Using Spectrum Fragmentation to Better Allocate Time-Varying Connections in Elastic Optical Networks

Pouria Sayyad Khodashenas, Jaume Comellas, Salvatore Spadaro, Jordi Perelló, and Gabriel Junyent

***Abstract*—Elastic optical network (EON) technology arises as a promising solution for future high-speed optical transport, since it can provide superior flexibility and scalability in the spectrum allocation for seamlessly supporting diverse services, while following the rapid growth of Internet traffic. This work focuses on lightpath adaptation under time-variable traffic demands in EONs. Specifically, we explore the possibility of utilizing the spectral fragmentation to increase the spectrum allocation (SA) capabilities of EONs. In this context, a heuristic SA algorithm, which intentionally increases the spectral fragmentation in the network, is proposed and validated. In our proposal, the spectrum assigned to each new connection is in the middle of the largest free spectral void over the route, aiming to provide considerable spectral space between adjacent connections. These free spectral spaces are then used to allocate time-varying connections without requiring any lightpath reallocation. The obtained simulation results show a significant improvement in terms of network blocking probability when utilizing the proposed algorithm.**

***Index Terms*—Elastic optical network; Routing and spectrum allocation; Time-varying connections.**

I. Introduction

Triggered by emerging services such as high-definition video distribution or social networking, IP traffic volume has been exponentially increasing to date. Furthermore, the traffic growth rate will not stop here due to the day-by-day technology advances. For example, new hardware advances, such as multicore processing, virtualization, and network storage will support a new generation of e-science and grid applications requesting data flows of 10 Gb/s up to terabit level. The predictable consequence is that network operators will require a new generation of optical transport networks in the near future, so as to serve this huge and heterogeneous volume of traffic in a cost-effective and scalable manner [1]. In response to these large capacity and diverse traffic granularity needs of the future Internet, the elastic optical network (EON) architecture has been proposed [2]. By breaking the fixed-grid spectrum allocation (SA) limit of conventional wavelength division multiplexing (WDM) networks, EONs increase the flexibility in the connection provisioning. To do so, depending on the traffic volume, an appropriately sized optical spectrum portion is allocated to a connection in EONs. Furthermore, unlike the rigid optical channels of conventional WDM networks, a lightpath can expand or contract elastically to meet different bandwidth demands in EONs [3]. In this way, incoming connection requests can be served in a spectrum-efficient manner.

This technological advance poses additional challenges on the networking level, specifically on the efficient connection establishment. Similar to WDM networks, an elastic optical connection must occupy the same spectrum portion between its end nodes, that is, ensuring the so-called spectrum continuity constraint. In addition, the entire bandwidth of the connections must be contiguously allocated, also referred to as the spectrum contiguity constraint. The routing and spectrum allocation (RSA) problem in EONs has been widely studied, putting more emphasis on dynamic network scenarios [4–7]. There, connection arrival and departure processes are random and the network has to accommodate incoming traffic in real time. Considering the near future technology advances (e.g., high-capacity bandwidth variable transponders) and the exponential increase of network traffic, it is foreseeable to have large intervals between the establishment and the release of connections. Thus, during such relatively long periods, the bit-rate demand of any connection may fluctuate following short- and mid-term traffic variations. Although the EON technology enables flexible adaptation of connections to such time-varying traffic demand changes, the spectrum fragmentation issue prevents high resource utilization in EONs. In light of this, various spectrum defragmentation proposals, requiring the periodic or on-demand reallocation of connections, have been introduced [8,9]. The basic idea behind these proposals is to consolidate fragmented spectrum bands, aiming to increase the probability of finding sufficient contiguous spectrum to accommodate future connection requests over the network. Despite the fact that spectrum defragmentation can improve the network efficiency under time-varying traffic, connection reallocation increases the complexity and cost of the network. Moreover, advanced optical devices such as wavelength converters and tunable optical

Manuscript received November 18, 2013; revised January 10, 2014; accepted March 13, 2014; published April 11, 2014 (Doc. ID 201518).

The authors are with the Advanced Broadband Communications Center (CCABA), Universitat Politècnica de Catalunya (UPC), Jordi Girona 1-3, 08034 Barcelona, Spain (e-mail: pkhodashenas@tsc.upc.edu).

http://dx.doi.org/10.1364/JOCN.6.000433

switches are necessary for realizing such proposals, further increasing the complexity and also the cost of the network. Finally, it is worth mentioning that a significant amount of control plane signaling messages are required to set up and release lightpaths in the network, which burdens the control plane in the network with additional overhead.

While spectral voids between adjacent connections (due to fragmentation) are traditionally considered as a problem, this statement does not always apply. Indeed, as long as connections' required bandwidth may grow over time, free spectral voids are crucial to accommodate additional bandwidth demands without requiring the reallocation of the already established connections (an existing connection can easily adapt to transmission rate fluctuations if it has free spectral voids around it). Therefore, an alternative approach could consist in deliberately leaving space between connections. In light of this, we propose a heuristic SA algorithm that intentionally increases spectral fragments over the network in order to boost time-varying connection allocation. In our proposal, each new elastic connection is allocated in the middle of the largest spectral void over the selected end-to-end route. As a result, the possibility of serving short- to mid-term bit-rate fluctuations increases significantly. The rest of paper is organized as follows. In Section II we initially review the principles of serving time-varying connections in EONs. Section III details the proposed algorithm. Simulation results are presented in Section IV. Finally, Section V concludes the paper.

