UNIVERSITY OF TRENTO

DOCTORAL THESIS

Resource Allocation and Modeling in Spectrally and Spatially Flexible Optical Transport Networks

Author:
Federico Pederzolli

Advisor:
Dr. Domenico Siracusa

Department of Information Engineering and Computer Science
School of Information and Communication Technology

2018 - 30° Cycle

Abstract

The world's hunger for connectivity appears to be endlessly growing, yet the capacity of the networks that underpin that connectivity is anything but endless. This thesis explores both short and long term solutions for increasing the capacity of the largest and most capacious of these networks, the backbones upon which the Internet is built: optical transport networks.

In the short term, Flexi-grid technology has emerged as the evolution of fixed-grid WDM optical networks, providing higher potential throughput but suffering from an aggravated form of the spectrum fragmentation problem that affects fixed-grid networks. A novel path-based metric to better evaluate the fragmentation of spectral resources in flexi-grid networks is presented, which considers both the fact that free spectrum slices may not be available on all the links of a path, and the likelihood that an end-to-end spectral void is usable to route incoming connections, and tested by means of simulations, finding that it outperforms existing ones from literature.

For the longer term, Space Division Multiplexing (SDM) is a promising solution to overcome the looming fiber capacity crunch, and, perhaps more importantly, can offer a beneficial ratio between the expected capacity gains and the resulting increase in the cost of the network thanks to Joint and Fractional Joint Switching architectures and integrated transceivers and amplifiers. A model for such network is presented, and multiple heuristics for solving the Routing, Space and Spectrum Allocation problem are described, studied via simulations and iteratively improved, with the objective of quantifying the likely performance of several SDM architectures under multiple traffic scenarios. In addition, possible improvements to joint switching architectures, and an experimental SDN control plane for SDM networks, are presented and characterized, again by means of simulations. SDM is shown to be an attractive technology for increasing future transport networks capacity, at a reasonable cost.

Acknowledgements

Deep thanks go to my family, for supporting my studies for many, many years, and to my supervisor and all my colleagues and co-authors, without whose knowledge and willingness to share it none of this work would have been possible.

Activity Report

Research/Study Activities

The activities performed during the doctoral programme included the following:

- Three years of research in several topics connected with optical networks with the Future Networks group at FBK CREATE-NET (formerly CREATE-NET), within EU (see below) and Industry backed research projects.
- Participation in the research activities of the EU FP7 ICT INSPACE project (grant agreement n. 619732);
- Participation in the submission and research activities of the EU H2020 ACINO project (grant agreement n. 645127);
- FBK Mobility period (01/03/2017 - 14/04/2017) at Aston University (Birmingham, UK), with the initial objective of working on the applications of the Probabilistic Shaping of optical modulation symbols.
- FBK Mobility period (07/06/2017 - 25/09/2017) at KTH Royal Institute of Technology (Stockholm, Sweden), to work on optical network security topics.
- Participation in the writing of multiple project proposals, both EU and industrial.

List of Publications

During the doctoral programme's time frame the following peer-reviewed conference and journal papers were published:

1. D. Siracusa, **F. Pederzolli**, D. Klonidis, V. Lopez and E. Salvadori, *"Resource allocation policies in SDM optical networks (Invited paper),"* 2015 International Conference on Optical Network Design and Modeling (ONDM), Pisa, 2015, pp. 168-173.
2. D. Siracusa, **F. Pederzolli**, P. S. Khodashenas, J. M. Rivas-Moscoso, D. Klonidis, E. Salvadori, I. Tomkos, *"Spectral vs. spatial super-channel allocation in SDM networks under independent and joint switching paradigms,"* 2015 European Conference on Optical Communication (ECOC), Valencia, 2015, pp. 1-3.
3. V. Lopez, J. M. Gran, J. P. Fernandez-Palacios, D. Siracusa, **F. Pederzolli**, O. Gerstel, Y. Shikhmanter, J. Mårtensson, P. Sköldström, T. Szyrkowiec, M. Chamania, A. Autenrieth, I. Tomkos, D. Klonidis, *"The role of SDN in application centric IP and optical networks,"* 2016 European Conference on Networks and Communications (EuCNC), Athens, 2016, pp. 138-142.

4. **F. Pederzolli**, D. Siracusa, J. M. Rivas-Moscoso, B. Shariati, E. Salvadori and I. Tomkos, *"Spatial group sharing for SDM optical networks with Joint Switching,"* 2016 International Conference on Optical Network Design and Modeling (ONDM), Cartagena, 2016, pp. 1-6.

5. **F. Pederzolli** et al., *"SDN application-centric orchestration for multi-layer transport networks,"* 2016 18th International Conference on Transparent Optical Networks (ICTON), Trento, 2016, pp. 1-4.

6. P. S. Khodashenas, J. M. Rivas-Moscoso, D. Siracusa, **F. Pederzolli**, B. Shariati, D. Klonidis, E. Salvadori, I. Tomkos, *"Comparison of Spectral and Spatial Super-Channel Allocation Schemes for SDM Networks,"* in Journal of Lightwave Technology, vol. 34, no. 11, pp. 2710-2716, June1, 1 2016.

7. **F. Pederzolli**, M. Gerola, A. Zanardi, X. Forns, J. F. Ferran and D. Siracusa, *"Experimental Evaluation of YAMATO, a SDN Control Plane for Joint and Fractional-Joint Switched SDM Optical Networks,"* ECOC 2016; 42nd European Conference on Optical Communication, Dusseldorf, Germany, 2016, pp. 1-3.

8. B. Shariati, D. Klonidis, D. Siracusa, **F. Pederzolli**, J. M. Rivas-Moscoso, L. Velasco, I. Tomkos, *"Impact of Traffic Profile on the Performance of Spatial Superchannel Switching in SDM Networks,"* ECOC 2016; 42nd European Conference on Optical Communication, Dusseldorf, Germany, 2016, pp. 1-3.

9. D. Siracusa, **F. Pederzolli**, M. Gerola, A. Zanardi, D. La Fauci and G. M. Galimberti, *"Demonstration of a Hybrid SDN/GMPLS Control Plane for Optical Virtual Private Networks with Restoration Capabilities,"* ECOC 2016; 42nd European Conference on Optical Communication, Dusseldorf, Germany, 2016, pp. 1-3.

10. M. Savi, **F. Pederzolli** and D. Siracusa, *"An application-aware multi-layer service provisioning algorithm based on auxiliary graphs,"* 2017 Optical Fiber Communications Conference and Exhibition (OFC), Los Angeles, CA, 2017, pp. 1-3.

11. **F. Pederzolli**, D. Siracusa, E. Salvadori and R. Lo Cigno, *"Energy Saving Through Traffic Profiling in Self-Optimizing Optical Networks,"* in IEEE Systems Journal, vol. 11, no. 2, pp. 752-761, June 2017.

12. **F. Pederzolli**, D. Siracusa, B. Shariati, J. M. Rivas-Moscoso, E. Salvadori and I. Tomkos, *"Improving performance of spatially joint-switched space division multiplexing optical networks via spatial group sharing,"* in IEEE/OSA Journal of Optical Communications and Networking, vol. 9, no. 3, pp. B1-B11, March 2017.

13. **F. Pederzolli**, M. Gerola, A. Zanardi, X. Forns, J. F. Ferran and D. Siracusa, *"YAMATO: The First SDN Control Plane for Independent, Joint, and Fractional-Joint Switched SDM Optical Networks,"* in Journal of Lightwave Technology, vol. 35, no. 8, pp. 1335-1341, April15, 15 2017.

Material from papers 1, 2, 4, 6, 7, 8, 12, and 13 is described in this thesis.

In addition to those above, the following papers were published before the formal start of the doctoral programme:

14. D. Siracusa, **F. Pederzolli**, E. Salvadori, R. L. Cigno and I. T. Monroy, *"Proactive restoration of slow-failures in optical networks,"* 2014 16th International Conference on Transparent Optical Networks (ICTON), Graz, 2014, pp. 1-4.

15. D. Siracusa, **F. Pederzolli**, R. Lo Cigno and E. Salvadori, *"Energy saving through traffic profiling and prediction in self-optimizing optical networks,"* OFC 2014, San Francisco, CA, 2014, pp. 1-3.

Furthermore, the following papers, published after the end of the third year of the PhD period, have been presented at the 2018 Optical Fiber Communications Conference and Exhibition:

16. **F. Pederzolli**, M. Furdek, D. Siracusa, L. Wosinska, *"Towards Secure Optical Networks: A Framework to Aid Localization of Harmful Connections,"* in Optical Fiber Communication Conference, OSA Technical Digest (online) (Optical Society of America, 2018), paper Th2A.42.

17. **F. Pederzolli**, M. Chamania, M. Santuari, T. Szyrkowiec, C. Matrakidis, C. Rozic, D. Klonidis, V. Lopez, D. Siracusa, *"Disaggregating Optical Nodes in a Multi-Layer SDN Orchestrator for the Integration of an In-Operation Planning Tool,"* in Optical Fiber Communication Conference, OSA Technical Digest (online) (Optical Society of America, 2018), paper Th2A.35.

Lastly, the following paper, giving an overview of the overall results of the INSPACE European project, was accepted for publication at the 2018 Advanced Photonics Congress:

18. S. Ben-Ezra, M. Gerola, D. Siracusa, **F. Pederzolli**, D. Marom, M. Blau, J. Macdonald, N. Psaila, C. Sanchez-Costa, A. Ellis, X. Forns, J. Ferran, F. Jimenez, N. Christodoulia, B. Shariati, D. Klonidis, I. Tomkos, *"First WDM-SDM Optical Network with Spatial Sub-Group Routing ROADM Nodes Supporting Spatial Lane Changes,"* APC 2018.

Contents

List of Figures

List of Tables

List of Algorithms

List of Abbreviations

ABP	**A**ccess **B**locking **P**robability
AoD	**A**rchitecture **o**n **D**emand
AWG	**A**rray **W**aveguide **G**rating
BER	**B**it **E**rror **R**ate
BP	**B**locking **P**robability
CDF	**C**umulative **D**istribution **F**unction
CM	**C**onnection **M**anager
DaaS	**D**atabase **a**s **a** **S**ervice
DGD	**D**ifferential **G**roup **D**elay
DP	**D**ual **P**olarization
DSP	**D**igital **S**ignal **P**rocessing
DSpaF	**D**egenerate-**Spa**ce-**F**irst
EDFA	**E**rbium-**D**oped **F**iber **A**mplifier
FDM	**F**requency **D**ivision **M**ultiplexing
FEC	**F**orward **E**rror **C**orrection
FF	**F**irst-**F**it
FJoS	**F**ractional **Jo**int **S**witching
FMF	**F**ew-**M**ode **F**iber
FR	**F**ragmentation **R**atio
FRR	**F**ast **R**e-**R**oute
GB	**G**uard-**B**and
GN	**G**aussian **N**oise
IEEE	**I**nstitute of **E**lectrical and **E**lectronic **E**ngineers
IETF	**I**nternet **E**ngineering **T**ask **F**orce
ILP	**I**nteger **L**inear **P**rogramming
InS	**In**dependent **S**witching
ITU	**I**nternational **T**elecommunications **U**nion
JoS	**Jo**int **S**witching
KSP	**K**-**S**hortest **P**ath
LAN	**L**ocal **A**rea **N**etwork
MAN	**M**etropolitan **A**rea **N**etwork
MCF	**M**ulti-**C**ore **F**iber
MIMO	**M**ultiple **I**nput **M**ultiple **O**utput
NAM	**N**etwork **A**straction **M**odule
NCM	**N**orth-bound **C**ommunications **M**anager
OA	**O**ptical **A**gent
OBS	**O**ptical **B**urst **S**witching
OCS	**O**ptical **C**ircuit **S**witching
OEO	**O**pto-**E**lectro-**O**ptical
OLA	**O**ptical **L**ine **A**mplifier
OLT	**O**ptical **L**ine **T**erminal
ONU	**O**ptical **N**etwork **U**nit
OPS	**O**ptical **P**acket **S**witching

OSA	**O**ptical **S**ociety of **A**merica
OSNR	**O**ptical **S**ignal to **N**oise **R**atio
OSPF-TE	**O**pen **S**hortest **P**ath **F**irst-**T**raffic **E**ngineering
OTN	**O**ptical **T**ransport **N**etwork
OXC	**O**ptical **C**ross-**C**onnect
PCE	**P**ath **C**omputation **E**lement
PLI	**P**hysical **L**ayer **I**mpairments
PON	**P**assive **O**ptical **N**etwork
PRE	**P**CE/**R**SSA **E**ngine
PSK	**P**hase-**S**hift **K**eying
QAM	**Q**uadrature **A**mplitude **M**odulation
RMLSA	**R**outing, **M**odulation **L**evel and **S**pectrum **A**llocation
RMLSSA	**R**outing, **M**odulation **L**evel, **S**pace and **S**pectrum **A**llocation
ROADM	**R**econfigurable **O**ptical **A**dd/**D**rop **M**ultiplexer
RSA	**R**outing and **S**pectrum **A**llocation
RSSA	**R**outing, **S**pace and **S**pectrum **A**llocation
RSVP-TE	**R**esource **R**eser**v**ation **P**rotocol-**T**raffic **E**ngineering
RWA	**R**outing and **W**avelength **A**ssignment
SDM	**S**pace **D**ivision **M**ultiplexing
SDN	**S**oftware-**D**efined **N**etworking
SMF	**S**ingle **M**ode **F**iber
SP	**S**hortest **P**ath
SPM	**S**outh-bound **P**rotocol **M**anager
SpaF	**Spa**ce-**F**irst
SpeF	**Spe**ctrum-**F**irst
SRLG	**S**hared **R**isk **L**ink **G**roup
SSS	**S**pectrum **S**elective **S**witch
TDM	**T**ime **D**ivision **M**ultiplexing
TM	**T**ED **M**anager
WAN	**W**ide **A**rea **N**etwork
WDM	**W**avelength **D**ivision **M**ultiplexing
WSS	**W**avelength **S**elective **S**witch

Dedicated to my late maternal grandfather, whose brilliant example and unshakable optimism have been my anchor and reference in difficult times.

Dedicato al mio defunto nonno materno, il cui brillante esempio ed irremovibile ottimismo sono stati la mia ancora e riferimento nei momenti difficili.

Chapter 1

Optical Networking Primer

This short chapter presents, in the interest of accessibility, some basic information and terminology pertaining optical network components, operations and relationship with other network layers and technologies. Readers familiar with optical networks and related concepts can safely skip to Chapter 2, which outlines the motivation, research problems and contributions to the State of the Art of this thesis, as well as the methodology used to verify them.

1.1 Overview of Optical Network Components

At a very high level, an optical network is a network in which information is exchanged encoded onto streams of photons, often (but not always) guided by means of optical fibers.

