Resource Allocation Policies in SDM Optical Networks

(Invited Paper)

Domenico Siracusa*, Federico Pederzolli*, Dimitrios Klonidis‡, Víctor López†, and Elio Salvadori*

*CREATE-NET, Trento, Italy
‡AIT, Athens, Greece
†Telefónica I+D, Madrid, Spain

Abstract—**The capacity benefits of flexi-grid optical networks are limited by the finite amount of bandwidth of standard fibers, and will not be able to scale indefinitely. Space Division Multiplexing (SDM) has emerged as a promising technology to overcome this limitation. In this preliminary work we give an overview of current SDM technologies, covering coupled and non-coupled SDM networks, and of the possibilities and limitations they bring to the problem of Routing, Space and Spectrum Allocation (RSSA). We focus on the trade-offs between spectral efficiency and amount of transmission devices enabled by SDM technologies, presenting a number of preliminary heuristic policies to solve RSSA, and evaluating them in the context of dynamic traffic by means of simulations. We show that different heuristics can optimize different aspects of RSSA, and that one strikes a reasonable balance.**

I. INTRODUCTION

As Internet traffic continues to increase, the usable bandwidth of standard single-core, single-mode fibers is starting to approach its limits [1]. The introduction of flexi-grid technology partly addresses this problem, by enabling the instantiation of spectrally efficient super-channels, where multiple co-routed carriers are placed near the Nyquist condition (thus avoiding the need to use switching guard-bands between said carriers), and the use of adaptable modulation schemes, where more efficient modulations (i.e., with a tighter spectral footprint, or one that does not align well with the 50 GHz channels of fixed-grid DWDM) can be used [2] [3]. Despite these improvements, the capacity growth potential of flexi-grid optical networks is still limited by the finite usable bandwidth of fibers [4].

In order to overcome this limitation, the simplest form of Space Division Multiplexing (SDM) proposes to deploy multiple fibers in parallel, resulting in an n-fold increase in the amount of usable spectral resources and the introduction of a new "*space*" dimension, orthogonal to the spectral domain. The simple deployment of bundles of fibers, resulting in multiple identical optical networks operating in parallel, is, however, economically un-scalable, as the amount of supporting hardware, which includes optical switches, amplifiers, transmitters and receivers, would also need to increase by n-fold. This issue can be overcome by the introduction of shared resources: bundles of single-core, single-mode fibers can be replaced by more advanced, "integrated" types of fibers, such as multi-core fibers (MCFs) and few-mode fibers (FMFs), which can then be compounded, such as in SDM over multi-core, few-mode-per-core fibers. Likewise, integrated optical switches and amplifiers can be used to switch and amplify multiple parallel transmission media. Finally, integrated SDM transponders can be used to transmit multiple signals in parallel over SDM fibers [4]–[8].

In this work we give an overview of the various types of fiber, transmission, switching and amplification technologies for SDM, focusing on how they relate to the resource allocation problem, and, for the first time, on the trade-off between spectral efficiency and amount of transmission devices enabled by some of these technologies. We also present some preliminary heuristic allocation policies that take into account the presence of the additional space dimension and optimize different aspects of trade-off outlined above, but do not consider physical impairments. We present an initial evaluation of these policies, performed using simulations.

The rest of this work is structured as follows: Section II presents an overview of the State of the Art of SDM technologies, after which Section III describes the problem of resource allocation in SDM networks. Afterwards, Section IV presents a number of heuristic allocation schemes for SDM networks, which are evaluated using the simulations detailed in Section V. Finally, Section VI summarizes and reviews our findings.

II. OVERVIEW OF SDM TECHNOLOGY

The capabilities and constraints of a SDM network are largely a function of the technologies it employes. This section describes the current State of the Art of various SDM optical components.

A. SDM Optical Fibers

SDM networks may operate over a number of different media types:

- **Single-Mode, Single-Core Fiber Bundles** (Fig. 1 (A)): standard fibers, often deployed in bundles (to offset the costs of digging trenches).
- **Multi-Core Fibers** (Fig. 1 (B)): fibers with multiple cores within a single fiber cladding, forming multi-core fibers (MCFs). They can currently contain up to 19 cores [5], and offer an increase in available bandwidth equal to their core count (assuming each core only supports a single spatial mode).

