

Proactive Defragmentation in Elastic Optical Networks under Dynamic Load Conditions

Jaume Comellas, Laura Vicario and Gabriel Junyent

Optical Communications Group, Universitat Politècnica de Catalunya,
UPC Campus Nord, c/Jordi Girona 1, 08034, Barcelona, Spain,
Email: comellas@tsc.upc.edu, Telephone: +34 934 016 441, Fax: +34 934 017 200
ORCID (Jaume Comellas): 0000-0002-9129-0562

ABSTRACT

The main weakness of Elastic Optical Networks (EON), under dynamic traffic conditions, stems from spectrum fragmentation. A lot of research efforts have been dedicated during recent years to spectrum defragmentation. In this work, a thorough study about proactive defragmentation is carried out. Effects of the different defragmentation parameters on the EON performance are analyzed and appropriate values of the defragmentation period, which guarantee suitable network performance while keeping the network control complexity at reasonable values, are obtained by means of extensive simulations. Benefit obtained by applying different defragmentation strategies, in terms of increase of the supported load at a given bandwidth blocking probability, is also reported. Different traffic conditions and network topologies are simulated to assess the validity of the obtained results.

Keywords: Elastic Optical Networks; Network Performance; Spectrum Defragmentation; Bandwidth Blocking Probability

1. Introduction

Dynamic establishment and tear-down of connections in Elastic Optical Networks (EON) leads inevitably to spectrum fragmentation [1]. Spectrum voids left after connections ending are scarcely re-used, as their size and position in the different network links are random. These unused spectrum voids lead to higher bandwidth blocking probability (BBP) at any network load value [2]. Spectrum defragmentation [3] becomes therefore a crucial process aiming to soften this situation by consolidating the used links spectrum as much as possible. Spectrum defragmentation strategies are usually classified as,

- Proactive: when a given parameter (usually time or network fragmentation) reaches a predefined threshold, the resources in use are re-arranged looking for a more spectrally optimal distribution of the existing connections. In this case, the defragmentation procedure is usually run periodically (e.g. every time that 100 connections previously established through the network have ended). It can be seen as a proactive measure to keep fragmentation (and therefore BBP) under certain level. Obviously, keeping BBP low depends ultimately on the load offered to the network, so when the load is too high, it becomes impossible since the network resources (nodes transmission capacity and links spectrum) are limited.
- Reactive: defragmentation is done on-demand when one connection, which has to be allocated in the EON, finds no available resources [4]. In this case, the defragmentation is reactive and takes place when a potential connection blocking is devised. It is advantageous in view of the fact that connections reallocation is only done when needed, so the number of changes in the EON and thus the complexity is reduced. On the other hand, its main objective is not related with the compactness of the spectrum but with just finding room for the next connection to be allocated, which otherwise would be blocked. This strategy can therefore lead to less-optimal spectrum usage.

As it is detailed in next section, previous works on EON defragmentation have dealt with path based reallocation of connections, which implies end-to-end rearrangement of the spectrum. This imposes additional constraints, and hence complexity to the defragmentation, because the problem to solve does not consist in isolatedly compacting the link spectrum but the whole sequence of links in the path. It has been demonstrated that the spectrum continuity and contiguity constraints [3] add considerable complexity to path based procedures in EONs. This complexity is observed in the exemplary 3 links path shown in Fig. 1, where 3 different connections exist (C1: 4 slots (FS) from 2 to 3 using FS 1-4, C2: 3 FS from node 1 to node 3 using FS 5-7 and C3: 2 FS from 1 to 2 using FS 8-9).

A key idea in defragmentation studies consists in avoiding connection disruption while consolidating the spectrum (hitless defragmentation, [3]). It is clear that keeping this restriction makes the problem extremely complex and unsolvable in many situations (such as the simple one shown in Fig. 1, where transition from the current state to the optimum one is impossible without disrupting at least one connection). Therefore, the use of extra resources to transitorily maintain connections which have to be re-allocated seems reasonable.

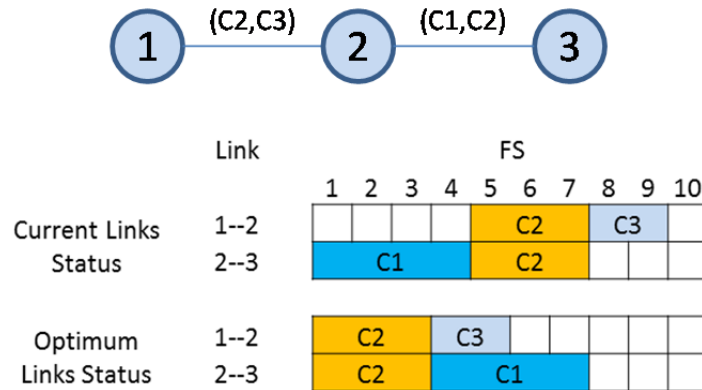


Fig. 1. Exemplary 3 links path (each link has 10 FS) where 3 connections (C1, C2 and C3) have been sequentially established using First-Fit spectrum allocation. Optimum allocation of the existing connections is shown in the lower table.

