Flex-Grid/SDM Backbone Network Design with Inter-Core XT-limited Transmission Reach

Jordi Perelló, Joan M. Gené, Albert Pagès, Jose A. Lazaro, and Salvatore Spadaro

***Abstract*—Spatial Division Multiplexing (SDM) has been presented as a key solution to circumvent the nonlinear Shannon limit of standard single core fibers. To implement SDM, Multi-Core Fiber (MCF) technology becomes a top candidate, leveraged by the very low inter-core crosstalk (XT) measurements obtained in real laboratory MCF prototypes with up to 22 cores. In this work, we concentrate on the design of MCF-enabled optical transport networks. To this goal, we present a methodology to estimate the worst-case transmission reach of the optical signals (at different bit-rates and modulation formats) across MCFs given real laboratory XT measurements. Next, we present an optimal Integer Linear Programming (ILP) formulation for the design of a Flex-Grid/SDM optical transport network that makes use of the transmission reach estimations. Additionally, an effective Simulated Annealing (SA)-based heuristic able to solve large problem instances with reasonable execution times is presented. Once the proposed heuristic is adequately tuned and validated, we use it to compare the resource utilization in MCF-enabled network scenarios against currently available multi-fiber link solutions. Numerical results reveal very close performances with up to 19 cores/fibers in National backbone network scenarios, and with up to 12 cores/fibers in long-haul continental ones.**

***Index Terms*—Optical network design; Flex-Grid; Spatial division multiplexing; Inter-core XT.**

I. Introduction

Ever since the nonlinear Shannon limit of single-mode optical fibers was unveiled [1], most of the research efforts have been oriented towards the so-called "next frontier" of fiber optics, that is Spatial Division Multiplexing (SDM) [2]. The two SDM flavors, namely, Mode-Division Multiplexing (MDM) and Multi-Core Fibers (MCF) are under direct competition to become the enabling technology. MDM takes advantage of the mutual orthogonality among the propagation modes of a waveguide, which can theoretically be used as independent channels [3]. This allows the total capacity to be potentially increased by two orders of magnitude. However, its main drawback is the need for Multiple-Input Multiple-Output (MIMO)-based channel equalization to undo the inherent mode coupling that, with the current State-of-The-Art (SoTA) technology, limits the capacity to just a few modes (typically referred as Few-Mode Fibers, FMFs). Indeed, the maximum number of modes proven in laboratory experiments is 15 [4].

In contrast, MCF technology is based on the insertion of several single-mode cores inside a single fiber cladding. The key parameter here is the inter-core crosstalk (XT). Interestingly, SoTA MCFs have shown extremely low inter-core XT measurements in fibers with up to 22 cores [5]-[8], thus removing the need for MIMO-based equalization. The number of cores seems difficult to be further increased due to fundamental space availability inside the fiber cladding. The use of both concepts together, referred to as few-mode-multicore fibers (FM-MCF) has also been studied, with a record spatial channel count of 114 (19 cores * 6 modes) [9]. This work focuses on the design (planning) of single-mode MCF-enabled optical transport networks that employ Flex-Grid technology to maximize the utilization of each core's spectrum when allocating lightpaths of heterogeneous bit-rates. We do not consider the FM-MCF technology in this work since, despite making the most of SDM toward ultra-high network capacity, it still requires MIMO equalization per core, thus increasing receivers' complexity.

The Flex-Grid/SDM network scenario is becoming increasingly attractive nowadays (e.g., see [10][11]), given its huge bandwidth capacity, as well as its efficiency in accommodating low bit-rate lightpaths and high bit-rate super-channels together. Specifically, we consider that lightpaths are contiguously and continuously allocated over a single core of every MCF that they traverse from source to destination (i.e., no spatial super-channels are contemplated in this work). Moreover, we also consider that any spectral portion of any input core to a node can be freely switched to any output core (spatial de-multiplexing of incoming MCFs occurs at each node). As mentioned in [11], this scheme is ideally equivalent to a Flex-Grid over network links with as many standard Single Core Fibers (SCFs) as cores in the MCFs. However, the impairments resulting from the coupling among cores (inter-core XT) can degrade the transmission reach of the optical signals, imposing the need for less efficient (but more robust) modulation formats.

Therefore, in this scenario, we want to answer the following question: *is resource efficiency of Flex-Grid over*

Manuscript received May 13, 2016.

Jordi Perelló (perello@ac.upc.edu), Joan M. Gené, Albert Pagès, Jose A. Lazaro and Salvatore Spadaro are with the Advanced Broadband Communications Center (CCABA), Universitat Politècnica de Catalunya (UPC) – Barcelona Tech, 08034 Barcelona, Spain.

MCFs really similar to that of multi-fiber links? Answering this question is key to determine if operators will be able to take profit of cost-effective integrated system components for MCFs, like transponders, amplifiers, ROADMs, etc., as envisioned in [2], or this cost reduction will be counteracted by the extra expenses required to equip more network resources to carry the same amount of traffic.

The remainder of this paper continues as follows. Section II reviews related work on the topic. Section III presents our methodology to obtain the worst-case transmission reach of the optical signals, at different bit-rates and using different modulation formats, across several SoTA MCF prototypes reported in the literature. Section IV presents our Flex-Grid/SDM network design strategies. Section V presents the evaluated backbone network scenarios and shows the obtained numerical results. Finally, Section VI concludes the paper and outlines potential future work directions.

II. Related work

Design and fabrication of MCFs with minimal XT has been possible thanks to the use of trench-assisted cores [12]. This technique provides a better confinement of the core's propagation modes. Intuitively, the reduction of the overlapping among modes coming from different cores explains why the XT is lower. The use of heterogeneous core characteristics (e.g., slight refractive index or diameter differences) is also a very effective way for further improvement. Finally, the core layout is critical because every micron that the core pitch (separation) is reduced has a huge impact on XT (~3dB/μm [12]). Estimation of XT is in general a complex task to be performed, as all these aspects need to be taken into account, plus some statistical properties of the fiber [12]. Simplified analytical expressions can be used in the case of single-mode homogeneous MCFs [13]. Besides, XT measurements have also been recently conducted in experimental MCF laboratory prototypes. The assumed XT values in this work are not based on estimations but on real measurements.

Spectral super-channels have been extensively studied given their numerous advantages [14]. SDM lays the foundation of a new paradigm, the spatial super-channel, in which several (or all) spatial channels are assigned to an end-to-end connection. The key advantage is resource sharing and integration (amplifiers, transponders, etc.) [15]. FMFs require MIMO equalization for mode uncoupling, and so spatial super-channels are a must. Conversely, MCFs are completely flexible given the low XT levels, thus both spectral and spatial super-channel flavors are possible [11].

Many works have addressed the planning of Flex-Grid (elastic) optical networks (e.g., see [16], [17]), most of them assuming SCFs. But SDM introduces a new degree of freedom into play (i.e., the space) [10], [11]. Given the novelty of the scenario, only a few works exist to date on the design of Flex-Grid/SDM networks. For example, the work in [18] addresses this goal by proposing optimal (based on Integer Linear Programming, ILP) and heuristic methods for the route, modulation format, core and spectrum assignment problem. However, the study is limited to only a 3-core MCF and it relies on estimations of the inter-core XT. The work in [19] proposes heuristics for the same problem with optical white-box (programmable architecture on demand) and black-box (hard-wired ROADM) devices, showing the substantial benefits of the former devices against the latter ones. Such a work considers a 6-core MCF and the same inter-core XT estimation expression as in [18]. Finally, the work in [20] shows the experimental evaluation of a 4-node programmable multi-granular SDM switching network using 7-core MCFs.

In this work, we consider 3 different MCFs of 7, 12 and 19 cores, whose inter-core XT characteristics have been proved in real laboratory experiments [5]-[7]. These measurements allow us to derive worst-case transmission reach estimations across such MCFs for different signal bit-rates and modulation formats, as detailed in the following section. These transmission reach estimations are used later on by the proposed Flex-Grid/SDM network design strategies.

III. Transmission reach estimation over SoTA MCFs

To estimate the transmission reach of an optically-amplified single-core fiber-optic link, several factors need to be taken into account. Nowadays, digital signal processing (DSP) capabilities of coherent receivers provide electronic compensation of both chromatic dispersion and polarization-mode dispersion. Intra-channel nonlinear effects can also be corrected by using nonlinear channel backpropagation which leaves inter-channel nonlinearities as the limiting factor [1]. As a first approximation, the optimum optical transmitted power per channel can be assumed constant for a given baud rate and channel spacing (independent of the modulation format). In a typical transport network, the optical power per channel is limited to avoid entering the nonlinear regime. Under these assumptions, noise arises as the ultimate limiting impairment, being Amplified Spontaneous Emission (ASE) the most relevant source. Depending on the bit-rate and modulation format, the minimum Signal-to-Noise Ratio (SNR) that guarantees a given Bit Error Ratio (BER) is determined. The maximum transmission reach limited by ASE noise using erbium-doped fiber amplifiers (EDFA) can be estimated as [1]:

$$L_{max,SNR} = \frac{P_S \cdot L_{span}}{SNR_{min} \cdot h \cdot f \cdot G \cdot NF \cdot R_S} \quad (1)$$

where P_S is the average optical power per channel at the transmitter, L_{span} is the distance between (equally-spaced) amplifiers, SNR_{min} is the required SNR at the receiver side (see Tab. I below), h is the Planck constant, f is the optical signal frequency, G is the gain of the amplifiers (fully compensating the losses across the associated span), NF is the noise factor of the amplifiers, and R_S is the symbol rate (including the coding overhead).

