Impact of Aggregation Level on the Performance of Dynamic Lightpath Adaptation under Time-Varying Traffic

A. Asensio[1], M. Klinkowski[2], M. Ruiz[1], V. López[3], A. Castro[1], L. Velasco[1], J. Comellas[1]

[1] Universitat Politècnica de Catalunya (UPC), Barcelona, Spain
[2] National Institute of Telecommunications, Warsaw, Poland
[3] Telefónica Investigación y Desarrollo, Madrid, Spain
Email: adrian.asensio@tsc.upc.edu

***Abstract*—In this article we focus on lightpath adaptation under time-varying traffic in a dynamic elastic optical network (EON) implementing flexgrid optical technology. In the considered scenario, a number of IP/MPLS metro area networks performing traffic aggregation are connected through a core EON. We explore the elastic spectrum allocation (SA) capability of EON and, in this context, we study the effectiveness of three alternative SA policies, namely *Fixed*, *Semi-Elastic* and *Elastic*. For each elastic SA policy, we develop a dedicated algorithm which is responsible for adaptation of spectrum allocated to lightpath connections in response to traffic changes. The evaluation is performed for a set of network scenarios, each one characterized by a different level of traffic aggregation, and hence different traffic variability. As simulation results show, the effectiveness of SA policies highly depends on both the aggregation level and maximum lightpath capacity. In particular, in our experiments up to 21% more traffic is served with the proposed elastic SA than with the fixed SA in a network with low aggregation and high lightpath capacity.**

***Key words*—dynamic lightpath adaptation, elastic optical network, IP/MPLS network, time-varying traffic.**

I. Introduction

The advent of the flexible spectrum grid (*flexgrid*) technology [1], [2] brings new opportunities to next-generation transport networks since it allows for elastic, adaptive, highly-scalable, and on-demand bandwidth (bit-rate) provisioning in optical networks. The key technologies that are paving the way to devise these novel elastic optical network (EON) architectures are: 1) the availability of flexgrid ready Wavelength Selective Switches (WSS) to build Bandwidth-Variable Optical Cross-Connects (BV-OXC), 2) the development of advanced modulation formats and techniques, both single-carrier (such as k-PSK, k-QAM) and multi-carrier (such as O-OFDM), to increase efficiency and being capable of extending the reach of optical signals avoiding expensive electronic regeneration (3R); 3) Multi-Flow Transponders (MF-TP) that are able to deal with several flows in parallel, thus adding even more flexibility and reducing costs [3]. BV-OXCs allow creating an optical routing path (*lightpath*) through the network by switching transmitted signals within their frequency bandwidth to appropriate switch output ports. Concurrently, the role of MF-TP is to adapt the client data to be sent to/received from the optical network using just enough frequency resources. ITU-T has recently revised the G.694.1 recommendation and included the definition of a flexible wavelength division multiplexing (WDM) grid [4]. According to [4], an optical *channel* has flexibly (ad-hoc) assigned spectrum, which covers both the frequency range occupied by the optical signal and the guard band required for the roll-off filters. For more details on EON architectures and proof-of-concept EON experiments we refer to [2], [5], [6].

Recently, in [7] and [8] we have investigated how aggregation can be done, in part, at the optical layer, reducing thus IP/MPLS network size. In particular, IP/MPLS routers equipped with MF-TPs are connected to BV-OXCs resulting in a single-layer network approach in which a number of IP/MPLS aggregation networks are connected through a flexgrid-based core network (see Fig. 1). The results obtained in [8] show that if the core network is extended towards the edges and, by these means, the number of metro areas is increased, then significant Capital Expenditure (CAPEX) savings (of above 20-30%) can be obtained. However, having a larger number of relatively small metro areas implies a reduction in the traffic aggregation at the IP/MPLS layer, thus resulting in higher variability of the traffic flows data-rate offered to the optical layer along the day.

One of the main advantages of EON is the capability to allocate spectrum resources elastically, according to traffic demands. Indeed, the resources may be used efficiently, firstly, because of the higher granularity of the flexgrid which allows fitting closely the allocated spectrum and the signal bandwidth, secondly, due to the elastic (adaptive) spectrum allocation (SA) in response to traffic variations. On the contrary to WDM networks, in which the width of optical channels is fixed and equal, in EON the channel allocated to a lightpath may be expanded/reduced when the required bit rate of a demand increases/decreases. In this context, adaptive SA with a known a priori 24-hour traffic patterns has been addressed in [9]-[11] and spectrum adaptation under dynamic traffic demands was studied in [12]-[13]. Concurrently, different policies for elastic SA were proposed, including symmetric [9]-[11] and asymmetric [9], [10], [13] spectrum expansion/reduction around a reference frequency as well as the entire spectrum re-allocation policy [9]-[10].

