Distance-adaptive online RSA algorithms for heterogeneous flex-grid networks

Jose-Luis Izquierdo-Zaragoza, Pablo Pavon-Marino, Maria-Victoria Bueno-Delgado

***Abstract*—Flex-grid elastic optical networks are an enabling technology for future heterogeneous on-demand optical bandwidth services. Heterogeneous means that optical connection requests for different services requiring different data rates, would coexist in the network. In this context, Routing and Spectrum Assignment (RSA) algorithms face the challenge of allocating fairly incoming connections, that is providing a similar blocking performance to all the services.**

In this paper, we review existing RSA proposals applicable to heterogeneous flex-grid networks, observing (i) its blocking performance averaged among services, and (ii) the fairness in the blocking observed by each individual service. In addition, we propose the Partial-Sharing-Partitioning (PSP), a scheme to balance both aforementioned metrics. We concentrate on a distance-adaptive scenario, where the same connection request can be carried with different modulations, associated to different spectral efficiencies and optical reaches. Our simulation results in the Net2Plan tool explore the interplay between average blocking and fairness. We observe that many classical RSA algorithms produce unfair allocations, while PSP permits tuning the balance between both metrics. The algorithms developed are publicly available in the open-source Net2Plan repository.

***Index Terms*— Flexible optical networks, admission control, routing and spectrum assignment, heterogeneous services, spectrum fragmentation, fairness**

I. Introduction

ELASTIC optical networks, or flex-grid networks, are foreseen to provide dynamic on-demand high-bandwidth services for a variety of applications (e.g. video streaming and cloud computing) in the near future. The introduction of flexible and reconfigurable switching architectures [1][2] and new modulation formats like OFDM [3] will enable such a major breakthrough. This emergent heterogeneous optical bandwidth market poses new challenges under dynamic scenarios, since network operators must serve connection requests in an efficient and fair manner.

In flex-grid networks, lightpaths are allocated into a contiguous set of low-rate OFDM subcarriers that are continuously available on every link along the path. Unfortunately, these technological constraints (*spectrum contiguity* and *spectrum continuity*, respectively) lead to the *spectrum fragmentation* problem. During the network operation, connections of different bandwidth are set up and tore down, dividing the spectrum into small-sized blocks of available slots (*vertical fragmentation*), which may also be misaligned among neighbor links (*horizontal fragmentation*). Spectrum fragmentation degrades the blocking performance: the links may have unused bandwidth, but the lack of contiguity/continuity in the spectrum may lead to dropping requests.

Fig. 1 illustrates the fragmentation problem. A connection is to be established between nodes A and C, requiring four spectrum slots. Link A-B has four available slots, but they are not contiguous, because of vertical fragmentation. Link B-C has four contiguous slots, but they do not match with four contiguous slots in link A-B, which we refer to as horizontal fragmentation.

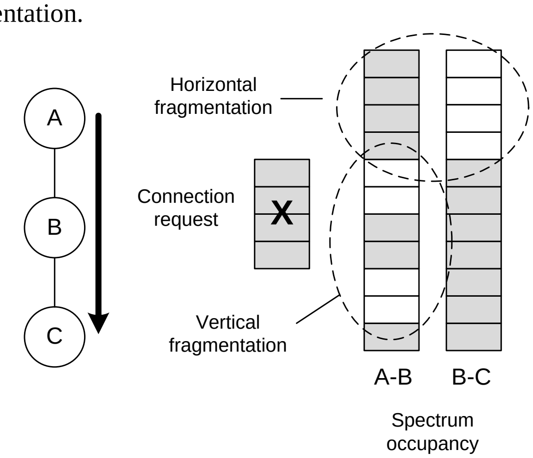

Fig. 1. Fragmentation problem in flex-grid networks.

A prominent advantage of flex-grid is its distance-adaptive nature. Whereas in wavelength-routed networks modulation is decided for the worst-case scenario (i.e. in terms of physical impairments), in flex-grid networks modulation format can be decided individually for each connection during the allocation process. Distance-adaptive concept adds to the picture a new trade-off between optical reach and spectral efficiency. For instance, robust modulations like BPSK have a long optical reach, at a cost of a low spectral efficiency. In its turn, other modulations like 16-QAM can transmit the same bit rate occupying one fourth of the bandwidth, but with a much shorter optical reach.