TABLE I
THEORETICAL SNR_{MIN} AT BER OF 10^{-2} [1]

BPSK	QPSK	16-QAM	64-QAM
4.2 dB	7.2 dB	13.9 dB	19.8 dB

Transmission through MCFs is also affected by inter-core XT, which may become a limiting factor. Worst aggregate inter-core XT values (measured at 1550 nm and referenced to 1 km of fiber) for SoTA 7, 12, and 19-core MCFs are shown in Tab. II. Given the inter-core XT wavelength dependency, 1550 nm has been considered as the most representative case. Assuming a 30 nm wavelength window (4 THz) imposed by the optical amplifiers a XT oscillation of ±2dB is systematically observed [5]-[7]. As can be seen, the inter-core XT levels (accounting only for the one coming from the fiber) are extremely low. Note that the negative sign is used in the inter-core XT definition.

TABLE II
WORST AGGREGATE INTER-CORE XT

7 cores [5]	12 cores [6]	19 cores [7]
-84.7 dB	-61.9 dB	-54.8 dB

The maximum transmission distance limited by inter-core XT reads [6]:

$$L_{max,XT} = 10^{\frac{XT_{dB,max} - XT_{dB,1km}}{10}} \quad [km] \tag{2}$$

where $XT_{dB,max}$ and $XT_{dB,1km}$ refer to the maximum XT limit and to the fiber's unitary inter-core XT (referenced to 1km), respectively. Both quantities are given in dB. $XT_{dB,max}$ depends on the modulation format used, as illustrated in Tab. III. Generally speaking, the higher the number of bits per symbol, the lower its tolerance.

TABLE III
IN-BAND XT VALUES FOR 1DB-PENALTY [2]

BPSK	QPSK	16-QAM	64-QAM
-14 dB	-17 dB	-23 dB	-29 dB

Tab. IV summarizes the estimated transmission reach values, obtained as $min\,(L_{max,SNR}, L_{max,XT})$, for the considered scenarios in terms of modulation format and bit rate. Calculation parameters are also provided below the table. As most operators do, a penalty margin has been added (4 dB in our case) to both ASE and XT limit values from Tab. II and III (e.g., as in [21]). Polarization multiplexing (PM) and 20%-overhead forward-error correction (FEC) are assumed to determine R_s. In Tab. IV, Noise-limited and XT-limited cases are differentiated by white and grey cells, respectively.

TABLE IV
TRANSMISSION REACH (IN KM)

Bit Rate	No. of Cores	BPSK	QPSK	16 QAM	64 QAM
40 Gb/s	7	13851	13851	5937	2289
	12	13851	12190	3062	769
	19	4755	2383	599	150
100 Gb/s	7	5540	5540	2375	916
	12	5540	5540	2375	769
	19	4755	2383	599	150
400 Gb/s	7	1385	1385	594	229
	12	1385	1385	594	229
	19	1385	1385	594	150

P_s = 1 mW	L_{span} = 100 km	G = 20 dB
NF = 5.5 dB	λ = 1550 nm	20% FEC

As can be seen, the higher the bit rate, the more limited the transmission reach is by noise. Despite the considered optimistic value of transmitted power and SoTA FEC [22], the XT-limited cases are clearly a minority. As one could expect from expressions (1) and (2), the transmission reach limitation imposed by ASE noise is inversely proportional to the bit rate, while the one imposed by XT is independent of it.

IV. FLEX-GRID/SDM NETWORK DESIGN

In this section, we present the common nomenclature that we will use to address the Flex-Grid/SDM network design, followed by a formal statement of the targeted problem. We firstly model the problem as a novel ILP formulation that reduces the required number of decision variables and constraints by several orders of magnitude compared to the previously proposed ILP formulation in [18]. Given the inherent complexity of the ILP techniques, the proposed formulation can fail when designing large Flex-Grid/SDM network instances. With this in mind, we finally present a heuristic approach based on Simulated Annealing (SA) meta-heuristic techniques to solve the same problem stated before.

A. Common Nomenclature

We model the network as a graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, where $\mathcal{N}$ represents the set of Bandwidth Variable Optical Cross Connects (BV-OXCs) that we assume able to switch any spectral portion from any input core to any output core (ensuring the spectral continuity constraint), and $\mathcal{E}$ the set of unidirectional MCFs connecting neighboring nodes. We assume that MCFs are of C cores, with their available spectrum discretized as an ordered set of Frequency Slots (FSs), denoted as $\mathcal{S} = \{s_1, s_2, \ldots, s_{|\mathcal{S}|}\}$. In this scenario, we pre-compute the set of physical paths between all source-destination node pairs that we denote as $\mathcal{P}$, being t_p and h_p the physical distance and number of hops of path $p \in \mathcal{P}$. Moreover, we denote as $\mathcal{P}_e$the subset of physical paths that include MCF $e \in \mathcal{E}$. Regarding the Bandwidth Variable Transponders (BV-TXPs) equipped at network nodes, we assume that they can operate at a set of bit-rates $\mathcal{B}$. Moreover, they can employ a set of modulation formats $\mathcal{M}$ at any bit-rate $b \in \mathcal{B}$.

We offer a set of unidirectional demands $\mathcal{D}$ to this network that we assume, for simplicity, of any bit-rate $b \in \mathcal{B}$. In addition, we do not consider traffic grooming in the lightpaths, meaning that a BV-TXP can be assigned to one demand at most, and each demand requires the allocation of exactly one lightpath.

We denote as $\mathcal{L}_d$ the set of candidate lightpaths eligible to support demand $d \in \mathcal{D}$. Specifically, a candidate lightpath at a certain bit-rate $b \in \mathcal{B}$ employing modulation format $m \in \mathcal{M}$ is defined as a subset of adjacent FSs along a path $p \in \mathcal{P}$ ensuring enough spectrum to allocate it. The number of such adjacent FSs can easily be computed as $\lceil (b/ef_m + \Delta G)/W \rceil$ where ef_m is the efficiency of modulation format m (in bits/s/Hz), ΔG the required spectrum guard bands between adjacent lightpaths (in GHz) and W the FS spectral width (also in GHz).

Note that for any demand $d \in \mathcal{D}$ we can easily pre-compute its $\mathcal{L}_d$. Firstly, we obtain from $\mathcal{P}$ the subset of physical paths from the source to the destination node of the demand that we denote as $\mathcal{P}_d$ (i.e., $\mathcal{P}_d \subseteq \mathcal{P}$). Next, for each physical path $p \in \mathcal{P}_d$ we find the most efficient modulation format $m \in \mathcal{M}$ at the same bit-rate of the demand, whose transmission reach $\geq t_p$ (enabling a communication over p). Thus, taking its efficiency into account, we can calculate the number of adjacent FSs necessary for the communication using the expression presented in the paragraph above.

This allows us to generate all possible candidate lightpaths for that demand over that path, which are all at the same $b \in \mathcal{B}$ and employ the same $m \in \mathcal{M}$, but use different spectrum portions. As an example, imagine that 2 FSs are needed to carry the demand. Such candidate lightpaths would be supported over FSs $\{s_1, s_2\}$, $\{s_2, s_3\}$, $\{s_3, s_4\}$, ..., $\{s_{|\mathcal{S}|-1}, s_{|\mathcal{S}|}\}$ in all MCFs along that path, thus enforcing spectrum continuity and contiguity constraints. Note, however, that we do not tie them to specific cores of the MCFs. Instead, we give the freedom to allocate them in one core or another of the MCFs along the path given their available spectral resources. This is supported by the core-switching flexibility assumed for the BV-OXCs in our Flex-Grid/SDM network. Lastly, we denote as $\mathcal{L}_d^e \subseteq \mathcal{L}_d$ the set of candidate lightpaths for demand $d \in \mathcal{D}$ that traverse MCF $e \in \mathcal{E}$, and as $\mathcal{L}_d^s \subseteq \mathcal{L}_d$ the set of candidate lightpaths for demand $d \in \mathcal{D}$ that employ FS $s \in \mathcal{S}$. For instance, any candidate lightpath in $\mathcal{L}_d$ supported over FS s_1 (e.g., supported over FSs $\{s_1\}$, $\{s_1, s_2\}$, $\{s_1, s_2, s_3\}$, etc., regardless of the physical path $p \in \mathcal{P}_d$ it traverses) is included in $\mathcal{L}_d^{s_1}$.

