

Exploiting Spectrum Sharing for Soft Failure Recovery

I. Sartzetakis, N. Sambo, K. Christodoulopoulos, P. Castoldi, and E. Varvarigos

Abstract—Elastic optical networks (EONs) offer several optimization dimensions that can be harvested to increase network efficiency and survivability. More specifically, the adaptation of the modulation format to a more robust one can recover the quality of transmission (QoT) of a lightpath affected by a soft failure (i.e., a physical layer degradation resulting in QoT deterioration). In this scenario, increasing the robustness of a transmission may require a spectrum increase. Thus, the spectrum also might have to be reconfigured. In this paper, we propose spectrum sharing between lightpaths of classes requiring different rate guarantees. In the case of a soft failure, a high-priority, rate-guaranteed lightpath can expand its spectrum to recover its QoT at the expense of spectrum allocated to a lower-priority, guaranteed minimum rate lightpath. We propose an integer linear programming (ILP) routing and spectrum assignment (RSA) algorithm for the planning phase and an online RSA for the operation phase of the network that optimize spectrum sharing among lightpaths of different classes. Using these proposed techniques can effectively reduce the network costs to guarantee the rate of high-priority connections.

Index Terms—Network optimization; Network survivability; Rate guarantees; Reconfiguration; Routing and spectrum allocation; Soft failures; Spectrum sharing.

I. INTRODUCTION

Elastic optical networks (EONs) offer several optimization dimensions, such as spectrum, modulation format, and symbol (baud) rate adaptation. Depending on the transceivers used, these parameters can be configured independently per lightpath to maximize the utilization of the network's equipment and the overall efficiency. As a result, EONs promise significant benefits, which include higher spectral efficiency, increased capacity [1], and reduced network costs [2].

The degrees of freedom that the aforementioned parameters offer, when combined also with the routing (space

dimension) options available, are vast. The effective definition of each of these parameters is not straightforward. For example, a given net bit rate can be achieved using a high-order modulation format and low baud rate or a lower-order modulation format and higher baud rate. The determination of the transmission parameters depends on, among other things, several effects appearing during the propagation of light in the fiber, such as noise, interference, nonlinear effects, equipment aging, and failures.

The adaptation of the network's parameters can be used to regulate the quality of transmission (QoT) of the lightpaths, but requires careful examination. More specifically, the adaptation of the modulation format or the forward correction code (FEC) is not straightforward because of bandwidth constraints. Indeed, if the modulation format is changed to a more robust one (e.g., from PM-16QAM to PM-QPSK), the bit rate is reduced. Therefore, a new subcarrier or a new lightpath should be established to keep the initial bit rate, assuming a constant baud rate. Similarly, in the case of an adaptation of a low data overhead FEC to a high one, the baud rate should be increased if the net bit rate is to be retained. In most cases, the adaptation of the parameters to more robust ones requires extra bandwidth to maintain a given net bit rate. In [3], a toolkit was developed that enabled the dynamic adaptation of the network to current conditions to always ensure acceptable QoT for all lightpaths.

In this paper, we propose spectrum sharing between lightpaths of different rate-guaranteed classes to recover the QoT in the case of a soft failure. We assume that the optical network supports traffic that is carried by different classes of connections. These classes are defined with respect to the guarantees provided for the rate and also the recovery time, and can thus be also viewed as different "survivability" classes. In particular, we focus on the case of two generic survivability classes: the guaranteed rate and the nonguaranteed rate, or the minimum guaranteed rate class. The QoT (thus BER) of both classes is ensured and, in case of a soft failure, it is recovered so that it is above a specific threshold. However, the rate of only the first survivability class is always fully guaranteed. For the second survivability class, only a minimum rate is guaranteed, and we can have several subclasses distinguished by the different minimum values. Note that the above classification is made on the basis of rate maintenance. It is complementary to the typical classification that is made with respect to the restoration time.

During normal network operation, all the lightpaths transmit at their full requested bit rate, usually using

Manuscript received February 21, 2018; revised May 28, 2018; accepted May 30, 2018; published June 25, 2018 (Doc. ID 323461).

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, and with the School of Electrical and Computer Engineering, National Technical University of Athens, Zografou, Athens 15773, Greece. E. Varvarigos is also with the Electrical and Computer Systems Engineering department of Monash University, Clayton VIC, Australia.

N. Sambo and P. Castoldi are with the Scuola Superiore Sant'Anna, Pisa 56127, Italy.

<https://doi.org/10.1364/JOCN.10.000653>

the most spectrally efficient modulation format. In the case of a soft failure, it is not necessary to maintain the rate of a rate-nonguaranteed priority connection [4,5]. Thus, part of its spectrum can be redistributed and used to maintain the rate of a guaranteed rate connection. In this way the modulation format of the guaranteed rate connections can be adapted to a more robust one to increase the QoT. At the same time the baud rate can be increased or a new lightpath or subcarrier can be established to maintain the original net rate. After the failure is repaired (if it is repairable), the network can return to its original transmission configurations.

