

RESEARCH ARTICLE

Correlation between traffic granularity and defragmentation periodicity in elastic optical networks

Pouria Sayyad Khodashenas*, Jaume Comellas, Jordi Perelló and Salvatore Spadaro

Advanced Broadband Communications Center (CCABA), Universitat Politècnica de Catalunya (UPC), Jordi Girona 1-3, Barcelona 08034, Spain

ABSTRACT

While elastic optical network technologies have emerged as promising solutions for future ultra-high-speed optical transmission, the unavoidable spectral fragmentation problem that appears in such networks significantly degrades their performance. In light of this, spectral defragmentation technologies have been introduced in elastic optical networks, aiming to increase the spectrum utilisation. During the defragmentation operation, the available fragmented spectrum bands are consolidated by reallocating existing connections, either re-routing them along alternative routes and/or re-tuning them onto different spectrum portions. Obviously, spectral defragmentation increases network complexity and cost. Therefore, it is highly desirable to limit its application as much as possible, while keeping network performance within acceptable margins. In this paper, we focus on analyzing the correlation between the optimal (i.e. minimum) spectrum defragmentation periodicity in the network with the granularity of the supported traffic. For this purpose, we initially introduce a novel algorithm for efficient spectrum defragmentation. The proposed algorithm aims to consolidate the available fibre spectrum as much as possible, while limiting the number of reallocated active connections. Then, supported on extensive simulation results, we show how spectral defragmentation periodicity can be effectively configured by having knowledge of the offered traffic granularity. Copyright © 2014 John Wiley & Sons, Ltd.

*Correspondence

P. S. Khodashenas, Advanced Broadband Communications Center (CCABA), Universitat Politècnica de Catalunya (UPC), Jordi Girona 1-3, Barcelona 08034, Spain.

E-mail: pkhodashenas@tsc.upc.edu

Received 15 June 2013; Revised 22 October 2013; Accepted 27 December 2013

1. INTRODUCTION

Triggered by emerging services such as high-definition video distribution or social networking, the Internet protocol traffic volume has been exponentially increasing to date [1]. Furthermore, the traffic growth rate will not stop here because of the day by day technology advances. For example, new hardware advances such as multi-core processing, virtualisation and network storage will leverage innovative 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 [2]. In response to these large capacity and diverse traffic granularity needs of the future Internet, the elastic optical network architecture has been proposed [3]. By breaking the fixed-grid spectrum allocation limit of conventional wavelength division multiplexing (WDM) networks [4], elastic optical networks

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 elastic optical networks. Moreover, unlike the rigid optical channels of conventional WDM networks, an elastic optical path can also expand or contract dynamically to meet different traffic loads [5]. As a result, incoming traffic demands can be served in a spectrum-efficient manner.

This technological advance poses additional challenges on the networking level and more 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, which is typically referred as the spectrum contiguity constraint. Upon tear down of connections, their allocated spectral resources are released and could be assigned to new connection requests. In a dynamic traffic scenario, the random connection arrival and departure process in the network leads to the fragmentation of

the spectral resources. Therefore, in medium-loaded and high-loaded network scenarios, the possibility of finding sufficient contiguous spectrum to establish a connection over such fragmented spectrum can be very low, leading to connection request blocking despite having enough but non-contiguous spectrum resources. In light of this, there is an increasing demand from network operators to be able to periodically reconfigure their networks, aiming to improve the spectrum utilisation [6]. This operation is called spectrum defragmentation.

During the defragmentation operation, the available fragmented spectrum bands are consolidated by reconfiguring active connections, that is, changing their routes or assigning them a different portion of spectrum or both, while maintaining the imposed continuity and contiguity constraints [7]. However, the traffic granularity, namely, the offered load to the network and the average bit rate of the connections, has a direct effect on the spectrum fragmentation experienced at network links. Such effect can also affect the efficacy of a periodic spectrum defragmentation to keep an acceptable network performance. In particular, a deep analysis of the relationship between the traffic granularity and the defragmentation periodicity can allow a network operator to find the optimal defragmentation interval yielding the desired network performance but requiring the minimum active connection disruptions and network control and management burden.

