

Using Spectrum Management Techniques to Differentiate Services in Elastic Optical Networks

J. Comellas, A. Asensio, M. Ruiz, G. Junyent, and L. Velasco

*Optical Communications Research Group (GCO), Universitat Politècnica de Catalunya (UPC),
Jordi Girona 1-3, 08034 Barcelona (Spain), Tel: (+34) 93 401 6441, Fax: (+34) 93 401 7200*

Email: comellas@tsc.upc.edu

ABSTRACT

Elastic Optical Networks (EONs) have emerged as a promising solution to cope with the growth and heterogeneity of networks data traffic. Nevertheless, several problems arise in dynamic scenarios, being spectrum fragmentation the most important one. Under a fragmented spectrum, connections with large bandwidth experience higher problems in finding enough free bandwidth to be set-up. In this work, creating spectrum sub-bands dedicated to different types of connections is used to provide different blocking probability values to these sets of connections. By properly adjusting the design and use of these sub-bands, it is possible to tune the network performance attaining different blocking probability (BP) values as a function of the spectrum band division and the priority given to the different connection requests sizes. The purpose of this paper is to demonstrate that different BP values can be obtained for different classes of traffic which, in turn, are assigned to different connections sizes. The strategies proposed are assessed by means of simulation.

Keywords: elastic optical networks, routing and spectrum assignment, spectrum partition.

1. INTRODUCTION

Owing to emerging services such as high-definition video distribution or social networking, the IP traffic volume has shown an exponential increase in the recent years. Furthermore, the traffic growth rate will not stop here thanks to the endless technology advances [1]. The predictable consequence is that network operators will require a new generation of optical transport networks, so as to serve this huge and heterogeneous volume of traffic in a cost-effective and scalable manner. 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 limit of conventional wavelength division multiplexing (WDM) networks, such EONs increase the flexibility in terms of connection provisioning. To do so, depending on the traffic volume, an appropriate-sized optical spectrum is allocated to each connection. The total available spectrum is divided into constant spectrum units referred as frequency slots (FS), whose typical width is 12.5 GHz. Each FS can carry some amount of traffic mainly depending on the modulation format used and the required reach [2]. For the sake of simplicity, this work is only focused on the spectrum consumed (number of FS), avoiding the spectral efficiency problem. Each connection is served by allocating a set of contiguous FS.

Nevertheless, the EON model has some functional drawbacks, and the so called spectrum fragmentation [3] is the most important one. The randomness in the connection setup and tear down processes leads to fragmentation of the spectral resources in the network. As the number of spectrum slots assigned to each connection can take random values (ranging from 1 to 10 FS in our model), the available spectrum in the network links is fragmented into small non-contiguous spectral bands. As a result, the probability of finding enough contiguous spectrum resources for serving incoming traffic demands, especially those traversing multi-hop paths and/or requesting large amounts of bandwidth (BW), decreases significantly.

This work provides a simple proposal to alleviate the fragmentation effects. By creating special spectrum sub-bands (partitions) where selected connections sizes have priority, it is demonstrated that the blocking probability (BP) suffered by these connections is modified, at the cost of slightly increasing the average network BP. Using this strategy, different scenarios are simulated and the final conclusion is that it can be useful in protecting selected parts of the traffic carried by the EON. The idea is to assign high priority traffic to one of the possible connections size values, which is somehow protected by using a dedicated spectral window larger than the statistically needed.

The rest of the paper is organized as follows. In Section 2, previous work about creating spectrum sub-bands to improve spectrum assignment in EONs is reviewed, and the strategy followed in this work is described. The simulation scenarios considered, as well as the results obtained are presented in Section 3. Finally, Section 4 concludes the work.

2. CREATING SPECTRUM SUB-BANDS

Finding ways to alleviate fragmentation in EONs has become a really crowded research topic during last years. Most of the works concerning routing and spectrum assignment (RSA) in EONs must appropriately tackle this issue.