The idea of spectrum sharing can effectively reduce the required network costs to satisfy the rate-guaranteed connections. The common approach would be to preserve the rate of high-priority, rate-guaranteed connections using additional spectrum. With the proposed method, we harvest the spectrum of lower priority rate-nonguaranteed connections and reduce the total spectrum used. Spectrum savings result in cost savings in terms of the expenditures for elastic network equipment [6]. The objective function of the optimization algorithm can control the degree to which the spectrum assigned to rate-nonguaranteed connections is redistributed to the rate-guaranteed connections. More specifically, a weight parameter W can regulate the emphasis (relative importance) the ILP algorithm gives on the total spectrum utilization and on the spectrum dropped by the rate-nonguaranteed connections experiencing the soft failure. In a sense, the parameter W controls the spectrum and thus the cost savings that can be achieved. For small values of W , guaranteed rate connections take full advantage of the spectrum of the nonguaranteed rate connections, decreasing network costs. For large values of W , the guaranteed rate connections require additional free resources to maintain their rate, thereby increasing costs.

In [7], the authors proposed an online heuristic RSA algorithm that can serve demands dynamically to maximize spectrum sharing between connections of different priority classes. In this paper, we extend this work and also present an ILP RSA algorithm to plan/allocate network resources to a set of connections of different classes. The proposed algorithms achieve significant spectrum savings, as we demonstrate in the simulations.

II. RELATED WORK

Optical networks typically use protection or restoration mechanisms [8] to recover from QoT degradations or a complete loss of signal. In protection, network resources are reserved for a lightpath in the case of a failure. No resource sharing is possible, which means protection is an expensive solution. In restoration, each lightpath affected by a failure dynamically searches for a different route, which may or may not be available. Both restoration and protection are necessary to recover from a complete loss of signal (e.g., fiber cut). In the case of QoT degradation, an adaptation of the transmission parameters of the affected lightpaths may be enough to compensate for the soft failure.

EONs offer several optimization dimensions that can be used to restore the QoT of a lightpath. The modulation format adaptation has been investigated in [9] as a means to recover the QoT of lightpaths with a degraded optical signal-to-noise ratio (OSNR). In [3], the authors considered the dynamic adaptation of the FEC, modulation format, and baud rate to recover from QoT degradation. The proposed solution examines the most appropriate combination of transmission parameters to improve the BER while preserving, if possible, the original rate. Spectrum sharing has been previously considered in optical networks. Along similar lines, the work in [10] considered connections with negatively correlated (in time) rates. These connections were placed spectrally next to each other to obtain statistical multiplexing gains (in terms of spectrum utilization). Spectrum sharing has also been considered for traffic fluctuations at EONs [11,12]. More specifically, the spectrum of the lightpaths was dynamically allocated to adapt to current traffic demands. Cooperative spectrum sharing along with spectrum defragmentation was considered in [13] to cope with the dynamicity of the traffic demands. The results offered significant benefits in terms of the blocking rate of new demands. The authors in [14] evaluated capacity sharing (in terms of spectrum) when considering base and peak rates for the demands. The authors also assumed that the base rate of the connections should be survivable in the case of any single link failure. Any spare spectrum can support the peak rate of certain connections. However, the transition from base to peak rate did not assume contiguous spectrum slots. The results indicated reduced blocking probabilities. In [7], the authors considered an in-operation route and spectrum allocation algorithm to assign the route and appropriate configurations for a new demand to achieve spectrum sharing between neighboring connections of different classes. The algorithm strived to route the high-class connections spectrally next to the low. In the case of a failure, the spectrum of the low-priority-class connections can be reallocated to the neighboring high-priority-class connections, as shown in Fig. 1. This approach can allow a high-priority connection to attain an acceptable QoT while simultaneously preserving the original rate. In this paper, we extend this work, and also present what we believe is a novel, optimal ILP RSA algorithm to plan the network. The objective is to maximize the spectrum sharing among the connections and minimize used spectrum. This approach can also result in significant capacity savings. The planning problem at hand is a combinatorial optimization problem. A simpler spectrum allocation problem would be to consider only normal network operation, in the absence of QoT degradations, where spectrum has to be allocated to minimize total spectrum utilization. If simplified, this issue would be equivalent to a graph coloring problem (the nodes represent the demands, the edges represent the shared links of the demands, and the color represents the assigned spectrum). Since graph coloring is a known NP-complete problem, our problem is also NP-complete since even its simplified version is NP-complete. We model this problem with an ILP formulation whose exact solution, when tractable, provides an optimal solution to the problem.