This work was partially supported by the Spanish project grant TEC2010-21405-C02-02/TCM (CALM) and the FPU predoctoral fellowship program of the Spanish Ministry of Education, Culture and Sport (reference no. FPU12/04571). It was also developed in the framework of project “Programa de Ayudas a Grupos de Excelencia de la Región de Murcia” funded by F. Séneca (Plan Regional de Ciencia y Tecnología 2007/2010).

Authors are with the Department of Information Technologies and Communications, Universidad Politécnica de Cartagena, Spain (e-mail: {josel.izquierdo, pablo.pavon, mvictoria.bueno}@upct.es).

In this paper, we focus on heterogeneous distance-adaptive networks. This means, multiservice networks which receive connection requests of different bandwidths (in Gbps), and that exploit the modulation degree of freedom to optimize the optical reach vs. spectral occupation trade-off. We are interested in online Routing and Spectrum Assignment (RSA) algorithms, which allocate connection requests when they are received, deciding on its route, modulation, and occupied band. Note that networks under both the heterogeneous and distance-adaptive paradigm are more prone to suffer from spectrum fragmentation, since the motivations to allocate connections with diverse spectral size are doubled.

Fragmentation has two degrading effects in network performance. First, the blocking performance averaging all the services is degraded, reducing the effective network capacity. Second, fragmentation affects the fairness among connections, since high-bandwidth connections are more likely to be blocked by vertical and horizontal fragmentation than low-bandwidth connections. This can lead to starvation of high-bandwidth services.

Despite of its importance in actual networks, RSA algorithms have been traditionally tailored for optimizing the network capacity, leaving aside the fairness issue. Actually, most of the works did not even address this performance metric. In this work, fairness is put in a more prominent place. The work presented in this paper makes a review of online RSA algorithms, and evaluate them in terms of (i) average blocking, and (ii) fairness among the blocking probabilities perceived by different services. Then, we propose a novel and simple mechanism (PSP algorithm) that permits finding a compromise between network capacity and fairness. Our evaluation study compares for the first time the network capacity and fairness performances for a moderate set of previously proposed algorithms, together with our novel proposal, exposing their merits in both dimensions. PSP algorithm has shown their strength for optimizing the network capacity, enforcing fairness among services.

The rest of the paper is organized as follows. Section II reviews the state of the art on (distance-adaptive) online RSA algorithms. Section III describes the new proposed mechanism. Section IV describes the case study and methodology. Section V presents the simulation results. Finally, Section VI concludes the paper.

II. Literature Review

In this section, we review different proposals of online RSA algorithms found in the literature. We concentrate on "classical" online RSA algorithms, leaving out of the scope of the paper topics like defragmentation [4] or split-spectrum routing [5]. Defragmentation would require the ability to rearrange existing connections to reduce the spectrum fragmentation, something that can cause traffic disruptions the operators may not be willing to face. In its turn, split-spectrum routing requires transponders able to allocate connections in non-contiguous spectrum, and it is left as a long-term alternative. A complete survey including these additional topics can be found in [6].

TABLE I
Number of slots required for each Service and Modulation Format

Feature		References
One-step algorithm	Greedy algorithm	[7]
	Auxiliary graph	[8]
Two-step algorithm	Static routing	[7][9][10][22]
	Dynamic routing	[11][12]
Distance-adaptive algorithm		[13][14]
Fragmentation-aware algorithm		[15][16][17]
Fairness-aware algorithm		[16]

In dynamic scenarios, online RSA algorithms are executed in real-time and must be as fast and simple as possible. Since RSA is a NP-hard problem [9], most existing works propose heuristic algorithms. Table I summarizes and categorizes the heuristic proposals for online RSA algorithms found in the literature. Selected algorithms for evaluation are briefly described.

A. *One-step algorithms*

Heuristic algorithms are often classified into two categories, depending on how they tackle the routing and spectrum assignment: jointly or separately. We use the term one-step algorithm for those schemes that solve the two subproblems simultaneously. Two approaches are found: greedy algorithms [7], and auxiliary graphs [8]. While the former achieve a sub-optimal solution in a rapid manner, the latter gets the optimal at cost of high computational complexity.

