

Dual-Stage Planning for Elastic Optical Networks Integrating Machine-Learning-Assisted QoT Estimation

Original

Dual-Stage Planning for Elastic Optical Networks Integrating Machine-Learning-Assisted QoT Estimation / Salani, M; Rottondi, C; Cere, L; Tornatore, M. - In: IEEE-ACM TRANSACTIONS ON NETWORKING. - ISSN 1063-6692. - ELETTRONICO. - 31:3(2023), pp. 1293-1307. [10.1109/TNET.2022.3213970]

Availability:

This version is available at: 11583/2981206 since: 2023-08-23T13:21:29Z

Publisher:

IEEE - ACM

Published

DOI:10.1109/TNET.2022.3213970

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2023 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Dual-Stage Planning for Elastic Optical Networks Integrating Machine-Learning-Assisted QoT Estimation

Matteo Salani¹, Cristina Rottondi², Leopoldo Cerè¹ and Massimo Tornatore³

Abstract—Following the emergence of Elastic Optical Networks (EONs), Machine Learning (ML) has been intensively investigated as a promising methodology to address complex network management tasks, including, e.g., Quality of Transmission (QoT) estimation, fault management, and automatic adjustment of transmission parameters. Though several ML-based solutions for specific tasks have been proposed, how to integrate the outcome of such ML approaches inside Routing and Spectrum Assignment (RSA) models (which address the fundamental planning problem in EONs) is still an open research problem.

In this study, we propose a dual-stage iterative RSA optimization framework that incorporates the QoT estimations provided by a ML regressor, used to define lightpaths' reach constraints, into a Mixed Integer Linear Programming (MILP) formulation. The first stage minimizes the overall spectrum occupation, whereas the second stage maximizes the minimum inter-channel spacing between neighbor channels, without increasing the overall spectrum occupation obtained in the previous stage. During the second stage, additional interference constraints are generated, and these constraints are then added to the MILP at the next iteration round to exclude those lightpaths combinations that would exhibit unacceptable QoT. Our illustrative numerical results on realistic EON instances show that the proposed ML-assisted framework achieves spectrum occupation savings up to 52.4% (around 33% on average) in comparison to a traditional MILP-based RSA framework that uses conservative reach constraints based on margined analytical models.

Index Terms—Machine Learning; Routing and Spectrum Assignment; QoT estimation;

I. INTRODUCTION

To cope with the never-ceasing growth of Internet traffic [1], optical network design and management tools must constantly evolve and incorporate new technical advances. In particular, the new paradigm of *Elastic Optical Networks* (EONs) is characterized by the adoption of coherent transmission technologies [2] (enabling higher spectral efficiency, e.g., by means of adjustable modulation formats) and of a flexible spectrum grid [3] that partitions the optical spectrum in fine-grained frequency slots to better adapt channel bandwidths to traffic requests. Overall, EONs achieve a much more effective spectrum utilization than traditional Wavelength Division

Multiplexed (WDM) networks [4], but they come with a much larger number of configurable transmission parameters that sternly increase complexity of network management. The most relevant of these parameters are:

- modulation format, baud rate and forward error correction (FEC) code to be used for transmission.
- optical channel width and inter-channel spacing (i.e., spectral width between spectrally-adjacent channels), which can now be finely tuned thanks to the flex-grid paradigm.

The fundamental EON planning problem is the Routing and Spectrum Assignment (RSA), which consists in assigning to every traffic demand a lightpath (i.e., an all-optical end-to-end circuit) connecting its source node to its destination node over a commensurate spectral bandwidth. To check the feasibility of an EON RSA solution, the Quality of Transmission (QoT) of each candidate lightpath must be estimated prior to its deployment, to evaluate which modulation format can be safely used. In fact, the QoT of a lightpath is affected by physical-layer impairments such as amplifier spontaneous noise emission (ASE) and non-linear impairments (NLI) along the route from source to destination. Hence, QoT estimation is a fundamental preliminary phase, that is typically conducted before running the RSA algorithm. State-of-the-art techniques for QoT estimation usually rely on analytical models (such as the Gaussian Noise (GN) model [5]), whose estimation accuracy substantially depends on how accurate is the knowledge of transmission parameters that characterize the physical layer, such as the type of used fiber, the noise figure of amplifiers, the amount of interference generated by spectrally-adjacent channels etc. Unfortunately, several of these parameters are unlikely to be precisely known in real network deployments, e.g., due to an incomplete equipment inventory, to deviations from datasheet specs caused by aging, or to the unavoidable ripple effect in amplifiers [6]. Inaccurate estimations then translate in the need of adding conservative estimation margins to account for uncertainties that, in turn, lead to significant underutilization of network resources¹.

Alternative QoT estimation approaches (as those based on the Split Step Fourier Transform [8]) rely on direct simulation of the signal propagation along the optical path and can

¹M. Salani and L. Cerè are with the Dalle Molle Institute for Artificial Intelligence, University of Lugano (USI), University of Applied Science and Arts of Southern Switzerland (SUPSI) – Switzerland matteo.salani@supsi.ch

²C. Rottondi is with the Dept. of Electronics and Telecommunications, Politecnico di Torino, Italy cristina.rottondi@polito.it

³M. Tornatore is with the Dept. of Electronics, Information and Bioengineering, Politecnico di Milano, Italy massimo.tornatore@polimi.it

¹Ref. [7] reports that throughput gains from 25% up to 300% - depending on the network topology - can be achieved by avoiding margins traditionally introduced to account for the uncertainty of physical-layer parameters.

achieve accurate estimations, but incur too heavy computational burden to ensure scalable and real-time decision making.

Recently, Machine Learning (ML) has been extensively investigated as alternative methodology for QoT estimation, potentially capable of avoiding the shortcomings of existing techniques. ML directly leverages the knowledge extracted from field data (e.g., the pre-FEC Bit Error Rate (BER), or the Optical Signal to Noise Ratio (OSNR) of existing lightpaths acquired by means of Optical Performance Monitors (OPMs) [9]). Once trained with such data, supervised ML models are then able to predict the QoT associated to previously unseen lightpath configurations with very high accuracy and negligible computational time. Low prediction timings are particularly desirable when a wide range of what-if scenarios with large network topologies need to be evaluated in constrained time, e.g. for dynamic resource allocation or fault restoration purposes.

Though several studies in the field of ML-based QoT estimation and RSA have already appeared (see next Section), the two tasks have been typically addressed separately, while frameworks that integrate QoT prediction tools in RSA algorithms have started appearing only recently. In particular, coupling black-box ML models with explicit mathematical formulation of combinatorial problems is an open research topic (see [10]). This study moves towards the integration of ML-based QoT estimation inside the optimization tools used to solve RSA by explicitly adding constraints to avoid infeasible solutions, by considering different features for the ML-based QoT estimator, and by modeling secondary objective functions that are eliciting promising solutions among equivalent ones, for the sake of QoT satisfaction. We propose a dual-stage iterative framework that ingests the QoT predictions of a ML regressor as input for the definition of iteratively added interference constraints in a Mixed Integer Linear Program (MILP) formulation for RSA, while considering multiple modulation formats. During the first stage, a preliminary solution is found with the aim of minimizing the overall spectrum occupation. At the second stage, an enhanced version of the regressor (that considers features characterizing the neighbor channels of each lightpath, information which was not available initially) is invoked to improve the QoT estimation by taking into account inter-channel interference. This allows to generate additional interference constraints that are added to the MILP formulation to prevent the deployment of neighboring channels that would lead to unacceptable QoT levels. Moreover, in the second stage, the objective of the optimization procedure is the maximization of the minimum inter-channel spacing between two spectrally adjacent lightpaths, without increasing the spectrum occupation value achieved in the first stage. We show that, in our case study, the proposed framework saves up to 52% in spectrum occupation with respect to a traditional RSA solution with margined reach computations.

The rest of the manuscript is structured as follows. Related scientific literature is overviewed in Section II, whereas Section III reports some background notions. In Section IV the proposed iterative optimization framework is described, whereas the MILP model that solves the Routing, Modulation format and Spectrum Assignment (RMSA) problem and the

integration of MILP formulation and QoT classifier are described in Section V and assessed in Section VI. Finally, the last Section offers some conclusive remarks.

II. RELATED WORK

A. Machine Learning for RSA

In the last few years, ML has been widely investigated as alternative methodological approach to accomplish a variety of tasks required in optical network design, automation and management. For a thorough review, the reader is referred to the survey papers [11]–[13]. Focusing on RSA applications, ML approaches have been devised to address the following tasks:

a) *QoT estimation in support of RSA*: QoT estimation can be applied either to forecast the transmission quality of unestablished lightpaths by relying on observations collected from the already deployed ones, or to monitor QoT excursions of existing lightpaths for fault detection/identification scopes. In this paper, we focus on the former scenario, where the prediction problem can be formulated either as a binary classification problem (i.e., the ML tool outputs a yes/no answer, depending on whether the quality metrics of the lightpath candidate for deployment are expected to satisfy a predefined system threshold), or as a regression problem (i.e., the ML tool outputs the predicted numerical value for the considered quality metric). The ML outputs can be ingested by a network planning framework to take informed decisions. The ML framework adopted in this study formulates the QoT estimation task as a regression problem.

b) *Lightpath routing*: finding the best path in a graph is a well-known optimization problem and several ML frameworks devoted to such scope have recently appeared. The vast majority of them adopt graph-based neural networks, which are specifically designed to embed the topological characteristics of a graph in a latent space. Chapter 5 of the survey paper [14] offers a complete overview on such techniques.

c) *Spectrum assignment*: once lightpath routing is defined, spectrum/wavelength assignment can still be performed by means of ML approaches: Ref. [15] proposes a ML-based wavelength reconfiguration method aimed at assigning wavelengths to minimize channel power excursions during optical circuit switching. The trained neural network is capable of recommending wavelength assignments exhibiting power excursions below 0.5 dB with above 99% precision.

d) *Joint routing and spectrum assignment*: The two tasks of identifying a suitable route and a commensurate spectrum portion for a traffic demand can be addressed by a ML framework even jointly. Authors in [16] propose to convert the standard Routing and Wavelength Assignment (RWA) problem in a supervised classification problem. Given as input traffic matrix in an optical WDM network, either a logistic regressor or a deep neural network maps each traffic demand to a sequence of links and wavelength assignments. The study shows that nearly optimal solutions can be obtained, while drastically reducing computational time in comparison to a traditional ILP approach.