In this paper, we focus on this correlation between traffic granularity and defragmentation interval. To this end, we initially introduce a novel hitless defragmentation algorithm. The proposed algorithm maximises the spectrum utilisation in elastic optical networks. Then, we explore the impact of traffic granularity on the defragmentation interval through extensive simulation studies. The rest of the paper is organised as follows. In Section 2, we initially review the related work on defragmentation operation. Section 3 details the proposed defragmentation algorithm. Simulation results are presented in Section 4. Finally, Section 5 concludes the paper.

2. RELATED WORK

As mentioned earlier, in a dynamic scenario, the spectral resources allocated to connections are released for future requests upon tear down. The randomness in the connection set-up and tear-down processes leads to fragmentation of the spectral resources in the network. Figure 1 shows scattered spectral fragments in four links of a typical elastic optical network. As shown, the available spectrum in the network links is fragmented into small non-contiguous spectral bands. The spectrum fragmentation significantly decreases the probability of finding enough contiguous spectrum for establishing new incoming connections, especially those traversing multi-hop paths (e.g. traversing links 1 to 4 in Figure 1) and/or requesting large amounts of bandwidth. In fact, new connection requests can be blocked in spite of having enough spectral resources if these are non-contiguous. Because an efficient utilisation




Figure 1. Scattered spectral fragments in four links of a typical elastic optical network.

of the limited spectral resources in the network is an important issue for network operators, this problem has grabbed the research community attention. To address it, the proposal of spectral defragmentation strategies has appeared in the literature [5–13]. In general, spectral defragmentation means arranging the spectral resources in order to consolidate the available spectral fragments contiguously so that they can be used for serving new incoming connection requests. Spectral defragmentation reduces connection blocking probability and maximises the service capacity of the network. In addition, it enables better network maintenance, enables more efficient network restoration and enhances the quality of service for a given spectral capacity.

For the first time, authors in [8] introduced the spectral defragmentation problem. They formulated it and proposed heuristic algorithms for spectral defragmentation operation. A scalable node architecture for realising the spectral defragmentation proposal was also presented in [9] by the same authors. They showed a proof-of-principle experimental demonstration of spectral defragmentation in their studies. These studies clearly illustrated the benefits of spectral defragmentation in reducing the blocking probability of elastic optical networks. However, when the network has to be reconfigured to perform defragmentation, existing active connections may be disrupted. Achieving a disruption-minimised spectrum optimisation is a major challenge in spectral defragmentation. To address this, a disruption-minimised spectrum defragmentation algorithm in distance-adaptive elastic optical networks was proposed [10], which uses a make-before-break re-routing scheme.

In addition, dynamic on-demand lightpath provisioning using spectral defragmentation was studied in [11]. A spectral defragmentation algorithm, as well as the corresponding node architectures with wavelength conversion to support on-demand lightpath provisioning, was proposed. The same authors continued their work in [12] by introducing new heuristic defragmentation algorithms based on an auxiliary graph approach for elastic optical networks. Besides, a new routing and spectrum assignment algorithm that maximises spectrum utilisation in elastic optical networks was presented in [13]. Authors introduced a metric

that quantifies the consecutiveness of the common available spectrum slots among relevant fibres. The first test-bed demonstration of an elastic optical network was presented in [14]. Authors used a real-time adaptive control plane for their test-bed that adjusts the modulation format and spectrum positioning to maintain high quality of service and spectral efficiency. Finally, a field trial experiment of flexible spectrum switching and spectral defragmentation was presented in [15].

The aforementioned studies on the benefits of spectral defragmentation of elastic optical networks show a significant improvement in terms of spectrum efficiency and number of interrupted connections during the reconfiguration operation. However, there is no clear vision on the effects of traffic granularity on the defragmentation operation in a dynamic scenario. Authors in [16] focused on the problem from the defragmentation operation interval perspective. They proposed a defragmentation algorithm based on a factor that indicates the occupation of spectrum in a link or in the network. In this sense, by comparing the proposed factor with a given threshold, their algorithm is able to link the traffic granularity with the defragmentation interval. However, because the threshold is an arbitrary value in this study and the whole available spectrum of network links is not considered for calculating the defragmentation factor, there is still a long way to go before really finding a correlation between traffic granularity and spectral defragmentation. In this work, we target at deeply analyzing this correlation. For this purpose, we initially propose a novel heuristic defragmentation algorithm, called iterative-defragmentation algorithm (IDA), which aims at increasing the network spectrum utilisation, while minimising the number of reallocated connections during the defragmentation operation.