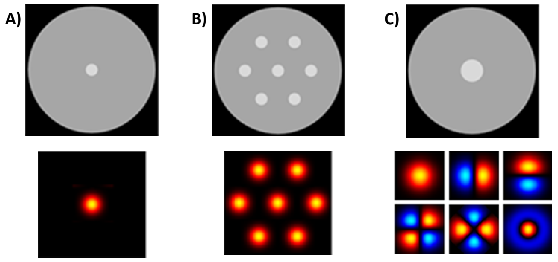

Fig. 1. Different types of SDM fibers.

- **Few-Modes Fibers** (Fig. 1 (C)): fibers with a single, large core, which can carry additional optically-guided spatial modes. Like for the case of MCFs, these few-mode fibers (FMFs) offer a potential capacity multiplier equal to the mode count.

The key difference between SDM fibers is the amount of spatial *coupling* they induce. Coupled transmission means that different spatial modes intermix, yet the amount of information being carried is retained within the set of modes, and can be unravelled at the receiver (using MIMO digital signal processing techniques) provided that all coupled modes are routed and received together. FMFs, where the modes are spatially overlapped, are inherently prone to coupling, while bundles of standard fibers are basically uncoupled as the fibers are separate.

MCFs can be constructed with coupled or uncoupled cores, depending on the core count, the distance between cores and their geometric arrangement within the cladding. The induced cross-talk in coupled MCFs can be corrected using MIMO DSP techniques similar to those used for FMFs. Uncoupled MCFs show ultra-low crosstalk properties with respect to distance (typically $< 30\,\mathrm{dB}$ over $100\,\mathrm{km}$). Finally, recent MCF designs consider few mode fiber cores [9], supporting the transmission of uncoupled spatial groups of coupled modes, where each group can be handled (e.g. routed) independently to other groups, while MIMO DSP techniques need to be applied only to the contents of each group. This case can be considered as a combination of the cases depicted in Fig. 1(B) and Fig. 1(C). For MCFs, connections to switches can be done be separating/aggregating the cores using a core fan-out/in device, then attaching to a high port-count switch using standard fibers. For FMFs, devices known as mode lamps can be used to separate the different modes.

B. SDM Optical Transponders

Current optical transmission hardware is optimized for transmitting and receiving a single beam of coherent light. In the context of flexi-grid DWDM multiple beams, from different transponders, can be placed, as long as they are co-routed through the network, at (or very near to) the Nyquist limit for increased spectral efficiency [3]. These structures, such as the one depicted in Figures 2(A) and 2(B), are known as super-channels, but with the introduction of SDM are better characterized as "*spectral super-channels*".

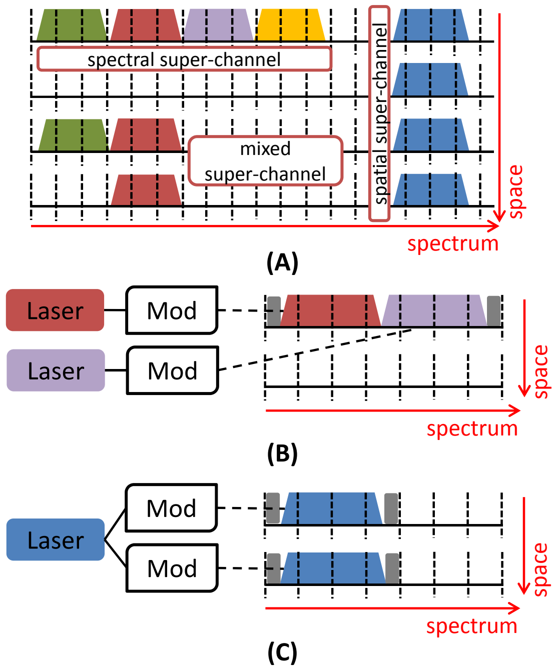

Fig. 2. Spatial vs. Spectral vs. Mixed Super-Channels (A). Spectral super-channels are more spectrally efficient, avoiding switching guard-bands (grey squares) between the carriers (B), while spatial super-channels can be generated using fewer transponders (C) but need more switching guard-bands.