B. Problem statement

Once the common nomenclature has been presented, we formally state the Flex-Grid/SDM network design problem that we aim to address, where the Route, Modulation format, Core and Spectrum Assignment (RMCSA) is decided for each offered demand. Particularly, we aim to:

Find: the candidate lightpaths to be allocated over the network, subject to the following *constraints*:

1) *Successful demand allocation*: every offered demand $d \in \mathcal{D}$ must be assigned a feasible candidate lightpath among those in $\mathcal{L}_d$.
2) *Multi-core fiber capacity*: a given FS $s \in \mathcal{S}$ can be used at most $\mathcal{C}$ times in any MCF $e \in \mathcal{E}$.

with the *objective* to minimize the number of FSs used in any core of any MCF in the network, i.e., the minimum $|\mathcal{S}|$ value to be available in the network. In fact, while this is a typical optimization target in the related Flex-Grid literature, it becomes a coarse measure in MCF-enabled networks. For example, different designs of the same network can require $|\mathcal{S}| = 1$ but differ in up to $\mathcal{C} \cdot |\mathcal{E}| - 1$ FSs allocated (that FS allocated in only one core of one MCF vs. in all cores of all MCFs in the network). Hence, we also contemplate a secondary optimization objective in this work to be the total number of FSs allocated in the network.

Note here that by assigning a candidate lightpath to a demand we are implicitly deciding its route, modulation format and spectrum allocation. Then, the core in each MCF along the candidate lightpath route is decided taking the occupation of the FSs into account, as mentioned before.

C. Optimal ILP formulation

In this sub-section we present a novel ILP formulation for the identified Flex-Grid/SDM network design problem, hereafter referred as ILP-RMCSA. To this end, the following decision variables are introduced:

$x_{d,l}$: binary; 1 if demand $d \in \mathcal{D}$ employs candidate lightpath $l \in \mathcal{L}_d$, 0 otherwise.

$y_{e,s}$: binary; 1 if FS $s \in \mathcal{S}$ is being utilized in any core of MCF $e \in \mathcal{E}$; 0 otherwise.

z_s: binary; 1 if FS $s \in \mathcal{S}$ is being utilized in any core at any MCF in the network; 0 otherwise.

The details of the proposed ILP-RMCSA formulation are the following:

$$minimize\ F = \sum_{s \in \mathcal{S}} z_s + \epsilon \sum_{d \in D} \sum_{l \in \mathcal{L}_d} h_l S_l x_{d,l} \quad (3)$$

$$\sum_{l \in \mathcal{L}_d} x_{d,l} = 1, \forall\, d \in \mathcal{D} \quad (4)$$

$$\sum_{d \in D} \sum_{l \in \mathcal{L}_d^e \cap \mathcal{L}_d^s} x_{d,l} \leq |\mathcal{C}| \cdot y_{e,s}, \forall\, e \in \mathcal{E}, s \in \mathcal{S} \quad (5)$$

$$\sum_{e \in \mathcal{E}} y_{e,s} \leq |\mathcal{E}| \cdot z_s, \forall\, s \in \mathcal{S} \quad (6)$$

Objective function (3) minimizes the number of FS used in any core of any MCF in the network (i.e., minimum required $|\mathcal{S}|$). Additionally, it minimizes the total number of FS allocated in the network as a secondary optimization goal. For this, a second term is added in (3), being ϵ a very small real-valued positive number, h_l the number of hops of candidate lightpath $l \in \mathcal{L}_d$ and S_l the number of contiguous FSs that it requires for the communication. Thus, while multiple solutions can lead to the same $|\mathcal{S}|$ value, the model selects the one requiring the allocation of the lowest total number of FSs. As for the constraints, constraint (4) ensures that a unique candidate lightpath is assigned for every demand in the demand set. Constraint (5) guarantees that at most $\mathcal{C}$ lightpaths can employ FS s at MCF e. We remind the reader that we do not tackle the specific core assignment for a particular MCF link due to the core-switching flexibility of the BV-OXCs nodes mentioned before. Finally, constraint (6) assigns the proper value to variables z_s, accounting for such FSs used.

Although reference [18] proposes an ILP formulation for a similar problem, we advocate for our proposal as it achieves the same optimization purpose but with drastically reduced number of decision variables and constraints. Indeed, the number of decision variables of the ILP formulation in [18] is in the order of $(K \cdot |\mathcal{D}| + 3 \cdot |\mathcal{D}|^2 \cdot |\mathcal{E}| \cdot |\mathcal{C}|)$, being K the number of candidate paths for each demand, while the number of constraints is in the order of $\mathcal{O}(3 \cdot |\mathcal{D}| + K \cdot |\mathcal{D}| + 2 \cdot |\mathcal{D}|^2 \cdot |\mathcal{E}| \cdot |\mathcal{C}|)$. In contrast, the number of variables and constraints of our proposed ILP formulation are in the order of $\mathcal{O}(K \cdot |\mathcal{D}| \cdot |\mathcal{S}| + |\mathcal{E}| \cdot |\mathcal{S}|)$ and $\mathcal{O}(|\mathcal{D}| + |\mathcal{E}| \cdot |\mathcal{S}|)$, respectively. It can be appreciated that our proposal significantly reduces the number of variables and constraints, most probably reducing the time complexity of the problem as well.

To better highlight this complexity reduction, Tab. V depicts the exact number of decision variables and constraints of our ILP-RMCSA formulation, as well as those of the ILP in [18]. To this end, we have particularized the above expressions considering the 6-Node TEST network presented later on in section V, a set of 1000 offered demand requests, 320 FSs initially available per core and 7-core MCFs. Moreover, the K=3 physically shortest paths (in km) have been considered when generating $\mathcal{L}_d$ for every demand $d \in \mathcal{D}$. As seen, several orders of magnitude less variables and constraints can be appreciated, thus, highlighting the reduced complexity of our proposed ILP formulation.

TABLE V
NUMBER OF DECISION VARIABLES AND CONSTRAINTS

	Num. of decision vars.	Num. of constraints
ILP in [18]	$3{,}4 \cdot 10^8$	$2{,}2 \cdot 10^8$
ILP-RMCSA	$9{,}7 \cdot 10^5$	$6{,}1 \cdot 10^3$
Relative reduction	**350 times**	**36065 times**

D. Heuristic approach

As previously mentioned in the literature, even the simpler Routing and Spectrum Assignment (RSA) problem in elastic optical networks is of NP-hard complexity [17]. In this work, we address the RMCSA problem, which adds additional complexity to RSA, hence making ILP-RMCSA not solvable to optimality for large problem instances in a reasonable amount of time.

With this in mind, in this subsection we propose a heuristic approach, called SA-RMCSA, which allows solving large instances of the targeted problem with practical execution times. SA-RMCSA runs an SA-based meta-heuristic that guides a simpler and fast candidate lightpath selection greedy heuristic. SA is a well-known probabilistic meta-heuristic method inspired from the annealing processes in metallurgy. In SA, a temperature parameter is initialized and cooled-down as the algorithm evolves. The higher this temperature, the more probable is to accept non-improving solutions, which allows SA escaping from local optima. This happens more frequently at the early stages of the algorithm, while only better solutions are generally accepted at the end (temperature values become low). The interested reader can find additional information on the SA metaheuristic method in [23].

Fig. 1 details the SA-RMCSA heuristic pseudo-code, whose goal is to find those candidate lightpaths allowing the successful delivery of all offered demands in $\mathcal{D}$ over $\mathcal{G}$, while minimizing the same objective function F as in ILP-RMCSA. We name the eventually selected set of candidate lightpaths as *BestSol*. To this end, the heuristic starts with a pre-computation stage where the set of candidate lightpaths eligible to support every demand $d \in \mathcal{D}$, i.e., $\mathcal{L}_d$, are firstly computed (line 1) and subsequently sorted by physical distance in increasing order (candidate lightpaths over the same physical path are sorted by position in the spectrum, also increasingly). Secondly, the minimum number of FSs required by every demand $d \in \mathcal{D}$ is set as the number of FSs required by the first candidate lightpath in $\mathcal{L}_d$. Thirdly, demands in $\mathcal{D}$ are sorted according to their required number of FSs, in descending order (line 3).

