

Multi-Period Planning With Actual Physical and Traffic Conditions

P. Soumplis,  K. Christodoulopoulos, M. Quagliotti, A. Pagano, and E. Varvarigos

Abstract—The common practice in optical networks is to deploy them statically in a future proof approach and to operate them with only traffic/fault driven changes. Lightpaths are overprovisioned with high quality of transmission (QoT) margins to anticipate future degradations, due to interference, aging, and maintenance operations, and to account for inaccuracies in QoT estimation. Such assumptions decrease network efficiency and increase the cost. Moreover, transponders and regenerators are assigned once to demands, yielding low efficiency if traffic varies significantly. Elastic optical technology, feedback from software monitors, and software defined networking enable the dynamic operation of the network. We propose an algorithm that takes into account the actual physical layer and traffic conditions and provisions new or adapts existing lightpaths with actual (just-enough) margins, optimizing placement and transmission parameter decisions for the transponders and regenerators. Using the proposed algorithm in a multi-period scenario, we quantify the benefits that can be achieved when planning with actual margins and varying levels of traffic dynamicity.

Index Terms—Elastic optical network (EON); Mixed line rate (MLR); Multi-period network planning; Physical layer impairments (PLIs); Quality of transmission (QoT); Routing and spectrum allocation (RSA); System and design margins; Traffic churn.

I. INTRODUCTION

The continuous growth of IP traffic and the emergence of new applications and services [1] require increased capacity but also a dynamic network operation. Coherent transceivers that combine coherent reception with high-speed electronics and digital signal processing (DSP) play a crucial role in the development of next generation optical networks, offering increased transmission rates and spectral efficiency [2]. Although coherent transceivers are able to compensate several physical layer impairments (PLIs), the physical layer still poses major limitations that can be addressed to achieve even higher efficiency.

Manuscript received August 2, 2017; revised November 5, 2017; accepted November 6, 2017; published December 22, 2017 (Doc. ID 303859).

P. Soumplis (e-mail: soumplis@ceid.upatras.gr) and K. Christodoulopoulos are with the Computer Technology Institute and Press—Diophantus, Rion 26504, Greece.

M. Quagliotti and A. Pagano are with Telecom Italia, Turin, Italy.

E. Varvarigos is with the School of Electrical and Computer Engineering, National Technical University of Athens, Greece. He is also with the Electrical and Computer Systems Engineering Department, Monash University, Melbourne, Australia.

<https://doi.org/10.1364/JOCN.10.00A144>

Optical connections (lightpaths) need to have acceptable quality of transmission (QoT) throughout their lifetime. The QoT changes—typically deteriorates—as time advances, due to equipment (fiber, transceivers, amplifiers, filters/switches) aging, accumulation of maintenance operations (e.g., reparation of fiber cuts), and increased interference coming from new established connections. The current practice is to operate the optical network statically and overprovision lightpaths with high physical layer margins to anticipate such future degradations and ensure uninterrupted operation until the end-of-life (EOL). However, in reality, the network operates most of the time far away from the EOL conditions; aging and interference play a considerable role only in the latest years, while predicting their effect after several years introduces uncertainties, requiring even higher margins [3,4]. The high margins translate to reduced optical reach, requiring the deployment of regenerators and more robust transponders than are strictly necessary at the time of (and several years after) the first installation.

Lowering the margins directly results in significant cost savings, by avoiding or postponing the purchase of equipment, the price of which decreases with time. Moreover, new and more advanced equipment becomes available with time, and thus postponing the purchase indirectly achieves a gradual upgrade of the network.

As time advances, additionally to the varying physical layer, traffic also changes. Traffic in optical transport networks is typically incremental; established lightpaths are never torn down and new lightpaths are added over the lifetime of the system, until a decision is made for the upgrade of the whole system using some new technology. However, the emergence of new applications can create traffic *churns* that deviate from the traditional incremental traffic model. For example, the disaggregation of network functions and their distribution in datacenters exhibits major variations in volume and traffic direction [5]. Moreover, network virtualization can change the service paradigms, requiring the dynamic slicing of the optical infrastructure and the dynamic re-optimization of lightpaths and equipment [6].

To harvest the advantages of lower margins and also account for the short- and long-term evolution of traffic, a more dynamic network is required. The network needs to be re-optimized over multiple periods to operate efficiently during each period for the specific physical and traffic conditions. Operating lightpaths with reduced margins

requires accurate knowledge of the physical layer and new control and management mechanisms. The ORCHESTRA project [7] proposes to close the control loop between the physical and the network layers, enabling unprecedented dynamicity and efficiency. ORCHESTRA relies on information from coherent receivers that operate as software optical performance monitors (soft-OPMs) and develops a monitoring and control plane to support and use such data. OPMs can be used to obtain accurate estimates of the interference and aging factors of the network so that new or re-optimized lightpaths have *just-enough* margins. OPMs and an active control plane are also helpful for anticipating and restoring the QoT problems that could occur later.