Work is being done towards developing integrated "spatial" SDM transponders for parallel transmission over SDM fibers, a summary of which can be found in [10]. Such devices share a single laser source or oscillator, and allow the simultaneous transmission/reception of multiple signals using splitters/couplers. This entails that all signals must share the same frequency (i.e., position in the spectral domain). An efficient use of such devices would be the one depicted in Figures 2(A) and 2(C), where the multiple signals are transmitted over different spatial dimensions, creating a "*spatial super-channel*". Observe that such a structure is less spectrally efficient than a spectral super-channel of equivalent capacity, since switching guard-bands need to be placed on every spatial dimension. However, it can be supported with a fraction of the transmission devices (which have a very large impact on CAPEX). *Mixed "spectral-spatial" super-channels* (Fig. 2(A)), where multiple spatial super-channels are placed at the Nyquist condition creating a spectra super-channel extending over multiple spatial dimensions, are also possible. Alternatively, a technology known as an optical "comb" could be used (e.g., [11]), which would allow the creation of spectral super-channel with a single "spectral" SDM transponder. However, DSP processing may be more cumbersome in such devices than in "spatial" SDM transponders.

C. SDM Optical Switches

The combinations of the spectral and spatial domainsleads to four alternative switching scenarios:

1) **Independent Spatial/Spectral Switching**: SDM fibers bandwidth can be switched independently for every combination of spatial mode and spectral range. This is the most fine-grained scenario, but also the most complex to implement (it needs the largest number of switch ports), and is only feasible for uncoupled SDM networks.

2) **Spectral Switching**: SDM fibers bandwidth can only be switched in the spectral domain, and the filtered spectra of all spatial modes are jointly switched, forming spatial super-channels. Joint wavelength switching across all modes is mandatory for strongly coupled SDM fibers (in order to preserve the mixed information and extract it at the receiver), but can also be applied to uncoupled SDM fibers.

3) **Spatial Switching**: SDM fiber bandwidth can only be switched a whole spatial mode at a time, i.e., the entire spectrum of a spatial mode is jointly switched (an extreme form of spectral super-channel). This scenario offers a very coarse switching granularity. Furthermore, it is only feasible over uncoupled networks, although it is simple to realize.

4) **Grouped Spectral Switching**: SDM fibers bandwidth can only be switched in the spectral domain, like in scenario 2, but the joint switching of spatial modes in restricted to an entire mode subgroup. In this way, spatial (or possibly mixed spectral-spatial) super-channels extend over all spatial dimensions of one (or more) subgroups. This applies naturally to MCFs with coupled subgroups, where the switching groups can be mapped to the groups in the fiber layout, but can also benefit uncoupled fibers, as it needs far less ports than scenario 1, albeit by increasing the minimum bandwidth granularity.

III. RESOURCE ALLOCATION IN SDM NETWORKS AND PREVIOUS WORKS

In fixed-grid DWDM, before establishing a lightpath, the control plane of the network has to find a solution to the Routing and Wavelength Assignment problem (RWA), i.e., select a path and an available 50 GHz channel (lambda) on an ITU-defined grid. In this context the only available degree of freedom is the channel selection, i.e., the selection of which spectral resources to assign to the lightpath. The introduction of flexi-grid networks, where a service is not limited to a single fixed-size channel but may use a number of contiguous 12.5GHz slots in a much tighter grid, transformed RWA in the more difficult Routing and Spectrum Allocation problem (RSA) [16]. This technology enables both an increased freedom in the selection of spectral resources (due to being based on a tighter grid) and a new degree of freedom in choosing the spectral width of the channel, enabling the use of novel modulations that do not fit in the 50 GHz grid [2] [3] and the introduction of spectrally efficient "spectral" super-channels.