B. *Two-step algorithms*

In contrast to the previous case, two-step algorithms decompose the problem into two subproblems. First, the algorithm finds a set of candidate paths. Routing may be static or dynamic, depending on whether paths are pre-computed (static) or may vary according to the network state (adaptive). Then, the RSA tries to allocate the request, starting with the top of the candidate paths. Some proposals stops when a feasible allocation is found, while others evaluate every single one, and get the one that is considered better.

The first algorithm to evaluate, and the most simple one, is the *k*-shortest path routing with first-fit allocation [7][9]. In this algorithm, a set of candidate paths are pre-computed. For each candidate path, the algorithm tries to allocate the request starting the search in the lowest-index spectrum void. If not possible, it moves to the next candidate path.

A refinement of this algorithm is presented in [22]. The *k*-shortest path routing with path-priority (Load-Balancing) evaluates the set of pre-computed candidate paths on a dynamic order, given by the spectrum availability on the path. The path with the highest availability (not necessarily contiguous) is tried first, then the second one, and so on. If two paths have the same availability, the one with less hops is examined.

C. *Distance-adaptive algorithms*

These algorithms include the possibility of adjusting the modulation format for each connection [13][14]. Most of them are based on two-step algorithms with static routing. The main difference with non-distance-adaptive algorithms, is that distance-adaptive ones map each request to a modulation format, and thus the spectrum bandwidth, according to the requested data rate and the path length. Note that different paths may be associated to different modulation formats.

D. Fragmentation-aware algorithms

These algorithms try to allocate connections in such a way that fragmentation is minimized. The main idea is to provide a metric to evaluate how good (or bad) is a routing and spectrum allocation with respect to fragmentation. Then, the best allocation according to this metric is selected.

From this case, we highlight the algorithm *k*-shortest path routing with fragmentation-aware allocation (Fragmentation-aware) [15]. Here, the algorithm evaluates all possible routing and spectrum allocations according to the expressions presented in [15], and chooses the best allocation among them. Finally, although work in [16] is not strictly a fragmentation-aware algorithm, authors provide an interesting metric to quantify how fragmented the spectrum is, which can be used to construct fragmentation-aware algorithms.

E. Fairness-aware algorithms

Fairness in circuit-switched heterogeneous (or multi-service) networks was widely studied for legacy technologies (i.e. B-ISDN) more than two decades ago [20][21]. In flex-grid networks, the work presented in [16] is an approach to enforce fairness by different schemes that dedicate part of the link bandwidth to each service.

From these algorithms, we introduce in our study the *k*-shortest path routing with path-priority and dedicated-partitioning (DP). DP algorithm partitions all the bandwidth in the link, assigning one band to each service. A mechanism for dimensioning the band size is provided. The basic idea is to add DP to the Load-Balancing algorithm. The shared-partitioning scheme proposed in [16] is not included in our study, since the needed details to compute the partitions are not provided.

III. Proposal: Partial-Sharing Partitioning

In this paper, we propose the Partial-Sharing-Partitioning (PSP) scheme for managing the spectrum in flex-grid networks. The operation of the PSP model for a heterogeneous network with S services is quite simple: we separate the bandwidth in each link in $S+1$ partitions, that is, as many as services, plus one. The i-th partition within the first S ones is dedicated to allocate connections of the i-th service. The last partition ($S+1$) is used as an overflow partition shared among all the services. This means that, upon reception of a connection request for service s, the PSP controller tries to allocate it into the dedicated partition for service s. If not possible, then it tries to allocate it into the shared partition. If it fails to do so, the request is blocked. Fig. 2a illustrates a 25%-PSP scheme, where 25% of the link bandwidth is shared, while the rest of the bandwidth is equally distributed among the services (10, 40, 100, 400 Gbps). Note that a 0%-PSP corresponds to the DP algorithm (Fig. 2b), where the link spectrum is distributed into S dedicated partitions, one per service, with no shared resources. In its turn, a 100%-PSP (Fig. 2c) means that all the bandwidth is shared, and no dedicated bandwidth exists for each service.

