Optimal Route, Spectrum, and Modulation Level Assignment in Split-Spectrum-Enabled Dynamic Elastic Optical Networks

Albert Pagès, Jordi Perelló, Salvatore Spadaro, and Jaume Comellas

***Abstract*—The spectrum fragmentation effect in elastic optical networks is one of their main limitations. Multiple techniques have been proposed to address this problem, with the split spectrum approach (SSA) being a very interesting candidate among them. This technique is based on splitting a demand into smaller sub-demands when a blocking situation arises. In split-spectrum-enabled networks, the route, spectrum, and modulation level assignment (RSMLA) problem that appears in elastic optical networks is further complicated due to the signal splitting operation. In this paper we present novel mechanisms to optimally attack this problem; various possible implementations of the SSA are also discussed. We highlight the benefits of the proposed mechanisms through illustrative results and compare the various implementation solutions in terms of average network cost.**

***Index Terms*—Elastic optical networks; Optimization; RSMLA; Split spectrum.**

I. Introduction

Elastic optical networks [1,2] have been proposed to overcome the inefficiencies of fixed-grid wavelength division multiplexing (WDM) optical networks. In elastic optical networks, the available spectrum of a fiber link is discretized into a set of frequency slots (FSs) with substantially smaller spectral width than a traditional wavelength (i.e., FS widths are 6.25 or 12.5 GHz, against the 50 GHz ITU-T WDM grid). Instead of a full wavelength, connections are allocated over a number of contiguous FSs tightly adjusted to their bandwidth needs. Moreover, it is forecast that elastic optical networks will also allow ultrahigh bit-rate super channels at 400 Gb/s and 1 Tb/s, which are not currently supported in a 50 GHz fixed grid.

Although elastic optical networks present large benefits when compared to traditional WDM optical networks, they also pose important challenges, the main one being the so-called spectrum fragmentation effect [3], which can greatly affect the performance [e.g., the bit-rate blocking ratio (BBR)] of such networks. This effect refers to the fact that, due to the random arrival and departure process of the demands over the network, the available spectral resources (the FSs) become highly fragmented, forming spectral gaps (i.e., bunches of free contiguous FSs) of heterogeneous sizes scattered along the spectrum, thus making more complex the allocation of future demands. This happens because, in elastic optical networks, demands request a contiguous spectrum portion tailored to their own needs, instead of a full wavelength. Hence, their chances to find a contiguous spectrum portion decrease, potentially leading to blocking situations.

Looking at the literature, there are many works that try to mitigate this effect by means of spectrum defragmentation techniques, performing them preventively when a demand is allocated [3] or reactively when blocking occurs [4,5]. The operation behind most of these works focuses on the rearrangement of the active connections in order to compact the occupied spectrum and release as many contiguous FSs as possible. In this regard, we can differentiate between techniques that reroute the existing connections (disruptive techniques) and techniques that rearrange only the utilized spectrum without changing connection routes (hitless techniques). In the latter case, the established connections are not disrupted when performing the spectrum defragmentation.

As an alternative way to deal with the spectrum fragmentation effect and, hence, increase the performance of elastic optical networks, the split spectrum approach (SSA) has been recently proposed [6]. In SS-enabled elastic optical networks, demands that face a blocking situation may undergo a splitting process in order to fit in the available FSs. Nevertheless, this splitting process complicates the assignment of resources to the incoming demands.

In this paper, we present novel mechanisms to address the problem of optimally assigning resources for incoming demands in a dynamic SS-enabled elastic optical network. Section II reviews the related work on the SSA problem. Section III describes the assumed scenario and the problem under consideration. Section IV presents various mechanisms to tackle the presented problem, whose performance

Manuscript received June 11, 2013; revised November 19, 2013; accepted December 4, 2013; published January 15, 2014 (Doc. ID 192095).

The authors are with the Advanced Broadband Communications Center (CCABA), Universitat Politècnica de Catalunya (UPC), Barcelona, Spain (e-mail: albertpages@tsc.upc.edu).

http://dx.doi.org/10.1364/JOCN.6.000114

is thoroughly evaluated in Section V. Finally, Section VI draws up the main conclusions.

II. Related Work

The SSA [6–13] is arising as a promising way to further increase the efficiency of elastic optical networks, mitigating the effect of spectrum fragmentation and lowering the BBR. The main rationale behind the SSA is founded on on dividing a high bit-rate demand into a set of lower bit-rate sub-demands if a blocking situation arises, making the resulting sub-demands more easily allocable in the available FSs.

This approach can basically be triggered in two kinds of situations: 1) when there are not enough contiguous FSs to serve the demand due to the spectrum fragmentation effect or 2) when the modulation format used to transmit the desired bit rate becomes infeasible due to transmission impairments that limit the transmission reach. In both cases, the demand can be split to better take advantage of the spectral gaps. Moreover, more robust modulation formats can be used, since the resulting sub-demands are transmitted at lower bit rates.

Once the original demand is split, the resulting sub-demands, hereafter referred to as parts, have to be routed and spectral resources assigned to them, signaled, and managed. Concerning the routing process, the different parts can be either routed over the same physical path (single-path approach) or over multiple physical paths (multi-path approach). The first one provides the advantage of simplifying the signaling and routing process, as only one route has to be signaled and established. Additionally, the differential delay among the parts is kept very low, making the hardware needed at the destination node (electronic buffers) for the reconstruction of the original demand less complex. In contrast, in the multi-path approach, the different parts can be routed over different physical paths. Although this second approach may achieve better performance in terms of BBR, the routing process is more complex since a routing instance must be triggered for each part. Furthermore, the size of the electronic buffers at the destination node must be large enough to mitigate the propagation delay differences between the parts.

Concerning the enabling technologies that make the SSA possible, there are basically two implementations regarding transponder architecture and utilization, namely, the bandwidth variable transponder (BV-TSP)-based and the multi-flow transponder (MF-TSP)-based implementation. In the BV-TSP-based implementation [12], once the demand has gone through the splitting process, the resulting parts are transmitted using independent BV-TSPs. That is, as many BV-TSPs as parts into which the demand has been split are employed, as shown in Fig. 1. To do so, the mechanism takes advantage of the BV-TSPs at the nodes that are not in use when splitting the demand. In this implementation the number of parts into which a demand is split takes a capital role in the overall performance of SSA, since too many parts can rapidly exhaust available

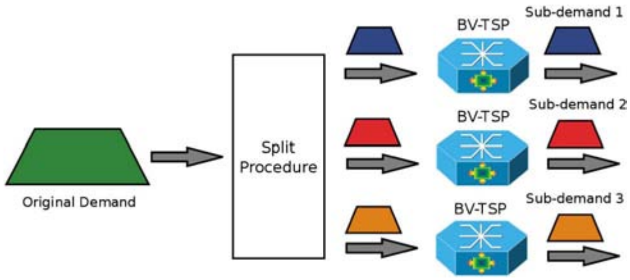

Fig. 1. Example of BV-TSP-based SSA implementation.

TSPs at nodes and lead to blocking due to the lack of them. For this reason, any allocation mechanism employing such an implementation has to carefully take into account this issue. On the other hand, this implementation does not impose additional hardware complexity with the sole SSA purpose, as it basically relies on the hardware already deployed in the network, keeping the capital expenditures (CAPEX) within reasonable limits.

As for the MF-TSP-based implementation, as its name suggests, it employs MF-TSPs to transmit the split demand [6,7]. A MF-TSP is capable of transmitting and receiving multiple elastic optical channels (flows), operating at independent bit rates, that can be routed independently. Therefore, SSA can be realized employing a unique MF-TSPs to transmit all the parts resulting from the split procedure. A simplified practical architecture of a MF-TSP, as proposed in [6], is depicted in Fig. 2.

Essentially, it would consist of an array of tunable lasers and blocks of modulators to support different types of modulation formats and bit rates, terminated with an optical combiner at the transmitter side, while an optical splitter followed by an array of detectors is placed at the receiver side. The advantage that the MF-TSP implementation presents over BV-TSP is that it only employs a TSP per demand, being split or not, so the potential lack of TSPs is not as pronounced in this scenario. On the other hand, the complexity and hardware requirements of this type of TSP are higher, which may lead to higher overall network CAPEX.

