

Dynamic Traffic Grooming in Joint Switching (JoS)-enabled Flex-Grid/SDM Optical Core Networks

R. Rumipamba-Zambrano⁽¹⁾, J. Perelló⁽¹⁾, and S. Spadaro⁽¹⁾

⁽¹⁾ Universitat Politècnica de Catalunya (UPC), Barcelona, Spain, e-mail: rrumipam@ac.upc.edu

Abstract We propose and evaluate a dynamic traffic grooming strategy enhancing the performance of JoS-enabled Flex-Grid/SDM networks. Results disclose up to 120% and 44% additional carried network traffic vs. no grooming and existing end-to-end traffic grooming solutions, respectively.

Introduction

Flex-Grid¹ and Space Division Multiplexing² (Flex-Grid/SDM) are key technologies to cope with the rapid growth of the Internet traffic, thanks to their superior capacity and spectrum utilization flexibility. As pointed-out in the SDM literature, its near-term realization will rely on the joint switching (JoS)³ technique, switching one spectrum portion across all spatial channels at once. Using this technique, all spatial channels are treated as a single entity and, therefore, the space dimension can be considered as fixed. Meanwhile, a long-term SDM realization will introduce flexibility in both spectral and space dimensions, thus coming up with spectrally and spatially flexible optical networks requiring space and spectrum switching granularity, known as independent switching (InS)³.

Flex-Grid/SDM networks realizing JoS reduce the complexity and cost of nodes by requiring a lower amount of spectrum selective switches⁴, forcing to allocate sub-channels in the form of spatial super-channels (Spa-SChs). Spa-SChs generated/detected by flexible transponders further reduce costs by avoiding frequency combs, sharing laser sources and facilitating the joint digital signal processing at receivers⁴. Alternative SCh configurations⁵ are spectral (Spe-SChs) and spectral-spatial super-channels (S2-SChs), commonly related to InS.

However, JoS can lead to spectrum wastage because of its spatial rigidity. As a result, the performance of JoS-enabled Flex-Grid/SDM networks can be negatively affected⁶. This spectrum wastage (the spatial-spectral resources left unused after the allocation of demands) can be leveraged by co-routed demands, performing end-to-end spatial traffic grooming (e2e-grooming)⁶. This work explores the fittest SCh configuration during the e2e-grooming operation. We name this strategy as *dynamic* e2e-grooming and we evaluate its benefits under several traffic conditions.

End-to-End Spatial Traffic Grooming

Thanks to the spectrum flexibility given by Flex-Grid, for the same demand d , different SCh

configurations are valid to allocate it. Specifically, given a fiber with S spatial channels different (n_s, n_{fs}) tuples are possible to serve d , being n_s the required number of spatial channels and n_{fs} the required number of frequency slots (FSs) per spatial channel. For instance, for a demand with bit-rate $r_d = 400$ Gb/s using Dual-polarized (DP) QPSK modulation format with Spectral Efficiency (SE) at the Nyquist limit (4 b/s-Hz), $S=9$, FS width (W) = 12.5 GHz and 7.5 GHz guard-band (GB), SCh candidates (Ω), expressed as a set of (n_s, n_{fs}) tuples are: $\Omega = \{(1,9), (2,5), (3,4), (4,3), (6,2)\}$, where $n_{fs} = \lceil (r_d / (n_s \cdot SE) + GB) / W \rceil$.

The typical goal in JoS-enabled Flex-Grid/SDM networks is to allocate the finest spectrum portion across n_s spatial channels in order to maximize the network Grade-of-Service⁷. Therefore, among the SCh candidates in the previous Ω example, the last tuple (6,2) would be the most suitable one, which would also be obtained following the procedure known as Partial Core (or Space) Assignment (PCA)⁷. Depending on the traffic profile, it has been demonstrated that JoS performance can be significantly worse than that of InS⁸ due to its spectrum wastage. In order to enhance it, some strategies have been proposed^{6,9}. For example, e2e-grooming aims to reuse lightpaths for other demands having common source (s) and destination (t) nodes, always targeting a n_{fs} minimization when allocating them. This procedure can be called as *predefined* e2e-grooming and causes that some lightpaths cannot be reused because the free spatial-spectral resources are not able to accommodate other demands⁶. To address this issue, a possibility can be to test other unexplored SCh configurations during the e2e-grooming (i.e., other tuples in Ω including Spe- or S2-SChs). If one lightpath has free spatial-spectral resources to groom another demand, we could start testing from the least demanding SCh configuration (in terms of n_s) to the most demanding one. This procedure can be called as *dynamic* e2e-grooming. For example, Fig. 1 shows a Spa-