Fig. 1. Spectrum usage of a bronze ($p1$ from A to B) and a gold ($p2$ from A to C) connection before and after a failure.

III. EXPLOITING SPECTRUM SHARING

A. Network Scenario

We assume an EON [1] with configurable transceivers that can adapt the spectrum, modulation format, and baud rate used. The nodes consist of reconfigurable optical add/drop multiplexers (ROADMs) with flex-grid capabilities connected through uncompensated fiber links. Each fiber link consists of a number of fiber spans, each terminating at an erbium-doped fiber amplifier (EDFA) that compensates the span loss. We define a connection as the logical path between a source and a destination. A lightpath is an optical path that is established between two nodes, created by the allocation of the same spectrum on all links throughout the path. When needed (e.g., for long connections), regenerators are added, and each segment between a regenerator is considered a separate lightpath that may use a different spectrum.

In general, optical networks support a number of different traffic classes. For example, five classes are reported in [15]. The classification is done with respect to the guarantees provided for the rate and also the restoration time. Assuming a requested rate R , each class can have different requirements on the minimum guaranteed rate (R_{\min}). The higher classes are rate-guaranteed and require that $R_{\min} = R$ at all times. We can further distinguish these higher classes with respect to the recovery time. Moreover, there can be several nonguaranteed rate or minimum guaranteed rate classes, and the lowest class could be the one with R_{\min} equal to zero (best effort).

In the case of a failure, the rate of the rate-guaranteed connections can be maintained either by using additional spectrum resources with a higher cost, or by using spectrum provided by the minimum-rate-guaranteed connections. Thus, R_{\min} determines how much spectrum each low-priority connection can provide to higher-priority connections in the case of a soft failure. Moreover, the objective function of the optimization algorithm can control through a parameter W , and we will shortly introduce the degree to which the spectrum assigned to rate-nonguaranteed connections is redistributed to the rate-guaranteed connections. Therefore, the amount of spectrum, and thus the cost savings, that can be achieved depend on the parameters W and R_{\min} .

In the following section, we will consider only two classes of connections: i) *gold*, corresponding to rate-guaranteed

traffic and ii) *bronze*, corresponding to minimum-rate-guaranteed traffic. During normal network operation (in the absence of soft failures), all the connections use the combination of the highest modulation format and the lowest baud rate that ensure acceptable QoT. This yields efficient usage of the network resources. In the case of a soft failure, the QoT (e.g., BER) of a lightpath may become unacceptable. In that case, actions such as the adaptation to a more robust modulation format, the increase of FEC, or the addition of another channel can restore the QoT. The transmission rate can be maintained by increasing, if necessary, the respective utilized bandwidth. Thus, we propose to redistribute spectrum from the bronze to the gold connections whenever a failure occurs, and a gold connection requires additional spectrum. In this way, the rate of the gold connections is always maintained, while the rate of bronze connections is temporarily downgraded. If the failure is repairable, the original transmission configuration can be restored.

B. ILP Formulation

In this section, we present the proposed ILP planning algorithm. Note that the formulation we present considers only two classes of service. More classes with different minimum guaranteed rates R_{\min} can be supported by including the respective constraints. Since the algorithm is related to the network planning phase, we assume it will always preserve the rate of all the gold-class connections under a specific failure, inducing an SNR penalty of X dB. The objective is to minimize the total occupied spectrum, which can be partly achieved by routing bronze connections spectrally next to gold ones, so that the latter can borrow spectrum from the former in case of a failure. We assume that each demand is routed over its shortest path. For a connection C , based on the required rate, we find the transmission tuples (set of modulation format, baud rate) whose rate can satisfy the demand. From the feasible tuples, we select the one with the higher modulation format, which requires smaller baud rate/fewer spectrum slots. If none is available, then we place regenerators. We assume a soft failure of X dB per connection and calculate again the feasible tuples by considering proper QoT modeling (more details on QoT estimation will be provided in the next section).

Inputs:

- C : Set of all (bronze and gold) connections,
- B : Set of bronze connections,

- G : Set of gold connections,
- L : Set of network links,
- F : Total number of 12.5 GHz spectrum slots,
- X : Soft failure in dB that each connection should be resilient,
- Z_c : Required number of slots for connection c ,
- Z_c^X : Required number of slots for connection c under failure of X dB,
- M : Minimum guaranteed number of slots for bronze connections,
- Q_c^X : Set that contains the valid slot options for a given connection c in the case of soft failure/degradation of X dB. This set is equal to:
 - $\{Z_c^X\}$: for the guaranteed rate (gold) connections,
 - $\{M, \dots, Z_c^X\}$: for the nonguaranteed rate (bronze) connections,
- W : Weight to control the relative importance of different optimization objectives.

