

Provisioning and Restoration With Sliceability in GMPLS-Based Elastic Optical Networks [Invited]

M. Dallaglio, A. Giorgetti, N. Sambo, F. Cugini, and P. Castoldi

Abstract—Multiflow bandwidth variable transponders supporting *sliceability* enable effective use of node and spectrum resources in elastic optical networks (EONs). Specifically, such transponders are capable of creating multiple optical flow units (i.e., *subcarriers*) that can be aggregated or independently routed according to the traffic requirements. This paper investigates the utilization of sliceability during provisioning and restoration in EONs. Specifically, a scheme is proposed to exploit the possibility of establishing/recovering an optical connection as a single *superchannel* or as a number of independent subcarriers. To this end the *spectrum overbuild* introduced by the sliceability application has to be taken into account. Both centralized and distributed implementations of the proposed schemes are evaluated in a generalized multiprotocol label switching (GMPLS)-based EON. Simulations show that despite the introduced spectrum overbuild, the utilization of sliceability permits us to significantly increase the amount of established/recovered traffic. Moreover, simulations show that the blocking probability reduction is achieved with only a slight increase in the setup time.

Index Terms—Elastic optical networks; GMPLS; PCE; Provisioning; Restoration; Sliceability.

I. INTRODUCTION

Optical transport networks are gradually evolving from the wavelength switched optical network (WSON) architecture, where all established optical connections (i.e., *lightpaths*) have to fit in a single channel of the fixed WDM grid [1], toward the elastic optical network (EON) architecture, where the spectrum is exploited by means of a flexible grid where each channel is adaptable to the effective bandwidth requirement [2]. Therefore, in EONs, high spectral efficiency is achieved, because each lightpath can use a different amount of spectrum depending on the bit rate and the exploited modulation format [3–5].

The most important upgrades needed in the data plane for supporting the EON architecture are related to

switching and transponder technologies. Nodes are required to support the switching of an arbitrary portion of spectrum [6]. Transponders are required to support the generation of multiple optical flow units (i.e., *subcarriers*) [7]. Using such transponders, subcarriers can be merged in high-rate *superchannels* (i.e., single connections composed of multiple subcarriers using a contiguous portion of spectrum) or can be *sliced*, i.e., independently routed along different paths toward the same or different destinations, and not necessarily on contiguous portions of spectrum. Those transponders are typically named *multiflow* or *sliceable* bandwidth variable transponders (SBVTs) [7,8]. Operators expect that the cost of an SBVT supporting N optical flows will be smaller than the cost of N bandwidth variable transponders (BVTs) [9,10]. The main reason is related to the production process; indeed, the same functionality of N BVTs, in the case of an SBVT, can be built in a single integrated platform, reducing production costs and transponder dimensions [9,10].

On the control plane, the needed generalized multiprotocol label switching (GMPLS) protocol extensions in support of flex-grid are under investigation [11–13]. Moreover, the path computation element (PCE) is evolving from a pure stateless condition to an active stateful architecture in which lightpaths' state information is stored and used to directly trigger networking operations, e.g., defragmentation [14,15]. Several experimental demonstrations have been recently deployed showing the feasibility of EONs with a properly extended GMPLS/PCE control plane [13,16–18].

Considering the huge amount of traffic traversing each link, in EONs, as in WSONs, it is fundamental to provide an effective recovery mechanism [19]. Focusing on dynamic restoration, the EON scenario poses an additional important challenge related to *spectrum fragmentation*. Indeed, due to the different spectrum occupation of established and released lightpaths, the spectrum of the links typically becomes fragmented in noncontiguous portions, thus preventing effective provisioning of new lightpaths and recovery of disrupted lightpaths. Since failures occur at unpredictable times and disrupted traffic should be recovered as fast as possible, defragmentation methods, which may further delay the recovery process, are not an

Manuscript received July 1, 2014; revised October 27, 2014; accepted November 3, 2014; published December 16, 2014 (Doc. ID 214649).

M. Dallaglio (e-mail: m.dallaglio@sssup.it), A. Giorgetti, N. Sambo, and P. Castoldi are with Scuola Superiore Sant'Anna, Pisa, Italy.

F. Cugini is with CNIT, Pisa, Italy.

<http://dx.doi.org/10.1364/JOCN.7.00A309>

option to be applied upon failure [20]. Therefore, dynamic restoration schemes for EONs have to cope with a highly fragmented spectrum. In this regard, sliceability exploitation provides an important opportunity: Disrupted superchannels can be sliced to independently recover each subcarrier; indeed, the probability of finding a number of narrow slots in a fragmented spectrum is typically higher than the probability of finding a single wide slot [21].

This paper assesses the potential benefits of sliceability during both provisioning and restoration in EONs in order to derive useful indications in support of the BVT implementation. Sliceability benefits are not obvious because, on the one hand, slicing subcarriers increases the probability of finding available spectrum resources, but on the other hand it requires a wider total spectrum with respect to the utilization of a single superchannel.