In [17], a deep neural network is used to implement RSA strategy selection for a given traffic demand, under preliminary

calculation of a predefined number of shortest paths. The ML framework also takes into account a spectrum fragmentation index. Results show that the achieved blocking probability and spectrum fragmentation are both lower than those obtained by a traditional combination of shortest path selection plus first fit spectrum allocation.

Moreover, reinforcement learning frameworks for routing, modulation format and spectrum assignment have been devised for dynamic network scenarios [18], also implementing survivable design approaches [19].

All the references mentioned so far (except [16]) consider heuristic design approaches for dynamic traffic scenarios, which do not provide any guarantee of optimality. Differently, this study focuses on a static traffic scenario and ensures minimum spectrum occupation, assuming the availability of multiple modulation formats. With respect to [16], which considers a traditional WDM grid, we model a flexi-grid and use dedicated constraints to capture the interference caused by spectrally-adjacent channels: the QoT estimator is called iteratively to enrich a MILP formulation for RMSA, whose constraints gradually gain knowledge on interference caused by neighbor channels, based on the optimal network configuration calculated at the previous round. Moreover, all the studies above do not consider the optimization of inter-channel spacing between spectrally adjacent optical channels, whereas in our dual-stage framework we optimize the inter-channel spacing with the aim of minimizing inter-channel interference without increasing the index of the rightmost occupied slice in the whole network. Note that, though advanced exact algorithms based on Dantzig-Wolfe decomposition and column generation have been proposed in [20] and [21], the purpose of the present study is not to provide a faster and more scalable algorithm for RSA, but rather to move some steps towards the integration of ML-based models in exact or heuristic algorithms for RSA.

B. Integration of QoT predictors in RSA frameworks

Several approaches to solve RSA in EONs have been investigated in the past decade (see [22], [23] for a recent survey). A consistent literature body on impairment-aware RSA is also available, though the vast majority of the QoT estimation methods therein adopted do not leverage ML, but analytical formulas such as those based on the GN model. Indeed, only a few studies have investigated how to integrate the outputs of ML-based QoT estimators in traditional RSA frameworks. Table I provides a taxonomy of QoT estimation and RSA methodologies and categorizes illustrative references, that are briefly reviewed in the remainder of this Section. In the following, we will adopt the acronyms used in the Table to identify the six categories therein reported, i.e., ML-HA, AM-HA, ML-LP, AM-LP, ML-RL, AM-RL. Note that, depending on the specific QoT estimation method and RSA algorithm being adopted, the interaction between the two methodological components may vary significantly. More in detail, from the functional point of view we can categorize the types of interactions in three broad groups:

- *Group 1*: QoT-aware RSA approaches where QoT estimations (mainly, signal reaches) are pre-calculated and

leveraged as inputs to RSA algorithms (either heuristics or linear programs) with no further involvement of the QoT estimation tools during the RSA computation phase;

- *Group 2*: QoT-aware reinforcement learning frameworks for RSA, in which QoT estimations are instrumental for the learning phase of a reinforcement learning algorithm to compute the agent's rewards, but are not considered during the execution of the RSA algorithm itself;
- *Group 3*: iterative approaches where QoT estimations are given as inputs to RSA heuristic algorithms or linear programs and the algorithms' outputs are further exploited to refine a next round of queries to the QoT estimation tool, in a closed-loop fashion, until a given optimality condition is met or no further improvements are obtained.

To the best of our knowledge, this study represents the first attempt of iterative integration of a ML-based QoT estimator in a linear program for RSA (i.e., it falls in the category ML-LP of Table I and in *Group 3*), which couples the advantages of avoiding over-conservative predictions (and, hence, under-utilization of network resources) typical of margined analytical approaches, with the optimality (or close-to-optimality with known gap) of solutions found by linear program solvers. It is worth mentioning that closed-loop approaches falling in *Group 3* have already been adopted to tackle other problems in optical networking, e.g., the optimization of launch powers [24] and the modeling of gain ripple and filter penalties in erbium-doped fiber amplifiers [25].

Among the studies that belong to category ML-HA, in [26] the output of a neural-network-based Q-factor estimator is exploited by a heuristic algorithm for dynamic routing and spectrum assignment in a multicast scenario. The algorithm first identifies a lightpath/lighttree from the source to the set of destinations and a feasible wavelength assignment, then queries the classifier to learn whether the predicted Q-factor is above a predefined threshold. If yes, the traffic demand is allocated, otherwise a new routing and spectrum allocation is identified and the classifier is queried again, until a feasible allocation is found or the request is blocked. Note that an infeasible outcome does not trigger any rearrangement of the already deployed traffic.

Authors of [27] propose a ML-based approach for inter-core crosstalk estimation in optical networks with multicore fibers. The proposed regressor is queried by a heuristic algorithm for core, route and spectrum assignment. In [29], a fuzzy C-means clustering algorithm is adopted to evaluate candidate lightpath configurations in an optical network with multicore fibers: for a newly incoming traffic request, once the set of feasible options in terms of route, spectrum and core allocation are identified, the ML algorithm chooses the most suitable option for the service's transmission needs, depending on its associated service level, which is then exploited by a heuristic algorithm for resource allocation. In both studies, the heuristic algorithms adopt working principles analogous to those of the algorithm in [26].

In [28], a ML-based traffic predictor and a QoT estimator leveraging deep neural networks are integrated in a Routing, Modulation format and Spectrum Assignment (RMSA) heuristic algorithm implemented in the broker plane for multi-

TABLE I: Taxonomy of illustrative studies integrating QoT estimation techniques in RSA routines

QoT estimation approach	Machine Learning (ML) Analytical Models (AM)	RSA approach		
		Heuristic Algorithm (HA)	Linear Program (LP)	Reinforcement Learning Agent (RL)
		[26] [27] [28] [29]	This study	[30]
		[31]	[32] [33] [34]	[35]

domain EONs. Again, the QoT estimator is leveraged in an open-loop fashion and its performance is benchmarked against a traditional transmission-reach-based RMSA scheme. Note that, from the functional point of view, all the above mentioned studies belong to *Group 1*.

Considering now category AM-RL, in [35] a QoT-aware dynamic resource allocation scheme in EONs based on deep reinforcement learning is proposed. The reinforcement learning agent is fed with information about the incoming traffic request and the overall network state. The feasibility of the attribution of a given modulation format is determined based on the GN model and used to calculate the agent’s rewards. A similar reinforcement learning approach is adopted in [30] for impairment-aware modulation format and wavelength assignment in a WDM network, but in this study the QoT prediction model leverages neural networks (i.e., it falls in category ML-RL). Note that, in both studies, the outcome of the reinforcement learning model is not a QoT prediction but an action such as deploying/dropping a channel. The QoT estimation phase is instead used to determine the agent’s rewards. From the functional point of view, the two latter studies belong to *Group 2*.

Moving to category HA-LP, in [31] analytical formulas are leveraged in a 3-stages RSA procedure for path computation and modulation format assignment (stage 1), followed by path selection (stage 2) and spectrum assignment (stage 3) performed via heuristic approaches. QoT estimations are used as inputs to stage 1 for modulation format assignment.

For what concerns category AM-LP, authors of [32] leverage precomputed crosstalk and non linear impairments values, based on the GN model, as input parameters to a mixed integer linear program for RSA in EONs. Functionally speaking, these two references belong again to *Group 1*.

In [33], [34], an iterative approach is proposed to solve a physical-impairment-aware linear program for routing and wavelength assignment in WDM optical networks. The first fundamental methodological difference with our study is that in those papers the iterative approach is aimed at fixing and rounding variables, and not at refining QoT estimations. Physical layer impairments are taken into account by dedicated variables and constraints in the mathematical formulation, where margined QoT thresholds (obtained using analytical formulas) are adopted as bounds to the maximum amount of tolerable interference/noise. So, despite adopting an iterative resolution approach, since the iterations do not involve the QoT estimation phase, from the functional point of view these studies also belong to *Group 1*. Differently, in our proposed approach we adopt a ML-based QoT estimation method, which reduces margins due to its accurate prediction capabilities. Moreover, the iterative approach adopted in the second stage of our dual-stage framework is instrumental to

adding new constraints with the aim of excluding lightpaths with unacceptable QoT, evaluated by means of a different ML model that adopts a wider set of input features with respect to that adopted in the first stage, thus taking into account also the impact of spectrally-adjacent lightpaths. Since every iteration implies a new set of queries to the QoT estimator, our approach falls into *Group 3*.

Note that a preliminary version of this study appeared in [36]. With respect to [36], we move from a classification-based to a (more practical) regression-based ML framework for QoT estimation and we evolve the iterative optimization framework therein proposed in a two-stage procedure, which drastically reduces the number of iterations required to achieve a feasible solution. Furthermore, we consider a slice-based RSA MILP formulation, which explicitly enumerates spectral slots, allowing to decrease the complexity of the MILP.

III. BACKGROUND

A. QoT ML Regressor

In this study we adopt two ML regressors for BER estimation inspired by the QoT classifier described in [37], which predicts the probability that the BER of a candidate lightpath will remain below a predetermined acceptability threshold (note that reference [37] also provides a detailed discussion on the choice of the learning algorithm to be adopted by the classifier and a thorough performance evaluation of the achievable classification accuracy). While in [36] we consider a ML classifier, in this study we prefer ML regression that possesses some advantages over classification: since a classifier outputs the probability that a lightpath configuration belongs to a given class, in order to convert this probability in a binary decision (i.e., whether the configuration is feasible), an acceptance threshold must be defined. Such threshold is provided as input during the training phase and must be properly tuned to achieve the best trade-off between false negatives and false positives. Conversely, when adopting a regressor, it is sufficient to specify a desired BER threshold “ex-post” as input to the MILP model to define the acceptability of a lightpath configuration based on the predicted BER value. Therefore, changing the BER threshold value does not impose to retrain the regression model, whereas a classification model would require a new training phase.