Due to timing constraints in the configuration and tuning of commercial lasers, most modern optical networks are circuit-switched (Optical Circuit Switching, OCS); an optical circuit is known as a **lightpath**.

All optical networks are based on one to four fundamental components:

- **Optical Transceivers**, which convert signals between the electronic and optical domains;
- **Optical Fibers**, which provide the low-loss pipe-work through which optical signals are transported;
- **Optical Switches**, which enable the switching (and routing) of signals directly in the optical domain;
- **Optical Amplifiers**, which enhance the distance that can be covered by optical signals.

While a pair of transceivers may suffice for short point-to-point links over free space, the other components can allow optical signals to form networks that efficiently traverse nations and cross oceans.

1.1.1 Optical Transceivers

Optical Transceivers are, at the simplest level, devices that convert between electrical and optical signals. Modern examples use high quality tunable coherent lasers, which exhibit excellent stability in signal frequency and power distribution, and can be configured to transmit and receive at one of a range of different frequencies. While early transceivers relied on simple On-Off Keying to encode data, modern

hardware employs advanced digital modulation formats, such as variants of Phase-Shift Keying (PSK) and Quadrature Amplitude Modulation (QAM), to encode several bits over complex symbols and achieve, for state of the art device, throughputs in the order of 400 Gbps [1]. Furthermore, modern transceivers transparently include **Forward Error Correction** (FEC) codes in the transported data, which, while imposing an overhead on the usable throughput, in combination with the low Bit Error Rate (BER) induced by optical fibers ensure virtually error-free transmission.

A couple of precisely aligned (via specialized lenses) transceivers are sufficient to create a point-to-point optical link over free space, a technique known as **free-space optics**, which is a mature but still active field of research. Such technology presents several drawbacks compared to more standard fiber-based deployments, chiefly a strong dependence on weather conditions: fog or heavy rain all but prevent communication. It can, however, be useful for short-reach communications where digging trenches to lay fibers would be impossible or prohibitively expensive, or for very specific low-latency applications. Examples of the former include crossing mountainous areas or rivers in the absences of adequate roads and bridges, while the latter include privately built networks for high frequency trading, where the small delay added by the twists and bends of traditional fiber deployments would be a competitive disadvantage.

1.1.2 Optical Fibers

Optical Fibers are thin transparent cables made of plastic or glass, composed of an external **cladding** protecting an inner **core**, which "steers" light waves from one of end to the other. They exhibit extremely low **Physical Layer Impairments**, resulting in exceptionally low levels of signal distortion, at least in certain regions of the frequency domain.

Typically, a modern Single Mode Fiber (SMF) for core transport networks works well in the so called C band, between 1530 and 1565 nm, although fibers which also support the L band, from 1565 to 1625 nm, are starting to become more common. While fibers themselves are relatively cheap, deploying them tends to be extremely expensive, since they tend to be placed in trenches running underground along roads, or under the sea, both of which require specialized equipment and personnel.

Cheaper Multi-Mode Fibers (MMFs), with larger cores allowing the co-propagation of several modes of light, are instead used for shorter range, lower bandwifth applications.

Since fibers support a relatively large transmission spectrum, much larger than most optical signals require, a single fiber can be used to carry multiple independent optical signals, typically via a Frequency-Division Multiplexing (FDM) scheme known as **Wavelength Division Multiplexing** (WDM). The International Telecommunications Union (ITU) defined multiple WDM standards over the years, following the evolution of commercial grade transceivers. The current standard is known as Dense-WDM (D-WDM), and specifies up to 96 (40 or 80 in most installations) 50 GHz-wide optical channels in the C-band. To prevent signal overlap, and account for imperfect filtering at intermediate switches and the receiver, the channel width is larger than the spectral footprint of a WDM optical signal; the difference is called the optical **Guard-Band** (GB), as shown in Figure 1.1 (a). Time-Division Multiplexing (TDM) schemes have also been proposed and demonstrated, in the form of Optical Packet Switching/Optical Burst Switching (OPS/OBS) technologies, but

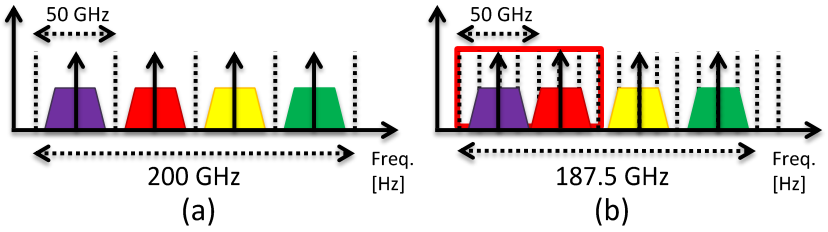

FIGURE 1.1: Schematic of Wavelength Division Multiplexing (a) and the more efficient Flexi-grid Super-Channels (b) carrying multiple independent signals in a single fiber.

have not found commercial success in the context of transport networks (more so in the access domain).

A recent advance [2] to increase the spectral efficiency of WDM networks, which is now approaching the technological maturity required for commercial deployment, is the introduction of the so called **Flexible Optical Networks**, also known as Flexi-grid (the term used throughout this thesis), Elastic or Sliceable networks. In these networks, channels are no longer mandated to be 50 GHz wide, but can be multiples of a 12.5 GHz unit (provisions are in place to reduce this number in the future, as technology improves) called, using Internet Engineering Task Force (IETF) terminology, "**frequency slots**". For lack of a better word ("slot" being already taken), throughout this thesis the 12.5 GHz unit of spectrum will be called a "**spectral slice**", for conformity with existing works on the subject. This division into small slices has two benefits: it allows the use of combinations of modulation formats and symbol rates that do not fit well in a 50 GHz channel, and it allows to place multiple independent signals nearly adjacent in the frequency domain, without placing switching guard-bands between them, when they share the same source and destination (and therefore need not be filtered and separated in intermediate switches), creating a **super-channel** (Figure 1.1 (b)). Unfortunately, this also has the drawback of complicating the resource allocation problem in flexi-grid networks, as described in Section 1.2.1.

With the use of fibers, optical transceiver can successfully communicate at distances of kilometers (~80 Km for commercial grade transceivers and fibers), making them ideal to support both Local Area Networks (LANs), within a large office or between nearby buildings when high capacity or longer direct transmission distances than what is feasible using copper are needed, as well as simple Wide Area Networks interconnecting different buildings within a single city. In these scenarios optical links are relatively short and point-to-point, realized with cheap transceivers and often multi-mode fibers (or even free-space optics), and stitched together into an actual network by means of electronic (packet/frame) switches.

By using passive devices known as **Optical Splitters/Combiners**, which split the optical power coming from an input port into two or more output ports, fibers can be organized trees. These are often used in Access Networks, generally built using Passive Optical Network (PON) technology, where a single Optical Line Terminal (OLT) in a central office sits at the root of a fiber tree, realized with passive optical splitters/combiners, connecting to several Optical Network Units (ONUs) in cabinets/homes. Both TDM and FDM based PONs exist, but a detailed description of the techniques they employ is beyond the scope of this short introductory chapter, and thesis.

1.1.3 Optical Switches

Optical Switches are systems that separate incoming signals traveling together in a single fiber, then route them towards one of multiple possible outputs, connecting to either other nodes or local Add/Drop ports (i.e., optical transceivers). Various configuration of switching nodes are possible; modern architectures allowing the dynamic reconfiguration of the routing of signals via software are known as a **Reconfigurable Optical Add/Drop Multiplexers** (ROADMs) if they have Add/Drop ports, or simply **Optical Cross-Connects** (OXCs) otherwise.

Any ROADM design operates in stages: first, it filters power around the spectra of incoming signals, separating them into discrete switchable units, using devices known as **Array Waveguide Gratings** (AWGs, for WDM) and, more recently, **Wavelength/Spectrum Selective Switches** (WSSs/SSSs, for Flexi-grid). This filtering absorbs most of the power being transmitted near the edge of the channels, hence the need for guard-bands to protect channel edges.

The second stage routes signals towards a desired output port using essentially one of two schemes:

- Broadcast & Select: an array of passive 1-to-N optical splitters generates N copies of each incoming signal (each carrying $1/N^{th}$ of the original power, minus losses in the connections of components), each of which is directed (via short fiber patch cables) to a different output port. A filter in the output ports allows only desired signals to go through.

- Route & Select: an array of, functionally, small mirrors is used to direct nearly all the power of each incoming signal towards its destination port.

Lastly, multiple signals are multiplexed into the same output port and potentially re-amplified before exiting the node. Modern Colorless, Directionless, Contentionless (CDC) designs do not pose any restriction on switching, other than preventing overlapping signals from exiting on the same port.

Larger Metropolitan Area Networks (MANs), where links can run for several kilometers, are often organized as unamplified optical rings, where each node is an optical switch with add/drop capabilities, meaning it can either "capture" and locally terminate an incoming signal, or let it through for the next node. These networks can be operated using TDM, FDM, or even both.

1.1.4 Optical Amplifiers

Optical Line Amplifiers (OLAs) are devices used to inject power into signals directly in the optical domain (at the cost of adding some noise), therefore extending their transmission reach, which, for a commercial transport-grade transceiver is in the order of 80-100 Km. They are crucial for achieving the kind of transmission distance that make international and inter-continental networks feasible.

The most widely used type of amplifier is the **Erbium-Doped Fiber Amplifier** (EDFAs), which consists of a strand of fiber doped with erbium atoms and a pump laser, which excites the erbium which, in turn, releases power at several frequencies, which adds to the power of incoming signals.

Large national, international and especially intercontinental and oceanic networks with links crossing several hundred or even thousands of kilometers would be extremely expensive to build and operate without direct optical amplification, since Opto-Electro-Optical (OEO) is a very power hungry process, and transport

grade optical transceivers are expensive. Typically, optical amplifiers are placed every 80-100 Km to reinforce optical signals, even in undersea cables (which comprise a small number of fibers and a layer of conductive copper to provide power), enabling direct optical links that cross even the whole Pacific ocean.

These large **Optical Transport Networks** aggregating and carrying traffic from all smaller networks, generally organized as some sort of mesh whose topology is largely dictated by the geographical features of the territory they cross, and whose links require amplification, are the focus of this thesis.

1.2 Overview of Optical Transport Networks Control

While different types of optical network, such as LANs, MAN rings and Access PONs all have their own control schemes and related research problems, these are outside the scope of this thesis, whose focus is on transport networks.

Most modern optical transport network deployments are controlled using the **Generalized Multi-Protocol Label Switching** (GMPLS) [3] suite of protocols, which comprises distributed protocols to flood and maintain the consistency of a **Traffic Engineering Database** (Open Shortest Path First-Traffic Engineering, OSPF-TE), and to provision optical circuits (Resource Reservation Protocol-Traffic Engineering, RSVP-TE), among others. In addition, large-scale networks using OLAs run vendor-specific distributed protocols to automatically tune them, which, it should be noted, can be a time-consuming process (in the order of several minutes for unloaded OLAs).

Many vendors are also slowly pivoting towards the deployment of (logically or physically) centralized **Software-Defined Networking** (SDN) controllers, controlling devices either via recent versions of the OpenFlow [4] protocol, or using semi-custom information models on top of standard protocols such as the venerable TL1 or more modern NETCONF [5].

1.2.1 Resource Allocation

Irrespective of the technology used, in order to establish an optical circuit (i.e., a **lightpath**), the control plane of the network must solve several problems:

1. **Routing**: that is, finding one or more paths between source and destination; this is generally done by employing variants of well-known routing algorithms, such as Dijkstra's [6] Shortest Path (SP) and Yen's [7] K-Shortest Paths (KSP) algorithms.

2. (Optionally) **Modulation Level, Symbol Rate and FEC selection**: for transceivers that support multiple formats, it may be possible to set these parameters, which ultimately affect the data rate and Optical Signal to Noise Ratio (OSNR) of a circuit.

3. (Optionally) **Feasibility evaluation**: by either pre-sharing information about Physical Layer Impairments (PLI), or collecting it during the signaling phase, the control plane can compute, using an appropriate model (such as the Gaussian Noise (GN) model [8]) whether the OSNR over a certain path will be sufficient to enable communications using the parameters selected at Step 2, in order to avoid performing several time-consuming connection attempts.

4. **Spectral resource assignment**: depending on the FDM technology being used (WDM or Flex-grid), the control plane must select which WDM channel or flexi-grid frequency slot to use for the connection.

Networks operating in this way are called **transparent**, as signals travel between their source and destination exclusively in the optical (i.e., analogue) domain.

When the longest optical-only reach is insufficient to reach the destination, or the frequency of a signal needs to change along its path (perhaps because of a very busy network), then a signal must either be terminated locally and routed through the higher layers, or pass through a signal **regenerator**, which performs direct OEO conversion. Mostly transparent networks employing one or more of these techniques are called **translucent**.

Finally, networks dropping optical signals at every hop and performing switching in the electronic domain, such as networks using an **Optical Transport Network** (OTN) layer, are called **opaque**. Note that electronic (packet) switching has both significant benefits with respect to direct optical circuit switching in terms of ability to perform grooming; however, the operating cost of converting between the electrical and optical domains and that of high performance electronic switches capable of handling the very large flows of core networks is much larger than that of all-optical switching and amplification equipment.

The combination of the first and last steps (and, optionally, the other ones) of establishing a lightpath is a well-known problem whose specifics depend on the technology being used, with different versions known by slightly different names.

Routing and Wavelength Assignment

For WDM networks, it is known as the **Routing and Wavelength Assignment** (RWA) problem, which can be summarized as follows.

A WDM optical network can be described as a directed graph, where edges are optical fibers labeled with binary arrays representing the availability of individual wavelengths (also called λs), and potentially also include a description of PLIs, both from propagation through the fiber and from traversed OLAs, and the nodes are either OXCs or ROADMs, i.e. transit optical switches or switches with one or more transceivers attached. More complex views are also possible, such as detailing intermediate amplifiers and exploding switching nodes into their internal components.

Since no intermediate device can change the frequency of an optical signal while in the optical domain (that is, without OEO conversion,) one such signal is constrained to use the same channel on all the links it traverses; this is known as the **Wavelength Continuity Constraint**.

Solving RWA therefore means finding a path from source transceiver to destination transceiver and selecting a single channel that is free on all the fibers traversed by that path. Furthermore, such solution should then be checked for feasibility, i.e. the control plane should check that the accumulated impairments on the path do not lead to an excessively low (i.e., unrecoverable) OSNR at the receiver, although this step is often neglected in the research world.

RWA is a well researched problem, studied throughout the nineties, and is known to be NP-complete [9]. Many works in literature have provided optimal or heuristic solutions to this problem, either using Integer Linear Programming (ILP) formulations or heuristics, solving the two sub-problems jointly or sequentially, and optimizing various practical sub-cases, such as adding a single connection to an already established network, jointly optimizing a set of connections to be added, computing the jointly-shortest disjoint primary and backup paths, and more [10].

