On the Benefits of Elastic Spectrum Management in Multi-Hour Filterless Metro Networks

Jose-Juan Pedreno-Manresa*, Jose-Luis Izquierdo-Zaragoza†, Filippo Cugini‡ and Pablo Pavon-Marino*
*Universidad Politécnica de Cartagena, Plaza del Hospital, 1, 30202 Cartagena, Spain
†Universitat Politècnica de Catalunya, Jordi Girona, 31, 08034 Barcelona, Spain
‡Consorzio Nazionale Interuniversitario per le Telecomunicazioni, Via Moruzzi 1, 56124 Pisa, Italy
E-mail: josej.pedreno@upct.es

***Abstract*—The dawn of 5G is pushing operators to deploy high-capacity, agile networks capable of adapting to time-varying traffic patterns, especially into metro sections. ROADMs are key enablers for agility in the optical layer, however the benefits of this agility do not always compensate for increased costs. As such, filterless optical networks are emerging as a cost-effective and reliable solution compared to active photonics, thanks to a winning combination of coherent transponders and passive splitters/couplers. However, spectrum allocation management policies are of paramount importance to maximize the overall network throughput. In this paper, we focus on a filterless metro network where the hourly variation of the demands traffic is known, coming from historic data estimations. Then, we observe how the knowledge of the traffic profiles can be exploited. To assess this, we evaluate the performance, in terms of throughput, of three different spectrum management approaches: (i) fixed, where lightpaths remain static along time once allocated; (ii) semi-elastic, where lightpath-bandwidth vary according to current traffic requirements, but central frequency remains fixed; and (iii) hitless full-elastic, where any lightpath parameter may be reconfigured without disrupting the traffic. Besides, we consider two transponder types equipped with (i) shared or (ii) independent tunable lasers for transmission and reception, which affects to spectrum allocation of bidirectional connections. According to our results, the semi-elastic approach clearly outperforms the fixed approach (23-33% more throughput) with a reduced gap to the hitless full-elastic case (10-24% less throughput), especially considering that the latter is not commercially available yet. Interestingly, using dual-laser transponders only yields a 10% gain with respect to single-laser transponders for the semi-elastic scenario, and thus may not justify the extra hardware.**

***Index Terms*—Filterless optical networks, time-varying traffic, elastic spectrum allocation.**

I. Introduction

Filterless optical networks (FONs) have been introduced to target a significant cost reduction compared to existing optical networks. Optical nodes based on reconfigurable optical add/drop multiplexers (ROADMs) and wavelength selective switches (WSS) provide agility in optical transport but incurring in additional capital expenditures (CAPEX), which does not always compensate the benefits of agility. Leveraging coherent transponders (TXPs) and passive splitters/couplers, filterless nodes are deemed to be a cost-effective and reliable alternative to active optical nodes [1]. In FONs several optical trees are built over the same physical topology. In each of these trees, optical connections are based on *light-trees* and optical signals reach all nodes. Thanks to the coherent technology, each signal can be properly selected at the receiver (RX) of the corresponding destination node, following a drop and waste (D&W) strategy.

Despite of cost savings and simple control and maintenance operations, FONs present some drawbacks: (i) there is a significant waste of spectrum resources since each signal occupies the entire associated fiber tree; (ii) simple topologies such as horseshoes or trees with no physical loops must be used to avoid power recirculation and lasering effects of amplified spontaneous emission (ASE) noise; and (iii) careful control of the overall power entering each TXP is required, as all light channels enter each of them, provided that power levels must be above sensitivity.

As of today, filterless solutions have not been considered for large scale deployments [2]. The main motivation is that coherent technologies have been mainly adopted in backbone networks, where large ROADM-based mesh topologies are employed and the aforementioned filterless drawbacks have practically prevented the deployment of such D&W solutions.

However, the traffic growth occurring in metro networks is driving the replacement of traditional direct-detection 10G cards with more advanced solutions at 100G. In this context, where network operators usually adopt simple topologies (i.e., horseshoes, rings…) over relatively short distances (up to 150 km) such transmission constraints become less critical, and coherent TXP over filterless transport may represent a suitable and cost-effective option. Nonetheless, the potential application of these solutions in the metro opens the way to additional and yet undiscussed design and technological aspects.