SDM introduces a further degree of freedom in the new "space" dimension, one that is further complicated by the different coupling characteristics of SDM fibers and the multiple possible switching scenarios outlined earlier. In this new context, we call the resource allocation problem the "*Routing, Space and Spectrum Allocation problem*" (RSSA). The additional space dimension, being orthogonal to the frequency domain, can give great freedom in the placement of signals. However, the feasibility of placements strongly depends on the characteristics of the underlying SDM fibers and switches: for fibers and switches that cause no coupling, such as bundles of standard fibers, or MCFs with large spacings between cores, spatial dimensions are practically independent (Fig. 3 (A)). However, for strongly coupled fibers or switches, such as networks employing FMFs, the signal on each spatial dimension is spread over all others (at the same frequency), requiring the joint reception of all spatial modes for signal reconstruction (Fig. 3 (C)). In such cases the space dimension is therefore degenerate. Finally, for more complex fibers and switches which exhibit group coupling, the groups may be treated as independent while the space dimension within each group is to be treated as degenerate (see e.g. Fig. 3 (B)).

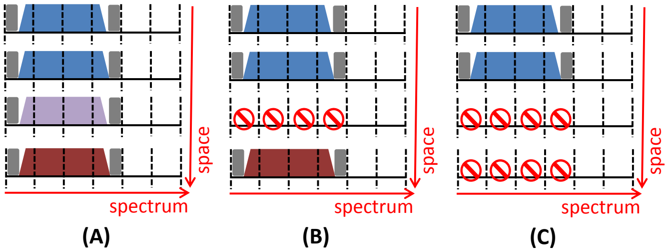

Fig. 3. Feasible allocations with no coupling (A), low coupling (B) and high coupling (C).

In addition to greater freedom in placing signals (modulo the coupling characteristics of the network), RSSA must also carefully consider the requirements that a given resource allocation places on the transmission hardware. The relation is no longer as simple as it was for RWA and RSA, as there is now, for "spatial" SDM transponders, a trade-off between spectral efficiency (favoring the use of spectral super-channels, which require more devices but less spectral resources, Fig. 2(B)) and device count (favoring the use of spatial super-channels, which require less devices but more spectral resources, Fig. 2(C)).

Another consideration is that in flexi-grid networks the resource allocation problem can be further extended to consider the transmission system impairments induced to different types of signal formats, as a function of the transmission distance. Thus, higher order modulation formats with increased spectral efficiency can be allocated over shorter distances, while longer distances can rely on the use of simpler formats to serve the requested demands. These so called impairment aware (IA) resource allocation approaches are mainly beneficial for large scale (e.g. the Pan-European or US) networks. Similar approaches can also be applied to SDM networks, however in this context an additional impairment must be considered, related to the induced crosstalk among the spatially multiplexed channels. The amount of crosstalk is a function of the transmitted distance and affects both the MIMO DSP performance (in the case of coupled SDM fibers) and the accumulated cross-talk level (in the case of uncoupled SDM fibers, particularly over long distances).

In this work we ignore the effect of impairments in optimizing the resources as well as the choice of modulation format, focusing purely on the comparison of the different allocation schemes as these are determined by the type of SDM fiber and the switching capabilities of the underlying network.

A. *Previous Works*

To the best of our knowledge, only few previous works addressed the problem of resource allocation in SDM optical networks: [12] provides an ILP formulation for the allocation problem over SDM networks employing multi-core fibers, and takes into account various parameters modelling the undesired interaction between adjacent cores. Ref. [13] proposes instead an heuristic policy to provide an approximate solution to the same problem (i.e., restricted to multi-core networks), based on avoiding the allocation of the same spectral resources on adjacent cores. Ref. [14] provides another heuristic to solve the same problem, based on maximizing the distance of utilized cores, at least at low network loads, by avoiding every other core when they are arranged in a ring. Finally, [15] provides a heuristic policy for assigning modes in a network employing FMFs, based on assigning wavelengths (the work is based on fixed-grid DWDM) on unused modes for which at least one mode is already in use by a connection to new connections that share the same source and destinations (i.e., based on "merging" small demands to better exploit the coarse granularity of FMF-based networks).

While these works propose interesting policies that address particular cases, none address the more general question of device vs. spectral efficiency in generic (multi-core, multi-mode or bundle-based) SDM networks employing spatial SDM transponders addressed in this work.

IV. RSSA POLICIES

In this section we describe a number of RSSA heuristics policies. In light of the trade-off between spectral efficiency and number of required transmission devices resulting from the use of "spatial" SDM transponders, we defined two extensions to the well-known flexi-grid First Fit policy, favoring spectral or spatial super-channels, respectively.