Specifically, EONs with a GMPLS control plane are considered in two different scenarios: a fully distributed scenario where each network node performs path computation for the locally originating lightpaths and a centralized scenario where a PCE is used for routing purposes. Specifically, the considered sliceability schemes allow the subcarriers to be established/recovered using a single superchannel or a number of independent subcarriers routed along different paths or using a noncontiguous spectrum along the same path.

The remaining part of this paper is organized as follows: Section II reviews the previous work on GMPLS-based EONs. Section III details the sliceability concept, its application in EONs, and its technological implementation aspects. Section IV describes the considered control plane scenarios and illustrates the proposed sliceability schemes. Section V describes the simulation scenario and presents the results obtained. Finally, Section VI concludes the paper.

II. PREVIOUS WORK

This section reviews several studies on EONs, with particular attention to fragmentation issues, sliceable transponders, and recovery mechanisms.

Several works proposed routing and spectrum assignment (RSA) schemes with the aim of reducing spectrum fragmentation. Specifically, the authors of [22] proposed an RSA scheme introducing the *expected capacity* parameter characterizing each spectrum slot; this allows us to allocate the slots that statistically minimize fragmentation. In [23], an RSA scheme adaptable to the lightpath rate and modulation format is proposed to achieve low fragmentation. In [24], the authors proposed a spectrum assignment scheme named *exact fit* that provides a lower fragmentation ratio compared to the common first-fit scheme. Other works considered defragmentation solutions exploiting rerouting [20] and frequency conversion [25] of previously established lightpaths to reoptimize the spectrum utilization once the fragmentation level becomes unsustainable. These solutions require additional spectrum resources, because they use a make-before-break approach. The

authors of [26] proposed a defragmentation technique based on spectrum retuning without implying any traffic disruption. However, the aforementioned defragmentation techniques may require a considerable amount of time and therefore cannot be applied during restoration.

Sliceability has been discussed to overcome spectrum fragmentation during provisioning in EONs. The work in [27] proposed to use SBVTs to implement a provisioning scheme based on path diversity, where sliceability is used for distributing subcarriers over multiple paths. The work in [28] exploited sliceability to fit spectrum fragments and to reduce the transmission limitation problem, and the results obtained show that sliceability enables lower bandwidth blocking. The authors of [29] proposed a provisioning scheme using sliceability to decrease the blocking probability experienced by high-data-rate connections. The authors of [30] applied sliceability to reduce fragmentation and unfairness among demands by means of integer linear programming (ILP) formulation, minimizing the fragmentation ratio. An experimental demonstration of SBVT architecture has been presented in [31], where subcarrier routing, bit rate, and optical reach can be adjusted based on the traffic requirements.

Regarding the recovery phase, several studies have been carried out to take advantage of the flexible nature of EONs, without considering sliceability. The work in [32] proposed a bandwidth-squeezed restoration scheme that employs class-based bandwidth reduction to increase the number of surviving lightpaths in the case of failure. Lightpaths with lower service level agreements (SLAs) can be throttled to allow some space to recover other connections. This scheme works fine in a scenario with lightpaths with different SLAs, while performance degrades when all connections have similar SLAs. In [33] a span restoration technique is investigated with different frequency conversion capabilities. Results showed that frequency conversion gives some benefit by improving spare capacity redundancy and spectrum efficiency. The authors of [34] proposed a genetic algorithm to solve an offline RSA problem aiming to maximize the recovered bit rate in the case of a single link failure.

To the best of our knowledge the only work considering sliceability during recovery in EONs is [35], where, in order to achieve a trade-off between recovery time and spectrum efficiency, the authors proposed a multipath recovery scheme that combines protection and restoration. However, sliceability still needs to be deeply investigated with the aim of reducing the restoration blocking probability, and also considering the introduced spectrum overbuild.

The work in [36] evaluates sliceability during dynamic restoration in a centralized control plane scenario considering only the achievable blocking probability. This paper investigates the potential of sliceability considering both centralized and distributed control plane scenarios. Moreover, the impact of sliceability is also evaluated during the provisioning of new connections. In addition to the blocking probability, this paper also considers the time required to establish and recover a connection in evaluating the sliceability performance.

III. SLICEABILITY APPLICATION AND IMPLEMENTATION

In this section, a possible network scenario is illustrated, where the application of sliceability can provide benefits during restoration. Then, the most relevant practical implications for supporting sliceability in SBVTs are discussed.