SCh allocation for the 400 Gb/s demand ($d1$), previously introduced. After allocating $d1$ there would be 3×2 free spatial-spectral resources that might be used to allocate a co-routed demand, e.g., $d2$ of 100 Gb/s with DP-QPSK. The SCh candidates (Ω) of $d2$ are: (1,3), (2,2) and (5,1). The *predefined* e2e-grooming would select the (5,1) tuple, which does not suitably fit into the free spatial-spectral resources. Nevertheless, following our *dynamic* e2e-grooming proposal, the first candidate tuple (1,3) is discarded because it requires higher spectral resources, whereas the second one (2,2) fits properly and even 1×2 spatial-spectral resources are still available to groom other demand in s_9 .

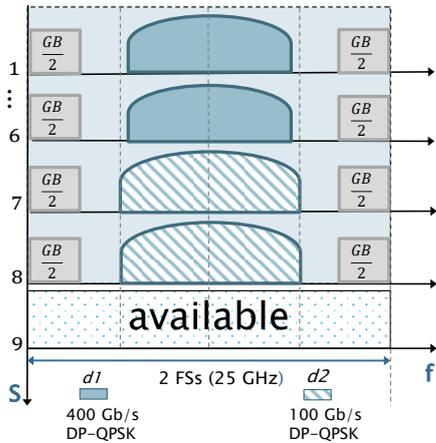


Fig. 1: *Dynamic* e2e-grooming for two co-routed demands

RMSSA heuristic

In this section, we elaborate on the Route, Modulation Format, Space and Spectrum Assignment (RMSSA) greedy heuristic for JoS-enabled Flex-Grid/SDM networks with *dynamic* e2e-grooming. In general, when a new demand d arrives at the network, it should be firstly checked if the e2e-grooming operation is feasible. To this end, all lightpaths sharing the same s and t nodes as d are found. Then, for each of these lightpaths the most spectrally efficient (MSE) modulation format is selected according to its physical length, in order to compute Ω from $n_s=1$ to $n_s=S$. i -th tuples (n_s, n_{f_s}) that do not yield lower n_{f_s} than their ($i-1$)-th are trivially discarded at the very beginning. The available spatial-spectral resources per lightpath are contrasted with those required by each SCh candidate $\in \Omega$. If one SCh candidate suitably fits into the available resources, d is considered as served; otherwise, a new lightpath has to be established. For this purpose, the $k=3$ shortest paths are used as candidate physical paths. Starting with the shortest one, the following procedure is

executed: 1) Select the MSE modulation format with an optical reach equal or higher than the physical path length; 2) Compute the (n_s, n_{f_s}) tuple by using the PCA strategy, i.e., the SCh configuration minimizing n_{f_s} ; 3) Find if continuous and contiguous n_{f_s} FSs are available in the spectrum of the first spatial channel along the candidate path. If so, directly reserve them across n_s spatial channels and consider d as served. Otherwise, subsequent candidate paths are analysed. If no candidate path can allocate d , it is finally considered as blocked.

Simulation Results and Discussion

To evaluate the proposed *dynamic* e2e-grooming, we consider the DT12 National backbone network² with 12 nodes and 40 unidirectional fiber links (average link length 243 km). Each fiber link consists of 22 spatial channels (bundle of standard single-mode fibers –SSMF–), each one with 320 FSs of 12.5 GHz (4 THz C-band). The optical reach estimation per Spa-SCh is computed by following the procedure detailed in ref. ⁷ for SSMF with up to 32 GBaud DP-BPSK, DP-QPSK, DP-16QAM, DP-64QAM signals.