In this paper, we present an algorithm that incrementally plans the network and provisions lightpaths based on actual physical layer conditions—that is, with actual (just-enough) margins—and also takes into account any changes in traffic conditions. Note that we target the reduction of the long-term margins while there are still margins (design margin) to account for short-term QoT variations. Using the actual margins, the algorithm serves the (changing or not) demands and optimizes the decisions on the placement and reuse of the equipment (from previous periods), and on the reconfiguration of the transmission parameters of transponders and regenerators. We study the incremental planning of a network over multiple periods, where new demands are added/removed at specific intervals that correspond to the end of a period, when the network operator performs its maintenance/upgrade tasks. The addition of new equipment needs to be carefully planned, and is done by the proposed algorithm. The deployment of new or the relocation of the deployed equipment can be done in scheduled cycles, in a span of a few days, while the (re)configuration with the existing control plane protocols can be done in the order of seconds [8]. Note that the proposed algorithm is quite generic, and can be used with configurable [also referred to as bandwidth variable (BVT)] as well as non-configurable (fixed) transponders and for periods of any duration.

Using the proposed algorithm, we evaluate the gains of provisioning lightpaths with actual margins in an incremental multi-period planning scenario, for an elastic optical network (EON) and a mixed line rate (MLR) network. For comparison purposes, we also consider the related scenarios that use worst-case margins. The studies are performed under various traffic dynamicity scenarios, regulated by choosing appropriate churn rate values. We also examine the effect of equipment relocation, which is important in a network that operates with reduced margins and exhibits traffic churns.

Our results indicate that provisioning with reduced margins can yield significant cost savings that accumulate up to 36% for an EON and 27% for an MLR network, at the end of 10 periods with incremental traffic (zero churn rate and considering the relocation cost). As the traffic churn rate increases, the cost savings increase compared to the case in which new equipment is purchased. For a churn rate equal to 50%, the total cost under the reduced margins case

improves by 8% and 14% for the EON and MLR networks, respectively.

The rest of the paper is organized as follows. Section II reports on related work. Section III discusses QoT and margins. In Section IV, we present the heuristic algorithm for network planning with actual margins. Finally, in Section V we present performance comparison results, followed by our conclusions in Section VI.

II. RELATED WORK

Provisioning lightpaths for the actual network conditions, instead of making worst-case assumptions and using EOL margins, has received considerable attention over the years [9–14]. Initially, studies were performed for WDM networks, with the authors of Ref. [9] presenting a hierarchical routing and wavelength assignment (RWA) algorithm that accounts for the current optical signal-to-noise ratio (OSNR) and polarization mode dispersion (PMD) when establishing the lightpaths. In Ref. [10], a set of wavelength assignment techniques that take into account the crosstalk on the new and existing connections were proposed. The authors of Ref. [11] proposed optimal RWA algorithms based on integer linear programming (ILP) RWA formulations that account for the PLIs, including both path-dependent (ASE, SPM, etc.) and interference-dependent (XPM and XT) ones, through additional ILP constraints.

As EONs gained ground, routing and spectrum allocation (RSA) algorithms that use reduced margins were developed. The gains in spectrum and capacity obtained by avoiding the consideration of future losses were examined in Ref. [12]. Interference reduction through spectrum guard bands, examined through a wider set of transmission options, was studied in Ref. [13] to plan the network with reduced cost. The authors of Ref. [14] quantify the cost savings obtained by using elastic transponders to fit the aging degradation dynamically as opposed to planning with EOL aging margins.

To provision lightpaths close to the actual/current network conditions, we need to use a QoT estimation model (Q-tool) that is accurate for the specific network conditions. QoT estimation models range from analytical to simulations. Recently, the Gaussian noise (GN) model [15] has been introduced and shown to be quite accurate. Feeding this (or some other) QoT model with physical layer parameters that represent the current network aging and interference conditions, the required accuracy is obtained as described in Ref. [14]. Alternatively, feedback-based QoT estimation models, such as the one presented in Ref. [16], that use monitoring information and estimate QoT for the given network conditions can be used.

The key contribution is the development of an RSA algorithm that can be used to incrementally plan the network over multiple periods with evolving physical and traffic conditions. The proposed RSA algorithm considers the actual PLIs, taking into account both actual aging and interference. It also accounts for the evolution of traffic

and decides how to provision the lightpaths, reusing or purchasing new equipment, so as to minimize the total network cost. The proposed algorithm is quite generic and can be used for both fixed- and flex-grid switches with fixed or configurable (BVT) transponders, and extends the one presented in Ref. [17]. In Ref. [17], there were no constraints on traffic rerouting and relocation of equipment, constraints that are introduced in the extended algorithm presented here. Using the proposed algorithm, we performed extensive studies for two network topologies and traffic patterns of varying churn and quantified the benefits that margin reduction and equipment relocation can achieve.

III. QUALITY OF TRANSMISSION AND MARGINS

When planning or upgrading the network, the QoT of the lightpaths must be estimated, using a “Q-tool” based on some physical layer model. In QoT estimation, *system margins* are used to anticipate future degradation due to equipment aging, interference from increases in load, and failures until the EOL [3,18]. A typical assumption is that of worst-case interference, where interference is accounted for at its maximum value, as if the network operates at full load [19]. Moreover, to account for inaccuracies in the QoT estimation model, another margin, referred to as the *design margin*, is used on top of the EOL system margins [3,18]. The high margins result in low efficiency, expressed by deploying equipment that is not strictly necessary at the initial setup time. The reduction of margins makes equipment operate at high efficiency, closer to their limits, yielding significant cost savings, but requires a more dynamic network [14,20].