II. Elastic Optical Networks and Time-Varying Traffic

In this section, we initially introduce the problem of serving time-varying connections in EONs and review some previous contributions on this topic. Then, applicable lightpath adaptation policies are also introduced.

A. Time-Varying Traffic

To accommodate traffic in EONs, the total available spectrum is divided into constant spectrum units with a granularity finer than the typical 50 GHz grid used in WDM systems (e.g., 12.5 GHz), referred to as frequency slots (FSs). For example, 1, 1.5, and 2 THz spectra correspond to 80, 120, and 160 FSs, respectively. Each FS can carry some bit rate depending on the modulation format used [10]. For the sake of simplicity, we will assume a modulation efficiency of 1 b/s/Hz in this work. A connection is served by assigning a route and allocating a set of contiguous FSs on all links along it.

Initially, the problem of static RSA received the research community's attention. In this scenario, all the requested connections are given, as a traffic matrix known in advance. Then, connections are served under the constraint that no spectrum overlapping is allowed among them. The solution gives the route and the allocated FSs to every connection, while minimizing a given performance metric, such as the total utilized spectrum. Some studies such as those in [11–13] focus on the routing, modulation level, and spectrum allocation (RMLSA) in EONs. They formulated the problem and proved that it is NP-complete. Integer linear programming (ILP) formulations, as well as heuristic algorithms for solving such a complex planning problem, were presented in these studies.

While minimizing the total required spectrum for serving a static traffic matrix is important in the planning phase of EONs, there is an increasing need for dynamic lightpath provisioning, as traffic fluctuates over time in current networks. Thus, the community's attention also focused on the dynamic SA problem. The authors of [4] triggered this effort by proposing a general algorithm to assign FSs in EONs under dynamic traffic. Meanwhile, they proposed the first distance adaptive dynamic routing and FS assignment in [5], concluding that the necessary bandwidth for serving connections can be significantly reduced by employing a distance adaptive proposal. In [14], the dynamic allocation and release of connections ranging from 10 to 400 Gb/s were investigated. This study shows that by introducing technologically advanced devices, such as high-capacity flexible optical transponders, the holding time (HT) of the optical connections in EONs increases significantly. During such a long period, the connection bit rate may vary as a function of time. Thus, the spectrum allocated to the connection (lightpath) has to flexibly change, so as to adapt to the variation in the requested bit rate. Figure 1 presents the spectrum utilization of an exemplary link in an EON with time-varying connections at two different moments. As shown in Fig. 1(a), connections over the exemplary link carry traffic between different end nodes. A specific set of FSs (which are highlighted in the figure) is assigned to each connection. Figure 1(b) shows the spectrum utilization of the same link at another instant, where the bandwidth required by connection 2 has doubled. As shown, the number of FSs allocated to this connection has changed to adapt to the required transmission rate. It is worthwhile to note that the free FSs between adjacent connections are used to accommodate the mentioned traffic change. There exist different policies for such a lightpath adaptation in the literature [15–17]. We will describe the one used in this paper in the next subsection.

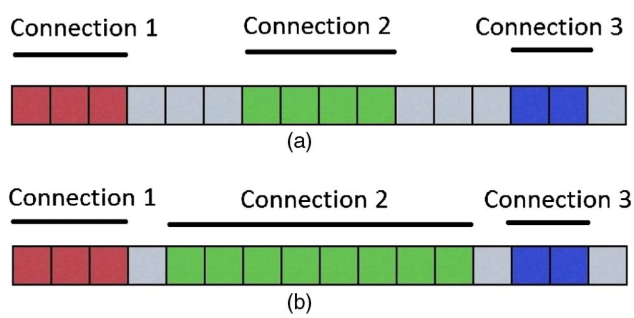

Fig. 1. Spectrum allocation of an exemplary link with time-varying connections. Two different time instances are displayed in (a) and (b). Free frequency slots between adjacent connections are used to accommodate time-varying traffic changes.

B. Lightpath Adaptation Policies

The way connections adapt their spectrum to the instantaneous required bit rates is called the lightpath adaptation policy. The authors of [15] and [16] studied for the first time different policies through extensive simulation results. They continued their work in [17], by introducing some mechanisms to the previously proposed policies, aiming to strike a balance between network blocking probability (BP) and the number of reallocated connections. These policies can be summarized as follows:

- Constant spectrum allocation (CSA): A fixed number of FSs are reserved around each connection. No spectrum sharing is permitted among adjacent connections. No connection reallocation is permitted.
- Dynamic alternate direction (DAD): A connection wishing to increase its bit rate alternates between using its higher and lower FSs, until it reaches an already occupied slot. If additional slots are needed, the symmetry of the expansion is lost, but the connection continues to expand toward the only possible direction. No connection reallocation is permitted.
- Avoid close neighbors (ACN): Each connection expands toward the opposite direction of its closest neighboring connection on any of the links along its path. Keeping the same principle, contraction is conducted from the direction of the closest neighboring connection. No connection reallocation is permitted.
- Shift-ACN: This is an enhanced version of ACN that allows connection reallocation. If there are no available FSs in the maximum allowable expansion region, neighboring connections are shifted in a way that maximizes the minimum available FS among all of its neighbors. Note that connection reallocation is only allowed to direct neighboring connections of the connection requesting a spectrum expansion.
- Float-ACN: This is an enhanced version of Shift-ACN. In this version, all connections are free to float in the spectrum as they are pushed by their neighbors. Thus, connection reallocation is not only restricted to the direct neighboring connections of the connection requesting a spectrum expansion, as it was in the previous case.

It is worthwhile to mention that, to avoid traffic interruption during the extension process, hitless techniques such as the push–pull technique in [18,19] are applied.

Considering the aforementioned cases, we propose an enhanced version of the DAD policy, namely the Shift-DAD policy, in this article. Note that, since the connections can contract without any extra efforts, here we explain only the extension part of this policy.

Shift-DAD policy:

1) Calculate the necessary number of FSs to satisfy the connection's bandwidth change. Set this number as the required FSs.
2) Count the number of free FSs at both sides of the connection that needs to expand. In the case of having fewer free FSs than the required FSs go to Step 3; otherwise go to Step 4.
3) Check the spectral status of the path to find the first void with enough FSs to accommodate the connection. If found, reallocate it. Otherwise, do not extend the connection and add the amount of bandwidth change into the total bandwidth dropped. **Finish**.
4) Check the left side of the connection for a free FS; if it exists, extend the connection by one FS toward the left. Decrease the required FSs by one.
5) If the required FSs are higher than zero, check the right side of connection for a free FS; if it exists, extend the connection by one slot toward the right. Decrease the required FSs by one.
6) If the required FSs are still higher than zero, go back to Step 4. Otherwise, **Finish**.

Similar to the DAD policy, a connection wishing to increase its bit rate alternates between using its higher and lower FSs, until reaching an already occupied slot. If additional slots are needed, the symmetry of the expansion is lost but the connection continues to expand toward the only possible direction. When a connection needs more FSs than those available around it, the connection is reallocated to the spectrum void with enough bandwidth along the connection's path. Note that, if a void large enough to reallocate the connection is not found along the path, the entire bandwidth expansion operation is blocked. An alternative to this could have been to permit partial satisfaction of the bandwidth expansion. Nonetheless, this has been left for future study.

In practical cases that use conventional SA algorithms such as First Fit, it happens that connections are established very close to each other over the network, aiming to increase the available spectrum for new arriving connections. With these approaches, it is hard to find enough bandwidth around a lightpath to adapt it to its traffic fluctuation. Meanwhile, since the occupied portion of the spectrum is compacted in one side of the available spectrum (lower side of spectrum), it is possible to have enough FSs for serving time-varying connections in the upper side of the spectrum. Therefore, by reallocating connections over the network in a hitless fashion, a significant improvement in time-varying traffic allocation can be achieved.

Despite its potential benefits, connection reallocation increases the complexity and the cost of the network. As was stated before, many advanced photonic devices are needed, as well as a large amount of control plane overhead. In light of this, it is important to make a good trade-off between the successfulness in allocating time-varying connections and network complexity and cost efficiency.

Moreover, it is worthwhile to highlight that SA algorithms have a great impact on the efficacy of serving time-varying connections in EONs. Indeed, SA algorithms can achieve good arrangement of the connections in the space and frequency domains, thus increasing the possibility of serving and expanding time-varying connections in the network. An appropriate SA algorithm serves the connection requests, by assigning FSs in a way that minimizes the average network BP, taking into account the requirements of the lightpath adaptation policy. To achieve the previously mentioned goals, the next section details a novel SA algorithm called Mid Fit.

III. Spectrum Assignment Algorithm

There exist various SA approaches to allocate incoming connections with different bandwidth granularities in EONs [20]:

- First Fit: The connection request is placed in the first spectral gap along its route fitting the requested bandwidth. This algorithm is used for benchmarking purposes in this work.
- Smallest Fit: The connection request is placed in the smallest available spectral band along its route.
- Random Fit: Any available spectrum portion along the route, with enough space to allocate the connection request, can be randomly selected to allocate it.

In contrast, in this paper we propose a novel heuristic called the Mid Fit SA approach. Its main objective is to maximize the spectral voids between adjacent connections, thus increasing the possibility to successfully serve their potential bandwidth increments. To achieve this goal, after calculating the candidate path between source and destination nodes of the connection request, this one is established in the middle of the largest contiguous spectrum void along the calculated path. Specifically, if multiple same-sized spectral bands exist, the one placed in the lowest part of the spectrum is selected. In this way, the spectral distance between adjacent connections is maximized. The available spectrum left between connections can later on be used to serve connection bandwidth fluctuations. It is worthwhile to mention that lightpath bit-rate adaptation can then be realized without any spectrum conversion or reallocation. This implies a significant reduction in the network cost, complexity, and energy consumption. In fact, bit-rate fluctuations are allocated by just increasing the bandwidth (number of FSs) assigned to the connection (assuming that the total connection bit rate is lower than the whole capacity of the transmitter).