If connections disruption is allowed, C1 could be temporarily torn down (it has been previously allocated in an alternative path), then C2 and C3 moved to left, and finally re-establish C1 using FS 4 to 7 in link 2-3. Allowing only hitless defragmentation leads to a lock-in situation which is common when randomly generated connections have to be set-up in such a multi-hop scenario. Therefore, the usefulness of hitless defragmentation is limited, and alternative solutions using backup resources and temporarily disrupting some connections are unavoidable to keep BBP values under certain threshold. The focus of this work is not on achieving hitless defragmentation but on quantifying the effects of proactive defragmentation on the overall network performance. Indeed, the impact of allowing only hitless defragmentation is given as a noteworthy result that justifies the use of disrupting defragmentation procedures.

The study of different strategies for proactive defragmentation, as well as their quantification in terms of benefits obtained and added complexity, is the main contribution of this work. As traffic generation is dynamic, the links spectrum status during EON operation becomes completely random. Hence, the defragmentation algorithm designed is very simple [5], and its main objective consists in relocating connections from the right to the left side of the spectrum (First-Fit spectrum assignment is used during connections establishment). No routing changes are allowed, so only spectrum re-allocation is done throughout this work. Rerouting of existing connections usually implies the use of longer paths and therefore a penalty in the BBP is incurred in average. It is important to highlight that although there is an evident dependence on the EON performance regarding the topology and the offered traffic profile, the main objective of this work is to find significant trends that hold regardless of the diverse scenarios under study.

The remaining of this paper is organized as follows. In Section 2, a review of previously reported results from defragmentation studies is given. The defragmentation algorithm used in this work, as well as the different strategies used to select the set of connections to be reallocated, is detailed in Section 3. The main contributions of this

work stem from extensive simulations using typical conditions of EONs; hence, the simulation framework, as well as the most significant results obtained, is reported in Section 4. Finally, Section 5 concludes the paper.

2. Background

As commented above, spectrum fragmentation is one of the main drawbacks of EONs. Lots of research efforts have been therefore devoted to different ways to alleviate its effects. A first approach consists in preventing fragmentation during the routing and spectrum allocation phase (a.k.a. fragmentation aware RSA, [6, 7]). These works typically propose a RSA cost function which takes into account the instantaneous fragmentation of the links to be used at the path selection stage. As an example, results obtained in [6] by using fragmentation-aware RSA show that the load supported by the network at BBP=1% is increased in about 30%, compared to the traditional shortest-path first-fit RSA. It is therefore demonstrated that considering fragmentation at the traffic allocation phase is effective. Further elaboration about the appropriate metrics to measure EON fragmentation can be found in [8], where time is shown also to be an important aspect of the problem. Multi-path routing [9] has also been proposed as a strategy to prevent fragmentation in EONs. By splitting the connection to be allocated in different pieces and taking advantage of multi-path allocation, an obvious gain is obtained in terms of spectrum efficient utilization. This gain stems from the fact that the need to find large free spectral voids is eliminated. Another interesting idea aiming to avoid fragmentation was introduced in [10]. It consists in partitioning the available spectrum and using each band to allocate the different types of generated connections. By applying this strategy spectrum entropy (and hence fragmentation) is reduced as a single sized type of connections is allocated in each spectrum band.

Nevertheless, fragmentation is intrinsic to the EON concept due to the cohabitation of heterogeneous connections sharing the links available spectrum. Many works have focused on defragmentation strategies during EON operation, reactive as well as proactive, and using integer linear programming (ILP) as well as heuristic approaches. A comprehensive collection of the different approaches proposed in the literature can be found in [11]. This work provides an interesting qualitative comparison among the different defragmentation proposals found in the literature.

Focusing on time based proactive defragmentation algorithms, which is the concrete work topic of this work, some previous contributions have also been found. Some ILP formulations were proposed and solved in [5]. In this work, after arranging the existing network connections, a complete reallocation is done with or without re-routing allowance. Obviously, obtained results are better when connections routes can be changed but the number of connection disruptions in this case is also higher. Indeed, advantages of connections re-routing are reduced when the traffic in the network is highly dynamic. Our proposal will consider only spectrum reallocation (without re-routing of connections). A make-before break scheme was proposed in [12] which allows reducing the number of connection disruptions. It is demonstrated there that avoiding connections disruption provides considerable network performance gain.