First in [9] and later in [10] we showed that in EON with

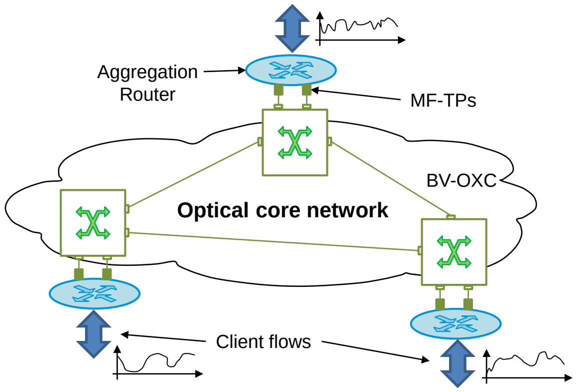

Fig. 1. A flexgrid-based optical core network connecting aggregation IP/MPLS networks.

time-varying traffic and under a priori known (deterministic) traffic demands, elastic SA policies performing lightpath adaptation (either symmetric, asymmetric, or reallocation-based) allow to serve more traffic than when fixed SA is applied. In this paper, we extend the work to a dynamic network scenario. Our main objective is to assess the efficiency of elastic SA policies for networks with different levels of traffic aggregation and, consequently, with different traffic variability. In particular, we investigate the relation between aggregation and the performance of the proposed elastic SA policies. To this end, we evaluate the performance of elastic SA assuming different dimensions of the flexgrid-based core network, which in turn translates into the level of aggregation of the traffic to be carried by lightpaths in the core network [7]-[8]. As a result of extending the core network towards the edges lightpaths' length could introduce reachability problems, especially for those with the largest capacity (e.g. 400Gb/s) [14]; this can be solved by managing the modulation format and introducing inverse-multiplexing (e.g. 4x100 Gb/s). To the best of our knowledge, the impact of aggregation level and traffic variability on the efficiency of adaptive spectrum allocation in supporting time-varying traffic demands has been not studied so far in a dynamic EON scenario. Note that this is one of issues that should be investigated in order to recognize true benefits of the flexgrid technology and its ability to be a successor of current fixed-grid technologies in optical transport networks.

The remainder of the paper is organized as follows. In Section 2, we define the problem of dynamic lightpath adaptation, present the considered spectrum allocation policies, and describe dedicated lightpath adaptation algorithms. In Section 3, we perform simulation experiments to study the performance of the proposed SA policies. Performance evaluation is studied using several realistic networks of different size and considering two different scenarios were the capacity of the lightpaths is limited to 400Gb/s and 100Gb/s. Finally, in Section 4 we conclude the work.

II. Dynamic Lightpath Adaptation

In this Section, we address the problem of lightpath adaptation in EON under time-varying traffic demands. First, we discuss the assumption that we take in our study and we describe the considered spectrum allocation policies. Then we formulate the problem of lightpath adaptation and, for each SA policy, we propose an algorithm to solve it.

We focus on an EON which implements a discrete frequency grid defined in [4]. In particular, a flexgrid consists of a set of nominal *central frequencies* (CF) and a set of frequency *slices*, where a slice corresponds to a spectrum segment that lays between two CFs. Flexgrid [4] requires to have the spectrum allocated symmetrically around a CF and, therefore, each lightpath occupies an even number of slices. In [4], nominally, the CF granularity, i.e., the spacing between neighboring CFs, and the slice width is equal to Δ_f=6.25GHz, although it was recently shown that finer granulates, such as Δ_f=3.125GHz, could be needed [8].

