Defragmentation-Based Capacity Enhancement for Fixed to Flexible-Grid Migration Scenarios in DWDM Networks

António Eira[(1,2)], João Pedro[(1,2)], João Pires[(2)], Daniel Fonseca[(1,2)]
Juan Fernández-Palacios[(3)], Víctor López[(3)], Stefan Spaelter[(4)]

[(1)] Nokia Siemens Networks S.A., R. Irmãos Siemens 1, 2720-093 Amadora, Portugal, (e-mail: antonio.eira.ext@nsn.com)
[(2)] Instituto de Telecomunicações, Instituto Superior Técnico, 1049-001 Lisbon, Portugal
[(3)] Telefónica I+D SAU, Don Ramón de la Cruz 84, Madrid, 28006, Spain
[(4)] Nokia Siemens Networks Optical, St.-Martin Strasse 13, 81541 Munich, Germany

Abstract *This paper investigates the most favorable conditions for a full defragmentation procedure when evolving from a fixed to a flexible-grid scenario. Results show that migrating legacy channels to flexible-grid formats can yield up to 50% more capacity for future traffic, provided these legacy channels are also spectrally re-planned.*

Introduction

The flexible-grid concept has emerged in optical networking as a means of extending the capacity of metro/core networks in the short to mid-term. This is accomplished not only by a control-plane support for channel bit-rates above 100 Gb/s (e.g. 400 Gb/s) that extend beyond a 50 GHz channel width, but also by enabling a trade-off between spectral efficiency and optical performance with a granularity defined at the spectral slot level (e.g. 12.5 GHz)[1].

This slot-level spectrum management is essential to fully exploiting the available resources towards increased network capacity. However, practical constraints may imply that the deployment of flexible-grid enabled hardware and management systems is not sufficient in itself to fully take advantage of the network resources. Specifically, one has to consider that in many cases new services employing flexible-grid functionalities will need to co-exist with legacy services still based on the 50 GHz grid. It has been shown that a network supporting channels with heterogeneous spectral widths induces more spectrum fragmentation than one with a unique channel width[2]. In a fixed-grid environment the concern about fragmentation relates to not being able to find a single wavelength channel available on all links traversed by an optical path. Diverse channel widths can aggravate the problem of stranded bandwidth since certain allocation patterns create free spectral windows too small to be used by some or all channel types.

The aim of this paper is to evaluate the impact on network capacity of deploying a flexible-grid solution over a network partially loaded with fixed-grid channels. Concretely, we study the merits of re-planning the network in such a way that new services based on a flexible-grid can efficiently exploit the available spectrum. Moreover, we discuss possible migration strategies from fixed to flexible-grid scenarios and also evaluate how the spectrum planning affects the usable capacity.

Framework

Our testbed scenario is the Spanish Backbone reference network in Fig. 1, comprising 14 nodes for traffic add/drop and 16 nodes for switching only. The add/drop nodes are organized in seven node-pairs enforcing the dual-homing constraint imposed by the lower layer. Aggregated IP traffic is sent to a node-pair and routed through the WDM layer to another node-pair through two link-disjoint paths.

In this study, we assume a network life cycle with two distinct periods. The “legacy” state includes 40 and 100 Gb/s channels employing CP-QPSK modulation in a 50 GHz grid. On a second phase, we assume the availability of flexible-grid technology through WSSs with switching granularities below 50 GHz and a flexible-grid enabled GMPLS control plane. This setup supports both (i) the transport of 40 and 100 Gb/s channels in 25 and 37.5 GHz spacings respectively[3], whenever optical performance is not compromised by the increase in spectral efficiency, and (ii) the support of 400 Gb/s channels in 75 or 100 GHz spacing.