Input: $\mathcal{G}$, $\mathcal{D}$, Φ, φ, τ, Λ, *maxIter*
Output: *BestSol*

1: Compute and sort $\mathcal{L}_d$ for every offered demand $d \in \mathcal{D}$
2: Set the minimum number of required FS by every $d \in \mathcal{D}$ as the number of FSs of the first $l \in \mathcal{L}_d$
3: Sort $\mathcal{D}$ according to the minimum number of required FS in descending order
4: *BestSol* ← **C-RMCSA** ($\mathcal{G}$, $\mathcal{D}$, $\cup_d \mathcal{L}_d$)
5: $T = \frac{-\Phi}{\ln(\varphi)}$
6: *iter* = 0
7: **while** *iter* < *maxIter* **do**
8: Π_1 ← Select Λ different demands randomly from $\mathcal{D}$
9: Π_2 ← Select Λ demands randomly from $\mathcal{D}$ not in Π_1
10: **swap** ($\mathcal{D}$, Π_1, Π_2)
11: *Sol* ← **C-RMCSA** ($\mathcal{G}$, $\mathcal{D}$, $\cup_d \mathcal{L}_d$)
12: $\Omega = F_{Sol} - F_{BestSol}$
13: **if** $\Omega < 0$ **then**
14: *BestSol* = *Sol*
15: **else**
16: **if** $e^{-\Omega/T} \le$ rnd ≤ 1 **then** //rnd = random [0,1)
17: **swap** ($\mathcal{D}$, Π_1, Π_2) //Restore original order
18: $T = T \cdot \tau$ //Decrease temperature
19: *iter*++
20: **return** *BestSol*

Fig. 1. SA-RMCSA heuristic pseudo-code

An initial solution of the problem is obtained by running the simpler C-RMCSA greedy heuristic (line 4) previously presented in [24], whose operation is reviewed later on in Fig. 2. This initial solution becomes the *BestSol* so far. Next, before initiating the SA iterative procedure, we also set the initial temperature (T) value (line 5). To understand the expression used for this purpose, note that the SA meta-heuristic typically accepts non-improving solutions with probability equal to $e^{-\Omega/T}$, i.e., the Boltzmann function, where Ω is the difference between the objective function value of the current solution (*Sol*) minus that of *BestSol* found until the moment (a positive value in non-improving solutions). Therefore, $T = -\Phi/\ln(\varphi)$ initially permits SA to accept with probability φ those non-improving solutions that increase by Φ the objective function value of the best solution found so far. Both Φ and φ parameters should be appropriately configured depending on the specific problem instance. For example, by setting $\Phi = 1$ and $\varphi = 0.2$ in our scenario, SA would start accepting with probability around 0.2 non-improving solutions that require up to 1 additional FS used in any core of any MCF, compared to those in *BestSol*. Here, we should point out that even accepted, a non-improving solution does not turn to be *BestSol*. However, it may allow SA escaping from a local optimum. Moreover, T is decreased by the cooling rate factor τ per SA iteration, so the probability to accept non-improving solutions reduces as the heuristic evolves.

From lines 7-19 the SA iterative procedure is shown. In order to effectively explore the solution space, up to Λ demands are swapped per iteration (lines 8-10) before running again the C-RMCSA greedy heuristic with the reordered $\mathcal{D}$. Note that the capacity of a MCF-enabled transport network can be huge, and so the number of demands that can be offered (i.e., $|\mathcal{D}|$). Hence, swapping only a single pair of demands may produce solutions that differ only slightly from one another, preventing SA to escape from local optima. Specifically, Π_1 and Π_2 are the two ordered sub-sets of demands with size Λ swapped, $\Pi_1, \Pi_2 \subset \mathcal{D}$, $|\mathcal{D}| \gg \Lambda$. After this demand swapping procedure, C-RMCSA is run again to obtain a new *Sol* (line 11). Being this one better than *BestSol* found so far ($\Omega < 0$), *Sol* is stored as *BestSol* (lines 12-14) Otherwise, the non-improving solution can also be accepted with probability $e^{-\Omega/T}$. If so, the new order of $\mathcal{D}$ is kept. Conversely, the swapping operation is undone and the previous $\mathcal{D}$ order is restored (lines 16-17). The value of T is lowered by the cooling rate τ per iteration (line 18). Lastly, the *BestSol* encountered along the *maxIter* iterations is returned as the output of SA-RMCSA.

For better comprehension of the overall SA-RMCSA heuristic performance, Fig. 2 illustrates the performance of the greedy C-RMCSA heuristic used in SA-RMCSA to rapidly solve the RMCSA problem instances with different orderings of $\mathcal{D}$. This greedy heuristic takes $\mathcal{G}$ and $\mathcal{D}$ as inputs, as well as the pre-computed sets of candidate lightpaths to serve all demands in $\mathcal{D}$. Basically, C-RMCSA runs an iterative process, where at each iteration it sets the highest allocable FS in any core of any MCF in the network (*maxFS*) to its value in the previous iteration (initially *maxFS* =0, line 1) plus the number of required FSs of the first pending demand in $\mathcal{D}$ (line 3). The first time when C-RMCSA is run, this demand is the largest pending one (in terms of required FSs), although this can change in the following executions due to the swapping of demands in $\mathcal{D}$. Then, for each pending demand $d \in \mathcal{D}$, an available candidate lightpath to support it not exceeding *maxFS* is searched on a first-fit basis. Recall that candidate lightpaths in $\mathcal{L}_d$ are sorted by physical distance and spectral position increasingly. Thus, the shortest ones in the lowest spectral parts are tried first. Having found an available candidate lightpath (line 6), the required FSs are reserved in the MCFs composing the end-to-end path and the demand is considered as served (line 7). A first-fit core assignment strategy is followed when reserving the required FSs. Once all demands are served, the resulting solution *Sol* is returned as the output of the heuristic (line 9).

Input: $\mathcal{G}$, $\mathcal{D}$, $\cup_d \mathcal{L}_d$
Output: *Sol*

1: *maxFS* = 0
2: **while** any pending demand in $\mathcal{D}$ **do**
3: *maxFS* += num. of required FSs by the first pending demand in $\mathcal{D}$
4: **for each** pending demand $d \in \mathcal{D}$ **do**
5: $l \leftarrow$ First available candidate lightpath in $\mathcal{L}_d$ that uses FSs $\in \{1, .., maxFS\}$
6: **if** l found **then**
7: Reserve the FSs supporting l in the cores of the MCFs throughout its route
8: Consider d as served
9: **return** *Sol*

Fig. 2. C-RMCSA greedy heuristic pseudo-code

V. Numerical Results

This section aims to assess the spectral efficiency of Flex-Grid/SDM backbone networks taking into account the limitation that inter-core XT introduces on the transmission reach of the optical signals. To this end, we start presenting the details of the scenario used for this purpose, as well as the assumptions made. Next, we find the most interesting values of the key parameters involved in the operation of our SA-RMCSA heuristic, subsequently validating its goodness against the optimal ILP-RMCSA formulation previously presented in section IV.B. To solve the latter, the commercial CPLEX v12.5 optimization software has been used. All executions are run in a commercial 8-core Intel i7 PC at 3.4 GHz with 16 GB RAM. Once the SA-RMCSA performance is validated, we use it to compare the spectral utilization in several Flex-Grid/SDM backbone network scenarios employing MCFs against their equivalent MF scenario, which does not suffer from inter-core XT as signals travel over separated parallel fibers per link.

A. Scenario details and assumptions

In order to generalize our findings as much as possible, we contemplate three different network topologies in our studies (Fig. 3): a small TEST network (6 nodes and 16 unidirectional links) where ILP-RMCSA incurs reasonable execution times, as well as the National Deutsche Telekom (DT) network (12 nodes and 40 unidirectional links) and a European-wide (EON) backbone network (11 nodes and 36 unidirectional links). In all these networks, we assume that the entire C-Band 4 THz spectrum is initially made available per core, discretized in FSs of 12.5 GHz (as recommended by the ITU-T in [25]), thus resulting into 320 FSs/core. Regarding the BV-TXPs equipped at the network nodes, we assume that they can operate at 40, 100 and 400 Gb/s. Moreover, they can employ any of the following modulation formats (PM is always assumed): BPSK, QPSK, 16-QAM and 64-QAM, namely, the same bit-rates and modulation formats previously considered in the transmission reach estimations in section III. As for the spectral guard bands, we consider 10 GHz between adjacent connections, a typical assumption in the literature (e.g., in [26]-[28]).

The aforementioned network scenarios are loaded with a set of offered uniformly distributed unidirectional demands, following either Traffic Profile (TP)-1 or TP-2. Specifically, TP-1 simulates a short-term network scenario where 30% of the offered demands are of 40 Gb/s, 50% of 100 Gb/s, and the remainder 20% of 400 Gb/s. In contrast, TP-2 represents a longer-term network scenario where the offered demands are only of 100 Gb/s (40%) and 400 Gb/s (60%). K=3 physically shortest paths (in km) have been considered when generating $\mathcal{L}_d$ for every demand $d \in \mathcal{D}$. Note that the very limited transmission reach of the signals at 400 Gb/s can prevent demands reaching their destination in the considered backbone networks, no matter the modulation format used. Therefore, during the candidate lightpath pre-computation, when a 400 Gb/s demand $d \in \mathcal{D}$ is unfeasible due to transmission reach even with the least efficient modulation format over path $p \in \mathcal{P}_d$, we try to build its candidate lightpaths as 4x100 Gb/s lightpaths, contiguously allocated and jointly switched from source to destination. In this case, the number of FSs required is 4 times that of a candidate lightpath at 100 Gb/s employing the selected modulation format, as they must include 4x10 GHz guard-bands.

