Cross-Layer Adaptive Elastic Optical Networks

Ippokratis Sartzetakis, Konstantinos Christodoulopoulos, and Emmanuel Varvarigos

***Abstract*—Optical networks are designed to be operated statically: lightpaths are provisioned for uninterrupted operation for several years using high margins to anticipate the deterioration of their quality of transmission (QoT) from various factors (equipment aging, malfunctioning, and maintenance operations). Operating the network dynamically and closer to its actual capabilities increases efficiency and reduces capital expenditure. We develop a cross-layer QoT-aware toolkit that leverages monitoring information and the flexibility dimensions of elastic optical networks. It adapts the network's parameters and regulates the QoT to achieve high efficiency. The toolkit can be used in a plethora of use cases in the deployment or during the operation phase of the network, e.g., to harvest the excessive margins when lightpaths are initially deployed or failures are repaired, to adapt the network to changing traffic demands, and to restore margins when soft failures such as equipment malfunction or aging render the QoT of certain lightpaths unacceptable. In the last case, the toolkit can be used to appropriately reconfigure the lightpaths to restore their QoT and postpone the deployment of regenerators, as indicated by our simulations.**

***Index Terms*—Cross layer optimization; Failures; Margins; Network optimization; Network survivability; Reconfiguration.**

I. INTRODUCTION

In an optical transport network, lightpaths must have an acceptable quality of transmission (QoT) at all times. The current practice is to operate optical networks statically, i.e., lightpaths are provisioned to have acceptable QoT until their end-of-life (EoL) without having to be reconfigured. Certain inevitable events such as equipment aging or malfunction (e.g., of an amplifier), maintenance operations (e.g., using splices to fix fiber cuts), and interference from new lightpaths degrade the QoT during the lifetime of the network. To anticipate such possible future degradations, the current practice is to provision lightpaths with high margins (also referred to as power budgets) [1]. Such margins result in the deployment of regenerators or more robust transponders that are not strictly necessary at the initial set-up time.

Lower margins result in higher efficiency because resources are utilized closer to their capabilities, thereby reducing the overall required equipment. This in turn yields significant cost savings [2–4]. To enable the reduction of margins, new mechanisms are required to i) understand the network state so as to appropriately reduce the margins, and ii) anticipate, identify, and take appropriate actions to resolve any QoT issue that may arise [5,6]. This requires new feedback-based control mechanisms that can rely on the use of optical performance monitors (OPMs), to observe the state of the network. During the past few years, optical coherent transceivers are being installed in core optical networks, and they are finding their way into metro/regional networks. The ORCHESTRA project [7] proposes to exploit the digital signal processing (DSP) capabilities of coherent receivers and make them function as OPMs. Information from OPMs combined with a scalable monitoring and control plane can enable dynamic network operation and provide the mechanisms needed for lowering the margins through real cross-layer optimization.

In this paper, which is an extension of [8], we present a novel algorithmic toolkit designed to enable the efficient operation of elastic optical networks (EONs) [9]. The toolkit exploits the configuration capabilities of elastic transponders (also referred to as bandwidth variable transponders [BVTs]) and feedback information from the physical layer. It can be used to decide on appropriate reconfiguration actions aiming to regulate the QoT margins, keep them low but always acceptable, and achieve overall high efficiency. In particular, the toolkit considers the combination of three reconfiguration techniques to regulate the QoT of a lightpath: (i) modifying its forward error correction (FEC) overhead, (ii) adjusting the spectrum guard band to control the interference from neighboring lightpaths, and (iii) adapting its modulation format. Depending on the capabilities of the transponder and the related class of service and service layer agreements (SLAs), these adaptations can be combined with the modification of the baud rate of the respective lightpath. Thus, these techniques result in a set of network adaptation actions that can involve spectrum reconfiguration of one or more lightpaths [8]. Because the available optimization combinations in an EON are vast, the toolkit searches possible solutions using various criteria such as the control plane overhead (i.e., the actions that have to be performed and the number of lightpaths that are adapted) or the required additional spectrum. It can also leverage a QoT estimation model such as [6] in order to estimate whether an adaptation action will result in an

Manuscript received July 6, 2017; revised November 7, 2017; accepted November 8, 2017; published January 17, 2018 (Doc. ID 301291).

I. Sartzetakis (e-mail: isartz@mail.ntua.gr), K. Christodoulopoulos, and E. Varvarigos are with the Computer Technology Institute and press (CTI), Rion 26504, Greece. They are also with the School of Electrical and Computer Engineering, National Technical University of Athens, Zografou, 15773 Athens, Greece.

https://doi.org/10.1364/JOCN.10.00A154

acceptable QoT. Alternatively, we can gradually apply the changes, monitor their effect, and adjust the configurations accordingly until we reach the targeted margins levels.

We examined the application of our proposal in the recovery from *soft failures*. We define a soft-failure as a QoT degradation, as opposed to a hard failure, which causes total loss of the signal, e.g., due to a fiber cut or a complete breakdown of some equipment. The current practice is to use high margins that can absorb soft failures up to a certain amount. When these margins are exceeded, operators rely on protection/restoration mechanisms [10]. A network operated with low margins, as the one proposed in this paper, has lightpaths whose QoT is close to the acceptable threshold; it achieves high efficiency and requires less equipment but is susceptible to soft failures. Our toolkit can be used to adapt the transmission parameters of the problematic lightpaths in order to restore their QoT to acceptable levels. If the failure is repaired and the QoT is improved, then we can again use the toolkit to harvest the excessive margins and restore the previous or another appropriate configuration. In this way, we always keep the utilization of the resources and the efficiency of the network high.

We evaluate the benefits of our proposal by simulating various soft-failure scenarios and comparing them with the traditional practice of planning with high margins. Our results indicate that significant savings can be achieved while always guaranteeing acceptable QoT.

The paper is organized as follows: In Section II, we provide an overview of previous related work. In Section III, we define the network scenario that we assume; in Section IV, we present the proposed toolkit. In Section V, we present the simulation results obtained, and Section VI describes our conclusions.

II. RELATED WORK

Dynamic network optimization has received significant research attention. In [11,12] certain adaptive routing and wavelength assignment (RWA) algorithms are presented. They perform cross-layer optimization in that they take into account the physical layer impairments in order to increase the QoT and reduce the blocking rate of new lightpaths. As EONs appeared, the optimization dimensions increased and new problems emerged. In [13,14] the authors consider dynamic spectrum defragmentation for EONs according to the current routing network state, and [14] further investigates the reconfiguration of the network after a failure. The authors in [15] propose dynamic guard-band creation in EONs according to physical layer impairment/QoT estimation. The results indicate significant reduction in the number of dropped connections and subsequent increase in throughput. Finally, [16,17] consider dynamic spectrum allocation in order to adapt to fluctuations in traffic demands. Network efficiency is increased and blocking probability is subsequently decreased.