Authors in [13] proposed to execute the defragmentation routine each time that a certain number of connections are finished. The proactive defragmentation algorithms used in this work are executed every N time units, and the first objective consists in finding an appropriate N (defragmentation period) value. Small N values imply having

always the spectrum as compacted as possible but at the expense of running the algorithm too frequently. This would be translated into huge complexity in the control and management of the EON. Indeed, frequent defragmentation algorithm runs prevent from using complex solutions, usually based on ILP solutions, whose runtimes (see for example [14]) become prohibitive under highly dynamic traffic conditions. On the other hand, if the period is too high, the effects of defragmentation are negligible. As has been stated previously, although hitless defragmentation is desirable, it imposes so severe conditions to the process then limiting its effectiveness. Therefore, many previous contributions have skipped this requirement and have concentrated on just increasing the defragmentation effectiveness. Parallelization [14] has been introduced as an efficient way to alleviate the complicated traffic migration and latency reconfiguration. A collateral consequence of defragmentation stems from the time consumed while rearranging the existing connections, and its effect on the network capacity to accept new requests during this period. In summary, the analysis and proposal of solutions to avoid the spectrum fragmentation in EONs has deserved lots of research efforts during the recent years, and it has been shown that some key parameters (defragmentation frequency, traffic disruption, complexity of the solution) have a great effect on the whole EON performance. These parameters are taken into consideration and discussed in the solutions evaluated in this work.

3. Proactive defragmentation proposal description

The proactive defragmentation scheme applied in this work will be time based (it is run periodically) and end-to-end (reallocation of connections is done in the whole path). It is worth noting that the defragmentation algorithm (as well as the different strategies to select the connections to be reallocated) will be maintained through simulations regardless of the network and the traffic profile considered. The algorithm used (further details are provided in next paragraphs) is represented below:

Defragmentation algorithm

- 1: decide the set of movable connections (M).
 - 2: order the movable connections (according to the selected strategy)
 - 3: **for** each movable connection M_i in M **do**:
 - 4: check the spectrum on every link along the path and find the First-Fit void capable
 - 5: **if** void is found
 - 6: move the connection to the void
 - 7: Moved=Moved+1
 - 7: **if** no connections on any link between initial position and final position
 - 8: Hitless_moved=Hitless_moved+1
 - 9: **end if**
 - 10: **end if**
 - 11: **end for**
-

The set of movable connections will be decided according to the final objective, which consists in reducing the BBP while minimizing the number of moved connections (as an example, connections lasting soon could be discarded since their potential effects will disappear as well soon). On the other hand, this algorithm will be implemented by

considering different strategies, which are given below and whose effects cannot easily be deduced:

- Older first (OF): this is the simplest one and is used for benchmarking purposes. Connections are ordered according to its birth, the older the first. The rationale behind that consists in the fact that the older a connection is, the more likely it is that some connections sharing links with it could have ended. Indeed, connections which have been allocated recently have been suitably placed (the SA algorithm looks for the best option available).
- Bigger (smaller) bandwidth first (BF/SF): the advantages obtained by re-allocating first the biggest connections are two-fold. First, the bigger voids will be filled by big sized connections (filling them with small connections would imply increasing fragmentation); second, larger voids are left for future arriving connections. Nevertheless, starting with the smaller connections can also be advantageous owing to the fact that the chances of reallocating each one of them increase (smaller voids in the left part of the spectrum are needed).
- Longer lasting first (LLF): in this case, intuition tells us that short lasting connections will disappear soon, so their potential damage will finish soon. Therefore, connections are ordered according to their ending time in descending order and then spectrum re-allocations start. By setting a threshold for the remaining duration that a connection needs so as to be considered for reallocation, the total number of moved connections can be notably reduced. This strategy has been proposed in [8] as a good way to prevent future fragmentation in the network.
- Longer (shorter) path first (LPF/SPF): it is not clear which one of the strategies is better. On the one hand, by starting with the long-routed connections, spectrum is released in more links. On the other hand, finding voids is more difficult as the spectrum continuity constraint affects more to longer paths, so fewer reallocations are allowed and thus the effect of defragmentation is less significant.
- Higher (lower) slot index first (HSF/LSF): was the one implemented in [5]. At first sight, the more on the right side of the spectrum a connection is, the more likely it should be to find an alternative spectrum void on its left side. Although starting with lower index connections seems in principle a bad idea, it has also been tested as it could be a simple way to increase the number of hitless reallocations. Indeed, results will show that the number of hitless reallocations in this case is the highest one.