A dynamic scenario where demands arrive at the network following a negative exponentially distributed inter-arrival time (IAT) is assumed. Each request asks for a bidirectional lightpath between uniformly distributed s and t nodes with bit-rate r_d during a certain holding time (HT), also following a negative exponential distribution. The r_d value follows several traffic profiles (TPs) with different mean (μ) and standard deviation (σ) bit-rate. Different offered loads (L) are obtained by fixing the IAT and varying the HT accordingly ($L=HT/IAT$). To get statistically relevant results, we offer 5×10^5 bidirectional requests per execution.

Fig. 2(a-f) show the Bandwidth Blocking Probability (BBP) vs. L (in Pb/s) for 6 TPs (μ, σ) from TP1 to TP6, respectively. Three curves are shown per graph disclosing results without e2e-grooming (No-Grooming, plots in purple), with *predefined* e2e-grooming (Predef-Grooming, plots in blue) and with *dynamic* e2e-grooming (Dyna-Grooming, plots in red). In the first set of simulations –Fig. 2(a-c) on top– the μ value per connection is the same, while the σ one is increased, namely, TP1 (2.0,0.0), TP2 (2.0,0.8) and TP3 (2.0,1.6), where μ and σ values are in Tb/s. Conversely, in the second set of simulations –Fig. 2(d-f) at the bottom– the σ value is the same, while the μ one is increased, namely, TP4 (0.5,0.4), TP5 (1.0,0.4) and TP6 (2.0,0.4).

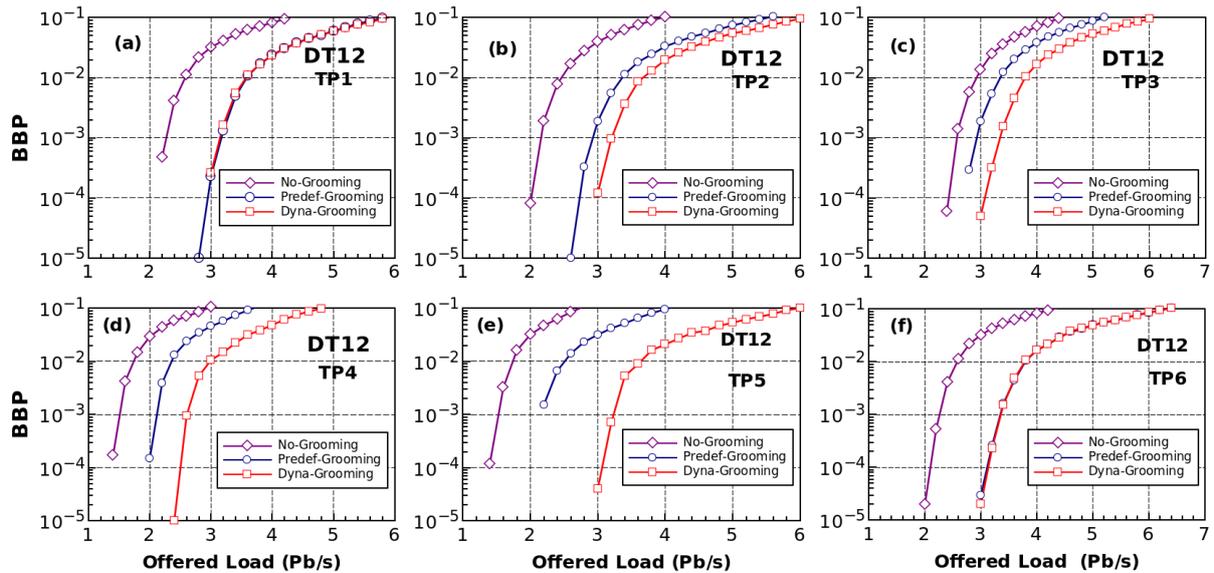


Fig. 2: BBP vs. Offered Load (in Pb/s) for different TPs, (top): fixed μ and increasing σ , (bottom): fixed σ and increasing μ