3. ITERATIVE-DEFRAGMENTATION ALGORITHM

The primary objective of the IDA algorithm is to maximise the consolidation of spectrum, followed by the secondary objective to minimise the number of reallocated connections during the operation. To achieve the first objective, after establishing a certain number of connections over the network, we try to confine all existing connections in the lowest part of the spectrum. Specifically, we process them one by one in descending order according to their position in the spectrum (i.e. connections allocated in the highest parts of the spectrum are processed first), also meeting the second objective in this way. The proposed algorithm is as follows:

Iterative-defragmentation algorithm:

- (1) Make a list of all the existing connections in the network and arrange them in descending order on the basis of the connection operating spectrum.
- (2) Pick up the first connection from the ordered list and record its route in the network.

- (3) Check the spectrum status of the route in step 2.
- (4) If a lower available portion of spectrum is found, reallocate the connection (in a make-before-break fashion) in the lowest possible spectrum. Otherwise, do not reconfigure the connection and go to step 5.
- (5) Pick the next unprocessed connection in the list and repeat steps 2–4.
- (6) Repeat steps 1–5 for a number of iterations.

All possible existing connections in the network are reallocated in the lowest available portion of spectrum on their route after the first iteration. However, because the generation of spectral fragments is a random process in a dynamic network scenario, which depends on many factors such as the offered load to the network, average bit-rate demand per connection or average path length, applying only one defragmentation iteration may not result in the higher spectrum consolidation. For example, consider the three-node simple network shown in Figure 2(a) with the initial link spectral state shown in Figure 2(b). As seen, there exist some connections in the network with different path length and bandwidth. The defragmentation operation tries to reallocate existing connections in the lowest available portion of the spectrum on their routes. Because some of these connections have smaller spectral requirements and short path length, there is more chance for them to be reallocated (e.g. connections D or F). Meanwhile, as mentioned before, both continuity and contiguity spectrum constraints should be met when reallocating the connections, making some connections, such as E and G, to have lower defragmentation chances. In the event of defragmentation, existing connections are ordered as follows: G, F, E, D, C, A, B. According to the spectral status, connections F, E, D and C can be reconfigured on their paths because all links in their paths [e.g. (A, B) for connection D and (B, C) for connection C] have sufficient continuous spectrum available at a lower portion of the spectrum. However, connection G cannot be reallocated because of the spectral conflicts. As shown in Figure 2(c), the existing connections are confined after the first iteration; however, it is still possible to consolidate the connections even more. In the second iteration, the ordered list is updated as follows: G, E, C, D, A, F, B. Figure 2(d) shows the spectral state of network, in which connections G and E are reallocated to a lower spectrum portion. As shown, the second iteration optimally confines the existing connections in terms of spectral resources.

The number of reallocations in the first iteration [Figure 2(c)] is 4, while in the second iteration [Figure 2(d)], it is 2. In general, by allowing more iterations in IDA, a more consolidated spectrum may be achieved but at expenses of increasing the network control and management burden. Moreover, because the reallocation of connections follows a make-before-break approach to avoid connection disruption, which is performed sequentially, the number of iterations impacts directly on the total defragmentation time. With this in mind, the effect of number of iterations on the performance of IDA will be investigated