Variables:

- u_{cfs} : Boolean variable equal to 1 if connection c uses slots $(f, f + s - 1)$ where $s = Z_c$
- u_{cfs}^X : Boolean variable equal to 1 if connection c uses slots $(f, f + s - 1)$ under failure where $s \in Q_c^X$
- y : Integer variable ($\leq F$) that denotes the maximum slot number used in normal operation (without soft failure).
- y^X : Integer variable ($\leq F$) that denotes the maximum slot number used under soft failure of X dB.

Here, the objective is

$$\min \left(y - W \cdot \sum_{C \in B} \sum_{f \in [1, F]} \sum_{s \in \{Q_c^X\}} s \cdot u_{cfs}^X \right), \quad (1)$$

which is subject to:

- Spectrum slot utilization constraints,

$$y^X \leq y, \quad (2)$$

- Connection serving constraints under normal operation,

$$\forall c \in C, \quad s = Z_c: \sum_{f|f+s-1 \leq F} u_{cfs} = 1, \quad (3)$$

$$\forall c \in C: \sum_{s \neq Z_c} \sum_{f|f+s-1 \leq F} u_{cfs} = 0, \quad (4)$$

- Connection serving constraints under failure,

$$\forall c \in C: \sum_{s \in Q_c^X} \sum_{f|f+s-1 \leq F} u_{cfs}^X = 1, \quad (5)$$

$$\forall c \in C: \sum_{s \notin Q_c^X} \sum_{f|f+s-1 \leq F} u_{cfs}^X = 0, \quad (6)$$

- Non-overlapping slot assignment constraints under normal operation,

$$\forall l \in L, \quad \forall i \in [1, F]: \sum_{c \in C | l \in c, s=Z_c} \sum_{f|f \leq i \text{ and } f+s-1 \geq i} u_{cfs} \leq 1, \quad (7)$$

- Non-overlapping slot assignment constraints under failure,

$$\forall l \in L, \quad \forall i \in [1, F]: \sum_{c \in C | l \in c} \sum_{s \in Q_c^X} \sum_{f|f \leq i \text{ and } f+s-1 \geq i} u_{cfs}^X \leq 1, \quad (8)$$

- Maximum slot used constraints under normal operation,

$$\forall l \in L, \quad \forall f \in [1, F]:$$

$$y \geq \sum_{c \in C | l \in c, s=Z_c, f+s-1 \leq F} (f + s - 1) \cdot u_{cfs}, \quad (9)$$

- Maximum slot used constraints under failure,

$$\forall f \in [1, F], \quad s \in Q_c^X,$$

$$f + s - 1 \leq F: y^X \geq (f + s - 1) \cdot \sum_{c \in C | l \in c} u_{cfs}^X, \quad (10)$$

- Connection ordering constraints under normal operation and under failure,

$$\forall c \in C, \quad s = Z_c, \quad s' \in Q_c^X,$$

$$\forall f \in [1, F] | f + s - 1 \leq F, \quad : u_{cfs} = u_{cfs'}^X.$$

$$\forall f' \in [f - 2, f + 2] | f' + s' - 1 \leq F \quad (11)$$

The minimization objective of Eq. (1) includes two terms: i) the spectrum utilization under normal operation (denoted by y), and ii) the number of spectrum slots used for the bronze connections under failure conditions. The triple summation for the second term accounts for all possible combinations of the number of spectrum slots that all bronze connections can occupy for all valid spectrum slots. The objective of Eq. (1) includes a weight W that controls the emphasis (relative importance) that the algorithm gives on the total spectrum utilization as opposed to the slots used by bronze connections under failure. If W is sufficiently small, then the overall spectrum usage under normal operation and under failure will be minimized. This implies that the rate of the bronze connections will be minimized in the case of a failure to provide the gold connections with the required spectrum to continue their operation. Remember that bronze connections are much less valued than the gold ones, and all of their spectrum, beyond the minimum guaranteed one, can be reassigned to gold connections. If W is sufficiently large, then the spectrum allocation will ensure that all the bronze connections will also retain their original rate under failure. In that extreme scenario, bronze connections are valued equally with gold connections. In the latter case, to restore the gold connections, some extra free spectrum slots would be provisioned. Thus, W can also be viewed as a means to control

the degree of priority that gold connections have in borrowing spectrum slots from bronze ones. The minimum rate of the bronze connections is always guaranteed.