For all of these, the SSA has triggered the interest of standardization organizations [13] and the research community [6,8–12]. For example, the Internet Engineering Task Force (IETF) is currently working on extending control plane architectures and protocols of fixed-grid optical networks to be able to also control and manage elastic optical networks. In this framework, the SSA is also

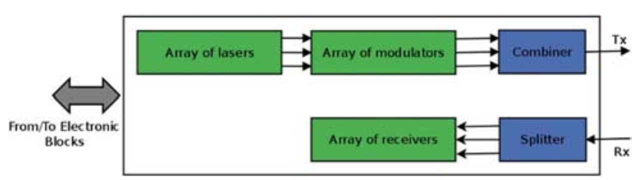

Fig. 2. Architecture of a MF-TSP.

being considered as one of the choices for establishing a connection in flex-grid networks. However, the standardization proposals in this regard are at their very early stage (e.g., [13]).

On the other hand, much effort has already been made to evaluate SSA as a solution to increase the performance of elastic optical networks. For instance, the authors of [12] proposed a simple algorithm for splitting in a SS-enabled elastic optical network demands that face blocking situations due to the spectrum fragmentation effect. The authors demonstrated that substantial improvements in terms of blocking can be achieved. Nonetheless, the authors considered only the single-path routing approach and did not tackle the choice of the best modulation format according to the demands' needs.

In addition, the authors of [10,11] investigated the multi-path routing approach in SS-enabled optical networks, proposing multiple algorithms to intelligently split the demands, and showed that the multi-path routing approach, or even a hybrid single-/multi-path approach, can lead to further gains in terms of blocking when compared to the single-path approach. Although the authors of [11] introduced the selection of the modulation format according to the specific bit rate of the demands and the transmission reach of the modulation formats, they did not tackle the optimal selection of the route, the spectrum portion, or the modulation format to serve a demand.

Similar work has been done by the authors of [6,9], in which multiple algorithms are proposed to choose the route, spectrum, and modulation format of demands in a SS environment. Additionally, the authors of [6] introduced the concept of MF-TSP as a way to realize the SSA in elastic optical networks. However, none of these works is devoted to finding the optimal solution to the route, spectrum, and modulation level assignment (RSMLA) problem in SS-enabled networks.

The authors of [8] proposed an integer linear programming (ILP)-based mechanism to tackle the optimal route and spectrum assignment under multi-path SSA considering the differential delay issue. They showed that correct dimensioning of the electronic buffers at the receiver end can compensate such delay and provide better performance in terms of blocking than the single-path approach. Despite providing a mechanism capable of finding the optimal resource assignment, the authors did not consider the selection of the modulation format, restricting the assignment to selection of only paths and FSs. Moreover, to the best of our knowledge, there is no work to date proposing an optimal RSMLA mechanism for SS-enabled optical networks that considers the implications of working under different transponder architectures.

For all these reasons, in this work we provide mechanisms to optimally assign resources (paths, FSs, and modulation formats) to a demand in a SS-enabled optical network for a dynamic scenario, considering both single- and multi-path routing approaches, as well as BV-TSPs or MF-TSPs. The following section elaborates on the scenario that we are considering and states the problem that we are targeting.

III. Problem Statement

Let us consider a transparent SS-enabled elastic network where nodes are not equipped with either regeneration or spectrum conversion capabilities. In a dynamic traffic scenario, demands randomly arrive at the network with heterogeneous holding times (HTs) and bit-rate needs. Given the current state of the network resources, the RSMLA problem [14] aims at finding the most appropriate physical route, spectrum portion, and modulation format to fully satisfy each demand's needs upon arrival, at minimum cost for the operator. Furthermore, being a dynamic network scenario, additional considerations that favor the establishment of future demands, such as load balancing or minimum spectrum fragmentation, are also desirable.

Moreover, let us assume that optical nodes can be equipped with BV-TSPs or MF-TSPs. Depending on the case, the RSMLA optimization objectives can differ. With BV-TSPs, as every part into which a demand is split employs an independent TSP, RSMLA should encourage the minimization of parts in order to avoid the exhaustion of BV-TSPs at nodes. With MF-TSPs, as described in Section II, every transponder is equipped with a set of tunable lasers to produce the independent flows that each supports one part of the demand. In such a case, RSMLA is constrained to avoid splitting a single demand in more flows (parts) than equipped lasers in the MF-TSPs.

To also introduce the procedure to translate the bit rate of a demand into a FS's need, we define R_d and R_m as the bit rates of the demand and that of the candidate modulation format used to allocate the demand, respectively, so that $R_d \leq R_m$. In addition, B_m denotes the spectral width [in gigahertz (GHz)] of the candidate modulation format and F_w the spectral width (in GHz) of a single FS. Ideally, the number of FSs needed to serve a demand would be the ceiling of B_m divided by F_w. However, bandwidth variable wavelength cross connects (BV-WXCs) require guard bands between demands to properly operate [15]. With G as the guard band size (in GHz), the number of FSs needed to allocate a demand, denoted as S, can be calculated as

$$S = \lceil (B_m + G)/F_w \rceil. \tag{1}$$

Guard band requirements imposed by BV-WXCs can increase the initial value of S, which may be more noticeable in a SS-enabled network. Indeed, the number of FSs needed to allocate a demand split into H parts is

$$S = \sum_{i=1}^{H} \lceil (B_{m_i} + G)/F_w \rceil, \tag{2}$$

where B_{m_i} is the spectral width of the modulation format employed in the ith part, with R_{m_i} being its bit rate, so that $R_d \leq \sum_{i=1}^{H} R_{m_i}$.

In a SS-enabled network, the RSMLA problem still applies, but in a more complex way, due to the splitting procedure that some demands may undergo in order to be successfully allocated. Indeed, the RSMLA has to optimally decide into how many parts a demand is split, jointly with the optimal physical route, spectrum portion, and modulation format for each of the resulting parts. This adds extra complexity to the problem.

With such conditions and scenario, the following section states the problem that we are targeting.

A. Problem Statement

In this section, we state the on-line problem of optimal RSMLA in a SS-enabled elastic optical network. The objective of the optimization problem is to find the most suitable resources (paths, spectrum portion, and modulation formats) to allocate the incoming demands (in one or multiple parts) in the network, according to the current resources state, so that the overall BBR is minimized. To encourage the minimization of BBR, two main sub-objectives are pursued: 1) minimize the number of resources to serve the incoming demand, and 2) minimize the spectrum fragmentation in the network. The specific actions to pursue these two sub-objectives will be discussed later. With all this in mind, the considered optimization problem can be formally stated as

Given:

1) a transparent elastic optical network represented by the directed graph $G_n = (N, E)$, where N denotes the set of network nodes and $E = \{(i,j),(j,i): i,j \in N, i \neq j\}$ the set of physical links;
2) an ordered set of FSs per physical link denoted as F;
3) an incoming demand d to be served between a source and a destination node (s_d and t_d), characterized by its requested bit rate (R_d);
4) a currently available number of TSPs in both source and destination nodes, denoted by TSP_s and TSP_t, respectively;
5) a set of modulation formats M supported by TSPs at nodes;
6) a maximum number of elastic flows that a MF-TSP is capable to produce, denoted as $\mathrm{FL}_{\max}$.

Find the end-to-end paths, FSs, and modulation formats for the potential parts of demand d subject to the following *constraints*:

1) *traffic constraint:* the combination of modulation formats employed by TSPs used to serve demand d must produce, at least, a bit rate equal to R_d;
2) *spectrum continuity:* as transparent lightpaths are established, the subset of allocated FSs should be the same for each link on every particular path selected to route demand d. Note that such constraint is applied to every path independently;
3) *spectrum contiguity:* for every particular part into which the demand is split, FSs should be allocated each next to the other in the spectrum domain;
4) *spectrum clashing:* a particular FS in a physical link can be allocated to one part of one demand at most;
5) *transponder utilization:* the number of parts into which d is divided should not exceed the number of available TSPs at the nodes (BV-TSP implementation) or the maximum number of allowed flows (MF-TSP implementation)

with the *objective* to minimize the use of network resources and the spectrum fragmentation.