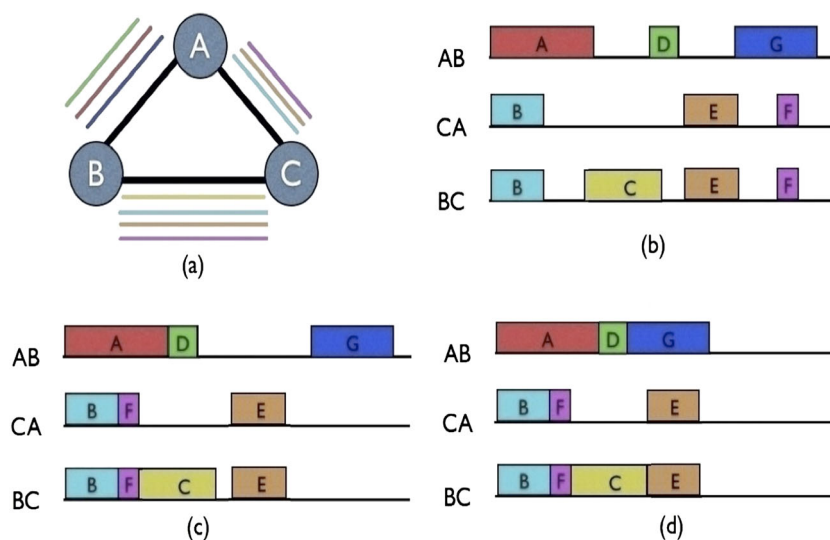


Figure 2. Illustrative example: (a) three-node simple network, (b) initial link spectral state, (c) link after the first iteration, and (d) link spectral state after the second iteration.

in the next section, together with the analysis of the correlation between traffic granularity and defragmentation periodicity in elastic optical networks.

4. SIMULATION RESULTS

To analyze the correlation between traffic granularity and defragmentation in elastic optical networks, we simulate a network scenario where IDA defragmentation is performed on a periodic basis (hereafter referred as periodic defragmentation). Moreover, two additional scenarios are simulated for benchmarking purposes: (1) no-defragmentation, namely, a basic elastic optical network scenario without defragmentation capabilities, and (2) on-demand defragmentation, where defragmentation is reactively executed upon blocking of an incoming connection request. Regarding the Routing and Spectrum Assignment (RSA) algorithm for allocating new incoming connections, all three scenarios run a standard k shortest path with first-fit slot assignment [17]. In the event of blocking, on-demand defragmentation tries to defragment the optical spectrum along the k shortest paths iteratively. More specifically, it starts with the shortest path, running the proposed IDA algorithm, but constrained to only those connections supported on any link along it. If a positive outcome is obtained from IDA, the spectrum defragmentation operation is triggered, the incoming connection is allocated and the on-demand defragmentation ends. Otherwise, the same actions are performed on the next shortest path and this until the last of the k shortest paths. Note here that connections must be reallocated end to end. Hence, although on-demand defragmentation focuses on one path only per iteration, changes can also affect other network links.

The 14-node NSFnet topology (Figure 3) has been selected for the simulations. We assume a total optical

spectrum of 1 THz per link. This spectrum is discretised in units of 12.5 GHz, referred as frequency slots [3]. In addition, according to the asymmetric nature of today's Internet traffic, unidirectional connections between source–destination nodes are considered. The offered load to the network follows a Poisson process so that different offered loads are obtained by keeping the mean holding time (HT) of the connections constant to 200 s, while modifying their mean inter-arrival time (IAT) accordingly (i.e. offered load = HT/IAT). Bit-rate demands for each source–destination pair are randomly generated following an exponential distribution over the range from 12.5 Gb/s (1 frequency slot) to 125 Gb/s (10 frequency slots). Hence, we are implicitly considering a signal modulation format with spectral efficiency of 1 b/s/Hz, such as binary phase-shift keying [18]. To start, we fix the defragmentation interval to 70 in the periodic defragmentation scenario. This means that a network-wide defragmentation is periodically performed every 70 connections established. Nonetheless, we will also evaluate the effect of this parameter on the network blocking probability later on in this section.

As a first study, we focus on periodic defragmentation and initially investigate the effect of the number of iterations performed by the IDA algorithm. The aim of this study is to find such a number of iterations in IDA that effectively reduces the blocking probability in the network, while keeping reasonable the number of reallocated connections. As mentioned before, more reallocated connections imply additional network control and management burden, which may not be translated into actual blocking probability reduction.