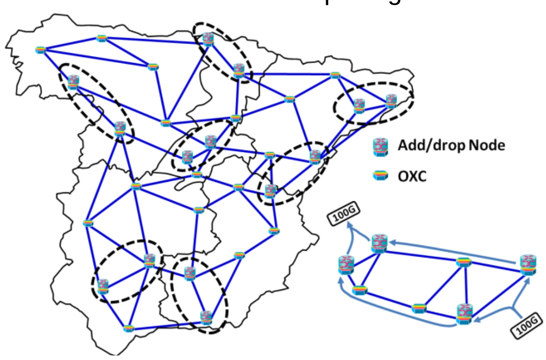

Fig. 1: Spanish Backbone reference network with dual-homing constraint

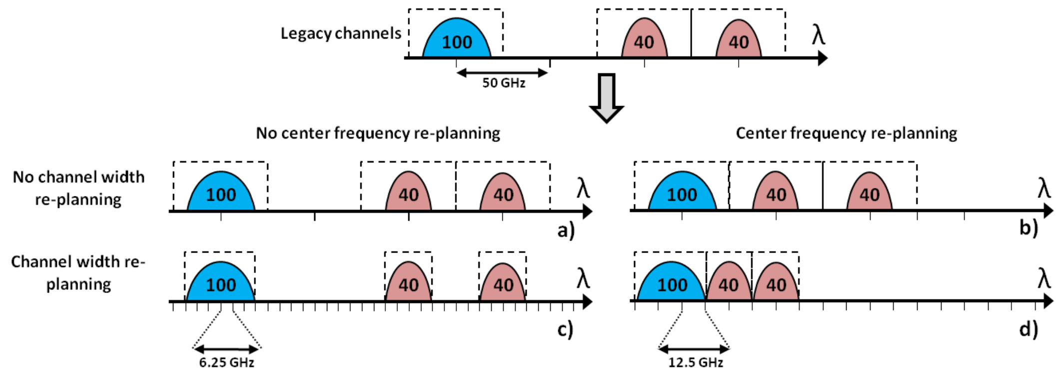

Fig. 2: Migration scenarios for re-planning legacy channels

While demands in the second stage are served with the least possible spectrum, this setup leaves open the question of how to manage the legacy connections as the flexible-grid is introduced in the network. The co-existence of multiple channel widths in the same network may magnify the problem of spectral fragmentation. This is particularly relevant if the co-existing formats have “mismatched” widths, that is, if their joint spectral allocation produces windows of free spectrum that are unusable by all or some of the available formats. Because of this, it has been suggested that format widths that are multiples or sub-multiples of each other should be spectrally grouped to reduce fragmentation[4,5]. Our goal is to identify the migration scenarios where the legacy channel allocation most hinders the efficiency of new channels using the flexible-grid, and whether a one-time defragmentation of these legacy channels can help flexible-grid formats utilize the network capacity to its fullest.

Given a network state before the introduction of flexible-grid channels, we specify four migration scenarios, depicted in Fig. 2, for handling the legacy channels. In the baseline case (Fig. 2a) the legacy channels’ center frequency and width is maintained. If a re-planning procedure is triggered on the legacy channels, two scenarios are possible: the center frequency of the channels can be re-tuned to a new position still in the 50 GHz grid (Fig. 2b), or the center frequency is moved to a new position while the channel width is simultaneously reduced to 25 or 37.5 GHz (Fig. 2d). This action is applied to all channels where the spectrum squeeze does not compromise the channel’s performance, requiring additional 3Rs. However, re-planning with the flexible-grid formats can produce biased results since there is an implicit capacity gain from the spectrum squeezing, and not just due to re-planning the demands. For a fairer comparison, the scenario in Fig. 2c considers that the demands are not spectrally re-planned, but their channel widths are squeezed to 25/37.5 GHz when possible. For 40 Gb/s channels, one 12.5 GHz slot can be removed on each side of the optical channel. For 100 Gb/s channels, only 6.25 GHz can be removed on either side of the channel, which implies the use of a 6.25 GHz grid granularity. This scenario is the most troublesome in terms of performance, as the legacy channels’ lasers may not have the stability required to operate outside the 50 GHz spacing. However, it provides a useful benchmark to measure how much of the capacity gains are due to the re-planning and how much is due to using more efficient formats.

Planning Workflow

The setup for evaluating the re-planning efficiency (described more thoroughly in [3]) is the following: firstly, the network is loaded with traffic demands using either 40 or 100 Gb/s channels in the 50 GHz grid. Demands are sequentially added to the network up to a desired capacity value. In order to emulate the spectrum churn and fragmentation arising from a realistic network operation, a fraction of the total demands are then torn down. At this point, the workflow splits into the four previously described scenarios, depending on if, and how, the legacy channels are re-planned. In the second stage, all four scenarios receive the same set of new demands to be deployed using the flexible-grid. The re-planning efficiency is measured by the amount of load that the network supported in the second stage relative to the case where no re-planning took place.

Simulation Results and Discussion

For simulation purposes, the assumed reach values are 2500 km for 40 Gb/s channels and 2000 km for 100 Gb/s channels. Additionally, 100 Gb/s channels at 37.5 GHz incur in a 80 km penalty per traversed node due to mismatched filter cascades[6]. For 400 Gb/s, the reach is 600 and 1200 km for 75 and 100 GHz widths respectively. In order to normalize the offered load, each demand is weighted by its payload and the amount of links crossed in its shortest path. Routing is performed using the shortest available path, and first-fit spectrum assignment is used. The format for each demand is decided based on which interface(s) yield the lowest cost to carry the demand. We considered relative costs for 40, 100 and 400 Gb/s interfaces of 1, 2 and 5 respectively. The demand size is uniformly distributed in multiples of 10G, up to 100G in the first planning stage, and up to 400G in the second stage.