In our study, we assume that the profile of the client flows arriving to the aggregation router is known. A traffic profile is characterized by a guaranteed bit-rate and a range of variation, given for a considered set of time periods. Therefore, lightpaths can be pre-planned, using for example the models in [10], and established on the network. We assume that the lightpath capacity is limited and, therefore, several lightpath connections may exist between a pair of aggregation networks. Note that the optical signal reach may be increased by using such techniques as inverse-multiplexing without a need for excessive 3R regeneration. Once in operation, the established lightpaths must elastically adapt their capacity so as to convey as much bit-rate as possible from the demanded by the aggregation networks. In particular, each change in aggregated traffic resulting in a change in the amount of optical spectrum of a lightpath (i.e., spectrum expansion/reduction) requires a request to be sent to the control plane of the flexgrid core network to find the appropriate SA for that lightpath (see problem definition in Sec. II.B). In response to this request, a dynamic SA algorithm (see Sec. II.C) which is implemented in a Path Computation Element (PCE) is in charge of adapting the spectrum allocation in accordance to the SA policy which is used in the network (see Sec. II.A).

A. Spectrum Allocation Policies

In [10] we have discerned three alternative spectrum allocation policies for time-varying traffic demands (Fig. 2).

The SA policies put the following restrictions on the assigned CF and the allocated spectrum width, in particular:

1. *Fixed* (Fig. 2a): both the assigned CF and spectrum width do not change in time. At each time period, demands may utilize either whole or only a fraction of the allocated spectrum to convey the bit-rate requested for that period.
2. *Semi-Elastic* (Fig. 2b): the assigned CF is fixed but the allocated spectrum may vary. Here, spectrum increments/decrements are achieved by allocating/releasing frequency slices symmetrically, i.e., at each end of the already allocated spectrum while keeping invariant the CF. The frequency slices can be shared between neighboring demands, but used by, at most, one demand in a time interval.

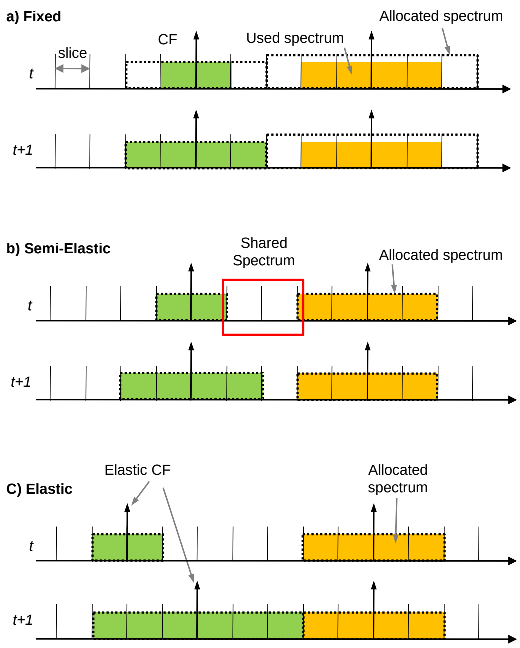

Fig. 2. Three spectrum allocation policies for time-varying traffic in a flexgrid network. Two time intervals are observed: t before and t+1 after spectrum adaptation has been performed.

3. *Elastic* (Fig. 2c): asymmetric spectrum expansion/ reduction (with respect to the already allocated spectrum) is allowed and it can lead to short shifting of the central frequency. Still, the relative position of the lightpaths in the spectrum remains invariable, i.e. no reallocation in the spectrum is performed.

B. *Problem statement*

The problem of dynamic lightpath adaptation addressed in this paper can be formally stated as:

<u>*Given*</u>:

- a core network topology represented by a graph G(N, E), being N the set of BV-OXC nodes and E the set of bidirectional fiber links connecting two BV-OXC nodes; each link consists of two unidirectional optical fibers;
- a set S of available slices of a given spectral width for every link in E;
- a set L of lightpaths already established on the network; each lightpath l is defined by the tuple $\{R_l, f_l, s_l\}$, where the ordered set $R_l \subseteq E$ represents its physical route, f_l its central frequency and s_l the amount of frequency slices.
- a lightpath $p \in L$ for which spectrum adaptation request arrives and the required number of frequency slices, $(s_p)^{req}$.

<u>*Output*</u>:

- the new values for the spectrum allocation of the given lightpath p: $\{R_p, f_p, (s_p)'\}$ and $\{R_p, (f_p)', (s_p)'\}$, respectively, if the *Semi-Elastic* and *Elastic* policy is used.

<u>*Objective*</u>: maximize the amount of bit-rate served.