During the training phase, a set of lightpath instances constituted by a number of features is provided as input

to the regression algorithm². More in detail, the features characterizing a given lightpath k include:

- the lightpath total length, \mathcal{L}_k ;
- the length of the longest link included in the lightpath, \mathcal{L}_k^{max} ;
- the number of links traversed by the lightpath, l_k ;
- the amount of traffic to be served, d ;
- the modulation format to be adopted for transmission, m .

Each training instance is associated to a BER value. In the rest of the paper, we will call a ML-regressor fed only with features of the incoming lightpath (or “*ML regressor partial knowledge*”). Optionally, additional features characterizing the spectrally nearest left and right neighbor channels, co-propagating along at least one link of lightpath k may be provided as input. Those include:

- the size of the left/right inter-channel spacing separating lightpath k from its nearest neighbor channels, g_l, g_r (note that g_l, g_r are integer multiples of the slice width and must be at least as large as the guardband size, which is equal to one spectrum slice);³
- the traffic volume served by the left/right neighbors, d_l, d_r ;
- the modulation format adopted for transmission along the left/right neighbors, m_l, m_r .

In the rest of the paper, we will call a ML regressor trained also with features pertaining to interferers as “*ML regressor with complete knowledge*”).

For a newly incoming traffic request, multiple queries can be issued to the regressor to evaluate the suitability of different deployment configurations in terms of route, spectrum allocation and modulation format.

B. Etool for Framework Validation

Due to the unavailability of public datasets containing a large enough amount of BER measurements gathered by operators from lighthpaths deployed in their optical network infrastructures, in this study we leverage data synthetically-generated by means of the Etool presented in [37]. The Etool simulates a linear optical communication system affected by chromatic dispersion and additive white Gaussian noise (AWGN) and outputs the value of the uncoded BER at the input of the FEC soft decoder (pre-FEC BER). Note that, under the AWGN assumption, the pre-FEC BER depends on the pre-FEC signal-to-noise ratio (SNR) and on the modulation format.

Given a lightpath, the pre-FEC SNR is estimated by means of the link budget, which accounts for the transmitted power,

²Without loss of generality, in this study we assume that the regression algorithm is well-trained, i.e., that the size of the training dataset is adequate to achieve satisfactory prediction accuracy and that the training samples are uniformly distributed across the whole feature space. A discussion on how to deal with training datasets of reduced size is out of the scope of this paper: the interested reader is referred, e.g., to [38], [39] for a discussion on active/transfer learning techniques that can be adopted to cope with training data scarcity.

³Note that, due to spectrum fragmentation, two spectrally-adjacent channels sharing one or multiple links could be separated by more than one unoccupied spectrum slice.

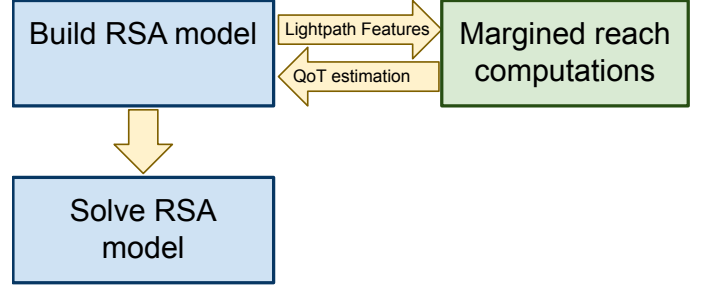


Fig. 1: Margined framework.

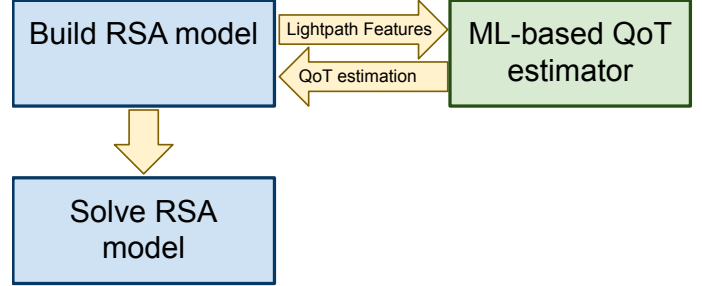


Fig. 2: ML-based framework.

gains and losses along the route. In particular, the Etool assumes transparent links constituted of dispersion uncompensated standard single-mode fibers and that the signal power is restored by identical optical amplifiers, equally spaced over the links (100 km), with gain 20 dB and noise figure 5 dB. On top of this link budget, a negative-exponential random additional penalty with 1 dB average is added, mimicking uncertainty in transmission parameters knowledge. According to [40], in the remainder of the paper we will set the pre-FEC BER threshold to $th = 4 \cdot E^{-3}$. Readers are referred to [37] for additional insights on the Etool assumptions and calculations. The calculation of the QoT of a deployed lightpath with the Etool requires about 60 to 80 seconds. Hence, the evaluation of the QoT of a single solution provided by our proposed framework would require hours (about 3 hours for the Japan network topology considered in Section VI) rendering it an impractical choice as a direct solution exploration tool.

IV. INTEGRATED ITERATIVE PLANNING FRAMEWORK

In the following subsections, we first formally state the RSA problem version solved in this study. Then we provide a high-level view of our proposed planning framework and describe benchmark frameworks that will be considered for comparison in the performance assessment reported in Section VI. Those include a baseline MILP-based RMSA solving method and two novel frameworks to include interference constraints with ML intelligence, first proposed in [36].

A. Problem Statement

The RSA problem can be defined as follows: given a directed graph representing a network topology and a number of traffic requests between source-destination nodes belonging to the graph, for every request a lightpath must be identified

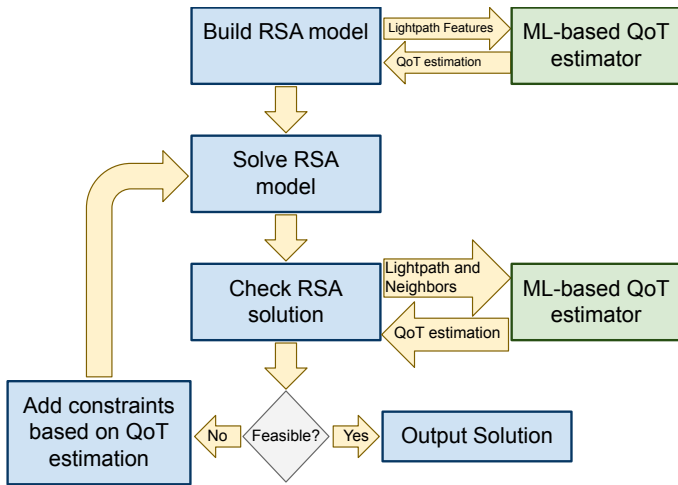


Fig. 3: ML-based iterative framework.




Fig. 4: The proposed two-stage ML-based iterative framework.

to convey the traffic from the source node to the destination node and a spectral channel must be attributed to the lightpath to allocate the optical signal in a commensurate spectral bandwidth. Spectrally neighbouring channels must be separated by guardbands to allow for proper signal filtering and switching. Note that the severity of the inter-channel interference between spectrally adjacent channels depends on the inter-channel spacing: for this reason, interference constraints need to be inserted to prevent configurations that would lead to unacceptable degradation of the lightpaths' QoT. Each channel must be equipped with a number of optical transceivers, which depends on the amount of traffic to be served and on the spectral efficiency of the modulation format adopted for transmission. If multiple modulation formats can

be used, the problem evolves in RMSA and must incorporate reach constraints that define the feasibility of the deployment of a lightpath along a given route⁴.

B. Baseline Design Framework with Margined Reach Constraints

In existing MILP models for RMSA, reach constraints are defined based on pre-computed reach values calculated by means of margined analytical formulas (see Fig. 1). Such margins are required to account for a number of unavoidable uncertainties discussed in Section I. Margined reach computations ensure an acceptable QoT, but the added margins may be too conservative, thus leading to underutilization of network resources. In this case, there is not need for interference constraints, as margined formulas ensure reachability also in case of maximum interference ("full load assumption"). In the following, this design framework will be named as *margined*.

C. ML-Based Reach Constraint Design Framework

A simple integration of a QoT estimation algorithm in a MILP model for RMSA [36] is depicted in Fig. 2. It consists in replacing the margined reach computations with a *ML-based regressor with partial knowledge* that, on input of a query for every possible combination of traffic volume, route and modulation format, estimates whether the considered combination is feasible. Based on such prediction, reach constraints can be defined and included in the MILP model. Unfortunately, also this integration does not consider interference between lightpaths, which may lead to infeasible solutions. We report this possible integration for the sake of completeness but we do not use it in our computational study. Indeed, in [36] we showed the superiority of an iterative framework, described in the next subsection, at finding feasible solutions with substantial spectrum savings with respect to the margined approach.

D. ML-Based Iterative Reach Constraint Design Framework

The *ML-based regressor with partial knowledge* discussed in the previous subsection may be subject to errors, due to lack of information about the value of some features not considered by the learning model. Such errors could be reduced, if the *ML-based regressor with complete knowledge* is used, as it is fed with additional information on the spectrally adjacent channels co-propagating along the links traversed by the considered lightpath (such as, e.g., their traffic volume, modulation format and inter-channel spacing). Unfortunately, even considering only the spectrally nearest left and right neighbor channels, an a-priori ML-based estimation for all possible configurations of a lightpath and its neighbor channels is computationally infeasible in realistic scenarios. As an example, let us assume that traffic demands may have 8 different sizes (e.g., from 50 Gbps to 400 Gbps with 50 Gbps granularity), that 6 modulation formats are available and that inter-channel spacings may have 10 different widths, for

⁴Note that higher traffic means a larger number of transceivers in the same optical channel, hence higher intra-channel interference.

each lighthpath the feasibility of $6^3 \cdot 8^3 \cdot 10^2 = 11,059,200$ configurations should be pre-computed. Therefore, integrating information on interference of neighboring channels in RMSA models introduces huge scalability issues⁵.

