahmad, and Erkki Harjula. "IoT Service Orchestration in Edge–Cloud Continuum With 6G: A Review". In: *IEEE Internet of Things Journal* 13.10 (2026), pp. 20339–20358. DOI: 10.1109/JIOT.2026.3663918.

[7] Qiang He, Jinqiu Lin, Hui Fang, Xingwei Wang, Min Huang, Xiushuang Yi, and Keping Yu. "Integrating IoT and 6G: Applications of Edge Intelligence, Challenges, and Future Directions". In: *IEEE Transactions on Services Computing* 18.4 (2025), pp. 2471–2488. DOI: 10.1109/TSC.2025.3586152.

[8] S. Bandiri, R. Braga, T. Pimenta, and D. Spadoti. "Energy consumption improvement based on distance adaptive modulation in elastic optical network". In: *IEEE International Conference on Communications, Control, and Computing Technologies (ICCDCS)*. June 2017, pp. 29–32. DOI: 10.1109/ICCDCS.2017.7959702.

[9] A. Valkanis, G. Papadimitriou, G. Beletsioti, E. Varvarigos, and P. Nicopolitidis. "Efficiency and fairness improvement for elastic optical networks using reinforcement learning-based traffic prediction". In: *IEEE/OSA Journal of Optical Communications and Networking* 14 (2022), pp. 25–42.

[10] Optical Internetworking Forum (OIF). *400ZR*. https://www.oiforum.com/technical-work/400zr/. Accessed on 10/05/2026. 2023.

[11] OpenZR+. *Multi-Vendor 400G Coherent Optical Transceiver Interoperability Testing*. https://www.openzrplus.org/documents/. Accessed on 10/05/2026. 2023.

[12] Qiaolun Zhang, Annalisa Morea, Patricia Layec, Memedhe Ibrahimi, Francesco Musumeci, and Massimo Tornatore. "Power-consumption analysis for different IPoWDM network architectures with ZR/ZR+ and long-haul muxponders". In: *Journal of Optical Communications and Networking* 16.12 (2024), pp. 1189–1203.

[13] Qiaolun Zhang, Annalisa Morea, and Massimo Tornatore. "A Pragmatic Power-Consumption Analysis for IPoWDM Networks with ZR/ZR+ Modules". In: *2022 European Conference on Optical Communication (ECOC)*. 2022, pp. 1–4.

[14] Serge Melle, Thierry Zami, Oriol Bertran-Pardo, and Bruno Lavigne. "Comparing IP-optical architectures & WDM transport technologies in metro, regional and long-haul networks". In: *Optical Fiber Communications Conference (OFC)*. IEEE. 2021, pp. 1–3.

[15] 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. DOI: 10.1109/DRCN51925.2021.9428437.

[16] João Pedro, Mohammad M. Hosseini, and Antonio Napoli. "Extended network applications of coherent pluggable transceivers [Invited]". In: *Journal of Optical Communications and Networking* 17.2 (2025), A210–A223. DOI: 10.1364/JOCN.537601.

[17] Paul Wright, Russell Davey, and Andrew Lord. "Cost Model Comparison of ZR/ZR+ Modules Against Traditional WDM Transponders for 400G IP/WDM Core Networks". In: *2020 European Conference on Optical Communications (ECOC)*. 2020, pp. 1–4. DOI: 10.1109/ECOC48923.2020.9333420.

[18] 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". In: *J. Opt. Commun. Netw.* 14.3 (2022), pp. 127–139. DOI: 10.1364/JOCN.444943.

[19] 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: *2021 European Conference on Optical Communication (ECOC)*. 2021, pp. 1–4. DOI: 10.1109/ECOC52684.2021.9606050.

[20] R. P. Davey, M. A. Iqbal, and P. D. Wright. "ZR 400 Gbit/s and 800 Gbit/s use cases, trials, deployments, and future prospects [Invited]". In: *Journal of Optical Communications and Networking* 16.1 (2024), A33–A39. DOI: 10.1364/JOCN.499365.