Another significant research subject is the reduction of margins. It further complicates the cross-layer optimization problem because it adds another dimension: the evolution of the network's parameters in time. Optical networks are traditionally planned using QoT estimation models and high margins. *System margins* are used to account for time-varying physical conditions [1]: future degradations due to equipment aging, interference from increases in load, and failures until the EoL. Design margin is used to account for inaccuracies of the QoT estimation model. Overall, high margins ensure future degradations (up to a certain amount, corresponding to the expected QoT at the EoL) will not render a lightpath unacceptable. Planning the network to take advantage of the evolution of margins is done in multiperiod studies [18,19] as opposed to the typical one shot (single period) cross-layer optimization problems with fixed network parameters discussed in the previous paragraph. Moreover, reduced margins make a network susceptible to QoT degradations (soft failures) and thus affect the operation phase of the network as well.

Recovery from severe soft (QoT degradation) or hard (complete loss of signal) failures is typically performed by protection or restoration mechanisms [10]. In the first case, certain resources (e.g., transponders, spectrum) are dedicated and reserved for the failure. Thus, protection is an expensive solution and is mainly used for gold connections where the SLAs require guaranteed availability and rate. On the other hand, restoration is typically more affordable because it enables the sharing of resources but is slower because the resources have to be reserved after the failure has occurred.

EONs offer several flexibility dimensions, which can be used to regulate the QoT and to recover from QoT degradations according to actual conditions. In [20], an EON testbed with a real-time adaptive control plane was demonstrated that adjusted modulation format and spectrum positioning to recover the QoT of lightpaths with degraded OSNR. Similarly, the authors in [21,22] investigated the combination of modulation format adaptation and lightpath rerouting to restore impaired connections. Significant blocking and spectrum savings were reported in all cases. Finally, [4] investigated the reduction of margins and the adaptation of the modulation format when aging occurs. The results indicate a capacity increase of up to 63%.

To the best of our knowledge, there is no previous work that considers dynamic cross-layer optimization through QoT regulation in EONs. We propose a toolkit to maintain QoT close enough to the acceptable threshold so as to achieve high efficiency. This implies that the toolkit can be used not only to decrease the QoT when there are excessive margins (e.g., when the lightpath is initially deployed or when a failure is restored) but also to increase it (e.g., in the case of a soft failure). Our toolkit can also be used to adapt the rate of the lightpaths to meet changing demands while always ensuring appropriate QoT.

III. NETWORK SCENARIO

We assume EONs [9] with configurable transceivers that can adapt a number of transmission parameters, e.g., modulation format, baud rate, FEC, and spectrum used. The nodes are comprised of reconfigurable optical add/drop multiplexers (ROADMs) with flex-grid capabilities, and they are connected through uncompensated fiber links. Each fiber link consists of a number of fiber spans that terminate at an erbium-doped fiber amplifier (EDFA) that compensates the span loss. We assume there is no wavelength (or spectrum) conversion; thus, the wavelength (or spectrum) continuity constraint holds for each lightpath when it is established or reconfigured. In long connections, regenerators are placed, and each segment between regenerators is considered a separate lightpath that may use a different wavelength or spectrum.

We assume that the network conforms to the ORCHESTRA approach (Fig. 1) [7]. The vision of ORCHESTRA is to close the loop between the physical layer and the control plane, by using real-time impairment measurement capabilities of coherent optical transceivers. Coherent receivers deployed today employ DSP and already monitor certain physical layer impairments and compensate them. However, this information remains local. ORCHESTRA proposes to use/extend the monitoring capabilities of the coherent receivers, but the main contribution is that it exploits the monitoring information in global optimization decisions. This enables real cross-layer optimization and the ability to lower the network margins, thereby using the transceivers' capabilities to the fullest extent. Thus, in this study we consider that the coherent receivers function as OPMs; further, OPMs are located at the termination point of each lightpath (receiver) and can provide information about various physical parameters (e.g., residual dispersion, OSNR, Q factor) of the lightpath. Also OPMs are programmable so as to trigger alarms on specific measurable events. Alarms and monitoring information are transferred to the controller where they can be used to estimate QoT [6,23], localize problems [24], and make dynamic optimization decisions, which is the topic of this paper. ORCHESTRA develops a novel hierarchical monitoring plane to efficiently process alarms, monitor information, and avoid overwhelming the central controller [25]. Active control capabilities of monitoring elements and preprogramming are studied to enable ultrafast restoration [26]. The toolkit proposed in this paper is independent of the management and control plane, as long as monitoring information for margin calculation is provided and appropriate alarms are triggered when QoT becomes critical.

The ORCHESTRA network enables unparalleled network efficiency. ORCHESTRA novel feedback-based mechanisms make the lowering of margins possible; its dynamic optimization capabilities guarantee the adaptation of the network to the actual traffic and physical layer state. ORCHESTRA harvests the elastic capabilities of the transceivers and reconfigures their transmission parameters, thus regulating the QoT and trading off capacity and spectrum. In this paper, we present a toolkit that plays the role of the optimization logic in the ORCHESTRA dynamic network operation.

The toolkit examines three reconfiguration actions: i) FEC adaptation, ii) spectrum repositioning to control the interference from neighboring lightpaths, and iii) modulation format adaptation. Depending on the capabilities of the transponder and the related classes of service and SLAs, these adaptations can be combined with the adaptation of the baud rate in order to preserve or not the channel data rate. For example, FEC adaptation from 12% to 28% overhead results in respective reduction in the useful net rate, which can be subsequently compensated by an appropriate increase of the baud rate. Similarly, baud rate adaptation can be applied when the modulation format is changed.

The proposed toolkit exploits the push–pull technique [27] in order to move lightpaths in the spectrum domain hitlessly (without traffic interruption). The first step of the push–pull technique is to reserve contiguous free spectrum slots from the initial frequency f to the new frequency f', over the whole path. The second step is to retune the laser at the transmitter: the central frequency of the lightpath is slowly pushed from f to f'. Finally, the resulting unused spectrum slots are released.

The proposed algorithmic toolkit combined with the push–pull technique can provide the necessary functions to dynamically optimize the network in a number of use cases, which are discussed in the following.

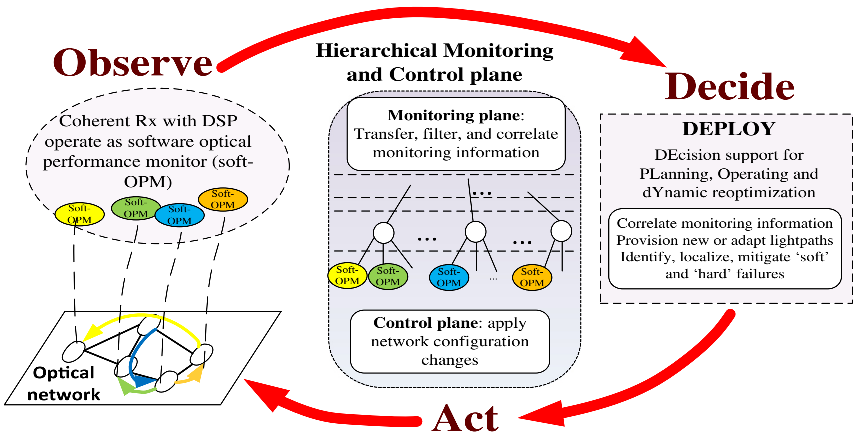

Fig. 1. ORCHESTRA observe–decide–act control cycle.

