A Comparison of Elastic and Multi-Rate Optical Networks Performance

P. S. Khodashenas, D. Pomares, J. Perelló, S. Spadaro, and J. Comellas
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
While Elastic Optical Networks (EONs) have recently emerged as a promising solution to cope with the growth and heterogeneity of data traffic, there are some drawbacks that have attracted the researchers' attention. One of such flaws is spectrum fragmentation, which has generated many controversial as it imposes huge number of extra actions during network operation. Some intermediate proposals have been disclosed, such as semi-elastic schemes that approach the performance of EONs while keeping the spectral entropy restrained. The purpose of this paper is to present a comparison between EONs and semi-elastic networks, where all the offered connections are allocated using only three different channel rate options. Different spectrum management strategies are introduced and evaluated by means of simulation considering both scenarios.
Keywords: elastic optical networks, network optimization, spectrum fragmentation.

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 continuous technology advances [1]. 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. 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 elastic optical networks 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 in EON. In this way, incoming network connection requests can be served in a spectrum efficient manner. Nevertheless, this EONs spectrum tailored model has some functional drawbacks, being the most important the so called spectrum fragmentation [3]. 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 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, significantly decreases.

To mitigate this problem a pseudo-elastic network is considered in this work. The traditional EON scenario is compared to a semi-flexible multi-rate network (MON) model where connections are adapted to a kind of fixed rate transmitters with different bit rate values (only 1, 4 or 8 transmitted spectrum slots are allowed in our model). By losing some flexibility (the incoming traffic demands are yet asking for a random number of slots between 1 and 10), some spectral entropy reduction is achieved. Both, EON and MON models are compared by means of simulation and some conclusions about their performance are obtained.

2. MON NETWORK MODEL

As stated in the previous section, the main purpose of this work consists in reducing the fragmentation problem introduced by EONs. To do so, instead of employing fully bandwidth variable transponders at network ingress nodes (able to support arbitrary number of slots), a set of transmitters able to generate 1, 4 or 8 slots are utilized in MONs. It is worth to note that, as stated previously, incoming traffic demands may ask for any number of slots from 1 to 10. Therefore, in order to solve the mismatch problem between the requested bandwidth and the transponders' capacity a mapping scheme has to be introduced. As shown in Table 1,we proposed two different mapping strategies:

Table 1. Connections splitting used for the MON scenario. Two different cases are considered: fit splitting (Down mapping) and loose splitting (Up mapping).

Incoming traffic demand	Down mapping	Up mapping
1	1	1
2	1,1	1,1
3	1,1,1	4
4	4	4
5	4,1	4,1
6	4,1,1	8
7	4,1,1,1	8
8	8	8
9	8,1	8,1
10	8,1,1	8,4

978-1-4799-5601-2/14/$31.00 ©2014 IEEE

1) *fit splitting* (down mapping) where the incoming traffic demand is split into smaller sub-demands that adjust the available transponders capacity; and 2) *loose splitting* (up mapping) which has the objective of reducing the number of transponders used even at the cost of wasting some bandwidth. As a matter of fact, since down mapping slices an incoming traffic demand into smaller pieces, it may need more transmitters to allocate it comparing to the Up mapping case, which sacrifices some spectrum to achieve transponders savings. This can be observed in Table 1 where an incoming traffic demand of 7 slots uses 4 transmitters for down mapping while and only 1 for up mapping, at the cost of wasting 1 slot. It is important to note that, in order to guarantee the QoS of network, all the sliced sub-demands have to be accommodated in the same path between connection's end nodes.