C. *Spectrum Adaptation Algorithms*

For the *Fixed* SA policy, the allocated spectrum does not change in time, therefore, any fraction of traffic that exceeds the capacity of the established lightpath is lost. Regarding the *Semi-Elastic* and *Elastic* policies, the corresponding lightpath adaptation algorithms are presented in Table I and Table II, respectively. In the following, we discuss the details of these algorithms.

- *Semi-Elastic* algorithm: the elastic operation is requested for a lightpath p and the required amount of frequency slices to be allocated, maintaining f_p invariant, is given. Since flexgrid [4] is implemented, $(s_p)^{req}$ must be an even number. If elastic spectrum reduction is requested, the tuple for lightpath p is changed to $\{R_p, f_p, (s_p)^{req}\}$ (lines 1-2). In the opposite, when an elastic expansion is requested, the set of spectrum adjacent lightpaths at each of the spectrum sides is found by iterating on each of the links of the route of p (lines 4-7). The greatest value of available spectrum without CF shifting, s_{max}, is subsequently computed and the value of spectrum slices actually assigned to p, $(s_p)'$, is computed as the minimum between s_{max} and the requested one (lines 8-9). The tuple representing lightpath p is now $\{R_p, f_p, (s_p)'\}$.
- *Elastic* algorithm: here, the CF of p can be changed so the difference with the *Semi-Elastic* algorithm explained above is related to that issue. Now, the value of s_{max} is only constrained by the amount of slices available between the closest spectrum adjacent paths. Then, s_{max} is

TABLE I *SEMI-ELASTIC* SPECTRUM ALLOCATION ALGORITHM

INPUT: $G(N,E)$, S, L, p, $(s_p)^{req}$	
OUTPUT: $(s_p)'$	
1:	**if** $(s_p)^{req} \le s_p$ **then**
2:	$(s_p)' \leftarrow (s_p)^{req}$
3:	**else**
4:	$L^+ \leftarrow \varnothing, L^- \leftarrow \varnothing$
5:	**for each** $e \in R_p$ **do**
6:	$L^- \leftarrow L^- \cup \{l \in L: e \in R_l$, adjacents(l, p) , $f_l < f_p\}$
7:	$L^+ \leftarrow L^+ \cup \{l \in L: e \in R_l$, adjacents(l, p) , $f_l > f_p\}$
8:	$s_{max} \leftarrow 2^*\min\{\min\{f_p - f_l - s_l, l \in L^-\},$ $\min\{f_l - f_p - s_l, l \in L^+\}\}$
9:	$(s_p)' \leftarrow \min\{s_{max}, (s_p)^{req}\}$
10:	**return** $(s_p)'$

TABLE II *ELASTIC* SPECTRUM ALLOCATION ALGORITHM

INPUT: $G(N,E)$, S, L, p, $(s_p)^{req}$	
OUTPUT: $(f_p)'$, $(s_p)'$	
1:	**if** $(s_p)^{req} \le s_p$ **then**
2:	$(s_p)' \leftarrow (s_p)^{req}$
3:	$(f_p)' \leftarrow f_p$
4:	**else**
5:	$L^+ \leftarrow \varnothing, L^- \leftarrow \varnothing$
6:	**for each** $e \in R_p$ **do**
7:	$L^- \leftarrow L^- \cup \{l \in L: e \in R_l$, adjacents(l, p) , $f_l < f_p\}$
8:	$L^+ \leftarrow L^+ \cup \{l \in L: e \in R_l$, adjacents(l, p) , $f_l > f_p\}$
9:	$s_{max} \leftarrow \min\{f_p - f_l - s_l, l \in L^-\}$ + $\min\{f_l - f_p - s_l, l \in L^+\}$
10:	$(s_p)' \leftarrow \min\{s_{max}, (s_p)^{req}\}$
11:	$(f_p)' \leftarrow$ findSA_MinCFShifting $(p, (s_p)', L^+, L^-)$
12:	**return** $(f_p)'$, $(s_p)'$

the sum of the minimum available slices along the links in the left side and the minimum available slices in the right side of the allocated spectrum (line 9). Finally, the returned value (f_p)' is obtained by computing the new CF value so as to minimize CF shifting (line 11).

III. Results

In this Section, we present and discuss the performance results of a dynamic EON connecting aggregation networks and operating with elastic spectrum allocation policies under time-varying traffic demands.