Spectrum-First: this policy (SpeF) uses spectral super-channels exclusively, that is, it tries to place the additional carriers of a connection at the Nyquist condition on the same spatial dimension of the first carrier of that connection. In practice, after computing the number of carriers required to serve a connection, it looks for an adequate available frequency range (i.e., a continuous spectral void large enough to accommodate all needed carriers) on the spectrum of the first spatial dimension, from left to right, moving to higher spatial dimensions if no suitable void can be found on the current dimension. This leverages the increased spectral efficiency of spectral super-channels, but also severely limits laser re-use, as each connection must be served with as many lasers as the number of carriers it requires (Fig. 2(B)). Using this heuristic, laser re-usage (assuming SDM transmitters) is limited to those cases where signals randomly happen to spectrally overlap exactly on different spatial dimensions at the same node, or when the same node originates two signals at the same frequency bound to different links, both of which are very unlikely in practice.

Space-First: this policy (SpaF) uses spatial super-channels exclusively, that is, it places the additional carriers of a connection at the same frequency of the first carrier of that connection, using a different spatial dimension for each. In practice, it iterates over all spatial dimensions in ascending order, looking for an available frequency range large enough to host one carrier. When one such range is found, the higher dimensions are also checked for a number of matching free frequency ranges equal to the remaining carriers of the connection (observe that no contiguity constraint on the spatial dimension is assumed). This leverages the laser re-use of SDM (Fig. 2(C)), but also limits spectral efficiency, as switching guard-bands are needed on all spatial dimensions.

Both of these heuristic where developed assuming SDM networks with negligible coupling. In the case where strong coupling is present, then spatial super-channels are clearly the best choice, but any spatial dimension left unused (at a certain frequency) cannot be re-used by other connections (unless a matching mechanism like the one discussed in [15] is employed, which we do not consider).

Degenerate-Space-First: this policy (DSpaF) is a variant of SpaF that models the case of networks with strong coupling (typically based on FMFs), where space is a degenerate dimension. This implies that once a portion of the spectrum is selected for a particular connection on a certain spatial dimension, the same portion of spectrum cannot be used by any other connection on any spatial dimensions on the links traversed by it. In practice it works as SpaF does, but only checks the availability of the initial frequency range on the first spatial dimension.

Finally, in a bid to optimize both spectral efficiency and number of transmission devices, we developed the following heuristic, again for uncoupled networks:

Align-Strict: this policy (AS) uses spectral super-channels; it assumes foreknowledge of the possible capacities of the incoming service requests (in order to pre-compute possible super-channels widths), and at least an approximate knowledge of their relative arrival distribution. It works by globally partitioning the spectrum into X regions, each one reserved for one of the X possible service classes (e.g. four classes, for 100, 200, 300 and 400 Gb/s). The spectral width of each region is chosen proportionally to the expected frequency of arrival of that class of requests and the size of the respective super-channel. For example, with two classes with the same arrival frequency, one of which requires twice as much spectral resources as the other, the resulting spectral bands would cover $1/3$ and $2/3$ of the available spectrum. In this way, each band can exactly contain one or more super-channels of the appropriate class. As a consequence, spectral super-channels placed on additional spatial dimensions are forced to spectrally overlap exactly with those placed on other dimensions, thus enabling far more laser re-uses than naïve SpeF. In practice, this re-use is further encouraged by placing the second, third and so on super-channels of a certain class on the additional spatial dimensions not used by the first one (going space-first in the placement of entire spectral super-channels), before verifying the feasibility of using a new spectral region on the first spatial dimension.

All of the proposed policies are limited to using only spatial or spectral super-channels, but not mixed ones. Furthermore, all policies assume the need of continuity in the spatial domain (i.e., once a spatial dimension is selected it is maintained for all links on the path), even if this may not be strictly necessary depending on the fiber and switch technology (e.g. it is unnecessary for standard fiber bundles with independent switching). These aspects will be investigated in future works.