In this paper, we focus on filterless solutions in the context of metro networks, specifically addressing the following two aspects: (i) the impact of traffic dynamicity, which is significantly more intense than the one experienced in backbone networks, in the design of spectrum assignment (SA) solutions for filterless networks; and (ii) the availability of either a single or two tunable lasers within the coherent TXP, which may limit flexibility on bidirectionality.

The contribution of this work is two-fold. First, an SA algorithm to (re-)allocate demands adapting to the new traffic conditions is presented, considering three different variants, namely *fixed*, *semi-elastic* and *hitless full-elastic* [3], depending on the tunable parameters on the transponder, that is,

978-3-903176-07-2 © 2018 IFIP

central frequency (CF) and bandwidth (BW). Second, we evaluate whether the investment in TXP with dual-laser is beneficial compared to single-laser variants. As a benchmark, we use a horseshoe metro FON subject to a *multi-hour traffic profile*, being throughput the metric under consideration.

The rest of the paper is organized as follows. In Section 2, we review some previous works in filterless networks. In Section 3, we provide with a background on the role of TXPs in filterless metro networks. In Section 4, we describe our SA algorithm. In Section 5, we report and discuss the results of our case study. Finally, Section 6 concludes the paper.

II. Related Work

The concept of FONs was first introduced in [4]. Aside of considerations from the photonics (i.e., power control, coherent detection) [5] or control plane [6] perspectives, in this section we review the efforts in the last decade in topics related to network planning and resource allocation.

The first problem arising in FONs is connectivity. In fact, these networks present several differences in terms of planning and operation with respect to ROADM-based networks. Actually, the routing and spectrum assignment (RSA) problem is augmented to a topology, routing and spectrum assignment (TRSA), where several physical sub-topology (or fiber trees, onwards) are built on top of a shared fiber topology (topology subproblem) and demands are alternatively assigned to some of them (routing subproblem). In addition, the broadcast nature within a tree means that a more careful SA is needed [7], aiming to take advantage of time-dependent spectrum sharing to optimize its utilization.

To summarize, as a result of the TR subproblems, we must ensure connectivity between all node pairs while avoiding laser loop and fulfilling reachability constraints. Finally, the SA provides the spectrum allocation across the corresponding fiber tree.

Authors in [8] propose a TRSA algorithm (with static traffic and focus on fixed-grid) to perform a techno-economic analysis comparing ROADM-based networks and FONs. Their results demonstrate cost savings up to two orders of magnitude, but lack of a throughput analysis. In [7], they extend their analysis to a dynamic scenario, where lightpaths are setup upon request over a previously deployed filterless network. Here, they provide a comparison of ROADM-based, static filterless and dynamic filterless in terms of wavelength usage over a given demand set but, still, there is nothing about a throughput analysis varying the overall network load. By contrast, authors in [9] present a multi-goal optimization solution to this problem in a pilot network, and they observed that up to 2.5 times more traffic can be supported by their reference network considering active photonics instead of filterless solutions. Finally, authors in [10] provide a comparison in terms of spectrum utilization similar to [7], but focused on flex-grid. This latter work also considers a multi-period scenario, where traffic grows year-over-year.

To the best of knowledge, our work is the first considering the potential implications the SA problem in the context of the particular constraints imposed by FONs for multi-hour traffic, where rate and spectrum occupation vary along time. The benefits of having such information about variations in traffic patterns, and its exploitation for improved network efficiency, have been studied in the past for fixed-grid [11] and flex-grid [3]. Here, we approach this problem in order to analyze the implications of different SA schemes and TXP technologies in terms of throughput and identify potential techno-economic aspects.

III. Filterless Technologies in Dynamic Metro Networks

Next generation metro networks are expected to be driven by technological solutions where all cost contributions (e.g., hardware components within the coherent TXP) need to be carefully considered. Moreover, compared to backbone networks, high traffic dynamicity will be experienced, leading to the potential adoption of cost-effective highly reconfigurable solutions. These two aspects are discussed in the following subsections.