3. SPECTRUM ALLOCATION POLICIES

To improve the performance of both EON and MON networks, it is possible to consider different spectrum allocation policies. Authors in [4] proposed the following policies:

- **First Fit:** Connections are established over the lowest available part of spectrum, so this strategy would result in a perfectly compacted spectrum in case of static traffic. This policy is widely used in EONs.
- **Pseudo partitioning:** It is a simple variation of the First Fit policy. The difference is that depending on their size, connections are allocated in either left-to-right direction or right-to-left direction over the whole available spectrum. To do so, small connections (with size of 1 slot) are allocated in left-to-right direction, while big connections (with size of bigger than 1 slot) are allocated in right-to-left direction. In this sense, the Pseudo partitioning policy provides two separate spectrum segments, reserved only for the small or big connections, which can be functionally considered as different traffic classes. It is worth to note that for high network load values, the gap between both segments shrinks to zero, and they start to overlap.
- **Dedicated partitioning:** Connections are allocated in predefined segments of the spectrum (partitions). The size of each partition is calculated respecting to the size of connections in terms of number of slots and the total offered load to the network. In the event that the partition of the target traffic class has not enough free slots to perform the allocation, the connection is dropped, regardless of having spectrum availability in other partitions.
- **Shared partitioning:** It is a variation of the previous policy. The difference is that once the partition of the target traffic class is full, instead of dropping the connection, there will be an attempt to allocate it on the partition assigned to any other traffic class whose occupation is inferior to the target traffic class.

Simulations about the performance of both network models (EON and MON) as well as the different spectrum management strategies considered are given in next section.

4. SIMULATION RESULTS

The performance of EON and MON networks are evaluated through extensive discrete event simulation studies. Both cases use a k-Shortest Path routing algorithm with spectrum assignment, starting with the shortest computed path. In addition, the different explained spectrum allocation policies in the previous section have been considered for both cases. The well-known 14-node NSFnet topology has been selected for the simulation purposes. A total optical spectrum of 1.5 THz per link and a spectrum slot size of 12.5 GHz are assumed. For the sake of simplicity, the modulation format selected yields a spectrum efficiency of 1 bit/s/Hz, so each spectrum slot has a bit rate capacity of 12.5 Gb/s. As for the traffic characteristics and according to the asymmetric nature of today's Internet traffic, unidirectional connections between end nodes are considered. The traffic generation follows a Poisson distribution process, so that different offered load values are obtained by keeping the mean Holding Time (HT) of the connections constant to 200s, while modifying their mean Inter-Arrival Time (IAT) accordingly (*i.e.*, offered load = HT/IAT). Traffic demands for each source-destination pair are randomly generated by normal distribution ranging from 12.5 Gb/s (1 frequency slot) to 125 Gb/s (10 frequency slots). The average traffic demand is used to study the relationship between aggregation policy efficiency and service granularity.

A traffic load in the range of 13 up to 16 Erlang per node (which provides a total offered network traffic ranging from 182 to 224 Erlang) has been used in the simulations. The average demand of each connection request is assumed to be 55 Gb/s. Hence, the total traffic generated per node ranges from 715 Gb/s to 880 Gb/s in this study. In addition, an initially unlimited number of transmitters per node has been considered; the effect of the number of transmitters on the performance of each case is investigated later.

First results obtained (which are related to the different spectrum management strategies utilized) concluded that the *shared partitioning* scheme outperforms the rest when MONs are considered. For example, when the load per node is around 800 Gbit/s, the blocking probability value for *First Fit* is 0.04 while it is below 0.02 for *shared partitioning*. So, taking that into account, this strategy has been chosen for the remaining MON model simulations.

Next results, once selected the spectrum management strategy, are related to the mapping of the flexible connections to the semi-flexible (1, 4 or 8 slots) transponders of MON networks. As it is shown in Fig. 1, the MON case with down mapping outperforms both other cases in the whole range of study. The reason is that, when compared to the EON case, it reduces the level of fragmentation over the network links. However, the high number of transponders needed in this scenario is a clear drawback when compared to the other cases as it is shown later. It has to be taken into account that in the worst case (MON with Up mapping), there is some overprovisioning of bandwidth that could be useful in case of time varying connections. Some allocated spectrum slots (*e.g.* 8 are allocated when the connection bandwidth is 6) would be really useful if the connection bandwidth is increased during its HT.