The proposed partition dimensioning for the PSP scheme is as follows. First, the size of the shared spectrum partition should be decided, balancing the trade-off between network capacity and blocking fairness, e.g. using simulations like the ones to be shown in this paper. Then, the remaining part of the spectrum can be partitioned using the same method as the one proposed in [16] for the DP algorithm.

It is interesting to remark that the PSP scheme can be combined with different routing policies and different spectrum assignment policies, as long as they comply with the spectrum partitions assigned to each service. In particular, in this paper we test the PSP approach with a load balancing policy for selecting the connection route, and a first fit scheme for allocating the spectrum within the dedicated or shared band.

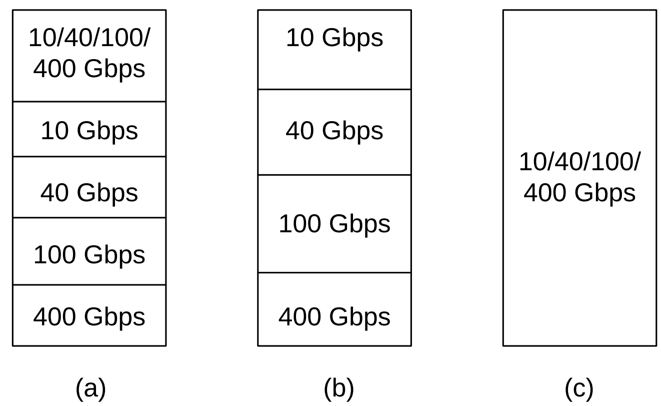

Fig 2. Partial-sharing partitioning (PSP) with different percentages of shared spectrum: (a) 25%, (b) 0%, and (c) 100%.

IV. Simulation Study

Our study is focused on a vertically-integrated operator [25] offering end-to-end connectivity via all-optical connections over its own flex-grid infrastructure. Clients make connection requests from a given set of bandwidth services. Each lightpath request is assumed to be handled by a centralized network controller (i.e. a Path Computation Element, PCE), with global information of the network state (routing and spectrum allocation of active lightpaths), which tries to allocate the request according to the RSA algorithm.

The objective is to analyze the performance of different online RSA algorithms according to two different simulation models: (i) blocking model (or long-run), and (ii) incremental (or first-passage) model [26]. In the blocking model, connection requests arrive and depart at random. The assumption is that eventually some requests may be blocked, while most of them should be successfully allocated. In contrast, the incremental model starts with an empty network (no connection is established), and requests, of infinite holding times (and thus permanent), are arriving randomly. Upon the first blocking event, simulation stops, signaling that the network would need a capacity upgrade to allocate the request.

In our opinion, both simulation models are of interest for operators and service providers. The blocking model is the main target in this paper, consistent with the future bandwidth-on-demand market, where some rare rejecting events (i.e. 1%) may be admissible. The incremental model is added to reflect present-day operation of networks not offering on-demand optical connection services: permanent connections are set up, and never/seldom modified.

A. Performance metrics

Each simulation model has its own set of performance metrics. For the incremental model, the amount of traffic carried (or network throughput) up to the first blocking event is the parameter of interest. With respect to the blocking model, we use as figures of merit the bandwidth blocking probability (BBP) [7], and the coefficient of variation of the BBP for fairness. Following, we define them formally.

Let S be a set of end-to-end services with an average offered traffic volume for service $s \in S$ equal to h_s, and an average blocking probability equal to BP_s. Then, BBP and coefficient of variable can be computed as follows.

$$BBP = \frac{\sum_{s \in S} h_s \cdot BP_s}{\sum_{s \in S} h_s} \quad (1)$$

$$CV = \frac{std(BP_s)}{avg(BP_s)} \quad (2)$$

where $std(\cdot)$ and $avg(\cdot)$ are the standard deviation and average operator, respectively.

In heterogeneous networks, since each service has different bandwidth requirements, BBP is more representative than the average BP. However, since all services are assumed to have the same quality-of-service requirements (i.e. blocking probability), we need to introduce fairness measure that capture whether all services perceive the same blocking probability. Whereas several fairness indices can be found in the literature (e.g. Jain's fairness index [19]), we choose the coefficient of variation, to weight the typical deviation among the different blockings, respect the average blocking. An algorithm is fairer when CV approaches to zero.