A. Transponder technologies

Coherent TXPs are built in hardware as *transceivers*, that handle within the same card both directions, i.e. transmitter (TX) and RX. According to the adopted technology and equipped hardware capabilities, transmission parameters may or not be configured in the same way both directions. For example, a connection between nodes A and B may or not be capable of supporting different CF from A to B (i.e., TX at A) and from B to A (i.e., RX at A). Such capability depends on the availability of either one or two tunable lasers and related electronic circuits within the card. If just one tunable laser is present, both TX and RX have to be operated over the same CF, i.e., the TX laser is also used as local oscillator at the RX side. Instead, if two tunable lasers are available in the same card, TX and RX can be operated over independent frequencies.

A single-laser provides cost savings, but introduces assignment constraints in spectrum allocations. Such constraints may be particularly relevant in next-generation metro networks, where significant traffic asymmetry is expected. For example, huge amount of traffic flows from a content delivery network (CDN) network to the access (i.e., downstream traffic to end-users) while the reverse direction (i.e., upstream traffic from users) is significantly less utilized.

Fig. 1 shows two asymmetrical bidirectional connections. From nodes A to B and C to D a large amount of spectrum is occupied to guarantee the required bandwidth capacity. Fig. 1a shows the case where a single-laser is shared between the TX and local oscillator at RX. In Fig. 1b two independent laser sources are utilized, enabling a free allocation of each central frequency. As it can be noticed, in the first scenario, the constraint on the same CF for both TX and RX leads to high fragmentation and, in turn, to relevant wasting of spectrum resources, which is avoided when each card is equipped with two independent laser sources.

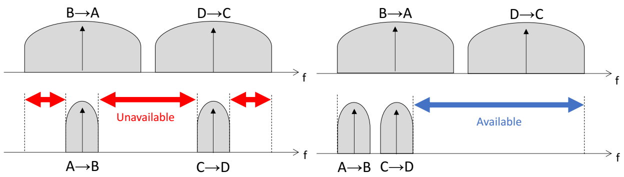

(a) Single-laser TXPs (b) Dual-laser TXPs

Figure 1. Resource allocation with different TXPs technologies.

B. Traffic dynamicity

Regardless the bidirectional tunability issue described in the previous subsection, of special interest in this work is taking advantage of bandwidth-variable transponders (BV-Ts) to follow traffic variations adapting rate and spectrum occupation.

Fig. 2 shows two neighboring connections from nodes A to B and C to D (for simplicity, here just one direction is shown). Both connections experience time dependent bandwidth utilization. For example, the first connection serves a business district where most of the bandwidth is requested in working hours (time t=0) with scarce use, e.g., in the evening (t=1). Conversely, the second connection serves a residential area where limited bandwidth is needed during working hours but larger service requests are experienced in the evening. To this respect, different SA policies can be applied potentially considering time-dependent spectrum sharing.

In Fig. 2a, the resource allocation is performed in a fixed manner, considering peak-hour traffic volumes. In Fig. 2b, the allocation is performed accounting for multi-hour variations and not on the peak-hour values. This way, spectrum resources used by connection A to B at t=0 can be reused by connection C to D at t=1. In this case, fixed CF are assumed. Compared to Fig. 2b, in Fig. 2c we consider re-tuning of CF as an additional degree of flexibility. Such re-tuning can be implemented by adopting the *push-pull defragmentation* technique presented in [12], where the automatic frequency control of the coherent RX is properly exploited to track the gradual shift deliberately applied to the TX, without affecting high-layer traffic. Such adaptation is further simplified in FONs networks given the absence of filters.

It is worth mentioning that, as discussed in the subsequent sections, the combination of common/independent CF assignment in bidirectional connections (as discussed in Section III-A and such fixed/multi-hour assignment with/without central frequency adaptation may lead to remarkably different network utilization performances, with potential high impact on deployment costs and use of resources before fiber exhaustion, even for multi-period optimization where traffic only “grows”.