Figure 2 illustrates a simple four-node network, whose current spectral status is shown in Fig. 2(a). There exist some connections with different path lengths (a 2 FS connection from D to B and a 1 FS connection from B to C). If a new connection request arrives, for example a 1 FS connection between nodes A and C, the candidate path for establishing the connection is first calculated. Let us assume that this path consists of links AB and BC. The spectral status of this candidate path is shown in Fig. 2(b). As illustrated, there exist two spectral voids in the spectrum of the candidate path, G1 and G2. According to our proposal, the Mid Fit algorithm assigns the central FS of G2 to the connection, as shown in Fig. 2(c). The free FSs between existing connections are used to serve data rate fluctuations. By applying this spectral allocation strategy the bit rate of the connection can be doubled and even tripled

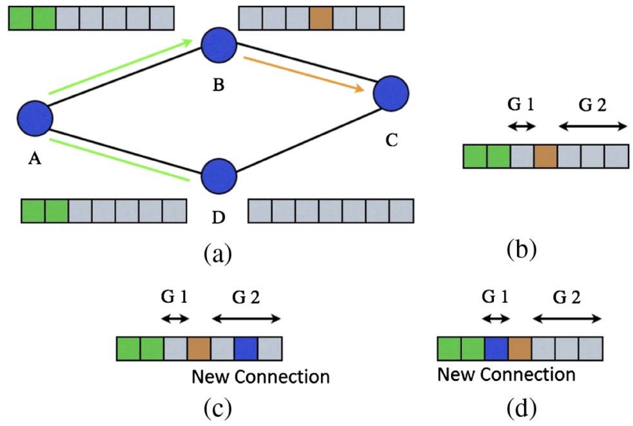

Fig. 2. Illustrative example: (a) simple four-node network, (b) spectral status of the candidate path, (c) FS status using the Mid Fit SA approach, and (d) FS status using conventional First Fit SA.

in the future without any need for connection reallocation. That is, bit-rate fluctuations are accommodated in the network by only increasing the bit rate of transmitters (assuming that the final bit rate is less than the whole capacity of the transmitter).

In contrast, let us consider a conventional SA algorithm such as First Fit for SA in this example. After calculating the candidate path, the first gap with enough FSs (G1) is selected for establishing the new connection request, as shown in Fig. 2(d). Hence, since there is no space between adjacent connections, there is no chance for serving the connection data rate fluctuations without reallocating it. Indeed, it could be possible to accommodate the connection in the network upon expansion (i.e., over spectral void G2). However, reallocation would have to be unavoidably employed in this case.

IV. Simulation Results

We evaluate the performance of the proposed Mid Fit SA approach through extensive discrete event simulation studies and compare its performance with a First Fit SA [10] in a scenario with time-varying elastic connections. In both cases, a k-shortest path routing strategy, with $k = 3$, is used. During the simulations, existing connections are allowed to change their bit rate, so that the spectrum they use can be dynamically expanded and contracted accordingly. Time-varying connection selection starts after an initial transient phase, in which 10^4 connection requests are simulated (i.e., enough to achieve steady-state network operation). Once this initial transient phase ends, a total number of 5×10^6 connection requests are simulated, which allows us to get statistically relevant BP results. Furthermore, a percentage of the established connections are randomly selected and triggered to change their bandwidth. This percentage of time-varying connections with respect to the total of established connections allows observing the relationship between traffic fluctuation frequency and network performance. In particular, we have assumed that 15% of the served connections change their bandwidth during their lifetime. To increase the fairness of the comparison, the Shift-DAD policy proposed in Subsection II.B is considered along with the First Fit SA scenario for benchmark purposes (i.e., a scenario in which connection reallocation is allowed). That is, if it is not possible to accommodate the traffic growth in the spectrum band originally assigned to a connection, it can be reallocated to the first spectrum void with enough capacity. It is worthwhile to mention that, since avoiding connection reallocations is the main purpose of Mid Fit SA, only a basic DAD adaptation policy as in Ref. [17] is allowed in this case. In addition, in this paper, we consider BP as one of the evaluation metrics. It represents the overall failed attempts to set up new connection requests and bandwidth changes of already established connections. The 14-node with 21 bidirectional links NSFnet reference topology has been selected for the evaluation (Fig. 3). We assume a total optical spectrum of 1.5 THz per link, which is discretized in FSs of 12.5 GHz. In addition, according to the asymmetric nature of today's Internet traffic, unidirectional connections between source and destination nodes are considered. To appropriately model the offered traffic granularity (i.e., average number of offered connections per node and average bit rate per connection), the traffic generation follows a Poisson distribution process, so that different offered load values (average number of connection requests per node) are obtained by keeping the mean HT of the connections constant to 200 s, while modifying their inter-arrival time (IAT) accordingly (i.e., offered load = HT/IAT). Connection bit rates are randomly generated following a log-normal distribution over the range from 12.5 (1 FS) to 125 Gb/s (10 FSs) [8]. As mentioned before, for the sake of simplicity, a spectral efficiency of 1 b/s/Hz has been considered (realizable with a simple BPSK modulation format).