[21] Amsterdam Internet Exchange (AMS-IX). *AMS-IX Traffic Statistics*. 2024. URL: https://stats.ams-ix.net/index.html (visited on 05/10/2026).

[22] M. Hadi and M.R. Pakravan. "Energy-Efficient Fast Configuration of Flexible Transponders and Grooming Switches in OFDM-Based Elastic Optical Networks". In: *Journal of Optical Communications and Networking* 10.2 (2018), pp. 90–103. DOI: 10.1364/JOCN.10.000090.

[23] Filip Idzikowski, Sebastian Orlowski, Christian Raack, Hagen Woesner, and Adam Wolisz. "Dynamic routing at different layers in IP-over-WDM networks — Maximizing energy savings". In: *Optical Switching and Networking* 8.3 (2011), pp. 181–200. DOI: 10.1016/j.osn.2011.03.007.

[24] Edoardo Bonetto, Luca Chiaraviglio, Filip Idzikowski, and Esther Le Rouzic. "Algorithms for the Multi-Period Power-Aware Logical Topology Design With Reconfiguration Costs". In: *Journal of Optical Communications and Networking* 5.5 (2013), pp. 394–410. DOI: 10.1364/JOCN.5.000394.

[25] Nisanur Camuzcu, Qiaolun Zhang, Memedhe Ibrahimi, Ruikun Wang, Lu Zhang, and Massimo Tornatore. "Minimizing power consumption in ipowdm networks with zr/zr+ using network reconfiguration". In: *2024 Asia Communications and Photonics Conference (ACP) and International Conference on Information Photonics and Optical Communications (IPOC)*. IEEE. 2024, pp. 1–6.

[26] Jayant Baliga, Robert Ayre, Kerry Hinton, Wayne V. Sorin, and Rodney S. Tucker. "Energy Consumption in Optical IP Networks". In: *J. Lightwave Technol.* 27.13 (July 2009), pp. 2391–2403.

[27] F. Musumeci, M. Tornatore, and A. Pattavina. "A power consumption analysis for IP-over-WDM core network architectures". In: *IEEE/OSA Journal of Optical Communications and Networking* 4 (2012), pp. 108–117.

[28] Gangxiang Shen and Rodney S. Tucker. "Energy-Minimized Design for IP Over WDM Networks". In: *Journal of Optical Communications and Networking* 1.1 (2009), pp. 176–186. DOI: 10.1364/JOCN.1.000176.

[29] Serge Melle, Rick Dodd, Chris Liou, Drew Perkins, Marco Sosa, and Mark Yin. "Network Planning and Economic Analysis of an Innovative New Optical Transport Architecture: The Digital Optical Network". In: *Optical Fiber Communication Conference and Exposition and The National Fiber Optic Engineers Conference*. Optica Publishing Group, 2005, NTuA1. URL: https://opg.optica.org/abstract.cfm?URI=NFOEC-2005-NTuA1.

[30] Abhijeet Deore, Onur Turkcu, Satyajeet Ahuja, Steven J. Hand, and Serge Melle. "Total cost of ownership of WDM and switching architectures for next-generation 100Gb/s networks". In: *IEEE Communications Magazine* 50.11 (2012), pp. 179–187. DOI: 10.1109/MCOM.2012.6353699.

[31] S. Nedevschi, L. Popa, G. Iannaccone, S. Ratnasamy, and D. Wetherall. "Reducing Network Energy Consumption via Sleeping and Rate-adaptation". In: *NSDI*. Vol. 8. 2008, pp. 323–336.

[32] Giuseppe Rizzelli, Annalisa Morea, Massimo Tornatore, and Olivier Rival. "Energy efficient Traffic-Aware design of on-off Multi-Layer translucent optical networks". In: *Computer Networks* 56.10 (2012), pp. 2443–2455. DOI: 10.1016/j.comnet.2012.03.013.