Specifically, we offer a load of 14 per node (i.e. 14 offered connections per node in average), which makes up a total offered load to the network of 196. Then, we extract the results for three different average bit rates per connection, namely, 35, 40 and 45 Gb/s, leading to an average




Figure 3. NSFnet topology.

traffic offered per node equal to 490, 560 and 630 Gb/s respectively. Figure 4(a) shows the blocking probability in the network against the number of iterations in IDA. As seen, by increasing the number of iterations, defragmentations are more effective, which helps in decreasing the blocking probability in the network. This is particularly true for 1 and 2 iterations under all average connection bit rates. From then on, however, blocking probability reduction becomes marginal with the number of iterations. In Figure 4(b), we quantify the percentage of reallocated connections over the total number of active connections in the entire network. As shown there, for instance, by changing the number of iterations from 2 to 3, 10% more connections have to be reallocated, while almost no improvement in terms of blocking probability is experienced. This trend also holds for larger number of iterations. Hence, from now on, we fix the number of iterations in IDA to 2, which allows to significantly reduce the blocking probability in the network with reasonable network control and management overhead.

Figure 5 compares periodic defragmentation against no defragmentation and on-demand defragmentation in terms of network blocking probability as a function of the offered load per node. In this study, the offered load per node ranges from 11 to 14.5, with an average bit rate per connection equal to 40 Gb/s. This is translated into an offered traffic per node ranging from 440 to 580 Gb/s. From the results, we observe that on-demand defragmentation outperforms the other two solutions along the entire offered load range. However, as it is well highlighted in Figure 5, only minor differences exist between on-demand and periodic defragmentation curves. Although the network-wide periodic defragmentation yields better spectrum utilisation than the on-demand k-shortest-path-constrained defragmentation in general, the configured defragmentation interval of 70 makes periodic defragmentation behave slightly worse than the on-demand one. Indeed, applying smaller

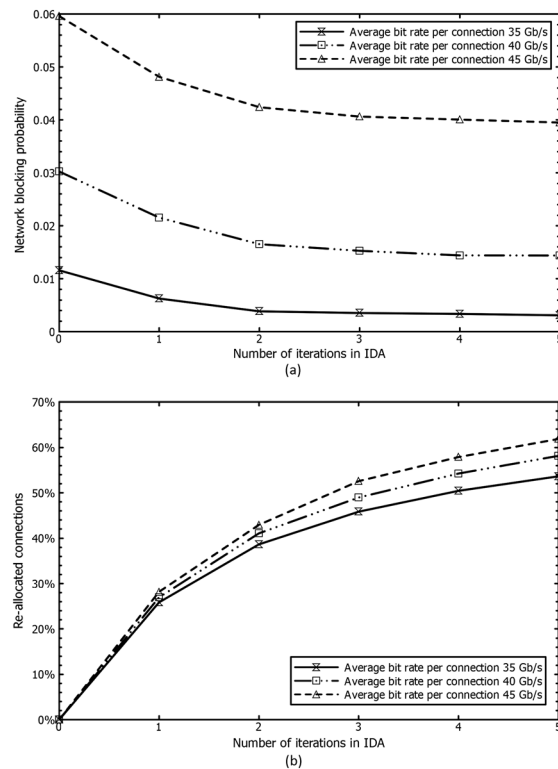


Figure 4. Effect of the number of iterations on the performance of the iterative-defragmentation algorithm (IDA) algorithm.

defragmentation intervals can make the former behave similarly or even better than the latter, as will be shown later on in Figure 8.

To give further insight into the effects of the offered traffic characteristics on the benefits of network defragmentation, Figure 6 compares the network blocking probability achieved by periodic, on-demand and no

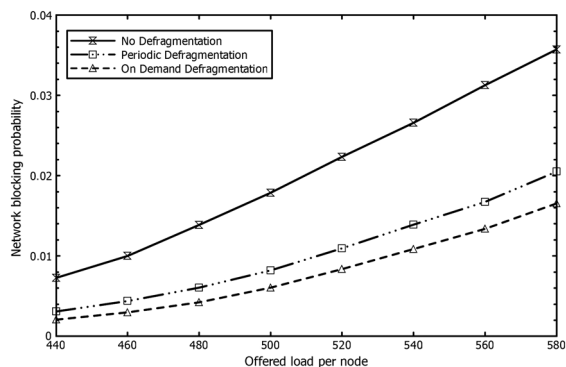


Figure 5. Network blocking probability for different traffic granularities. The offered load per node varies from 11 to 14.5, with an average bit rate per connection equal to 40 Gb/s.