Routing and Spectrum Allocation

In the context of flexible optical networks RWA becomes slightly more complex, as signals can be mapped to arbitrarily large ranges of spectral slices (frequency slots), may be grouped into super-channels, and may only partially overlap on successive links. Additionally, the choice of channel width allows greater freedom in modulation level and FEC selection, which respectively improve bit-rate at the cost of reach and vice-versa (assuming fixed baud rate and, hence, spectral width).

In this context, the resource allocation problem is therefore known as the **Routing and Spectrum Allocation** (RSA) problem, again subject to the continuity constraint but now applied to frequency slots, and therefore known as the **Spectrum Continuity Constraint**.

In addition, in the case where transceivers support multiple modulation formats, the selection of the most suitable (typically the highest performing capable of producing enough OSNR at the receiver, given the usual trade off between capacity and reach) is an additional sub-problem generally solved right after routing, making the complete problem sometimes known as **Routing, Modulation Level and Spectrum Allocation** (RMLSA).

Optical Fragmentation

A well known issue resulting from the Wavelength (Spectrum) Continuity Constraint is known as Wavelength (Spectrum) **Fragmentation**.

The problem can be visualized as follows: if, for a certain path, a certain wavelength (frequency slot) is free on one of its links but occupied by a different connection on another, then that wavelength (frequency slot) is unavailable on that particular path. As the number of connections grows, the probability of such a situation occurring increases. This leads to connection requests being refused because, despite there being enough available resources on all the links of the path, the same piece of spectrum is not. In the case of flexi-grid, the issue of spectrum fragmentation is further compounded by the fact that end-to-end sequences of available spectral slices may or may not be large enough to support a super-channel of a particular size (i.e., capacity) [11].

As a corollary, it is generally unfeasible to fully utilize all spectral resources in a generic meshed optical network without regenerators.

1.2.2 Reliability Mechanisms

Optical networks are highly reliable, despite their potential to suffer from a number of failures, including the common fiber cut (often due to careless digging during road works), component (transceiver, amplifier) failure or control software faults.

High reliability against component failure is guaranteed through two main family of techniques, one preventive and the other reactive:

- Optical **Protection** involves transmitting a signal twice on link- and possibly node-disjoint paths, using either the same (with Y cables) or a second pair of transceivers. It is expensive, but guarantees error-free transmission even if one of the paths (or even, potentially, transceivers) fails.

- Optical **Restoration** [12] involves dynamically re-routing an optical circuit upon detecting a failure in one of its links. Since this triggers OLA power equalization (irrespective of the path computation and signaling times), the process

is extremely long from the point of view of the packet network (potentially several minutes) and must be complemented by packet-layer reliability mechanisms such as IP Fast Re-Route (FRR). Variants of this technique include pre-computed (and possibly shared) backup paths vs. just-in-time backup path computation.

In addition, the reliability of commercial-grade control software is often ensured by means of general distributed system techniques, including redundant control stacks and databases in every node. Note that, since these networks are circuit switched, most control software failures prevent the setup of new lightpaths (or modification of existing ones), but do not prevent existing connections from working (unless OXC configurations are overwritten in some bizarre byzantine failure case), at least for as long as RSVP-TE soft reservations hold, thus giving a window of potentially up to a few minutes to have backup control instances take over or reboot without experiencing catastrophic service interruptions.

1.2.3 Relationship with other Networks

While different networks (e.g. metro, access) use different underlying technologies, and other approaches are available in literature (e.g. OBS/OPS), most modern transport networks have a stacked design involving two or three layers, with an IP or MPLS (packet) layer operating on top of an optical (circuit) layer, as depicted in Figure 1.2, with an optional **Optical Transport Network** (OTN) layer for coarse electronic grooming in the middle.

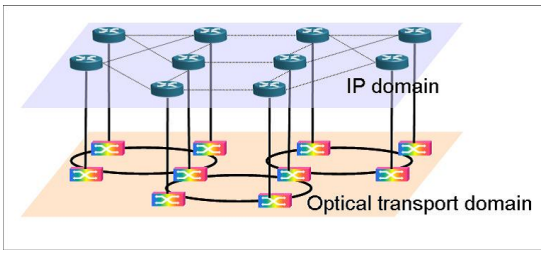

FIGURE 1.2: Schematic protocol stack for an Optical Transport Network: a packet (IP) logical topology on top of an optical physical topology (figure source: [13]).

As shown in Figure 1.2, the topology of the packet (and optional OTN) layer(s) is a logical entity (an overlay) realized by the circuit layer. In other words, an optical (or OTN) circuit appears as a direct adjacency (a link) between two routers in the packet layer.

Chapter 2

Introduction

2.1 Motivation

Optical transport networks are the backbone upon which the Internet, that fantastically successful telecommunications experiment that enables the modern hyper-connected society, is built.

These networks interconnect and carry traffic originating in all smaller networks (LANs, WANs, Access Networks, MANs), but, while they offer tremendous capacity, in the order of multiple hundred Gbps per optical connection, and potentially multiple Tbps per fiber, thanks to multiplexing technologies such as WDM and Flexi-grid, this capacity is not infinite [14]. In fact, we are getting close to exploiting the maximum theoretical information capacity of the commonly used spectrum band in optical fiber links (the so-called C-band, 1530-1565 nm), while the total network load on transport networks seems to continue its historical trend of exponential growth [15]. There are several factors leading to this growth, a detailed analysis of which is beyond the scope of this thesis. Nonetheless, example sources of this increase include improvements in access network speeds which in turn enable more bandwidth-hungry services, such as the distribution of high quality (HD, 4K) video both on-demand and as IPTV, the still increasing number of subscriptions for Internet Service Providers, the emergence of remote game rendering which requires transferring high-quality video streams produced on-the-fly (hence no caching is possible), etc.

In order to cope with this expected increase in traffic and the finite capacity of optical fibers, the introduction of Flexi-grid networks [2] has been proposed. This technology enables two important capacity gains: the use of aggressive modulation formats (which would waste spectrum in fixed-width WDM channels), and spectrally efficient super-channels (see Figure 1.1 (b) back in Chapter 1) which forgo the intermediate guard-bands between multiple co-routed optical signals. Unfortunately, Flexi-grid networks suffer from a more acute version of the well-known WDM channel fragmentation problem, known as the **Spectrum Fragmentation** problem [11], which, if not addressed, largely nullifies the efficiency gains of Flexi-grid in busy meshed networks. Minimizing its impact is the focus of Chapter 3.

This is, however, only a short-term solution, with a relatively limited impact on the capacity of future networks. For the longer term, several approaches to improve optical network capacity are possible and are being investigated:

1. extending the usable frequency band within fibers [16], [17] by changing the amplification technology from current EDFAs to e.g. Raman amplifiers [18], [19];

2. transmitting several non-interfering optical beams at the same frequency using Optical Angular Momentum [20]–[22];

3. transmitting several interfering beams at the same frequency at vastly different power levels so that they can be reconstructed at the receiver by treating anything but the strongest signal as noise, then recursively subtracting its reconstructed original waveform from the received one [23];

4. (reverse) multiplexing signals on multiple parallel fibers, fiber cores, or modes of light, collectively called the "space" or "spatial" dimension, orthogonal to the spectral domain [24], [25].

The last point is, in its simplest, most basic level, the idea behind optical **Space Division Multiplexing** (SDM), which will be further detailed in Chapter 4.

Given that so many alternatives exists to solve the same incumbent problem, it is only natural to ask what makes SDM worth studying. The answer to this is threefold. Firstly, it should be noted that while at first glance all these technologies may appear to be competing to solve the same problem, they are ultimately complementary, and could one day be combined to achieve even greater overall network capacity. Furthermore, the actual adoption of any of these technologies is conditional on it offering a beneficial ratio between the resulting increase in capacity and deployment costs, an area where SDM, as better explained in Section 4.1, offers a range of solutions catering to different investment levels, as well as migration plans to slowly ramp up capacity as needed. Lastly, from the perspective of the control plane, the routing and switching constraints of each of these technologies are a subset of those identified for SDM networks, as detailed in 4.1.3. Therefore, models and algorithms developed for the control of SDM networks can easily be re-purposed for use with these other technologies.

2.2 Contributions and Structure of the Thesis

At a high level, the goal of this thesis is to **contribute to the realization of the needed increase in the performance of future transport networks, by means of improved allocation resource algorithms**, that is, by mostly proposing tweaks to the control plane of both existing and already proposed data plane technologies rather than proposing new ones.

More concretely, the work focuses on two aspects: *i*) dealing with the issue of spectrum fragmentation via appropriate resource allocation algorithms, and *ii*) devising algorithms to exploit the emerging optical SDM equipment, i.e., on efficiently and efficaciously solving the resource allocation problem in the context of future SDM optical transport networks. The former will maximize the benefits of the upcoming flexi-grid technology, while the latter will be a necessary software tool to turn the disparate components outlined in Section 4.1 into an usable telecommunications network.

The contributions of this thesis to the State of the Art thus cover three main areas:

1. Reducing Spectrum Fragmentation.

2. Modeling SDM Networks.

3. Allocating Resources in SDM Networks.

2.2.1 Reducing Spectrum Fragmentation

There is no consensus in the scientific literature on how to measure spectrum fragmentation, let alone minimize it. Chapter 3 describes a number of existing fragmentation metrics from the available literature, highlights some of their shortcomings, and proposes a novel metric to quantify fragmentation, along with two simple RSA heuristics to minimize it according to some metric and simulative comparisons of the proposed metric against others available in literature.

2.2.2 Modeling SDM Networks

Traditional optical network graphs are not designed to encode the characteristics and peculiar constraints of SDM transceivers, links and optical switching equipment, all of which must be known to design an optical connection over such networks. A model to encode these SDM constraints, as well as a showcase of its flexibility, is described in Section 4.2 in Chapter 4.

2.2.3 Allocating Resources in SDM Networks

Due to the addition of the space domain, and the multitude of possible constraints deriving from combinations of different SDM transceivers, fibers and ROADM designs, the RWA/RSA problem of WDM and Flexi-grid networks requires yet another revision and expansion. In this setting an optical connection is associated to a path (route), i.e., a sequence on aggregate SDM links connecting optical switching nodes, optionally a modulation level, symbol rate and FEC selection, and a spectral, spatial or mixed super-channel, optionally with a per-link space assignment, making in the **Routing, Space and Spectrum Allocation** (RSSA) problem, sometimes also called the **Routing, Modulation Level, Space and Spectrum Allocation** (RMLSSA) problem. A more complete description of the emerging SDM technology and constraints, the RSSA problem and its variants, as well as proposals and simulative evaluations of several heuristics to solve it are detailed in Chapter 4.

2.3 Methodology

With respect to contribution 1 (reducing spectrum fragmentation), the many available metrics in literature were studied and, having identified some of their weakness, a new one designed to better capture what was perceived to be the relevant parameters leading to blocked connection requests due to fragmentation. Based on existing literature and best current practices in the field, the performances of the proposed metric were assessed through the use of synthetic computer simulations, trying to quantify its effectiveness in terms of a number of network performance metrics and comparing it against multiple existing metrics from literature.

With respect to contribution 2 (modeling SDM networks), a high level network model capable of representing SDM components was devised and implemented as both a YANG model and in a relational database (albeit, in the interest of full transparency, the thesis' author only supported this implementation work, which was mainly performed by a colleague), and example use cases were used as informal "induction steps" to show how it could be used to support weird configuration of SDM components.

Lastly, with respect to contribution 3 (allocating resources in SDM networks), existing and well known RWA/RSA algorithms were extended to take into account

the space dimension and its many constraints, and progressively refined to support additional constraints and maximize the benefits of the most promising approaches to SDM. Due to the high computational complexity of solving RWA/RSA exactly, the expectation that such complexity would increase with the introduction of the space dimension, and experience in the types of algorithmic solutions employed by commercial systems, priority was given to the development of practical and efficient heuristics rater than producing exact solutions via e.g. complex and computationally expensive Integer Linear Programming (ILP) formulations. Like before, these heuristics were tested using simulations to quantify their effectiveness in terms of a number of metrics, including network performance indicators such as total network throughput and connection blocking probability, and cost-related metrics such as required number of WSS ports, with the aim of assessing the benefits and costs of different SDM solutions. Multiple heuristics were iteratively developed and tested, with a secondary aim of lessening the performance penalty of the inherently cheaper (and hence more promising for realizing future deployments) SDM technologies while retaining their competitive edge in terms of cost.

Details of the evaluation of each proposed heuristic, including parameters and specific metrics used, are detailed for every simulation described in the thesis in their dedicated sections.

Chapter 3

Fragmentation-Aware Resource Allocation in Flexi-Grid Optical Networks

3.1 The Spectrum Fragmentation Problem

Flexi-grid optical networks [2], where each connection can be assigned a variable number of contiguous spectral "slices" defined on a much tighter grid than traditional WDM channels, have been proposed to increase the capacity of optical transport networks, by eschewing the intermediate guard-band separating multiple co-routed optical signals, thus creating "super-channels" with improved spectral efficiency, therefore increasing the total achievable network throughput.

The flexi-grid paradigm, however, also exacerbates the problem of spectral fragmentation [11], especially when connections may be deactivated (i.e., dynamic traffic scenarios). In fixed-grid optical networks fragmentation happens when a channel (λ) is available only on some of the links of a desired path. Flexi-grid networks also suffer from this type of fragmentation, applied to individual spectral slices on a path, but the problem in this context is more complex, since the continuous ranges of free spectral slices on which a connection is supposed to be allocated may:

A) be too small to support an incoming connection, and/or

B) overlap only partially on the links of a path, thus creating a smaller end-to-end usable range of free slices.

Additionally, fragmentation may also happen, even on a single link, when there are enough free spectral slices to support an incoming connection, but they are not contiguous (i.e., they belong to multiple smaller free ranges unsuitable for serving that connection).

Fragmentation is introduced when instantiating connections for two reasons:

1. Firstly, the new connection may cut, on some links of its path, a relatively large sequence of contiguous free spectral slices in two; this may affect both the end-to-end sequences of free slices in its own path, and those on other paths that intersect it.

2. Secondly, even when it does not split any contiguous free slice sequence, it may leave tiny sequences that are too small to serve some or any new connections. This problem also happens when removing a small (in terms of spectral slices) connection flanked by other lightpaths on either side of the spectrum in some links, thus leaving a usable but relatively small spectral void that is suitable only for a small subset of possible connection requests.

In other words, fragmentation may happen because spectral voids do not align on subsequent links in a path (inter-link fragmentation, which also exists in WDM, were different channels may be available on subsequent links), and because spectral voids on a link (or also path) are too small to support all possible super-channel sizes in use in the network.

3.2 State of the Art

The research field of optical wavelength/spectrum fragmentation is not new, however there is no consensus in the available scientific literature regarding how to measure fragmentation in optical networks, and more specifically flexi-grid networks: different authors propose different metrics, with no clear indication regarding which of these best captures the dynamics of this process.