A. Use Cases

We assume that the network is operated close to its actual condition, in the sense that the lightpaths have reduced (i.e., just enough and not worst case) margins [1] in order to obtain high network efficiency.

One use case is the recovery from *soft failures*. In a “soft-failure” event, we want to avoid rerouting or adding new regenerators because both are considered expensive actions. Thus, our goal is to find the set of reconfiguration actions that solve the QoT problem at hand, while regenerator placement is considered a measure of last resort.

We classify the soft failures into two types: predictable and sudden. A predicable soft failure is when the QoT deteriorates slowly, e.g., due to aging of equipment or after a proper repair of a fiber cut, while a sudden soft failure is related to the malfunctioning of some equipment, e.g., EDFA pump problem, and results in a sharp deterioration of the QoT. When a soft failure renders the QoT unacceptable or degrades it close to the FEC threshold, we need to restore it. The current practice for predictable soft failures is to account for them in high margins, while sudden soft failures are partially covered by the margins. A sudden failure might be reported (or not) through the network management system and trigger the repair or replacement of the related equipment, but the network is bound to operate for a certain period (hours/days) with it. Thus, for sudden failures not covered by the high margins, the operators rely on the traditional protection/restoration mechanisms.

In an optical network operating with low margins, as the one proposed here, we can use the proposed toolkit to account for both types of soft failures. For predictable soft failures, we apply the toolkit as the network operates and according to the scheduled network maintenance cycles. The toolkit can decide on the full set of reconfiguration actions to restore the QoT when it approaches the FEC threshold. If the situation cannot be dynamically restored, then the lightpath can be rerouted or a regenerator can be deployed. All such decisions can be timed according to the next scheduled network maintenance window. As discussed, traditionally, we would deploy all the necessary equipment from the beginning-of-life (BoL), anticipating these predictable failures. Postponing the purchase of equipment has been shown to provide significant capital expenditure savings [18,19].

In the case of a sudden soft failure, a proactive approach is more appropriate. To be more specific, we can run the toolkit before we provision the lightpaths and decide proactively on the reconfiguration actions to execute or place regenerators and use the toolkit to survive from the sudden soft failure when it occurs. Such actions can also be preprogrammed [26] to achieve ultrafast restoration. If these operations are implemented in a real network, the concept and gains that can be achieved are close to those of protection/proactive restoration of hard failures. In proactive restoration, we take advantage of the independence of the failure events (shared risk groups), e.g., for single fiber cuts, we reuse equipment for different cuts [28]. In the sudden soft failure use case, protection is similar to planning with high margins, while proactive restoration enabled by the proposed toolkit can reuse the spectrum and the regenerators for different soft failures.

Apart from restoring the QoT, the proposed toolkit can be used to harvest the excess QoT. For example, prior to the establishment of a lightpath, the QoT is usually estimated using a physical layer model and high margins. To reduce the margins, we can use an accurate QoT estimator, such as [6], although such an estimator requires several lightpaths to be established and monitored. An alternative approach is to deploy the lightpaths with high margins (not only a high design margin but also high EoL system margins), monitor the actual margins after the deployment and use the proposed toolkit to dynamically reconfigure the parameters. For example, once the lightpath is established, and we observe that margins allow it, we can use a higher-order modulation format to support extra capacity and save in the installation of equipment by means of traffic grooming either in the current or in a future state of the network. A similar use case where we can harvest excessive margins is when malfunctioned equipment is replaced and the QoT improves.

Finally, another use case is when the traffic requirements change either periodically (e.g., typical daily traffic fluctuations, scheduled data-center backup) or suddenly (e.g., certain lightpaths need to be upgraded due to traffic increase/upgrades at the lower network segments). In these cases, the toolkit can be used to adapt the spectrum and the rate of the lightpaths while regulating (in both directions: positively and negatively) their QoT to ensure high efficiency.

IV. ADAPTATION TOOLKIT

We propose a toolkit that takes into account three reconfiguration actions in order to regulate the QoT of lightpaths: FEC, spectrum guard-band, and modulation format adaptation.

A. FEC Adaptation

The first technique we consider is FEC adaptation. FEC relies on transmitting redundant information that can be used to correct errors at the receiver. The performance of FEC (in terms of post-FEC BER) depends on the amount of the redundant information (overhead). In order for the useful net data rate to be preserved, a higher overhead requires either higher baud rate and therefore spectrum assuming Nyquist shaping, or a higher constellation modulation format. The baud rate can be adapted in a much finer granularity; thus, it is usually more suitable to compensate for small variations in the data rate. Therefore, FEC adaptation can be used to trade-off QoT with spectrum. For example, assuming lightpath provisioning with reduced margins, there are cases where the most robust FEC available is not used, e.g., the selected lower FEC yields acceptable QoT, while the most robust FEC requires an additional slot (e.g., for 25 Gbaud net baud rate, using 12% or 28%

FEC results in three or four slots for 28 or 32 Gbaud, respectively, assuming we want to keep the interference low and/or we cross more than six filters [29]). The most robust FEC has a lower pre-FEC BER threshold; as a result, a lightpath with unacceptable QoT with the lower FEC may have acceptable QoT when the higher FEC is employed. Thus, in this case the BER is not actually changed but the related threshold is.

The authors in [30] demonstrated hitless adaptation of the baud rate, while [31] demonstrated FEC adaptation without traffic interruption. When the extra spectrum is not available in all the links that the lightpath crosses, the push–pull technique [27] can be used to hitlessly shift the neighboring lightpaths in frequency. If a certain lightpath should not be moved (could, e.g., depend on the class of service), then other options can be investigated (e.g., pushing lightpaths at the other side or perform an adaptation where no shifting is required). Note that, to the best of our knowledge, current commercial transponders cannot perform baud-rate and FEC adaptation hitlessly and require the lightpaths to be instantly switched off, in which case a make-before-break approach can be used for the lightpath at hand. However, if its neighbors are required to be shifted, then no traffic interruption is necessary for them because the push–pull technique can be supported in commercial transponders. A limitation might exist for the maximum frequency range in which a lightpath can be shifted without traffic interruption. This can be taken into account by the shifting algorithm that we describe in Subsection IV.D.

B. Spectrum Guard-Band Adaptation

The second technique we consider is spectrum guard-band adaptation in order to control interference from neighboring lightpaths. In the case of a low QoT, the developed technique reduces interference by using spectrum as a guard band, that is, creating spectrum space between lightpaths. This can decrease the noise (e.g., due to crosstalk) or cross- and multichannel nonlinear interference, which in turn leads to reduced BER. Note that the considered technique reduces the out-of-band crosstalk for both add/drop and pass-through traffic that is created by neighboring channels. It does not account for the in-band crosstalk (between signals of the same frequency), which may occur, e.g., in the add/drop stages. Also, this technique is particularly useful in cases where there is misalignment of the filters (wavelength selective switches) due to aging or malfunction, a problem that gets greater the higher the number of filters/nodes crossed by a lightpath (e.g., in metro networks). Depending on the severity of a filter malfunctioning, it may or may not be self-reported. In this case, the spectrum guard band will not actually solve the root cause of the problem, but it can be a temporary solution until the filter is replaced. Spectrum guard-band creation involves shifting (using, e.g., push–pull) in frequency one or both spectrum neighbors from each side of the affected lightpath. This can result in a high number of cascading lightpaths needed to be shifted in frequency, which incurs high control plane overhead. The recursive lightpath spectrum shifting algorithm that we developed and is described in Subsection IV.D examines how to create the required guard band while minimizing the related control-plane overhead at the same time.