IV. Proposed Algorithm

As described in Section II, the TRSA problem in filterless networks can be decomposed into three subproblems: (i) generate the set of fiber trees (T), (ii) associate end-to-end demands to one of the trees (R), and (iii) allocate spectrum for each of the demands into the trees (SA).

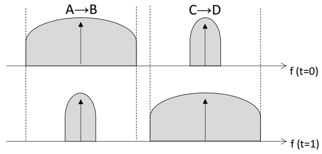

(a) Fixed reservation considering peak-hour traffic

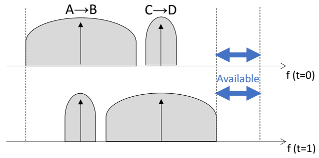

(b) Dynamic reservation accounting for multi-hour variations with fixed CF

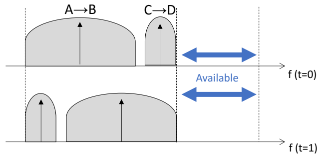

(c) Dynamic reservation accounting for multi-hour variations and CF re-tuning

Figure 2. Resource allocation at different time of the day/week.

In this paper, we focus on horseshoe topologies, containing one bidirectional fiber tree, targeting the SA problem for all the demands within a tree. Due to the nature of FONs, topology and routing subproblems are implicitly solved for horseshoes.

The aim of our algorithm is to guarantee the correct provisioning of all demands while minimizing spectrum utilization. The input data of the algorithm are the following:

- FON topology containing a single fiber tree ($G(N, E)$), composed of a set of nodes with BV-Ts and bidirectional fibers interconnecting them, each supporting up to S slices at a line rate of R Gbps/slice.
- A list containing 24 different traffic matrices ($M^t_{(i,j)}$ where (i, j) is the source-destination node pair and $t \in [0, 23]$ the time period), containing per-hour traffic volume.

Our multi-hour spectrum allocation (MH-SA) initially considers a *semi-elastic* approach [3]: each lightpath uses a fixed CF whereas the number of occupied slices is adapted according to the offered traffic on each period. Next subsection is devoted to describe the implementation of our MH-SA algorithm, whereas further subsections will describe variations of the algorithm according to the technologies and techniques explained in Section III.

A. Implementation

Because of the $\mathcal{NP}$-complexity of the SA problem, we resort to a heuristic algorithm based on biased random-key genetic algorithm (BRKGA) [13], a well-known variant of the genetic algorithm (GA) meta-heuristic. The algorithm (pseudo-code shown in Fig. 3) consists on several phases that are executed sequentially within a loop until a feasible solution is found or the maximum number of iterations (given as an input parameter) is reached.

- *GeneratePopulation*. Each candidate solution (also called *chromosome*) is encoded as an array, where each index (key) represents the demand identifier and its value the CF. To generate a chromosome, all demands are randomly sorted, and one-by-one we try to find a valid CF (using a first-fit approach) that ensures enough slices available for each hour of the day. To guarantee all demands can be initially serviced, we assume unlimited slice capacity (that is, an arbitrarily large *virtual spectrum*), whose rationale is further explained. The total number of chromosomes generated in this phase is given as an input parameter of the algorithm.
- *Crossover*. New chromosomes (*offspring*) are generated in this phase. Each new chromosome inherits its values from two 'parent' chromosomes: one of them is always selected among the *elite* population, whereas the other can be *elite* or not. Without loss of generality, in the first iteration, there is no *elite* population and both parents are selected at random. We iterate over each key (i.e., demand) selecting randomly one value (i.e., CF) among the two parents. The offspring chromosomes will then become part of the general population and the parents are discarded. The number of offspring generated in this phase is given as an input parameter.
- *CostComputation*. In order to test the feasibility of each candidate solution, we iterate over the 24 sets of demands. For each demand we select the appropriate CF from the chromosome and try to allocate the required number of slices. In case of overlapping, we consider the solution as infeasible. Once we have iterated over all sets of the demands, we select those slices that have never been used and we delete them from the virtual spectrum (shifting the remaining ones to the left). If the new shrinked spectrum size is lower or equal than the *real spectrum* size (input parameter), we consider the solution feasible and the algorithm ends. Each unfeasible solution is assigned a cost, equal to the percentage of blocked traffic.
- *Mutation* To avoid stagnation, in each iteration a portion of the general population is renovated, introducing newly created chromosomes (as explained in the first phase).