V. SIMULATION SETUP & RESULTS

We evaluated the performance of the heuristic RSSA policies described in the previous section in a dynamic traffic scenario using a purpose-built simulation tool. The simulation uses a Poisson process for connection arrivals and an exponential holding time chosen so to obtain a desired average network load (expressed as the fraction of in-use spectral slots over the total), computed on the assumption that all connections are accepted on their shortest path. We use the well known Spanish National Backbone topology [17], and K-Shortest Path (KSP) routing with K=3, where the two additional paths are spare paths used when the resource allocation on the first one fails (due to lack of resources). We assume 384 12.5 GHz spectral slots per spatial dimension (i.e., a standard C-band fiber with 96 50 GHz WDM channels), and 4 independent spatial dimensions. We also assume up to 4 independently modulated signals per transmitter. We assume 100 Gb/s carriers, modulated using DP-QPSK, requiring 32 GHz of spectrum (colored areas in Fig. 2 (B) and (C)), and 9 GHz of switching guard-band on each side of each spectral super-channel (gray areas in Fig. 2 (B) and (C)), generated using spatial SDM transponders. In our simulations, connection requests can be for 100, 200, 300 or 400 Gb/s, uniformly distributed, which, using the modulation described above, entails spectral super-channels of 4, 7, 10 and 12 slots, and spatial super-channels of 4, 8, 12 and 16 slots (one 4-slot channel per spatial dimension), respectively. The proposed policies were evaluated in terms of Blocking Probability (BP), i.e., the ratio between refused and total connection requests at a given load, the resulting total network throughput (TP), and the average number of active transmitters needed to support them. No restriction was put on the number of total transmitters, nor on the number of transmitters per node. Each experiment simulates 10^6 bidirectional connection requests (plus an initial 10^4 requests not considered in the results, but used to reach an initial steady state), which, for a confidence level of 99%, leads to an interval of ±0.01% for the smallest values of BP (the relatively highest error margin).

As a benchmark, we used a fictional policy based on a slightly modified flexi-grid First Fit, which we called **First Fit*S** (FF*S), applied to a single virtual spatial dimension with an amount of spectral slots equal to the sum of the slots available to the other RSSA policies on all spatial dimensions. This enables us to understand the effect that the limitations posed by SDM have on the efficacy of the proposed policies.

The results for measured BP vs. Input Load are depicted in Fig. 4 (the results of the measured Bandwidth Blocking Ratio are very similar on logarithmic scale and not included in this work for space reasons); as can be expected, none of the proposed policies outperform the idealized FF*S with S=4 (represented by the red line underneath the purple AS one), since they all have more constraints or are based on less efficient spatial super-channels. The performances of SpeF are, however, quite similar to the benchmark, owing to the fact that its only difference from FF*S is that it cannot place a super-channel across one of the three boundaries between the spectra of different spatial dimensions. Likewise, AS performs rather well in terms of BP, albeit starting to block slightly earlier due to its rigid partitioning of the available spectrum. Using spatial super-channels appears to have a detectable negative effect on BP, with SpaF blocking about one order of magnitude more than SpeF at low loads, before slowly converging towards similar values as the load increases, and DSpaF blocking many orders of magnitudes more than the rest. The bad performances of DSpaF are somewhat expected, and are partly a function of the adverse conditions in which the simulations were carried out: since connections only require an average of 2.5 carriers each, and DSpaF is constrained to reserving space for 4 (one per each spatial dimension), about 37% of the spectrum it reserves is wasted. It can be expected that under more favorable conditions, such as using modulations with reduced spectral width but which exploit all of the available spatial dimensions or applying a mechanism like the one described in [15], the performances of DSpaF would be greatly increased, up to being about on par with SpaF. The performance difference between SpaF and SpeF can instead be explained taking into account that (*a*) SpaF, being space-oriented, is less spectrally efficient than SpeF and AS, and (*b*) since our experiments consider only 4 spatial dimensions (more will be considered in future works), all of which are needed for the largest spatial super-channels, a larger proportion of large requests is denied.