Constraint (2) limits the maximum number of slots used under failure to be less than or equal to the number of slots used under normal operation. Constraint (3) ensures that a connection c under normal operation is allocated resources, described by the starting frequency f and the spectrum slots s for its required number of slots ($s = Z_c$), while constraint (4) prohibits the use of inappropriate variables in any case. Constraints (5) and (6) similarly ensure the serving of the demands under failure. The only difference is that the required number of slots depends on whether the connection is gold or bronze. Constraints (7) and (8) ensure that the connections that share at least one link do not use the same slot under normal operation or failure, respectively. Constraints (9) and (10) define the maximum slots used under failure and normal operation, respectively, for every possible starting slot f . The value $f + s - 1$ represents the ending slot of a connection that requires s total number of slots. In (9) and (10), the values y and y^x have to be higher or equal to the ending slot of all connections. Constraint (11) limits the spectrum repositioning of the connections in the case of a failure to, at most, two slots of 12.5 GHz. This ensures that the order of the connections (in terms of spectrum) is not changed, assuming that the minimum number of slots that a connection can occupy is three. The fixed order of the connections makes it possible to use the push-pull technique [16] to modify the central frequency of each connection without traffic interruption. At the same time, the number of actions that have to be performed to reconfigure the connections is limited.

C. Dynamic Spectrum Sharing RSA

In the previous subsection, we considered a network planning approach where a set of demands were jointly served. In this subsection, we assume demands arrive dynamically and are served as they arrive. We present an RSA algorithm to leverage spectrum sharing between different connection classes. A centralized control plane [e.g., a software-defined network (SDN)] is assumed [17]. The centralized controller consults the network databases, including traffic engineering information and the state of lightpaths (e.g., traversed links and ports). This information is provided to the RSA algorithm.

Figure 2 shows the flow chart of the proposed spectrum sharing RSA (SS-RSA). Upon request for a demand with rate R from source s to destination d , the centralized controller computes a set of k -paths between s and d using the most efficient modulation format that ensures acceptable QoT. Then, spectrum sharing is evaluated for each of the aforementioned paths; spectrum sharing is enabled if spectrum satisfying the continuity constraint in all the links is available close to a lightpath belonging to a different class (as in Fig. 1). After performing the above calculations for all k -paths, paths not enabling spectrum sharing are removed. On this new set, the least congested path [18] is selected.




Fig. 2. SS-RSA flow chart.

Then, spectrum assignment is performed on the first available portion of spectrum that enables spectrum sharing. If we do not find an RSA solution that is capable of spectrum sharing, we choose the least congested path on the initial set of k paths with a first-fit spectrum assignment policy.

IV. PERFORMANCE SIMULATIONS

We evaluate the benefits of our proposed sharing and priority scheme by conducting a series of simulation experiments. We assumed an Italian backbone topology with 27 nodes and 43 bidirectional links, SSMF fiber with an attenuation coefficient 0.22 dB/km, dispersion parameter 16.7 ps/nm/km, and nonlinear coefficient 1.3 1/W/km. The span length was set at 80 km and the EDFA noise figure at 5 dB. We assumed 100 and 200 Gb/s connections. In normal operation, the 100G connections are served using PM-QPSK and three spectrum slots of 12.5 GHz and the 200G using PM-16QAM and the same number of slots. In the case of a failure, the 100G connections will (if needed) be served using PM-BPSK and six spectrum slots, while the 200G connections use PM-QPSK and again six spectrum slots.

A. ILP RSA Planning Algorithm

We consider two traffic cases of 10 and 15.8 Tb/s. The connections are derived from realistic traffic scenarios.

TABLE I
CONNECTION STATISTICS

		10 Tb/s	15.8 Tb/s
Number of Connections	100G	80	102
	200G	10	28
Length (km)	Mean	333	389
	Standard Deviation	264	263
Number of Hops	Mean	3.7	4.1
	Standard Deviation	2	1.9

Further connection statistics can be found in Table I. We assumed the GN model [19] to calculate the feasible tuples for the demands, accounting for worst case interference. The FEC threshold was set at 1.32×10^{-2} (which equals to -1.88 dB in a logarithmic scale and a Q^2 -factor of 6.93 dB). We set the maximum running time of the ILP algorithm at 1 h. Simulations are performed with a quad-core processor at 4 GHz. We used an IBM CPLEX with MATLAB to design and solve the ILP problem. We define the spectrum utilization as the maximum spectrum utilization among all the network links, assuming normal operation. As we mentioned in Section III.B, we assume that the ILP planning algorithm always maintains the rate of the gold connections for a given degradation. We consider a $X = 3$ dB degradation in the failure scenario. In Fig. 3 we present the Pareto efficiency graph obtained for the 10 Tb/s traffic, assuming 40% randomly assigned bronze connections. This graph is obtained by executing the ILP algorithm for various values of $W \in [0, 10^3]$. The spectrum utilization metrics obtained for these values are compared to the scenario where no spectrum sharing is considered. Therefore, the gold connections should occupy more spectral resources to be survivable. As a result, the total spectrum utilization of the network is increased. We plot the percentage savings in spectrum utilization versus the percentage of bronze slots dropped. The largest savings is observed for $W = 0$ and $W = 10^{-13}$. In these cases, the objective of the algorithm is to minimize the total spectrum without trying to maintain the rate of the bronze connections. As W increases, the objective gives more weight to the bronze connections, and when the W is sufficiently large (e.g., larger than the total number of slots in the network) then there is no possibility of spectrum sharing