For instance, denoting as **spectral void** a set of contiguous free slices on a certain link or path, the authors of [26] propose the **Fragmentation Ratio** metric, defined as the ratio of the maximal data rate provisionable using the available spectral voids and the one provisionable if the same number of free slices were contiguous, which can be greater due to improved efficiency of large super-channels, which eschew intermediate spectral guard-bands, and is easily computed using the well-known 0-1 knapsack algorithm (see e.g. [27]). More formally, denoting as G_i the i^{th} block of contiguous free spectral slices and as $v(G_i)$ the maximal data rate provisionable using block G_i, then the Fragmentation Ratio (FR) is defined as:

$$FR = 1 - \frac{\sum_i v(G_i)}{v(\sum_i G_i)} \tag{3.1}$$

The authors of [28] use, instead, a metric based solely on the number of spectral voids, while the authors of [29] present a fragmentation metric based on the well-known Shannon's Entropy [30] (which considers both voids size and number). More formally, quoting from [29], they compute their fragmentation index H_{frag} as:

$$H_{\text{frag}} = -\sum_{i=1}^{N} \frac{D_i}{D} ln \frac{D_i}{D}. \tag{3.2}$$

Where D is the total number of spectral slices in the spectrum of a link or path, and D_i is the size, in slices, of the i^{th} block of contiguous used or unused spectrum.

The authors of [31] use instead a metric, called Access Blocking Probability (ABP), which computes the ratio between the number of possible demands that can be placed on the available spectral voids in a certain link and the number that could be placed if all those free slices were contiguous, as encoded in Equation 3.3 (quoted from [31]):

$$\text{ABP} = 1 - \frac{\sum_i \sum_k f_i \; DIV \; G_k}{\sum_j all\_free\_slices \; DIV \; G_j}. \tag{3.3}$$

Where f_i is the number of spectral slices in void i, G_x is the set of possible demand sizes, and DIV is the integer division operator. It should be noted that this metric is very similar to FR, differing only in the way the capacity of spectral voids is computed (knapsack for FR against total number of fittable connections of each size for ABP).

In [32], the authors compute differential fragmentation as the number of spectral voids that would be cut (as opposed to just being shrunk) by the placement of a new

optical circuit. A different approach is taken by the authors of [33], who define an analytical model for predicting fragmentation, which is however limited to networks employing a single modulation level and a random fit spectrum allocation strategy.

While, as shown above, several metrics have been proposed to measure spectrum fragmentation, a related and possibly more practical problem pertains how to deal with it. Three main families of techniques have been described in literature to address fragmentation, which can be broadly classified as reactive, proactive or preventive:

- **Reactive** de-fragmentation techniques are triggered only when an incoming connection request is blocked, resulting in an attempt to re-organize the resources assigned to some or all active connections.

  An example of the latter is described in [34], where a blocked incoming connection may trigger the iterative shifting of some of the active connections closest to the spectral resources that are to be assigned to blocked connection (which potentially may trigger further shifts, up to some recursion depth or the absence of spectral resources to shift further). Another example is described in [35], where blocked connections may trigger the shifting of already established ones already occupying needed spectral resources using a Make-Before-Break approach. Another approach, described in [36], and later extended in [37] for entire super-channels, is known as "push-pull", and involves shifting an active connection in the spectrum without service interruptions. This technique is, however, only applicable when the shifting connection is not re-routed, and when there are no other connections active on any intermediate spectral slice between the starting and final positions of the shifting connection.

- **Proactive** de-fragmentation techniques are similar to reactive ones, but are run either periodically or as soon as a threshold of total network fragmentation, as measured by some metric, is reached.

  An example of the former can be found in [38], where each optical connection periodically checks whether it can shift itself towards smaller frequencies, resulting in an overall tendency of compacting in-use spectral resources towards the left side of the spectrum. Another can be found in [39], where active connections are re-organized every time a certain number of them are terminated.

- **Preventive** techniques do not perform resource de-fragmentation, but rather carefully consider the placement of new connections during the computation of solutions to the RWA/RSA problem in order to minimize fragmentation (e.g. by attempting to exactly fill the gaps left by a terminated connection), and therefore reduce or eliminate the need to execute complex and potentially disrupting de-fragmentation processes later. Note that this is the only feasible approach for static network scenarios, where connections can be added (even in large groups) but never removed.

  There are many examples of such fragmentation-aware RSA schemes in literature, such as [28], [29], where the authors propose a heuristic RSA scheme based on minimizing an entropy-based fragmentation metric. Another work using an entropy-based metric is [40], where the authors use the "utilization entropy" of a path (the average entropy of each slice on that path, computed on the sequence of links of the path) to measure fragmentation, and adaptively change the minimum size of connections based on the overall fragmentation on the network, reducing it as fragmentation increases. This implies that

large demands are increasingly split and inversely multiplexed into multiple smaller ones to increase their probability of fitting into the available resources as fragmentation increases. Another heuristic can be found in [32], based on the number of cuts fragmentation metric, which is also used in [41], in a distributed GMPLS setting, and [42], in a centralized SDN controller. Yet another approach for preventing fragmentation, and in particular lessening its effects for larger demands, which are more vulnerable to it, is to separate the available spectral resources into partitions, and assigning them either forcedly or just preferably to connections of a certain size, as described by e.g. [43]. A variation of this approach is described in [44], where partitions are strictly assigned to a particular demand size, with an extra shared partition to accommodate any demand that does not otherwise fit. The authors of [45], instead, use an approach based on always leaving spectral voids that can contain an integer number of new connections, on the assumption that their possible sizes are known in advance. Another example can be found in [46], where the authors propose to use a metric very similar to that proposed in [26], and to use an heuristic which finds the assignment that minimizes such metric on all the links of the network, potentially also giving preferences to utilizing slices that are already utilized in some other part of the network.

One issue with proactive and especially reactive de-fragmentation techniques is that of the long time needed to shift and-or re-route potentially many optical connections in a serial manner, which may cause the network to be unresponsive to new connection demands in the interim. The authors of [47] tackle this issue, by proposing several algorithms that attempt to parallelize the de-fragmentation process, albeit limited to the case of spectral shifting and not connection re-routing, since the latter generally requires the re-equalization of the power of all affected amplifiers, which is typically a long process that can take several minutes.

Lastly, the problem of fragmentation is not always limited to just the spectrum domain. The authors of [48] describe the problem of time fragmentation, which occurs when advance-reservation connection demands are considered, and propose a technique to preemptively re-arrange scheduled (but not yet active) connections to make room for new advance-reservation demands. Furthermore, the introduction of the space dimension in Space Division Multiplexing (SDM) optical networks, further described in Chapter 4, can lead, for some types of SDM hardware, to the problem of space fragmentation, as discussed in [49], as well as alternative solutions, such as using spare, high cross-talk cores to perform de-fragmentation using Make-Before-Break even in highly loaded networks [50], [51].

All of these works measure fragmentation based on either a binary vector representation of the spectral occupation of each link, thus failing to recognize misalignments of spectral voids on successive links, or a binary vector representation of the spectral occupation of each path, which captures the end-to-end spectral voids but loses information pertaining free slices that are only available on some of the links of a path.

3.3 The Wasted Slices & Usability Factor (WSUF) Fragmentation Metric

To improve on existing fragmentation metrics, a novel metric which considers the ratio between the unusable or even partially usable slices on a path and the total

free slices is proposed in this section. In order to provide a formal definition, a few auxiliary concepts need to be defined first.

An **Available Frequency Range on a Link** (AFR_L) is defined as a range of frequencies that is not allocated to any connection in that link, i.e., a group of contiguous free spectral slices in a link ($\neq \emptyset$, i.e., containing ≥ 1 slices). This is equivalent to what was earlier called a spectral void (also known as spectral hole or spectral gap in literature), but restricted to the case where the considered spectrum models a single link.

An **Available Frequency Range on a Path** (AFR_P) is the corresponding concept applied to paths, that is, defined instead as a frequency range that is not allocated to any connection in any link of that path. In other words, an AFR_P on a certain path represents a group of contiguous spectral slices that describe the same free frequencies in all the links of the path and can therefore be used to support a new connection on that path.

Given these definitions, the **Wasted slices & Usability Factor** (**WSUF**) fragmentation metric can be defined, for a certain path, as the ratio between the sum of:

1. the **Wasted Slices** factor W, which counts slices that are free on a strict subset of the links of the path (more formally: which belong to at least one AFR_L but not an AFR_P);

2. the **Usability** factor U, which accounts for small AFR_Ps, i.e., end-to-end spectral voids that are only big enough to support a strict subset of the possible super-channels supported by the network

This sum is divided by the **Number of Free Slices** (F) on the path (i.e., the sum of all free slices on each of the links of the path), as shown in Equation 3.4:

$$\text{WSUF} = (W + U)/F \tag{3.4}$$

More in detail, the spectrum availability on a link (an AFR_L) is described, without loss of generality, using a binary vector, where 0 means that a slice is free and 1 that it is currently being used by an active optical connection. Note that this choice is completely arbitrary and was done for compatibility with existing metrics; a flipped representation would just require adjusting the following equations (which would, in fact, be even simpler). The integer vector S resulting from the sum of all the binary vectors ("casted" to integer as described above) of the links of a certain path therefore represents the slice availability on that path. The values of each element $s \in S$ are, by construction, confined to the range $[0, N_L]$, where N_L is the number of links traversed by the path under study.

Given this representation, the W, U and F terms of that path can be computed using the following equations:

$$W = \sum_{s \in S} \begin{cases} N_L - s & \text{if } 0 < s < N_L \\ 0 & \text{otherwise} \end{cases} \tag{3.5a}$$

$$U = \sum_{h \in H} \left(|h| \cdot N_L \cdot (1 - p\left(|h|\right)) \right) \tag{3.5b}$$

$$F = \sum_{s \in S} (N_L - s) \tag{3.5c}$$

The Wasted Slices term, W, is defined in equation 3.5a as the sum of those elements of the vector that are neither 0 nor N_L, or, in more intuitive terms, the count

of those slices that are free only on some of the links of this path, and are therefore unusable for serving a connection on it.

The Usability Factor term, U, is defined in equation 3.5b, where the set H is the set of all AFR$_P$s on the path; the cardinality of one of its elements $|h|$ is the size, in spectral slices, of the AFR$_P$ h. The p function encodes the probability that an incoming connection could fit in that particular AFR$_P$. In other words, equation 3.5b factors as unusable a fraction of the perfectly usable slices of an AFR$_P$ proportional to the probability of not being able to fit an incoming connection in it. Computing p requires foreknowledge of the possible sizes of the incoming connection requests (i.e., the spectral width of the associated super-channels), and at least approximate knowledge of their relative distribution, so that more frequent sizes can be given more weight. Note that both of these parameters should be at least approximately known to operators, since they depend on their own policies and available hardware. If such information is not available, super-channel dimensions can be assumed to be uniformly distributed. Also note that in the degenerate case of uniformly distributed connection sizes, p reduces to computing the fraction of possible super-channel sizes that fit into a certain AFR$_P$.

Finally, the free slices term F, defined in equation 3.5c, is the count of free slices on the links of the path, and can be obtained from S simply by summing, over the elements $s \in S$, the number of links in the path (N_L) minus s (which, by itself, counts the links in which a particular slice is occupied).

Observe that, by construction, the WSUF metric for any path can exhibit values in the range $[0, 1]$, where 0 means no fragmentation (i.e., no wasted or unusable free slices), for example when there are no active connections on any link of the path, or when all slices are in use, and where 1 means that all available (i.e., free) slices on the links of the path cannot be used to establish connections, for example because no two AFR$_L$s align on different links. WSUF is purely a fragmentation metric as, unlike some other metrics, it contains no information about the number or ratio of free and occupied slices.

3.3.1 Examples of Computing Fragmentation

As an example, imagine a simple network with four nodes, A, B, C, and D and three links, A-B, B-C, and C-D, whose spectrum occupation is depicted in Figure 3.1, and where incoming connections can require 4 or 7 slices, corresponding to 100 and 200 Gbps super-channels (with the latter being more spectrally efficient due to avoiding the intermediate guard-band between two 100 Gbps signals), with equal probability. The path vector for path A-B-C-D ($N_L = 3$) can be obtained by simply summing the vectors S_{AB}, S_{BC} and S_{CD}, obtaining the path spectrum vector S.

Then, the count of wasted slices W can be computed as detailed in equation 3.5a, which corresponds to summing N_L minus the values in the yellow boxes (neither zeroes nor threes), obtaining: $W = 11$. These are the free slices that cannot be use on this path because they are not free on all of its links. Computing U is a little more involved: start with $U = 0$; the first AFR$_P$ (green group) is 3 slices wide, but 3 is too small to support any connection, so $p(3) = 0$, and therefore:

$$U \leftarrow U + 3 \cdot 3 \cdot (1 - 0) = 0 + 9 = 9.$$

That is, all the slices of the first AFR$_P$ group, on all the links of the path, are counted as unusable. Continuing with the second AFR$_P$, which has size 6, only half of the

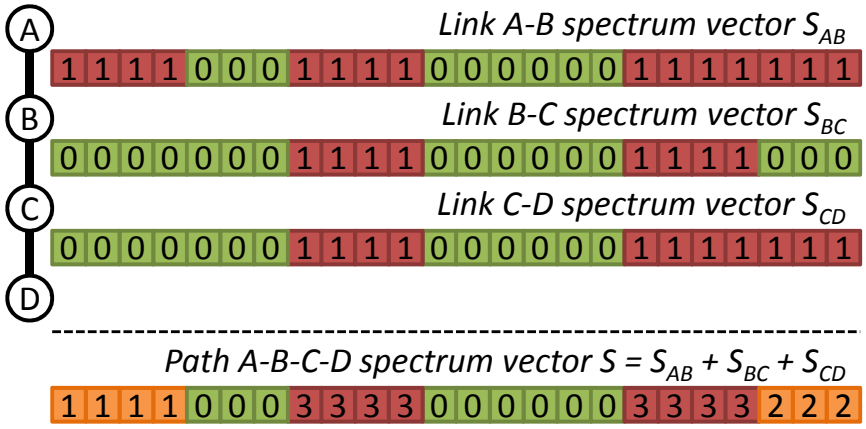

FIGURE 3.1: Examples of Link and Path Spectrum Vector representations on a simple linear topology, exhibiting 7 AFR_Ls and 2 AFR_Ps.

possible connections can fit in it, therefore:

$$U \leftarrow U + 6 \cdot 3 \cdot (1 - 0.5) = 9 + 9 = 18.$$

In other words, the slices in this AFR_P are considered as only half usable because it has a 50% chance of being unable to fit an incoming connection. Finally, F is computed by simply applying equation 3.5c, obtaining $F = 38$. Therefore, the fragmentation index WSUF for path A-B-C-D is:

$$\text{WSUF} = (W + U)/F = (11 + 18)/38 = 0.763.$$

Observe that the output of WSUF has a **probabilistic interpretation**: there is a 76.3% chance that a free slice on this path may not be useful, due to the fragmentation of available resources, to establish an optical connection on it. Note that a similar argument can be made for the ABP metric when applied to a vector representing availability on a paty, but only for the subset of slices that are free on the entire path.