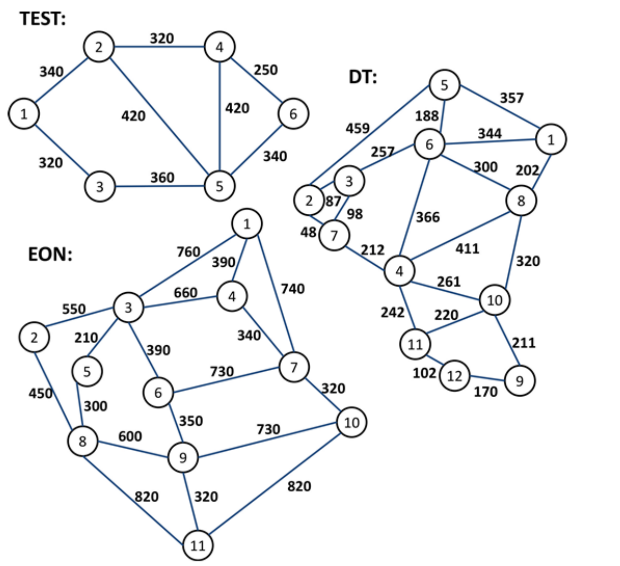

Fig. 3. Network topologies used. Link distances are shown in km.

B. SA-RMCSA tuning and performance validation

As discussed in previous section, the performance of the SA-RMCSA heuristic depends on the proper configuration of several parameters (Φ, φ, τ, Λ, and *maxIter*). Therefore, we start discussing how do we configure them to ensure good SA-RMCSA heuristic performance. Next, we validate its performance against the optimal results obtained by solving the ILP-RMCSA formulation.

After numerous SA-RMCSA executions to solve many heterogeneous network scenarios we have decided to set *maxIter* = 10000. Indeed, further SA iterations do not generally translate into objective function improvements, while unnecessarily increasing execution times. Moreover, the cooling rate per iteration has been set τ=0.9999. Regarding the number of demand pairs swapped per SA iteration (Λ), it should allow the heuristic escaping from local optima when exploring neighboring solutions, but without compromising the proper evolution of SA-RMCSA toward good solutions. To achieve this, we configure Λ according to the size of $\mathcal{D}$ in the specific problem instance. The rationale behind this is that swapping a demand pair may effectively explore the solution space when having to allocate few hundreds of demands. However, it may be completely ineffective when the number of demands increases to several thousands, which can easily happen in MCF-enabled optical networks given their huge capacity. For SA-RMCSA, we configure $\Lambda = \lfloor |\mathcal{D}|/500 \rfloor + 1$, leading to $\Lambda = 1$ if $|\mathcal{D}| \in [0,499)$, to $\Lambda = 2$ if $|\mathcal{D}| \in [500,999)$, etc.

The remainder SA-RMCSA parameters to be configured are Φ and φ, which influence the initial temperature (T) value. To investigate their adequate values, we initially set φ = 0.2 and run several SA-RMCSA executions with different Φ value. For this, we employ the small TEST network with 7-core MCFs, where we offer a set of 1000 offered demands following TP-1. Fig. 4 depicts the total number of FSs allocated in the network for Φ = 1, 2, 3 and 4 along the SA iterations, where markers show the improving solutions found. Please recall from previous section that such Φ values allow SA to initially accept with probability around 0.2 non-improving solutions that require up to $|\mathcal{S}|_{BestSol}$ +1, +2, +3 and +4, respectively.

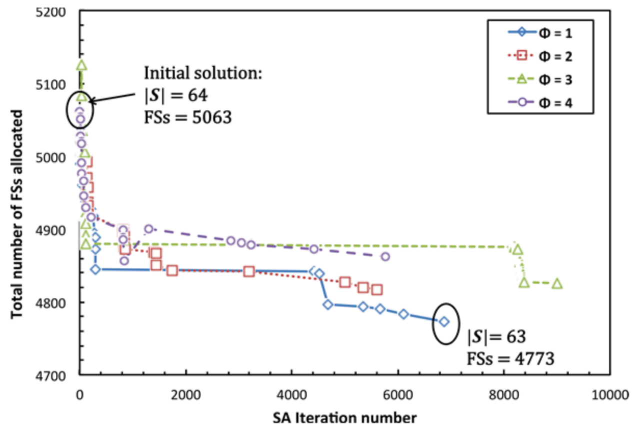

Fig. 4. Total number of FSs allocated in the 6-node TEST network along the SA iterations. The initial and final $|\mathcal{S}|$ values are shown.

From the results in Fig. 4, we can see that the largest considered Φ values do not provide the most efficient network design (in terms of resources) but the opposite. This is because SA-RMCSA starts accepting non-improving solutions very frequently and loses effectivity toward its optimization goal. Conversely, the smaller Φ values permit SA-RMCSA to still accept non-improving solutions, so as to escape from local optima, but only those of relatively high quality. Among them, Φ = 1 seems to find the best network design, reducing $|\mathcal{S}|$ from 64 to 63 and the total number of FSs allocated from 5063 to 4773, compared to the initial solution found (initial C-RMCSA heuristic execution). An interesting observation from the figure is that the total number of FSs allocated does not necessarily have to monotonically decrease along the SA executions. Indeed, when $|\mathcal{S}|$ is not a minimization target, demands can follow the shortest paths (in terms of hops) from their source to destination nodes if resources are available, which typically minimizes the total number of FSs allocated. This does not happen when minimizing $|\mathcal{S}|$ as in this work, since demands may have to traverse longer routes, thus requiring the allocation of extra FSs. This can be appreciated for Φ = 4 curve at approximately iteration 1000, where SA-RMCSA finds a solution reducing $|\mathcal{S}|$ from 64 to 63, but at expenses of increasing the total number of FSs allocated in the network. Although not explicitly shown in the figure, all Φ values eventually find solutions where $|\mathcal{S}|$ = 63, but resulting in substantial differences in terms of total number of FSs allocated. From these results, we configure φ = 0.2 and Φ = 1 in all SA-RMCSA executions from now on.

After finding the most appropriate values for all SA-RMCSA parameters, we validate its goodness against the results provided by the optimal ILP-RMCSA formulation. For this, we also employ the small TEST network with 7-core MCFs. Indeed, ILP-RMCSA is not a valid option for planning the larger DT and EON backbone networks in realistic times, especially as the number of cores in the MCFs grows, and so the number of demands that can be offered until substantially filling the network. Moreover, we consider randomly generated sets of 250, 500, 750, 1000 and 1500 offered demands, all following TP-1. For better comparison, ILP-RMCSA and SA-RMCSA take the same sets of offered demands. Note that a 2% optimality gap and a maximum execution time of 12 hours have been configured in the CPLEX solver in each ILP-RMCSA execution. In such situations where the CPLEX solver exceeds the 12h time limit without reaching the optimality gap of 2%, the estimated optimality gap at the end of the execution is kept. The obtained results are shown in Tab. VI, where execution times are shown in seconds.

TABLE VI
SA-RMCSA PERFORMANCE VALIDATION

	ILP-RMCSA			SA-RMCSA				
$\lvert\mathcal{D}\rvert$	$\lvert\mathcal{S}\rvert$	*FSs*	Exec. time (s)	$\lvert\mathcal{S}\rvert$	*FSs*	Exec. time (s)	G_S(%)	G_{FS}(%)
250	17	1160	>12h (7.48%)	17	1162	2	0	0.17
500	36	2548	406.1	36	2586	12	0	1.49
750	45	3482	894.5	46	3605	31	2.2	3.53
1000	63	4659	>12h (2.09%)	63	4773	78	0	2.44
1500	101	7392	>12h (3.55%)	101	7644	261	0	3.41

As observed, almost identical $|\mathcal{S}|$ value is reached by ILP-RMCSA and SA-RMCSA no matter the size of $\mathcal{D}$. Only when $|\mathcal{D}|$=750, SA-RMCSA requires 1 additional FS used in any core of any MCF of the TEST network, which implies a 2.2% error gap (denoted as G_S in the table). Regarding the secondary optimization target, i.e., the total number of FSs allocated in the network, SA-RMCSA also achieves very close results to those of ILP-RMCSA, resulting in error gaps (G_{FS}) below 3.55% in all executions. Finally, as for the total execution time note that, even in the small TEST network, ILP-RMCSA required more than 12 hours to find the optimal solution, which highlights its limited scalability. Conversely, SA-RMCSA spends 261s at maximum in the scenario with 1500 offered demands.

The results in the previous table show the good performance of SA-RMCSA providing close to optimal results in the small TEST network. The reader might question, however, about its added value against the simpler C-RMCSA greedy heuristic previously proposed in [24]. To answer this request, Tab. VII compares SA-RMCSA and C-RMCSA when designing the TEST, EON and DT networks with 7-core MCFs. Specifically, 1000 demands are offered to the TEST network, while 3000 demands are offered to EON and DT networks. Moreover, such offered sets of demands follow either TP-1 or TP-2, as indicated in each scenario accordingly.