To examine the effect of the margin reduction, we developed an aging model, which we integrated into the GN model described in Ref. [15], and it is assumed in our studies to describe the physical layer. This model was used by our RSA algorithm (presented in Section IV) for estimating the QoT of established and new lightpaths with actual system margins.

We assume a dispersion uncompensated optical network, whose links consist of spans of single-mode fiber (SMF) followed by an erbium doped fiber amplifier (EDFA) that fully compensates the span losses. Figure 1(a) shows an example of the multi-span link model. Network nodes consist of optical cross connect (OXC) switches, implemented as reconfigurable optical add-drop multiplexers (ROADMs) based on wavelength selective switches (WSSs). The main network components whose performance deteriorates with time are the transponders, the fibers, the optical switches, and the amplifiers.

To determine the QoT of a lightpath (p, λ) according to the GN model, we need to calculate the total power spectral density $G_{ASE,p,\lambda}$ and $G_{NLI,p,\lambda}$ of the ASE noise and the NLIs accumulated across the path. The ASE noise depends on the traversed amplifiers (EDFA), which in turn depend on the traversed spans and nodes. The ASE noise contribution of a span depends on the length and attenuation coefficient $\alpha_{loss}(\tau)$ of the fiber segment, and the number of traversed connectors $c_c(\tau)$ and splices $s_s(\tau)$, which contribute




Fig. 1. OSNR model for (a) a multi-span link and (b) a ROADM for add, drop, and pass-through signals.

losses $c_{loss}(\tau)$ and $s_{loss}(\tau)$, respectively. All these introduced losses are assumed to be fully compensated by the EDFA at the end of the span, with noise figure $N_e(\tau)$. The above parameters are modeled as a (linear or nonlinear) function of time τ . Splices are used to repair the fiber cuts and are modeled as counting random processes following a specific distribution, while connectors are used to connect/disconnect the components and typically do not change with time. All the above losses increase as time advances, except for the number of connectors, which can be taken as constant.

Regarding the ROADMs, Fig. 1 shows their contribution to the noise added by the node amplifiers for add, drop, and pass-through traffic. Although these various types of traffic pass a different number of amplifiers and WSSs, to simplify our model, we assume that all traffic types accumulate equal loss $A_n(\tau)$. This value is assumed to increase with time τ (filters deteriorate and node amplifiers worsen their noise figure). Adding the ASE contributions of all traversed EDFAs (spans) and switches (ROADMs), we calculate the total power spectral density $G_{ASE,p,\lambda}$ of the ASE noise accumulated on the lightpath.

Taking into account the established lightpaths in the network, we can also calculate the power spectral density of the NLIs (both self- and cross-channel) for each span, for the actual wavelength utilization of the span. Then, assuming incoherent noise accumulation, we accumulate the NLI noise power over a link, and over the whole path $G_{NLI,p,\lambda}(\tau)$, for the current utilization of the links/network. $G_{NLI,p,\lambda}(\tau)$ is also a function of time, as the network load changes (typically increases) with time.

The OSNR of lightpath (p, λ) under the assumption of not suffering from intersymbol interference (ISI) and mismatched filters [21] is then given by

$$\text{OSNR}_{p,\lambda}(\tau) = \frac{T_{p,\lambda}}{G_{ASE,p,\lambda}(\tau) \cdot B + G_{NLI,p,\lambda}(\tau) \cdot \bar{B}}, \quad (1)$$

where $T_{p,\lambda}$ is the signal power at the receiver (assumed to be equal to the launch power for EDFAs that fully compensate the losses) and B is the equivalent noise bandwidth over which OSNR is evaluated.

We also model as a function of time the aging of the transponder with a margin $M_T(\tau)$ on its sensitivity. Finally, we

use a separate design margin $M_d(\tau)$ to account for QoT model inaccuracies for all lightpaths and to also avoid operating right at the limit, which could introduce ping pong effects. All aging parameters are assumed to increase with time, while the design margin can be either constant or decreasing with time. The assumption of a decreasing design margin is based on [3,22] and the fact that as time advances we obtain more accurate knowledge of the network and thus a more accurate QoT estimation.

To decide whether a lightpath (p, λ) is acceptable or not, we calculate its bit error rate:

$$\text{BER}_{p,\lambda}(\tau) = \Psi(\text{OSNR}_{p,\lambda}(\tau) - M_T(\tau) - M_d(\tau)), \quad (2)$$

where Ψ is a suitable function depending on the modulation format of the specific lightpath, as in Ref. [23]. The lightpath is established if and only if it has acceptable QoT:

$$\text{BER}_{p,\lambda}(\tau) \leq \text{BER}_{\text{limit}}, \quad (3)$$

where $\text{BER}_{\text{limit}}$ is the FECs limit.

IV. INCREMENTAL MULTI-PERIOD PLANNING WITH REDUCED MARGINS

We now present an RSA algorithm for provisioning lightpaths with just-enough margins in an incremental multi-period network planning scenario. The algorithm takes into account the positions of already deployed transponders and regenerators and the evolving traffic and is quite generic; although it is described for an elastic network (flex-grid and BVTs), it can also be used for fixed-grid and fixed transponders.