The effectiveness of these strategies is measured in terms of extra network supported load at a given BBP, but also other parameters such as total number of connections reallocated (as well as fraction of hitless reallocations), which give an idea about the added complexity, are important to gauge their performance. The theoretical analysis of fragmentation is very complex and almost unsolvable [15], since the number of possible combinations explodes assuming typical values of hundreds of FS per link and partial mesh topologies. Indeed, link fragmentation is usually defined [2] by just taking into account the number of isolated spectrum available bands. Nevertheless, this parameter does not necessarily measure the link capacity to accept novel connection requests [13]. Therefore, with the aim of finding some clear performance trends, a simulation environment has been deployed and some conclusions are extracted from extensive simulations under different scenarios. The obtained results will be assessed for different network topologies and offered traffic statistical distributions. It is worth noting that the

main objective of this work consists in finding the optimal defragmentation period when truly random (obviously with some reasonable statistical properties) offered connections are generated. Hence, results will give insight into optimum defragmentation periodicity values for EONs under dynamic traffic conditions.

According to line 1 of the defragmentation algorithm depicted above, it is important to decide the set of connections which are to be potentially moved and limit this set as much as possible. By reducing M , the number of operations during the defragmentation run is reduced as well. Nevertheless, if the set of connections eligible for reallocation is too limited, the effects of defragmentation become negligible. Therefore, it is important to find an appropriate threshold which allows saving reallocation operations while guaranteeing that defragmentation is still effective.

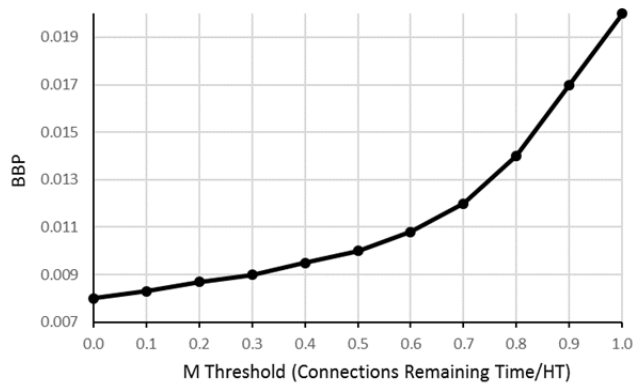


Fig. 2. Effects of limiting the set of movable connections (M) on the defragmentation efficacy. M includes all the connections whose remaining lifetime is higher than $\text{Threshold} \cdot \text{HT}$.

Figure 2 shows the effects of limiting M according to the remaining lifetime of the connections existing in the network. Values in Fig. 2 have been obtained by averaging many simulations for the different topologies used in this work (detailed in next Section), with uniform distribution of source-destination (S-D) pairs, and uniform distribution of connections bandwidth. It has been found that by including only in M those connections whose remaining time is higher than $0.2 \cdot \text{HT}$ (where HT corresponds to the average connections holding time), the penalty on the defragmentation efficacy is very low while the number of reallocated connections at each defragmentation run is reduced by about 25%. It is inferred from Fig. 2 that, if the threshold value is 0.2 (only connections ending later than $0.2 \cdot \text{HT}$ time units after the current time are included in M), the BBP only increases from 0.008 (obtained when all the existing connections are included in M) to 0.0087. On the other hand, it has been measured that the number of reallocated connections in a complete EON simulation run ($5 \cdot 10^4$ generated connections) is reduced from 24000 to less than 18000. If the threshold is further increased, the defragmentation benefits are reduced and it has been found that, when its value is 1.5, the BBP value approaches that of the network with no defragmentation ($\text{BBP} = 0.028$ in this case). The 0.2 threshold value will be used throughout the different simulations carried out in this work (reported in next Section), and corresponds to the first remarkable result found throughout this study.

4. Simulation Framework and Results

In order to evaluate the impact that the defragmentation conditions previously

mentioned have on the performance on the EON, an ad-hoc Matlab® based simulator has been deployed. Its main characteristics are,