To overcome such limitations, in [36] we proposed an iterative procedure that, after finding a solution to the initial RMSA formulation, *for each lighthpath in the solution, issues a query to the ML-based QoT estimator containing the exact features of that lighthpath and of its neighboring channels (hence, there is no need to account for all possible cases as in the previous example)*. If the estimator returns a negative outcome (i.e., the lighthpath QoT is not acceptable), an additional constraint is added that excludes from the RMSA solution the unacceptable deployment of the lighthpath and of its neighbors. The process is repeated iteratively, until either a feasible solution for all lighthpaths is found, or a maximum number of iterations is reached. In the following sections, this framework will be referred to as *iterative*. Fig. 3 provides a graphical representation of this *iterative* procedure.

E. ML-Based Dual-stage Iterative Design Framework

In this study, we enhance the iterative procedure in [36], using a dual-stage framework, as shown in Fig. 4. The goal of the dual-stage framework is to reduce the number of iterations required to obtain a valid solution, with the scope of improving the scalability to larger network topologies. To this aim, the idea is to introduce an additional optimization stage after the initial resolution of an RMSA instance, to increase the probability that every lighthpath belonging to the selected network configuration exhibits sufficiently low BER. While the framework proposed in [36] leads to a degradation of the quality of the solution as it assigns lower modulation formats to the lighthpaths with unacceptable BER, the proposed framework tries to re-optimize the spectral assignment of those paths. In this way, since larger inter-channel spacings reduce crosstalk between spectrally-adjacent channels, the duration of the subsequent iterative phase required to find a valid solution is expected to be reduced, thus shortening the execution time and preserving the quality of the solution. In case the new spectral assignment is still inadequate to provide a valid solution, also the new framework resorts to less efficient modulation assignments as done in [36].

F. Complexity discussion

All four frameworks illustrated in Figs. 1 to 4 rely on commercial solvers to optimize the MILPs associated with the RMSA problem. RMSA is a problem belonging to the class of NP-Hard problems. Unless $P = NP$, the solution of RMSA requires, in the worst case, an exponential time in the size of the instance. Anyway, modern solvers (such as Gurobi [41], used in the following experiments) show acceptable behavior

⁵In our considered scenarios, for the case of the Japan network topology with 14 nodes and 22 links reported in fig. 5, the overall number of QoT pre-computations considering an all-to-all traffic matrix would be in the order of 10^9 . Since the evaluation of the QoT of a deployed lighthpath via the Etool requires about 60-80 seconds, that would lead to a computational time of roughly 1900 years.

for moderate sized instances. Furthermore, for the purpose of this study, reaching proven optimality is not strictly necessary. We thus accept solutions with an optimality gap proven by the solver below 2%.

Note that the RMSA problem solved by the iterative frameworks, illustrated in Figs. 3 and 4, is a *relaxation* of the RMSA solved by the margined framework illustrated in Fig. 1, because all modulation and path selection pairs that are acceptable for the margined framework are also acceptable for the iterative frameworks. As a consequence, if a feasible solution for the margined framework exists, it is guaranteed that by eliminating modulation and path selection pairs that result infeasible in terms of QoT with the addition of constraints, the iterative frameworks eventually converge to the same feasible solution of the margined framework, in the worst case. Anyway, as already shown in [36], we expect that these frameworks find feasible and more efficient solutions than the margined framework adding just a few constraints. We also expect that iterative frameworks are capable of finding feasible solutions for problems that cannot be solved by the margined framework due to the limited optical bandwidth.

The iterative frameworks eliminate a different modulation and path selection pair per iteration. Thus, the theoretical number of iterations required to convergence is bounded by a function of order $O(|M|^2 \cdot |K|^2)$. While this number is finite, the size of the constraint set can grow quite dramatically with the size of the instance. For this reason, reducing the number of iterations is of paramount importance for the practical application of the proposed frameworks. The purpose of the proposed dual stage framework of Fig. 4 is to minimize the number of additional constraints necessary to converge to a feasible solution.

Despite the above mentioned theoretical shortcomings, the numerical assessment provided in Section VI shows that, in all our practical case studies, integrating a fast but accurate QoT estimation in the RMSA results in large spectrum savings and feasible solutions in terms of QoT.

V. THE RMSA PROBLEM FORMULATION

A. Assumptions and Notations

We consider an arbitrary network topology represented by an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} is the set of nodes and \mathcal{E} is the set of bidirectional links. Each link $e \in \mathcal{E}$ is characterized by a length L_e . A static traffic matrix $D = [d_{sd}]$ of traffic requests between each node pair $(s, d) \in T = \{\mathcal{V} \times \mathcal{V} : s \neq d\}$ defines the traffic volume generated by source node s and directed to destination node d . To simplify the notation, in the following a traffic request between the (s, d) pair will be denoted as an element t of set T . For each request t , a set of candidate lighthpaths K_t is predefined. We assume that such set contains the \bar{k} shortest paths (i.e., $|K_t| = \bar{k}$). The length of a lighthpath is defined as $\mathcal{L}_k = \sum_{e \in k} L_e$. For each lighthpath, the length of its longest link is denoted as $\mathcal{L}_k^{max} = \max_{e \in k} L_e$. The number of links traversed by lighthpath k is denoted as $l_k = |\{e \in \mathcal{E} : e \in k\}|$. We assume that optical fiber spectrum is subdivided in a flexible grid with standard slice width of F GHz and elastic transceivers with

optical bandwidth of B GHz (where B is an integer multiple of F). Note that in the formulation we adopt a slice-based model to represent spectrum occupation. Superchannels with multiple adjacent transceivers are used to serve traffic demands exceeding the capacity of a single transceiver. Spectrally neighboring (super)channels are separated by an optical inter-channel spacing of G GHz (where G is an integer multiple of F). Transceivers operate at one of the modulation formats within a given set M , which results in a capacity of r_m Gbps.

B. MILP Formulation

Let us introduce the MILP formulation of the slice-based RMSA problem.

Sets:

- T , set of (s, d) pairs
- K_t , set of feasible lightpaths between (s, d) pair $t \in T$
- $K = \cup_{t \in T} K_t$, set of all lightpaths
- M , set of modulation formats
- $O \subseteq K \times K$ as the set of path pairs such that $(k1, k2) \in O$ if and only if $k1$ and $k2$ share at least one link. We assume $k1 < k2$, that is O is a set of ordered pairs of paths.

Parameters:

- Q , positive constant (greater than the maximum number of transceivers that could be needed to serve a demand)
- G , inter-channel spacing size
- B , bandwidth of a transceiver
- λ_k^m , 1 if modulation $m \in M$ is compatible with path $k \in K$, 0 otherwise
- d_t , traffic demand for (s, d) pair $t \in T$
- r_m , transceiver capacity operating at modulation $m \in M$
- S , the available bandwidth as a multiple of the slice size F .

Variables:

- b_k^m , integer, number of transceiver pairs (one installed at node s , another at node d) with modulation format $m \in M$ serving traffic $t \in T$ on path $k \in K_t$ ⁶
- β_k^m , binary, 1 if modulation $m \in M$ is used to serve traffic $t \in T$ on path $k \in K_t$
- $0 \leq S_{max} \leq S$, index of the rightmost slice occupied along any network link
- $f_k \geq 0$, starting frequency of transmission on path $k \in K_t$ as a multiple of the slice size F .
- $d_{k1, k2}$, binary, 0 if starting frequency on path $k1$ is lower than the starting frequency on path $k2$, $(k1, k2) \in O$

Objective Function:

$$z_1 = \min \alpha_1 \cdot S_{max} + \alpha_2 \cdot \frac{1}{|T|} \cdot B \cdot \sum_{m \in M} \sum_{k \in K} b_k^m \quad (1)$$

The objective function (1) minimizes a weighted sum of two contributions: the first component is the index of the rightmost slice occupied along any network link (this objective is commonly employed with the goal of leaving room for future connections), the second summation computes the average spectrum occupation per lightpath.

⁶Note that, as path k is uniquely associated to a source-destination pair $t = (s, d)$, subscript t is omitted to simplify the notation in all the variables.

Constraints:

$$\sum_{m \in M} \sum_{k \in K_t} \beta_k^m = 1 \quad \forall t \in T \quad (2)$$

$$b_k^m \leq \lambda_k^m \cdot Q \cdot \beta_k^m \quad \forall m \in M, \forall k \in K \quad (3)$$

$$\sum_{m \in M} \sum_{k \in K_t} r_m \cdot b_k^m \geq d_t \quad \forall t \in T \quad (4)$$

$$f_k + \sum_{m \in M} B \cdot b_k^m + G \leq S_{max} \quad \forall k \in K \quad (5)$$

$$\begin{aligned} f_{k1} + \sum_{m \in M} B \cdot b_{k1}^m + G - f_{k2} \\ \leq S \cdot d_{k1, k2} \quad \forall (k1, k2) \in O \end{aligned} \quad (6)$$

$$\begin{aligned} f_{k2} + \sum_{m \in M} B \cdot b_{k2}^m + G - f_{k1} \\ \leq S \cdot (1 - d_{k1, k2}) \quad \forall (k1, k2) \in O \end{aligned} \quad (7)$$

Constraints (2) state that each traffic demand uses only one modulation format; constraints (3) impose that all the transceivers installed on lightpath k use the same modulation format m and that modulation format m can be used on lightpath k only if compatible with its length; constraints (4) state that traffic demand t must be satisfied by the total capacity of the transceivers installed at the two end nodes; constraints (5) state that all lightpaths must be routed within the available bandwidth; constraints (6) and (7) state that two lightpaths sharing at least one link should not spectrally overlap.

In [36] we proposed a ML classifier to set the value of the parameter λ_k^m . As described in section III-A, in this paper we instead use a ML regressor. More in detail, for the margined framework, we compute λ_k^m with state-of-the-art margined formulas where the reach (i.e. maximum distance) coverable for a given modulation format and traffic volume is obtained as in [40], whereas in the ML-based iterative frameworks, we compute λ_k^m based on the output of the ML regressor.

In particular, given a reference value BER_{max} , λ_k^m is set as follows:

$$\lambda_k^m = \begin{cases} 1 & \text{if } \text{BER}(\mathcal{L}_k, \mathcal{L}_k^{max}, l_k, d_t, m) \leq BER_{max} \\ 0 & \text{otherwise} \end{cases}$$

In Section VI we explore different values of BER_{max} .