A lot of efforts have been directed to de-fragmentation strategies [4], but using RSA algorithms that prevent fragmentation [5], and even further fragmenting the traffic [6] can also help in alleviating its effects. Spectrum partitioning was introduced in [7] as a simple way to efficiently manage some problems arising in EONs, which are there classified as:

1. Fairness issues, due to the fact that small BW connection requests have higher probability of finding enough contiguous FS than larger ones.
2. Fragmentation, already defined above, which is worsened as more heterogeneous traffic is offered.

Authors in [7] proposed *Dedicated partitioning*, where connections are allocated in predefined segments of the spectrum (partitions, only used by a class of connections), and *Shared partitioning*, which allows to allocate a connection on a partition assigned to other traffic classes when its partition is full.

In this work, a variation of *shared partition* is used. Some connection sizes have priority in using a given partition. When this class of connections (class and connections BW are equivalent throughout this work) do not find room to be allocated, existing connections of other classes are re-allocated to another partition or even torn-down. In other words, there is a spectrum band where these connections have priority, so they finally experience an altered blocking probability (BP) value. Using this strategy it is possible to change the BP characteristics of EONs, which, as stated before, always show higher BP for higher BW connections [7]. As it is shown in Section 3, BP experienced by small connections is almost zero for a typical EON, while it takes considerable values for large connections. Therefore, a simple strategy which consists in assigning a window bigger than the statistically expected to large connections provides fairness to the network as BP values for different connections sizes are equalized. An alternative approach, also tested during simulations Section, consists in penalizing small connections. By assigning a partition smaller than needed to connections with small sizes it is foreseeable that larger connections will have more chances to find enough free spectrum slots to be accommodated.

Nevertheless, the proposals introduced in this work, based on partitioning the spectrum, have a drawback in the fact that fragmentation is somehow duplicated. When compared to work presented in [7], we are considering a fully elastic scenario where connections can take any BW value from 1 to 10 FS, in front of the multi-rate case (only three BW values were allowed) supposed there. Therefore, as the policies here presented consist in assigning different bands of the spectrum to different types (BW values) of connections, fragmentation will indeed appear in any of the partitions, thus arising as a clear weakness of our system. The structure assumed can lead to situations where, in a given partition, there is a fragment that cannot be used by connections assigned to that partition but could be utilized by connections that are not allowed to use that part of the spectrum. At the end, this poses our system under more stringent conditions and implies that the benefits of using this spectrum partition policy are limited, and the partition sizes have to be carefully selected.

3. SIMULATION FRAMEWORK AND RESULTS

The performance of the abovementioned proposals are evaluated through extensive simulations using an ad-hoc Matlab® simulator whose main characteristics are:

- Maximum number of FS per link equal to 160. This is a typical value used in previous works, and it has no special relevance on the results obtained.
- Number of FS per connection (BW) ranging from 1 to 10, with uniformly distributed probability so each connection can take any of the allowed BW values.
- Connections inter-arrival time (IAT) follows a Poisson process with average value equal to 1. Connections holding time (HT) follows a negative exponential distribution. The average value is adjusted to obtain the appropriate Offered Load values (those that led the link bandwidth Blocking Probability to the desired values).
- Network topology. The well-known 14-nodes 21-links NSFnet topology has been selected. Source-Destination pairs are uniformly distributed among all the networks' nodes. Connections are considered unidirectional.
- Regarding the routing and spectrum allocation algorithm for allocating new connections, simulations run a typical k -Shortest Path (with $k = 3$) routing and First-Fit SA [8].

It has to be highlighted that the Offered Load values shown in the results represent load per link, considering that the maximum load would be given by 160 divided by the average number of hops per connection. Taking into account that every connection spreads over N_{hops} links, a perfectly loaded network would carry $160 \times N_{\text{links}} / N_{\text{hops}}$ FS, so the Offered Load values are normalized to this maximum load.