Typically, established connections in optical networks are not removed (and not reconfigured since they are established with EOL margins) and new connections are added at intermediate upgrade periods. To model a more dynamic traffic evolution, we assume that at the start of each period we are given the demands that remain, are added, and are removed with respect to the previous period. Each demand is expressed as a source–destination pair with a specific client rate and is served by a BVT transponder configured to have equal line rate. At each period, in the first step the proposed algorithm frees the transponders and regenerators that were used to serve the removed demands. In the second step, the algorithm checks the actual QoT of the remaining demands, and if it estimates that they will become unacceptable during the next period, it makes appropriate changes (e.g., reconfigures the BVT or adds regenerators). Finally, in the third step, lightpaths are established to serve the new demands. Note that transponders and regenerators deployed in previous periods that are freed at each period are considered to be reusable. Reuse of transponders and regenerators at their installed location is considered to be cost free, while moving (relocating) them is allowed but induces a penalty (which, if higher than the cost of buying a new one, cancels out the benefits of relocation).

In the following, we formally describe the multi-period planning problem and the proposed algorithm.

A. Formal Problem Description

The optical network is represented by a graph $G = (V, E)$, where V denotes the set of optical nodes and E the set of optical links. The spectrum is divided into spectrum slots of z GHz each, with one spectrum slot corresponding to the finer switching granularity of the flexible network elements (flex-grid switches and BVTs), and the network supports a total of F slots at each link. For each source–destination pair (s, d) we pre-calculate k paths. We let P_{sd} be the set of candidate paths for (s, d) . Then, for each demand (s, d, r) in $\Lambda^r(\tau_i)$ and $\Lambda^n(\tau_i)$, i.e., of client rate r , source s , and destination d for period τ_i , we have a number of candidate tuples t , those with effective data rate DR_t equal to r . Note that a demand is served by either a transparent lightpath or a translucent (i.e., requiring regenerators) connection. If regenerators are used, two consecutive regeneration points define a transparent sub-lightpath. Since the transmission reach depends on the current network state (including aging and interference) and in the pre-processing phase the spectrum has not been allocated yet, we calculate the possible node locations where regenerators can be placed, taking into account only the actual aging. Therefore, we pre-calculate the best and the worst transmission reaches for each transmission tuple, assuming zero and worst-case interference, respectively, and in the following we identify the possible placements for the regenerators. A selection of regeneration nodes is represented by a set M . Thus, in the pre-processing phase, for period τ_i and for each demand (s, d, r) we calculate a set $Q_{s,d,r}(\tau_i)$ of candidate (path, transmission tuple, regeneration points) triplets, each represented by (p, t, M) , taking into account the current aging of the network.

We assume a traffic scenario where at the start of period τ_i we have three sets of client demands: (i) the set of remaining client demands $\Lambda^r(\tau_i)$, (ii) the set of new client demands $\Lambda^n(\tau_i)$, and (iii) the set of deleted client demands $\Lambda^d(\tau_i)$. A demand from source node s to destination node d , asking for client rate r , is represented by the triplet (s, d, r) , and multiple demands of the same rate for a specific s – d pair may exist. Note that the remaining or deleted demands [types (i) and (iii)] in period τ_i were remaining or added demands in period τ_{i-1} , and all demands in the first period τ_0 are considered as new demands [type (ii)].

Traffic is served by BVTs that configure some or all of the following parameters: (a) modulation format, (b) baud rate, and (c) FEC. A possible transmission configuration is described by a tuple $t = (\text{MF}_t, \text{BR}_t, \text{OV}_t)$, where the modulation format MF_t (bits/symbol) describes the number of bits encoded in a symbol, the baud rate BR_t (symbols/s) describes the number of transmitted symbols per sec (including the FEC), and the FEC overhead OV_t is expressed as a decimal. Therefore, the effective data rate is $\text{DR}_t = \text{MF}_t \cdot \text{BR}_t \cdot (1 - \text{OV}_t)$. Assuming Nyquist WDM transmission and an overhead factor y to account for non-ideal pulse shaping and filters and spectrum slots, the transmission of tuple t requires $\lceil \text{BR}_t \cdot (1 + y) / z \rceil$ spectrum slots.

Given the set T that includes all possible transmission tuples t for a transponder, the RSA algorithm has to choose one of these options to serve a demand. The above description remains valid even for fixed transponders; a fixed transponder is characterized by a single tuple. Also different fixed transponders with different capabilities can be modeled for a so-called MLR scenario.

The proposed algorithm, at the start of a given period τ_i , takes as input the three sets of demands $\Lambda^r(\tau_i)$, $\Lambda^n(\tau_i)$, $\Lambda^d(\tau_i)$ and the previously installed (up to and including $\tau_i - 1$) lightpaths as well as their equipment. It also takes as input the available transponders and regenerators in this period and their costs (that may have changed from the previous period). We denote by $C_T(\tau_i)$ and $C_R(\tau_i)$ the cost of a new transponder and regenerator, respectively, at period τ_i , which typically falls as time advances. We also denote by $R(\tau_i)$ the relocation cost for both transponders and regenerators, which can evolve independently from the costs of the new equipment, and we use a weighting coefficient W to represent the relative cost of the utilized spectrum. We assume that the algorithm can estimate the actual system (aging and interference) margins for the current period with a specific design margin using an accurate Q-tool (see Section III).