C. Iterative Slice-based RMSA Model

When the solution check step illustrated in Fig. 3 returns a negative outcome for a given solution (i.e., at least one lightpath exhibits a non acceptable QoT), the framework adds additional constraints. Let β^* be the vector of optimal path selection and modulation assignment variables. Let K' be the set of infeasible lightpaths. For each pair $(m', k') \in M \times K' | \beta_{k'}^{*m'} = 1$ (i.e. each pair of modulation format and path selected in the solution) let kl' and kr' be the closest spectrally adjacent deployed lightpaths thus forming an infeasible triplet (k', kl', kr') and their respective modulation formats (m', ml', mr') . The framework adds the following constraints:

$$\beta_{k'}^{m'} + \beta_{kl'}^{ml'} \leq 1 \quad (8)$$

$$\beta_{k'}^{m'} + \beta_{kr'}^{mr'} \leq 1 \quad (9)$$

These cuts prevent the selection of the same solution, fundamentally by reducing the spectral efficiency of the allotted modulation format on a given pair of network paths.

D. Two-stage Slice-based RMSA Model

Let z^* be the optimal solution of problem (1) s.t. (2) - (7) and β_k^{*m} , b_k^{*m} and S_{max}^* be the optimal values of the variables. Let $w_{k1,k2}$, for $(k1, k2) \in O$, be new variables representing the distance between assigned frequency slots for paths $k1$ and $k2$ (i.e., the inter-channel spacing size, expressed as multiple of the slice size F). The second stage model aims at designing new starting frequencies for the deployed paths with the guarantee of achieving the same performances in terms of spectrum occupation (i.e. value of the objective function (1)). We therefore want to maximize a function of the distance between pairs of paths sharing at least one link. The formulation of the second stage model is defined as follows:

$$z_2 = \max h(w_{k1,k2})$$

$$f_k + \sum_{m \in M} B \cdot b_k^{*m} + G \leq \xi \cdot S_{max}^* \quad \forall k \in K \quad (10)$$

$$f_{k1} + \sum_{m \in M} B \cdot b_{k1}^{*m} + G - f_{k2} \leq S \cdot d_{k1,k2} \quad \forall (k1, k2) \in O \quad (11)$$

$$f_{k2} + \sum_{m \in M} B \cdot b_{k2}^{*m} + G - f_{k1} \leq S \cdot (1 - d_{k1,k2}) \quad \forall (k1, k2) \in O \quad (12)$$

$$w_{k1,k2} \leq f_{k1} - f_{k2} \sum_{m \in M} B \cdot b_{k2}^{*m} - G + S \cdot (1 - d_{k1,k2}) \quad \forall (k1, k2) \in O \quad (13)$$

$$w_{k1,k2} \leq f_{k2} - f_{k1} - \sum_{m \in M} B \cdot b_{k1}^{*m} - G + S \cdot d_{k1,k2} \quad \forall (k1, k2) \in O \quad (14)$$

where $h(\cdot)$ is a function that will be defined later and $\xi \geq 1$ is a design parameter. When $\xi = 1$ the objective function z_1 cannot be worsened by the optimization of z_2 under constraints (10) - (14), indeed β_k^{*m} , b_k^{*m} and S_{max}^* are constants for this model and for the same reason constraints (2), (3) and (4) are no longer needed.

Constraints (10)-(12) guarantee that only the available spectrum is used and that lightpaths sharing at least one link do not spectrally overlap. Constraints (13) and (14) compute the spectrum distance in slices between any pair of lightpaths sharing at least one link. Note that at optimality, as variables $w_{k1,k2}$ are not constrained elsewhere in the model and the objective function is a maximization, constraints (13) and (14) will be satisfied at equality, thus $w_{k1,k2}$ models the inter-channel spacing size that separates paths $k1$ and $k2$ along their common link(s).

In this paper we explore two possible objective functions h for the second stage model. The first one is defined as:

$$z_2^1 = \max \sum_{(k1,k2) \in O} \eta_{k1,k2} \cdot w_{k1,k2}$$

which is a simple linear function of the spectrum *inter-channel spacing*, where the coefficient $\eta_{k1,k2}$ equals 0 for feasible lightpath pairs in the solution of (1) - (7) and is proportional to the estimated BER violation for infeasible ones. More in detail, let $\gamma_{k1,k2}$ be the BER violation normalized between -0.5 and 0.5 with respect to the worst BER violation found. The parameter $\eta_{k1,k2}$ is set to:

$$\eta_{k1,k2} = 1 - \gamma_{k1,k2}$$

That is, all paths violating the BER threshold get a positive weight, which is higher for those that are more problematic. We refer to this model as MIN-SUM.

The alternative objective function we consider is:

$$z_2^2 = \max \min_{(k1,k2) \in O} w_{k1,k2}$$

which aims at maximizing the smallest *inter-channel spacing* between any pair of paths. It is a non-linear function that can be trivially linearized as follows, with the help of an auxiliary variable and a set of constraints:

$$z_2^2 = \max w$$

$$w \leq w_{k1,k2} \quad \forall (k1, k2) \in O \quad (15)$$

We refer to this model as MAX-MIN.

VI. NUMERICAL RESULTS

A. Simulation Framework

We consider the Japan network topology depicted in Fig. 5, with 14 nodes and 22 links. We assume the usage of a flexible grid with standard slice width of $F = 12.5$ GHz [4] and elastic transceivers operating at 28 Gbaud with optical bandwidth $B = 37.5$ GHz (i.e., 3 slices). The minimum inter-channel spacing size, i.e., the guardband size, is set to one slice, i.e., 12.5 GHz. The available spectrum over each link is 4 THz (i.e., 320 slices). The set M of modulation formats includes dual polarization (DP)-BPSK, DP-QPSK and DP- n -QAM, with $n = 8; 16; 32; 64$, resulting in capacities of 50, 100, 150, 200, 250, 300 Gbps, respectively. We consider 8 all-to-all traffic matrices (i.e., $|T| = 14 \times 13 = 182$) with random traffic requests uniformly distributed among the sampling set $\mathcal{M} = \{50, 100, 150, 200, 250, 300, 350, 400\}$ Gbps. It follows that, in our instances, superchannels consisting of up to 8 adjacent transceivers can be deployed (e.g., in the case of a 400 Gbps request served with DP-BPSK). Margined transmission reaches for each combination of traffic volume and modulation format are reported in Table II. It should be noted that the usage of multiple spectrally-adjacent transceiver pairs in a superchannel leads to a significant increase in the intra-channel interference due to nonlinear interactions, which is numerically evaluated using the approach provided in [42]. For example, a 100 Gbps request can be served using one transceiver pair operating with QPSK modulation format (which occupy 37.5 GHz), or two transceiver pairs operating with BPSK modulation format (which occupy $37.5 \cdot 2 = 75$ GHz). However, with the former configuration, the maximum

achievable transmission reach is 3300 km, whereas with the latter only 1700 km.

Instances are listed in Table III, characterized by their maximum and average traffic per request and total traffic. The instances we consider are representative of a nationwide network with link distances in order of hundreds of kilometers (e.g., from 40 to 320) and traffic matrices ranging from low congested cases, with average traffic of 76.65 Gbps per request, to high congested cases of up to 264.0 Gbps per request. Note that, since reaches depend not only on modulation formats but also on traffic volumes, traffic requests are generated so that, for the corresponding shortest path, there exists at least one feasible modulation format able to carry the traffic considering the margined reach values. In other words, for long distance pairs, large traffic volumes are excluded from set \mathcal{M} (as an example, according to Table II, for traffic generated by node 1 and destined to node 13, the shortest path is 1120 km long, which imposes that at most 150 Gbps can be transmitted using BPSK).

As candidate paths to satisfy the end-to-end demand, we selected either the set of shortest paths (named 1PATH) or the set of shortest and second shortest paths (2PATHS) for each pair of source/destination nodes. Furthermore, for ML-based frameworks, five different thresholds of acceptance BER_{max} are defined: $\{4.00E^{-3}, 3.50E^{-3}, 3.10E^{-3}, 3.05E^{-3}, 3.01E^{-3}\}$. Therefore, an overall of 80 ($8 \cdot 2 \cdot 5$) different instances are evaluated for the ML-based frameworks while 16 ($8 \cdot 2$) instances are evaluated for the *margined* framework, as it is insensitive to BER_{max} .

In the MILP objective function we set $\alpha_1 \ll \alpha_2$, to 1 and 1000 respectively, i.e., we privilege the minimization of the bandwidth occupation of installed transceivers, but, if multiple optimal solutions exist, the one minimizing the overall spectrum usage is selected. The parameter ξ is set to 1, i.e., we do not allow any worsening of the objective function z_1 . The commercial solver Gurobi [41] has been used to solve the related MILPs. Computation of the first stage is not limited in time but an acceptance threshold for the percentage duality gap (mipgap) has been set at 2%. For the second stage both a time limit of 30 seconds and a mipgap of 2% have been set. Furthermore, the solver has been tuned to work more on improving feasibility of solutions. Iterations of both iterative frameworks have been limited to 30.

To verify the feasibility of the identified solutions, we calculate the received BER of the lightpaths included in the optimal solution (i.e., the ground truth) using the Etool described in Section III-B the validation of a solution requires about 3 hours of computation.

RMSA optimization runs are carried out on the margined framework, the ML iterative framework and the two-stage ML framework according to the variants of the objective function adopted in the second stage. Table IV clarifies the abbreviations used in the current section.

The iterative phase of the two iterative frameworks is interrupted if convergence is not reached within 30 iterations, i.e. in the case the solution still includes one or multiple lightpaths exhibiting an estimated BER above BER_{max} .