- NSFnet (14 nodes, 21 bi-directional links), European Optical Network (EONet, 19 nodes, 37 bi-directional links) and USBackbone (24 nodes, 48 bi-directional links) topologies [3, 16] have been simulated. Connections S-D pairs are uniformly distributed among all the network's nodes. The transmission capacity assumed at every node has been adjusted to ensure that, when BBP=1% (under ideal defragmentation conditions), approximately half of the blocked connections are due to lack of spectral resources while the other half are due to lack of transceivers [17].
- Regarding the routing and spectrum allocation (SA) algorithm for allocating new connections, simulations run a typical adaptive k-shortest path routing (with k=3) and First-Fit SA [18].
- Number of FS (each one corresponding to 12.5 GHz) per link is equal to 360, so a 4.5 THz spectrum band is used in each fiber. This value corresponds approximately to the optical fibers C-band. Only 320 FS are used to allocate requested connections, and the remaining 40 FS are reserved for auxiliary temporal connections needed during the reallocations that may affect neighbor traffic (moved connections not counted as hitless in the defragmentation algorithm). It is important to highlight that reserving only 40 FS as backup resources during reallocations has a clear impact on the maximum number of connections that can be reallocated. Concretely, potential lack of spectrum in this band appears when the number of connections to be allocated in it ranges from 60 to 80 depending on the topology considered. In other words, blocking will appear in the backup band if the number of connections to be reallocated in a defragmentation run reaches these values.
- Number of FS per connection (BW) ranging from 1 to B, where B has been set to 8 and 10 during the different simulations. Concerning the connections BW statistics and aiming to obtain results as generic as possible, both Uniform and Normal distributions of connections have been tested, and no remarkable differences between them have been found.
- The connection's inter-arrival time (IAT) follows a Poisson process with an average value of IAT=1 t.u. (time units). The connection's holding time follows a negative exponential distribution. The average value (HT t.u.) is adjusted to obtain the appropriate offered load values (those that led the BBP to the desired values).
- The defragmentation algorithm introduced above is executed every N t.u. The objective of this work is focused on finding out the efficiency of the different strategies proposed as well as selecting the best strategy and presenting results about its performance. Finding the optimal N values is a key contribution of this work.

As stated in the previous Section, results will give insight on the gain obtained by using the defragmentation algorithms. This gain is measured in terms of reduction of the BBP at a given offered load or, conversely, increase in the load supported by the network at a given BBP. Additionally, measuring the complexity introduced by the defragmentation process is also very important, specially to assess the feasibility of the proposed schemes in real scenarios. This complexity is determined by the fraction of connections that are reallocated during the simulation. The larger this fraction is, the more complex the network management will result. In fact, the number of connections that have to be reallocated at each defragmentation run is a key limiting factor as has

been explained above. Indeed, the hitless reallocations fraction is also a key parameter since these movements are quite easier (no transitory connections have to be set-up and torn-down to ensure that the connection is not disrupted during reallocation). Hence, the flawlessness of the different strategies is also defined by using these parameters: Reallocated Connections and Hitless Reallocated Connections.

Fig. 3 shows by means of an example (NSFnet topology, uniform distribution of source-destination (S-D) pairs, and uniform distribution of connections bandwidth, whose values range from 1 FS to 8 FS) the generic effect of the defragmentation frequency on BBP. It is observed that network performance reaches its optimal values for defragmentation periods close to 50 for lightly loaded networks, and around 20 for the heavy loaded case.

Of course, the defragmentation effects depend on the temporal characteristics of the offered load, so values in the horizontal axis of Fig. 2 correspond to the ratio between the defragmentation period and the inter-arrival time (IAT) of connections. On the other hand, it also depends on the network size (number of S-D pairs). In the NSFnet case, there are 182 possible pairs. This implies that it is statistically expected that a new connection will correspond to a given S-D pair every 182 new connections in average. When network links are examined, and considering the usage of k-SP routing, it is obviously found that some of them (those placed in geographically central points) are more frequently used than others and thus more crowded. Indeed, their spectrum is usually more fragmented, because the number of connections allocated and torn down that pass through these central links is considerably higher. Therefore, the influence of these links on the network BBP is bigger.




Fig. 2. Effects of defragmentation periodicity on the network bandwidth blocking probability (BBP) at different network load values: (a) Light load, $BBP(w/o \text{ defrag.})=0.005$. (b) Heavy Load, $BBP(w/o \text{ defrag.})=0.108$.

Further analyzing Fig. 3, the effect of the offered load on the defragmentation efficacy is observed. When the load is low (Fig. 3.a) the maximum slope on the curve is found for values on the x-axis ranging from 1000 to 100. The BBP tends to stabilize to its minimum value and only marginal gain is achieved when defragmentation period is reduced below 100. As for the BBP gain attained by spectrum defragmentation, it corresponds to more than one order of magnitude (from 0.005 to less than 0.0005). On the other hand, when the offered load value is high (Fig. 3.b, where the BBP without defragmentation is slightly over 10%), different trends are observed: first, the BBP

relative gain attained is less than 50% (from 0.108 to a minimum value of 0.069); second, very frequent defragmentation is required to reach this minimum BBP (the BBP stabilizes when values on the x-axis are around 20). In summary, the defragmentation effects on the EON performance depend on the load offered to the network. This trend has been obviously shown regardless of the topology considered. Therefore, simulations from now on will focus on networks whose BBP is around 1% since this corresponds to a manageable value in realistic scenarios and the defragmentation benefits are higher.