A. Sliceability Spectrum Overhead

An SBVT typically generates N subcarriers. Let F_s be the International Telecommunication Union Telecommunication Standardization Sector (ITU-T) frequency slot occupied by a superchannel composed of N subcarriers, i.e., the portion of the spectrum with bandwidth $|F_s|$ is expressed as an integer number of *frequency slices* of width 12.5 GHz [2]. Alternatively, if subcarriers are sliced, it is assumed that F_i is the frequency slot required by a single subcarrier. Slicing subcarriers typically introduces a spectrum overhead because, in general, $|F_s| < N \cdot |F_i|$. As an example, a polarization-multiplexed quadrature phase-shift keying (PM-QPSK) superchannel composed of four subcarriers can be accommodated in eight slices, i.e., $|F_s| = 100$ GHz; see Fig. 1. However, if the subcarriers are sliced, each F_i needs 37.5 GHz to fit the ITU-T flex-grid [9].

B. Sliceability Effectiveness

Figure 2 shows how sliceability can be exploited during restoration to increase the amount of recoverable traffic. A superchannel composed of $N = 4$ subcarriers is established along the path $s - b - d$ [Fig. 2(a)]. Upon failure on link $b - d$, the superchannel cannot be recovered as a whole because no path has enough contiguous spectrum [Fig. 2(b)]. If sliceability is applied, the different spectrum slots available along each path can be used to recover the four subcarriers. In Fig. 2(c) two subcarriers are recovered




Fig. 1. Spectrum required by a 400 Gbps superchannel, two adjacent 200 Gbps superchannels, and four adjacent 100 Gbps subcarriers.




Fig. 2. Example of sliceability utilization during restoration.

using a 200 Gbps superchannel along path $s - a - d$, one subcarrier is recovered along path $s - c - e - d$, and the last subcarrier is recovered along path $s - f - g - d$.

C. Sliceability Implementation

A typical SBVT architecture is illustrated in Fig. 3, including three modules: the transmission module, the adaptation module, and the client module [8,9]. The client module receives the tributary traffic in the form of $N \cdot M$ flows and performs the required signal processing, e.g., optical transport network framing (in Fig. 3, $N = M = 4$). The adaptation module provides flexibility between the $N \cdot M$ client traffic flows and the N subcarriers that will build up the superchannel generated by the SBVT. The utilization of this adaptation module can be particularly useful when sliceability is used. For instance, if upon network failure only some subcarriers are recovered, the client traffic flows with more stringent SLAs can be dynamically remapped to the recovered subcarriers. Finally, the transmission module is in charge of generation, modulation, and aggregation of the optical signal. This module is composed of three submodules. The laser source submodule generates N optical signals. The modulation submodule is composed of N modulators (e.g., implemented using photonic integrated circuits), and each of them modulates one optical signal using M traffic tributaries coming from the adaptation module (e.g., $M = 4$ in the case of PM-QPSK modulation). The aggregation submodule multiplexes the N subcarriers in a single output fiber.

The laser source can be implemented currently using two alternative technologies. The first technology is a multiwavelength source (i.e., a single source generating multiple subcarriers with a single laser), and the second one is the utilization of N independent laser sources. The former enables better frequency stability among subcarriers, thus reducing the risk of subcarrier interference. Moreover, one multiwavelength source is expected to be cheaper, and with a smaller footprint and energy consumption than N laser sources. The multiwavelength source supports the

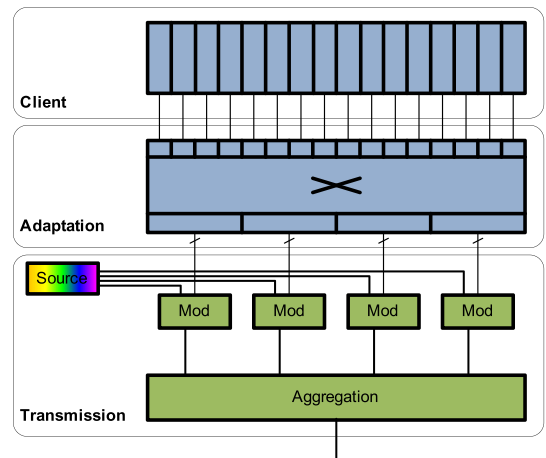


Fig. 3. SBVT architecture example. $N \cdot M = 16$ client flows are mapped to a superchannel composed of $N = 4$ subcarriers.

tunability of the whole comb; however, the relative tunability among subcarriers is typically not supported, and the contiguity among subcarriers is required, thus imposing a further constraint to the RSA. On the other hand, using N independent lasers does not introduce a contiguity constraint, providing full and independent tunability of each subcarrier [8].

In this study we evaluate the utilization of sliceability without a subcarrier contiguity constraint, in order to understand if the SBVT supporting sliceability can effectively provide benefits during network operation.

IV. GMPLS SCENARIOS IN EONS

Two different control plane scenarios are considered for evaluating the effectiveness of sliceability: the GMPLS scenario, where path computation during both provisioning and restoration are locally performed at the network nodes, and the GMPLS/PCE scenario, where a centralized PCE is employed to perform all the required path computations.