On the other hand, if there are excessive margins on the QoT of a lightpath, then the spectrum guard band can be reduced in order to save the spectrum, as long as the additional interference does not render the QoT unacceptable. A spectrum defragmentation algorithm such as [13] can be used for this purpose.

C. Modulation Format Adaptation

The third technique we consider is modulation format adaptation, which can be combined with baud-rate adaptation to maintain the original transmission rate if needed. In order to increase the QoT, we can reduce the modulation format in half (e.g., from PM-16QAM to PM-QPSK). If the original transmission rate should be preserved, then the baud rate can be increased as well (e.g., from 28 Gbaud to 56 Gbaud). The increase of the baud rate requires additional spectrum that can be acquired using the push–pull technique. The modulation format adaptation results in minor traffic disruption because the lightpath has to be switched off instantaneously. It also may require several other lightpaths to be shifted in frequency because doubling the baud rate requires more spectrum when compared with what the FEC adaptation requires. If the objective is to harvest excessive margins, then we can consider a higher-order modulation format with an appropriate baud-rate adaptation. Figure 2 depicts a network scenario where a modulation format adaptation is used

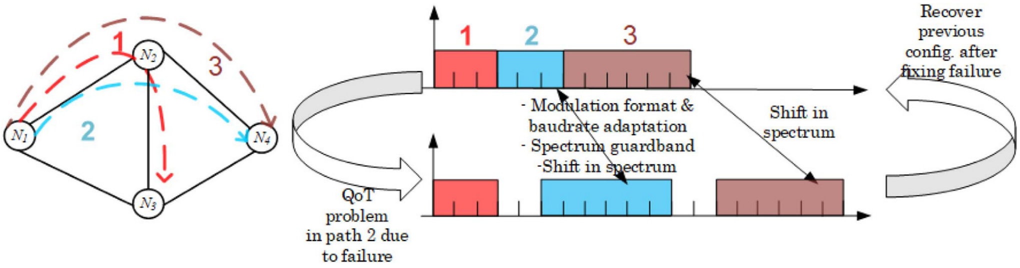

Fig. 2. Paths 2 and 3 are shifted in order to create guard bands and assign more slots to increase baud rate and decrease mod format of path 2. After the failure is identified and fixed, the network can return to the original state.

in order to restore the QoT of a lightpath affected by a soft failure. When the failure is repaired, then the toolkit can be used again to restore either the previous or another more appropriate configuration based on the current routing and physical layer conditions of the network.

D. Recursive Lightpath Spectrum Shifting Algorithm

The above techniques may result in the allocation to a lightpath of a certain number of additional slots, which could serve as a guard band or allow a baud-rate increase. This in turn might result in certain spectrum reconfigurations of some neighboring lightpaths if the required spectrum space is not available. To carry out such reconfigurations, we developed a recursive heuristic algorithm assuming the use of the push–pull technique [27]. The algorithm examines all the links of the considered lightpath, and it *recursively* shifts the neighboring lightpaths toward the same direction (in frequency) so as to clear the required number of slots. In doing so, it tries all the possible combinations of slots both higher and lower to the occupied spectrum of the considered lightpath; it calculates the total number of shifted lightpaths and chooses the combination resulting in the lowest number of recursively shifted lightpaths. Note that we do not assume circular shifting. This means that, if a connection, e.g., reaches the lower end of the spectrum, it will not be able to move circularly to the higher end of the spectrum, and another solution will be searched. The algorithm can also take into account any lightpaths that should not be shifted at all (due to certain policies) and also the maximum frequency range in which a lightpath can be practically shifted without traffic interruption. If any of the above constraints is violated, another solution is searched. This algorithm is a variation of the heuristic algorithm presented in [32], where an optimal algorithm assuming both spectrum shifting and reroutings is also given.

E. Toolkit Workflow

The aforementioned techniques can regulate the QoT of a lightpath, trading off QoT and margins for capacity and spectrum. The order in which they are investigated and executed can depend on various criteria. One criterion is the amount of spectrum that each technique requires. Under such a standard, a lowest spectrum utilization policy will lead to better use of the network resources. This results in a better ability to adapt the parameters of other lightpaths in the future and in a lower blocking rate. If it is imperative to not have any traffic interruption, then the spectrum guard band and the FEC adaptation (assuming transponders capable of hitlessly adapting the FEC) can be considered. The disadvantage of the spectrum guard-band adaptation method is that it can have a high control plane overhead because of the large number of lightpaths that may have to be shifted. In general, the control plane overhead is proportional to the amount of spectrum required because the latter results in shifting lightpaths, which in turn is the main contributor to the control plane overhead in the techniques we consider. Note that, apart from the adaptation actions we have described, other actions can be considered as well, such as power adaptation or reroutings. The only requirement is to define the relative cost of each action and the policy under which each action will be applied.

Having described the order in which the solutions can be searched, we need to estimate in a fast and accurate way the impact of each action on the QoT of a lightpath and the related margins. For this purpose, a QoT estimator such as [6] can be used. Alternatively, we can gradually apply the actions, monitor their effect, and adjust the configurations accordingly until we reach the targeted QoT margins.

V. SIMULATIONS

To evaluate the efficiency of the proposed toolkit, we performed a number of simulation experiments in MATLAB. We decided to focus on the soft-failure use cases that relate to the deterioration of the QoT of the lightpaths. In the considered scenarios, we assumed that the rate of the connections should be preserved (i.e., that all connections are gold class connections). The application of the proposed toolkit to the other use cases is part of our future endeavors.

We assumed a topology inspired by the Telecom Italia European backbone network (Fig. 3) with 49 nodes and 66 bidirectional links (we always assumed one pair of bidirectional links between two nodes, as opposed to the original network), SSMF fiber with attenuation coefficient 0.25 dB/km, dispersion parameter 16.7 ps/nm/km, and nonlinear coefficient 1.3 1/W/km. The span length was set at 100 km and the EDFA noise figure to 6 dB. We used the GN model [33] to approximate the behavior of the physical layer. We assumed 100 and 200 Gbps connections that are served using the following options: modulation format: PM-16QAM and PM-QPSK with transmission reaches of 1000 and 4600 km, respectively, under low load and 800 and 3800 km under high load, baud rate: 28, 32, 56,

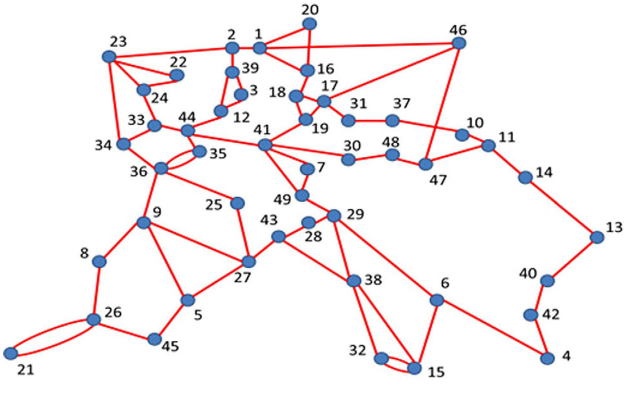

Fig. 3. Extended Pan-European backbone network.