B. Network and traffic

Fig. 3 shows the network topology, NSFNet [18], which we used in the simulations. Table II summarize other simulation parameters. Client requests fall into four possible services, of 10, 40, 100 and 400 Gbps lightpaths. Table III shows the bandwidth requirements and optical reach for each service and modulation [22]. Connection requests are generated following a Poisson process with an average rate λ, and holding time following a negative exponential distribution with mean time equal to one second. Inter-arrival times (1/λ) are adjusted so that all services from a source-destination pair have the same intensity, and the total intensity between two nodes matches the values given by a traffic matrix *M*, which is a scaled version of the reference traffic matrix in [18] for NSFNET network. A normalization process is completed to obtain the traffic matrix to be associated to a fully-loaded network. We scale the original traffic matrix so that, if routed over the shortest path (in hops), the total in-network is equal to the total network capacity. Note that in this process the traffic matrix associated to 100% load is an upper bound to the maximum traffic that can be carried in the network: we neglect the effects of fragmentation, we do not enforce fiber capacity constraints, and we use the more spectrum-efficient modulation available.

TABLE II
SIMULATION PARAMETERS

Parameter	Value
Number of nodes	14
Number of unidirectional fiber links	42
Total spectrum available per fiber	4.5 THz
Slot granularity	12.5 GHz
Number of frequency slots per fiber	360
Total average offered traffic volume at load = 100%	124.54 Tbps
Capacity of a frequency slot with BPSK	12.5 Gbps
Maximum number of candidate paths for each source-destination pair	5

TABLE III
NUMBER OF SLOTS REQUIRED FOR EACH SERVICE AND MODULATION FORMAT

Modulation format	Optical reach (km)	Spectral efficiency (bps/Hz)	Bandwidth requirements (Gbps)			
			10	40	100	400
BPSK	9600	1	1	4	8	32
QPSK	4800	2	1	2	4	16
8-QAM	2400	3	1	2	3	11
16-QAM	1200	4	1	1	2	8

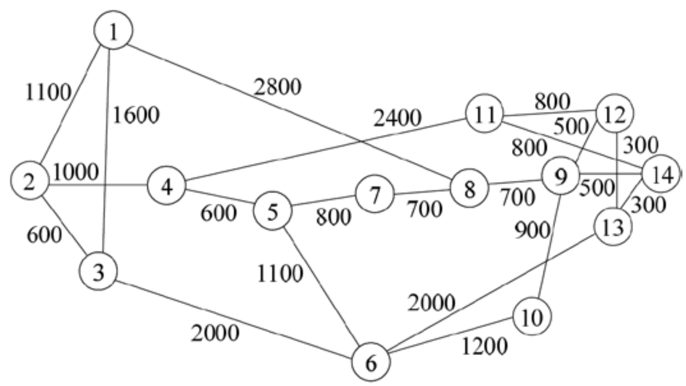

Fig. 3. NSFNet topology with 14 nodes and 21 bidirectional links (distances measured in kilometers).

C. Net2Plan open-source simulation framework

Simulations in this paper were completed in the Net2Plan Connection-Admission-Control (CAC) simulator, one of the available tools within the open-source Net2Plan framework [24]. For the blocking model, 10^6 connection requests are simulated for each load factor. For the first-passage model, simulations are running until the first blocking event happens, and results are averaged over 1000 runs.

All RSA algorithms have been developed as Net2Plan CAC algorithms, and its code is publicly-available on the website

[24] for review and reuse of interested readers. Results in this paper can be easily reproduced in Net2Plan.

D. Implemented algorithms

Due to the space limitation, we select a reduced set of algorithms to compare with our proposal. All algorithms are assumed to implement static routing, that is, a set of k-loopless shortest paths (in number of hops) are pre-computed in the initialization phase. Given an end-to-end path, the modulation format with the highest spectral efficiency is chosen, among all whose optical reach is greater or equal than the path length. Now, we enumerate all the evaluated algorithms (code names for results are put into parentheses):

- K-shortest path routing with first-fit allocation (First-FF)
- K-shortest path routing with fragmentation-aware allocation (Fragmentation-aware)
- K-shortest path routing with path-priority (Load-Balancing)
- K-shortest path routing with path-priority and dedicated-partitioning (DP)
- K-shortest path routing with path-priority and partial-sharing partitioning (PSP): This algorithm employs also the path-priority mechanisms as DP, but implementing the partial-sharing model proposed in this work. We execute simulations for this algorithm using three different sizes of the shared pool of spectrum (25%, 50% and 75% of the fiber bandwidth).