Since WSUF is a path-specific fragmentation metric, and optical network control planes generally attempt to place connections on the shortest possible path (to minimize the effect of PLIs), in order to obtain a network wide fragmentation value it is possible to simply compute the average WSUF among the shortest paths between each possible source-destination couple (thus preserving the probabilistic interpretation of the metric's output). If additional information is available pertaining the likelihood that an optical connection will be requested between two specific nodes or the specific routing policies used, this simple average can be straightforwardly replaced with a weighted one, to give more importance to the busiest communicating source-destination couples, and/or ignore the effect of fragmentation on unfeasible or unused paths.

Conversely, the amount of fragmentation on path A-B-C-D according to the FR [26] and H_{frag} [29] metrics can be computed using equations 3.1 and 3.2 on the binary path vector S' derived by setting all $s \in S | s > 0$ to 1. By applying equation 3.1, it is possible to find that G_1, i.e., what was earlier called the first AFR_P, has a potential provisionable data rate of 0, since it cannot support any connection, while G_2, which corresponds to the second AFR_P, can support a single 100 Gbps super-channel. If they were not disjoint, they could support either two 100 Gb/s super-channels or a

single 200 Gb/s super-channel, therefore:

$$FR = (0 + 100)/200 = 0.5$$

This captures the fragmentation on the end-to-end spectrum, but cannot, since it is based on a binary vector, capture that of slices that are only available on certain links of the path (hence the lower value). Note that this metric could be applied directly to links rather than paths, yet this would still leave it unable to capture information related to partially available spectral slices on paths.

Along the same lines, H_{frag} can be computed by applying equation 3.2 to S', where the outcome is:

$$H_{\text{frag}} = -\left(\frac{4}{24}ln\left(\frac{4}{24}\right) + \frac{3}{24}ln\left(\frac{3}{24}\right) + \frac{4}{24}ln\left(\frac{4}{24}\right) + \frac{6}{24}ln\left(\frac{6}{24}\right) + \frac{7}{24}ln\left(\frac{7}{24}\right)\right) = 1.56$$

Which, in addition of being unable to capture information related to partially available spectral slices on the path, is somewhat difficult to interpret due to lacking a clear upper bound. Please note that both FR and H_{frag} could also be computed on link vectors, whereas this does not make much sense for WSUF, for which the link would simply be a single-hop path.

3.4 Fragmentation-Minimizing RSA Heuristics

This section describes two heuristic RSA algorithms that minimize a fragmentation metric, both reliant on the well-known and widely used K Shortest Paths routing algorithm [7]:

- **Minimize-Fragmentation-Routing** (**M-F-R**), which analyzes all feasible slice assignments for an incoming connection over all pre-computed K shortest paths;
- **Minimize-Fragmentation-NoRouting** (**M-F-NR**), which analyzes all feasible slice assignments for an incoming connection over just the shortest usable path (i.e., the shortest among the pre-computed paths with at least one suitably large AFR_P).

In both algorithms "Fragmentation" represents and must be specialized into any fragmentation metric defined on the whole network. Both are defined over the first K shortest paths used for routing, but while M-F-R analyzes all the usable ones (up to K), M-F-NR restricts its search to the first (i.e., shortest) usable one. In both cases, requests are blocked if and only if no path among the K-shortest has an AFR_P of sufficiently large size.

Note that the proposed RSA schemes, particularly M-F-R, can be computationally expensive, since there can be a large number of possible assignments to analyze. Perversely, this is especially true in lightly loaded scenarios, where the number of possible assignments is close to the upper bound of $K \cdot (N_S - N_C + 1)$, where K is the K parameter of KSP (it is just 1 in the case of M-F-NR), N_S is the number of elements of the path vector (i.e., the number of spectral slices), and N_C is the number of slices required by the incoming connection. Nonetheless, in practical tests, running a simple (i.e., not particularly well optimized) implementation of M-F-R with a relatively complex metric (WSUF) on a realistic scenario (see Section 3.5.1) took

in the order of tens of seconds using commodity hardware (a not particularly fancy laptop). This is still smaller than the time needed to equalize power in the traversed OLAs, and may even be low enough to be implemented on top of an optical node, but it is definitely feasible to run even such simple and unoptimized algorithms on a centralized Path Computation Element (PCE) [52] or SDN controller, which are generally based on powerful commodity servers. Furthermore, since for both proposed algorithms each candidate assignment can be checked in isolation, it is trivial to partition the search space and parallelize the search on multiple CPU cores (which in fact was done to speed up the simulations described in Section 3.5.2, reducing the worst-case execution time to a few seconds).

Another problem related to the implementation of the proposed metric and algorithms in real networks is that of having access to the information necessary to perform the computation. In fact, it is necessary to know:

A) The state of occupation of every slice of every link in the network (for most metrics);

B) The weights associated to the p function of equation 3.5b (for WSUF specifically).

Existing distributed control planes may or may not distribute A), and definitely do not distribute B) without appropriate extensions (note that each node should already be able to compute the list of shortest paths between all nodes in distributed settings). Both pieces of information can easily be made available in a centralized setting.

3.5 Simulative Analysis of Fragmentation Metrics Performance

3.5.1 Simulation Setup and Parameters

The performance of the proposed metric and RSA heuristics was evaluated using a custom purpose-built event-driven simulation tool, employing the German national backbone network topology consisting of 14 nodes and 46 directed links (Figure 3.2, [53]). Links were modeled using 384 12.5 GHz spectral slices (i.e., as standard C-band fibers).

Connection demand sizes were chosen with a uniform distribution from a set of {100, 200, 300, 400} Gbps corresponding to {4, 7, 10, 12} spectral slices of 12.5 GHz, assuming DP-QPSK modulation at 32 Gbaud, which also has enough reach to connect any two nodes in the chosen topology without needing signal regeneration. Neither impairment validation for lightpath feasibility evaluation nor signal regeneration are considered.

Simulations involve a dynamic traffic scenario, where connection request are generated according to a Poisson process, with a fixed arrival time and an exponentially distributed holding time chosen so as to have a target ideal input network load, expressed in terms of the fraction of spectral resources available in the network that would be in use if those connections were accepted on their shortest path. Each simulation involved generating 10^6 connection requests (plus an additional 10^3 to achieve a desired initial steady state and not accounted for in the results) in a dynamic traffic scenario at a desired target ideal input network load.

The performance of the proposed RSA algorithms (M-F-R and M-F-NR) was evaluated using three fragmentation metrics: WSUF, Fragmentation Ratio (FR) from [26], and Entropy (ENT) from [28], [29], leading to six different RSA strategies: three

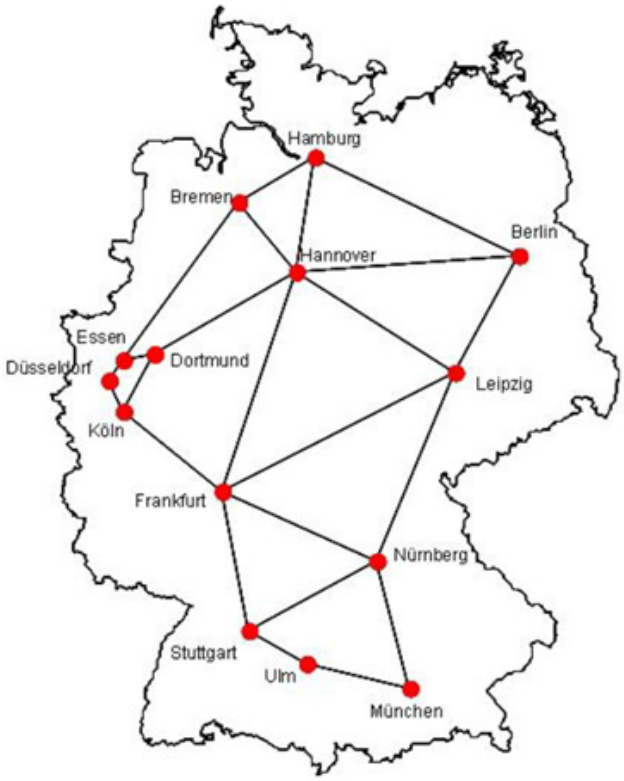

FIGURE 3.2: German national backbone network topology.

based on Routing (M-WSUF-R, M-FR-R and M-ENT-R) and three on NoRouting (M-WSUF-NR, M-FR-NR and M-ENT-NR).

Please note that for these results the network-wide values for WSUF and FR where averaged over the path vectors of all shortest paths, while those for Entropy where averaged over the vectors describing all network links, as this is suggested to be the most effective solution in [28], [29]. FR and Entropy were also tested by averaging, respectively, their value for each link and shortest path vectors, obtaining broadly similar results, indistinguishable from those presented as far as the ranking of these algorithms is concerned.

The performance of these algorithms was also compared to that of the the well known First Fit (FF) strategy in the same conditions, namely employing KSP routing, with $K = 3$, where paths were ordered based on the number of hops they contained. FF, described in more detail in Algorithm 1 in Section 4.4, chooses the first (i.e., left-most) available AFR_P in the first (i.e., shortest) path with a suitably large AFR_P.

3.5.2 Simulation Results

The effectiveness of the proposed and existing metrics was evaluated in terms of resulting **Blocking Probability** (BP), defined as the ratio between rejected and total connection requests, measured **Network Load**, defined as the ratio between in-use and total spectral slices, and **Network Throughput**, defined as the average capacity of all active optical circuits over time, against the input network load. Furthermore, results concerning the amount of fragmentation present in the network (using the proposed metric, WSUF) and the effect of changing the K parameter of KSP are also reported. Since plotting all results for a given performance metric in a single graph produces figures that are very difficult to read, they are split between Routing and

NoRouting algorithms, showing FF in both cases to have a common baseline, to improve readability.

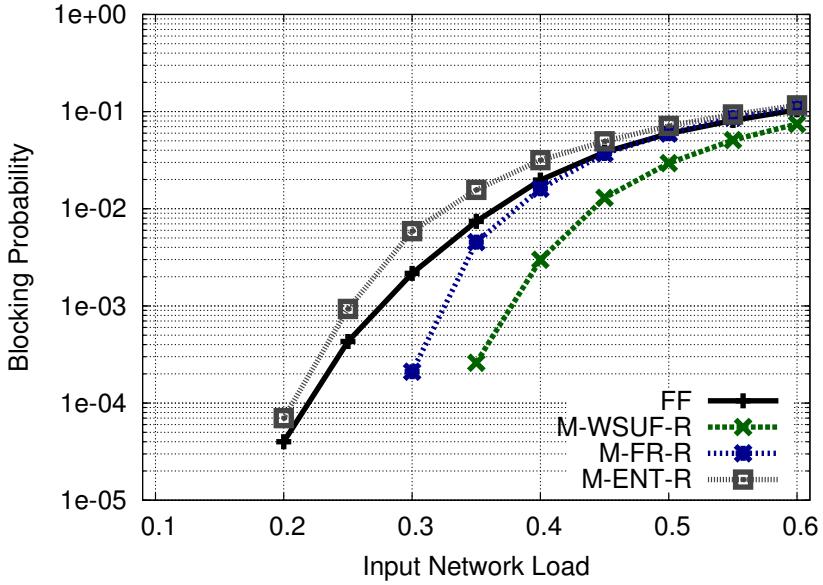

FIGURE 3.3: Blocking Probability vs. Input Load for FF, M-WSUF-R, M-FR-R, and M-ENT-R.

Figure 3.3 depicts the measured BP against the input network load for FF and the Routing based strategies, namely M-WSUF-R, M-FR-R, and M-ENT-R. Here M-FR-R performs significantly better than the baseline FF, starting to block later (at about 30% load rather than 20%), and exhibiting lower overall blocking, at least until 40-45% load, when it's BP converges to that of FF. Even better is the performance exhibited by M-WSUF-R, which starts to block even later (at about 35% load) and blocks visibly less all the way to 60% load, even when plotted in logarithmic scale. Lastly, M-ENT-R performs consistently worse than FF; the probable reason behind this behavior are discussed later, after further evidence has been shown.

Figure 3.4 shows instead the measured BP against the input network load for FF and the NoRouting strategies, i.e., M-WSUF-NR, M-FR-NR, and M-ENT-NR. Like in the previous case, M-FR-NR and M-WSUF perform slightly better than the baseline, with the latter exhibiting slightly lower blocking than the former, while M-ENT-NR appears to be, again, slightly worse than FF.

These figures show that, while M-F-Routing algorithms are more computationally expensive than M-F-NoRouting ones, they exacerbate the benefits (or drawbacks) of the chosen fragmentation metric. Please note that both families of algorithms average fragmentation over all shortest paths, so in both cases small AFR_Ls on "busy" links (i.e., links that are used by a larger number of paths) are counted proportionally more than those on nearly unused ones. M-F-Routing algorithms, however, unlike M-F-NoRouting ones are allowed to act on this system of accounting for fragmentation, by selecting alternative (and therefore longer when using KSP routing) paths which, despite consuming more spectral resources than strictly necessary, is shown to potentially bring significant benefits in terms of blocking when done to avoid creating fragmentation on "busy" links. This matches the intuition that creating a large amount of fragmentation on a nearly unused link is likely to be less of a problem than wasting even a small amount of spectral resources on a very busy link.

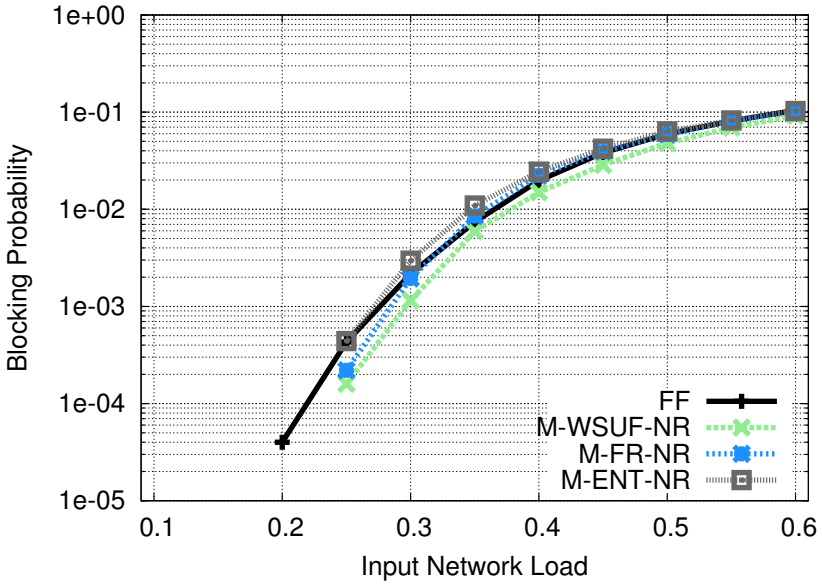

FIGURE 3.4: Blocking Probability vs. Input Load for FF, M-WSUF-NR, M-FR-NR, and M-ENT-NR.