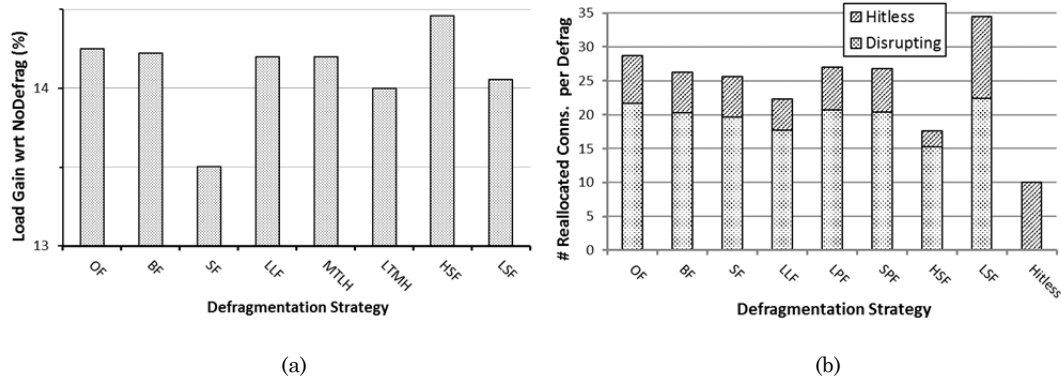


Fig. 4. Quantifying the performance of the different defragmentation strategies. (a) Gain in network throughput at BBP=1%. (b) Hitless and Disrupting reallocated connections for each defragmentation strategy.

Results obtained by applying the different ordering strategies proposed in previous Section are shown in Fig. 4, where the generated connections bandwidth values are uniformly distributed from 1 FS to 8 FS ($B=8$) and the NSFNet topology is assumed, with any-to-any nodes uniformly distributed connections. The worst case is obtained, as expected, when only hitless reallocations are allowed. The number of reallocations in this case is considerably lower than in any of the remaining cases. Indeed, the gain obtained in network throughput terms is 4.2%, while as seen in Fig. 4.a, it takes values ranging from 13.5% to 14.5% for any of the strategies considered, with slight differences among the different cases.

Concerning the number of connections reallocated in each case, differences are more significant. Worst case (maximum number of reallocations) is obtained when connections are ordered according to their position in the spectrum and starting with those placed more on the left (LSF). In this case, the connections generated during the simulation are reallocated in average slightly more than once during their holding time. These results have been obtained using a defragmentation period $N=32$ t.u. which, as shown in Fig. 3, guarantees a network performance close to the optimum value at BBP around 1%. Taking into account that the simulation time is about 50000 t.u. and that, as shown in Fig. 4.b, a maximum of 35 connections are reallocated at each defragmentation run, the total number of reallocations is almost 55000 in this worst case. Concerning the fraction of hitless reallocations, it ranges from 35% in LSF to around 15% in the HSF case. Analogous results have been obtained considering the alternative network topologies (EONet and USBackbone) with respect to the performance of the different defragmentation strategies. The obtained results can therefore be generalized to any of these typical backbone topologies.

Assuming that the disrupting reallocations imply an increase of network operation

complexity, selecting the strategy where this number is minimized seems to be the right decision. Therefore, the subsequent simulations will use the HSF strategy which presents a good tradeoff between gain achieved and number of reallocated connections. These simulations are focused on finding the appropriate defragmentation period regarding the gain achieved and the number of connections which have to be reallocated at each defragmentation run, and results are shown in Fig. 5. These results intend to be as generic as possible, so three different topologies (NSFNet, EONet and USBackbone, [3, 16]) have been simulated. It is worth noting that Fig. 5 results show very similar trends regardless of the topology used. The validity of these results is thus assessed, and could be extended to further scenarios (partially meshed topologies with about 20 nodes and nodal degree around 3). Evolution of network throughput gain as well as reallocated connections, with respect to defragmentation period, are depicted, respectively, in Fig. 5.a and Fig. 5.b, which offer some interesting conclusions:

- Concerning throughput gain, it increases steadily from $N=1000$ t.u. to $N=32$ t.u., irrespectively of the network topology considered. It is worth commenting that $N=32$ t.u., and $N=16$ t.u. when the EON topology is considered, are the periods that allow maximizing the gain values. Further reducing the time between successive defragmentation runs lead to saturation and no additional gain is observed. It can be concluded that the throughput gain obtained by applying periodic defragmentation ranges from 14% to 18%. This range holds when changing the traffic profile characteristics. Very similar results have been found when using connections bandwidth sizes normally distributed from 1 to 10 FS ($B=10$). The values shown are normalized to the respective network throughput values obtained in each topology, which are obviously different (the absolute throughput is approximately 10% higher in the EONet and the USBackbone than in the NSF network).
- As for the number of reallocations, it is surprising how the performance is almost identical for the three topologies considered. This result allows again assessing the generality of the results obtained. Fig. 5.b shows the results when the defragmentation period is lower than 128 t.u. (differences between the different networks at higher period values are negligible). It is seen that when the defragmentation period is below 32 t.u. the total number of reallocated connections is over 0.9 regardless of the topology considered (values for NSF topology are slightly higher in this case). This means that, in average, 90% of the connections are reallocated at least once during their lifetime. These results have been assessed simulating other traffic profiles (connections size ranging from 1 FS to 10 FS, and Normal instead of Uniform distribution); no significant differences have been observed.