E. Results: blocking simulation model

Fig. 4 shows the bandwidth blocking probability of the different RSA algorithms. Offered load is varied from 40% to 100% in steps of 10%. As can be observed, Load Balancing algorithm offers the best average blocking results, and DP the worse, given its rigid bandwidth partitioning. PSP results are in the middle of both: the higher the sharing the better the blocking. This is logical since PSP is equivalent to DP when sharing ratio is equal to zero, and is equivalent to Load-Balancing when sharing ratio is equal to 100%. At 100% load, when the blocking is very high (10%), DP and PSP 25% improve the results. This is because at such high loads, the shared bandwidth benefits are exceeded by the excess vertical fragmentation when all the services share the bandwidth. We appreciate that DP achieves a slightly better than PSP, however, we will see how this does not translate into a fair provisioning.

Regarding to fairness, results presented in Fig. 5 indicates that at higher loads, only DP and PSP with 25-50% sharing are able to provide fairness. While CV for DP remains almost stationary at every load, at the cost of higher blocking probability at low loads, PSP is fairer as the load is increasing. In its turn, Load Balancing algorithm shows large differences in blocking probability for different services. Overall, PSP scheme provides a mean to achieve very good blocking probability figures, while enforcing fairness. The sharing ratio in PSP can be used to tune the desired balance between network capacity and fairness: the larger the sharing ratio, the better the network capacity, at a cost of worsening the fairness.

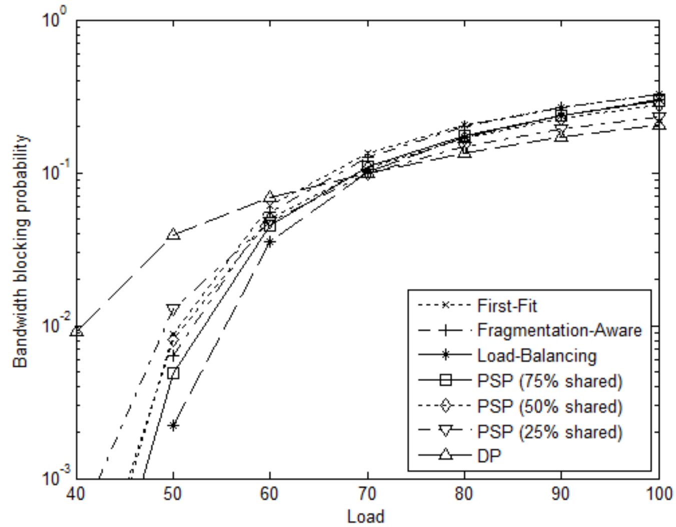

Fig. 4. Comparison of bandwidth blocking probability in logarithmic scale.

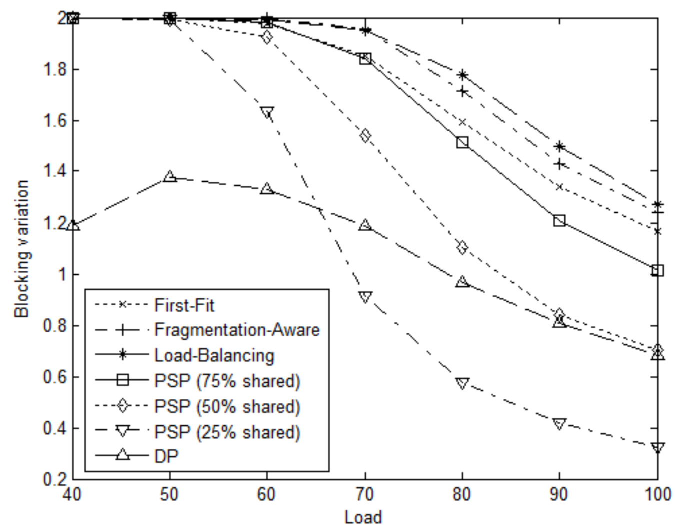

Fig. 5. Comparison of fairness measured from coefficient of variation

F. Results: incremental simulation model

Fig. 6 shows the results of the incremental model. Again, we observed the degraded performance of DP. Its rigid partitioning makes requests being blocked, even though the network is still lightly loaded. Intuitively, a strong admission control policy is acting before it is required. Our partial-partitioning schemes allows achieving a reasonable throughput before the first blocking event, validating our proposal.