A. Distributed GMPLS

In the distributed scenario, each network node stores its own traffic engineering database (TED), including network topology and spectrum availability information. The TED is updated by means of link state advertisement (LSA) information exchanged through open shortest path first traffic engineering (OSPF-TE) protocol.

During provisioning, upon lightpath request, the source node computes a path and, using resource reservation protocol with extensions for traffic engineering (RSVP-TE), triggers the establishment of a label switched path (LSP) for serving the specific request.

During restoration, upon failure, the detecting node sends an RSVP-TE Notify [37] to the source node of each disrupted LSP and generates an OSPF-TE router LSA identifying the failed link that is flooded on the network [38]. A source node receiving the RSVP-TE Notify performs the following actions. First, it sends an RSVP-TE Tear message along the working path of the disrupted LSP, to release the utilized resources; then it computes a new path by solving the RSA problem considering the locally stored TED and utilizing one of the sliceability schemes detailed in Subsection IV.C; finally, when path computation is completed, the source node triggers RSVP-TE signaling to activate the computed backup path.

B. Centralized GMPLS/PCE

In the centralized scenario, besides the TED, the considered stateful PCE also stores an LSP state database including information about all the LSPs currently established in the network [39].

During provisioning, upon lightpath request, the source node uses the PCE protocol (PCEP) PCReq message to seek

a path to the PCE. The PCE performs the path computation and replies with a PCEP PCRep message including the computed path and a suggested label indicating the spectrum slot to reserve. The source node will then trigger RSVP-TE to perform the LSP establishment.

Similar to the distributed scenario, when a failure occurs, the detecting node sends an RSVP-TE Notify [37] to the source node of each disrupted LSP and floods an OSPF-TE router LSA for advertising the failure [38]. When the source node receives the Notify message, it sends an RSVP-TE Tear message to release the resources used by the disrupted LSPs. Then a PCEP PCReq message is sent to request the computation of a backup path. When the Tear message reaches the destination node, a PCEP PCRep message is sent to the PCE to report the released resources along the working path of the disrupted LSPs.

With respect to the distributed scenario, the use of a PCE delays the recovery procedure. However, the PCE acts as a unique point that can effectively coordinate the recovery operations in order to avoid contentions among the several LSPs under recovery. Specifically, the PCE returns the computed backup path for each disrupted LSP accompanied by a suggested label indicating the spectrum slot to reserve. The PCE avoids blocking by keeping track of the already assigned spectrum slots [40–42].

C. Proposed RSA Schemes Using Sliceability

Three RSA schemes are applied in both the GMPLS and GMPLS/PCE scenarios to evaluate the effectiveness of sliceability:

- *NO slice* is the reference scheme where all lightpath requests are routed/recovered using a single LSP.
- The *MAX slice* scheme applies sliceability to all the lightpath requests, slicing them in a number of LSPs using a single subcarrier.
- The *ADAPTIVE slice* scheme iteratively applies sliceability only to those lightpaths that cannot be routed/recovered as a whole.

V. PERFORMANCE EVALUATION

A. Simulation Scenario

The described schemes are evaluated using OPNET Modeler [43] during both provisioning and restoration in two network scenarios. The developed model includes an accurate implementation of the RSVP-TE, OSPF-TE, and PCEP protocols, with all the extensions required for managing EONs supporting sliceability.

During provisioning, the network is loaded with 100 Gbps LSPs (requiring three slices; see Fig. 1) and 400 Gbps LSPs (requiring eight slices; see Fig. 1). An LSP between node pair (s, d) is routed along one of the pre-computed paths in the set $P_{s,d}$ including all the paths within one hop from the shortest path. In both network

scenarios, using TED information, the path with the largest number of available frequency slots capable of accommodating the LSP is selected (i.e., *least congested routing*). 100 Gbps LSPs are routed as a whole using a single sub-carrier, whereas 400 Gbps LSPs can be routed as a whole using a single superchannel or divided in sub-LSPs of 200 Gbps (requiring five frequency slices) and/or 100 Gbps depending on the sliceability scheme considered.

When restoration is considered, a series of single failures of data plane bidirectional links is considered, assuming that the control plane remains fully operational. Specifically, the network is provisioned without applying sliceability (i.e., NO slice scheme). Upon failure, disrupted 100 Gbps LSPs are restored as a whole using least congested routing on the set of paths $P_{s,d}^i$ computed considering the failure of link i . Conversely, 400 Gbps LSPs can be recovered as a whole or sliced in sub-LSPs of 200 Gbps and/or 100 Gbps depending on the sliceability scheme considered.