First simulations correspond to the described EON without applying any spectrum use improving strategy. They clearly show (see Fig. 1) that BP values are higher for large connections. Only connections sizes greater than 4 FS are shown because BP for smaller connections are below 10^{-3} for all the simulated load values. It has to be highlighted that the average BP coincides almost perfectly with the BP for connections with $BW = 8$ FS. Easy network monitoring could be implemented using this fact, but this is out of the scope of this work. The important conclusion from Fig. 1 is that, as expected, BP clearly rises with as connections BW values do.

A simple strategy to alleviate this unfair behavior arises: by creating a partition larger than the strictly necessary for big connections or, alternatively, limiting the spectrum band where small connections can be allocated, BP values for different connections sizes will be better balanced. Therefore, policies applied should follow this idea.




Figure 1. Blocking probability (BP) vs. Offered Load for different connections sizes in EON. Average network bandwidth BP is also represented. NSFnet topology is considered.

A first example is shown in Fig. 2. A special partition has been created where connections with $BW = 8$ FS have priority. By properly adjusting the size of this spectral band, the BP values of this connection class is tuned.




Figure 2. Effect of creating a spectrum partition of 54 FS where connections with $BW = 8$ FS have priority. Graphs show BP (in logarithmic scale) vs. Offered Load per link for different connections sizes (Left: without partition; Right: with partition).

In the example shown in Fig. 2, a 54 FS band has been created where 8 FS are preferably allocated. This band (named *DB*, from Dedicated Band in the Figures) is only used for other connections types when there is no available space on the remaining spectrum. Indeed, these different connections are removed from that band (either re-allocated in the regular part or torn-down) when new 8 FS connections arrive. The obvious result is that 8 FS connections experience a lower BP than before (it is kept almost constant around $10^{-2.7}$). The negative part is that the rest of the connections sizes BP values increase. If we look at the load that leads the average BP to 10^{-2} , it changes from $Load = 0.58$ (before applying this strategy) to $Load = 0.54$ (when this partition is created). In summary, one of the connection types is protected at the cost of deteriorating the whole network performance. It has to be noted that although 8 FS connections have been chosen in Fig. 2 example, the same strategy could be applied to any connections BW value.

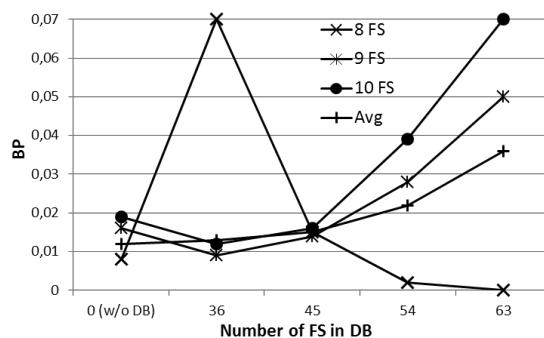


Figure 3. Effect of special partition size on the network performance.

A subsequent activity consisted in measuring the effect of the *DB* size. Results obtained are shown in Fig. 3, where the *Offered Load* value has been fixed to 0.56, which gives an average *BP* slightly lower than 10^{-2} . The size of the special partition (where 8 FS connections have priority) takes different values (the case without *DB* is also shown).

The effects of this *DB* size tuning can be clearly appreciated in Fig. 3, where 8, 9 and 10 FS as well as average *BP* values are represented. It can be concluded that assigning 45 FS to *DB* is a trade-off solution. Although the average *BP* is slightly increased at this point, *BP* values for all the connections in the network converge in a value near 0.015. For higher *DB* values 8 FS connections are protected (experiencing lower *BP*), but other sizes are penalized.

As abovementioned, small connections ($BW < 4$ FS) experience always $BP < 10^{-3}$ at the simulated load values. An alternative strategy to equalize *BP* values consists in limiting the spectrum band assigned to these connections. Figure 4 shows the effects of confining connections with $BW < 4$ FS to a dedicated partition. Once again, the case without *DB* is also shown for benchmarking purposes.