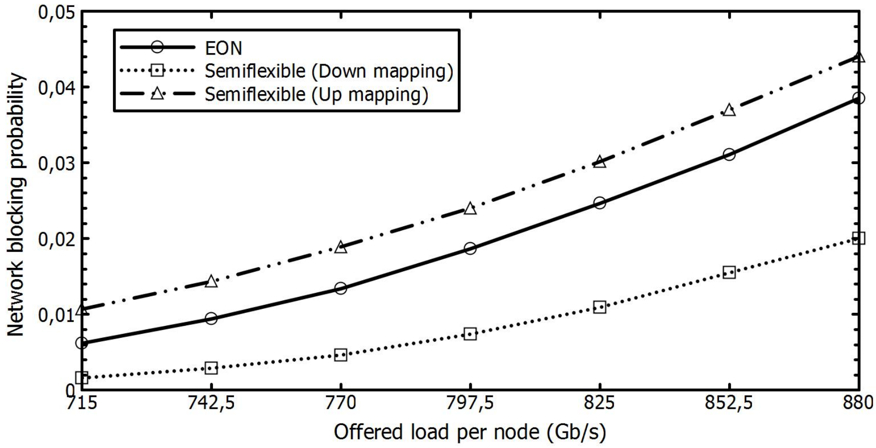

Figure 1. Blocking probability vs. offered load per node for EON and MON (with up and down mapping) networks.

The effect of traffic granularity on the performance of the aforementioned scenarios for a fixed offered load per node is shown in Fig. 2. In this case the average number of active optical connections from each source node is kept to 14.5 Erlang, but the average number of slots per connection is increased from 3.5 to 5.5. It can be observed that again the MON with down mapping outperforms the other cases. It is worth to note that the effectiveness of down mapping proposal is more significant comparing to the up mapping case when moving towards higher connections' bandwidth. The reason is that the possibility of accommodating smaller connections in a fragmented spectrum (due to an increased network load), is higher.

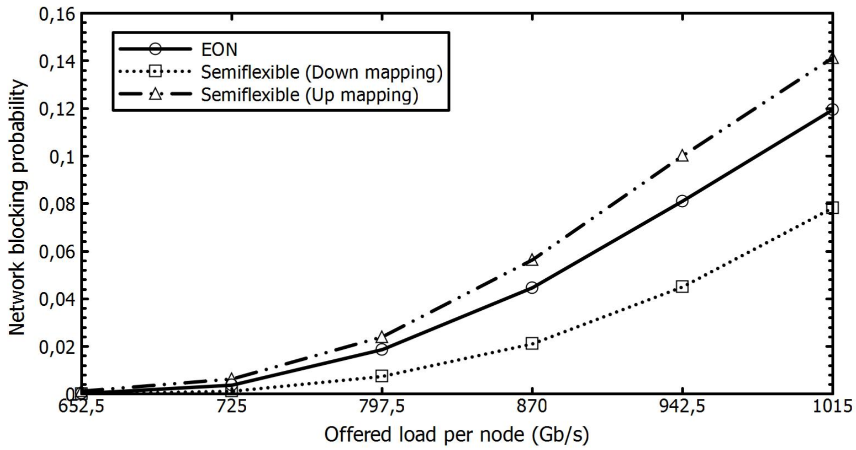

Figure 2. Effect of connections' bandwidth on the blocking probability. The average number of connections in the network is kept constant while their bandwidth grows (from 3.5 to 5.5 slots).

Once we have concluded, not surprisingly, that the MON with down mapping offers the best performance in terms of blocking probability, the number of available transponders per node has been limited. The effect of this limitation on the performance of the network is shown in Fig. 3. In this study, the fixed load of 14.5 Erlang with an average bit rate value per connection of 55 Gb/s (or 4.4 spectrum slots) is assumed. From the previous results, the corresponding blocking probability under this load value was lower than 0.01 for all the scenarios considered.