During both provisioning and restoration, using the ADAPTIVE slice scheme, if a 400 Gbps LSP cannot be served as a whole, it is sliced in two sub-LSPs at 200 Gbps; in turn, if a 200 Gbps sub-LSP cannot be served as a whole, it is sliced in two 100 Gbps LSPs. When the sliceability is applied, a separate instance of the least congested routing is applied for each sub-LSP on the set of precomputed paths so that load balancing is achieved. However, depending on the network load, different sub-LSPs can also be routed on the same path by using a different spectrum allocation.

The considered test network has a Spanish topology, with 30 nodes and 56 bidirectional links with 256 frequency slices per direction, as depicted in Fig. 4. In the GMPLS/PCE scenario the PCE is located in Madrid. The traffic is uniformly distributed among node pairs, and LSPs arrive following a Poisson process; the mean holding time is fixed to 1 h. Spectrum assignment is first-fit. For each traffic load, a simulation has been performed for 1000 days of network operation where a single failure is generated each day.

The time required to configure an optical cross connect is considered to be 30 ms [44]. The average path computation time has been estimated to be 0.15 ms running the



Fig. 4. Test network topology.

considered routing algorithm on a personal computer with the following characteristics: Intel Core i7-4770 at 3.4 GHz (RAM: 16 GB).

B. Simulation Results: Provisioning

During provisioning, the considered schemes are evaluated in terms of provisioning blocking probability (Pr^p) and provisioning time (T^p). Pr^p is defined as the ratio between the blocked traffic bandwidth and the overall requested traffic bandwidth. LSP requests can be blocked during the path computation for lack of resources (i.e., *routing blocking*), or during the signaling procedure for both lack of resources (i.e., *forward blocking*) and resource contention (i.e., *backward blocking*) [19]. T^p is defined only for effectively established LSPs as the time between the generation of the LSP request and the conclusion of the RSVP-TE signaling.

Figure 5 shows Pr^p as a function of the network load in the GMPLS scenario. The figure shows that applying sliceability to all LSP requests (i.e., MAX slice scheme) degrades the blocking probability due to the introduced spectrum overhead (see Fig. 1). However, the figure shows that the ADAPTIVE scheme is able to reduce the blocking achieved by the NO slice scheme. This proves that sliceability can be effectively exploited to provision more traffic in the network. Figure 6 depicts T^p as a function of the network load in the GMPLS scenario. This figure shows that the utilization of sliceability implies an increase of the provisioning time. Indeed, if sliceability is applied, multiple RSVP-TE instances are contemporarily triggered to establish several sub-LSPs, thus generating node configuration queueing [45]. The figure also shows that using the ADAPTIVE scheme, the provisioning time is kept at the level of the NO slice scheme for low and medium network loads, while it gradually moves to the MAX slice scheme level for high loads when sliceability is applied to the majority of LSP requests.

Figure 7 shows Pr^p as a function of the network load in the GMPLS/PCE scenario. With respect to Fig. 5, the utilization of a PCE is able to considerably reduce the blocking probability of both NO slice and ADAPTIVE schemes.




Fig. 5. GMPLS scenario provisioning blocking probability, Pr^p .




Fig. 6. GMPLS scenario provisioning time, T^p .




Fig. 7. GMPLS/PCE scenario provisioning blocking probability, Pr^p .




Fig. 8. GMPLS/PCE scenario provisioning time, T^p .

Indeed, in the GMPLS scenario there is a backward blocking floor at about 10^{-4} that is completely avoided if all the path computations are coordinated by the PCE [40]. T^p in the GMPLS/PCE scenario is depicted in Fig. 8; with respect to the GMPLS scenario (Fig. 6), there is only a negligible increase due to the PCEP communication between network nodes and the PCE.

C. Simulation Results: Restoration

During restoration, the considered schemes are evaluated in terms of restoration blocking probability (Pr^r) and recovery time (T^r). Pr^r is defined as the ratio between the unrecovered traffic bandwidth and the overall traffic bandwidth disrupted by the failure. T^r is defined only for effectively recovered LSPs as the time between the failure and the conclusion of the RSVP-TE signaling to establish the backup path.

Figure 9 depicts Pr^r as a function of the network load in the GMPLS scenario. The figure shows that the achieved blocking probability is not acceptable (i.e., higher than 20%). Indeed, during restoration, backward blocking is very likely, because a high number of RSVP-TE instances are triggered almost simultaneously [19]. In this case, multiple restoration attempts with the crankback procedure are typically used to achieve an acceptable restoration blocking probability; however, only one restoration attempt is considered in Fig. 9. T^r is depicted in Fig. 10 as a function of the network load in the GMPLS scenario. This figure shows that for all the schemes, T^r is significantly longer than T^p , mainly due to the high probability of node configuration contention [45].

Figure 11 depicts Pr^r as a function of the network load in the GMPLS/PCE scenario. The figure shows that at low




Fig. 9. GMPLS scenario restoration blocking probability, Pr^r .




Fig. 10. GMPLS scenario recovery time, T^r .