Figure 6. Network blocking probability for an offered traffic per node equal to 500 Gb/s. The average bit rate per connection varies from 30 to 60 Gb/s, changing the offered load per node accordingly.

defragmentation as a function of the average bit rate per connection. To get these results, we have fixed an offered traffic per node of 500 Gb/s in all cases. Hence, higher bit rates per connection mean lower offered loads per node values and vice versa. As illustrated in the figure, even under the same offered traffic, relative benefits of defragmentation are clearly affected by the actual bit rates of the connections. For instance, for low bit rates per connection, almost one order of magnitude blocking probability reduction is achieved by periodic defragmentation against no defragmentation, while such a difference decreases when bit rates increase. Comparing on-demand versus periodic defragmentation, we can also see that the former slightly outperforms the latter along the entire range of study. As mentioned before, network-wide defragmentation is generally more effective than the on-demand constrained one, but the applied defragmentation interval of 70 makes periodic defragmentation to perform slightly worse eventually. These differences decrease with the bit rate per connection, as spectrally wider connections are more difficult to be efficiently reallocated in all cases.




Figure 7. Percentage improvement in blocking probability for different bit rates per connection (blocking probability equal to 2% with no defragmentation).

As highlighted in the previous figure, defragmentation benefits clearly depend on the bit rate of those connections offered to the elastic optical network. To quantify these benefits, Figure 7 illustrates the profitability of defragmentation against no defragmentation for different bit rates per connection. Only periodic defragmentation is considered here, with three different defragmentation intervals, namely, 28, 70 and 280. Furthermore, such a defragmentation profitability is measured as follows: for the evaluated bit rate per connection, we adjust the offered load per node in the network so as to obtain a 2% blocking probability (BP_{ND}) with no defragmentation. Then, we measure the improvement in blocking probability, equal to the steady-state blocking probability with no defragmentation minus the achieved blocking probability with periodic defragmentation (BP_{PD}), all divided by BP_{ND} :

$$Improvement(\%) = [(BP_{ND} - BP_{PD})/BP_{ND}] * 100$$

Looking at Figure 7, the obtained results are very in line with those in Figure 6, showing a periodic defragmentation improvement decrease as the bit rate of the connections increase. Moreover, the same behaviour is observed no matter the applied defragmentation interval is, although smaller defragmentation intervals (i.e. more frequent defragmentation) seem to moderate this effect. For example, with a defragmentation interval of 280, improvement sharply falls from around 40% in a network scenario with average connection bit rate of 20 Gb/s to merely 5% when connection bit rate increases to 70 Gb/s, even though the offered load in both scenarios is adjusted to equally achieve 2% blocking probability. In contrast, with a defragmentation interval of 28, 10 times lower, improvement only decreases from 80% to around 70% when connection bit rates change from 20 to 70 Gb/s, respectively.

Another conclusion that can be extracted from Figure 7 is the paramount importance of the defragmentation

interval value on the improvements that periodic defragmentation can yield. Indeed, more frequent defragmentation fosters improvement in blocking probability, although introducing more control and management burden in the network, as already mentioned in the paper. An interesting outcome of this figure, however, is that taking the bit rate of the connections into account, a network operator can intelligently set the periodicity of the defragmentation in the network. For example, to achieve 40% improvement with a bit rate of 70 Gb/s, a defragmentation interval of 70 is required. Nonetheless, if the bit rate of the connections appears to be 20 Gb/s, the defragmentation interval can be increased to 280, as only 1/4 of the defragmentations are necessary in this scenario to achieve the same improvement. This could also open the possibility for operators to implement monitoring strategies of the incoming connections bit rate (e.g. during pre-defined time periods or a number of connection establishments) to dynamically adjust the defragmentation interval in the network, making a good trade-off between blocking probability and control and management burden.