B. Algorithm Description

The algorithm consists of three steps for allocating paths, transponders, and regenerators. In *step 1* (Fig. 2), the transponders, regenerators, and spectrum slots allocated to each deleted demand in the set $\Lambda^d(\tau_i)$ are released and become available for serving other demands.




Fig. 2. Flowchart of step 1 of the heuristic algorithm.

In *step 2*, we examine the set of remaining demands $\Lambda^r(\tau_i)$ in this period and do so in three rounds. In round 2.1 of step 2 [Fig. 3(a)] we estimate for each demand of the following period the QoT of the corresponding light-path(s) (single lightpath if no regenerator is used or series of lightpaths if regenerators are used). If the QoT is acceptable, we keep the same transponder and regenerators, while for the lightpaths that have unacceptable QoT we try to reconfigure the utilized transponder/regenerator(s). If this does not yield acceptable QoT, the algorithm frees the regenerators used by the demand and processes it in the next round (round 2.2). In round 2.2 [Fig. 3(b)], we check one-by-one the remaining demands from round 2.1 (demands with at least one lightpath of unacceptable QoT). For each such demand, we try to allocate (reuse) free regenerators at some intermediate nodes, using at most one more than those allocated before. Assuming that in the previous period the demand was served with the (p', t', M') triple, the algorithm checks the candidate (p, t, M) triplets of set $Q_{s,d,r}(\tau_i)$ with $|M| \leq |M'| + 1$. The regenerators that can be reused are those freed in step 1 or round 2.1 of the algorithm. The demands for which there do not exist free regenerators at the required locations are not served in round 2.2 and are left to be processed in round 2.3. Finally, in round 2.3 [Fig. 3(c)], we process the demands for which we were unable to find free regenerators at their locations. For each such demand, all triplets (p, t, m) in $Q_{s,d,r}(\tau_i)$ are examined and the related cost $C_{p,t,m}$ is calculated considering the relocation of any free regenerators left unused up to this point and the deployment of new regenerators. For the specific set of regenerator locations m , we allocate n_c available regenerators at the required position, and for the $|m| - n_c$ remaining ones we select to relocate n_r and deploy n_n new ones, where n_r and n_c are chosen to minimize the total demand cost:

$$C_{p,t,m}(\tau_i) = \min_{n_n + n_r = |m| - n_c} (n_n C_R(\tau_i) + n_r R(\tau_i)). \quad (4)$$

Having examined all the candidate triplets (p, t, m) , we select that with the minimum cost:

$$\min_{(p,t,m) \in Q_{s,d,r}(\tau_i)} (C_{p,t,m}(\tau_i) + W \cdot S_{p,t,m}(\tau_i)), \quad (5)$$

where $S_{p,t,m}$ is the spectrum used by choice (p, t, m) and W is a weighting coefficient transforming the used spectrum into cost.

At the start of step 3, we have a set of free transponders (from step 1) and remaining regenerators (from step 2) and process the new demands in $\Lambda^n(\tau_i)$, in two rounds. In round 3.1, we examine for each such demand to reuse transponders and regenerators (without relocation) for the transmission options that require the minimum number of regenerators (similar to round 2.2). In round 3.2, we examine the relocation of transponders and regenerators and the deployment of new ones, and select the solution with the lowest cost (similar to round 2.3 with the only difference that we need to account for the transponders in addition to regenerators).




Fig. 3. Flowchart for (a) round 2.1, (b) round 2.2, and (c) round 2.3 of step 2.

The above describes the allocation of transponders, regenerators, and paths to demands. For each chosen path, the algorithm also allocates spectrum. Once all lightpaths are provisioned and the spectrum is allocated, the algorithm checks again the QoT taking into account the actual interference (recall that the candidate triplets used above did not consider interference). Demands whose QoT is unacceptable due to interference are processed again in step 2, and this is repeated until all demands are served.

V. PERFORMANCE EVALUATION

In this section, we present our multi-period network planning studies. In particular, we carried out a multi-period analysis over 10 periods for two network topologies and different traffic profile scenarios. In the obtained results, each period corresponds to a year, but we decided to keep the term period for generality purposes. The results are presented with a step of two periods.

The first studied topology, shown in Fig. 4(a), was inspired by Telecom's Italia national network and consists of 29 nodes with link lengths up to 382 km. The second topology, shown in Fig. 4(b), was inspired by Telecom's Italia Pan-European backbone network and consists of 49 nodes with link lengths up to 1251 km. For both topologies, we assumed that the network employs ROADMs and uncompensated SMF links. The amplification span was taken equal to 100 km, with each span followed by an EDFA that fully compensated the attenuation of the previous span. The gain of the EDFAs was assumed flat over all spectrum (no gain ripple/tilting effects). To model the physical layer, we extended the GN model to consider network aging (see Section III). We used the margin parameters reported in Table I, which shows the values for the two extreme time instants, beginning of life (BOL) or τ_0 and end of life (EOL) or τ_{10} ; the values for intermediate periods were calculated through linear (in time and in decibels) interpolation. When a transponder is deployed at an intermediate period, its related margin starts with the BOL value from that specific period.