An additional investigation on blocking performance, based on using the Bandwidth Blocking Ratio (BBR) metric, i.e., the weighted average of the BP of each class of service (100 Gb/s, 200 Gb/s, etc.) weighted against its bandwitht, produced nearly identical results to those of simple BP, and is therefore not shown. In fact, measuring blocking using BBR instead of BP produced the exact same ranking of the various algorithms (i.e., the order of best to worst performing does not change at all).

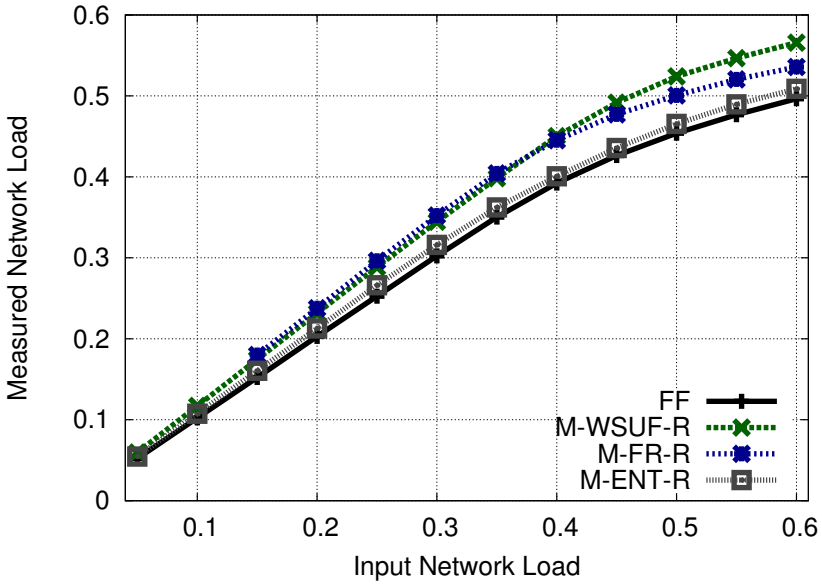

FIGURE 3.5: Measured Load vs. Input Load for FF, M-WSUF-R, M-FR-R, M-ENT-R.

The behavior of the various algorithm with respect to BP is also reflected in manner in which the measured load scales with the input load, as depicted in Figure 3.5

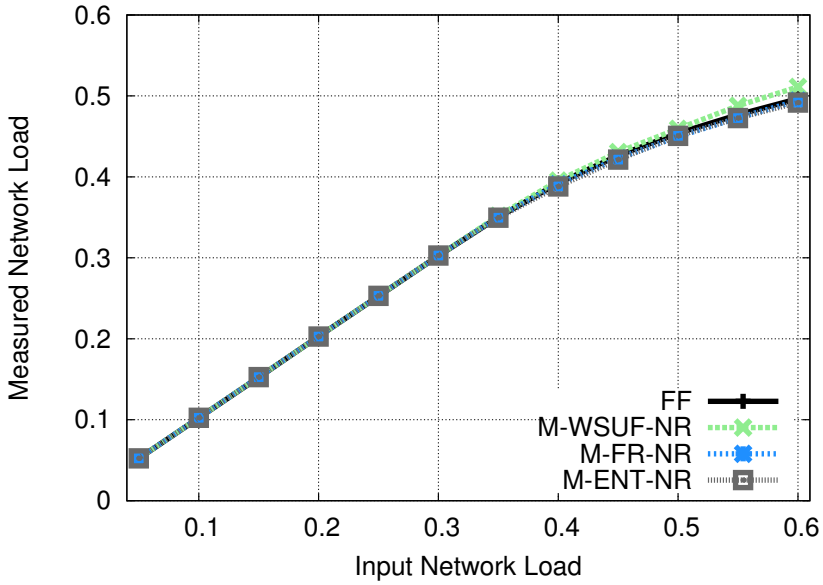

FIGURE 3.6: Measured Load vs. Input Load for FF, M-WSUF-NR, M-FR-NR, M-ENT-NR.

and Figure 3.6, which depict the measured vs. input network load for FF and the M-F-R and M-F-NR algorithms, respectively. As can be expected, while he input load is low, and therefore blocking is nearly or completely absent, the measured load scales linearly with the input load. More specifically, it scales exactly as the input load for FF and the M-F-NR family of algorithms, while M-F-R algorithms result in an higher effective load than requested. This is a direct result of the way the expected input load is measured: as the fraction of spectral resources (i.e., spectral slices) that should, on average, be occupied if all input connections were accepted on their shortest paths. Since M-F-R sometimes shun shortest paths in favor of longer ones, this naturally entails that the load they place on spectral resources is higher than their M-F-NR counterparts. At higher loads, where blocking starts to become significant, the measured load in the network starts to scale sub-linearly with respect to the input load for all algorithms. Of the M-F-NR algorithms, only M-WSUF-NR exhibits a detectably different load than the baseline FF, and this only at higher loads, where this difference can be attributed to slightly lower BP and therefore an higher number of active connections in the network. Of the M-F-R algorithms, M-ENT-R shows a slightly higher resource consumption than FF, despite a comparable or slightly worse BP. M-WSUF-R and M-FR-R both show markedly higher resource consumptions matching their increased rate of acceptance of connections, to the point that the former starts to outdistance the latter at about 40-45%, due to the significantly lower blocking exhibited by M-WSUF-R (in the order of 10^{-3} vs. 10^{-2} at 40% load).

The next set of results describes the measured network-wide throughput against the input load of the network, as is depicted in Figure 3.7 and Figure 3.8, which again group M-F-R and M-F-NR algorithms. Consistently with the earlier results, at low loads there is no difference between the various algorithms, since blocking is too infrequent to have significant impact. As the input load increases, M-WSUF-NR and especially M-WSUF-R start to exhibit significantly higher throughput than FF. The other M-F-R algorithms actually show a decrease in total throughput compared to FF at very high loads, while M-FR-NR and M-ENT-NR do not exhibit significant differences from the baseline.

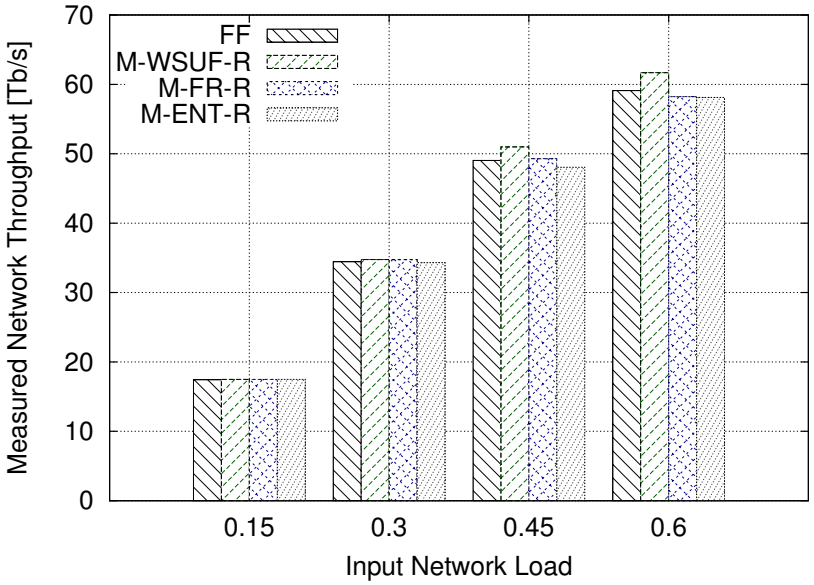

FIGURE 3.7: Measured throughput vs. Input Load for M-WSUF-R, M-FR-R, M-ENT-R and FF.

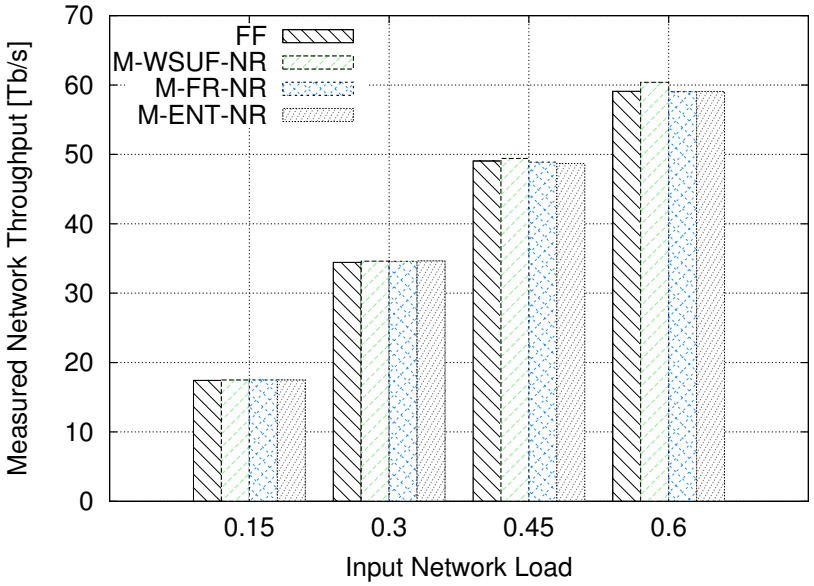

FIGURE 3.8: Measured throughput vs. Input Load for M-WSUF-NR, M-FR-NR, M-ENT-NR and FF.

Enough evidence has now been to discuss in greater detail and summarize the performance of the proposed algorithms: in short, the M-F-NR family of algorithms can offer limited to no increase in total throughput compared to a simple FF, by slightly reducing the BP experienced by the network. By using only the shortest path with a suitably large AFR_P, like FF does, these algorithms do not measurably increase the amount of spectral resources used by the network compared to FF, at least not until lower blocking justifies this increase with improved throughput. This is the case for M-WSUF-NR, which exhibits a 2.9% increase in used spectral resources compared to FF, but also sports a 2.2% increase in average network throughput thanks to accepting a greater percentage of incoming demands.

The M-F-R family of algorithms, instead, exhibits more dramatic results: it can offer larger improvements in terms of blocking and throughput (up to 4.5% more instantiated capacity for M-WSUF-R), which are strongly dependent on the efficacy of the fragmentation metric being used, but at the cost of using significantly more spectral resources (almost 14% more for M-WSUF-R). This increased usage of spectral resources is the consequence of two factors: 1) at higher loads, the increased network throughput exhibited by some of these algorithms implies that more and/or larger (and thus more efficient) connections are being established, and 2) at all loads, these algorithms have the ability to route requests on paths that are not strictly shortest even if there are enough resources on shorter ones, when doing so reduces the overall fragmentation they perceive. This is done whenever taking a detour moves enough fragmentation from the busiest links (i.e., the links traversed by the most paths) to the least used ones, since fragmentation on a link is counted as many times as the numbers of paths traversing it, so that the net amount of fragmentation in the network would be lower than if the detour was not taken. This implies that M-F-R algorithms exhibit an implicit weak link-load balancing mechanism performed jointly with fragmentation minimization, which can have a large impact in terms of BP at low loads. Observe that this is different from simply using a load-balancing mechanism that selects the best route (among the K-shortest) and then computes the best assignment to minimize fragmentation, which is exactly what is done by M-F-NR algorithms.

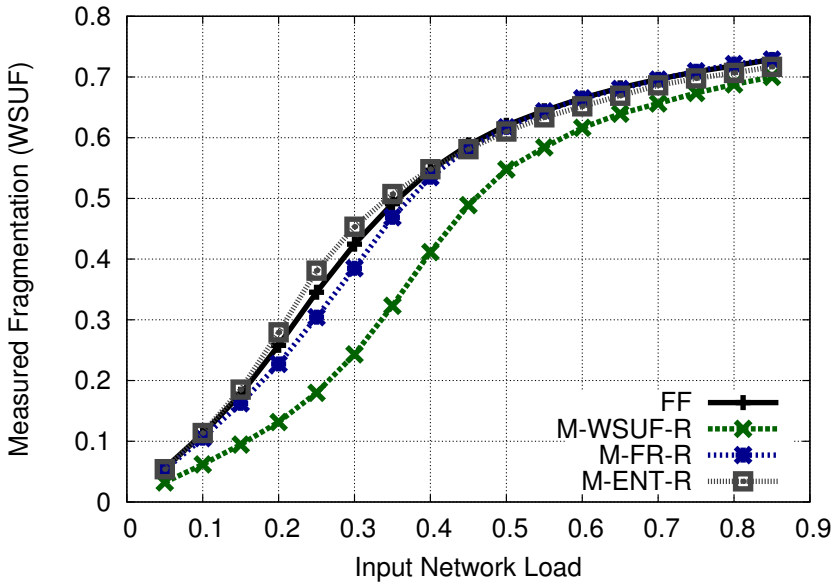

FIGURE 3.9: Measured fragmentation as WSUF vs. Input Load for FF, M-WSUF-R, M-FR-R, and M-ENT-R.

To corroborate these results, the rest of this section further investigates the effect of the choice of fragmentation metric and of the K parameter in KSP. Figure 3.9 shows the measured fragmentation, according to the WSUF metric, for FF, M-WSUF-R, M-FR-R and M-ENT-R, while Figure 3.10 shows the same for FF, M-WSUF-NR, M-FR-NR and M-ENT-NR. The WSUF metric was chosen because, as discussed in Section 3.3, it is the one that captures the most nuances about the issue of spectral fragmentation.

As expected, the WSUF-based algorithms clearly outperform the others when minimizing the fragmentation metric which they are designed to minimize, with the

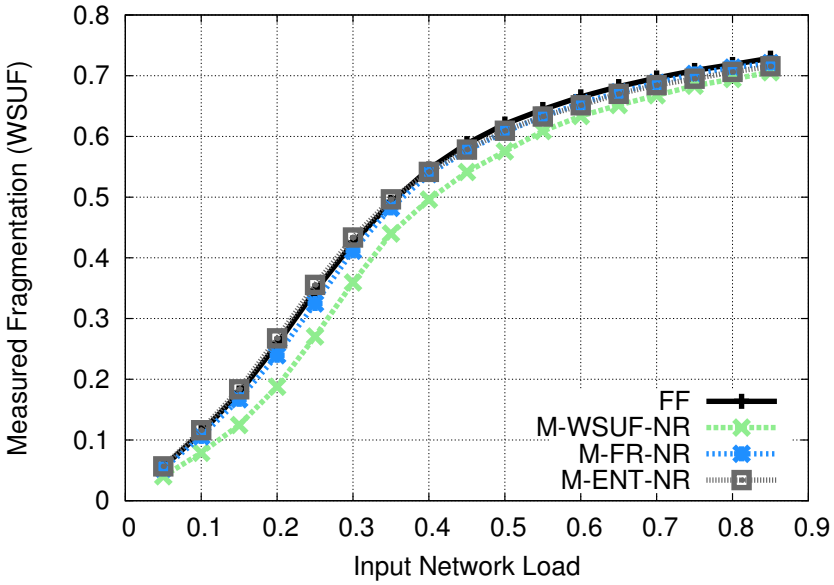

FIGURE 3.10: Measured fragmentation as WSUF vs. Input Load for FF, M-WSUF-NR, M-FR-NR, and M-ENT-NR.