For each migration case in Fig. 2, we evaluate three separate scenarios which differ in the array of formats available for deployment in the second stage: only 100G channels available after the re-plan (37.5 GHz), only 400G channels (75/100 GHz), or with both bit-rates available. In Fig. 3, we plot the results averaged over 50 runs for increasing values of the installed legacy load.

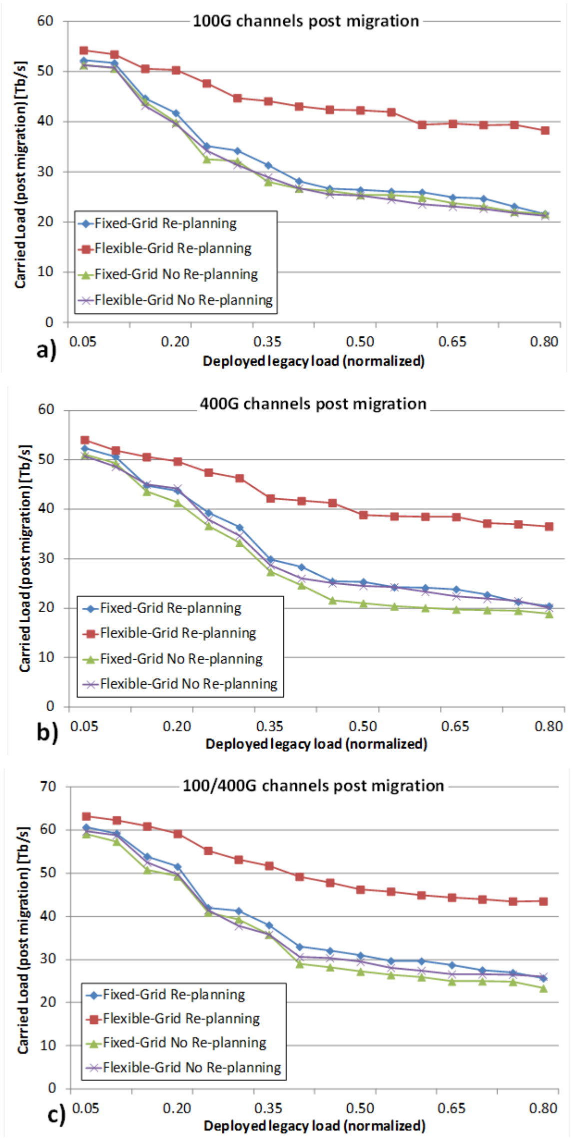

Fig. 3: Network capacity for: a) only 100G channels; b) only 400G channels; c) 100 and 400G channels

As one would expect, the flexible-grid re-planning efficiency increases with the value of the legacy load, as there is more fragmented spectrum to claim. The average capacity gain of defragmenting while still using the fixed-grid is 4% for 100G only, 12% for 400G only and 9% for both bit-rates. When the legacy channels are migrated to flexible-grid formats, these values rise to 52%, 47% and 46% respectively, even though the non re-planned case also benefited from more spectrally efficient formats. Although the spectrum squeeze on legacy channels saves bandwidth, it only provides a moderate capacity increase relative to the fixed-grid case, implying that most of the freed bandwidth cannot be reused without a spectral re-planning. An intuitive example of this is the squeezing of adjacent channels that creates, at best, a free 25 GHz window in between which cannot be exploited without re-planning.

The fixed-grid re-planning gains are more dependent on the formats that co-exist with the 50 GHz channels. When only 100G with 37.5 GHz are mixed with fixed-grid channels, they will fit into the 50 GHz windows left by fragmented fixed-grid spectrum. Re-planning can avoid the 12.5 GHz blocks stranded in this case, but overall there is a lower impact on total capacity. For 400G formats, 50 GHz windows cannot be reused at all, hence the re-planning gain is higher. For 100/400G, there is a mix of these effects resulting in an intermediate gain value.

Conclusion

This paper presented a framework to identify the scenarios that most benefit from a re-planning of legacy fixed-grid channels when migrating to a flexible-grid setting. We showed that if the legacy channels are kept on the 50 GHz grid, the gain of re-planning strongly depends on the spectral widths of the co-existing formats. On the other hand, if the legacy channels are also migrated to the flexible-grid, a re-planning operation is essential to reuse the spectral dividend and maximize network capacity.

Acknowledgements

This work was supported by the FP-7 IDEALIST project under grant agreement number 317999.

References

[1] Amaya et al., OFC'11, OMW3 (2011)
[2] Takagi et al., ECOC'11, Mo.2.K.3 (2011)
[3] Eira et al., "Defragmentation of fixed/flexible grid optical networks", Future Network & Mobile Summit (accepted for publication) (2013)
[4] Wang et al., ECOC'11, Mo.2.K.6 (2011)
[5] Tsardinakis et al., OFC'13, JW2A.65 (2013)
[6] Eira et al., OFC'13, OTu2A.5 (2013)