To achieve the first sub-objective, multiple action lines are defined. First, the number of parts into which the demand is split should be minimized, as dividing the demands into too many parts may exhaust rapidly the available TSPs in the nodes (BV-TSP case). Moreover, as each of the resulting parts must be surrounded by guard bands, using more parts to serve the demand implies occupying more spectral resources in the network.

Second, as a demand can be served by employing multiple paths and modulation formats, it may happen that different feasible solutions employing the same number of parts may result in different numbers of required FSs. Therefore, the explicit minimization of the total of FSs employed by the demand should also be encouraged.

Third, it is clear that the sum of the bit rates of all parts used to serve a demand has to be, at least, equal to the original bit rate of the demand. However, depending on the demand bit rate and the modulation formats supported by the TSPs, it may be infeasible to allocate exactly the requested bit rate due to the limitations imposed by the granularity of the TSPs. With this in mind, the disparity between these two values, stated as $\sum_{i=1}^{H} R_{m_i} - R_d$, should be minimized. The main reason for doing so is to minimize the amount of spectrum finally used by the demand, as higher bit-rate modulation formats usually require larger spectral widths.

As for the second sub-objective, its main goal is to facilitate the allocation of future incoming demands, as spectrum fragmentation may lead to undesired BBR due to demands not finding enough contiguous spectral resources, even if the SSA is applied.

In the following, we provide a mixed ILP (MILP)-based mechanism to attack the stated optimization problem, as well as a purely heuristic mechanism for the cases when the scalability of the MILP-based mechanism may be compromised.

IV. Proposed Mechanisms

A. MILP-Based Algorithm

In this section, we propose a novel MILP-based mechanism called split-spectrum-enabled RSMLA (SSRSMLA) to optimally address the problem presented during the previous section. The presented mechanism is valid for both multi-path and single-path approaches as well as BV-TSP-based and MF-TSP-based implementations. Note, however, that we do not tackle the differential delay issue among parts, which is left out of the scope of this paper.

Before depicting the details of the SSRSMLA mechanism, let us discuss how we address both the spectrum continuity and contiguity constraints. In this work, we consider that a set of candidate paths for demand d is given. Let us define P as the set of possible paths over G_n and $P_d \subseteq P$ as the set of candidate paths for demand d, with l_p and h_p denoting the physical distance and number of hops of path p, respectively. Additionally, we define $P_e \subseteq P_d$ as the set of candidate paths for d that traverse physical link $e \in E$. We also define $F_p^u \subseteq F$ as the set of end-to-end FSs in path p that are already occupied. Moreover, we define $M_p \subseteq M$ as the set of feasible modulation formats for demand d in path $p \in P_d$; that is, the modulation formats that support a transmission reach equal or greater than the physical length of the candidate path. Every modulation format m belonging to M_p is characterized by its transmission reach TR_m, bit rate R_m, spectral bandwidth B_m, and associated number of FSs S_m calculated using Eq. (1).

The definition of such sets will allow us to satisfy the spectrum continuity constraint more easily, as spectral resources will be reserved explicitly in end-to-end paths; hence, they will remain the same for each link forming the selected paths.

Reviewing the literature that tackles the optimal solution of the RSMLA problem with ILP-based formulations [14,16–18], it can be appreciated how the consideration of the spectrum contiguity constraint greatly increases the complexity of the formulations. Note that, in SS-enabled elastic optical networks, such a constraint still has to be ensured for each one of the parts into which the demand is split. To this end, we model the spectrum contiguity constraint following the same approach proposed by the authors in [19]. There, the authors use precomputed sets of contiguous FSs, named channels. Each channel consists of a subset of adjacent FSs that provide just enough spectrum to support the spectral needs of a particular demand. This approach makes ILP-based solutions for the RSMLA problem solvable in practical times [19].

With all this in mind, we define C as the set of precomputed candidate channels for demand d, $C_p \subseteq C$ as the set of channels for candidate path $p \in P_d$, $C_{p,m} \subseteq C_p$ as the set of channels employing modulation format $m \in M_p$ in path p, and $C_{p,f} \subseteq C_p$ as the set of channels that are mapped in FS f in path p. Note that all channels belonging to $C_{p,m}$ will have an associated number of FSs equal to S_m as they are the candidate sets of contiguous FSs that fit the spectral needs of modulation format $m \in M_p$. Specifically, every channel $c \in C$ has its own value of associated bit rate, number of FSs, and first FS of the channel, denoted as S_c, R_c, and f_c, respectively.

Finally, we also define $H_{\max}$ and $L_{\max}$ as the maximum number of parts and paths that a demand is allowed to use, respectively. These parameters will be used to avoid excessive splitting of a demand or employing too many paths, which may result in spectrum allocation inefficiency.

Additionally, let us describe the specific actions that the mechanism takes to minimize the spectrum fragmentation in the network. As an overall network defragmentation would be too costly in terms of time and complexity, we minimize the number of spectral gaps present in the candidate paths for demand d. By minimizing the spectral gaps, spectral resources are less scattered along the candidate paths, helping future demands in finding large portions of contiguous FSs.

With all of this said, Fig. 3 depicts the pseudo-code of the presented mechanism. Essentially, given an incoming demand, a MILP formulation is executed in order to find the optimal RSMLA depending on the network status and the demand needs. More specifically, the mechanism is structured in three phases. The first phase is devoted to obtaining the sets that will be used later in the MILP formulation, namely, the candidate path and channel sets. For the candidate path set, we employ a K link-distinct shortest path strategy, where K is an input parameter. Note that limiting the set of candidate paths is necessary since we are facing an on-line optimization problem, and, hence, a good balance between optimality of the solution and total execution time has to be reached. The metric of the paths is the physical distance, as it becomes a very relevant factor in order to maximize the possibilities of

```
Inputs: d,G_n,F,M,K,H_max,L_max,FL_max,α,β,γ,δ; Output: Sol

Phase 1: Pre-processing
if TSP_s = 0 or TSP_t = 0 then
    Demand blocked
P_d ← K shortest paths from s_d to t_d using l_p as the metric
for i = 1 to K do
    M_i ← set of modulations belonging to M with TR_m ≥ l_i
    if M_i = ∅ then
        Remove p_i from P_d
if P_d = ∅ then
    Demand blocked
C ← ∅
for i = 1 to |P_d| do
    C_i ← ∅
    for j = 1 to |M_i| do
        C_{i,j} ← ∅
        c ← ∅
        for k = 1 to |F| do
            if f_k ∈ F in path p_i is free then
                c ← f_k ∪ c
            else
                c ← ∅
            if |c| = S_j then
                C_{i,j} ← c ∪ C_{i,j}
                c ← ∅
        C_i ← C_{i,j} ∪ C_i
    if C_i = ∅ then
        Remove p_i from P_d
    else
        C ← C_i ∪ C

Phase 2: MILP solving
if P_d = ∅ then
    Demand blocked
else
    Sol ← output from MILP(d, P_d, C, H_max, L_max, FL_max,α, β,
    γ, δ)

Phase 3: Solution evaluation
if Sol = ∅ then
    Demand blocked
else
    Return Sol
    Update network resources according to Sol
    Demand allocated
```

Fig. 3. SSRSMLA mechanism pseudo-code.

successfully allocating a demand. That is, shorter paths allow more modulation formats than longer ones, hence supporting more candidate channels to serve a demand.

As for the channel calculation, first we determine for each of the previously calculated paths which are the modulation formats supported through that path, that is, the modulation formats so that $\mathrm{TR}_m \geq l_p$. If no path has at least one feasible candidate modulation, the demand is automatically blocked, as it will be impossible to reach the destination in a transparent way (we remind the reader that we assume an all-optical scenario, where intermediate nodes do not perform any kind of regeneration).