Fig. 5. Effects of defragmentation period on network performance. (a) Network throughput gain at BBP=1% with respect to the case where no defragmentation is used. (b) Number of reallocated connections during simulation (normalized to the total number of generated connections).

It is shown in Fig. 5.a how the slope starts to decrease when the defragmentation period is smaller than 64 t.u. It is worth noting that low N values imply frequent defragmentation runs. However, the number of connections moved per defragmentation run is also lower. As observed in Fig. 4.b, the number of connections moved per defragmentation run when $N=32$ t.u. is below 20 if HSF strategy is used. This value is increased up to 130 reallocated connections when $N=256$ t.u. As previously discussed, most of these reallocations are disrupting, so the connection must be temporarily allocated in a backup path, then the original one is torn down and then created again using the new assigned FS. As has been explained before, reserving only 40 FS for temporarily allocation of connections limits the maximum number of connections to be reallocated to about 70. On the other hand reserving a larger part of the spectrum for this purpose would impact on the network capacity (the spectrum available for regular connections would be reduced). Indeed, the higher the number of connections to be reallocated is, the more time this process will take. The defragmentation algorithms used in this work have been found to have runtime values below 0.5 s (for the worst case corresponding to USBackbone topology, whose number of nodes is higher). The processes of tearing down and establishing the new connection are expected to take few 10's of ms per connection. Let us assume that each connection establishment or torn down takes 50 ms [19], and that reallocations are performed sequentially. Each disrupting reallocation implies establishing the backup connection, tearing down the old one, establishing the new one and finally tearing -down the backup one. Therefore, the time consumed by each reallocation is $50 \times 4 = 200$ ms. If 20 connections need to be reallocated (as shown in Fig. 4, best case corresponding to HSF strategy, is about 15), the defragmentation duration would be $0.5 \text{ s} + 20 \times 0.2 \text{ s} = 4.5 \text{ s}$. The considered connection requests are Poisson distributed with IAT = 1 t.u. Therefore, considering that each time unit corresponds to 60 s (one connection request arriving every minute in average), the probability that no new connections are requested during the defragmentation process is higher than 92%. If each time unit corresponds to 10 minutes, this probability rises to 99.2 %. If the number of connections to be reallocated rises to 130, the time between connection requests must be increased to almost 1 h to guarantee that the probability of not receiving new requests during defragmentation is higher than 99%. In summary, selecting low N values is advantageous because the number of connections to reallocate at each defragmentation run is reduced and therefore the process is faster. Parallelization

of the defragmentation reconfigurations, as proposed in [14] and [20], would drastically reduce the defragmentation latency, thus relaxing these requirements. The main drawback of applying this so frequent defragmentation scheme comes from the communication overhead that it implies in a centralized (e.g. Software Defined Network based) control. Further study is needed to decide whether a less frequent defragmentation would be beneficial.

5. Conclusions

Spectrum defragmentation is a key issue in EON research activities. In this work, results about the main effects of proactive defragmentation on EON network performance have been obtained by means of extensive simulations. The optimum ratio between defragmentation period and rate of connections arrival is found to be in the range from 16 to 32. Using this defragmentation period, the network throughput at BBP=1% is kept close to its optimal value (that obtained when continuous defragmentation is applied). The minimum number of connections that need to be reallocated at each defragmentation run is below 20, corresponding to the case where the existing connections are ordered according to their current position in the spectrum from higher to lower FS occupied (HSF strategy). It has been demonstrated that, values obtained using only hitless defragmentation are far from those obtained allowing disrupting reallocations. The simulation results have been assessed using different network topologies as well as different statistical distributions in the connections requests sizes and time parameters (HT and IAT).

Acknowledgments

This work was partially supported by the EC through the METRO-HAUL project (G.A. 761727), and from the Spanish MINECO TWINS project (TEC2017-90097-R).