[33] Margita Radovic, Andrea Sgambelluri, Filippo Cugini, and Nicola Sambo. "Power-aware high-capacity elastic optical networks". In: *Journal of Optical Communications and Networking* 16.5 (2024), B16–B25. DOI: 10.1364/JOCN.514067.

[34] H. Khodakarami, B.S.G. Pillai, B. Sedighi, and W. Shieh. "Flexible Optical Networks: An Energy Efficiency Perspective". In: *J. Lightwave Technol.* 32.16 (2014), pp. 3356–3367. DOI: 10.1109/JLT.2014.2339637.

[35] Yanxia Tan, Rentao Gu, and Yuefeng Ji. "Energy-efficient routing, modulation and spectrum allocation in elastic optical networks". In: *Optical Fiber Technology* 36 (2017), pp. 297–305. ISSN: 1068-5200. DOI: https://doi.org/10.1016/j.yofte.2017.05.001. URL: https://www.sciencedirect.com/science/article/pii/S106852001730038X.

[36] Zhizhen Zhong, Nan Hua, Massimo Tornatore, Yao Li, Haijiao Liu, Chen Ma, Yanhe Li, Xiaoping Zheng, and Biswanath Mukherjee. "Energy Efficiency and Blocking Reduction for Tidal Traffic via Stateful Grooming in IP-Over-Optical Networks". In: *Journal of Optical Communications and Networking* 8.3 (2016), pp. 175–189. DOI: 10.1364/JOCN.8.000175.

[37] I. Turus, A.M. Fagertun, L. Dittmann, A. Morea, D. Verchere, and J. Kleist. "Energy savings in dynamic and resilient optical networks based on traffic-aware strategies". In: *2014 16th International Telecommunications Network Strategy and Planning Symposium (Networks)*. 2014, pp. 1–6. DOI: 10.1109/NETWKS.2014.6959230.

[38] N. Sambo, C. Castro, N. Costa, P. Castoldi, and A. Napoli. "Energy Efficiency in Next-generation Optical Networks". In: *2023 23rd International Conference on Transparent Optical Networks (ICTON)*. 2023, pp. 1–4. DOI: 10.1109/ICTON59386.2023.10207497.

[39] Romain Jacob, Jackie Lim, and Laurent Vanbever. "Does rate adaptation at daily timescales make sense?" In: *Proceedings of the 2nd Workshop on Sustainable Computer Systems*. 2023, pp. 1–7.

[40] Byrav Ramamurthy, Helena Feng, Debasish Datta, Jonathan P. Heritage, and Biswanath Mukherjee. "Transparent vs. Opaque vs. Translucent Wavelength-Routed Optical Networks". In: *CSE Conference and Workshop Papers, Computer Science and Engineering Department*. University of Nebraska - Lincoln, DigitalCommons@University of Nebraska - Lincoln. 1999. URL: https://digitalcommons.unl.edu/cgi/viewcontent.cgi?article=1028&context=cseconfwork.

[41] Thierry Zami, Bruno Lavigne, and Oriol Bertran-Pardo. "Benchmarking of Opaque Versus Transparent Core WDM Networks Featuring 400ZR+ QSFP-DD or CFP2 Interfaces". In: *2020 European Conference on Optical Communications (ECOC)*. IEEE, 2020, pp. 1–4. DOI: 10.1109/ECOC48923.2020.9333414.

[42] J. Zhang, Y. Ji, M. Song, Y. Zhao, X. Yu, J. Zhang, and B. Mukherjee. "Dynamic Traffic Grooming in Sliceable Bandwidth-Variable Transponder-Enabled Elastic Optical Networks". In: *IEEE Journal of Lightwave Technology* 33.1 (2015), pp. 183–191. DOI: 10.1109/JLT.2014.2383444.