We performed our studies for two network settings: (i) EONs and (ii) fixed-grid WDM. In the elastic network setting, we assumed two types of BVT transponders. Both types support baud rates up to 43 Gbaud and modulate using dual polarization (DP) up to DP-16QAM with (i) one subcarrier and transmission rate up to 200 Gbps and (ii) two subcarriers with transmission rate up to 400 Gbps. For the WDM network, we assumed an MLR case with three types of fixed transponders of 100, 200, and 400 Gbps transmission rates.

Table II presents the cost of the network equipment in the first period (τ_0), relative to the cost of a 100 Gbps fixed transponder (taken to be 1 C.U.).

These values correspond to $C_R(\tau_0)$ and $C_T(\tau_0)$, as described in Section IV. Using these as starting values, the cost of the network equipment for the following periods can be calculated using learning curves [24]. The basic idea of such models is that the cost of performing a task




Fig. 4. Topologies inspired by the Telecom Italia's European backbone network.

decreases with the number of the repetitions and thus they are based on the forecasted production number between the different periods. The different cost models we examined did not reveal significant differences in the achieved cost savings compared to the fixed 10% depreciation per period, which is used in our performance analysis. The cost of the equipment relocation, corresponding to the reuse of transponders or regenerators deployed during previous periods at different nodes, was taken to be fixed for all periods τ_i and equal to $R(\tau_i) = 0.1$ C.U.

In all cases, we assumed the transmission power of the transponders to be fixed at 0 dBm and the use of LDPC (4161, 3431, 0.825) FEC code, with 21.2% overhead and pre-FEC BER limit of 10^{-2} . The bandwidth overhead for non-ideal pulse shaping and filters was taken to be

TABLE I
SYSTEM AND DESIGN MARGINS BOL AND EOL VALUES

Margin		BOL τ_0	EOL τ_{10}	
System Aging margins	Fiber attenuation α_{loss} (dB/km)	0.22	0.23	
	Connector loss c_{loss} (dB)	0.20	0.30	
	Splice loss s_{loss} (dB)	0.30	0.50	
	Number of slices s_e (km^{-1})	0	0.027	
	Number of connectors per span c_e	2	2	
	EDFA noise figure N_e (dB)	4.50	5.50	
	OXC loss A_n (dB)	20.00	23.00	
	Transponder margin M_T (dB)	1.00	1.50	
		Interference	Empty	Full
	Design margin M_d (dB)		2.00	1.00

TABLE II
RELATIVE COSTS OF EQUIPMENT VALUES AT TIME τ_0

Network Equipment (Q)	Unitary Price (C.U.)
100 Gbps fixed transponder/regenerator	1.00
200 Gbps fixed transponder/regenerator	1.20
400 Gbps fixed transponder/regenerator (introduced at period τ_4)	1.36
Elastic transponder/regenerator 200 Gbps	1.30
Elastic transponder/regenerator 400 Gbps (introduced at period τ_4)	1.47
Transponder/regenerator relocation cost (fixed for all periods)	0.10

$y = 0.15$. The width of the spectrum slot was taken to be $z = 12.5$ GHz. Each fiber link supported $F = 320$ slots/fiber for the elastic network setting, while for the fixed-grid we used $z = 50$ GHz and $F = 80$ wavelengths/fiber. An extra fiber was utilized at specific links when the available spectrum was not enough to serve the demands.

Details regarding the modulation format and the baud rate configurations of the 400 Gbps BVT and the related reaches for the different margin cases are given in Table III. The transmission capabilities of the MLR transponders and the related reaches are presented in Table IV. Note that the reaches presented in these tables do not account for the node losses, which depend on the particular path, and thus are presented for reference.

Regarding interference, we examine two extreme cases. Initially, in the BOL state, we assume that the network is empty, meaning that few connections (three for the GN model to give accurate results) are established. On the other hand, at the EOL state we assume that all neighboring channels are active and densely spaced. When planning with reduced margins, the transmission reach varies with time between the first and the third column, while when planning with EOL margins the transmission reach corresponds to the second column of Tables III and IV. In the

TABLE III
400 GBPS BANDWIDTH VARIABLE TRANSPONDER (BVT) CONFIGURATIONS AND REFERENCE TRANSMISSION REACHES FOR THE DIFFERENT MARGIN CASES

Data Rate (Gbps)	Baud Rate (Gbaud)	Mod Format	Reach (km)		
			BOL Aging and BOL Interference and BOL Design	EOL Aging and EOL Interference and BOL Design	EOL Aging and EOL Interference and EOL Design
100	16	DP-16QAM	1200	600	700
100	32	DP-QPSK	4800	2000	2600
200	32	DP-16QAM	1000	400	500
200	43	8QAM	1600	700	800
200	2 × 32	DP-QPSK	4500	2000	2600
400	2 × 32	DP-16QAM	900	400	500
400	2 × 43	8QAM	1400	700	800

TABLE IV
MLR TRANSPONDERS AND REFERENCE TRANSMISSION REACHES FOR THE DIFFERENT MARGIN CASES

Data Rate (Gbps)	Baud Rate (Gbaud)	Mod Format	Reach (km)		
			BOL Aging and BOL Interference and BOL Design	EOL Aging and EOL Interference and BOL Design	EOL Aging and EOL Interference and EOL Design
100	32	DP-QPSK	4800	2100	2700
200	43	DP-8QAM	1600	700	800
400 (2 × 200)	2 × 32	DP-16QAM	1200	500	600

simulations the QoT of each lightpath is calculated taking into account the OXC that it passes, as well as the interference of the other lightpaths.