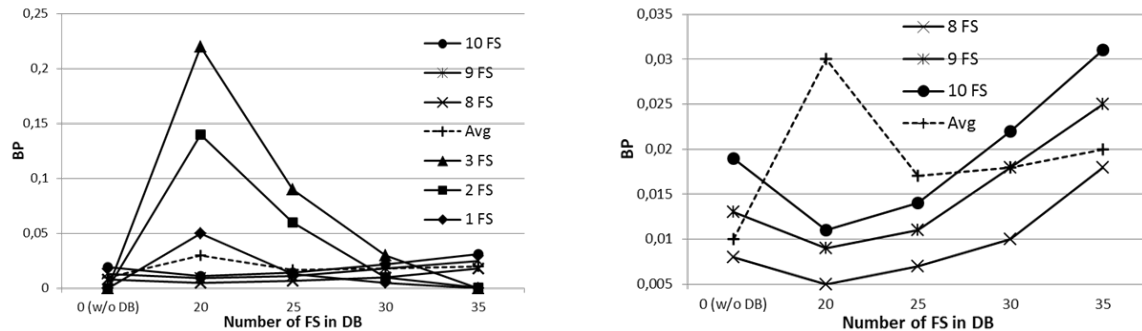


Figure 4. Effects of limiting the dedicated band (*DB*) assigned to connections with $BW < 4$ FS.

When *DB* is small, 1, 2 and 3 FS connections are clearly penalized and *BP* for these connections reaches unacceptable high values (Fig. 4, left). Right graph shows the evolution of *BP* for 8, 9 and 10 FS connections as well as the average *BP* values. When $DB = 20$ FS the *BP* values for large *BW* connections is clearly reduced but the cost is on the prohibitive *BP* values for 1 and 2 FS connections. For $DB = 25$ FS *BP* for larger connections is still better than without *DB* and it has been reduced to less than 0.1 for 1 FS connections. All in all, a scenario where best-effort traffic is assigned to penalized connection sizes while golden traffic uses the protected sizes could provide better *BP* results than a typical EON where there are no differences between different kinds of connections.

4. CONCLUSION

In this work, the use of spectrum partitions in EON to adjust the *BP* experienced by different connection *BW* sizes is proposed. Stemming from the fact that allocating large connections is more difficult in EONs, a scheme to solve this unfairness problem is provided. The obtained simulation results show that it is possible to adjust the *BP* of a selected connections size to any chosen value by appropriately sizing the spectrum partitions. Nevertheless, the global network performance is at some extent deteriorated when applying this strategy.

ACKNOWLEDGEMENTS

This work has been supported by the Spanish National project SYNERGY (TEC2014-59995-R).

REFERENCES

- [1] Cisco White Paper, "Cisco Visual Networking Index: Forecast and Methodology, 2013-2018".
- [2] 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, 2012.
- [3] Z. Zhu, W. Lu, L. Zhang, and N. Ansari, "Dynamic service provisioning in elastic optical networks with hybrid single-/multi-path routing," *IEEE/OSA J. Lightw. Technol.*, vol. 31, no. 1, pp. 15-22, 2013.
- [4] R. Wang and B. Mukherjee, "Provisioning in elastic optical networks with non-disruptive defragmentation," *IEEE/OSA J. Lightw. Technol.*, vol. 31, no. 15, pp. 2491-2500, 2013.
- [5] P. S. Khodashenas, J. Comellas, S. Spadaro, J. Perelló, and G. Junyent, "Using spectrum fragmentation to better allocate time-varying connections in elastic optical networks," *Journal of Optical Communications and Networking*, vol. 6, no. 5, pp. 433-440, 2014.
- [6] A. Pagès, J. Perelló, S. Spadaro, and 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.
- [7] R. Wang and B. Mukherjee, "Spectrum management in heterogeneous bandwidth optical networks," *Optical Switching and Networking*, vol. 11, pp. 83-91, 2014.
- [8] 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.