Fig. 3. Pareto efficiency graph.

between gold and bronze connections. Note that there are cases where we obtain the same performance for different values of W (e.g., for $W = 10^{-3}$ and $W = 10^{-5}$ the spectrum savings were the same). Generally, there is no standard way to predict the exact spectrum savings for different values of W . Therefore, a Pareto efficiency graph should be constructed. We noticed that the values used in Fig. 3 yield similar results for all the traffic scenarios considered next.

In Fig. 4 we demonstrate the spectrum savings that can be achieved when spectrum sharing between gold and bronze connections is considered, as opposed to the case where no spectrum sharing is allowed. To calculate the spectrum savings, we first set $W = 10^{-2}$. In that case, the ILP algorithm maintains the rate of the bronze connections under soft failure conditions, unless a gold connection can borrow the spectrum of the neighboring bronze. Then we set $W = 10^3$ so that the algorithm will never consider spectrum sharing. The difference in the spectrum utilization between the two cases represents in Fig. 4 the spectrum savings achieved by our proposal.

Three different percentages of bronze connections are assumed to be 25%, 40%, and 55%. We randomly decide which connections are gold or bronze. Two traffic scenarios of 10 Tb/s (red bars) and 15.8 Tb/s (blue bars) are assumed. We notice that for the smallest bronze percentage, the spectrum savings are relatively small. Also, in the 15.8 Tb/s scenario, the savings are greater than in the 10 Tb/s scenario. This result is expected since a larger amount of connections offers greater opportunities for spectrum sharing. Also, more connections are degraded and require additional spectrum in the case of a failure. When the percentage of bronze connections is 40%, we notice that the percentage savings are increased. This is because the additional number of bronze connections offers more possibilities for spectrum sharing. At low load, more spectrum savings are achieved than at high load. This is because in the high load scenario the ILP algorithm cannot find an optimal solution within 1 h of the execution time. The gap between the reported solution and the objective value was 20%. The reason is that the additional number of bronze connections increases the complexity of the problem. The algorithm has to investigate more spectrum allocation choices

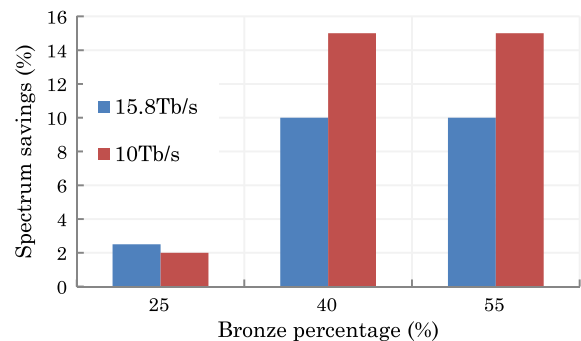


Fig. 4. Spectrum savings (%) of spectrum sharing case ($W = 10^{-2}$) over the no spectrum sharing case ($W = 10^3$) for different percentages of bronze connections.

between gold and bronze connections to find the one that minimizes the objective. After 2 h of running time the gap was decreased to approximately 10% and the savings for the 15.8 Tb/s were almost equal to the 10 Tb/s. When the percentage of the bronze connections is 55, we notice that there is no additional spectrum savings. For the 10 Tb/s, the ILP manages to find the optimal solution within the time limit. Therefore, the reason there is no additional spectrum savings is that the additional bronze connections do not offer additional sharing possibilities. This result also implies that the spectrum of the bronze connections has been preserved (since there is no spectrum sharing with gold class connections). For the 15.8 Tb/s, the respective optimality gap was 40% after 1 h of running time. After 2 h, the gap was 20% and the savings had increased by an additional 3%.

We also assumed the Deutsche Telecom topology with 12 nodes and 40 bidirectional links for the same 15.8 Tb/s total load and ratio of 100G and 200G. This time, the demands were randomly generated. The spectrum savings were 12%, 17%, and 20% for 25%, 40%, and 55% of the bronze connections, respectively. The ILP algorithm managed to find the optimal solution in the first two cases and an approximate solution in the third one. In the last case, the respective gap was 10%. The larger savings in this scenario can be attributed to the smaller network. This means that for the same amount of connections there are more possibilities for resource sharing. Also, the smaller network typically requires less time for an optimal solution to be found.