Regarding the configuration of the SA-RMCSA parameters (Φ, φ, τ, Λ, and *maxIter*) to address the design of the EON and DT networks, we keep their most appropriate values found before for the TEST network. Although a re-tuning of parameters is generally advisable for any heuristic when the characteristics of the scenario change, we did not find substantial SA-RMCSA performance improvements with other values of such parameters when designing the EON and DT networks. This is why we decided to keep them fixed for all the remaining experiments in the paper.

TABLE VII
SA-RMCSA IMPROVEMENT TO C-RMCSA

	SA-RMCSA			C-RMCSA		
Scenario	$\lvert\mathcal{S}\rvert$	*FSs*	Exec. time (s)	$\lvert\mathcal{S}\rvert$	*FSs*	Exec. time (s)
TEST-TP1	63	4773	78	64	5063	< 1
TEST-TP2	119	8729	229	120	8933	< 1
EON-TP1	104	20227	1007	108	21390	< 1
EON-TP2	181	38429	2818	196	41159	< 1
DT-TP1	118	20901	1068	130	20696	< 1
DT-TP2	197	34560	2904	215	34068	< 1

As observed, while C-RMCSA provides fast (in less than 1s) and significantly good results in terms of $|\mathcal{S}|$, i.e., its optimization goal, it cannot reach those of SA-RMCSA as the size of the scenario grows up. For example, in the EON network with 3k demands following TP-2, C-RMCSA requires 15 additional FSs used in any core of any MCF in the network, as well as around 2700 FSs more allocated in total. Similar differences are seen in the DT network with 3k demands following TP-2, increasing $|\mathcal{S}|$ by 18. These results illustrate the benefits of the SA procedure in SA-RMCSA, serving as a further motivation to using it.

C. Comparison of MCF against MF technologies

A key concern on the way to an extensive MCF deployment is the negative effects that inter-core XT can have on the transmission reach of the optical signals, making advanced modulation formats unfeasible and, thus, leading to poor resource utilization. This subsection aims to provide insight into this issue by comparing the MCF-enabled optical network designs resulting from the SA-RMCSA executions with the equivalent MF ones, where transmission reach is only limited by ASE noise. These studies are conducted for the EON and DT backbone network topologies with 7, 12 and 19 cores or fibers per link (depending on whether a MCF or a MF scenario is designed). Moreover, to reasonably fill such network scenarios 3000, 5000 and 8000 demands are offered to them, respectively, either following TP-1 or TP-2.

Fig. 5 depicts the number of FSs used (top) and total number of FSs allocated (bottom) in the EON network with

7, 12 and 19 cores/fibers. As seen, almost no difference is observed between the MCF/MF-enabled scenarios with 7 and 12 cores/fibers, in terms of neither $|\mathcal{S}|$ nor total FSs allocated, which highlights the good behavior of MCFs even with a moderately large number of cores. In fact, this outcome is completely expectable with 7-core MCFs, since the transmission reach of the optical signals is always limited by noise and never by inter-core XT. With 12 core-MCFs, inter-core XT limitation exists for most modulation formats at 40 Gb/s. However, transmission reach with QPSK at 40 Gb/s still remains longer than 10000 km. Moreover, with the assumed 10 GHz guard bands, 2 FSs are required by QPSK at 40 Gb/s, as well as by 16-QAM and 64-QAM at the same bit-rate. Hence, even though some 40 Gb/s lightpaths may have to employ QPSK in the MCF-enabled scenario, instead of 16-QAM or 64-QAM as in the equivalent MF one, this is not translated into additional resources, as reflected in the results.

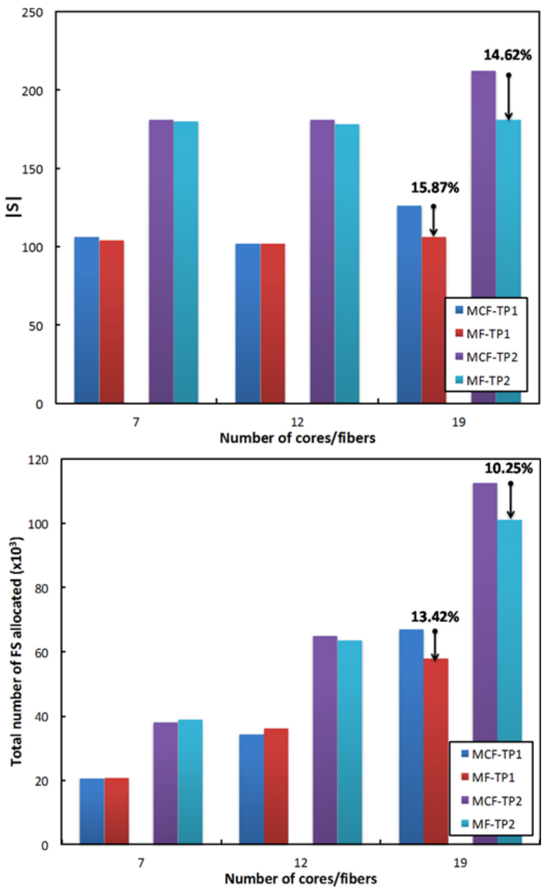

Fig. 5. Number of FSs used (top) and total number of FSs allocated (bottom) in a MCF/MF-enabled EON network with 7, 12 and 19 cores/fibers.

With 19-core MCFs, inter-core XT limitation not only applies to most lightpaths at 40 Gb/s but also at 100 Gb/s. This has a more pronounced effect, as depicted in the figure, where relative reductions of 14-15% in terms of $|\mathcal{S}|$ and 10-13% in terms of total FSs allocated are found in the MF scenario against the equivalent MCF one. Note that such differences are not only due to 40 and 100 Gb/s demands, but also due to the 400 Gb/s ones served as 4x100 Gb/s (because the end-to-end physical distance of the candidate lightpath exceeds 1385km, the longest transmission reach at 400 Gb/s).

Fig. 6 presents the same results but for the DT National backbone network. In this network, physical links are shorter than in the EON, which makes MCF-enabled scenarios to behave very closely to the equivalent MF ones, even with 19 cores/fibers, where relative differences stay below 7.5% and 9% in terms of $|\mathcal{S}|$ and total number of FSs allocated, respectively.

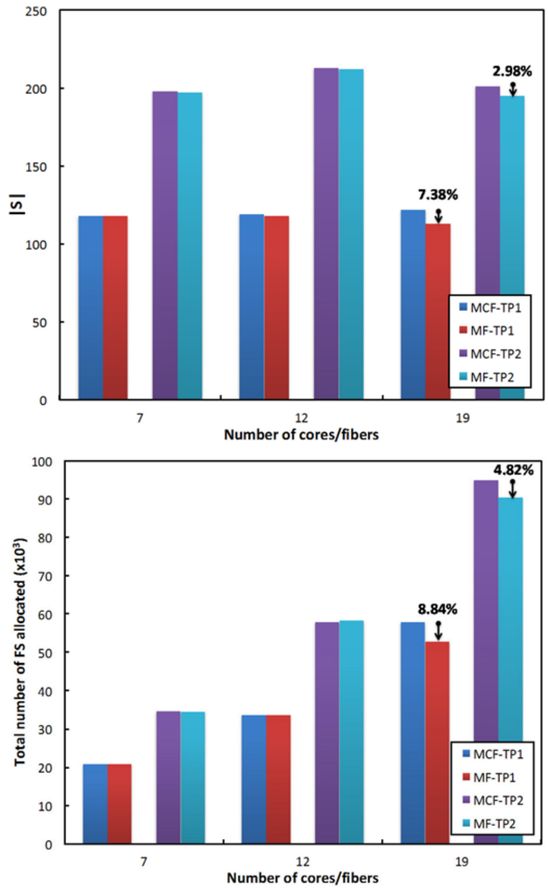

Fig. 6. Number of FSs used (top) and total number of FSs allocated (bottom) in a MCF/MF-enabled DT network with 7, 12 and 19 cores/fibers.

To better understand the reasons behind the differences observed between MCF and MF scenarios in Figs. 5 and 6, we analyze the bit-rate and modulation format employed by the operational BV-TXPs in the network. Specifically, we focus on the EON and DT network scenarios with 19 cores/fibers, where most significant differences have been identified. The results are presented in Figs. 7 and 8.

As seen in Fig. 7 (top), the long physical distances that optical signals have to traverse from source-destination over the EON backbone network, together with the transmission reach limitation imposed by inter-core XT across MCFs, prevent the utilization of advanced modulation formats (particularly 64-QAM) in most cases. Conversely, QPSK is the modulation format employed by the vast majority of operational BV-TXPs at any bit-rate, while only a few of them can employ up to 16-QAM. Moving to Fig. 7 (bottom), we can see that in the MF scenario, where the transmission reach is only limited by ASE noise, advanced modulation formats start gaining momentum. For example, under TP-1, all BV-TXPs operating at 40 Gb/s can employ 64-QAM (transmission reach is 2289 km vs. the 150km in the 19-core MCF-enabled scenario). Furthermore, BV-TXPs at 100 Gb/s can make extensive use of 64-QAM and 16-QAM. As advanced modulation formats significantly increase the efficiency of the lightpaths to be allocated over the network, this justifies the differentiated behavior of MCF and MF EON scenarios with 19 cores/fibers previously observed in Fig. 5. We shall mention that BPSK is never employed by the BV-TXPs, which is expectable as no transmission reach gain is obtained against QPSK at any bit-rate.