A. Scenario

We consider three scenarios representing different levels of traffic aggregation consisting in three different core networks: TEL21, TEL60, and TEL100, with 21, 60, and 100 nodes, respectively. These networks are based on the Telefónica (TEL) network presented in Fig. 3. The main characteristics of the networks are presented in Table III.

For the sake of fairness in further comparisons, each core network was designed to provide approximately the same blocking performance at a given total traffic load. To this aim, the models and methods proposed in [15] to design core networks under dynamic traffic assumption were adapted and used to fit our pre-planning needs. Note that, although the relation between traffic load and un-served bandwidth is kept pretty similar in all defined scenarios, the number of core node strongly affects the level of traffic aggregation. Namely, since the analyzed core networks cover the same geographical area and the offered traffic is the same for each network, the lower number of core nodes the larger aggregation networks as well as the higher the flow grooming on lightpaths.

For the preplanning phase, the flows uniformly distributed among core node pairs were generated following the time-variant bandwidth distribution used in [9]. In particular, three different traffic profiles (TP), defined by their h^{min} and h^{max} bit-rate values, were used to compute bit-rate fluctuations of flows in time. Here, a granularity of 1h was considered. Each flow belongs to one of those TPs (with the same probability) and bit-rate fluctuations are randomly chosen following a uniform distribution within an interval defined by h^{min} and h^{max}. The considered TPs with their bandwidth intervals are TP1 = [10,100], TP2 = [40,200], and TP3 = [100,400]. All traffic flows between a pair of core nodes are transported using the minimum number of lightpaths. Hence, more than one lightpath between each pair of nodes might be needed. Furthermore, the lightpaths' maximum capacity was limited to 100 Gb/s and to 400 Gb/s for the completeness of our study and the quadrature phase shift keying (QPSK) modulation format was assumed.

The preplanning problem of finding the initial route and spectrum allocation for each pair of generated traffic matrix and core network so as to minimize the amount of blocked bit-rate, was solved using the algorithms proposed in [10] for each of the proposed elastic policies. The solutions of the preplanning phase, i.e., the routing and the initial spectrum allocations of established lightpaths, were used as input data

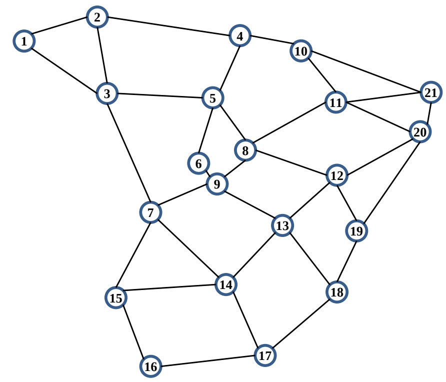

Fig. 3. TEL21 network topology used in this paper.

Table III Details of Network Topologies

	Nodes	Links	Nodal degree
TEL21	21	35	3.33
TEL60	60	90	3.0
TEL100	100	125	2.5

for our elastic network simulator.

For evaluation purposes we developed an ad-hoc event-driven simulator in OMNET++ [16] running the algorithms presented in Sec. II. Similarly as for the planning phase, a set of flows following the TP defined above where configured. For each flow, the requested bit-rate varies randomly ten times per hour, i.e. a finer and randomly-generated granularity both in time and bit-rate is considered in the simulator. Hence, bitrate variations arrive randomly to the system where a grooming module ensures that all the bit-rate can be served. If either more resources are needed for a lightpath or resources can be released, the grooming module of the aggregation network requests the core network to perform an elastic operation on that lightpath. To this end, algorithms in Table I and Table II are executed provided that either the *Semi-Elastic* or the *Elastic* policies are used; no lightpath adaptation algorithm is run when the *Fixed* policy is considered.

In our experiments, the optical spectrum width was set to 1.2 THz and the slice width was fixed to 3.125GHz (spectral granularity of 6.25GHz).

B. Efficiency of SA policies

Firstly, we analyze the efficiency of SA policies in the considered core network scenarios assuming different lightpath capacity limits (100 Gb/s and 400 Gb/s). We recall that for the largest core network (i.e., TEL100) we have the lowest level of traffic aggregation and, as a consequence, the highest variability of traffic offered to the core.

Fig. 4 illustrates the accumulated percentage of un-served bit-rate (UB) as a function of the offered traffic load, represented by the average load for each core network and lightpath capacity limit.