The ILP algorithm can provide a solution that results in significant spectrum savings for small and medium problem sizes. The problem at hand is NP-complete and there is no definite answer or prediction method to identify whether a problem instance is easily tractable or not. The problem depends on many parameters (size of the network, specific routing of the demands and shared links, magnitude of the soft failure, etc.). In cases where the ILP cannot provide a good solution, we could use the heuristic algorithm instead. We can order the set of demands and use the heuristic algorithm to serve them one-by-one. A metaheuristic could also be used to search among different orderings to obtain better solutions.

B. Dynamic SS-RSA

We assumed 100 Gb/s connection requests, which arrive dynamically according to a Poisson process with an average inter-arrival time of $1/\lambda$ time units and are uniformly distributed among node pairs and classes. The holding time of the connections followed a negative exponential distribution with mean $1/\mu = 500$ time units. X was set to 5 dB. The set of k paths consists of all the paths within one hop from the shortest hop path. Spectrum sharing is evaluated by randomly generating single-link soft failures and averaging the collected data. In all the examined cases, the adaptation requires one 12.5 GHz frequency slot more, assuming that the original frequency slots were over-dimensioned with respect to the actual signal bandwidth to limit filtering effects. QoT is estimated through the

signal-to-noise ratio model in [18]. This model uses an OSNR margin to account for nonconsidered effects (e.g., nonlinear effects). Since in the model of the previous section (GN model) we accounted for worst case interference, the two physical layer models provide similar results. The rate of the dynamically provisioned gold connections is guaranteed by exploiting spectrum sharing or by means of rerouting if spectrum sharing is not applicable. The rerouted connections could be also restored by allocating extra spectrum, as is done in the planning scenario. So, the percentage of gold connections that is restored through spectrum sharing, which is the main performance metric here, is directly connected to the spectrum utilization savings used as the main performance metric in the planning scenario.

We evaluate the benefits of our proposal in terms of the number of recovered gold lightpaths using spectrum sharing in the case of a soft failure. Such an evaluation is performed considering two RSA algorithms: the proposed SS-RSA and the (sharing unaware) RSA of [18] based on least-congested routing and first-fit spectrum assignment. Figure 5 shows the percentage of gold traffic recovered through modulation format adaptation, versus the provisioned traffic load. Provisioned load is defined as the product between the offered network load and the establishment probability $(1 - P_b)$: $\lambda/\mu \times (1 - P_b)$. Such load is considered instead of the offered load λ/μ for fairness. Indeed, the two adopted RSA strategies may affect provisioning blocking probability P_b , thus the number of lightpaths affected by a failure and, consequently, the recovery. The figure shows that, even with a simple RSA, up to 44% of gold traffic can be recovered by adapting transmission parameters, thus avoiding rerouting (path and frequency slot computation). If the proposed SS-RSA is used for routing, then up to 77% of gold traffic can be recovered through adaptation almost independently of the traffic load, but at the expense of bronze traffic. Such percentage decreases with the provisioned load because once the network is more loaded, it is more difficult to find lightpaths that facilitate spectrum sharing. In that

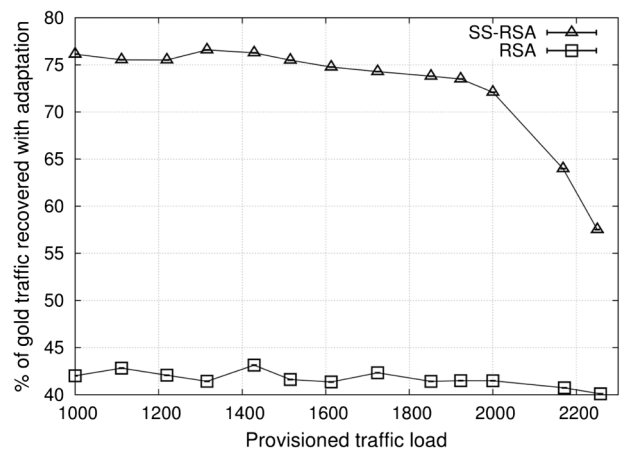


Fig. 5. Percentage of adapted gold traffic versus provisioned network load.




Fig. 6. Provisioning blocking probability versus offered network load.

case, simple least congested routing and first-fit spectrum assignment is applied (see right side of Fig. 5). In the case of the simple RSA, the percentage of recovered gold connections through adaptation is almost not affected by the provisioned load because spectrum sharing is not considered at all during provisioning, as opposed to the SS-RSA. To further compare SS-RSA and RSA, we also include Fig. 6, which shows the provisioning blocking probability versus the offered network load λ/μ . SS-RSA has a negligible impact on the provisioning blocking probability since the two RSAs experience almost the same performance during provisioning. The slightly higher provisioning blocking of SS-RSA is due to a more fragmented spectrum, since it prefers to enable spectrum sharing, instead of applying first-fit spectrum allocation. The latter tends to optimize spectrum utilization to achieve low blocking. In conclusion, the SS-RSA can be a very effective solution to recover from soft failures. It facilitates the transmission parameter adaptation of high-priority traffic without causing excessive spectrum fragmentation, and therefore without negatively affecting the provisioning blocking probability.