[43] Qingcheng Zhu, Xiaosong Yu, Yongli Zhao, Avishek Nag, and Jie Zhang. "Auxiliary-Graph-Based Energy-Efficient Traffic Grooming in IP-Over-Fixed/Flex-Grid Optical Networks". In: *Journal of Lightwave Technology* 39.10 (2021), pp. 3011–3024. DOI: 10.1109/JLT.2021.3057389.

[44] E. Pincemin and Y. Loussouarn. "Silicon Photonic ZR/ZR+ DCO-CFP2 Interface for DCI and Metro-Regional 400G Optical Communications". In: *Optical Fiber Communication Conference (OFC)*. 2021. DOI: 10.1364/OFC.2021.T2C.1.

[45] Tony Dicorato, Peter Landon, Pino G. Dicorato, Swamynathan Balasundaram, Matias Schneeberger, Luca Baragiola, and Victor Lopez. "Enabling IP-optical integration in core and metro networks [Invited]". In: *Journal of Optical Communications and Networking* 16.7 (2024), pp. C121–C132. DOI: 10.1364/JOCN.516740.

[46] Rasoul Sadeghi, Bruno Correia, André Souza, Nelson Costa, João Pedro, Antonio Napoli, and Vittorio Curri. "Transparent vs Translucent Multi-Band Optical Networking: Capacity and Energy Analyses". In: *Journal of Lightwave Technology* 40.11 (2022), pp. 3486–3498. DOI: 10.1109/JLT.2022.3168455.

[47] Rasoul Sadeghi et al. "Comparison of Transceiver and C+ L Band Upgrades: Network Traffic and Energy Assessment". In: *2021 International Conference on Electrical, Communication, and Computer Engineering (ICECCE)*. IEEE. 2021, pp. 1–6.

[48] Ricardo Martínez, Carlos Hernández-Chulde, Ramon Casellas, Ricard Vilalta, and Raül Muñoz. "Deep reinforcement learning for energy-efficient RMSA in IPoWDM networks with coherent ZR+ transceivers [Invited]". In: *Journal of Optical Communications and Networking* 18.2 (2026), A123–A133. DOI: 10.1364/JOCN.574251.

[49] Ruikun Wang, Qiaolun Zhang, Jiawei Zhang, Zhiqun Gu, Memedhe Ibrahimi, Hao Yu, Bojun Zhang, Francesco Musumeci, Yuefeng Ji, and Massimo Tornatore. "Multi-Failure Localization in High-Degree ROADM-Based Optical Networks Using Rules-Informed Neural Networks". In: *IEEE Journal on Selected Areas in Communications* 43.5 (2025), pp. 1738–1754. DOI: 10.1109/JSAC.2025.3543482.

[50] Chengzhi Song, Rongxi He, Ze Wei, and Xiaojing Chen. "Auxiliary-Graph-Based Traffic Grooming With Hybrid Single/Multipath Routing in Mixed-Grid Optical Networks". In: *Journal of Lightwave Technology* 43.12 (2025), pp. 5653–5668. DOI: 10.1109/JLT.2025.3553303.

[51] Menglin Liu, Massimo Tornatore, and Biswanath Mukherjee. "Survivable Traffic Grooming in Elastic Optical Networks—Shared Protection". In: *Journal of Lightwave Technology* 31.6 (2013), pp. 903–909. DOI: 10.1109/JLT.2012.2231663.

[52] Chaitanya S. K. Vadrevu, Rui Wang, Massimo Tornatore, Charles U. Martel, and Biswanath Mukherjee. "Degraded Service Provisioning in Mixed-Line-Rate WDM Backbone Networks Using Multipath Routing". In: *IEEE/ACM Transactions on Networking* 22.3 (2014), pp. 840–853. DOI: 10.1109/TNET.2013.2259638.

[53] Ramon Casellas, Ricardo Martínez, Raül Muñoz, Ricard Vilalta, Lei Liu, Takehiro Tsuritani, and Itsuro Morita. "Control and management of flexi-grid optical networks with an integrated stateful path computation element and OpenFlow controller". In: *Journal of Optical Communications and Networking* 5.10 (2013), A57–A65.