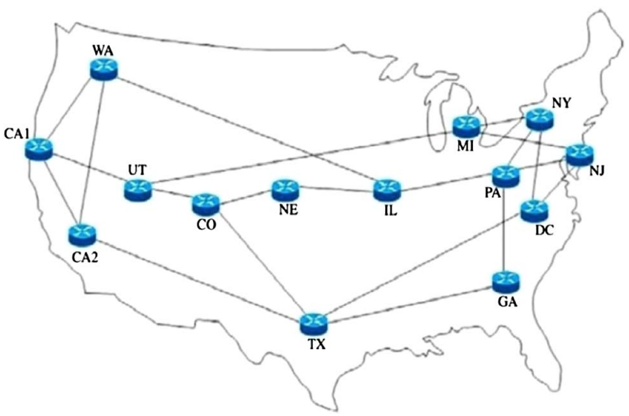

Fig. 3. NSFnet network topology: 14 nodes and 21 fiber links.

The first results are shown in Fig. 4, where the average number of offered connections per node ranges from 8 to 15. As the average bit rate per connection is 35 Gb/s, the total average traffic generated per node ranges from 280 to 525 Gb/s. Concerning the bandwidth variation of the selected time-varying connections, we assume that their bandwidth can either be doubled or halved with 50% probability (for the sake of simplicity). As shown in Fig. 4(a), the First Fit SA approach with connection reallocation outperforms the other approaches along the entire offered traffic range. The notable differences between First Fit SA with and without connection reallocation are due to the fact that with First Fit SA most of the connections need reallocation upon bandwidth expansion. Indeed, when the First Fit SA approach is used, connections are allocated very close to each other; thus the possibility of finding free spectral resources around them is very low. As a result, many of them have to be moved from their original band. To highlight this, the percentage of reallocated connections for each offered traffic value is shown in Fig. 4(b). As seen, more than 60% of the connections experiencing bandwidth expansion must be reallocated when First Fit SA is used. Therefore, the benefits of First Fit with connection reallocation on the BP are achieved at the expense of a large number of reallocations, which are both complex and costly processes. In contrast, an alternative way to achieve significant benefits with low complexity is the proposed Mid Fit SA approach. In this case, around one order of magnitude

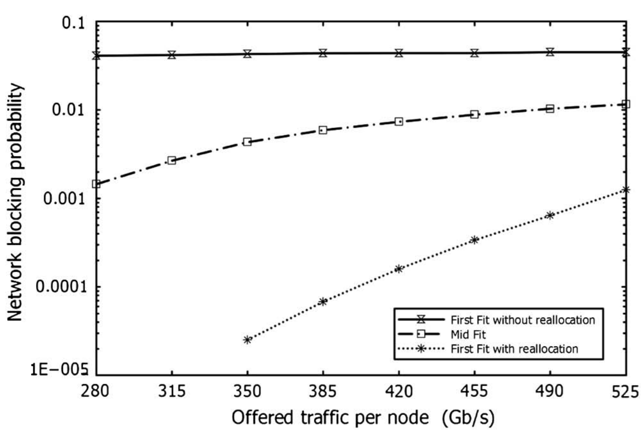

(a)

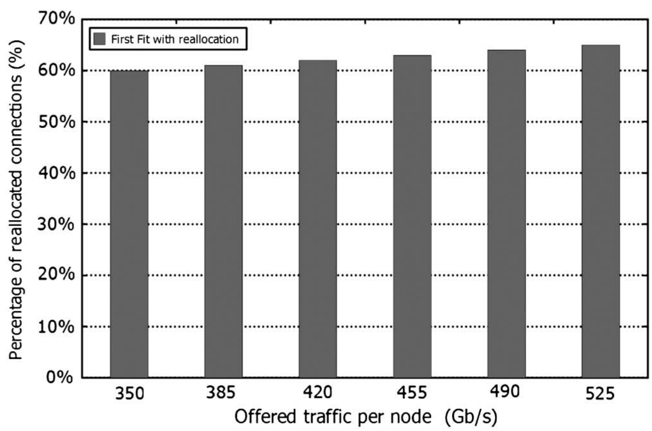

(b)

Fig. 4. (a) Network BP for different offered traffic per node values. The average number of offered connections per node changes from 8 to 15 while the average bit rate per connection is 35 Gb/s. (b) Percentage of reallocated connections versus offered traffic per node (only connections requiring expansion are taken into account to compute this percentage).