V. CONCLUSION

In this paper we evaluate a large set of RSA schemes for distance-adaptive heterogeneous flex-grid networks. We are interested in exploring the trade-off between network capacity (average blocking), and fairness in the blocking probability observed by different services. Many of the RSA schemes proposed in the literature have been designed without any concern on fairness, producing unfairly large blocking probabilities to connections requesting large bandwidth. We propose the PSP spectrum management scheme, as a method to balance network capacity and fairness metrics.

We include a large set of results comparing the most relevant previous RSA proposals, together with PSP. Algorithms have been implemented in the open-source Net2Plan framework, and its code is publicly available in the website [24].

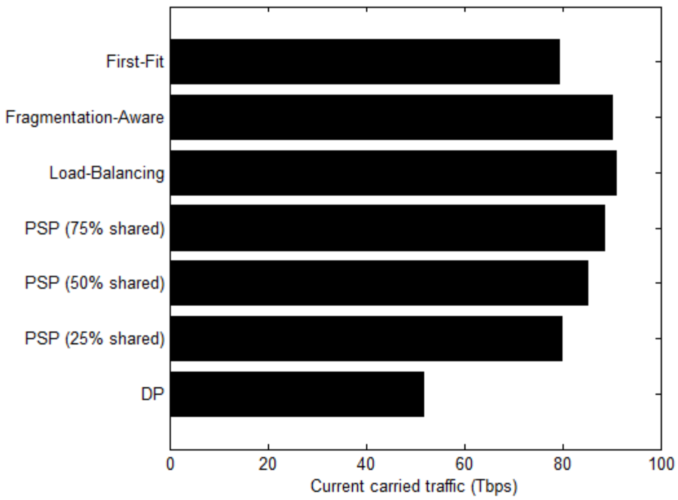

Fig. 6. Comparison of current traffic carried until the first blocking event

REFERENCES

[1] 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, November 2009.

[2] M. Jinno, H. Takara, Y. Sone, K. Yonenaga, A. Hirano, "Multiflow Optical Transponder for Efficient Multilayer Optical Networking," *IEEE Communications Magazine*, vol. 50, no. 5, pp. 56-65, May 2012

[3] O. Gerstel, M. Jinno, A. Lord, and S.J. Ben Yoo, "Elastic Optical Networking: A New Dawn for the Optical Layer?," *IEEE Communications Magazine*, vol. 50, no. 2, pp. 12-20, February 2012.

[4] A. Castro, L. Velasco, M. Ruiz, M. Klinkowski, J. P. Fernandez-Palacios, and D. Careglio, "Dynamic routing and spectrum (re)allocation in future flexgrid optical networks," *Computer Networks*, vol. 56, no. 12, pp. 2869-2883, August 2012.

[5] A. Pagès, J. Perelló, S. Spadaro, J. Comellas, "Optimal Route, Spectrum and Modulation Level Assignment in Split Spectrum-enabled Dynamic Elastic Optical Networks," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 6, no. 2, pp. 114-126, February 2014.

[6] S. Talebi, F. Alam, I. Katib, M. Khamis, R. Khalifah, and G. N. Rouskas, "Spectrum Management Techniques for Elastic Optical Networks," unpublished.

[7] X. Wan, L.Wang, N. Hua, H. Zhang, and X. Zheng, "Dynamic routing and spectrum assignment in flexible optical path networks," in *Proc. of OFC/NFOEC 2011*, March 2011.

[8] A. N. Patel, P. N. Ji, J. P. Jue, and T. Wang, "Dynamic routing, wavelength assignment, and spectrum allocation in transparent flexible optical WDM networks," in *Proc. of SPIE*, January 2011.