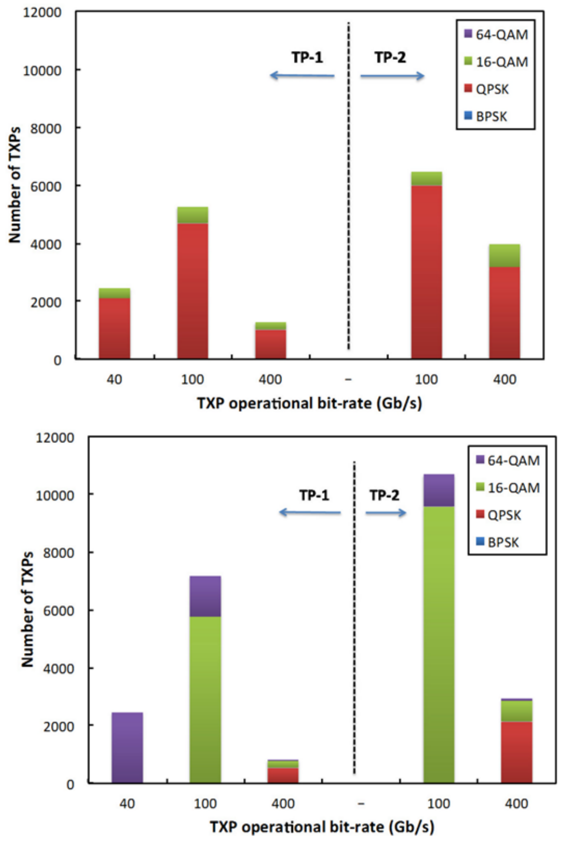

Fig. 7. Number of TXPs used at 40, 100 and 400 Gb/s in the EON with 19-core MCFs (top) and 19 fibers per MF link (bottom). Scenarios with demands following TP-1 and TP-2 are shown.

Fig. 8 depicts the same results, but in the National DT network, also with 19 cores/fibers. Since physical distances are shorter in this case, BV-TXPs can extensively employ 16-QAM in the MCF-enabled scenario, and even 64-QAM in some cases. As a result, resource efficiency can approach that of the equivalent MF scenario, as previously shown in Fig. 6. It is interesting to highlight here the lower number of 100 Gb/s BV-TXPs used compared to the EON network. In both networks, 8000 demands are offered following the same traffic profiles (TP-1 or TP-2). We have found, however, that no demand at 400 Gb/s must be served as 4x100 Gb/s due to unfeasible transmission reach in the DT network. This does not happen in the EON. For instance, in the EON with MCFs, around 300 and 800 demands at 400 Gb/s are eventually supported over 4x100 Gb/s lightpaths under TP-1 and TP-2, respectively. Therefore, around 1200 and 3200 additional BV-TXPs at 100 Gb/s are needed in each case (instead of 300 and 800 at 400 Gb/s). This issue is even more significant in the MF scenario, as SA-RMCSA sometimes decides to support a demand over physically longer 4x100 Gb/s lightpaths instead of a single 400 Gb/s one with aims to reduce $|\mathcal{S}|$, although requiring more BV-TXPs. This could be avoided by removing the longer paths that would require 4x100 Gb/s lightpaths to traverse them from $\mathcal{P}_d$ of any demand d ∈ $\mathcal{D}$ at 400 Gb/s, provided that any alternative shorter path exists (otherwise the demand would be directly blocked). We have not applied this, as our optimization goal in this paper has been to minimize $|\mathcal{S}|$. However, we discuss this effect if BV-TXP minimization comes into play in any related future work.

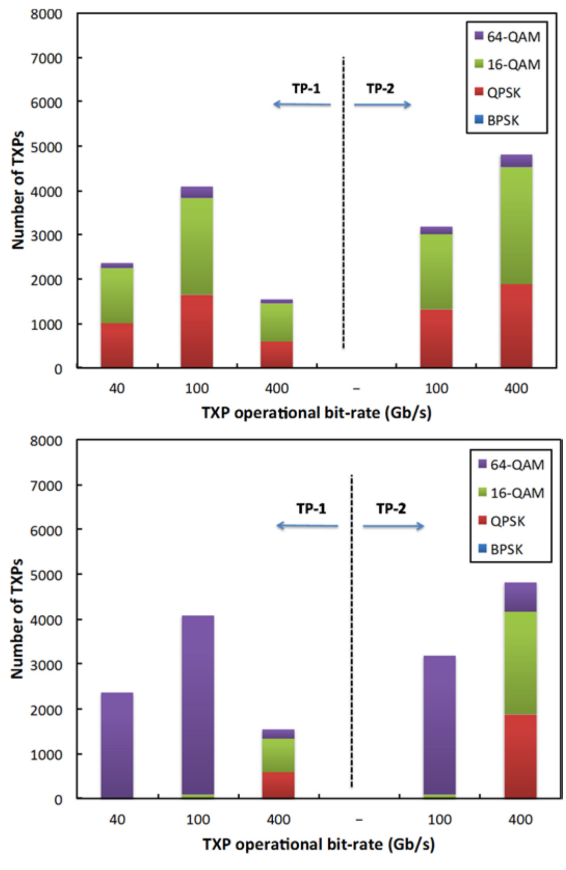

Fig. 8. Number of TXPs used at 40, 100 and 400 Gb/s in the DT with 19-core MCFs (top) and 19 fibers per MF link (bottom). Scenarios with demands following TP-1 and TP-2 are shown.

Lastly, we launch some additional executions to analyze the resulting Flex-Grid/SDM network design when we allow lightpaths to traverse longer physical paths from source to destination. Recall that in all the executions until this point we have considered K=3 physically shortest paths (in km) when generating $\mathcal{L}_d$ for every demand $d \in \mathcal{D}$, as mentioned at the beginning of section V.A. In Fig. 9 we depict the number of FSs used (top) and total number of FSs allocated (bottom) in the EON network with 19 cores/fibers, that is, the scenario most affected by inter-core XT among all the previously evaluated ones.

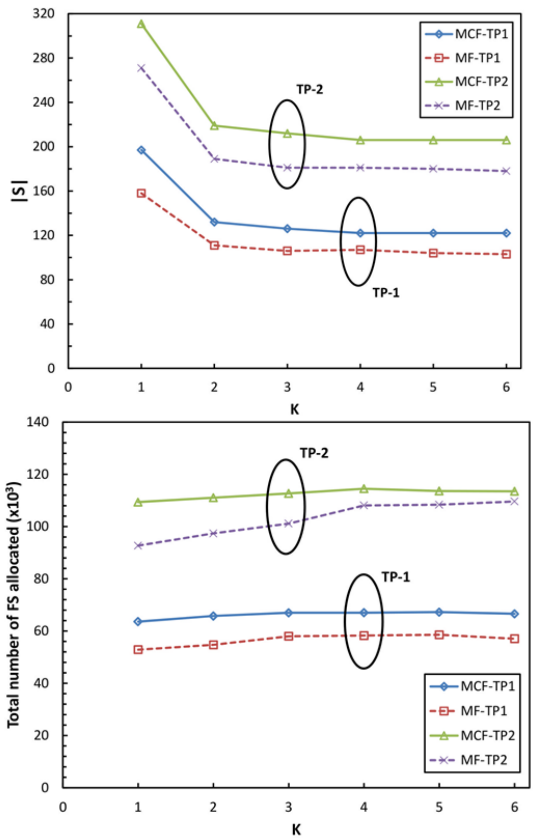

Fig. 9. Number of FSs used (top) and total number of FSs allocated (bottom) in the MCF/MF-enabled EON network with 19 cores/fibers for different values of $K \in \{1,..,6\}$.

Looking at Fig. 9 (top), we observe that a large $|\mathcal{S}|$ reduction can be achieved when allowing candidate lightpaths to traverse alternative physical paths than only the physically shortest path from the source to destination nodes of the demand (K=1), which leads to high congestion in the links in the central part of the network (i.e., a lot of FSs used) while underutilizing those in its borders. Nevertheless, such a reduction stops for K values larger to 2 or 3. This happens because longer physical paths typically traverse more hops, thus requiring the allocation of more FSs to carry a lightpath over them. Furthermore, longer distances also cause transmission reach issues, forcing the utilization of less efficient modulation formats. This outcome qualifies our K=3 assumption to obtain all previous results in the paper, as increasing the K value only translates into increased execution times of the SA-RMCSA heuristic. As a final observation, note in the figure that relative differences between the respective MCF and MF scenarios remain quite constant along the evaluated values of K, thus those previously identified in Fig. 5 (top) also apply here for $K \in \{1,..,6\}$.