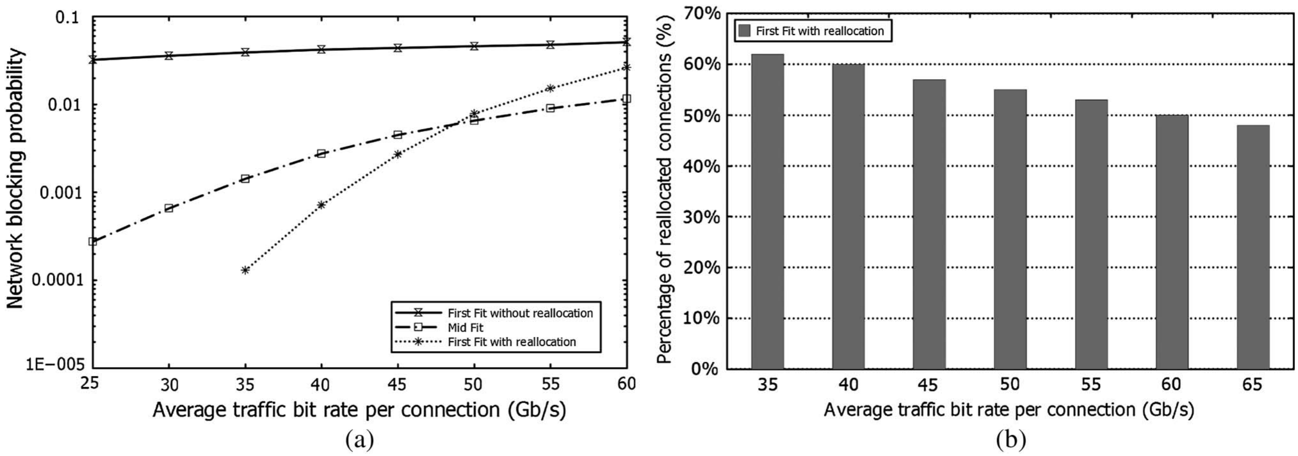

Fig. 5. (a) Network BP for different average bit rate per connection values. The average number of offered connections per node is fixed to 12, while the average bit rate per connection changes from 25 to 60 Gb/s. (b) Percentage of reallocated connections versus average bit rate per connection (only connections requiring expansion are taken into account to compute this percentage).

improvement can even be achieved for low loads with respect to the First Fit SA without reallocation, while no connection reallocation is needed. Not so pronounced but still significant benefits are observed for Mid Fit SA against First Fit SA without reallocation in highly loaded network scenarios.

To investigate more about the benefits of our proposal, similar studies have been done for a fixed offered load of 12 offered connections per node (168 connections offered to the entire network), while changing the average bit-rate demand per connection from 25 to 60 Gb/s (2–5 FSs). According to the mentioned load profile, the average traffic generated per node ranges from 300 to 720 Gb/s. Again the percentage of time-varying connections in the whole simulation is assumed to be 15%, and the bandwidth of each randomly selected connection can be either doubled or halved with 50% probability. As shown in Fig. 5(a), a significant reduction in BP can be achieved by allowing time-varying connections to freely shift in the spectrum (First Fit SA approach with connection reallocation). However, by increasing the average bit rate per connection, the possibility of finding enough spectrum resources for reallocating active connections is reduced. In fact, for average bit-rate values greater than 50 Gb/s the Mid Fit SA approach leads to better BP performance in the network. Moreover, this performance is achieved in a simpler and more cost-effective manner, since no control-plane-driven reallocation is triggered. Figure 5(b) shows the percentage of reallocated connections for each average bit-rate value. As shown, by increasing the connection bit rate, the percentage of reallocations connections decreases, which verifies the abovementioned effect.

In the next study, the impact of the percentage of time-varying connections on the EON performance is evaluated. The offered load is fixed to 12 connections per node, and the average bit-rate demand per connection is equal to 35 Gb/s. As illustrated in Fig. 6, Mid Fit SA improves the performance of the network in terms of BP in the whole range of the simulation when compared to First Fit SA without reallocation. Considering First Fit SA with reallocation, as illustrated, its efficacy in serving time-varying connections decreases as traffic fluctuation occurs more frequently in the network. For percentages greater than 30% the Mid Fit SA approach leads to better BP performance in the network. Indeed, by increasing the number of time-varying connections in the network, the possibility to successfully reallocate a connection decreases. Therefore, worse BP performance using First Fit SA with reallocation against Mid Fit SA is observed.

Looking for a way to approach the performance of Mid Fit SA to that of First Fit SA with connection reallocation, we target at providing some extra spectrum to the Mid Fit SA scenario. This would allow finding larger voids and, thus, decreasing the BP due to lack of spectrum upon bandwidth expansion. It is noteworthy here that a strategy like this does not necessarily entail higher network CAPEX, as a network operator may be under-utilizing the entire 5 THz C-band bandwidth. Furthermore, if the number of offered connections to the network remains constant, the number of devices that must be equipped in the network nodes (e.g., transponders) also remains unaltered. For this study, we consider an average number of offered connections per

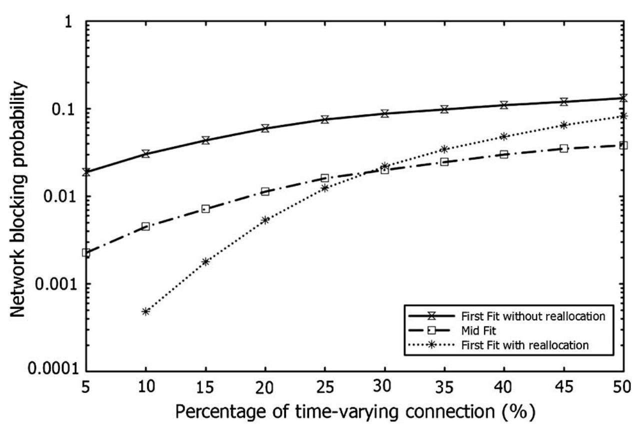

Fig. 6. Network BP versus percentage of time-varying connections.