For the candidate paths that have at least one feasible candidate modulation, the mechanism proceeds to calculate the candidate channels, given the modulation characteristics and FS availability over the path. The operation to determine the candidate channels is very simple, it consists of just finding spectral gaps of width equal to the spectral width of the modulation format to be associated to the channel.

Once the candidate path and channel sets are filled, the mechanism proceeds with the following step. In this phase, the mechanism executes a MILP formulation, whose details will be later explained, that will find the optimal RSMLA for the demand d given the aforementioned candidate sets and the state of the network. The parameters α, β, γ, and δ are weighting factors. Before executing the MILP, though, the mechanism checks if at least one of the candidate paths has a nonempty set of candidate channels; otherwise the demand is directly blocked as no combination of modulation format and contiguous spectral blocks that fulfill the requirements of the demand is found. By performing this check, the mechanism avoids executing unnecessary instances of the MILP formulation.

Finally, the mechanism checks if the solution of the MILP formulation is empty or not. If not empty, it means that there exists an optimal solution that fulfills the requirements of the demand. This being the case, the solution is adopted as the RSMLA for the demand and network resources are updated accordingly. We proceed now to detail the MILP formulation, whose decision variables are introduced below.

$x_{p,c}$: binary; 1 if path $p \in P_d$ and channel $c \in C_p$ are used to allocate the demand, 0 otherwise.

y_p: binary; 1 if path p is utilized to allocate the demand, 0 otherwise.

$z_{p,f}$: binary; 1 if FS f in path p is occupied, 0 otherwise.

P_p: integer; number of parts routed over path p.

T_p: integer; number of spectral gaps in path p.

$w_{p,f,f+1}$: real; auxiliary variables in the range [0,1].

A: real; difference between the requested bit rate and the allocated bit rate.

At this point, let us clarify the role of variables T_p and $w_{p,f,f+1}$. In order to discriminate a spectral gap in the network, we can take the state of the FSs in the candidate path. Specifically, a spectral gap can be discriminated by a free-to-allocated and another allocated-to-free FS transition, that is, a 0–1 and a 1–0 transition in variables $z_{p,f}$. Thus, we subtract $z_{p,i+1}$ from $z_{p,i}$. If the result is different than 0, it means that a transition is detected. Therefore,

$$T_p = 0.5 \sum_{i=1}^{|F|-1} |z_{p,i} - z_{p,i+1}|. \tag{3}$$

However, Eq. (3) does not detect transitions correctly before the initial and after the last FS of the spectrum. Then, we add $z_{p,1}$ and $z_{p,|F|}$ in Eq. (3) to properly compute T_p:

$$T_p = 0.5\left(z_{p,1} + z_{p,|F|} + \sum_{i=1}^{|F|-1} |z_{p,i} - z_{p,i+1}|\right). \tag{4}$$

The term within the sum in Eq. (4) can only take values equal to −1, 0, or 1. Then, we can substitute the absolute value function by the square function, which results in

$$T_p = 0.5\left(z_{p,1} + z_{p,|F|} + \sum_{i=1}^{|F|-1} (z_{p,i}^2 - 2z_{p,i}z_{p,i+1} + z_{p,i+1}^2)\right). \tag{5}$$

Thanks to the binary nature of the variables, quadratic terms can be simplified. Thus, rearranging Eq. (5),

$$T_p = \sum_{i=1}^{|F|} z_{p,i} - \sum_{i=1}^{|F|-1} z_{p,i} z_{p,i+1}. \tag{6}$$

Using Eq. (6) we obtain the number of spectral gaps in path p. However, a product of decision variables is encountered, making the expression nonlinear. To linearize it, auxiliary variables $w_{p,f,f+1}$ are introduced, together with some additional constraints in the proposed MILP formulation, which is detailed in what follows:

$$\min \alpha \frac{1}{H_{\max}|P_d|} \sum_{p \in P_d} h_p P_p + \beta \frac{2}{|F||P_d|} \sum_{p \in P_d} h_p T_p + \gamma \frac{1}{|C|} \sum_{p \in P_d} h_p \sum_{c \in C_p} S_c (1 + \epsilon f_c) x_{p,c} + \delta \frac{A}{\max\{R_m\}}, \tag{7}$$

subject to:

$$R_d \leq \sum_{p \in P_d} \sum_{c \in C_p} R_c x_{p,c}, \tag{8}$$

$$A = \sum_{p \in P_d} \sum_{c \in C_p} R_c x_{p,c} - R_d, \tag{9}$$

$$y_p \geq x_{p,c}, \quad \forall p \in P_d, c \in C_p, \sum_{p \in P_d} y_p \leq L_{\max}, \tag{10}$$

$$\sum_{p \in P_d} P_p \leq H_{\max}, \mathrm{TSP}_s/\phi, \mathrm{TSP}_t/\phi, \mathrm{FL}_{\max}/(1-\phi), \tag{11}$$

$$\sum_{p\in P_e}\sum_{c\in C_f} x_{p,c} \le 1, \quad \forall\, e\in E, f\in F, \tag{12}$$

$$P_p = \sum_{c\in C_p} x_{p,c}, \quad \forall\, p\in P_d, \tag{13}$$

$$T_p = \sum_{i=1}^{|F|} z_{p,i} - \sum_{i=1}^{|F|-1} w_{p,i,i+1}, \quad \forall\, p\in P_d, \tag{14}$$

$$\begin{aligned} z_{p,f} &= 1, \quad \forall\, p\in P_d,\ f\in F_p^u, \\ z_{p,f} &= \sum_{c\in C_f} x_{p,c}, \qquad \forall\, p\in P_d,\ f\in F\backslash F_p^u, \end{aligned} \tag{15}$$

$$\begin{aligned} &w_{p,i,i+1} \le z_{p,i}, z_{p,i+1}, \\ &w_{p,i,i+1} \ge z_{p,i} + z_{p,i+1} - 1, \quad \forall\, p\in P_d, i = 1, \ldots, |F|-1. \end{aligned} \tag{16}$$

Objective function (7) has multiple optimization goals: 1) minimize the number of parts into which the demand is split, 2) minimize the number of spectral gaps in the candidate paths, 3) minimize the number of FSs that the demand uses, and 4) minimize the difference between the allocated and the requested bit rate. Note that the model tries to prioritize the use of paths with fewer hops; this way, less network resources are occupied. Parameters α, β, γ, and δ are weighting factors used to put more or less weight to the terms in the objective function depending on the scenario. Factor ϵ is used as a convergence factor with value $\ll$1 to reduce the execution time of the MILP model. Finally, $\max\{R_m\}$ denotes the maximum value among all bit rates of the modulation formats supported by TSPs.

Regarding the constraints, Constraint (8) is the traffic constraint, ensuring that at least a bit rate equal to R_d will be assigned to demand d, while Constraint (9) accounts for the aforementioned difference between the allocated and requested bit rates. Constraints (10) set the limit of paths to be used by the demand to $L_{\max}$. Note that this value is independent of the value K used to obtain P_d. It could happen, for instance, that even having three candidate paths, the operator wants to limit the number of paths over which the demand will be routed to two. The single-path option can be easily implemented by setting $L_{\max} = 1$ with no need of further modifications in the formulation nor the mechanism itself. Constraints (11) set the allowed number of parts, limiting them to the number of available TSPs at source and destination and the maximum number of flows; parameter ϕ is set to 0 if MF-TSPs are employed and to 1 otherwise. Constraints (12) are the spectrum clashing constraints that avoid using more than one channel in a FS of a given path. Constraints (13) and (14) account for the number of parts and spectral gaps in the candidate paths, respectively. Constraints (15) give value to variables $z_{p,f}$ according to the actual status of the spectral resources and the decisions made by the model. Finally, Constraints (16) are used to give value to auxiliary variables $w_{p,f,f+1}$.

B. Heuristic Algorithm

As will be shown in the results section, the proposed SSRSMLA mechanism shows sub-second execution times in all tested network instances. However, its dependence on MILP may limit its scalability to address very large network instances. For this reason, we also developed a heuristic mechanism, called H-SSRSMLA, in order to provide still accurate results at lower computational cost, making it an option when scalability becomes challenging.