Finally, in Fig. 9 (bottom) we find that, in contrast to the observed $|\mathcal{S}|$ behavior, the total number of FSs allocated increases along with K, especially in the MF scenario with TP-2. This increasing pace can be justified by the larger number of hops of the physically longer paths available with higher values of K. This effect is more pronounced in the MF scenarios, as their longer transmission reach allows SA-RMCSA using such longer physical paths if this can lead to a reduction of $|\mathcal{S}|$ (i.e., SA-RMCSA prioritizes reducing $|\mathcal{S}|$ at expenses of increasing the total number of FSs allocated in the network). In the MF scenario under TP-2, a substantial number of 400 Gb/s demands are served over 4x100 Gb/s lightpaths, particularly when they have to traverse the longer physical paths, which increases even more the total number of FSs allocated in the network, approaching those allocated in the respective MCF scenario for K=4, 5, 6.

VI. Conclusions and Future Work

In this work, we have addressed the design of MCF-enabled Flex-Grid/SDM backbone networks. Once introduced the subject under study and related work in the literature, we have proposed a methodology to compute the worst-case transmission reach value over MCFs accounting for ASE noise and inter-core XT, taking the real inter-core XT laboratory measurements across some state-of-the-art MCFs available to date. Next, we propose effective optimal (ILP-RMCSA) as well as heuristic (SA-RMCSA) approaches for the design of MCF-enabled Flex-Grid/SDM backbone networks, making use of the worst-case transmission reach estimations. From the obtained results in two reference backbone network scenarios, we are in a position to advocate for up to 19-core MCF-enabled solutions in moderately large (National) backbone networks, as resource efficiency very close to todays' available multi-fiber link solutions is obtained, while taking benefit from cost-effective integrated system components envisioned for MCF-enabled networks, like TXPs, amplifiers, ROADMs, etc. In long-haul continental backbone networks, as inter-core XT effects are more pronounced, the maximum number of cores should, in contrast, be reduced to 12 to achieve similar performance to that of multi-fiber solutions. Note, however, that if more restricted values for the noise-limited transmission reach would have been assumed, even closer performance of MCF networks against the equivalent multi-fiber networks would have been achieved.

Future work can follow up in several research directions. For example, both ILP-RMCSA and SA-RMCSA could be extended to contemplate a translucent network scenario with sparse 3R regeneration, which would enable the use of advanced modulation formats even in long-haul lightpaths (i.e., eliminating the need for the 4x100 Gbps solution assumed in this work). Moreover, an analysis of the impact of alternative offered traffic patterns on the comparison between MCF and MF networks could also be of interest.

Acknowledgements

This work has been partially funded by the Spanish National Project SUNSET (TEC2014-59583-C2-1-R), which receives funding from FEDER.

References

[1] R.-J. Essiambre et al., "Capacity Limits of Optical Fiber Networks", *IEEE/OSA J. Lightwave Technol.*, vol. 28, no. 4, pp. 662-701, Feb. 2010.

[2] P. J. Winzer, “Spatial Multiplexing in Fiber Optics: The 10x Scaling of Metro/Core Capacities”, *Bell Labs Tech. Journal*, vol.19, pp. 22-30, Sept. 2014.

[3] Howard R. Stuart, “Dispersive Multiplexing in Multimode Optical Fiber”, Science 289 (5477), 281, July 2000.

[4] N.K. Fontaine, et al, "30×30 MIMO transmission over 15 spatial modes," in *Proc. OFC 2015* , March 2015.

[5] J. Sakaguchi et al., “Space Division Multiplexed Transmission of 109-Tb/s Data Signals Using Homogeneous Seven-Core Fiber”, *IEEE/OSA J. Lightwave Technol.*, Vol. 30, no. 4, p. 658-665, Jan. 2012.

[6] Sano et al., “409-Tb/s + 409-Tb/s crosstalk suppressed bidirectional MCF transmission over 450 km using propagation-direction interleaving”, *Opt. Express*, vol. 21, pp. 16777-16683, Jul. 2013.

[7] J. Sakaguchi et al., “19-core MCF transmission system using EDFA with shared core pumping coupled via free-space optics”, *Opt. Express*, vol. 22, pp. 90-95, Jan. 2014.

[8] B. J. Puttnam et al., “2.15 Pb/s Transmission Using a 22 Core Homogeneous Single-Mode Multi-Core Fiber and Wideband Optical Comb”, *Proc. ECOC 2015*, PDP.3.1.

[9] D. Soma, et al., “2.05 Peta-bit/s Super-Nyquist-WDM SDM Transmission Using 9.8-km 6-mode 19-core Fiber in Full C band”, *Proc. ECOC 2015*, PDP.3.2.

[10] G.M. Saridis, D. Alexandropoulos, G. Zervas, D.Simeonidou, “Survey and Evaluation of Space Division Multiplexing: From Technologies to Optical Networks”, *IEEE Commun. Surveys Tuts.*, vol. 17, no. 4, pp. 2136-2156, 2015.

[11] D. Klonidis, et al., “Spectrally and Spatially Flexible Optical Network Planning and Operations”, *IEEE Commun. Mag.*, vol. 53, no. 2, pp. 69-78, Feb. 2015.

[12] M. Koshiba, K. Saitoh, K. Takenaga, S. Matsuo, “Analytical Expression of Average Power-Coupling Coefficients for Estimating Intercore Crosstalk in Multicore Fibers”, *IEEE Photonics Journal*, vol.4, no.5, pp.1987,1995, Oct. 2012.

[13] F. Ye, J. Tu, K. Saitoh, T. Morioka, “Simple analytical expression for crosstalk estimation in homogeneous trench-assisted multi-core fibers”, *Opt. Express*, vol. 22, no. 19, pp. 23007-23018, 2014.

[14] “Advances in tb/s superchannels”, in Optical Fiber Telecommunication, 6th ed., I. P. Kaminow, T. Li, and A. E. Willner, Eds. Boston, MA, USA: Academic, 2013, ch. 3, pp. 83–119.

[15] P. J. Winzer, “Making spatial multiplexing a reality”, *Nature Photon.*, vol. 8, no. 5, pp. 345–348, Aug. 2014.

[16] K. Christodoulopoulos, I. Tomkos, and E. Varvarigos, “Elastic bandwidth allocation in flexible OFDM based optical networks”, *IEEE/OSA J. Lightwave Technol.*, vol. 29, no. 9, pp. 1354–1366, 2011.

[17] M. Klinkowski, K. Walkowiak, “Routing and spectrum assignment in spectrum sliced elastic optical path network”, *IEEE Commun. Lett.*, vol. 15, no. 8, pp. 884-886, Aug. 2011.

[18] A. Muhammad et al., “Routing, Spectrum and Core Allocation in Flexgrid SDM Networks with Multi-core Fibers”, *Proc. ONDM 2014*, Stockholm, Sweden, May 2014.

[19] A. Muhammad, G. Zervas, R. Forchheimer, “Resource Allocation for Space-Division Multiplexing: Optical White Box Versus Optical Black Box Networking”, *IEEE/OSA J. Lightwave Technol.*, vol.33, no.23, pp.4928-4941, Dec. 2015.

[20] N. Amaya, et al., “Fully-elastic multi-granular network with space/frequency/time switching using multi-core fibres and programmable optical nodes”, *Opt. Express*, vol. 21, pp. 8865-8872, 2013.

[21] F. Fidler, P. J. Winzer, M. K. Thottan, K. Bergman, "Impairment-aware optical networking using cross-layer communication”, *IEEE/OSA Journal of Opt. Commun. Netw.*, vol. 5, no. 2, pp. 144-158, Feb. 2013.

[22] T. Mizuochi, et al., “Evolution and status of forward error correction”, *Proc. Optical Fiber Communication Conference and Exposition (OFC/NFOEC), 2012 and the National Fiber Optic Engineers Conference*, pp.1-3, 4-8 March 2012.

[23] El-Ghazali Talbi, “Metaheuristics: From Design to Implementation”, Wiley Ed., 2009.

[24] J. Perelló et al., “Assessment of Flex-Grid/SDM Backbone Networks under Inter-Core XT-limited Transmission Reach”, *Proc. PS 2015*, Pisa, Italy, Sept. 2015.

[25] ITU-T recommendation G.694.1, “Spectral grids for WDM applications: DWDM frequency grid”, Feb. 2012.

[26] O. Gerstel, et al. “Elastic Optical Networking: A New Dawn for the Optical Layer?”, *IEEE Commun. Mag.*, vol. 50, no. 5, pp. s12-s20, Feb. 2012.

[27] A. Peters, et al. “Measuring Flexibility and Design Trade-offs of N x M SSS-based ROADMs and BVTs”, *Proc. OFC 2015*, W3J.5, Mar. 2015.

[28] 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 Opt. Commun. Netw.*, vol. 6, no. 2, pp. 114-126, Feb. 2014.