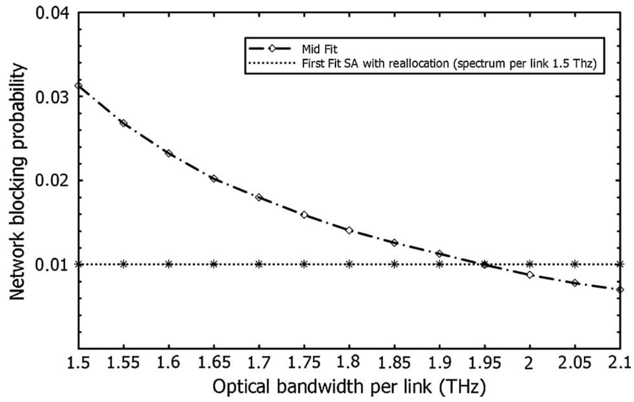

Fig. 7. Network BP versus total spectrum per link.

node equal to 20, with an average bit-rate demand of 35 Gb/s. Moreover, the percentage of time-varying connections in the whole simulation is assumed to be 15%, and the bandwidth of each selected time-varying connection can either be doubled or halved with 50% probability.

As shown in Fig. 7, such an amount of traffic leads to 1% BP with First Fit SA with reallocation if the total bandwidth per link is 1.5 THz. In contrast, BP rises up to 3% when Mid Fit SA is applied. Nevertheless, by increasing the spectrum from 1.5 to 1.95 THz the performance of Mid Fit SA shows no penalty compared to First Fit SA with reallocation. Hence, increasing the operational bandwidth of the fiber links by around 25%–30% allows Mid Fit SA to achieve the same BP performance as First Fit SA, but without any connection reallocation, which simplifies the network operation to a large extent. In view of this result, Mid Fit SA becomes an interesting option for EON operators that can reduce network complexity by dedicating some extra spectrum in their potentially underutilized fiber links.

V. Conclusion

EON technologies arise as promising solutions for future high-speed optical transmission, since they can provide superior flexibility and scalability in SA toward the seamless support of diverse services along with the rapid growth of Internet traffic. In this paper, we focused on lightpath adaptation under variable traffic demands in EONs. The possibility of utilizing spectral fragmentation for increasing the elastic SA capability of EONs has been explored. We proposed a heuristic SA algorithm, called the Mid Fit approach, to intentionally increase spectral fragmentation in the network. In this approach, the spectrum dedicated to a connection is in the middle of the largest possible free spectral void over the route, providing greater spectral resource between adjacent connections. These spectral spaces are used for dynamic expansions of lightpaths, so as to adapt them to the time-varying required transmission rate. By means of simulation, it has been demonstrated that such a proposal can serve time-varying connections in a simple and cost-efficient manner.

Acknowledgment

This work has been supported by the Government of Catalonia and the European Social Fund through a FI-AGAUR research scholarship grant and by the Spanish National project ELASTIC (TEC2011-27310).

References

[1] "Cisco Visual Networking Index: Forecast and Methodology, 2012–2017," Cisco White Paper [Online]. Available: http://www.cisco.com/c/en/us/solutions/collateral/service-provider/ip-ngn-ip-next-generation-network/white_paper_c11-481360.html.

[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] M. Jinno, H. Takara, and B. Kozicki, "Dynamic optical mesh networks: Drivers, challenges and solutions for the future," in *Proc. ECOC*, Sept. 2009, paper 7.7.4.

[4] K. Christodoulopoulos, I. Tomkos, and E. Varvarigos, "Dynamic bandwidth allocation in flexible OFDM-based optical networks," in *Proc. OFC*, 2011, paper OTuI5.

[5] T. Takagi, H. Hasegawa, K. Sato, Y. Sone, B. Kozicki, A. Hirano, and M. Jinno, "Dynamic routing and frequency slot assignment for elastic path networks that adopt distance adaptive modulation," in *Proc. OFC*, 2011, paper OTuI7.

[6] A. Castro, L. Velasco, M. Ruiz, M. Klinkowski, J. Fernández-Palacios, and D. Careglio, "Dynamic routing and spectrum (re)allocation in future flexgrid optical networks," *Comput. Netw.*, vol. 56, no. 12, pp. 2869–2883, 2012.

[7] A. Patel, P. Ji, J. Jue, and T. Wang, "Routing, wavelength assignment, and spectrum allocation in transparent flexible optical WDM (FWDM) networks," in *Proc. Photonics in Switching*, 2010, paper PDPWG1.