64 Gbaud, and FEC: 12% and 28%, with pre-FEC BER thresholds of $6.3 * 10^{-3}$ (which translates to –2.2 dB and Q^2-factor of 7.94 dB) and $1.32 * 10^{-2}$ (–1.88 dB and Q^2-factor of 6.93 dB), respectively. The 12% FEC option is combined with either 28 or 56 Gbaud data rates, while the 28% FEC is combined with 32 or 64 GBaud in order to maintain the same useful net data rate. We assumed two traffic scenarios of 20 and 49.2 Tbps. The connections are derived from realistic traffic scenarios for the topology, and we always assumed shortest path routing. Further information about the statistics of the connections is stated in Table I. The regenerators required in the BoL for these connections in order for their QoT to be acceptable are 42 and 325 for 20 Tbps and 49.2 Tbps, respectively. The large number of regenerators for the second scenario is due to the additional number of connections most of which are 200G (see Table I).

We first calculated the BER savings (which correspond to the increase of the margins) each reconfiguration action can yield. All the metrics below are derived using the GN model's assumptions and capabilities. The FEC adaptation modifies the blocking threshold by 0.32 dB; however, a slight baud rate increase is required in order to preserve the original transmission rate. This translates to a maximum penalty of 0.1 dB from increased NLIs, so the resulting BER saving of the FEC adaptation is 0.22 dB. The creation of spectrum guard band yields at most 0.2 dB when the two direct neighbors (one at each side) are pushed and 0.3 dB when the four neighbors are pushed (which results in large control-plane overhead, due to the pushing of many lightpaths, and is not used in the subsequent simulations). However, depending on the spectrum allocation, in most cases the benefit is much less. Note that, in a real network, these values may vary significantly. Also, in cases where there is a drift in the frequencies of filters, the spectrum guard band can provide larger benefits because it will avoid the distortion of the signal at its edges and also reduce the (out-of-band) crosstalk between neighboring lightpaths. However, in our simulations, we do not take such effects into account. The modulation format adaptation from PM-QPSK to PM-BPSK along with the baud-rate adaptation (from 28 to 56, or from 32 to 64 Gbaud) yields maximum benefit of 1.6 dB, while the adaptation of PM-16QAM to PM-QPSK yields 2.5 dB, and that of PM-16QAM to PM-8QAM 2.2 dB. Note that the PM-8QAM metrics are provided only for comparison purposes. This modulation format is not used in the subsequent simulations.

TABLE I
CONNECTION STATISTICS

		20 Tbps	49.2 Tbps
Number of connections	100G	160	220
	200G	20	136
Length (km)	Mean	1382	1676
	Std	819	908
Number of hops	Mean	5.98	6.52
	Std	2.68	2.63

A. Proactive Restoration

In this section, we assume a given network state and a *single link at a time* SNR degradation of 1, 2, or 3 dB, which could, for example, correspond to equipment malfunctioning. Our objective is to plan the network to be survivable from any such failure. In particular, for each single link failure crossed by a lightpath, we examine whether the proposed toolkit can absorb the created QoT problems by reconfiguring the lightpaths that fall below the QoT threshold, and if not we place regenerators. Note that, in this way, we calculate the reconfiguration actions that should be performed when such a failure occurs on that link. These can be pre-programmed [26] so as to achieve fast restoration time. When the failure is repaired, the toolkit can be used again to restore the initial or another appropriate configuration. Regarding the algorithm once we examine a single link, then we return to the initial state of the lightpaths and examine independently the next single link failure. Regenerators placed (in the same node) can be re-used when examining a different link failure (following the concept of backup multiplexing restoration [28]). We compare the proposed reconfiguration toolkit with the concept of planning with high margins to absorb the QoT problems: in this case, we decide on transmission configuration (using the appropriate combination of the highest baud rate and most resilient modulation format) and place regenerators from the BoL so as to be able to absorb the QoT problems without having to do post hoc any kind of reconfiguration. One single-link failure is examined at a time, and we decide the transmission configuration of the lightpaths and place regenerators. The decisions for each connection are kept when we examine the next link failure. In essence, we assign from the *BoL* the transmission options (for *all* the connections and for *all* possible single link failures), which would have been used by our toolkit *if one* single link failure occurred. Both scenarios offer the same robustness against single link soft failures. In planning with high margins, the network operator does not know where a soft failure may happen. Therefore, all single link soft failures that lead to unacceptable QoT have to be planned for. Let's assume, for example, that the operator planned for the failure of only the worst single link degradation. Then the degradation of a different link may render the QoT of the connection unacceptable because the regeneration sites are fixed for each planned failure. The total required number of regenerators is the ultimate comparison metric; we also consider the total spectrum utilization, which, in the proposed restoration case, is defined as the maximum spectrum slot used under any single link failure. The comparison of our proposal to the concept of planning with high margins is showing the equipment savings in regenerators (which are expensive) that can be obtained by using our toolkit to repair a QoT problem *if* and *when* it appears instead of assigning the most robust transmission options for *all* the connections from the BoL. In any case, for a correct evaluation of the relative cost reductions that our toolkit provides, the network operator should take into account the costs for all the necessary equipment for the operation of the network (e.g., BoL regenerators,

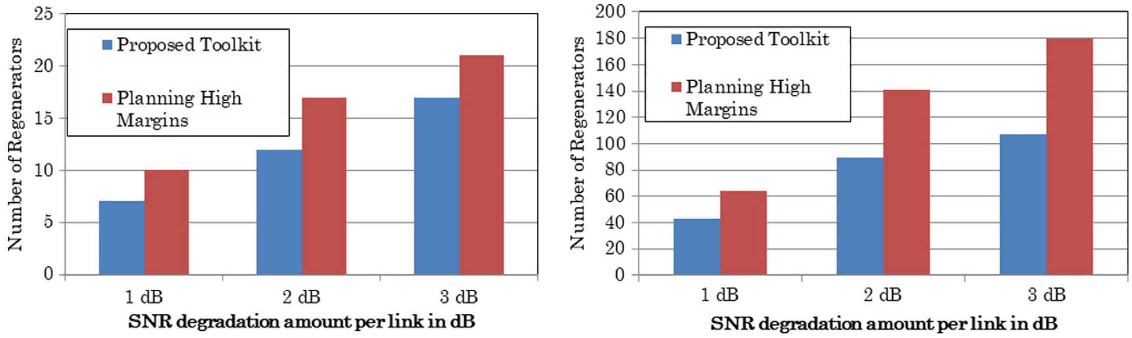

(a) (b)

Fig. 4. Number of regenerators required as a result of a single link soft failure of 1, 2, and 3 dB for (a) 20 Tbps and (b) 49.2 Tbps loads.

transponders, etc.). The spectrum utilization is of secondary importance so long as no blocking occurs.