H-SSRSMLA is based on a greedy iterative mechanism that selects at every iteration the best solution element from a precalculated set of candidate solution elements. This element is considered as a part of the whole solution constituting the RSMLA for the demand. Next, the candidate solution elements are updated and the mechanism proceeds with the next iteration, until the demand is fully allocated or the candidate set becomes empty. Let us define O as the candidate set, with o being an element inside this set. Every element o has an associated path p, a modulation format m, and an ordered set of spectral gaps g_p corresponding to the spectral gaps of path p, with g_p^1 as the first element in the set g_p. With these, Fig. 4 depicts the pseudo-code of H-SSRSMLA.

The first phase in H-SSRSMLA has a purpose similar to the same phase in SSRSMLA, namely, to build the candidate path, modulation, and spectral gaps sets. For the latter, the mechanism obtains the end-to-end spectral gaps for a candidate path and sorts them in descending size order, that is, starting with the gap of biggest size and ending with the gap of smallest size. The idea behind this ordering is to favor the allocation of a demand in fewer parts, as bigger spectral gaps will likely be able to fit more types of modulation formats. As for the modulation format sets, note that the mechanism removes from the sets any modulation format for which its spectral width in FSs, that is, S_m, is bigger than the first element in the candidate gap sets, that is, the biggest spectral gap, as it would be infeasible to employ these modulation formats to serve the demand. Note that the mechanism stops promptly if empty candidate sets are detected.

Once all the candidate sets have been built, the mechanism proceeds with the construction of the initial set of candidate solution elements, assigning to each one of them a candidate path, modulation format, and set of spectral gaps. Then, it orders the solution elements according to a comparison method, named **CompareProcedure** in the pseudo-code. The exact details of **CompareProcedure** will be discussed later on.

The next phase starts a loop where the first element in O is added to the complete solution in each iteration, until any of the following stop criteria is met: 1) all the demand's bit rate is served, 2) no more candidates exist, 3) the number of parts has reached its limit $H_{\max}$, 4) there are not enough free TSPs at source or destination to support such number of parts (BV-TSPs case), or 5) the number of parts has reached the maximum number of flows (MF-TSPs case). Parameter ϕ has the same purpose as in SSRSMLA. In every iteration the mechanism also updates the

```
Inputs: d,G_n,F,M,K,H_max,L_max,FL_max; Output: Sol

Phase 1: Pre-processing
if TSP_s = 0 or TSP_t = 0 then
  Demand blocked
P_d ← K shortest paths from s_d to t_d using l_p as the metric
for i = 1 to K do
  M_i ← set of modulations belonging to M with TR_m ≥ l_i
  if M_i = ∅ then
    Remove p_i from P_d
if P_d = ∅ then
  Demand blocked
for i = 1 to |P_d| do
  g_i ← spectral gaps of i_th path
  if g_i = ∅ then
    Remove p_i from P_d
  else
    Order gaps in g_i in descending size order
if P_d = ∅ then
  Demand blocked
for i = 1 to |P_d| do
  for j = 1 to |M_i| do
    S_{i,j} ← number of FSs for modulation j in path i
    if S_{i,j} > |g_i^1| then
      Remove m_{i,j} from M_i
  if M_i = ∅ then
    Remove p_i from P_d
if P_d = ∅ then
  Demand blocked

Phase 2: Construction
O ← ∅
for i = 1 to |P_d| do
  for j = 1 to |M_i| do
    O ← o{p_i, m_{i,j}, g_i} ∪ O
Sort O according to CompareProcedure

Phase 3: Execution
ResBr ← R_d, Parts ← 0, Paths ← 0, p_s ← ∅, m_s ← ∅, s_s ← ∅
while ResBr > 0 and O ≠ ∅ and Parts < H_max and
Parts < TSP_s/φ and Parts < TSP_t/φ and
Parts < FL_max/(φ − 1) do
  p_s ← path of first element in O
  m_s ← modulation of first element in O
  s_s ← first S_m FSs in g_p^1 of first element in O
  Sol ← {p_s, m_s, s_s} ∪ Sol
  Parts + +
  ResBr ← ResBr − o.R_m
  if Paths = 0 then
    Paths + +
  else if Paths < L_max then
    Paths ← number of different paths used currently in Sol
  else
    Remove all elements o ∈ O with o.p not in Sol
  Update and sort again g_p for every o ∈ O
  for i = 1 to |O| do
    if |o_i.g_p^1| = 0 or o_i.S_m > |o_i.g_p^1| then
      Remove o_i from O
  Sort O according to CompareProcedure

Phase 4: Evaluation
if ResBr > 0 then
  Demand blocked
else
  Return Sol
  Update network resources according to Sol
  Demand allocated
```

Fig. 4. H-SSRSMLA mechanism pseudo-code.

resources associated to the elements in O (basically the set of spectral gaps g_p) and sorts them again according to **CompareProcedure**. Moreover, it eliminates all the infeasible candidates once updated. Furthermore, in every iteration the mechanism checks if the path limit $L_{\max}$ is reached, eliminating all elements in O with paths not in the solution built so far once this limit has been reached.

Finally, the mechanism checks if the entire bit rate of the demand has been served and, hence, a feasible RSMLA has been found for that particular demand. This being the case, the demand is successfully allocated and network resources are updated accordingly.

As for **CompareProcedure**, essentially, it is a prioritization mechanism that, given two elements from the set O, determines the relative order between them in terms of quality. So, when executing **CompareProcedure**, the elements of higher quality will be at the beginning of the set. The prioritization criteria are as follows:

1) Element with R_m = ResBr. If both meet, then:
   a) Element with bigger $|g_p^1|/(S_m \times h_p)$.
   b) Element with smaller $S_m \times h_p$.
   c) Element with bigger $|g_p^1|$.
2) Element with R_m > ResBr. If both meet, then:
   a) Element with smaller R_m.
   b) Element with bigger $|g_p^1|/(S_m \times h_p)$.
   c) Element with smaller $S_m \times h_p$.
   d) Element with bigger $|g_p^1|$.
3) Element with R_m < ResBr. If both meet, then:
   a) Element with bigger R_m.
   b) Element with bigger $|g_p^1|/(S_m \times h_p)$.
   c) Element with smaller $S_m \times h_p$.
   d) Element with bigger $|g_p^1|$.

In this regard, candidate elements with R_m equal to the current residual bit rate are prioritized, followed by elements with R_m higher than the residual bit rate and, last, elements with R_m lower than the residual bit rate. Within the latter two categories, elements with R_m closer to the residual bit rate are prioritized among the rest. By doing so, the number of parts into which a demand may be split is minimized. Also, inside each category, elements are prioritized in the following order: 1) bigger $|g_p^1|/(S_m \times h_p)$, 2) smaller $S_m \times h_p$, and 3) bigger $|g_p^1|$. By doing so, the mechanism makes a good trade-off between combinations of path and modulation format that will employ less network resources (i.e., FSs) and paths with bigger spectral gaps, that will likely fit the demand in fewer parts and leave more contiguous resources for future demands.

Note, however, that H-SSRSMLA is intended only for the multi-path case, while some adjustments are necessary to obtain a good mechanism for the single-path case. In fact, only setting $L_{\max} = 1$ in H-SSRSMLA would limit the heuristic to only examine the first selected path among the candidate ones, potentially getting trapped in a local optimum. To avoid this, we introduce some minor modifications in H-SSRSMLA to also work properly in the single-path case. Due to the space constraints, however, we comment only on the differences against the presented H-SSRSMLA, as the basic mechanism remains unchanged.

Essentially, instead of feeding the mechanism with all candidate paths at once, we execute the mechanism independently for every candidate path and, at the end, we select the produced solution of higher quality in terms of $|g_p^1|$/(number of employed FSs $\times h_p$). By doing so, we make sure that all candidate paths are explored to obtain a single-path solution, selecting among these solutions the one that makes the best trade-off between contiguous and used spectral resources. Note in this case that the value of $L_{\max}$ is not relevant as the mechanism is executed for one path at a time, implicitly setting this parameter to 1.