M-F-R algorithm clearly outperforming its M-F-NR counterpart, consistently with the results shown so far.

The FR-based algorithms are much closer in performance, with M-FR-R outperforming M-FR-NR at middling loads but being slightly worse at higher loads (as best reflected in the throughput results). This suggests that the improvements in terms of BP and throughput exhibited by M-FR-R are mainly due to the load-balancing mechanism implicitly performed by M-F-R algorithms, and explains why there is a negligible improvement in exhibited throughput when using M-FR-NR, while M-FR-R is even worse than FF at very high loads (where the extra resources spent when taking detours have an impact on the overall performance).

Lastly, for the ENT-base algorithms, M-ENT-NR appears to be more or less comparable to FF, while M-ENT-R, in accordance with previous results, actually increases fragmentation (as perceived using WSUF), suggesting that entropy may not be an effective fragmentation metric in the dynamic traffic scenario used in this study (in [28], [29] a static network filling scenario was used).

Finally, to ensure that the previous results were not due to a lucky choice of parameters, the effect of changing the value of K, i.e., the effect of considering less or more candidate paths for resource assignment was investigated. Figure 3.11 shows the effect of changing the value of K on blocking probability for M-WSUF-R, M-WSUF-NR and FF, using an input load of 40%; similar results were obtained for the other loads and algorithms we tested. With $K = 1$, i.e., only the shortest path being allowed, M-F-R and M-F-NR are, by construction, equivalent, and only slightly better than FF. Adding alternatives to the shortest paths is initially strictly beneficial, for all algorithms. However, at some points the benefits of having an additional, longer alternative are sometimes outweighed by the fact that by using a path that is longer than normal the resources used for a connection are significantly more than what are strictly needed; this implies that those resources are then unavailable to serve other connections on more efficient (i.e., shorter) paths, therefore reducing overall efficiency and, ultimately, total throughput. Figure 3.11 suggests that for values of K higher than 4 or 5 the drawback described above more or less counterbalances

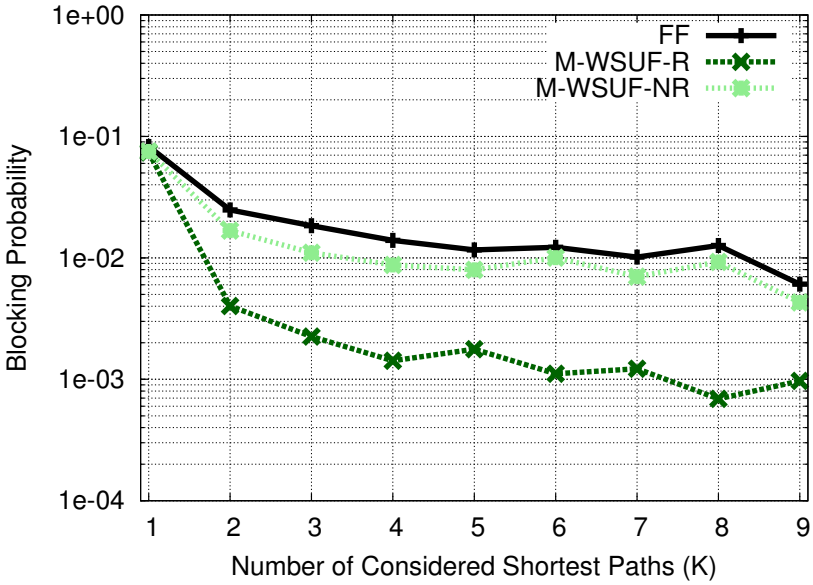

FIGURE 3.11: Blocking Probability vs. value of the K parameter with an input load of 40% for M-WSUF-R, M-WSUF-NR and FF.

the benefits of adding alternate paths, therefore leading to oscillations in the measured BP. Nonetheless, Figure 3.11 shows that the results described above remain consistent for sane values of K.

Chapter 4

Resource Allocation and Modeling for SDM Optical Networks

4.1 State of the Art

Space Division Multiplexing (SDM) and related sub-topics have emerged as an hot research subject in recent years. Back when work on this thesis started in late 2014 there was already a significant number of published works regarding SDM components, outlined in Sections 4.1.1 through 4.1.6, and point-to-point communications, yet very little in terms of networking, models and resource allocation.

In order to apply SDM to optical transport networks all the components outlined in Section 1.1 require slight modifications. While studying these modifications in detail is outside the scope of this thesis, an outline of the necessary advances as available in literature, with an emphasis on aspects related to modeling and resource allocation issues, is given in the rest of this section to provide appropriate context for the work discussed in the rest of this chapter.

4.1.1 SDM Fibers

A multitude of fiber technologies have been proposed to realize optical SDM:

- **Single-Mode, Single-Core Fiber Bundles** (Figure 4.1 (A)): standard Single-Mode Fibers (SMFs) for long-haul transmission, often already deployed in large bundles to offset the costs of digging trenches (about 8 fibers for oceanic cables, more than one hundred for roadside cables).
- **Multi-Core Fibers** (Figure 4.1 (B)): fibers with multiple cores within a single fiber cladding, forming multi-core fibers (MCFs), and offer an increase in available bandwidth equal to their core count (assuming each core only supports a single spatial mode). Examples with up to 19 cores can be found in literature [54].
- **Few-Modes Fibers** (Figure 4.1 (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.
- **Combinations** of the above (e.g. [55], [56]).

Individual fibers, cores or spatial modes (depending on the chosen technology) are generically called **spatial dimensions**.

While the "advanced" SDM fibers offer a potential capacity multiplier equal to their core $\times$ mode count (i.e., to the number of spatial dimensions) in the same

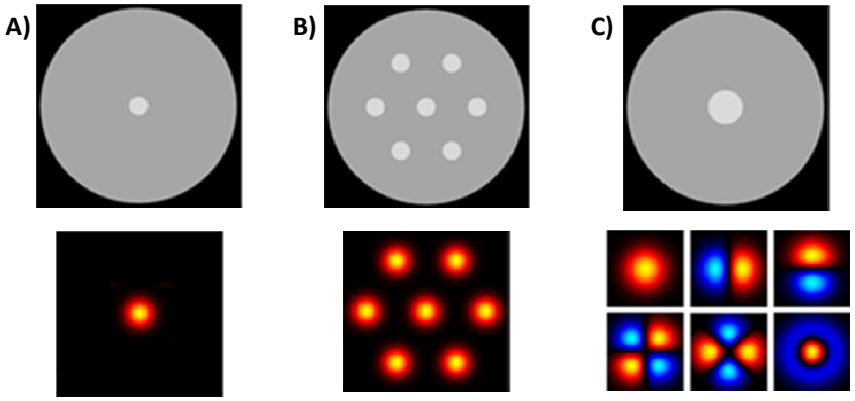

FIGURE 4.1: Different types of SDM fibers: Single Mode, Single Core Fiber (A), Multi-Core Fiber (B), and Few-Mode Fiber (C).

physical space, the key difference between them (and standard fiber bundles) is the amount of additional PLI they induce, such as mode **coupling**, core **crosstalk**, and issues related to Differential Group Delay (DGD). DGD issues are similar to those affecting old parallel (e.g. printer) cables, and occur because multiple parallel signals may travel slightly different paths (for MCFs, but also any bundle of fibers) or at slightly different speeds (for FMFs), arriving at the receiver at different enough times to be considered part of different symbols [57]. 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 unraveled at the receiver using Multiple Input Multiple Output (MIMO) Digital Signal Processing (DSP) techniques, provided that all coupled modes are routed and received together [58]. 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 different levels of core crosstalk, depending on the core count, the distance between cores and their geometric arrangement within the cladding. This can be corrected using MIMO DSP techniques similar to those used for FMFs, but in practice translates to a reduction in the usable reach of optical signals traversing such fibers when "leaky" cores are used at the same frequencies. Well isolated MCFs exhibit ultra-low crosstalk properties with respect to distance ($<$ -30 dB over 10.000 km [59], [60]).

Finally, recent MCF designs consider few mode fiber cores [55], [56], supporting the transmission of uncoupled spatial groups of coupled modes, where each group can be handled (i.e., routed) independently to other groups, while MIMO DSP techniques need to be applied only to the contents of each group (barring excessive crosstalk). This case can be considered as a combination of the cases depicted in Figure 4.1(B) and Figure 4.1(C). In addition, quite obviously any type of fiber, including MCFs and FMFs, can be deployed in bundles.

With respect to cost, while they offer the lowest ceiling in terms of potential capacity, bundles of SMFs have two significant advantages: *i*) in many cases they are already deployed, and *ii*) individual fibers can be lit one at a time to support new optical connections as traffic demand increases. MCFs and FMFs, while more expensive than SMFs, offer greater potential capacity density, as they can also be deployed in bundles but offer a larger number of spatial dimensions per fiber.

Lastly, observe that "advanced" fibers cannot directly interface with traditional optical switches. For MCFs, connections to switches can be realized 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, although this is generally only useful at the receiver.

4.1.2 SDM Transceivers

Modern optical transceivers are optimized for transmitting and receiving a single beam of coherent light. In the context of flexi-grid multiple beams, from different transceivers, 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 [2]. These structures, such as the one depicted in Figure 1.1(b) back in Chapter 1, are known in scientific literature as super-channels.

In order to efficiently exploit these structures, a new family of optical transceivers was devised, using a device known as an optical "comb" [61] to filter and split the beam coming from a single wide-band laser into a number of independently modulated signals. In this manner, a single laser/oscillator source can be shared to generate/receive an entire super-channel. These signals, which are constrained to be relatively close in frequency, can be used to generate a single large super-channel or multiple smaller ones, by placing appropriate guard-bands using e.g. filters or micro-ring resonators.

The introduction of SDM has spurred researchers to devise integrated transceivers to efficiently exploit the spatial domain without escalating costs. Another family of integrated transceivers for SDM has been proposed, based on the use of passive optical splitters, which separate a single laser beam into two or more beams at the same frequency (i.e., position in the spectral domain), each with a fraction of the original power, which can then be independently modulated [62].

4.1.3 SDM Switches and Switching Paradigms

The introduction of SDM has led to several switching paradigms for SDM being explored [63], [64].

The most significant new requirement posed by SDM is that ROADMs need to be able to interconnect a much larger number of incoming fibers, specifically an increase proportional to the count of spatial dimensions after fan-outs or mode lamps (which allow the interconnection between SDM fibers and bundles of standard fibers) are taken into account. Additionally, the spatial coupling experienced by FMFs requires all spatial modes to be jointly routed, and excessive crosstalk between cores in MCFs may do the same.

Three main approaches to SDM switching have thus emerged:

- **Independent Switching** (InS), where each spatial dimension can be switched independently from the others. This is the most flexible approach, but it requires an extremely high number of switch ports, which has a large impact on cost. Furthermore, it is only applicable to SDM networks employing SMFs or low-crosstalk MCFs.

- **Joint Switching** (JoS) is a much more cost-effective solution for networks using FMFs, where space is a degenerate dimension which must be jointly switched, since it involves switching signals at a given frequency on all spatial dimensions simultaneously. This can be done in a very cost-effective manner, as it functionally requires the same switching capabilities of a traditional switch

and is not significantly more costly than it [63]. However, this also greatly reduces flexibility in assigning resources, as it basically forces the use of spatial super-channels.

- **Fraction Joint Switching** (FJoS) is an intermediate solution (with respect to both cost and flexibility) between these two extremes, where spatial dimensions are mapped into spatial groups, and each member of a given group is switched jointly with the others of the same group, while groups are switched independently from each other. InS and JoS can be seen as degenerate cases of FJoS, where the groups size is uniformly 1 or the total number of spatial dimensions, respectively.

Technically, a variation of JoS which involves switching entire spatial dimensions (i.e., fibers, cores or modes) is also possible [63], but given the much smaller number of spatial dimensions compared to spectral channels, this would allow only a very limited number of (very capacious) routes, and hence is only suitable for relatively small mesh topologies.

Additionally, each of these approaches may or may not implement what is known in literature as **lane changes** [63], i.e., the ability of one input spatial dimension of index x to be routed towards one of many output spatial dimensions of index $y \neq x$. Another way to look at this property is to state whether a ROADM enforces what could be called the **space continuity constraint** (i.e., does not implement lane changes) or not.

Observe that these switching constraints are not intrinsically linked to SDM; indeed, they also model the restrictions imposed by some of the other possible technologies being investigated to increase the capacity of future transport networks discussed in Section 2.1: Raman amplification leads to a simple independently-switched network (but requires a large investment in terms of additional optical transceivers), while overlapping signals at different power levels leads to a jointly-switched network, as do OAM-multiplexed channels (which, in a way, can be seen as a kind of "non-overlapping spatial modes").

As partly investigated in this thesis (Section 4.5), the choice of switching paradigm has a major impact on ROADM architecture, both in terms of the number of elements and their complexity [63], [65]. While it is possible to implement InS using large numbers of commercially available WSSs with relatively low port count (e.g. 1×9), JoS requires higher port count WSSs configured to operate as $S \times (M \times N)$, i.e., M input ports, each carrying S spatial dimensions, directed toward N output ports [66]. Therefore realizing, for example, a 4-degree ROADM, with 9 spatial dimensions, would require using WSSs with a minimum port count of $9 \times (1 \times 4)$, i.e., 45 ports. Given this premise, in order for JoS to make economic sense, the increase in unitary cost associated with higher port count WSSs would need to be counteracted by a significantly smaller number of WSSs required in the ROADM architecture, as is in fact the case [67]. Taking into account that, historically, the per-port cost of single-mode WSSs has been inversely proportional to the port count [68], telecom operators can be expected to derive substantial benefits by exploiting (F)JoS in a SDM network scenario. An alternative approach, based on "Architecture on Demand" (AoD) for SDM rather than ROADMs, is also discussed in literature [69], however no hard comparisons have been made (to the best of the author's knowledge) regarding the relative cost and merits of AoD and (F)JoS solutions.

4.1.4 SDM Amplifiers

While it is technically feasible to individually amplify each spatial dimension, e.g. by using fan-out/in devices or mode lamps to separate and re-aggregate them in conjunction with multiple parallel OLAs, it would be exceedingly expensive. Fortunately, devices that can jointly amplify all spatial dimensions in a SDM link are already being developed, e.g. multi-core EDFAs [70].