Figure 4 presents the number of regenerators that are required as a result of a single link soft failure for traffic loads of 20 [Fig. 4(a)] and 49.2 Tbps [Fig. 4(b)]. Under both loads, our toolkit requires approximately at least 22% [Fig. 4(a) 3dB degradation] and at most 40% [Fig. 4(b) 3dB degradation] less regenerators than the high margin scenario. Note that, if we take into account the BoL regenerators required for both our toolkit and the high margin scenario, the relative savings are 7% and 16%, respectively. Our proposal requires at most 138 spectrum slots of 12.5 GHz each for the 20 Tbps load and at most 318 slots for the 49.2 Tbps. The spectrum utilization is 2% lower in the proposed solution compared to planning with high margins because we examine restoration for each single link failure instead of planning for all link failures. We do not provide relevant figures or tables for the spectrum utilization due to the small differences and space limitations. The spectrum savings are not large because the high margin scenario requires more regenerators, which relax the spectrum continuity constraint and result in shorter lightpaths, allowing a denser utilization of spectrum. The shorter lightpaths also allow the allocation of more efficient modulation formats, thereby further decreasing the utilization of the spectrum. We notice that using our toolkit results in significant equipment savings, as we harvest the available reconfiguration options whenever a failure occurs and not from the BoL and place regenerators when needed.

B. Dynamic Network Operation

In this section, we consider a dynamic network evolution scenario and present results from using the adaptation toolkit in this context. We assume a continuous time horizon and examine how the gradual deterioration of the physical layer (equipment aging, increasing number of splices to repair fiber cuts) affects the network. The QoT of the lightpaths deteriorate gradually, and, when the QoT of a lightpath falls beneath the related threshold, an alarm is triggered, and the proposed adaptation toolkit is used to reconfigure the problematic lightpath (and any other needed) to restore its QoT. In case the toolkit does not find a solution and cannot restore the QoT to acceptable levels, regenerators are placed. We assume that the alarm threshold is set appropriately so that enough time is given to allow the operation of the lightpath until the next network maintenance period, where the regenerator is placed. The maintenance period is considered to be six months.

As in the previous section, we again consider the same network topology, physical layer parameters, and realistic traffic profiles (20 and 49.2 Tbps). In a given period, we generate uniform soft failures for each link and a randomly generated soft failure to one (random) link at a time to account for uneven aging degradation. Note that, in this case, the intensity of the soft failure is quite lower than the soft failures modeled in the previous section. The effect of a failure may not immediately be noticed, and, as time passes and new (additional) failures arrive, a significant QoT degradation may occur. In particular, we consider a 10-year time horizon. In this time frame, we consider the following degradations according to [1]. We assume two fiber cuts for each span each contributing 0.3 dB loss at EoL. We also assume 0.7 dB EDFA noise degradation, again at EoL. Therefore, each span will have accumulated EoL degradation of 1.3 dB. We also consider 0.05 dB/filter EoL degradation for each node, and 0.5 dB EoL degradation for transponder aging. We assumed a linear (in dB) degradation over the 10 years, and we simulate the network at a month scale. Thus, for each month at each link we accumulate the related (constant) degradation. We also consider an additional 0.05 dB degradation per time unit (month) in one random link in order to introduce variation in the aging of the components. The maximum transmission reach at the EoL is 600 and 3100 km for PM-16QAM and PM-QPSK, respectively, under low load and 500 and 2700 km under high load.

We follow the state of the network over a total of 120 time-unit soft failures. After each soft failure, we examine the QoT of all the lightpaths. If the BER of a lightpath is below the acceptable threshold, we use the proposed adaptation toolkit to restore it. If no solution is feasible, then we place regenerators. We graph the number of regenerators that are required due to the failures, as a function of the number of time unit-induced soft failures. For comparison purposes, we also consider the case of planning with EoL margins. In this case, as in Subsection V.A, we assign the

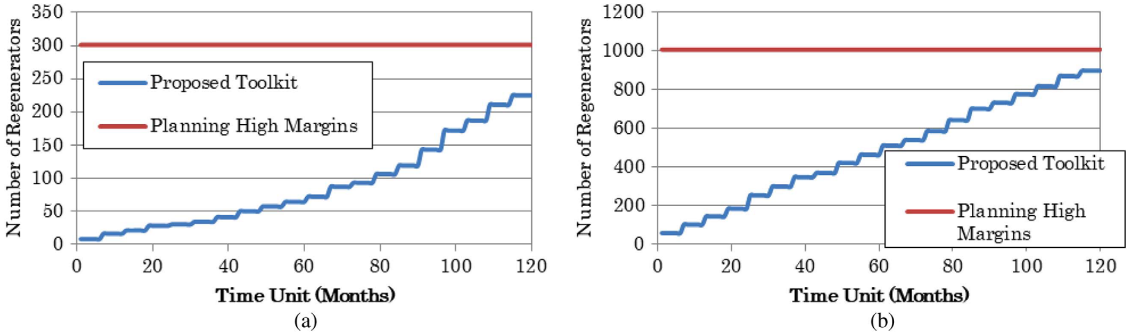

Fig. 5. Evolution of the number of regenerators required as a result of network aging for (a) 20 Tbps and (b) 49.2 Tbps loads.

appropriate combination of highest baud rate and most resilient modulation format for *all* the connections from the BoL. We estimate the QoT of all the lightpaths after 120 time-unit failures with the addition of a design margin, which is not necessary when our toolkit is used, as noted in Subsection III.A. We consider a 1 dB design margin that can be regarded as an appropriate value [2]. We calculate the required number of regenerators, and we assume that they are all placed from the BoL representing a high margin scenario. Therefore, in this section, we use our toolkit to *delay* the deployment of the necessary equipment and evaluate the respective benefits.

Figures 5(a) and 5(b) shows the number of deployed regenerators required due to aging as a function of time for two traffic scenarios (20 and 49.2 Tbps). Note that the number of regenerators increases in steps because, as previously mentioned, we have considered a six-month maintenance period. This means that, if our toolkit is used, then the regenerators for the 0 time unit actually correspond to the regenerators that will be needed at the sixth month. Also, note that, as we previously mentioned, both the high margin scenario and our toolkit require 42 and 325 BoL regenerators for 20 Tbps and 49.2 Tbps, respectively. These numbers are, however, small when compared with the 301 and 1008 regenerators that the operator should *additionally* place in *advance* in the high margin scenario to account for the future degradations. In Figs. 5(a) and 5(b), we note a significant difference in the number of regenerators between low and EoL margins. In particular, the low margins in the BoL, when combined with our proposed toolkit, require approximately 94% less regenerators when compared with the high margin scenario. If we take into account the initial regenerators required for the high margin scenario and our toolkit, then the difference is approximately 80%. Around year five, our toolkit requires approximately 50% less regenerators compared with planning the high margin scenario and 40% if we take into account the initial regenerators for both scenarios. Therefore, the purchase of a significant amount of equipment can be postponed for a considerable number of years. This results in substantial savings because their price generally decreases with time, and saved capital can be lent with interest. The benefits of this investment postponement have been demonstrated in [19]. We also observe that, in the 49.2 Tbps scenario, the total regenerators are much more than the low load 20 Tbps. This is, as we mentioned, due to the added set of connections, most of which are 200G, and also because the 200G connections use high modulation formats, operate closer to the FEC thresholds, and are more vulnerable to soft failures when compared with the 100G connections. We also note that, in the high load scenario, the number of regenerators after 10 years is closer to the high margin scenario when compared with the low load scenario. This can be attributed to the fact that the high network utilization results in less available spectrum for the adaptation toolkit to exploit. Note that the number of regenerators in both cases does not converge to that of the high margin scenario because the latter considers an additional design margin.