To round off the paper, Figure 8 shows in more detail the effect of the defragmentation interval on the performance of periodic defragmentation. To this end, we have offered a load of 10 per node with average bit rate per connection of 45 Gb/s, which leads to a blocking probability of around 1% with on-demand defragmentation. In such a scenario, we plot the blocking probability versus defragmentation interval for no defragmentation and on-demand and periodic defragmentation. Of course, the former two behave constant with the defragmentation interval, as their performance does not depend on this parameter. However, we use them in the figure as a benchmark. As shown, when increasing the defragmentation interval, the possibility of successfully reallocating connections in periodic defragmentation decreases. This could be explained from the fact that the randomness of the connection arrivals and departures creates such a bad network situation during each defragmentation period that the defragmentation periodically performed cannot re-arrange it in a better way. In any case, periodic defragmentation always stays




Figure 8. Network blocking probability against defragmentation interval in periodic defragmentation.

below no defragmentation for any of the evaluated intervals. Specifically, to get similar performance as on-demand defragmentation (or better), the defragmentation interval has to be set around 50 (or lower). Therefore, provided that defragmentation intervals below 50 are acceptable for a network operator, periodic defragmentation can be a good option against the on-demand one, which has to be forcefully simpler to avoid incoming demands experiencing high establishment delays.

5. CONCLUSION

The recently proposed elastic optical network that provides enhanced flexibility in spectrum allocation and data rate accommodation has opened up a new prospect to serve the future Internet demands more efficiently. However, the inevitable spectral fragmentation problem in such networks significantly degrades their performance. To address the problem, the idea of spectral defragmentation in elastic optical network has been proposed. Assuming a periodic defragmentation scenario, this paper discussed the effects of the offered traffic characteristics (offered load and bit rate per connection), as well as the defragmentation periodicity, on the profitability that defragmentation can yield to network operators. Indeed, while decreasing the defragmentation interval has direct impact on improving the blocking probability in the network, taking into account the bit rate of the connections can also lead to the same results but with more relaxed defragmentation intervals and, thus, lower network control and management burden. All these findings were supported by extensive simulation results, where no defragmentation and on-demand defragmentation scenarios were also contemplated as a benchmark. For future work, authors will derive analytical models to formally quantify the effect of traffic granularity on the periodicity of defragmentation operation.

ACKNOWLEDGEMENTS

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. Palkopoulou E, Merkle C, Schupke DA, Gruber CG, Kirstdter A. Traffic models for future backbone networks a service-oriented approach. *European Transactions on Telecommunications* 2010; **22** (4): 137–150. DOI: 10.1002/ett.1464
2. *Cisco white paper: cisco visual networking index: forecast and methodology 2011-2016*, 2012.
3. Jinno M, Takara H, Kozicki B, Tsukishima Y, Sone Y, Matsuoka S. Spectrum-efficient and scalable elastic