B. Transponder technologies

To reflect the effect of using different coherent TXP equipped with single or dual-lasers, we added an input parameter to the algorithm to control the bidirectional policy of the CF: (i) *different CF* with no restriction applied, and (ii) *same CF* which forces lightpaths to share the same CF as its bidirectional counterpart.

The former is implemented executing two independent instances of the MH-SA algorithm, one per direction of the horseshoe. In case a solution can be found for both, the scenario is feasible. Conversely, the latter is implemented in a tricky way: we concatenate traffic demands from each (i,j) in both directions in a virtual day of 48 hours. As such, we only execute an instance of the MH-SA algorithm.

Algorithm 1 Multi-Hour Spectrum Allocation

```
Require: G(V,E), M^t_(i,j), S, R
  GeneratePopulation
  iteration = 0
  while iteration ≤ maxIterations do
    Crossover
    CostComputation
    if feasible solution is found then
      end algorithm
    end if
    Select 'Elite' among the offspring
    Mutation
    iteration++
  end while
```

Figure 3. Pseudocode for the MH-SA algorithm.

C. Spectrum allocation approaches

To assess the quality of our algorithm, we compare the results with two boundary solutions in terms of achievable throughput. The lower bound is given by a *non-multihour, non-elastic* approach, where spectrum allocation is made for the worst case, that is, the daily maximum (peak) across all time intervals. The upper bound is given by a multi-hour aware *full-elastic with defragmentation* approach, where CF may also change, and therefore all connections are packed each time period using *push-pull defragmentation* techniques [12].

In order to model non-elastic/full-elastic, some minor modifications were added to our MH-SA algorithm. To implement the *non-elastic* option, we removed the dynamicity of the traffic information, considering only a single demand per node-pair where the offered traffic was equal to the maximum among the 24-hour time. For the *hitless full-elastic* approach, we consider all allocated demands can be packed, removing all spectrum between connections, for each single slot, but maintaining the same left-to-right order (in terms of slice indexes) of demands. Note that keeping such left-to-right order means that the variations in the CF can be accomplished using non-disruptive slow spectrum-shifting techniques.

V. Case Study

In this section, we report the results collected from testing our algorithm in a real-life scenario. We aim to analyze the total throughput achieved by combining different modulations and spectrum allocation approaches to find possible trade-offs.

The results were obtained using the offline network design tool Net2Plan [14]. With this tool, users can design and dimension networks assuming some static information (e.g. physical topology and traffic matrix). The algorithm was develop in Java, implementing public and well-documented interfaces. For the purpose of inspection and validation, both the source code of the algorithm and Net2Plan are available on the website [15].

A. Testing scenario

In order to test the algorithm we use the horseshoe topology in Fig. 4 as reference scenario. It is composed of 2 metro-core edge nodes (MCENs) providing connectivity toward core networks, 5 access-metro edge nodes (AMENs) as gateways for end-users (actual producers and consumers of traffic), and 3 nodes acting as CDN. For simplicity, we consider asymmetric, bidirectional traffic demands for node pairs MCEN-MCEN, MCEN-CDN, CDN-CDN, CDN-AMEN and AMEN-AMEN. MCEN-CDN and CDN-AMEN traffic is only considered for the closest node pairs. For example, there is only MCEN-1-CDN-1 traffic but not MCEN-2-CDN-1 traffic, as well as CDN-1-AMEN-1 traffic but not CDN-1-AMEN-5 traffic.

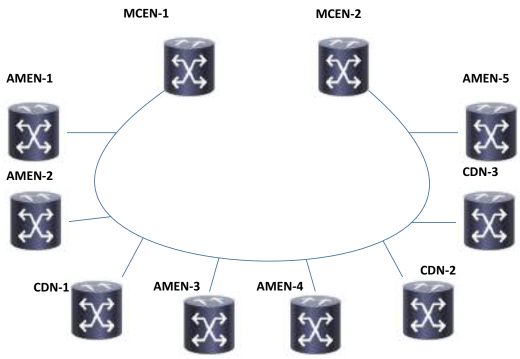

Figure 4. Reference horseshoe topology.