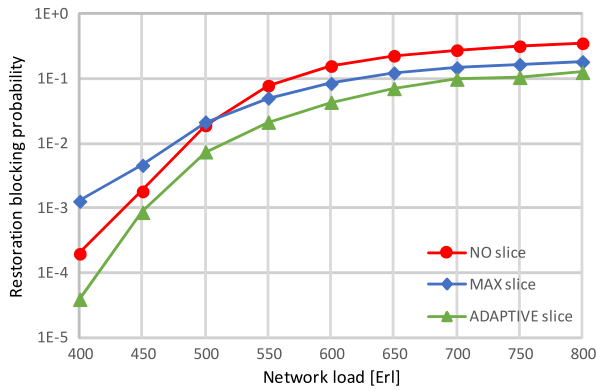


Fig. 11. GMPLS/PCE scenario restoration blocking probability, P_r^r .




Fig. 12. GMPLS/PCE scenario recovery time, T^r .

and medium loads, the MAX slice scheme degrades the blocking achieved when sliceability is not considered (i.e., NO slice scheme). This is because slicing an LSP when it could be recovered as a whole introduces a useless resource overbuild. Conversely, at high loads when the spectrum is highly fragmented and recovering LSPs as a whole is very unlikely, MAX slice provides a slight benefit with respect to NO slice. Finally, the figure shows that independently of the network load the ADAPTIVE slice scheme significantly decreases the restoration blocking probability, i.e., blocking is decreased by 72.5% at 600 erlang. Figure 12 depicts T^r as a function of the network load in the GMPLS/PCE scenario. The figure shows that the recovery time is considerably increased if the sliceability is applied to all disrupted LSPs (i.e., NO slice scheme); however, the increase in recovery time is marginal with the utilization of the ADAPTIVE slice scheme.

As mentioned in Subsection III.C, the use of sliceability introduces the possibility of partially recovering an LSP, i.e., only some of the subcarriers are effectively recovered. In those cases the adaptation module can be used to dynamically remap the client traffic flows with more stringent SLAs to the recovered subcarriers. Table I refers only to disrupted 400 Gbps LSPs with a network load of 600 erlang, and it shows the percentage of unrecovered, partially recovered (i.e., 100, 200, and 300 Gbps), and

TABLE I
PERCENTAGE OF RECOVERED BANDWIDTH OF 400 GBPS
DISRUPTED LSPs, WITH 600 ERLANG OF NETWORK LOAD

| Scheme | 0 | 100 | 200 | 300 | 400 |
|----------------|-------|------|------|------|-------|
| | [%] | [%] | [%] | [%] | [%] |
| NO slice | 19.53 | — | — | — | 80.47 |
| MAX slice | 5.01 | 2.77 | 2.63 | 2.82 | 86.76 |
| ADAPTIVE slice | 2.63 | 1.29 | 1.33 | 1.50 | 93.25 |

fully recovered LSPs. The table shows that, by using the ADAPTIVE scheme, also at high loads, most of the disrupted LSPs are totally recovered.

VI. CONCLUSION

This paper investigated the utilization of sliceability in GMPLS-based EONs. Two network scenarios have been considered: a fully distributed GMPLS control plane and a GMPLS/PCE control plane where a PCE is used for centralized path computation. Three sliceability schemes have been compared during both provisioning and restoration.

Simulation results showed that sliceability can provide significant benefits in terms of provisioning and restoration blocking probability if it is used only when there is not a possibility of serving the traffic request with a single superchannel. In terms of provisioning and recovery time, the utilization of sliceability introduces only a marginal degradation. Finally, it is proved that during restoration, sliceability can be effectively exploited only if a PCE is used to coordinate the recovery operations.

ACKNOWLEDGMENTS

This work has been partially supported by the project REACTION (Research and Experimental Assessment of Control plane architectures for In-Operation flexgrid Network re-optimization), approved within the first call for additional beneficiaries of the EC FP7 GÉANT (GN3plus) project. This paper is an extended version of the work presented in [36].

REFERENCES

- [1] "Spectral grids for WDM applications: DWDM frequency grid," ITU-T Recommendation G.694.1, June 2002.
- [2] "Spectral grids for WDM applications: DWDM frequency grid," ITU-T standard draft revised G.694.1 version 1.3, Unpublished ITU-T Study Group 15, Question 6, Feb. 2012.
- [3] 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.
- [4] O. Gerstel, M. Jinno, A. Lord, and S. Yoo, "Elastic optical networking: A new dawn for the optical layer?" *IEEE Commun. Mag.*, vol. 50, no. 2, pp. s12–s20, Feb. 2012.