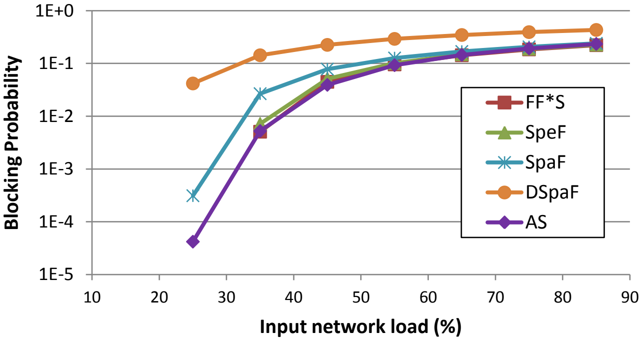

Fig. 4. Blocking Probability (BP) vs. Input Network Load for the proposed RSSA heuristics. No blocking occurs below 25% load.

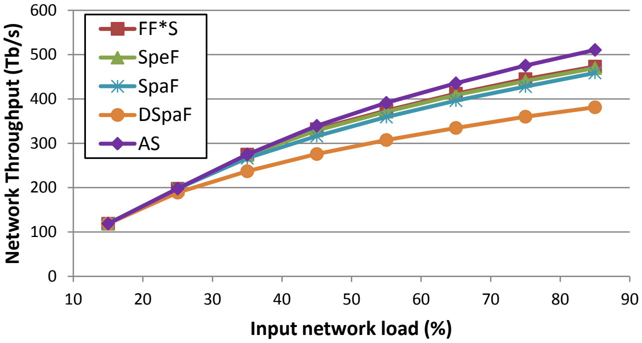

Fig. 5. Total Network Throughput (TP) vs. Input Network Load for the proposed RSSA heuristics.

Fig. 5 shows instead the measured average network throughput with respect to the input network load. With respect to this metric, the differences between SpaF and SpeF are minimal, with again a slight edge for the spectrum-oriented policy, while both are very close to the benchmark FF*S. DSpaF once again lags significantly behind all other policies, owing to its high inefficiency in the tested scenario. Of particular interest

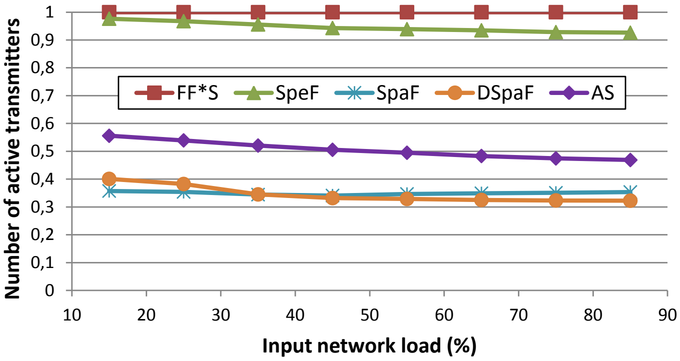

Fig. 6. Average number of active transmitters (normalized w.r.t. FF*S) vs. Input Network Load for the proposed RSSA heuristics.

is the data for AS: at low loads, where blocking is a non-issue, it performs as well as the other policies. As the load increases, it supports a level of throughput even higher than the benchmark FF*S. This is due to the fact that, by design, AS is almost perfectly fair; by this we mean that it denies requests for super-channels of different sizes with approximately the same probability, while all other policies, which attempt to fill the first usable set of resources (whether stretched over space or spectrum), tend to exhibit a significantly higher blocking rate for larger super-channels than for smaller ones, thus leading to lower spectral efficiency.

Finally, Fig. 6 shows the average number of active transmitters with respect to the input network load for all policies. All measures are normalized with respect to FF*S, which is the worst case; for reference, the number of transmitters used by FF*S at 15% load is more than 2300. While such a number is far higher than those of current real deployments, it is somewhat misleading, since it is computed on the assumption of using current transponders on top of vastly more capacious SDM networks, where the available spectrum is S times the current one (observe that, in Fig. 5, at 15% load the simulated network already carries more than 100 Tb/s). Therefore, here we focus at the ranking between the policies rather than the absolute numbers of active devices. Two clusters clearly emerge: the first one comprises FF*S and SpeF, both of which use a much larger number of transmitters than the other heuristics. The slight improvement of SpeF with respect to FF*S is due to the fact that, in the event that two carriers with the same origin but on different spatial dimensions just happen to share the same central frequency, then they can be generated by the same laser. The second cluster, comprising SpaF, DSpaF and AS, uses only a fraction of the transmitters used by the benchmark, which decreases as the load increases. This is expected for the two space-oriented policies, which prioritize the re-use of existing transmitters rather than efficient use of spectrum resources. It is, however, interesting that AS, which is spectrum-oriented, exhibits a very similar slope to that of the space-oriented algorithms, albeit with a higher starting point. This behavior is due to the ordered approach of AS, which ensures that signals on different spatial dimensions are always spectrally aligned, resulting in requiring less than 57% of the transmitters needed by SpeF, while using less than 55% more than those needed by SpaF, achieving a reasonable a balance between hardware and spectral efficiency.