V. Results and Discussion

To evaluate the performance of both SSRSMLA and H-SSRSMLA mechanisms we have run extensive simulations. In order to quantify the benefits of optimally splitting the blocked demands, we benchmark them against a scenario where no splitting is performed (hereafter referred to as NSRSMLA). To this end, we employ the presented SSRSMLA mechanism forcing $H_{\max} = 1$ and $\alpha = 0$ in the formulation. We have considered two network topologies, the EON16 (16 nodes, 23 links) and the DT (14 nodes, 23 links) topologies, as well as two situations regarding spectral resources, namely, 160 and 320 FSs per fiber link, with $F_w = 6.25$ GHz and $G = 10$ GHz. As for the modulation formats supported by TSPs, Table I depicts them, together with their characteristics in terms of bit rate (Br), spectral width (Bw), and transmission reach (TR).

Results have been extracted by generating 10^5 bidirectional demand requests per execution. Requests arrive at the network randomly following a Poisson process, with exponentially distributed HTs. The bit rate of the demands is uniformly chosen among 25, 50, 100, and 200 Gb/s. Additionally, source and destination nodes for the demands are chosen with equiprobability among all network nodes. Moreover, $K = 3$, $H_{\max} = FL_{\max} = 4$, and $\alpha = \beta = \gamma = \delta = 1$ are set in all cases. As for $L_{\max}$, we set its value to 3 for the multi-path case and to 1 for the single-path case. All simulations are run in standard PCs with i7-3770 CPUs at 3.4 GHz and 16 GB RAM using CPLEX v12.2 as optimization software.

Graphs in Fig. 5 depict the BBR of the mechanisms in the 320 FSs scenarios. Graphs are plotted as a function of the normalized HT of the demands, fixing their average inter-arrival time to 1 time unit. It can be appreciated that the optimal splitting of the demands leads to substantial reductions on the BBR when compared to a scenario where no splitting is performed, with the multi-path option performing better than the single-path option. Particularly, relative gains in the ranges of 21%–40% and 27%–50% can be appreciated in the DT network scenarios for SSRSMLA against NSRSMLA, in single and multi-path scenarios, respectively. As for EON16 scenarios, relative gains in the ranges of 22%–80% and 38%–85% can be appreciated for SSRSMLA against NSRSMLA in the single and multi-path scenarios, respectively. For the sake of space, we do not show the graphs for the 160 FSs scenarios as similar results are obtained in them.

In this regard, it can be observed that, on average, the mechanism achieves greater reductions in the EON16 scenarios. This can be explained due to its topological characteristics; particularly, it has a physical diameter bigger than DT, so fewer candidate modulation formats are available. Hence, giving the opportunity to be served as multiple parts has larger impact on the number of demands that can be successfully allocated. As for the DT network scenarios, since demands have more candidate modulation formats, it is easier to allocate them without splitting, so the splitting procedure has a slightly lower impact in terms of BBR in such scenarios. On average,

TABLE I
Supported Modulation Formats

Modulation	Br (Gb/s)	Bw (GHz)	TR (km)
28 Gbaud SP-BPSK	25	42	3000
28 Gbaud PDM-BPSK	50	42	2400
28 Gbaud PS-QPSK	75	42	2000
28 Gbaud PDM-QPSK	100	42	1200
28 Gbaud PDM-8QAM	150	42	500
28 Gbaud QPM-16QAM	200	42	300
56 Gbaud SP-BPSK	50	70	3000
56 Gbaud PDM-BPSK	100	70	2400
56 Gbaud PS-QPSK	150	70	2000
56 Gbaud PDM-QPSK	200	70	1200
84 Gbaud SP-BPSK	75	98	3000
84 Gbaud PDM-BPSK	150	98	2400
112 Gbaud SP-BPSK	100	126	3000
112 Gbaud PDM-BPSK	200	126	2400

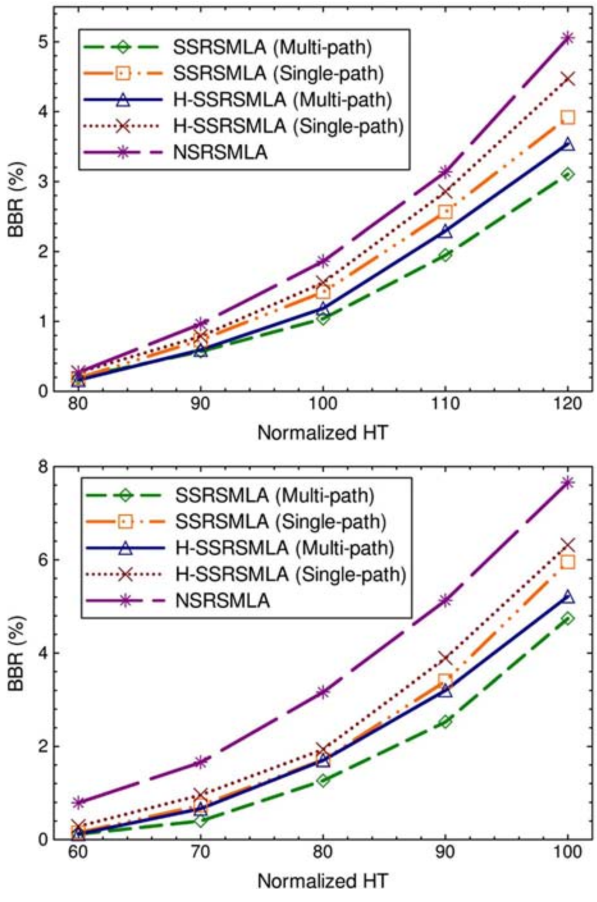

Fig. 5. BBR in DT/320 (top) and EON/320 (bottom) scenarios.

demands in EON16 network scenarios are split in 1.02 parts, while in DT network scenarios they are split in 1.0023 parts.

Concerning the comparison between single-path and multi-path cases with SSRSMLA, the multi-path case achieves relative gains in the ranges of 13%–22% and 20%–28% for the DT and EON16 network scenarios. Again, it can be seen that the relative difference is greater in the EON16 scenarios. The reason is the same as before: since EON16 has larger diameter, a single physical path has a limited set of candidate modulation formats, thus routing a demand through multiple physical paths increases the chances of serving it successfully. On average, for the multi-path case, demands in EON16 network scenarios are routed over 1.0082 paths, while in the DT network scenarios they are routed over 1.0009 paths.

In light of these results, let us analyze in more depth how often a demand exploits the SS opportunity and to what extent. Taking the average for the SSRSMLA mechanism in all tested scenarios, we can see that, in the DT network around 99.77% of the successfully allocated demands can be served without being split, while the remaining 0.23% is served employing two parts; as for the EON network, around 98.35% of the demands are not split, 1.51% employ two parts, 0.13% employ three parts, and 0.01% employ four parts. This means that the majority of demands that need splitting require only two parts. It can be concluded that in realistic traffic conditions, the SS opportunity is exploited only by a small percentage of the successful connections (roughly around 1%). However, it has to be highlighted that it is mostly used by the higher bit-rate demands that would be blocked otherwise. As a consequence, the gain provided in terms of BBR can be substantial, as depicted in Fig. 5.

Additionally, given the multi-objective nature of the MILP formulation employed in SSRSMLA, we have done additional simulations to determine the influence of the specific optimization objectives in the overall performance. To this end, we have focused on the scenario DT/320. There, for increasing values of the normalized HT, we have plotted the BBR of SSRSMLA for various configurations, namely, the case where all weighting parameters are set to 1 ("All one" case) and the cases where only one of these parameters is set to 1 and the rest to 0. Figure 6 depicts the obtained results for both single and multi-path cases.