TABLE II: Margined reaches (R) expressed in km and capacities (C) expressed in Gbps for different amounts of transceiver pairs and modulation format in use (computed as in [40])

		Number of transceiver pairs							
		1	2	3	4	5	6	7	8
BPSK	R	3400	1700	1200	900	700	600	500	400
	C	50	100	150	200	250	300	350	400
QPSK	R	3300	1700	1100	900	700	600	500	400
	C	100	200	300	400	500	600	700	800
8-QAM	R	1300	700	400	300	300	200	200	100
	C	150	300	450	600	750	900	1050	1200
16-QAM	R	1000	500	300	200	200	200	100	100
	C	200	400	600	800	1000	1200	1400	1600
32-QAM	R	500	200	100	100	100	100	0	0
	C	250	500	750	1000	1250	1500	1750	2000
64-QAM	R	300	100	100	100	0	0	0	0
	C	300	600	900	1200	1500	1800	2100	2400

TABLE III: Instances of the problem

Instance	Max [Gbps]	Avg [Gbps]	Total [Tbps]
1	100	76.65	13.95
2	200	128.02	23.30
3	200	146.70	26.70
4	300	185.44	33.75
5	300	235.71	42.90
6	400	204.67	37.25
7	400	211.26	38.45
8	400	264.01	48.05

B. Feasibility results

Considering instances with 1PATH and 2PATH routing choices, we report the aggregated number of feasible and infeasible instances in Table V. *MARGINED* fails to solve 3 instances out of 16, since it uses modulation formats with low spectral efficiency, and hence available spectrum is not sufficient to accommodate all the traffic in high load instances. *MLI* still fails to find the solution for two instances, even though it significantly improves the number of feasible instances with respect to *MARGINED*. *MLI2A* and *MLI2B* always manage to find a feasible solution for each considered instance.

C. Spectrum occupation and computational time results

Let us now compare the results obtained by *MARGINED*, *MLI* and *MLI2A*, *MLI2B* in terms of spectrum occupation and computational time. We consider *MARGINED* as a baseline and report the relative percentage savings obtained by the other frameworks in Table VI and Table VII, for 1PATH and 2PATH, respectively. Note that each table row reports results averaged over five runs for every considered instance, where the BER_{max} is varied as discussed in the beginning of this section. In the rows marked with an asterisk, *MLI* could not solve at least one instance. In the rows where the savings are reported with a dash, *MARGINED* could not find any feasible solution.

Tables VI and VII show that *MLI* and both *MLI2A*, *MLI2B* greatly reduce spectrum occupation with respect to *MARGINED*, with savings up to 48.1% in the best case and above 33% on average. The best results are obtained with

TABLE IV: Abbreviations

MARGINED	Margined framework of Fig. 1
MLI	ML iterative framework of Fig. 3
MLI2A	ML it. framework of Fig. 4 with z_2^1 (MIN-SUM)
MLI2B	ML it. framework of Fig. 4 with z_2^2 (MAX-MIN)




Fig. 5: Japan network topology

MLI2B. Indeed, the second stage allows to reduce the average number of iterations necessary to converge. Remarkably, *MLI2B* always converges in just one iteration.

Tables VIII and IX report the computational times required by the considered frameworks. *MARGINED* is faster, as it solves the MILP problem only once, but it fails to find a feasible solution for several instances and the solved ones exhibit higher spectrum occupation. Under the 1PATH assumption, *MLI*, *MLI2A* and *MLI2B* require longer time to converge (up to 37.5 s on average), but this is the price to pay for achieving lower spectrum occupation. Here, the second stage of *MLI2A* and *MLI2B* is the one requiring the highest computational time. Under the 2PATH assumption, the computational times of *MLI* and *MLI2A* further increases by about one order or magnitude. This is due to the higher complexity of the model and to the increased number of iterations to reach convergence. Instead, *MLI2B* always requires just one iteration to converge. Sometime *MLI2B* is even faster than *MARGINED* because the first stage model converges very quickly and this compensates the extra time necessary to solve the second stage model.

D. Comparison between MIN-SUM and MIN-MAX objective functions

In this section we compare *MLI2A* and *MLI2B* aggregating results per instance and varying the acceptance threshold BER_{max} , under the 2PATH assumption, in terms of overall spectrum occupation and computational time. Fig. 6 shows that, for instances 1-7, both objective functions achieve remarkable savings with respect to *MARGINED*, with *MLI2B* performing slightly better than *MLI2A* (note that instance 8 is excluded, since *MARGINED* could not find a feasible solution). As previously remarked, *MLI2B* is around one order of magnitude faster w.r.t. *MLI2A*. Fig. 7 reports the impact of the value of the BER acceptance threshold on the overall spectrum occupation and computational times. We observe that lowering the BER acceptance threshold, i.e. accepting more conservative lightpath configurations in the first MILP stage, does not affect the quality of the solutions. Indeed, savings with respect to *MARGINED* are always above 34%, on average. As for the computational time, the acceptance

TABLE V: Instance feasibility achieved with different frameworks.

Framework	Feasible Inst.	Infeasible Inst.	Infeasible Inst. [%]
MARGINED	13	3	18.8
MLI	78	2	2.6
MLI2A	80	0	0.0
MLI2B	80	0	0.0


TABLE VI: Average iterations and spectrum occupation savings achieved by *MLI*, *MLI2A* and *MLI2B* on instances with 1PATH, using the *MARGINED* framework as baseline.

Inst.	MLI Iter.	MLI2A Iter.	MLI2B Iter.	MLI Sav. [%]	MLI2A Sav. [%]	MLI2B Sav. [%]
1	3.0	1.8	1.0	25.0	25.4	25.4
2	9.0	4.0	1.0	33.9	34.2	35.0
3	5.0	6.4	1.0	39.0	38.7	40.5
4	6.6	5.6	1.0	40.8	41.0	42.4
5	5.0	2.6	1.0	22.3	22.5	22.7
6	5.0	3.6	1.0	42.8	42.7	43.5
7	4.6	5.2	1.0	-	-	-
8	5.6	2.8	1.0	-	-	-
AVG	5.5	4.0	1.0	33.9	34.1	34.9


threshold reduction favors the *MLI2A* framework only. The average computation time is reduced as a direct consequence of the reduction of the average number of iterations required to converge. Still, the improvement is not sufficient to reach the performances obtained with the MAX-MIN objective function implemented in *MLI2B*.

E. Solution validation

We finally verify the validity of the solutions obtained by *MARGINED*, *MLI*, *MLI2B*. *MLI2A* is not considered here because clearly dominated by *MLI2B*. To this aim, for each lightpath configuration included in the solution, we compute the BER using the Etool presented in [37] and compare it to the BER acceptance threshold value, which is set to $4E^{-3}$. If the computed BER is above threshold, the lightpath is considered invalid. This procedure constitutes an ‘‘a posteriori’’ validation that, in real scenarios, would be replaced by a post-deployment BER measurement along every installed lightpath. In order to provide a fair comparison, the number of solved instances should be considered. Note that it is sufficient that one lightpath shows a BER above the threshold to invalidate the solution. Table X reports the number of a-posteriori infeasible instances and the fraction of infeasible instances and infeasible lightpaths, for the instances that are considered feasible by the different frameworks, respectively. We observe that *MARGINED* never produces an infeasible solution among the solved instances but we recall that *MARGINED* is not able to solve 18.8% of the available instances. Instead, ML-based frameworks are less conservative. However, when adopting




(a) Spectrum Savings




(b) Computational Time

Fig. 6: Comparison of MIN-SUM and MIN-MAX objective functions per instance



(a) Spectrum Savings



(b) Computational Time

Fig. 7: Comparison of MIN-SUM and MIN-MAX objective functions per acceptance threshold

MLI, despite a few threshold violations (30 over more than 14000 evaluated paths), the large majority of deployed paths (above 99.7%) are within the prescribed threshold. As for the *MLI2B* framework, we remark that over more than 14000 evaluated lightpaths, only *one* lightpath ($< 0.01\%$) exhibits above threshold BER, thus making one instance infeasible.

The obtained distributions of BER values is plotted in Fig. 8, where we can observe that the *MARGINED* framework, being more conservative, exhibits lower BER values than the ML-based ones.

Fig. 9 reports the number of unacceptable lightpath configurations over all the solved instances for *MLI*, *MLI2A* and *MLI2B*, depending on the value of the acceptance threshold BER_{max} provided as input to the MILP model (we remark that the a posteriori validation via the Etool always assumes a BER threshold $th = 4E^{-3}$). Results show that,

with acceptance thresholds lower than $3.5E^{-3}$, *MLI2A* and *MLI2B* always produces acceptable lightpath deployments for all instances, while *MLI* still yields a few unacceptable ones.

F. Scalability assessment on larger networks

In order to evaluate the scalability of the considered frameworks we run some tests on larger networks obtained from the Survivable fixed telecommunication Network Design library SNDLib [43], [44]. We perform experiments on the optical networks of France and Germany, basic information on such networks is reported in table XI. As for the Japan network, we selected either the set of shortest paths (1PATH) or the set of shortest and second shortest paths (2PATHS) for each pair of source/destination nodes as candidate paths to satisfy each end-to-end demand. We used demands reported in the SNDLib

TABLE VII: Average iterations and spectrum occupation savings achieved by MLI, MLI2A and MLI2B on instances with 2PATHS, using the MARGINED framework as baseline.

Inst.	MLI Iter.	MLI2A Iter.	MLI2B Iter.	MLI Sav. [%]	MLI2A Sav. [%]	MLI2B Sav. [%]
1	7.8	7.0	1.0	25.8	25.8	26.2
2*	20.7	10.0	1.0	34.1	34.2	35.4
3	18.8	18.0	1.0	39.4	39.0	42.1
4	19.8	18.4	1.0	41.4	40.9	43.0
5	3.6	5.0	1.0	23.0	22.8	23.1
6*	23.7	8.4	1.0	43.2	43.0	43.9
7	19.4	15.8	1.0	47.0	46.7	48.1
8	4.2	2.8	1.0	-	-	-
AVG	14.8	10.7	1.0	36.3	36.1	37.4

TABLE VIII: Average execution time on instances with 1PATH for MARGINED, MLI, MLI2A and MLI2B.

Inst.	MARGINED Time [s]	MLI Time [s]	MLI2A Time [s]	MLI2B Time [s]
1	4.2	7.4	9.0	35.5
2	4.0	19.8	33.1	35.9
3	4.5	12.0	56.4	35.6
4	4.7	12.9	44.1	35.2
5	10.2	26.6	17.1	46.8
6	4.4	12.1	20.4	35.7
7	-	11.0	32.8	35.3
8	-	24.2	69.7	39.5
AVG	5.3	15.8	35.3	37.4

as a reference but scaled them so that transmission feasibility is guaranteed for the *MARGINED* method on the selected set of paths according to the reach constraints reported in table II. Summary statistics of the demand matrices are reported in Table XII.