[54] Sebastian Orlowski, Roland Wessäly, Michal Pióro, and Artur Tomaszewski. "SNDlib 1.0 — Survivable Network Design Library". In: *Networks* 55.3 (2010), pp. 276–286. DOI: 10.1002/net.20371.

[55] Luca Chiaraviglio, Antonio Cianfrani, Esther Le Rouzic, and Marco Polverini. "Sleep Modes Effectiveness in Backbone Networks with Limited Configurations". In: *Computer Networks* 57.15 (2013), pp. 2931–2948. DOI: 10.1016/j.comnet.2013.06.016.

[56] Rodolfo Alvizu, Xingyu Zhao, Guido Maier, Yajing Xu, and Achille Pattavina. "Energy Efficient Dynamic Optical Routing for Mobile Metro-Core Networks Under Tidal Traffic Patterns". In: *Journal of Lightwave Technology* 35.2 (2017), pp. 325–333. DOI: 10.1109/JLT.2016.2638739.

[57] Chengzhi Song, Rongxi He, Xiaojing Chen, and Ze Wei. "Make-before-break based spectrum defragmentation in mixed-grid optical networks". In: *Optical Fiber Technology* 88 (2024), p. 103838.

AUTHORS

QIAOLUN ZHANG received a Ph.D. from the Department of Electronics, Information and Bioengineering at Politecnico di Milano, where he is currently an assistant professor at Politecnico di Milano. His research interests cover a broad spectrum of topics in resource allocation for communications networks, with a focus on optimization and machine learning approaches.

NISANUR CAMUZCU received her M.Sc. degree in telecommunication engineering from Politecnico di Milano in 2024, where she worked on optimization in optical networks. She is currently pursuing a Ph.D. at Rochester Institute of Technology, with research interests in communication networks and localization.

XIAOYANG GUO is currently pursuing an M.Sc. degree in telecommunication engineering at Politecnico di Milano. Her research interests include optical network design and optimization, communication network resource allocation and management, and intelligent optimization for sustainable computing infrastructure.

JIAHENG XIONG received his M.Sc. degree from Politecnico di Milano and is currently pursuing a Ph.D. in information technology in the Department of Electronics, Information and Bioengineering at Politecnico di Milano. His research interests include distributed AI systems, optical network design and optimization, quantum key distribution networks, etc.

RUIKUN WANG received a Ph.D. from Beijing University of Posts and Telecommunications. He is currently an associate professor at Qingdao University of Science and Technology. His research interests are in the fields of machine-learning-based failure management and network optimization, which focuses on underwater communication networks.

TIANQU LUO is currently an undergraduate student in Mathematical Engineering at Politecnico di Milano (PoliMi), Italy. He is also a research assistant at the Broadband Optical Networks, Security and Advanced Internet (BONSAI) Lab in the Department of Electronics, Information and Bioengineering at Politecnico di Milano. His research interests focus on network design and optimization, with a particular emphasis on Quantum Key Distribution (QKD) networks.

MËMËDHE IBRAHIMI received his Ph.D. from Politecnico di Milano in 2022. He is currently an assistant professor in the Department of Electronics, Information and Bioengineering at Politecnico di Milano. His research interests include optical communications, optical network design and optimization, and machine learning for network management.

MASSIMO TORNATORE is a full professor in the Department of Electronics, Information and Bioengineering at Politecnico di Milano, and has also held positions as an adjunct professor at the University of California, Davis and as a visiting professor at the University of Waterloo, Canada. He is a fellow of the IEEE. His research interests include design and optimization of optical networks, network resilience, and machine learning for network management. He has co-authored over 500 peer-reviewed papers, earning 25 best paper awards, and serves on the editorial boards of IEEE Transactions on Network and Service Management, IEEE Transactions on Networking, IEEE Communication Surveys and Tutorials, and Elsevier Optical Switching and Networking.