[9] K. Christodoulopoulos, I. Tomkos, E. A. Varvarigos, "Elastic Bandwidth Allocation in Flexible OFDM-based Optical Networks," *IEEE/OSA Journal of Lightwave Technology*, vol. 29, no. 9, pp. 1354-1366, May 2011.

[10] Y. Wang, X. Cao and Y. Pan, "A Study of the Routing and Spectrum Allocation in Spectrum-sliced Elastic Optical Path Networks," in *Proc. of IEEE INFOCOM 2011*, April 2011.

[11] R. C. Almeida Jr., R. A. Delgado, C. J. A. Bastos-Filho, D. A. R. Chaves, H. A. Pereira, and J. F. Martins-Filho, "An Evolutionary Spectrum Assignment Algorithm for Elastic Optical Networks," in *Proc. of ICTON 2013*, June 2013.

[12] Y. Wang, J. Zhang, Y. Zhao, J. Wang, and W. Gu, "Routing and spectrum assignment by means of ant colony optimization in flexible bandwidth networks," in Proc. of OFC/NFOEC 2012, March 2012.

[13] M. Jinno, B. Kozicki, H. Takara, A. Watanabe, Y. Sone, T. Tanaka, and A. Hirano, "Distance-Adaptive Spectrum Resource Allocation in Spectrum-Sliced Elastic Optical Path Network," *IEEE Communications Magazine*, vol. 48, no. 8, pp. 138-145, August 2010.

[14] T. Takagi, H. Hasegawa, K. Sato, Y. Sone, B. Kozicki, A. Hirano, and M. Jinno, "Dynamic routing and frequency slot assignment for elastic optical path networks that adopt distance adaptive modulation," in *Proc. of OFC/NFOEC 2011*, March 2011.

[15] Y. Yin, H. Zhang, M Zhang, M Xia, Z. Zhu, S. Dahlfort, and S. J. B Yoo, "Spectral and spatial 2D fragmentation-aware routing and spectrum assignment algorithms in elastic optical networks," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 5, no. 10, October 2013.

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

[17] X. Chen, S. Ma, B. Guo, Y. Wang, J. Li, X. Chen, and Y. He, "A novel fragmentation-aware spectrum allocation algorithm in flexible bandwidth optical networks," Optical Switching and Networking, vol. 12, pp. 14-23, April 2013

[18] R. Ramaswami, and K.N. Sivarajan, "Design of Logical Topologies for Wavelength-routed Optical Networks," *IEEE Journal on Selected Areas in Communications*, vol. 14, no. 5, pp. 840-851, June 1996.

[19] R. Jain, D. Chiu, and W. Hawe, "A Quantitative Measure Of Fairness And Discrimination For Resource Allocation In Shared Computer Systems", DEC Research Report TR-301, September 1984.

[20] A. Girard, *Routing and Dimensioning in Circuit-Switched Networks*, Boston, MA: Addison-Wesley, 1990.

[21] K. W. Ross, *Multiservice Loss Models for Broadband Telecommunication Networks*, Secaucus, NJ: Springer-Verlag, 1995.

[22] Z. Zhu, W. Lu, L. Zhang, and N. Ansari, Dynamic Service Provisioning in Elastic Optical Networks with Hybrid Single-/Multi-Path Routing," in *Journal of Lightwave Technology*, vol. 31, no. 1, January 2013.

[23] J. W. Causey and H. S. Kim, "Comparison of Buffer Allocation Schemes in ATM Switches: Complete Sharing, Partial Sharing, and Dedicated Allocation," in *Proc. of ICC'94*, New Orleans, LA, United States, May 1994.

[24] Net2Plan – The open-source network planner [Online]. Available: http://www.net2plan.com/ [Last accessed: January 2014]

[25] E. L. Varma, S. Sankaranarayanan, G. Newsome, Z.-W. Lin, and H. Epstein, "Architecting the services optical network," *IEEE Communications Magazine*, vol. 39, no. 9, pp. 80-87, September 2001.

[26] R. Ramaswami, K. N. Sivarajan, and G. H. Sasaki, *Optical Networks: A Practical Perspective*, 3rd ed., Burlington, MA: Morgan Kaufmann, 2010.