It can be appreciated indeed, that the joint minimization of all the objectives provides lower BBR figures than focusing only on one of them. We can also see that minimizing only the number of allocated FSs performs very similarly to the joint minimization case, indicating that this objective takes the highest importance in the combined objective function. Focusing on the rest of the objectives, the minimization of the number of parts takes the second place, the minimization of the difference between allocated and requested bit rate takes the third one, and, finally, the minimization of the spectral gaps has the lowest importance. Focusing on this last objective, it can be seen that minimizing it solely results in a significant performance degradation when compared with the other cases. This can be

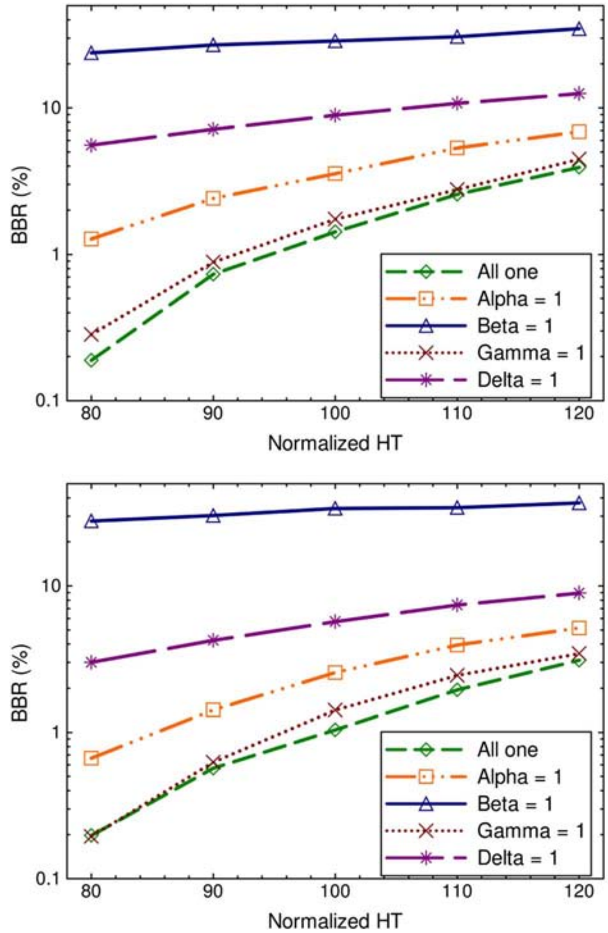

Fig. 6. BBR in single-path (top) and multi-path (bottom) cases.

explained as follows. Since SSRSMLA tries to minimize the spectral gaps present in the candidate paths, it divides the signal into many parts and chooses the most spectral-width-demanding modulation format in order to fill as many spectral gaps as possible, resulting in poor resource utilization.

Regarding the performance of the H-SSRSMLA mechanism, looking back at Fig. 5, it can be seen that it achieves BBR figures quite close to the ones obtained by the SSRSMLA mechanism. Table II depicts the performance comparison between both mechanisms, displaying for some of the tested scenarios the optimality gap of H-SSRSMLA against SSRSMLA and the average execution times, denoted as T_m and T_h for SSRSMLA and H-SSRSMLA, respectively. Indeed, H-SSRSMLA performs quite good when compared against SSRSMLA, with optimality gaps around 10%–14% on average. Note that for normalized HTs leading to BBR $<1\%$, a possible working range for these networks, the differences between H-SSRSMLA and SSRSMLA are very small. Regarding the execution times, as we mentioned in previous sections, although SSRSMLA can work in the sub-second range, when the number of FSs per link increases, the average execution time grows substantially. This was the main motivation for the development of H-SSRSMLA. In this regard, it

TABLE II
H-SSRSMLA VERSUS SSRSMLA

	Single-path			Multi-path		
Scenario	Gap (%)	T_m(ms)	T_h(ms)	Gap (%)	T_m(ms)	T_h(ms)
DT/160 HT = 35	0.24	180	16	1.44	160	15
DT/160 HT = 45	10.2	144	15	12.4	132	15
DT/320 HT = 90	7.27	472	16	4.14	438	16
DT/320 HT = 110	11.3	363	16	17.8	339	15
EON/160 HT = 25	9.34	147	14	6.43	133	14
EON/160 HT = 35	12.1	109	14	20.2	99	14
EON/320 HT = 70	12	326	14	10.2	313	14
EON/320 HT = 90	14.5	221	14	23	202	14

can be appreciated that its execution times remain steady for all the scenarios and more than 1 order of magnitude below those of SSRSMLA, highlighting the superior scalability of H-SSRSMLA.

Up to now we have shown the performance of the proposed mechanisms for single-path and multi-path routing solutions. Let us now focus on analyzing the implications on the results if a BV-TSP-based or a MF-TSP-based implementation is adopted for realizing SSA. To this end, we have conducted a cost analysis for various of the tested scenarios in order to determine which of the two implementation results is most beneficial cost-wise speaking. We assumed that nodes are equipped with an unlimited number of TSPs, so that we can extract which is the average number of TSPs needed in both implementations for a maximum benefit. Note that in the previous depicted results no blocking due to lack of TSPs occurs.

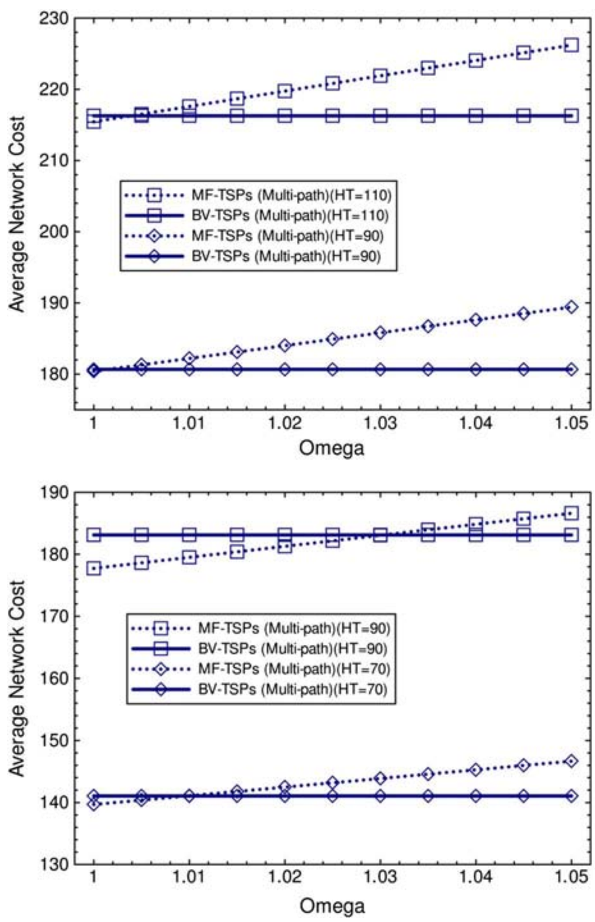

Fig. 7. Cost comparison in DT/320 (top) and EON/320 (bottom) scenarios.

For these results, we extract the average number of TSPs needed per node and multiply this value by the number of nodes in order to have the average number of TSPs used in the network. Then, we multiply this number by the cost of a TSP in order to obtain the average network cost of both implementations. Specifically, we have conducted a relative cost analysis, where we assume the cost of a BV-TSP as the normalized cost unit. Therefore, due to its higher complexity, the cost of a MF-TSP will be modeled as $\text{Cost}_{\text{MF-TSP}} = \omega \cdot \text{Cost}_{\text{BV-TSP}}$, with $\omega \geq 1$ accounting for this higher complexity and, hence, higher cost. Specifically, we have focused on the SSRSMLA mechanism, assuming the multi-path approach, as well as on two representative demand HTs for DT/320 and EON/320 scenarios.

Figure 7 depicts the average network cost as a function of ω for the mentioned scenarios. As observed, MF-TSPs are only cost effective with ω very close to 1 for all considered scenarios, since demand splitting is very infrequent in realistic working points (BBR <5%), as it has been pointed out before. Conversely, assuming that $\omega > 1$ due to the inherently higher complexity of MF-TSPs compared to BV-TSPs, it is generally more appropriate to deploy a slightly overprovisioned number of BV-TSPs per node to perform the SSA. Although not shown, the same discussion can be applied in the single-path approach, with the difference that it results in a slightly lower network cost for both BV-TSP and MF-TSP implementations. The reason behind this fact is that single-path leads to higher BBR figures; thus, on average, fewer TSPs are needed per node.