References

- [1] O. Gerstel, M. Jinno, A. Lord, B. Yoo, "Elastic Optical Networking: A New Dawn for the Optical Layer," *IEEE Commun. Mag.*, vol. 50, no. 2, pp. s12-s20, 2012.
- [2] 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, 2013.
- [3] R. Wang, B. Mukherjee, "Provisioning in Elastic Optical Networks with Non-Disruptive Defragmentation," *Journal of Lightwave Technology*, Vol. 31, No. 15, pp. 2491-2500, 2013.
- [4] Y. Yin, K. Wen, D. J. Geisler, R. Liu, S. J. B. Yoo, "Dynamic on-demand defragmentation in flexible bandwidth elastic optical networks," *Optics Express* Vol. 20, pp. 1798-1804, 2012.
- [5] A. N. Patel, P. N. Ji, J. P. Jue, T. Wang, "Defragmentation of Transparent Flexible Optical WDM (FWDM) Networks," in *Proc. of OFC 2011*, paper OTuI8, 2011.

- [6] Y. Yin, M. Zhang, Z. Zhu, S. J. B. Yoo, “Fragmentation-Aware Routing, Modulation and Spectrum Assignment Algorithms in Elastic Optical Networks”, in Proc. of OFC 2013, paper OW3A.5, 2013.
- [7] X. Chen, S. Ma, B. Guo, Y. Wang, J. Li, Z. Chen, Y. He, “A novel fragmentation-aware spectrum allocation algorithm in flexible bandwidth optical networks”, *Optical Switching and Networking*, Vol. 12, pp. 14–23, 2014.
- [8] S. K. Singh, A. Jukan, “Efficient Spectrum Defragmentation With Holding-Time Awareness in Elastic Optical Networks”, *Journal of Optical Communications and Networking*, Vol. 9, No. 3, pp. B78-B89, 2017.
- [9] A. Pagès, J. Perelló, S. Spadaro, J. Comellas, “Optimal Route, Spectrum, and Modulation Level Assignment in Split-Spectrum-Enabled Dynamic Elastic Optical Networks”, *Journal of Optical Communications and Networking*, Vol. 6, No. 2, pp. 114-126, 2014.
- [10] R. Wang and B. Mukherjee, “Spectrum management in heterogeneous bandwidth networks,” in Proc. IEEE GLOBECOM 2012, pp. 2907–2911, 2012.
- [11] B. C. Chatterjee, S. Ba, E. Oki, “Fragmentation Problems and Management Approaches in Elastic Optical Networks: A Survey”, *IEEE Communications Surveys & Tutorials*, vol. PP, no. 99, 2017.
- [12] T. Takagi, H. Hasegawa, K. Sato, Y. Sone, A. Hirano, M. Jinno et al., “Disruption minimized spectrum defragmentation in elastic optical path networks that adopt distance adaptive modulation,” in Proc. of ECOC 2011, pp. 1–3, Sept. 2011.
- [13] M. Zhang, W. Shi, L. Gong, W. Lu, Z. Zhu, "Bandwidth defragmentation in dynamic elastic optical networks with minimum traffic disruptions," in Proc. IEEE International Conference on Communications, ICC 2013.
- [14] M. Zhang, C. You, Z. Zhu, “On the Parallelization of Spectrum Defragmentation Reconfigurations in Elastic Optical Networks”, *IEEE/ACM Transactions on Networking*, Vol. 24, No.5, pp. 2819-2833, 2016.
- [15] H. Beyranvand, M. Maier, J. A. Salehi, “An Analytical Framework for the Performance Evaluation of Node- and Network-Wise Operation Scenarios in Elastic Optical Networks”, *IEEE Transac. on Comms*, vol.62, no. 5, pp. 1621-1633, 2014.
- [16] J. Comellas, G. Junyent, “Improving link spectrum utilization in Flexgrid optical networks,” *Journal of Optical Communications and Networking*, Vol. 7, No. 7, pp. 618-627, 2015.
- [17] J. Comellas, L. Vicario, G. Junyent, “Using Static Connections to Improve Elastic Optical Networks Performance,” *OSA Advanced Photonics Congress*, paper NeTu1B.5, 2017.
- [18] Y. Wang, X. Cao, and Y. Pan, “A study of the routing and spectrum allocation in spectrum-sliced elastic optical path networks,” in Proc. INFOCOM, 2011, pp. 1503–1511.
- [19] A. Pagès, F. Agraz, R. Montero, G. Landi, R. Monno, J. I. Aznar, A. Viñés, C. Jackson, D. Simeonidou, S. Spadaro, “Experimental Assessment of VDC Provisioning in SDN/OpenStack-based DC Infrastructures with Optical DCN”, in Proc. Of the 42nd European Conference on Optical Communication, 2016.
- [20] A. Castro, R. Martínez, R. Casellas, L. Velasco, R. Muñoz, R. Vilalta, and J. Comellas, “Experimental Assessment of Bulk Path Restoration in Multi-layer Networks Using PCE-based Global Concurrent Optimization”, *Journal of Lightwave Technology*, Vol. 32, No. 1, pp. 81-90, 2014.