A combination of several different methods were used to generate a realistic multi-hour traffic matrix. First, a reference traffic matrix was built using a population-distance model described in [16], using scaling factors coming from analysis, trends and forecast for the incoming years [17]. Then, to recreate a multi-hour scheme, an activity factor was applied to the aforementioned traffic to simulate variance on each hour of the day [18] using a bimodal distribution to model peak and idle periods. Parameters of these distribution like width of the peak period and peak-to-idle ratio where selected at random. More details can be obtained directly from the source code.

B. Results

Different tests were performed using two different modulations: (i) 32 GBaud PM-QPSK over 37.5 GHz, with a total bitrate of 100 Gbps (R=16.6 Gbps per 6.25 GHz slice) and (ii) PM-16QAM, leading to 200 Gbps over 37.5 GHz (R=33.3 Gbps per 6.25 GHz slice). We establish a spectrum of 2.5 THz (S=400 slices). Using the scheme explained in Section V-A, we scaled the total offered traffic increasingly until blocking occurs. Then, we determine spectrum requirements SU for traffic (i, j) in time period t according to Eq. (1):

$$SU_{(i,j)}^{t} = \left\lceil M_{(i,j)}^{t}/R \right\rceil \tag{1}$$

We tested the algorithm MH-SA for the three SA approaches, single/dual-laser TXPs and the two modulation formats. Throughput results, indicating the peak-hour carried traffic before resource exhaustion, are presented in Fig. 5.

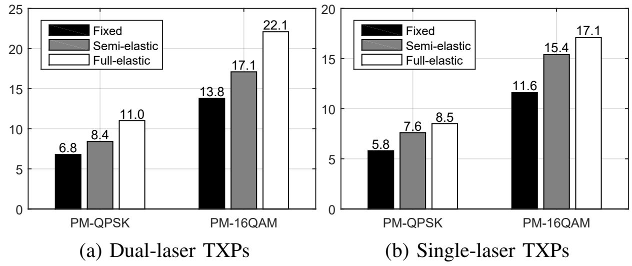

Figure 5. Maximum throughput (in Tbps).

We observed that the semi-elastic approach outperforms the fixed approach by a 23-33% in overall throughput. As expected, the algorithm performs poorly compared to the *full-elastic* scheme (with losses of 10-24% of throughput). These numbers are presented in Table I. Interestingly, even though it is clear that a full-elastic approach would be preferred in all instances, we would like remark that despite the fact that push-pull techniques [12] have been fully tested and demonstrated, unfortunately are not yet commercially available.

Table I
THROUGHPUT VARIATION (PERCENTAGE) WHEN USING SEMI-ELASTIC ALLOCATION COMPARED TO OTHER STRATEGIES

SA approach	dual-laser	PM-QPSK	PM-16QAM
Fixed	Yes	23.53%	23.91%
	No	31.03%	32.76%
Full-elastic	Yes	-23.64%	-22.62%
	No	-10.59%	-9.94%

From the results, it is worth discussing about the gain in terms of throughput using dual-laser TXPs instead of single-layer TXPs. Surprisingly, the use of the former only yields to an improvement of 10-11% of throughput and, as such, may not justify investment on these devices.

Table II
THROUGHPUT GAIN (PERCENTAGE) WHEN USING DUAL RESONATOR TRANSPONDERS INSTEAD OF SINGLE RESONATOR

SA approach	PM-QPSK	PM-16QAM
Fixed	17.2%	18.9%
Semi-elastic	10.5%	11%
Full-elastic	29.4%	29.2%

Before concluding, we illustrate the efficiency of our MH-SA algorithm in Fig. 6. This picture presents a partial snapshot of the spectrum utilization during a day in the case of using same CF for TX/RX. In Fig. 6a we can see the poor spectrum utilization of the fixed approach, since there is no change for spectrum sharing. Conversely, for the semi-elastic approach (see Fig. 6b) it is clear that because of bandwidth adaptation, slices can be shared by different demands in different time intervals, thus reducing the overall spectrum requirements. However, because of using same CF for both directions in this scenario, spectrum sharing cannot be totally exploited, thus reducing the overall throughput with respect to dual-laser TXPs. Finally, the full-elastic approach yields to the best spectrum utilization.