In the experiments, we compared the *MARGINED*, *MLI*, and *MLI2B* frameworks. After some preliminary experiments, we tuned the algorithm to be slightly more conservative for both




Fig. 8: Comparison of BER distributions obtained with the three considered frameworks

TABLE IX: Average execution time on instances with 2PATHS for MARGINED, MLI, MLI2A and MLI2B.

Inst.	MARGINED Time [s]	MLI Time [s]	MLI2A Time [s]	MLI2B Time [s]
1	121.9	610.6	578.0	107.0
2*	96.1	2141.0	1527.3	163.7
3	120.7	1734.7	2077.0	126.2
4	214.2	1985.0	2390.3	110.2
5	135.4	320.0	407.3	129.9
6*	97.7	2239.7	970.5	114.2
7	354.0	2134.7	2193.7	120.4
8	-	551.2	320.5	148.2
AVG	162.8	1464.6	1308.1	127.5

TABLE X: Infeasible instances and BER per framework.

Framework	Total Inst.	Inf. Inst.	Inf. Inst. [%]	Inf. Lightp. [%]	Average MAX BER	Average BER
MARGINED	13	0	0.00	0.00	1.91E-03	3.14E-05
MLI	78	19	21.59	0.21	5.59E-03	2.28E-04
MLI2B	80	1	1.11	0.01	3.32E-03	1.63E-04

MLI and *MLI2B* setting the acceptance threshold of unestablished lightpaths to $1E^{-3}$. Consequently, both frameworks showed very good convergence properties while still ensuring major spectrum savings. Furthermore, for all frameworks, we tuned the Gurobi solver to stop computation when the solution was proven to be less than 2% from the optimum.

As reported in Tables XIII and XIV, all instances were solved within the time limit set to 1800 seconds. As expected, the *MARGINED* framework is faster but too conservative. Savings between 47.8% to 52.4% are observed for the ML based frameworks with respect to the *MARGINED* approach for the




Fig. 9: Number of infeasible solutions depending on the BER acceptance threshold

France network. Slightly better savings are obtained by the *MLI2B* framework but more computational time is required. For the German network, spectrum savings are lower, between 15.5% and 22.3%. While computational times are larger than those of the French network, both ML frameworks show good convergence properties. The instance *Germany 2 Paths* proved to be very challenging to solve by all frameworks.

For both the French and German networks, the ML frameworks converged in just 1 iteration when routing the demand among the shortest and second shortest paths (2PATH), thus favoring the “basic” *MLI* framework in terms of computational time. Instead for the 1PATH case the *MLI2B* framework showed better convergence properties but, having to solve two optimization problems instead of one, the computational times are comparable to those of *MLI*.

As a general conclusion, we observe that the ML frameworks behave acceptably compared to the *MARGINED* one on large networks. Scaling to even larger networks or considering a larger routing choice, would imply the adoption of a faster, probably heuristic, optimization algorithm. We remark that this would be required for the *MARGINED* approach as well. The *MLI* frameworks, with little effort, can be adapted to work with heuristic solvers.

TABLE XI: Details of France and Germany networks from SNDLib.

Name	Nodes	Links	OD Pairs
France	25	45	300
Germany	50	88	662

TABLE XII: Details of France and Germany demand statistics.

Instance	Max [Gbps]	Avg [Gbps]	Total [Tbps]
France 1 Path	500	237.9	71.4
France 2 Paths	500	222.9	66.9
Germany 1 Path	500	137.4	91.0
Germany 2 Paths	500	133.9	88.6

VII. CONCLUSION

We propose a MILP-based planning framework for routing, spectrum and modulation format assignment in EONs that integrates ML-based QoT estimation to estimate lightpath feasibility while avoiding conservative QoT overestimations. A dual-stage iterative procedure is adopted, in which supplementary constraints are added to the MILP based on the outputs of iterated additional queries issued to the estimator, which include as features some characteristics of the neighbor channels. Results show spectrum savings up to 52% with respect to traditional frameworks based on margined reach computations. By properly tuning the value of the bit error rate acceptability threshold provided as input to the model, network configurations characterized by different risk levels - which are quantified by the number of lightpaths that exceed the threshold - can be obtained.

TABLE XIII: Iterations and spectrum occupation savings achieved by *MLI* and *MLI2B* on SNDLib instances, using the *MARGINED* framework as baseline.

Inst.	<i>MLI</i> Iter.	<i>MLI2B</i> Iter.	<i>MLI</i> Sav. [%]	<i>MLI2B</i> Sav. [%]
France 1 path	3.0	1.0	51.3	52.4
France 2 paths	1.0	1.0	47.8	47.8
Germany 1 path	3.0	1.0	21.7	22.3
Germany 2 paths	1.0	1.0	15.5	15.5

TABLE XIV: Execution time on SNDLib instances for *MARGINED*, *MLI* and *MLI2B*.

Inst.	<i>MARGINED</i> Time [s]	<i>MLI</i> Time [s]	<i>MLI2B</i> Time [s]
France 1 Path	12.5	58.4	99.6
France 2 Paths	66.7	99.9	166.8
Germany 1 Path	40.4	241.4	155.3
Germany 2 Paths	472.2	617.3	685.9

REFERENCES

- [1] Cisco, “Cisco Annual Internet Report (2018–2023) White Paper,” Accessed in Dec. 2020. [Online]. Available: <https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html>
- [2] O. Gerstel, M. Jinno, A. Lord, and S. B. Yoo, “Elastic optical networking: A new dawn for the optical layer?” *IEEE Communications Magazine*, vol. 50, no. 2, 2012.
- [3] S. Talebi, F. Alam, I. Katib, M. Khamis, R. Salama, and G. N. Rouskas, “Spectrum management techniques for elastic optical networks: A survey,” *Optical Switching and Networking*, vol. 13, pp. 34–48, 2014.
- [4] International Telecommunication Union, Telecommunication Standardization Sector, “Spectral grids for WDM applications: DWDM frequency grid,” *ITU-T Rec. G.694.1*, Feb. 2012. [Online]. Available: <http://www.itu.int>
- [5] P. Poggiolini, G. Bosco, A. Carena, V. Curri, Y. Jiang, and F. Forghieri, “The GN-model of fiber non-linear propagation and its applications,” *Journal of lightwave technology*, vol. 32, no. 4, pp. 694–721, 2013.
- [6] R. Proietti, X. Chen, A. Castro, G. Liu, H. Lu, K. Zhang, J. Guo, Z. Zhu, L. Velasco, and S. J. B. Yoo, “Experimental demonstration of cognitive provisioning and alien wavelength monitoring in multi-domain eon,” in *Optical Fiber Communication Conference. Optical Society of America*, 2018, p. W4F.7. [Online]. Available: <http://www.osapublishing.org/abstract.cfm?URI=OFC-2018-W4F.7>
- [7] Y. Pointurier, “Design of low-margin optical networks,” *IEEE/OSA Journal of Optical Communications and Networking*, vol. 9, no. 1, pp. A9–A17, 2017.
- [8] E. Ip and J. M. Kahn, “Compensation of dispersion and nonlinear impairments using digital backpropagation,” *Journal of Lightwave Technology*, vol. 26, no. 20, pp. 3416–3425, Oct 2008.
- [9] K. Christodoulopoulos, P. Kokkinos, A. Di Giglio, A. Pagano, N. Argyris, C. Spatharakis, S. Dris, H. Avramopoulos, J. Antona, C. Delezoide *et al.*, “Orchestra-optical performance monitoring enabling flexible networking,” in *Transparent Optical Networks (ICTON), 2015 17th International Conference on*. IEEE, 2015, pp. 1–4.
- [10] Y. Bengio, A. Lodi, and A. Prouvost, “Machine learning for combinatorial optimization: a methodological tour d’horizon,” *European Journal of Operational Research*, vol. 290, no. 2, pp. 405–421, 2021.
- [11] F. Musumeci, C. Rottondi, A. Nag, I. Macaluso, D. Zibar, M. Ruffini, and M. Tornatore, “An overview on application of machine learning techniques in optical networks,” *IEEE Communications Surveys & Tutorials*, vol. 21, no. 2, pp. 1383–1408, 2018.
- [12] J. Mata, I. de Miguel, R. J. Duran, N. Merayo, S. K. Singh, A. Jukan, and M. Chamania, “Artificial intelligence (ai) methods in optical networks:

- A comprehensive survey," *Optical switching and networking*, vol. 28, pp. 43–57, 2018.
- [13] Y. Zhang, J. Xin, X. Li, and S. Huang, "Overview on routing and resource allocation based machine learning in optical networks," *Optical Fiber Technology*, vol. 60, p. 102355, 2020.
- [14] R. Boutaba, M. A. Salahuddin, N. Limam, S. Ayoubi, N. Shahriar, F. Estrada-Solano, and O. M. Caicedo, "A comprehensive survey on machine learning for networking: evolution, applications and research opportunities," *Journal of Internet Services and Applications*, vol. 9, no. 1, p. 16, 2018.
- [15] C. L. Gutterman, W. Mo, S. Zhu, Y. Li, D. C. Kilper, and G. Zussman, "Neural network based wavelength assignment in optical switching," in *Proceedings of the Workshop on Big Data Analytics and Machine Learning for Data Communication Networks*, 2017, pp. 37–42.
- [16] I. Martín, S. Troia, J. A. Hernández, A. Rodríguez, F. Musumeci, G. Maier, R. Alvizu, and Ó. G. de Dios, "Machine learning-based routing and wavelength assignment in software-defined optical networks," *IEEE Transactions on Network and Service Management*, vol. 16, no. 3, pp. 871–883, 2019.
- [17] J. Yu, B. Cheng, C. Hang, Y. Hu, S. Liu, Y. Wang, and J. Shen, "A deep learning based rsa strategy for elastic optical networks," in *2019 18th International Conference on Optical Communications and Networks (ICOON)*. IEEE, 2019, pp. 1–3.
- [18] X. Chen, B. Li, R. Proietti, H. Lu, Z. Zhu, and S. B. Yoo, "Deeprrmsa: A deep reinforcement learning framework for routing, modulation and spectrum assignment in elastic optical networks," *Journal of Lightwave Technology*, vol. 37, no. 16, pp. 4155–4163, 2019.
- [19] X. Luo, C. Shi, L. Wang, X. Chen, Y. Li, and T. Yang, "Leveraging double-agent-based deep reinforcement learning to global optimization of elastic optical networks with enhanced survivability," *Optics express*, vol. 27, no. 6, pp. 7896–7911, 2019.
- [20] A. Mohammed and B. Jaumard, "Nested column generation algorithm for the routing and spectrum assignment problem in flexgrid optical networks," in *2021 IEEE Canadian Conference on Electrical and Computer Engineering (CCECE)*, 2021, pp. 1–5.
- [21] B. Jaumard, Y. Wang, and D. Coudert, "Dantzig-wolfe decomposition for the design of filterless optical networks," *J. Opt. Commun. Netw.*, vol. 13, no. 12, pp. 312–321, Dec 2021.
- [22] B. C. Chatterjee, N. Sarma, and E. Oki, "Routing and spectrum allocation in elastic optical networks: A tutorial," *IEEE Communications Surveys & Tutorials*, vol. 17, no. 3, pp. 1776–1800, 2015.
- [23] M. Klinkowski, P. Lechowicz, and K. Walkowiak, "Survey of resource allocation schemes and algorithms in spectrally-spatially flexible optical networking," *Optical Switching and Networking*, vol. 27, pp. 58–78, 2018.
- [24] A. Mahajan, K. K. Christodoulopoulos, R. Martínez, R. M. noz, and S. Spadaro, "Quality of transmission estimator retraining for dynamic optimization in optical networks," *J. Opt. Commun. Netw.*, vol. 13, no. 4, pp. B45–B59, Apr 2021. [Online]. Available: <http://opg.optica.org/jocn/abstract.cfm?URI=jocn-13-4-B45>
- [25] A. Mahajan, K. Christodoulopoulos, R. Martínez, S. Spadaro, and R. M. noz, "Modeling edfa gain ripple and filter penalties with machine learning for accurate qot estimation," *J. Lightwave Technol.*, vol. 38, no. 9, pp. 2616–2629, May 2020. [Online]. Available: <http://opg.optica.org/jlt/abstract.cfm?URI=jlt-38-9-2616>
- [26] T. Panayiotou, S. Chatzis, and G. Ellinas, "Performance analysis of a data-driven quality-of-transmission decision approach on a dynamic multicast-capable metro optical network," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 9, no. 1, pp. 98–108, Jan. 2017.
- [27] Q. Yao, H. Yang, R. Zhu, A. Yu, W. Bai, Y. Tan, J. Zhang, and H. Xiao, "Core, mode, and spectrum assignment based on machine learning in space division multiplexing elastic optical networks," *IEEE ACCESS*, vol. 6, pp. 15 898–15 907, 2018.
- [28] X. Chen, R. Proietti, H. Lu, A. Castro, and S. B. Yoo, "Knowledge-based autonomous service provisioning in multi-domain elastic optical networks," *IEEE Communications Magazine*, vol. 56, no. 8, pp. 152–158, 2018.
- [29] H. Yang, Q. Yao, A. Yu, Y. Lee, and J. Zhang, "Resource assignment based on dynamic fuzzy clustering in elastic optical networks with multi-core fibers," *IEEE Transactions on Communications*, vol. 67, no. 5, pp. 3457–3469, 2019.
- [30] M. Freire-Hermelo, A. Lavignotte, and C. Lepers, "Dynamic modulation format and wavelength assignment in optical networks using reinforcement learning," in *Photonic Networks and Devices*. Optica Publishing Group, 2021, pp. NeF2B–4.
- [31] H. Beyranvand and J. A. Salehi, "A quality-of-transmission aware dynamic routing and spectrum assignment scheme for future elastic optical networks," *Journal of Lightwave Technology*, vol. 31, no. 18, pp. 3043–3054, 2013.
- [32] S. Behera, A. Deb, G. Das, and B. Mukherjee, "Impairment aware routing, bit loading, and spectrum allocation in elastic optical networks," *Journal of Lightwave Technology*, vol. 37, no. 13, pp. 3009–3020, 2019.
- [33] K. Christodoulopoulos, K. Manousakis, and E. Varvarigos, "Offline routing and wavelength assignment in transparent wdm networks," *IEEE/ACM Transactions on Networking*, vol. 18, no. 5, pp. 1557–1570, 2010.
- [34] K. Manousakis, K. Christodoulopoulos, E. Kamitsas, I. Tomkos, and E. A. Varvarigos, "Offline impairment-aware routing and wavelength assignment algorithms in translucent wdm optical networks," *Journal of Lightwave Technology*, vol. 27, no. 12, pp. 1866–1877, 2009.
- [35] C. Shi, M. Zhu, J. Gu, T. Shen, and X. Ren, "Deep-reinforced impairment-aware dynamic resource allocation in nonlinear elastic optical networks," in *26th Optoelectronics and Communications Conference*. Optica Publishing Group, 2021, p. M4A.8. [Online]. Available: <http://opg.optica.org/abstract.cfm?URI=OECC-2021-M4A.8>
- [36] M. Salani, C. Rottondi, and M. Tornatore, "Routing and spectrum assignment integrating machine-learning-based qot estimation in elastic optical networks," in *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*. IEEE, 2019, pp. 1738–1746.
- [37] C. Rottondi, L. Barletta, A. Giusti, and M. Tornatore, "Machine-learning method for quality of transmission prediction of unestablished light-paths," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 10, no. 2, pp. A286–A297, Feb 2018.
- [38] D. Azzimonti, C. Rottondi, and M. Tornatore, "Reducing probes for quality of transmission estimation in optical networks with active learning," *Journal of Optical Communications and Networking*, vol. 12, no. 1, pp. A38–A48, 2020.
- [39] D. Azzimonti, C. Rottondi, A. Giusti, M. Tornatore, and A. Bianco, "Comparison of domain adaptation and active learning techniques for quality of transmission estimation with small-sized training datasets," *Journal of Optical Communications and Networking*, vol. 13, no. 1, pp. A56–A66, 2021.
- [40] G. Bosco, V. Curri, A. Carena, P. Poggiolini, and F. Forghieri, "On the performance of nyquist-WDM terabit superchannels based on PM-BPSK, PM-QPSK, PM-8QAM or PM-16QAM subcarriers," *IEEE/OSA Journal of Lightwave Technology*, vol. 29, no. 1, pp. 53–61, 2011.
- [41] L. Gurobi Optimization, "Gurobi optimizer reference manual," 2018. [Online]. Available: <http://www.gurobi.com>
- [42] R. Dar, M. Feder, A. Mecozzi, and M. Shtaif, "Accumulation of nonlinear interference noise in fiber-optic systems," *Optics express*, vol. 22, no. 12, pp. 14 199–14 211, 2014.
- [43] S. Orłowski, M. Pióro, A. Tomaszewski, and R. Wessäly, "SNDlib 1.0–Survivable Network Design Library," in *Proceedings of the 3rd International Network Optimization Conference (INOC 2007)*, Spa, Belgium, April 2007, <http://sndlib.zib.de>, extended version accepted in *Networks*, 2009. [Online]. Available: <http://www.zib.de/orlowski/Paper/OrlowskiPioroTomaszewskiWessaely2007-SNDlib-INOC.pdf.gz>
- [44] —, "SNDlib 1.0–Survivable Network Design Library," *Networks*, vol. 55, no. 3, pp. 276–286, 2010. [Online]. Available: <http://www3.interscience.wiley.com/journal/122653325/abstract>



Matteo Salani received the M.S. degree in computer science and the Ph.D. degree in computer science from Università degli Studi di Milano, Italy, in 2001 and 2006. Currently, he is a Senior Researcher at the Dalle Molle Institute for Artificial Intelligence ID-SIA USI-SUPSI, Lugano, Switzerland. His research interests deal with mathematical programming and optimization. In particular, he is interested in integrating Machine Learning in heuristic and exact algorithms for combinatorial problems arising in transportation, telecommunication, energy manage-

ment and logistics.



Cristina Rottondi is Associate Professor with the Department of Electronics and Telecommunications of Politecnico di Torino (Italy). Her research interests include optical networks planning and networked music performance. She received her Ph.D. in Information Engineering from Politecnico di Milano (Italy) in 2014. From 2015 to 2018 she had a research appointment at the Dalle Molle Institute for Artificial Intelligence in Lugano, Switzerland. She is co-author of more than 90 scientific publications in international journals and conferences. She served

as Associate Editor for IEEE Access from 2016–2020 and is currently Associate Editor of the IEEE/OSA Journal of Optical Communications and Networking. She is co-recipient of the 2020 Charles Kao Award, three best paper awards (Conference of Open Innovations Association (FRUCT)– International Workshop on the Internet of Sounds 2020, International Conference on Design of Reliable Communication Networks 2017, and IEEE Green Computing and Communications Conference 2014), and one excellent paper award (International Conference on Ubiquitous and Future Networks 2017).



Leopoldo Ceré holds an MS degree in Intelligent Systems and a BSc in Computer Science from University of Applied Sciences of Southern Switzerland, Lugano, Switzerland. He worked as Research Assistant at the Dalle Molle Institute for Artificial Intelligence and currently holds the position of Data Scientist at Artificialy SA, Lugano, Switzerland.



Massimo Tornatore is currently an Associate Professor with the Department of Electronics, Information, and Bioengineering, Politecnico di Milano. He also holds an appointment as Adjunct Professor at University of California, Davis, USA and as visiting professor at University of Waterloo, Canada. His research interests include performance evaluation, optimization and design of communication networks (with an emphasis on the application of optical networking technologies), network virtualization, network reliability and machine learning application for

network management. In these areas, he co-authored more than 400 peer-reviewed conference and journal papers (with 19 best paper awards), 2 books and 1 patent. He is a member of the Editorial Board of IEEE Communication Surveys and Tutorials, IEEE Communication Letters, IEEE Transactions on Network and Service Management, IEEE Transactions on Networking, and Elsevier Optical Switching and Networking.