Regarding the served traffic, each period is described by the three sets of remaining, new, and deleted demands denoted by $\Lambda^r(\tau_i)$, $\Lambda^n(\tau_i)$, and $\Lambda^d(\tau_i)$, respectively. We assumed that the total network traffic, expressed by the summation of remaining and new demands, increases with a 25% compound growth rate (CAGR) per one period. The number of changing demands (deleted and added as new with a possible increase of the rate) is described through the churn rate parameter ch . As discussed above, currently in optical transport networks the demands are not deleted, which corresponds to a churn rate $ch = 0$. In our studies, we consider churn rates ranging from 0% to 50% so as to examine the effect of traffic dynamicity that can appear with the emergence of new application and services.

We compare planning under worst and actual margins over 10 periods for the two network topologies and the two network settings—that is, elastic and fixed-grid MLR.

In Fig. 5, we present the total cost in C.U. considering the acquisition of transponders and regenerators for incremental traffic, without churn ($ch = 0\%$), where connections established at various periods remain until the last period. For both examined topologies, the elastic network scenario achieves the lowest cost at the end of the examined period, although in the initial periods the MLR scenario has lower cost. This happens because initially the demands are low (up to 200 Gbps) and the MLR takes advantage of the lower transponders' and regenerators' costs. As time passes, the elastic network case manages to overcome it, reconfiguring the already used equipment.




Fig. 5. Accumulated total cost (in C.U.) for (a) topology of Fig. 4(a) and (b) topology of Fig. 4(b).

Furthermore, for the topology of Fig. 4(b) (with the longer links), the difference in the last period increases significantly because the demands are high and the MLR uses more regenerators to serve the demands. For the topology of Fig. 4(a), the cost savings accumulate to 12% at the end




Fig. 6. Accumulated number of dislocations for the elastic and MLR networks for the different planning scenarios for churn rate (a) 0%, (b) 5%, (c) 20%, and (d) 50%.

of 10 periods for an elastic and 30% for an MLR network, while for the topology of Fig. 4(b) these savings accumulate to 36% and 27%, respectively. The elastic network seems to meet better the requirements for next generation optical networks, managing to overcome in both cases the accumulated cost of the MLR network for 5% and 29% for the topologies of Figs. 4(a) and 4(b), respectively.

The number of relocations of transponders and regenerators is presented in Figs. 6(a)–6(d). In Fig. 6(a), the churn rate is 0% and all established demands remain until the

end of the simulation period. In that case, the MLR reduced margins case does more relocations compared to the elastic case, which is able to satisfy QoT requirements by changing transponder configurations, moving to new adequate tuples each time. The number of relocations in the EOL cases for both network scenarios is equal to zero, as the demands are established at each period with EOL assumptions and do not require reconfigurations or equipment relocation.

For lower churn rates smaller than 20%, the reduced margins MLR makes the most relocations, while for churn rates greater than 20%, the MLR EOL margins perform the maximum number. Comparing the elastic and the MLR cases, the former requires fewer relocations for all examined churn rates, by exploiting the transmission reconfigurations supported by elastic transponders. In addition, for the elastic network, the EOL case underperforms compared to the reduced margins case. This is because EOL margins make elastic transponders unable to take significant advantage of transmission parameter reconfigurations.

In Fig. 7(a) we present the savings achieved when relocations are prohibited and new equipment has to be purchased to serve demands when there are no transponders/regenerators available at the source node of the light-path. In Fig. 7(b) we present the corresponding savings when relocation of equipment is allowed. As expected, the advantages obtained by the lower cost equipment relocation become more significant as the churn rate increases. Furthermore, the MLR network relocates more equipment compared to the elastic case and consequently has increased savings compared to the elastic case. On the other hand, the elastic case proves to be the most adequate for




Fig. 7. Total cost when (a) acquisition of new or (b) relocation of already established equipment is considered.

all the traffic variation cases, with steady performance, which proves to not be significantly affected by the traffic dynamicity.

VI. CONCLUSION

Lightpaths in optical networks are currently provisioned with worst-case margins, calculated under EOL system assumptions for the interference, the aging of equipment, and the maintenance tasks, while a design margin is also used to account for QoT estimation inaccuracies. The impairment measuring capabilities of the coherent optical transponders offer the ability to operate the network closer to its true capabilities, improving efficiency and postponing or avoiding expenditures. To evaluate such benefits, we proposed an RSA algorithm that provisions lightpaths taking into account the actual network conditions so as to establish them with just-enough margins. Our comparison study quantified the cost benefits of planning with actual as opposed to worst-case margins. In a multi-period planning scenario, some established demands may be terminated, freeing the used resources, while new demands may arrive. We observed savings up to 36% for the elastic network and up to 27% for the MLR when planning with reduced margins compared to the case where EOL assumptions are made. Finally, the reuse of equipment (with a relocation cost) offers extra savings, reducing up to 8% and 14% the cost of the elastic and the MLR network, respectively.