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

[9] P. S. Khodashenas, J. Comellas, J. Perelló, and S. Spadaro, "Correlation between traffic granularity and defragmentation periodicity in elastic optical networks," *Trans. Emerg. Telecommun. Technol.*, to be published.

[10] K. Christodoulopoulos, I. Tomkos, and E. A. Varvarigos, "Routing and spectrum allocation in OFDM-based optical networks with elastic bandwidth allocation," in *Proc. IEEE Globecom*, 2010, paper ONS06.

[11] T. Takagi, H. Hasegawa, K. Sato, T. Tanaka, B. Kozicki, Y. Sone, and M. Jinno, "Algorithms for maximizing spectrum efficiency in elastic optical path networks that adopt distance adaptive modulation," in *Proc. ECOC*, 2010, paper We.8.D.5.

[12] K. Christodoulopoulos, I. Tomkos, and E. Varvarigos, "Spectrally/bitrate flexible optical network planning," in *Proc. ECOC*, 2010, paper We.8.D.3.

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

[14] A. Patel, P. Ji, J. Jue, and T. Wang, "Defragmentation of transparent flexible optical WDM (FWDM) networks," in *Proc. OFC*, 2011, paper OTuI8.

[15] K. Christodoulopoulos and E. Varvarigos, "Routing and spectrum allocation policies for time-varying traffic in flexible optical networks," in *Proc. ONDM*, 2012, pp. 1–6.

[16] K. Christodoulopoulos, I. Tomkos, and E. Varvarigos, "Time-varying spectrum allocation policies and blocking analysis in flexible optical networks," *IEEE J. Sel. Areas Commun.*, vol. 31, no. 1, pp. 13–25, 2013.

[17] E. Palkopoulou, I. Stiakogiannakis, D. Klonidis, K. Christodoulopoulos, E. Varvarigos, O. Gerstel, and I. Tomkos, "Dynamic cooperative spectrum sharing in elastic networks," in *Proc. OFC*, Mar. 2013, paper OTu3A.

[18] F. Cugini, M. Secondini, N. Sambo, G. Bottari, G. Bruno, P. Iovanna, and P. Castoldi, "Push-pull technique for defragmentation in flexible optical networks," in *Proc. OFC*, 2012, paper JTh2A.40.

[19] R. Wang and B. Mukherjee, "Provisioning in elastic optical networks with non-disruptive defragmentation," *J. Lightwave Technol.*, vol. 31, no. 15, pp. 223–228, 2013.

[20] A. Rosa, C. Cavdar, S. Carvalho, J. Costa, and L. Wosinska, "Spectrum allocation policy modeling for elastic optical networks," in *Proc. HONET*, 2012, pp. 242–246.

Pouria Sayyad Khodashenas received his B.Sc. degree in 2005 from the University of Guilan, Rasht, Iran, and his M.Sc. degree in 2008 from the University of Tabriz, Tabriz, Iran, both in electrical engineering. Since October 2009, he has been working toward a Ph.D. in telecommunications engineering at Universitat Politècnica de Catalunya (UPC, BarcelonaTech), Barcelona, Spain, being advised by Prof. Jaume Comellas. His research interests include all-optical communication systems, all-optical switches, nonlinear optics, quantum optics, multi-layer structures, and semiconductor lasers, especially VCSEL.

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

Salvatore Spadaro received M.Sc. (2000) and Ph.D. (2005) degrees in telecommunications engineering from Universitat Politècnica de Catalunya (UPC). He also received his Dr. Ing. degree in electrical engineering from Politecnico di Torino (2000). 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. His research interests are in the fields of all-optical networks with emphasis on network control and management, resilience, and network virtualization.

Jordi Perelló received his M.Sc. and Ph.D. degrees in telecommunications engineering in 2005 and 2009, respectively, from UPC. Currently, he is an Assistant Professor in the Computer Architecture Department (DAC) at UPC. He has participated in various IST FP-6 and FP-7 European research projects, such as the FP-7 projects STRONGEST, EULER, DICONET, and NoE BONE, as well as the FP-6 projects NOBEL Phase 2, e-Photon/One+, and COST Action 291. He has published more than 60 articles in international journals and conferences. His research interests concern resource management, quality-of-service issues, and survivability of future optical transport networks.

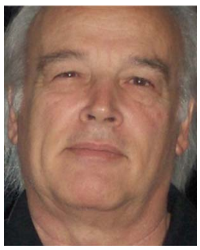

Gabriel Junyent graduated in telecommunication engineering from the Universidad Politécnica de Madrid, Spain, in 1973, and received a Ph.D. degree in communications from the Universitat Politècnica de Catalunya (UPC), Barcelona, Spain, in 1979. From 1973 to 1989, he was a Teaching Assistant and Associate Professor at UPC, where he has been a Full Professor since 1989. He founded the Advanced Broadband Communications Center (CCABA) at UPC in 1990. In the last 15 years he has participated in more than 30 national and international R&D projects. He is coauthor of more than 150 journal and conference papers.