VI. CONCLUSIONS

We introduced the SDM paradigm of optical networking, giving an overview of the related technologies, then focused on the problem of resource allocation for spatially-spectrally flexible networks, which in this context takes the form of the Routing, Space and Spectrum Allocation (RSSA) problem. We presented a number of novel heuristic policies for solving RSSA, and evaluated their performance in the context of dynamic traffic and spatial SDM transponders using simulations. We found that space-oriented strategies minimize the number of active lasers at the cost of (somewhat) higher blocking probability, while spectrum oriented strategies favor the latter at the cost of the former. One of the strategies presented, AS, manages to strike a good balance between the two, using only slightly more lasers than space oriented strategies, while blocking only slightly more than the other spectrum-oriented ones (and only at low loads). All of the proposed policies are limited to using only spatial or spectra super-channels, but not mixed ones, and they enforce spatial continuity; different combinations will be investigated in future works.

ACKNOWLEDGMENT

This research is partially funded by the FP7 EU INSPACE project, under grant agreement n. 619732.

REFERENCES

[1] D. Ellis et al., *Approaching the Non-linear Shannon Limit*, Journal of Lightwave Technology, 2010.

[2] M. Jinno et al., *Spectrum-efficient and scalable elastic optical path network: architecture, benefits, and enabling technologies*, IEEE Comm. Mag., 2009.

[3] O. Gerstel et al., *Elastic optical networking: a new dawn for the optical layer?*, IEEE Commun. Mag., Feb. 2012.

[4] P. J. Winzer, *Spatial Multiplexing: The next frontier in network capacity scaling*, ECOC, September 2013.

[5] J. Sakaguchi et al., *19-core fiber transmission of 19x100x172-Gb/s SDM-WDM-PDM-QPSK signals at 305 Tb/s*, OFC, March 2012.

[6] Takara et al., *1.01-Pb/s (12 SDM/222 WDM/456 Gb/s) Crosstalk-managed Transmission with 91.4-b/s/Hz Aggregate Spectral Efficiency*, ECOC, September 2012.

[7] R. Ryf et al., *Space-division multiplexing over 10km of three-mode fiber using coherent 6x6 MIMO processing*, OFC, March 2011.

[8] A. Li et al., *Space Division Multiplexed High-Speed Superchannel Transmission Over Few-Mode Fiber*, IEEE/OSA JLT, Dec. 2012.

[9] Y. Sasaki et al., *Dynamic multidimensional optical networking based on spatial and spectral processing*, Optics Express, Dec. 2012.

[10] S. Randel, *Space Division Multiplexed Transmission*, OFC, March 2013.

[11] N. Sambo et al., *Sliceable transponder architecture including multi-wavelength source*, JOCN, July 2014.

[12] A. Muhammad et al., *Routing, spectrum and core allocation in flexgrid SDM networks with multi-core fibers*, ONDM, May 2014.

[13] S. Fujii et al., *On-demand spectrum and core allocation for multi-core fibers in elastic optical network*, OFC, March 2013

[14] S. Fujii et al., *Dynamic resource allocation with virtual grid for space division multiplexed elastic optical network*, ECOC, Sept. 2013

[15] Y. Zhang et al., *Routing, wavelength and mode assignment algorithm for space division multiplexing transmission network*, IMCCC, 2012.

[16] M. Klinkowski et al., *Routing and Spectrum Assignment in Spectrum Sliced Elastic Optical Path Network*, IEEE Communications Letters, August 2011.

[17] F. Rambach et al., *A Multi-Layer Cost Model for Metro/Core Networks*, IEEE/OSA JOCN, March 2013.