- [5] N. Sambo, P. Castoldi, F. Cugini, G. Bottari, and P. Iovanna, "Toward high-rate and flexible optical networks," *IEEE Commun. Mag.*, vol. 50, no. 5, pp. 66–72, May 2012.
- [6] S. Frisken, S. Poole, and G. Baxter, "Wavelength-selective reconfiguration in transparent agile optical networks," *Proc. IEEE*, vol. 100, no. 5, pp. 1056–1064, May 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] N. Sambo, A. D'Errico, C. Porzi, V. Vercesi, M. Imran, F. Cugini, A. Bogoni, L. Poti, and P. Castoldi, "Sliceable transponder architecture including multiwavelength source," *J. Opt. Commun. Netw.*, vol. 6, no. 7, pp. 590–600, 2014.
- [9] V. Lopez, B. de la Cruz, O. G. de Dios, O. Gerstel, N. Amaya, G. Zervas, D. Simeonidou, and J. P. Fernandez-Palacios, "Finding the target cost for sliceable bandwidth variable transponders," *J. Opt. Commun. Netw.*, vol. 6, no. 5, pp. 476–485, May 2014.
- [10] A. Lord, "Core networks in the flexgrid era," in *Proc. ECOC*, Sept. 2014.
- [11] I. Hussain, A. Dhillon, and M. Sosa, "Super-channel optical parameters GMPLS signaling extensions," IETF draft, Oct. 2013.
- [12] I. Hussain, A. Dhillon, and M. Sosa, "Super-channel optical parameters GMPLS routing extensions," IETF draft, Oct. 2013.
- [13] R. Casellas, R. Munoz, J. Fabrega, M. Moreolo, R. Martinez, L. Liu, T. Tsuritani, and I. Morita, "Design and experimental validation of a GMPLS/PCE control plane for elastic CO-OFDM optical networks," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 1, pp. 49–61, Jan. 2013.
- [14] F. Paolucci, F. Cugini, A. Giorgetti, N. Sambo, and P. Castoldi, "A survey on the path computation element (PCE) architecture," *IEEE Commun. Surv. Tutorials*, vol. 15, no. 4, pp. 1819–1841, 2013.
- [15] R. Casellas, R. Martinez, R. Munoz, L. Liu, T. Tsuritani, and I. Morita, "Dynamic provisioning via a stateful PCE with instantiation capabilities in GMPLS-controlled flexi-grid DWDM networks," in *Proc. ECOC*, Sept. 2013.
- [16] F. Cugini, G. Meloni, F. Paolucci, N. Sambo, M. Secondini, L. Gerardi, L. Poti, and P. Castoldi, "Demonstration of flexible optical network based on path computation element," *J. Lightwave Technol.*, vol. 30, no. 5, pp. 727–733, Mar. 2012.
- [17] N. Sambo, G. Meloni, G. Berrettini, F. Paolucci, A. Malacarne, A. Bogoni, F. Cugini, L. Poti, and P. Castoldi, "Demonstration of data and control plane for optical multicast at 100 and 200 Gb/s with and without frequency conversion," *J. Opt. Commun. Netw.*, vol. 5, no. 7, pp. 667–676, July 2013.
- [18] R. Casellas, R. Munoz, J. Fabrega, M. Moreolo, R. Martinez, L. Liu, T. Tsuritani, and I. Morita, "Experimental assessment of a combined PCE-RMA and distributed spectrum allocation mechanism for GMPLS elastic CO-OFDM optical networks," in *OFC/NFOEC*, 2012.
- [19] A. Giorgetti, N. Sambo, I. Cerutti, N. Andriolli, and P. Castoldi, "Label preference schemes for lightpath provisioning and restoration in distributed GMPLS networks," *J. Lightwave Technol.*, vol. 27, no. 6, pp. 688–697, Mar. 2009.
- [20] T. Takagi, H. Hasegawa, K. Sato, Y. Sone, A. Hirano, and M. Jinno, "Disruption minimized spectrum defragmentation in elastic optical path networks that adopt distance adaptive modulation," in *Proc. ECOC*, Sept. 2011.
- [21] W. Shi, Z. Zhu, M. Zhang, and N. Ansari, "On the effect of bandwidth fragmentation on blocking probability in elastic optical networks," *IEEE Trans. Commun.*, vol. 61, no. 7, pp. 2970–2978, July 2013.
- [22] X. Chen, S. Ma, B. Guo, Y. Wang, J. Li, Z. Chen, and Y. He, "A novel fragmentation-aware spectrum allocation algorithm in flexible bandwidth optical networks," *Opt. Switching Netw.*, vol. 12, pp. 14–23, Apr. 2014.
- [23] N. Sambo, F. Cugini, G. Bottari, G. Bruno, P. Iovanna, and P. Castoldi, "Lightpath provisioning in wavelength switched optical networks with flexible grid," in *Proc. ECOC*, Sept. 2011.
- [24] A. Rosa, C. Cavdar, S. Carvalho, J. Costa, and L. Wosinska, "Spectrum allocation policy modeling for elastic optical networks," in *Proc. HONET*, Dec. 2012.
- [25] K. Wen, Y. Yin, D. Geisler, S. Chang, and S. Yoo, "Dynamic on-demand lightpath provisioning using spectral defragmentation in flexible bandwidth networks," in *Proc. ECOC*, Sept. 2011.
- [26] F. Cugini, F. Paolucci, G. Meloni, G. Berrettini, M. Secondini, F. Fresi, N. Sambo, L. Poti, and P. Castoldi, "Push-pull defragmentation without traffic disruption in flexible grid optical networks," *J. Lightwave Technol.*, vol. 31, no. 1, pp. 125–133, Jan. 2013.
- [27] M. Xia and S. Dahlfort, "Resilient provisioning for multi-flow elastic optical networking," in *Proc. ECOC*, Sept. 2013.
- [28] S. Dahlfort, M. Xia, R. Proietti, and S. Yoo, "Split spectrum approach to elastic optical networking," in *Proc. ECOC*, Sept. 2012.
- [29] A. Pages, J. Perello, S. Spadaro, and J. Comellas, "Optimal route, spectrum, and modulation level assignment in split-spectrum-enabled dynamic elastic optical networks," *J. Opt. Commun. Netw.*, vol. 6, no. 2, pp. 114–126, Feb. 2014.
- [30] R. Wang and B. Mukherjee, "Spectrum management in heterogeneous bandwidth networks," in *Proc. GLOBECOM*, Dec. 2012, pp. 2907–2911.
- [31] H. Takara, T. Goh, K. Shibahara, K. Yonenaga, S. Kawai, and M. Jinno, "Experimental demonstration of 400 Gb/s multi-flow, multi-rate, multi-reach optical transmitter for efficient elastic spectral routing," in *Proc. ECOC*, Sept. 2011.
- [32] Y. Sone, A. Watanabe, W. Imajuku, Y. Tsukishima, B. Koziacki, H. Takara, and M. Jinno, "Bandwidth squeezed restoration in spectrum-sliced elastic optical path networks (SLICE)," *J. Opt. Commun. Netw.*, vol. 3, no. 3, pp. 223–233, Mar. 2011.
- [33] Y. Wei and G. Shen, "Span restoration for flexi-grid optical networks under different spectrum conversion capabilities," in *Proc. DRCN*, Mar. 2013, pp. 79–87.
- [34] A. Castro, M. Ruiz, L. Velasco, G. Junyent, and J. Comellas, "Path-based recovery in flexgrid optical networks," in *Proc. ICTON*, July 2012.
- [35] A. Castro, L. Velasco, M. Ruiz, and J. Comellas, "Single-path provisioning with multi-path recovery in flexgrid optical networks," in *Proc. ICUMT*, Oct. 2012, pp. 745–751.
- [36] M. Dallaglio, A. Giorgetti, N. Sambo, F. Cugini, and P. Castoldi, "Impact of slice-ability on dynamic restoration in GMPLS-based flexible optical networks," in *OFC/NFOEC*, Mar. 2014.