VI. CONCLUSIONS

In this paper we have presented mechanisms to optimally address the RSMLA problem in SS-enabled elastic optical networks. The presented mechanisms can work efficiently in both single-path and multi-path routing approaches. Reductions in BBR up to 80% can be achieved when comparing the proposed SSRSMLA mechanism against a mechanism not employing SSA. Between the two routing approaches, the multi-path approach performs slightly better due to its higher degree of freedom. Moreover, we have also analyzed multiple implementations of SSA, namely, BV-TSP-based and MF-TSP-based implementations. We have observed that, in realistic working loads, it is more cost effective to use the BV-TSP-based implementation as, in such scenarios, the capabilities of MF-TSPs are highly underutilized and do not justify their higher cost.

Acknowledgments

The work in this paper has been supported by the Government of Catalonia and the European Social Fund through a FI-AGAUR research scholarship grant and the Spanish Science Ministry through the project ELASTIC (TEC2011-27310).

References

[1] O. Gerstel, M. Jinno, A. Lord, and S. J. B. Yoo, "Elastic optical networking: A new dawn for the optical layer?" *IEEE Commun. Mag.*, vol. 50, no. 2, pp. s12–s20, Feb. 2012.

[2] M. Jinno, H. Takara, B. Kozicki, Y. Tsukishima, Y. Sone, and S. Matsuoka, "Spectrum-efficient and scalable elastic optical path network: Architecture, benefits, and enabling technologies," *IEEE Commun. Mag.*, vol. 47, no. 11, pp. 66–73, Nov. 2009.

[3] R. Wang and B. Mukherjee, "Spectrum management in heterogeneous bandwidth networks," in *Proc. of GLOBECOM*, Dec. 2012.

[4] Y. Yin, K. Wen, D. J. Geisler, R. Liu, and S. J. B. Yoo, "Dynamic on-demand defragmentation in flexible bandwidth elastic optical networks," *Opt. Express*, vol. 20, no. 2, pp. 1798–1804, Jan. 2012.

[5] K. Wen, Y. Yin, D. J. Geisler, S. Chang, and S. J. B. Yoo, "Dynamic on-demand lightpath provisioning using spectral defragmentation in flexible bandwidth networks," in *Proc. of European Conf. and Exhibition on Optical Communication*, Sept. 2011, paper Mo.2.K.4.

[6] M. Xia, R. Proietti, S. Dahlfort, and S. J. B. Yoo, "Split spectrum: A multi-channel approach to elastic optical networking," *Opt. Express*, vol. 20, no. 28, pp. 29143–29148, Dec. 2012.

[7] M. Jinno, H. Takara, Y. Sone, K. Yonenaga, and A. Hirano, "Multiflow optical transponder for efficient multilayer optical networking," *IEEE Commun. Mag.*, vol. 50, no. 5, pp. 56–65, May 2012.

[8] X. Chen, A. Jukan, and A. Gumaste, "Multipath defragmentation: Achieving better spectral efficiency in elastic optical path networks," in *Proc. of INFOCOM*, Apr. 2013.

[9] S. Thiagarajan, M. Frankel, and D. Boertjes, "Spectrum efficient super-channels in dynamic flexible grid networks—A blocking analysis," in *Proc. of Optical Fiber Communication Conf.*, Mar. 2011, paper OTuI6.

[10] W. Lu, X. Zhou, L. Gong, M. Zhang, and Z. Zhu, "Dynamic multi-path service provisioning under differential delay constraint in elastic optical networks," *IEEE Commun. Lett.*, vol. 17, no. 1, pp. 158–161, Jan. 2013.

[11] Z. Zhu, W. Lu, L. Zhang, and N. Ansari, "Dynamic service provisioning in elastic optical networks with hybrid single-/multi-path routing," *J. Lightwave Technol.*, vol. 31, no. 1, pp. 15–22, Jan. 2013.

[12] A. Pagès, J. Perelló, and S. Spadaro, "Lightpath fragmentation for efficient spectrum utilization in dynamic elastic optical networks," in *Proc. of Optical Network Design and Modeling*, Apr. 2012.

[13] I. Hussain, Z. Pan, M. Sosa, B. Basch, S. Liu, A. G. Malis, and A. Dhillon, "Generalized label for super-channel assignment on flexible grid," IETF draft, Oct. 2013.

[14] K. Christodoulopoulos, I. Tomkos, and E. A. Varvarigos, "Elastic bandwidth allocation in flexible OFDM-based optical networks," *J. Lightwave Technol.*, vol. 29, no. 9, pp. 1354–1366, May 2011.

[15] N. Amaya, I. Muhammad, G. S. Zervas, R. Nejabati, D. Simeonidou, Y. R. Zhou, and A. Lord, "Experimental demonstration of a gridless multigranular optical network supporting flexible spectrum switching," in *Proc. of Optical Fiber Communication Conf.*, Mar. 2011, paper OMW3.

[16] Y. Wang, X. Cao, and Y. Pan, "A study of the routing and spectrum allocation in spectrum-sliced elastic optical path networks," in *Proc. of INFOCOM*, Apr. 2011.

[17] M. Klinkowski and K. Walkowiak, "Routing and spectrum assignment in spectrum sliced elastic optical path network," *IEEE Commun. Lett.*, vol. 15, no. 8, pp. 884–886, Aug. 2011.

[18] X. Wan, N. Hua, and X. Zheng, "Dynamic routing and spectrum assignment in spectrum-flexible transparent optical networks," *J. Opt. Commun. Netw.*, vol. 4, no. 8, pp. 603–613, Aug. 2012.

[19] L. Velasco, M. Klinkowski, M. Ruiz, and J. Comellas, "Modeling the routing and spectrum allocation problem for flexgrid optical networks," *Photonic Network Commun.*, vol. 24, no. 3, pp. 177–186, Dec. 2012.

Albert Pagès received the M.Sc. degree in Telecommunications Engineering from Universitat Politècnica de Catalunya (UPC), Spain, in 2010. He is currently working toward the Ph.D. degree with the Optical Communication Group (GCO) and the Advanced Broadband Communications Center (CCABA) of UPC. He has participated in the FP-7 European research projects DICONET, STRONGEST, and GEYSERS. His research interests include optimization and management issues in elastic optical networks, as well as virtualization techniques in all-optical networks and data center environments.

Jordi Perelló received the M.Sc. and Ph.D. degrees in Telecommunications Engineering in 2005 and in 2009, respectively, both from Universitat Politècnica de Catalunya (UPC), Spain. Currently, he is an Assistant Professor in the Computer Architecture Department (DAC) at UPC. He has participated in various FP-6 and FP-7 European research projects (LIGHTNESS, EULER, STRONGEST, DICONET, etc.). Dr. Perelló has published more than 60 articles in international journals, conference proceedings, and book chapters. His research interests concern resource management, quality of service issues, and survivability of next-generation optical transport networks.

Salvatore Spadaro received the M.Sc. (2000) and the Ph.D. (2005) degrees in Telecommunications Engineering from Universitat Politècnica de Catalunya (UPC), Spain. He also received the Dr.Ing. degree in Electrical Engineering from Politecnico di Torino (2000), Italy. He is currently an Associate Professor in the Optical Communications group of the Signal Theory and Communications Department of UPC. Since 2000 he has been a staff member of the Advanced Broadband Communications Center (CCABA) of UPC. He has participated in various FP-6 and FP-7 European research projects. Dr. Spadaro has published more than 100 articles in international journals and conference proceedings. His research interests are in the fields of all-optical network solutions with emphasis on network control and management, network virtualization, intra-/inter-data-center interconnects, flex-grid technologies, energy efficiency in transport networks, software defined networking (SDN) for optical transport networks, and virtual optical network embedding.

Jaume Comellas received the M.Sc. (1993) and Ph.D. (1999) degrees in Telecommunications Engineering from Universitat Politècnica de Catalunya (UPC), Spain. His current research interests focus on IP over optical networking topics. He has participated in many research projects funded by the Spanish government and the European Commission. He has co-authored more than 100 research articles in international journals and conferences. He is an Associate Professor in the Signal Theory and Communications Department of UPC.