Looking at the results, for a fixed $\mu=2$, the higher the σ , the higher the benefits (in terms of L) of *dynamic* versus *predefined* e2e-grooming. Specifically, for a homogeneous TP1 (i.e., $\sigma=0$) the benefits of *dynamic* e2e-grooming are inappreciable. However, for heterogeneous TPs (i.e., $\sigma>0$) these benefits are up to 6% and 10% for $\sigma=0.8$ –TP2– and $\sigma=1.6$ –TP3–, respectively, for a 1% target BBP. Meanwhile, for a 10% BBP these benefits raise up to 7% and 15% for both cases of $\sigma=0.8$ and $\sigma=1.6$, respectively.

Regarding the second set of simulations, fixing $\sigma=0.4$ and varying μ , *dynamic* e2e-grooming benefits can raise up to 25% and 44% for TP4 and TP5, respectively, for 1% target BBP. Moreover, Fig. 2(d-f) show the impact of the size of demands (in terms of bit-rate). As observed, the higher the bit-rate of demands, the higher the supported L (therefore, the network throughput) under the three strategies (no, *predefined* and *dynamic* e2e-grooming). For instance, for a 1% target BBP, L can increase up to 27%, 58% and 47% for no grooming, *predefined* and *dynamic* e2e-grooming, respectively, when μ value per TP ranges from 0.5 Tb/s –Fig. 2(d)– to 2 Tb/s –Fig. 2(f)–. Meanwhile, for a 10% BBP the same percentages raise up to 40%, 86% and 35%, respectively. Among the three strategies, the results show that *predefined* e2e-grooming is the most sensible to the size of demands.

Taking into account the joint impact of μ and σ values, the benefits of e2e-grooming are more significant for low/moderate bit-rate and heterogeneous demands, as observed for TP4 and TP5 –Fig. 2(d,e)–, whereas these benefits are reduced for high bit-rate and very

heterogeneous demands, such as the case of TP3 –Fig. 2(c)–. Specifically, for a 1% target BBP, for TP4 and TP5, the benefits of e2e-grooming are up to 50% for the *predefined* strategy, against an up to 120% for the *dynamic* one. Meanwhile, for TP3, the benefits are reduced by 15% and 27%, for *predefined* and *dynamic* e2e-grooming, respectively. Finally, according to the results, *dynamic* e2e-grooming does not yield benefits for high bit-rates and homogeneous or quasi-homogeneous TPs, as observed for TP1 and TP6.

Conclusions

Dynamic e2e-grooming strategy in JoS-enabled Flex-Grid/SDM networks yields high benefits compared to *predefined* e2e-grooming especially under moderate traffic profiles in terms of bit-rate and heterogeneity. Particularly, a performance enhancement of up to 44% has been showed. Future works can address complexity analysis related to control plane and transponders for *dynamic* e2e-grooming.

Acknowledgements

This work has been funded by the Spanish National Project (TEC2017-90034-C2-1-R and TEC2017-90034-C2-2-R) under FEDER contribution and by SENESCYT – Ecuador.

References

- [1] O. Gerstel et al., IEEE Commun. Mag. 50, S12, 2012
- [2] P. Winzer, Nat. Photonics 8, 345, 2014.
- [3] D. M. Marom et al., IEEE Commun. Mag. 53, 60, 2015
- [4] J. M. Rivas Moscoso et al., in ECOC 2016, M.1.F.4
- [5] D. Klionidis et al., IEEE Commun. Mag. 53, 69, 2015
- [6] R. Rumipamba-Zambrano et al., in OFC 2017, W4F.8
- [7] R. Rumipamba-Zambrano et al., Opt. Switch. Net. 27, 93, 2018
- [8] B. Shariati et al., in ECOC 2016, M.1.F.1
- [9] F. Pederzoli et al., J. Opt. Comm. Netw. 9, B1, 2017