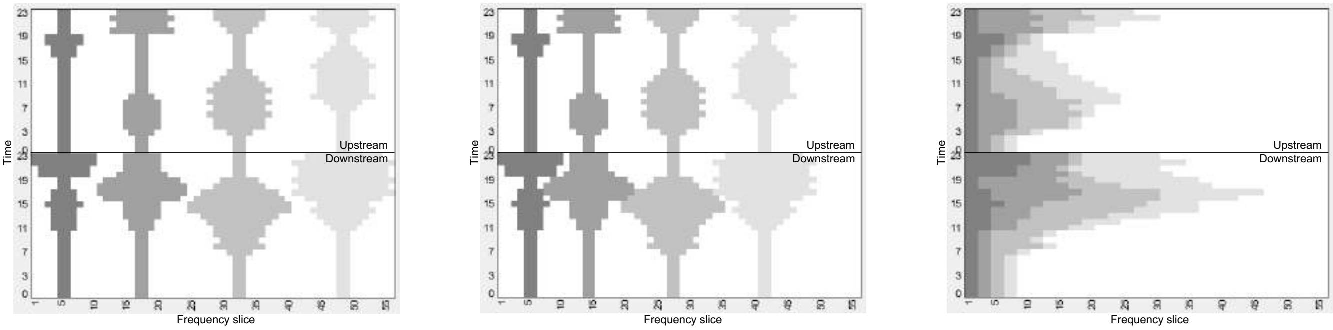

(a) Fixed SA (b) Semi-elastic SA (c) Full-elastic SA

Figure 6. Spectrum occupation for four example demands.

VI. Conclusions and Further Work

In this paper, we have presented an algorithm to determine the maximum throughput which can be achieved in filterless optical networks according to different spectrum management policies. In our results, a market-ready semi-elastic approach (considering state-of-the-art BV-Ts) represents a reasonable solution compared to a non-commercial hitless full-elastic scenario. Besides, we have illustrate the fact that single-layer TXPs may be a cost-effective alternative to dual-laser ones with limited throughput losses.

Extensions to this work may include mesh topologies, where fiber trees are also a decision variable, and/or techno-economic studies (e.g., independent tunability of TX/RX, comparison with ROADM-based solutions, and so on).

Acknowledgment

The research leading to these results has received funding from the EC through the METRO-HAUL project (G.A. nº 761727), from the Spanish MINECO TWINS (TEC2017-90097-R) and ONOFRE-2 (TEC2017-84423-C3-1-P), the FPU Spanish fellowship program (ref. FPU14/04227), and the "Programa de Formación del PDI a través de la Movilidad de la UPCT" (PMPDI-2017).

References

[1] C. Tremblay, É. Archambault, M. P. Bélanger, J.-P. Savoie, F. Gagnon, and D. V. Plant, "Passive filterless core networks based on advanced modulation and electrical compensation technologies," *Telecommunication Systems*, vol. 54, no. 2, pp. 167–181, Oct. 2013.

[2] C. Tremblay, P. Littlewood, M. P. Bélanger, L. Wosinska, and J. Chen, "Agile Filterless Optical Networking," in *Proceedings of the 21st Conference on Optical Network Design and Modeling (ONDM 2017)*, Budapest (Hungary), May 2017.

[3] M. Klinkowski, M. Ruiz, L. Velasco, D. Careglio, V. Lopez, and J. Comellas, "Elastic Spectrum Allocation for Time-Varying Traffic in FlexGrid Optical Networks," *IEEE Journal on Selected Areas in Communications*, vol. 31, no. 1, pp. 26–38, Jan. 2013.