Figure 6 presents the number of failed connections per time period and the number of connections that were successfully reconfigured by our toolkit. Note that the

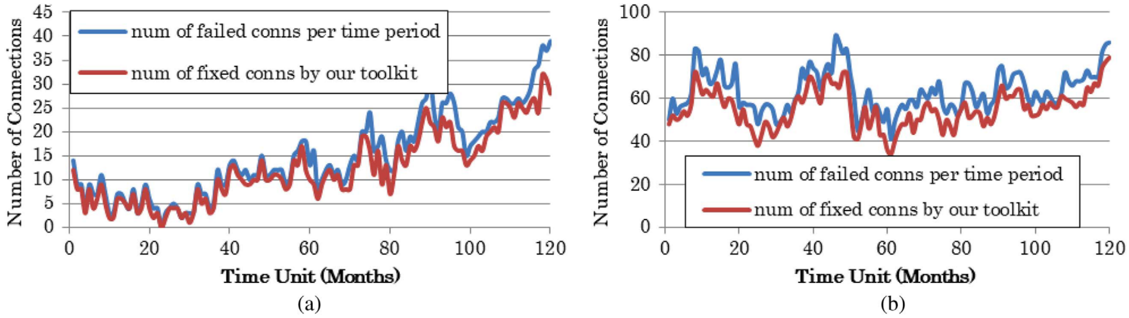

Fig. 6. Evolution of the number of failed and repaired connections due to aging in the network for (a) 20 Tbps and (b) 49.2 Tbps initial loads.

difference between these values is the number of regenerators (shown in Fig. 5). In the low load scenario [Fig. 6(a)], we note that, in the BoL, our toolkit can restore most of the failures that occurred. As the time passes and degradations accumulate, there are more failures that cannot be repaired. This can be explained as follows: as the degradations accumulate, they may increase over the maximum amount that our toolkit can absorb. Also, as time passes and reconfigurations are executed, spectrum utilization increases, which means there will be less available resources in the future to exploit in order to recover for additional failures. Also, we note that the number of failures does not increase monotonically. This can be explained as follows: as the degradations accumulate, they cause an increasing number of failures. At some point they reach a peak, where a large number of connections are affected and subsequently reconfigured. This causes future degradations to not cause many connection failures because most established connections are already fixed. As the degradations continue to accumulate, more failures start to appear for connections that may have been fixed in the past, which causes the next peak. In the high load scenario, we note that the number of failed connections is high and the percentage of connections that were repaired from our toolkit is less than in the low load scenario for the same reasons we explained for Fig. 5.

VI. CONCLUSION

We developed a toolkit that leverages the flexibility dimensions of elastic optical networks and regulates the QoT of the lightpaths to achieve high efficiency. The toolkit can be employed in a number of use cases and was used in this paper to restore the QoT of lightpaths in the case of soft failures. We performed simulation experiments where we modeled two soft failure scenarios. For sudden soft failures, we observed that we can save at least 22% and at most 40% in regenerators when compared with planning with high margins. We also observed significant postponement in the deployment of regenerators in a 10-year operating period of the network. Future work includes the application of the toolkit in other use cases such as the adaptation to traffic fluctuations and includes further actions such as power adaptation and rerouting.

ACKNOWLEDGMENT

I. Sartzetakis was supported by an IKY Greek State Ph.D. Scholarship, which was funded from the initiative "Support of the research human resources through the implementation of doctoral research" from the resources of Human Resources Development, Education and Lifelong Learning 2014–2020, co-funded by the European Social Fund and the Greek State. K. Christodoulopoulos was supported by an IKY post-doctoral scholarship co-funded by the European Social Fund and the Greek State. E. Varvarigos was supported by the ORCHESTRA project, funded by the EC (grant agreement 645360). The authors would like to thank M. Quagliotti for providing the realistic traffic scenarios and the anonymous reviewers for their constructive comments.

REFERENCES

[1] Y. Pointurier, "Design of low-margin optical networks," *J. Opt. Commun. Netw.*, vol. 9, no. 1, pp. A9–A17, Jan. 2017.

[2] J. L. Auge, "Can we use flexible transponders to reduce margins?" in *Optical Fiber Communication Conf.*, Anaheim, California, Mar. 2013, paper OTu2A.1.

[3] D. J. Ives, P. Bayvel, and S. J. Savory, "Assessment of options for utilizing SNR margin to increase network data throughput," in *Optical Fiber Communication Conf.*, Los Angeles, California, Mar. 2015, paper M2I.3.

[4] A. Mitra, A. Lord, S. Kar, and P. Wright, "Effect of link margin and frequency granularity on the performance of a flexgrid optical network," *Opt. Express*, vol. 22, no. 1, pp. 41–46, 2014.

[5] S. Oda, M. Miyabe, S. Yoshida, T. Katagiri, Y. Aoki, T. Hoshida, J. C. Rasmussen, M. Birk, and K. Tse, "A learning living network with open ROADMs," *J. Lightwave Technol.*, vol. 35, no. 8, pp. 1350–1356, Apr. 2017.

[6] I. Sartzetakis, K. Christodoulopoulos, C. P. Tsekrekos, D. Syvridis, and E. Varvarigos, "Quality of transmission estimation in WDM and elastic optical networks accounting for space–spectrum dependencies," *J. Opt. Commun. Netw.*, vol. 8, no. 9, pp. 676–688, Sept. 2016.

[7] K. Christodoulopoulos, P. Kokkinos, A. Di Giglio, A. Pagano, N. Argyris, C. Spatharakis, S. Dris, H. Avramopoulos, J. C. Antona, C. Delezoide, P. Jennevé, J. Pesic, Y. Pointurier, N. Sambo, F. Cugini, P. Castoldi, G. Bernini, G. Carrozzo, and E. Varvarigos, "ORCHESTRA—Optical performance monitoring enabling flexible networking," in *17th Int. Conf. on Transparent Optical Networks (ICTON)*, Budapest, Hungary, July 2015, pp. 1–4.

[8] I. Sartzetakis, K. Christodoulopoulos, and E. Varvarigos, "QoT aware adaptive elastic optical networks," in *Optical Fiber Communication Conf.*, Los Angeles, California, Mar. 2017, paper W4F.4.