- optical path network: architecture, benefits, and enabling technologies. *IEEE Communications Magazine* 2009; **47** (11): 66–73. DOI: 10.1109/MCOM.2009.5307468
4. Farjady F, Kim KS, El Dahdah N, Doran NJ. Cost-effective upgrade of WDM all-optical networks using overlay fibres and hop reduction links. *European Transactions on Telecommunications* 2010; **21** (6): 563–566. DOI: 10.1002/ett.1444
 5. Jinno M, Takara H, Kozicki B. Dynamic optical mesh networks: drivers, challenges and solutions for the future. In *Proceeding of 35th European Conference on Optical Communication (ECOC 2009)*, Vienna, Austria, September 2009; 1–4.
 6. Christodoulopoulos K, Tomkos I, Varvarigos E. Spectrally/bitrate flexible optical network planning. In *Proceeding of 36th European Conference on Optical Communication (ECOC 2010)*, Turin, Italy, September 2010; 1–3.
 7. Wang R, Mukherjee B. Spectrum management in heterogeneous bandwidth networks. In *IEEE Global Communications Conference (GLOBECOM)*, Anaheim, CA, USA, 2012; 2907–2911. DOI: 10.1109/GLOCOM.2012.6503558.
 8. Patel AN, Ji PN, Jue JP, Wang T. Defragmentation of transparent flexible optical WDM (FWDM) networks. In *Proceeding of Optical Fiber Communication Conference and Exposition (OFC)*, Anaheim, CA, USA, 2011; 1–3. DOI: 10.1364/OFC.2011.OTuI8.
 9. Geisler DJ, Yawei Y, Ke W, Fontaine NK, Scott RP, Shuo C, Yoo SJB. Demonstration of spectral defragmentation in flexible bandwidth optical networking by FWM. *IEEE Photonics Technology Letters* 2011; **23**(24): 1893–1895. DOI: 10.1109/LPT.2011.2171933
 10. Takagi T, Hasegawa H, Sato K, Sone Y, Hirano A, Jinno M. Disruption minimized spectrum defragmentation in elastic optical path networks that adopt distance adaptive modulation. In *Proceeding of 37th European Conference on Optical Communication (ECOC 2011)*, Geneva, Switzerland, September 2011; 1–3. DOI: 10.1364/ECOC.2011.Mo.2.K.3.
 11. Wen K, Yawei Y, Geisler DJ, Shuo C, Yoo SJB. Dynamic on-demand lightpath provisioning using spectral defragmentation in flexible bandwidth networks. In *Proceeding of 37th European Conference on Optical Communication (ECOC 2011)*, Geneva, Switzerland, September 2011; 1–3. DOI: 10.1364/ECOC.2011.Mo.2.K.4.
 12. Yin Y, Wen K, Geisler DJ, Liu R, Yoo SJB. Dynamic on-demand defragmentation in flexible bandwidth elastic optical networks. *Optics Express* 2012; **20** (2): 1798–804. DOI: 10.1364/OE.20.001798
 13. Sone Y, Hirano A, Kadohata A, Jinno M, Ishida O. Routing and spectrum assignment algorithm maximizes spectrum utilization in optical networks. In *Proceeding of 37th European Conference on Optical Communication (ECOC 2011)*, Geneva, Switzerland, September 2011; 1–3. DOI: 10.1364/ECOC.2011.Mo.1.K.3.
 14. Geisler DJ, Proietti R, Yin Y, Scott RP, Cai X, Fontaine N, Paraschis L, Ori Gerstel O, Yoo SJB. The first testbed demonstration of a flexible bandwidth network with a real-time adaptive control plane. In *Proceeding of 37th European Conference on Optical Communication (ECOC 2011)*, Geneva, Switzerland, September 2011; 1–3. DOI: 10.1364/ECOC.2011.Th.13.K.2.
 15. Amaya N, Irfan M, Zervas GS, Banias K, Garrich M, Henning I, Simeonidou D, Zhou YR, Lord A, Smith K, Rancano V, Liu S, Petropoulos P, Richardson D. Gridless optical networking field trial: flexible spectrum switching, defragmentation and transport of 10G/40G/100G/555G over 620-km field fiber. In *Proceeding of 37th European Conference on Optical Communication (ECOC 2011)*, Geneva, Switzerland, September 2011; 1–3. DOI: 10.1364/ECOC.2011.Th.13.K.1.
 16. Yu X, Zhang J, Zhao Y, Peng T, Bai Y, Wang D, Lin X. Spectrum compactness based defragmentation in flexible bandwidth optical networks. In *Proceeding of Optical Fiber Communication Conference and Exposition (OFC)*, Anaheim, CA, USA, 2012; 1–3. DOI: 10.1364/NFOEC.2012.JTh2A.35.
 17. Christodoulopoulos K, Tomkos I, Varvarigos E. Routing and spectrum allocation in OFDM-based optical networks with elastic bandwidth allocation. In *IEEE Global Communications Conference (GLOBECOM)*, Miami, FL, USA, 2010; 1–6. DOI: 10.1109/GLOCOM.2010.5684008.
 18. Xia M, Proietti R, Dahlfort S, Yoo SJB. Split spectrum: a multi-channel approach to elastic optical networking. *Optics Express* 2012; **20** (28): 29143–29148. DOI: 10.1364/OE.20.029143