[4] C. Tremblay, F. Gagnon, B. Châtelain, É. Bernier, and M. P. Bélanger, "Filterless Optical Networks: A Unique And Novel Passive WAN Network Solution," in *Proceedings of the 12th OptoElectronic and Communications Conference and 16th International Conference on Integrated Optics and Optical Fiber Communication (OECC/IOOC'07)*, Yokohama (Japan), Jul. 2007.

[5] F. Cugini, C. Porzi, N. Sambo, A. Bogoni, and P. Castoldi, "Receiver Architecture with Filter for Power-Efficient Drop&Waste Networks," in *Proceedings of the Optical Fiber Communications Conference and Exhibition 2016 (OFC 2016)*, Anaheim, CA (United States), Mar. 2016.

[6] G. Mantelet, C. Tremblay, D. V. Plant, P. Littlewood, and M. P. Bélanger, "PCE-Based Centralized Control Plane for Filterless Networks," *IEEE Communications Magazine*, vol. 51, no. 5, pp. 128–135, May 2013.

[7] G. Mantelet, A. Cassidy, C. Tremblay, D. V. Plant, P. Littlewood, and M. P. Bélanger, "Establishment of Dynamic Lightpaths in Filterless Optical Networks," *Journal of Optical Communications and Networking*, vol. 5, no. 9, pp. 1057–1065, Sep. 2013.

[8] É. Archambault, D. O'Brien, C. Tremblay, F. Gagnon, M. P. Bélanger, and É. Bernier, "Design and Simulation of Filterless Optical Networks: Problem Definition and Performance Evaluation," *Journal of Optical Communications and Networking*, vol. 2, no. 8, pp. 496–501, Aug. 2010.

[9] S. Krannig *et al.*, "How to design an optimized set of fibre-trees for filterless optical networks – The elegance of a multi-goal evolutionary Pareto optimization versus a deterministic approach," in *Proceedings of 17th ITG-Workshop on Photonic Networks*, Leipzig (Germany), May 2016.

[10] É. Archambault *et al.*, "Routing and Spectrum Assignment in Elastic Filterless Optical Networks," *IEEE/ACM Transactions on Networking*, vol. 24, no. 6, pp. 3578–3592, Dec. 2016.

[11] R. Aparicio-Pardo, N. Skorin-Kapov, P. Pavon-Marino, and B. Garcia-Manrubia, "(Non-)Reconfigurable Virtual Topology Design Under Multihour Traffic in Optical Networks," *IEEE/ACM Transactions on Networking*, vol. 20, no. 5, pp. 1567–1580, Oct. 2012.

[12] F. Cugini *et al.*, "Push-Pull Defragmentation Without Traffic Disruption in Flexible Grid Optical Networks," *Journal of Lightwave Technology*, vol. 31, no. 1, pp. 125–133, Jan. 2013.

[13] J. F. Gonçalves and M. G. C. Resende, "Biased random-key genetic algorithms for combinatorial optimization," *Journal of Heuristics*, vol. 17, no. 5, pp. 487–525, Oct. 2011.

[14] P. Pavon-Marino and J.-L. Izquierdo-Zaragoza, "Net2Plan: An Open Source Network Planning Tool for Bridging the Gap between Academia and Industry," *IEEE Network*, vol. 29, no. 5, pp. 90–96, Sep.-Oct. 2015.

[15] "Net2Plan - The open-source network planner," [Last accessed: January 2018]. [Online]. Available: http://www.net2plan.com

[16] R. S. Cahn, *Wide Area Network Design: Concepts and Tools for Optimization*, 1st ed., ser. The Morgan Kaufmann Series in Networking. Morgan Kaufmann, May 1998.

[17] "Cisco – The Zettabyte Era: Trends and Analysis," [Last accessed: January 2018]. [Online]. Available: https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/vni-hyperconnectivity-wp.htm

[18] J. Milbrandt, M. Menth, and S. Kopf, "Adaptive Bandwidth Allocation: Impact of Traffic Demand Models for Wide Area Networks," Institute of Computer Science, University of Würzburg, Würzburg (Germany), Research Report 363, Jun. 2005.