[9] V. Lopez and L. Velasco, Eds., *Elastic Optical Networks: Architectures, Technologies, and Control*, Springer, 2016.

[10] D. Zhou and S. Subramaniam, "Survivability in optical networks," *IEEE Netw.*, vol. 14, no. 6, pp. 16–23, Nov./Dec. 2000.

[11] Y. Pointurier, M. Brandt-Pearce, S. Subramaniam, and B. Xu, "Cross-layer adaptive routing and wavelength assignment in all-optical networks," *IEEE J. Sel. Areas Commun.*, vol. 26, no. 6, pp. 32–44, Aug. 2008.

[12] K. Christodoulopoulos, P. Kokkinos, and E. M. Varvarigos, "Indirect and direct multicost algorithms for online impairment-aware RWA," *IEEE/ACM Trans. Netw.*, vol. 19, no. 6, pp. 1759–1772, Dec. 2011.

[13] M. Zhang, C. You, H. Jiang, and Z. Zhu, "Dynamic and adaptive bandwidth defragmentation in spectrum-sliced elastic optical networks with time-varying traffic," *J. Lightwave Technol.*, vol. 32, no. 5, pp. 1014–1023, Mar. 2014.

[14] L. Velasco, A. Castro, D. King, O. Gerstel, R. Casellas, and V. Lopez, "In-operation network planning," *IEEE Commun. Mag.*, vol. 52, no. 1, pp. 52–60, Jan. 2014.

[15] N. Dharmaweera, L. Yan, M. Karlsson, and E. Agrell, "An impairment-aware resource allocation scheme for dynamic elastic optical networks," in *Optical Fiber Communication Conf.*, Los Angeles, California, Mar. 2017, paper Th2A.19.

[16] M. Klinkowski, M. Ruiz, L. Velasco, D. Careglio, V. Lopez, and J. Comellas, "Elastic spectrum allocation for time-varying traffic in flexgrid optical networks," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 1, pp. 26–38, Jan. 2013.

[17] K. Christodoulopoulos, I. Tomkos, and E. Varvarigos, “Time-varying spectrum allocation policies and blocking analysis in flexible optical networks,” *IEEE J. Sel. Areas Commun.*, vol. 31, no. 1, pp. 13–25, Jan. 2013.

[18] J. Pesic, T. Zami, P. Ramantanis, and S. Bigo, “Faster return of investment in WDM networks when elastic transponders dynamically fit ageing of link margins,” in *Optical Fiber Communication Conf.*, Los Angeles, California, Mar. 2016, paper M3K.2.

[19] P. Soumplis, K. Christodoulopoulos, M. Quagliotti, A. Pagano, and E. Varvarigos, “Actual margins algorithm for multi-period planning,” in *Optical Fiber Communication Conf.*, Los Angeles, California, Mar. 2017, paper W4F.1.

[20] D. J. Geisler, R. Proietti, Y. Yin, R. P. Scott, X. Cai, N. Fontaine, L. Paraschis, O. Gerstel, and S. J. B. Yoo, “The first testbed demonstration of a flexible bandwidth network with a real-time adaptive control plane,” in *European Conf. and Expo. on Optical Communications*, Geneva, Switzerland, Sept. 2011, paper Th.13.K.2.

[21] X. Cai, K. Wen, R. Proietti, Y. Yin, D. J. Geisler, R. P. Scott, C. Qin, L. Paraschis, O. Gerstel, and S. J. B. Yoo, “Experimental demonstration of adaptive combinational QoT degradation restoration in elastic optical networks,” *J. Lightwave Technol.*, vol. 31, no. 4, pp. 664–671, Feb. 2013.

[22] K. Wen, X. Cai, Y. Yin, D. J. Geisler, R. Proietti, R. P. Scott, N. K. Fontaine, and S. J. B. Yoo, “Adaptive spectrum control and management in elastic optical networks,” *IEEE J. Sel. Areas Commun.*, vol. 31, no. 1, pp. 39–48, Jan. 2013.

[23] L. Barletta, A. Giusti, C. Rottondi, and M. Tornatore, “QoT estimation for unestablished lighpaths using machine learning,” in *Optical Fiber Communication Conf.*, Los Angeles, California, Mar. 2017, paper Th1J.1.

[24] I. Sartzetakis, K. Christodoulopoulos, and E. Varvarigos, “On reducing optical monitoring uncertainties and localizing soft failures,” in *Int. Conf. on Communication*, Paris, France, May 2017, paper ONS-IS03.4.

[25] N. Sambo, F. Cugini, A. Sgambelluri, and P. Castoldi, “Monitoring plane architecture and OAM handler,” *J. Lightwave Technol.*, vol. 34, no. 8, pp. 1939–1945, Apr. 2016.

[26] M. Dallaglio, N. Sambo, F. Cugini, and P. Castoldi, “Pre-programming resilience schemes upon failure through NETCONF and YANG,” in *Optical Fiber Communication Conf.*, Los Angeles, California, Mar. 2017, paper W1D.3.

[27] F. Cugini, F. Paolucci, G. Meloni, G. Berrettini, M. Secondini, F. Fresi, N. Sambo, L. Potì, and P. Castoldi, “Push-pull defragmentation without traffic disruption in flexible grid optical networks,” *J. Lightwave Technol.*, vol. 31, no. 1, pp. 125–133, Oct. 2013.

[28] M. Ilyas and H. T. Muftah, *The Handbook of Optical Communication Networks*, CRC Press, 2003.

[29] A. Morea, J. Renaudier, T. Zami, A. Ghazisaeidi, and O. Bertran-Pardo, “Throughput comparison between 50-GHz and 37.5-GHz grid transparent networks,” *J. Opt. Commun. Netw.*, vol. 7, no. 2, pp. A293–A300, Feb. 2015.

[30] A. Dupas, P. Layec, E. Dutisseuil, S. Bigo, S. Belotti, S. Misto, S. Annoni, Y. Yan, E. H. Salas, G. S. Zervas, and D. Simeonidou, “Hitless 100 Gbit/s OTN bandwidth variable transmitter for software-defined networks,” in *Optical Fiber Communication Conf.*, Anaheim, California, Mar. 2016, paper Th3I.1.

[31] F. Cugini, F. Fresi, F. Paolucci, G. Meloni, N. Sambo, A. Giorgetti, T. Foggi, L. Potí, and P. Castoldi, “Active stateful PCE with hitless LDPC code adaptation,” *J. Opt. Commun. Netw.*, vol. 7, no. 2, pp. A268–A276, Feb. 2015.

[32] P. Soumplis, K. Christodoulopoulos, and E. Varvarigos, “Dynamic connection establishment and network re-optimization in flexible optical networks,” *Photon. Netw. Commun.*, vol. 29, no. 3, pp. 307–321, June 2015.

[33] P. Poggiolini, “The GN model of non-linear propagation in uncompensated coherent optical systems,” *J. Lightwave Technol.*, vol. 30, no. 24, pp. 3857–3879, Dec. 2012.