We observe that when the aggregation level is high (TEL21) and lightpath capacity is limited to 100 Gb/s, all the SA policies offer similar performance. The rationale behind that is in the low traffic variability in this scenario and, therefore, the network does not benefit from adaptive SA. Furthermore, since the volume of aggregated traffic is large

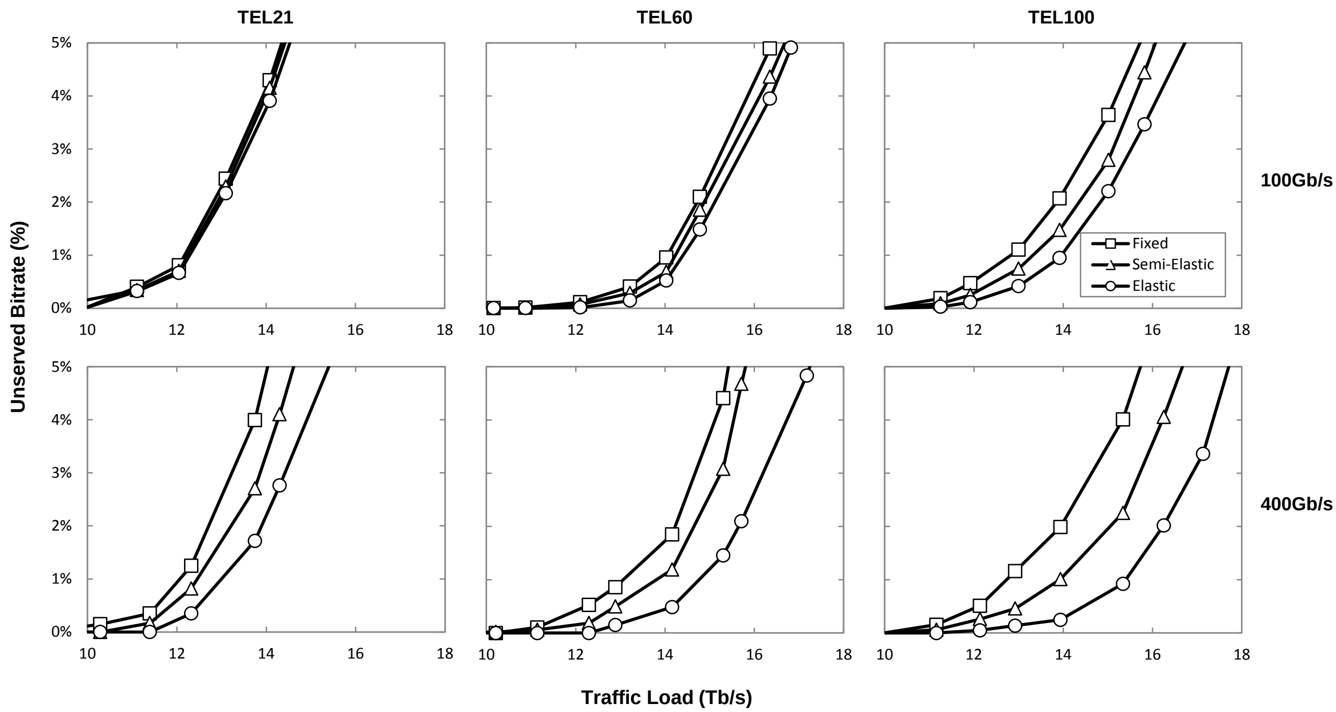

Fig. 4. Percentage of un-served bitrate against the traffic load for TEL21 (left), TEL60 (centre) and TEL100 (right) networks when the capacity of the lightpaths is limited to 100Gb/s (top) and 400Gb/s (bottom).

TABLE IV GAIN OF ADAPTIVE SA POLICIES VS. FIXED SA AT 1% OF UN-SERVED BIT RATE

Max. Capacity	SA Policy	TEL21	TEL60	TEL100
400 Gb/s	Elastic	7.73%	12.99%	20.99%
	Semi-Elastic	3.25%	5.67%	9.34%
100 Gb/s	Elastic	0.92%	2.47%	8.83%
	Semi-Elastic	0.60%	1.24%	3.85%

and the lightpath capacity is low, a large number of lightpath connections are established between each pair of aggregation networks where a relatively large fraction of these lightpath is always saturated. Hence, elasticity is not required in that scenario. As soon the aggregation level is gradually increased in TEL60 and TEL100, the benefits of elasticity also increase, although they are rather limited.