As it is illustrated, all cases perform almost equally while the number of transmitters per node is higher than 40. When this number is reduced, MON with up mapping approaches the performance of EON for

30 transmitters per node, while the MON with down mapping blocking probability is already higher than the acceptable values.

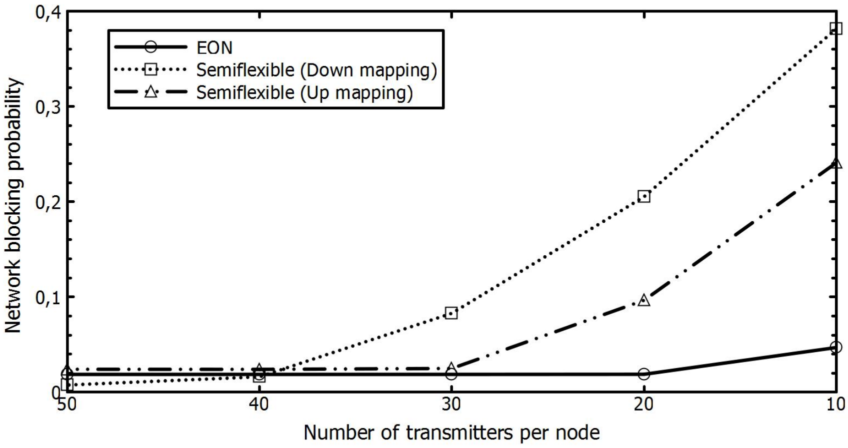

Figure 3. Effect of limiting the number of transponders per node on the network blocking probability.

It can be seen that, if this number is further reduced, only the EON network performs appropriately. The relative cost of flexible transponders in front of fixed ones would therefore give us light about the worthiness of EON networks. Preliminary results show that with current values (cost of 20 flexible transponders is today higher than cost of 30 fixed ones) EON is not yet cost effective. In addition, cost of switching nodes is higher if they have to adapt to fully flexible transponders. However, if future technology allows that the cost of flexible transponders [5] reduces to less than 1.5 times the cost of fixed ones, EON will be the most effective solution.

5. CONCLUSION

In this paper, a comparison between EON (transponders can use any number of spectrum slots) and semi-flexible (transponders can allocate only 1, 4 or 8 slots) networks performance has been done. Different strategies for spectrum occupation in this kind of networks have been evaluated by means of simulation. The obtained results show that semi-flexible can perform better in terms of blocking, due to the fact that spectrum fragmentation is reduced and more connections can be allocated for medium to high offered load values. However, when the number of transponders per node is limited, the EON networks show some advantages.

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] Cisco White Paper, "Cisco Visual Networking Index: Forecast and Methodology, 2011-2016".

[2] M. Jinno, H. Takara, B. Kozicki, Y. Tsukishima, Y. Sone, and S. Matsuoka, "Spectrum-efficient and scalable elastic optical path network: Architecture, benefits, and enabling technologies", *IEEE Communications Magazine*, vol. 47, no. 11, pp. 66-73, 2009.

[3] P. S. Khodashenas, J. Comellas, J. Perelló, and S. Spadaro, "Correlation between traffic granularity and defragmentation periodicity in elastic optical networks", *Transactions on Emerging Telecommunications Technologies*, Article first published online, DOI: 10.1002/ett.2795, 2014.

[4] R. Wang and B. Mukherjee, "Spectrum management in heterogeneous bandwidth optical networks", *Optical Switching and Networking*, vol. 11, pp. 83-91, 2014.

[5] K. Christodoulopoulos, M. Angelou, D. Klonidis, P. Zakynthinos, E. Varvarigos, and I. Tomkos, "Value analysis methodology for flexible optical networks", in *Proc. ECOC 2011*, Geneva, Switzerland, paper We.10.P1, 2011.