V. CONCLUSION

We considered spectrum sharing among connections belonging to classes with different rate guarantees to recover from soft failures. We presented an optimal ILP RSA algorithm to plan for the network to recover from soft failures. The ILP can control the amount of spectrum redistribution between high- and low-priority connections. It can therefore trade spectrum savings with the rate of lower-priority connections in case of a failure. We observed that using spectrum sharing (redistribution), we can obtain a spectrum savings of up to 20%. We also presented an online RSA algorithm that favors the routing of the high- and low-priority connections next to each other to increase the probability of spectrum sharing. Simulation results showed that up to 77% of high-priority connections can be recovered at the expense of low-priority ones. Future work

includes the application of the spectrum sharing concept to absorb traffic fluctuations.

ACKNOWLEDGMENT

I. Sartzetakis was supported by IKY Greek State PhD 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, cofunded by the European Social Fund and the Greek State. K. Christodoulopoulos was supported by a IKY post-doctoral scholarship cofunded by the European Social Fund and the Greek State. This work was also partially supported by the ORCHESTRA project, funded by EC (grant agreement 645360).

REFERENCES

- [1] V. Lopez and L. Velasco, Eds., *Elastic Optical Networks: Architectures, Technologies, and Control*, Springer, 2016.
- [2] O. Rival and A. Morea, "Cost-efficiency of mixed 10-40-100 Gb/s networks and elastic optical networks," in *Optical Fiber Communication Conf./National Fiber Optic Engineers Conf.*, 2011, paper OTuI4.
- [3] I. Sartzetakis, K. Christodoulopoulos, and E. Varvarigos, "Cross-layer adaptive elastic optical networks," *J. Opt. Commun. Netw.*, vol. 10, no. 2, pp. A154–A164, 2018.
- [4] 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, 2016.
- [5] ORCHESTRA deliverable D2.3, 2016 [Online]. Available: <http://www.orchestraproject.eu/index.php/downloads/public-documents>.
- [6] E. Palkopoulou, M. Angelou, D. Klonidis, K. Christodoulopoulos, A. Klekamp, F. Buchali, E. Varvarigos, and I. Tomkos, "Quantifying spectrum, cost, and energy efficiency in fixed-grid and flex-grid networks [Invited]," *J. Opt. Commun. Netw.*, vol. 4, no. 11, pp. B42–B51, 2012.
- [7] N. Sambo, K. Christodoulopoulos, P. Castoldi, and E. Varvarigos, "Spectrum sharing for elastic transmission parameter adaptation," in *42nd European Conf. Optical Communication (ECOC)*, Dusseldorf, Germany, 2016.
- [8] D. Zhou and S. Subramaniam, "Survivability in optical networks," *IEEE Netw.*, vol. 14, no. 6, pp. 16–23, Nov./Dec. 2000.
- [9] 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 Exhibition on Optical Communication*, Geneva, Switzerland, Sept. 2011.
- [10] G. Shen, Q. Yang, S. You, and W. Shao, "Maximizing time-dependent spectrum sharing between neighbouring channels in CO-OFDM optical networks," in *13th Int. Conf. Transparent Optical Networks*, Stockholm, Sweden, 2011.
- [11] 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.
- [12] 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.
- [13] I. Stiakogiannakis, E. Palkopoulou, D. Klonidis, O. Gerstel, and I. Tomkos, “Dynamic cooperative spectrum sharing and defragmentation for elastic optical networks,” *J. Opt. Commun. Netw.*, vol. 6, no. 3, pp. 259–269, Mar. 2014.
- [14] F. Khandaker, J. P. Jue, X. Wang, Q. Zhang, H. Cankaya, I. Kim, and T. Ikeuchi, “Sharing of primary and back-up capacity in survivable elastic optical networks,” in *Advanced Photonics Congress*, New Orleans, Louisiana, 2017, paper NeTu1B.6.
- [15] R. Ramaswami, K. Sivarajan, and G. Sasaki, *Optical Networks: A Practical Perspective*, Morgan Kaufmann, 2009.
- [16] 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. 2012.
- [17] D. King and A. Farrel, “A PCE-based architecture for application-based network operations,” IETF RFC 7491, Mar. 2015.
- [18] N. Sambo, F. Cugini, G. Bottari, G. Bruno, P. Iovanna, and P. Castoldi, “Lightpath provisioning in wavelength switched optical networks with flexible grid,” *37th European Conf. and Exhibition on Optical Communication*, Geneva, Switzerland, 2011.
- [19] 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.