ACKNOWLEDGMENT

P. Soumplis, M. Quagliotti, A. Pagano, and E. Varvarigos were supported in part through the Horizon 2020 ORCHESTRA project (grant agreement 645360).

REFERENCES

- [1] "Cisco Visual Networking Index: Forecast and Methodology, 2015–2020," Cisco White Paper, 2016 [Online]. Available: <http://www.cisco.com>.
- [2] "Nokia ushers in 100G transport services era with new programmable silicon chipset and next-generation optical networking systems," Nokia, 2016 [Online]. Available: <http://nokia.ly/2lwnMqQ>.
- [3] Y. Pointurier, "Design of low-margin optical networks," *J. Opt. Commun. Netw.*, vol. 9, pp. A9–A17, 2017.
- [4] J.-L. Auge, "Can we use flexible transponders to reduce margins?" in *Optical Fiber Communication Conf. (OFC)*, 2013.
- [5] Y. Chen, S. Jain, V. K. Adhikari, Z. L. Zhang, and K. Xu, "A first look at inter-data center traffic characteristics via Yahoo! datasets," in *IEEE INFOCOM*, Shanghai, China, 2011, pp. 1620–1628.
- [6] A. Fischer, J. Botero, M. Beck, H. De Meer, and X. Hesselbach, "Virtual network embedding: A survey," *IEEE Commun. Surv. Tutorials*, vol. 15, no. 4, pp. 1888–1906, 2013.
- [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 *Int. Conf. on Transparent Optical Networks (ICTON)*, 2015.
- [8] S. Gringeri, N. Bitar, and T. J. Xia, "Extending software defined network principles to include optical transport," *IEEE Commun. Mag.*, vol. 51, no. 3, pp. 32–40, Mar. 2013.
- [9] Y. Huang, J. P. Heritage, and B. Mukherjee, "Connection provisioning with transmission impairment consideration in optical WDM networks with high-speed channels," *J. Lightwave Technol.*, vol. 23, no. 3, pp. 982–993, 2005.
- [10] T. Deng, S. S. Subramaniam, and J. Xu, "Crosstalk-aware wavelength assignment in dynamic wavelength-routed optical networks," in *Int. Conf. Broadband Networks (BroadNets)*, 2004, pp. 140–149.
- [11] K. Christodoulopoulos, K. Manousakis, and E. Varvarigos, "Offline routing and wavelength assignment in transparent WDM networks," *IEEE/ACM Trans. Netw.*, vol. 18, no. 5, pp. 1557–1570, 2010.
- [12] A. Mitra, S. Kar, and A. Lord, "Effect of frequency granularity and link margin at 100G and beyond flexgrid optical networks," in *Nat. Conf. on Communications (NCC)*, 2014.
- [13] K. Christodoulopoulos, P. Soumplis, and E. Varvarigos, "Planning flexible optical networks under physical layer constraints," *J. Opt. Commun. Netw.*, vol. 5, no. 11, pp. 1296–1312, 2013.
- [14] 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 Communications Conf. (OFC)*, 2016.
- [15] P. Poggiolini, "The GN model of non-linear propagation in uncompensated coherent optical systems," *J. Lightwave Technol.*, vol. 30, no. 24, pp. 3857–3879, 2012.
- [16] 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, pp. 676–688, 2016.
- [17] P. Soumplis, K. Christodoulopoulos, M. Quagliotti, A. Pagano, and E. Varvarigos, "Actual margins algorithm for multi-period planning," in *Optical Fiber Communications Conf. (OFC)*, 2017.
- [18] J. Auge, "Can we use flexible transponders to reduce margins?" in *Optical Fiber Communications Conf. (OFC)*, 2013.
- [19] K. Christodoulopoulos, K. Manousakis, E. Varvarigos, and M. Angelou, "Considering physical layer impairments in offline RWA," *IEEE Netw.*, vol. 23, pp. 26–33, 2009.
- [20] "Transforming Margin into Capacity with Liquid Spectrum," Ciena White Paper, 2016, <http://www.ciena.com/insights/white-papers/Margin-Capacity-Liquid-Spectrum-prx.html>.
- [21] P. Poggiolini, G. Bosco, A. Carena, V. Curri, Y. Jiang, and F. Forghieri, "A detailed analytical derivation of the GN model of non-linear interference in coherent optical transmission systems," arXiv:1209.0394, 2012.
- [22] E. Seve, J. Pesic, C. Delezoide, and Y. Pointurier, "Learning process for reducing uncertainties on network parameters and design margins," in *Optical Fiber Communications Conf. (OFC)*, 2017.
- [23] M. Seimetz, *High-Order Modulation for Optical Fiber Transmission*, Springer Series in Optical Sciences. Berlin, Germany: Springer, 2009.
- [24] W. J. Morse, "Reporting production costs that follow the learning curve phenomenon," *Account. Rev.*, vol. 47, no. 4, pp. 761–773, 1972.