In contrast, the advantages of elasticity clearly appear when the capacity of the lightpaths is increased to 400Gb/s. Starting from the limited benefits observed when the aggregation level is high (TEL21), un-served bit-rate remarkably increases when that level is reduced in TEL60 and specially in TEL100.

Regarding SA policies, in all the analyzed scenarios the *Elastic* SA policy outperforms the *Semi-Elastic* SA policy which, on the other hand, performs better than *Fixed* SA.

In Table IV, the performance gain of elastic SA policies with respect to the fixed spectrum allocation for the network load which results in 1% of un-served bit-rate is reported in detail. Interestingly, *Elastic* SA allows serving up to 21% of traffic more than *Fixed* SA in the TEL100 scenario with 400 Gb/s lightpath capacity limit. In the same scenario, the gain when using the *Semi-Elastic* SA policy is just over 9%.

C. *Study of lightpath variability*

It is interesting analyzing the relationship between traffic aggregation level and the variability of each lightpath along the lime. To this end, the set of operating lightpaths is divided into three subsets, named L1 to L3, depending on the degree of bit-rate variations in time:

TABLE V DISTRIBUTION OF LIGHTPATHS ACCORDING TO THE LEVEL OF VARIABILITY

Max. Capacity	Lightpath type	TEL21	TEL60	TEL100
100 Gb/s	L1	19.0%	15.2%	14.3%
	L2	34.9%	41.3%	42.9%
	L3	46.2%	43.5%	42.9%
400 Gb/s	L1	0.0%	0.0%	0.0%
	L2	94.8%	98.0%	100.0%
	L3	5.2%	2.0%	0.0%

- L1 contains all those lightpaths whose bit-rate variations are almost unappreciable, and thus, the amount of spectrum resources is constant. This is mainly the case of full-loaded lightpaths grooming several demands, where statistical multiplexing keeps constant the total bit-rate at maximum.
- L2 includes lightpaths where bit-rate variations are dramatic within some range of minimum and maximum values. While aggregated traffic can reach the maximum bit-rate at a certain time, a minimum amount of traffic is always present.

- L3 contains those lightpaths where the highest relative bit-rate fluctuations are observed, changing instantly from not being used to a peak where several spectrum slices are required.

In Table V, the distribution of lightpaths belonging to each of the above defined subsets is detailed. These results are consistent with the discussion carried out in the analysis of the efficiency of SA policies in Sec. III.B and similar arguments can be used.

Firstly, there is a higher percentage of lightpaths of type L1 in the scenario with low lightpath capacity limit (100 Gb/s) compared to the scenario where 400 Gb/s lightpaths are allowed. This fact considerable reduces the amount of lightpaths needing elasticity, and thus the performance of the fixed is close to the elastic SA policies. That percentage decreases if the aggregation level is decreased (see TEL100 vs. TEL21).

Secondly, when 400 Gb/s lightpaths can be used, virtually all lightpaths belong to set L2, being just a small fraction of lightpath in set L3. Hence, lightpath capacity is always subject to changes, there are no saturated lightpaths, and only some lightpaths experience high capacity variations. On the other hand, the traffic fluctuations of lightpaths when the limit of 100 Gb/s is applied are much higher (see the percentage of lightpaths in set L3); especially, in the TEL100 scenario in which traffic aggregation is low and traffic variability is higher.

IV. CONCLUSIONS

In this paper, we have focused on dynamic adaptation of lightpath connections, by means of elastic spectrum allocation, in a flexgrid-based elastic optical network with time-varying traffic demands. In details, we have addressed a scenario in which an EON core network connects a number of IP/MPLS metro area networks performing traffic aggregation. As we found in our previous works, a large core network connecting a large set of small aggregation networks is a cost-effective solution. However, if the core network is extended towards the edges and the number of aggregation networks is increased then the level of traffic aggregation is decreased and the variability of traffic to be carried over the core network is higher.