- [37] Y. Rekhter and T. Li, "RSVP-TE extensions in support of end-to-end generalized multi-protocol label switching (GMPLS) recovery," IETF RFC 4872, May 2007.
- [38] D. Kats, K. Kompella, and D. Yeung, "Traffic engineering (TE) extensions to OSPF version," IETF RFC 3630, Sept. 2003.
- [39] E. Crabbe, I. Minei, J. Medved, and R. Varga, "PCEP extensions for stateful PCE," IETF draft, Feb. 2014.
- [40] A. Giorgetti, F. Cugini, N. Sambo, F. Paolucci, N. Andriolli, and P. Castoldi, "Path state-based update of PCE traffic engineering database in wavelength switched optical networks," *IEEE Commun. Lett.*, vol. 14, no. 6, pp. 575–577, June 2010.
- [41] D. Alvarez, V. Lopez, J. Anamuro, J. de Vergara, O. de Dios, and J. Aracil, "Utilization of temporary reservation of path computed resources for multi-domain path computation element protocols in WDM networks," in *Proc. NOF*, Nov. 2011.
- [42] A. Giorgetti, L. Valcarenghi, F. Cugini, and P. Castoldi, "PCE-based dynamic restoration in wavelength switched optical networks," in *Proc. ICC*, May 2010.
- [43] "OPNET Technologies" [Online]. Available: <http://www.opnet.com>.
- [44] L. Liu, R. Munoz, R. Casellas, T. Tsuritani, R. Martinez, and I. Morita, "OpenSlice: An OpenFlow-based control plane for spectrum sliced elastic optical path networks," in *Proc. ECOC*, Sept. 2012.
- [45] A. Giorgetti, F. Paolucci, F. Cugini, and P. Castoldi, "Fast restoration in SDN-based flexible optical networks," in *OFC*, Mar. 2014.