To deal with that variability, in this paper we have proposed to make use of elastic spectrum allocation, which translates to the adaptation of spectrum allocated to lightpath connection in response to changes in traffic demands. We have analyzed two spectrum allocation policies, namely, a symmetric SA policy (referred to as *Semi-Elastic* SA) and asymmetric SA policy (referred to as *Elastic* SA), and compared their performance vs. a *Fixed* SA policy, which does not allow for spectrum changes. For each elastic SA policy, we developed a dedicated lightpath adaptation algorithm. The evaluation has been performed in a simulation environment assuming different network scenarios, which are characterized by different levels of traffic aggregation and lightpath capacity limits. We have shown that the effectiveness of lightpath adaptation in dealing with time-varying traffic highly depends on both the aggregation level and the maximum lightpath capacity. In a network with low traffic aggregation, the best performance gap of about 21% (vs. the *Fixed* SA policy) has been achieved for *Elastic* SA operating with high lightpath capacity limits.

It was shown that when core networks are extended towards the edges, the capacity of the lightpaths could be limited to deal with longer distances. In such scenarios, the effectiveness of elasticity is rather limited. Therefore, as a final remark, those networks where the majority of lightpaths are some few hundred kilometers long (e.g. national European core networks) can currently take advantage of elasticity. On the opposite, additional research to extend the reach of 400 Gb/s (or even higher) lightpaths is needed so as to larger core networks can benefit also from elasticity.

ACKNOWLEDGMENTS

This work has been partially funded by the European Community's Seventh Framework Programme FP7/2007-2013 under grant agreement n° 317999 IDEALIST project and from the Spanish Science Ministry through ELASTIC (TEC2011-27310) project. The work of M. Klinkowski was supported in part by NCN under Grant DEC-2011/01/D/ST7/05884.

REFERENCES

[1] O. Gerstel *et al.*, “Elastic optical networking: A new dawn for the optical layer?,” *IEEE Comm. Mag.*, vol. 50, no. 2, 2012.

[2] M. Jinno *et al.*, “Spectrum-efficient and scalable elastic optical path network: Architecture, benefits, and enabling technologies,” *IEEE Comm. Mag.*, vol. 47, no. 11, 2009.

[3] M. Jinno *et al.*, “Multiflow Optical Transponder for Efficient Multilayer Optical Networking,” *IEEE Commun Mag.*, vol. 50, no. 5, 2012.

[4] ITU-T Recommendation G.694.1 (ed. 2.0), “Spectral grids for WDM applications: DWDM frequency grid,” Feb. 2012.

[5] D. J. Geisler *et al.*, “The first testbed bemonstration of a flexible bandwidth network with a real-time adaptive control plane,” in *Proc. of ECOC*, Geneva, Switzerland, Sep. 2011.

[6] F. Cugini *et al.*, “Demonstration of flexible optical network based on path computation element,” *J. Lightw. Technol.*, vol. 30, no. 5, 2012.

[7] L. Velasco *et al.*, “Designing National IP/MPLS Networks with Flexgrid Optical Technology,” in *Proc. of ECOC 2012*.

[8] L. Velasco *et al.*, “How National IP/MPLS Networks can benefit from Flexgrid Optical Technology?,” accepted in *OFC 2013*.

[9] L. Velasco, et al., “Elastic Spectrum Allocation for Variable Traffic in Flexible-Grid Optical Networks,” in *Proc. OFC 2012*.

[10] M. Klinkowski *et al.*, “Elastic spectrum allocation for time-varying traffic in flexgrid optical networks,” *IEEE J. Sel. Areas Commun.*, vol. 31, pp. 26-38, 2013.

[11] G. Shen *et al.*, “Maximizing time-dependent spectrum sharing between neighbouring channels in CO-OFDM optical networks,” in *Proc. ICTON 2011*.

[12] K. Christodoulopoulos *et al.*, “Dynamic bandwidth allocation in flexible OFDM-based networks,” in *Proc. OFC 2011*.

[13] K. Christodoulopoulos *et al.*, “Time-Varying Spectrum Allocation Policies in Flexible Optical Networks,” *IEEE J. Sel. Areas Commun.*, 2013 (to appear).

[14] A. Klekamp *et al.*, “Transmission Reach of Optical-OFDM Superchannels with 10-600 Gb/s for Transparent Bit-Rate Adaptive Networks,” in *Proc. of ECOC 2011*.

[15] L. Velasco *et al.*, “Modeling the routing and spectrum allocation problem for flexgrid optical networks,” *Phot. Netw. Commun.*, vol. 24, no. 3, 2012.

[16] M. Ruiz, et al., “A Traffic Intensity Model for Flexgrid Optical Network Planning under Dynamic Traffic Operation,” accepted in *OFC 2013*.

[17] OMNET++, http://www.omnetpp.org/