4.1.5 SDM Network Modeling

Once abstracted from implementation details, currently used optical network models, such as those used by GMPLS [71], [72], model an optical network as a relatively simple graph. Using this abstraction, nodes may represent either whole ROADMs and optionally EDFAs and transceivers, or present an "exploded" (technically "disaggregated") view of the internal components of each node, while edges represent fiber links between and in the second case within ROADMs. When the opaque node representation is used, internal node constraints, i.e., which ports can connect to which other ports, are modeled with an object called a **Switching Matrix**. Another object, called the **Connectivity Matrix**, instead records which of the possible cross-connections are active, and for which frequencies.

Spectral resources (WDM channels, Flexi-grid spectral slices) are represented as a property of edges (or, sometimes, of "port" objects, which interface edges and nodes), more specifically binary arrays where indexes refer to specific resource and the binary value represents its state (i.e. "free" or "in-use"). Alternatively, some old works in literature dealing with computing RWA solutions using Auxiliary Graphs (e.g. [73]) model spectral resources as individual links between nodes, each representing a single WDM channel (to the best of the author's knowledge no such representation has been published for Flexi-grid networks).

Unfortunately, such models are ill-suited to represent SDM networks, as they have no provisions to model the relations between spatial dimensions in SDM links, nor can they codify the routing constraints mandated by certain SDM ROADM designs outlined in Section 4.1.3. They may, however, be extended/retrofitted to represent them, as outlined in Section 4.2.1.

4.1.6 SDM Resource Allocation

As better detailed in Section 4.3, in the context of SDM the RWA/RSA problem requires a further revision due to the presence of the space dimension.

Among the early works on this topic, the authors of [74] described an ILP formulation for the resource allocation problem of SDM networks employing MCFs, and takes into account various parameters modeling the undesired interaction between adjacent cores. Ref. [75] (later expanded in [76]) proposed 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. [77] proposed 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, [78] provided a heuristic policy for assigning modes in a network employing FMFs, based on assigning wavelengths (the work is based on fixed-grid WDM) 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 proposed interesting policies that addressed particular cases, none address the more general problem of resource allocation in SDM networks, and none put much emphasis on the important factor of cost of the proposed SDM solution.

Since then, in addition to a number of works co-authored by the author of this thesis, on whose contents this work is largely based ([65], [79]–[85]), there has been a veritable explosion in the number of works published on the topic of SDM resource allocation, a comprehensive overview of which can be found in a recently published survey paper [86].

4.2 A Network Model for SDM

As outlined in Section 4.1.5, current network models [71], [72] are not suitable to represent the complexities of SDM networks, namely the constraints deriving from certain types of SDM fibers, transceivers and ROADMs.

Taking inspiration from the layered structure of recent models being discussed within the IETF (e.g. [87] and related drafts), a SDM optical network could be modeled as what could be called an "augmented directed graph with ports", formally $G = (V, P, node, ports, E, portTrees, spectrum, switchingMode)$, where:

- V is the set of vertexes, each $v \in V$ modeling an individual node, be it a transceiver or ROADM;
- P is the set of ports, each $p \in P$ modeling an optical port in a ROADM or transceiver;
- $node : P \to V$ is a function mapping each port to its owning node;
- $ports : v \in V \to P_v \subseteq P$ is a function mapping a node to its set of ports, and $\bigcup_{v \in V} P_v = P$;
- E is the set of edges, each $e \in E$ modeling some kind of spatial dimension linking two ports, and $e := (p', p'')$ where p' is the source and p'' is the destination port of edge e; note that $node(p') = node(p'')$ is admissible, i.e., edges can be self loops (a.k.a. intra-node links);
- $portTrees : P_v \to T_v$ is a function mapping all the ports of each node into a set of tree graphs T_v, such that each $t_v \in T_v$ is a tree $t_v = (V_t, E_t)$ rooted in some port p of node V such that:
  - $\bigcup t_v = T_v$;
  - $\forall t'_v, t''_v \in T_v, V'_t \cap V''_t = \emptyset$ and $E'_t \cap E''_t = \emptyset$, i.e., all t_vs are disjoint;
  - $\forall t_v \in T_v, E \cap E_t = \emptyset$, i.e., tree links are **not** links in the overall graph;
  - $\forall (p', p'') \in E$, either both are roots of some t'_v, t''_v, both are leaves, or both are intermediate nodes in the trees;
- $spectrum : \{p \mid p \text{ is a leaf of some } t_v\} \to \mathbb{B}^n$ is a function associating ports that are leaves to a binary array representing the availability of spectral resources in that port.
- $switchingMode : P \to \mathbb{B}$ is a function mapping ports to a binary label, specifying, for each port, whether it's child ports, if any, can be switched independently or must be jointly switched.

More intuitively, the network graph consists of nodes (ROADMs, transceivers) where ports model the properties of individual spatial dimension i.e., the SMF carrying a light from a specific core or mode of light after it was separated using a mode lamp or fan-out device). Edges connect ports (an alternative view could be consider ports as vertexes and nodes as disjoint sets of ports), and can be both intra- and inter-node. Intra-node links model how spatial dimensions can be routed within a node, which is similar to the connectivity matrix structure used by current models, while inter-node links model connections between remote ports.

With just these elements, the graph would show a large number of parallel links between (and within) nodes, but with no clear structure among them. In order to formalize it, a number of disjoint trees are used to organize the set of ports of each node. Each root port, which we call an **aggregation port**, models a SDM inter-node link point of attachment, while leaf ports, which we dub **switching ports**, model a representation of the spectral resources (e.g. flexi-grid spectral slices or WDM channels) of a single spatial dimension. Intermediate layers of ports, if any, can be used to complex model fiber and switching constraints, via the switchingMode parameter, which specifies whether immediate child ports are to be switched together (e.g. because the associated spatial dimensions are strongly coupled, or because they are attached to a JoS or FJoS ROADM) or can be switched independently. A similar approach can also be used to separately model ROADM constraints and a interactions between spatial dimensions (crosstalk, coupling, etc.). Indeed, by simply adding any other label function $f : p \rightarrow L$ mapping ports to a desired label space L, any additional per-port parameter can be recorded (observe that labeling functions can also be defined for objects other than ports, if desired). In practice, this corresponds to adding an appropriate label field to the port (or whatever else) object.

Using this structure, inter-node edges connecting aggregation ports represent whole SDM links, and allow to easily compute inter-node routing using standard, well-known routing algorithms, such as Dijkstra's or Yen's, without requiring unwieldy kludges. Notably, modeling this hierarchy of ports allows the seamless representation of any type of SDM link (based on SMFs, MCFs, FMFs, etc.) and ROADM (implementing JoS, FJoS, etc.). WSSs (functionally represented by some port objects) can be modeled as either "independent" or "joint", based on the switchingMode of their single input (output) port, which, in the case of Joint WSS, must be an aggregation port representing multiple lower ports in the hierarchy. Using this model, the connectivity switching matrix of a ROADM can be derived by relating its highest-level (i.e., closer to the root) "joint" ports in the hierarchy, if any, or by directly relating the switching ports, in the case of "traditional" independent switching. Some of the constraints in/due to SDM transceivers can partly be modeled by the hierarchy of ports (e.g. for splitter-based ones), but the model does not directly encode in the graph structure constraints related to frequency or bandwidth. While there exist in literature examples on how to solve this for WDM channels [73], this is an open problem in the context of Flexi-grid (and hence, SDM). In our implementation we solved this via simple labeling of transceivers and ad-hoc handling of such labels during resource allocation.

For example, Figure 4.2 depicts a simplified one-directional, two degree piece of a ROADM, where a single bundle of 4 fibers is attached to two Joint Switching WSSs in groups of two, and each of these groups can then be switched independently from the other to exit the node in either the first two or latter two fibers (more degrees were omitted in the interest of clarity). This is a Fractional Joint Switching ROADM with lane changes [63]. The first layer of aggregation ports (a) models the bundle and has independent children; the second (b) models the logical ports of the JoS WSSs

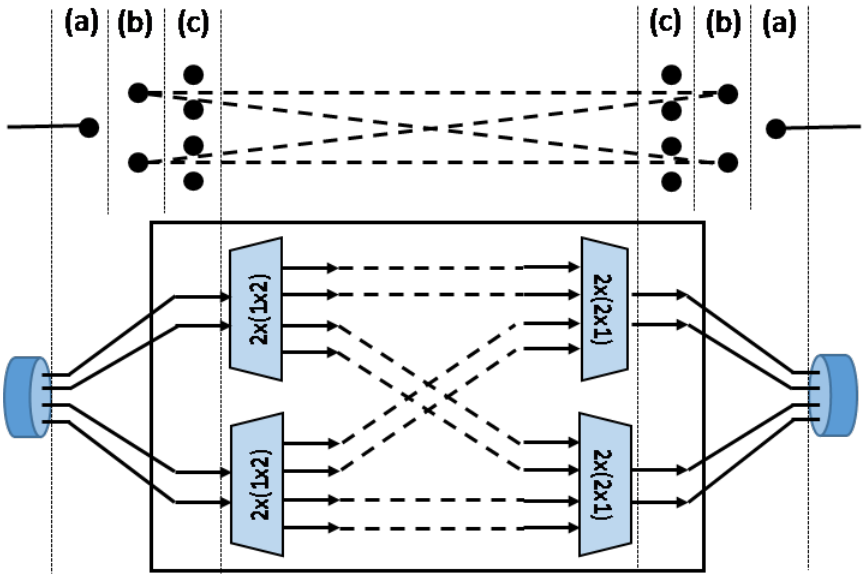

FIGURE 4.2: Example modeling of a one-directional, two-degree Fractional Joint Switching SDM "ROADM" with lane changes; aggregation ports (a) model points of attachment for inter-node SDM links, switching ports (c) model the spectra of each spatial dimension, while intermediate ports (b) model the Fractional Joint Switching constraint.

and therefore has joint children, while the individual switching ports (c) model the spectra of each fiber. Observe how the internal connectivity matrix is derived using the highest-placed joint ports in the hierarchy, i.e., layer (b), while the external link is mapped by the top level aggregation port (a). SDM transceivers with multiple outputs and/or inputs can be encoded in a similar manner.

An implementation of this model was presented at the 2016 European Conference on Optical Communication [82], as part of a larger work describing an experimental control plane for SDM networks, more on which in Section 4.8, and later expanded upon in an article published in the IEEE/OSA Journal of Lightwave Technology [85].

4.2.1 Retrofitting existing models for use with SDM

While existing models [71], [72] are not designed to incorporate an equivalent structure to the one proposed in Section 4.2, they could be extended to do so, e.g. by introducing an ad-hoc port labeling mechanism (i.e., a new Type-Length Value). Furthermore, should it be needed, such a structure may be retroactively superimposed on such models without changing the actual data model, by overloading the semantics of some other field, like the aforementioned port labels or even unrelated fileds, e.g. Shared Risk Link Groups (SRLGs).

Under such a scenario, individual fibers, cores or modes of light could be represented as separate links between nodes; then, to relate them, labels (SRLGs or something else) would need to be used to group them, mimicking the same structure described above. Additionally, by appropriately partitioning this label space, it would be possible to convey meanings such as "to be jointly switched", "MCF" or "FMF", therefore achieving the needed level of detail. Intra-node links would then

connect these equivalents of switching ports, modeling the presence or absence of lane changes, as done before.

Using this approach, the example in Figure 4.2 would only exhibit the (c) layer ports (i.e., 4 inter-node links per direction). The two topmost ports would be tagged with "joint group 1", the bottom ones as "joint group 2", and all of them marked with "bundle 1" and "bundle 2", for the incoming and outgoing ports, respectively. The downside of this approach is that some tweaks would be needed to the inter-node routing and resource allocation algorithms, so that bundles are treated as a single link for routing purposes and switching constraints are honored.

4.2.2 Modeling SDM Super-Channels

In the context of WDM the only available degree of freedom after a route is chosen is the channel selection, i.e., the choice of which spectral resources to assign to the lightpath. As outlined earlier, 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 12.5 GHz, 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 spectrally efficient **"spectral" super-channels**, such as the one depicted in Figure 4.3 (top left).

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, output modes of SDM transceivers and SDM switching paradigms outlined earlier.

The most natural way to represent this new dimension is to turn the usual spectrum array, where each index represents a channel (WDM) or spectral slice (Flexi-grid), into a two-dimensional matrix, where each row represents the spectral resources of an individual spatial dimension.

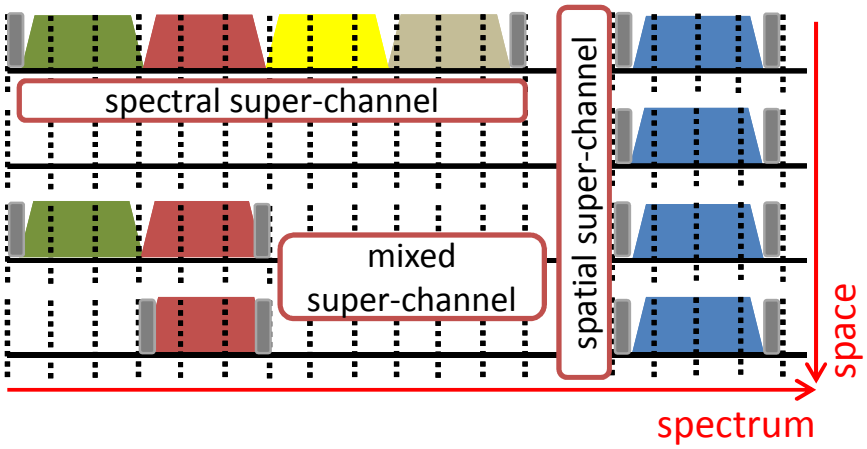

FIGURE 4.3: Spatial vs. Spectral vs. Mixed Super-Channels. Spectral super-channels are more spectrally efficient, avoiding switching guard-bands (gray squares) between individual optical carriers, while spatial super-channels can be jointly switched but require more switching guard-bands.

Given this representation, and the capabilities of certain SDM devices, such as splitter-based transceivers and JoS ROADMs, it is only natural to envision **"spatial" super-channels** spanning multiple spatial dimensions, such as the one depicted in Figure 4.3(right). 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.

Combining the two technologies (optical combs and splitters) it is also possible to envision **mixed "spectral-spatial" super-channels** (Figure 4.3(bottom left)), where multiple spatial super-channels are placed at the Nyquist condition creating a spectral super-channel extending over multiple spatial dimensions.

As further motivation for the need of multiple super-channel shapes due to coupling considerations between spatial dimensions, consider the example of Figure 4.4: for independent fibers and switches, such as bundles of standard fibers, or MCFs with large spacings between cores, super-channels may be placed freely with respect to space (Figure 4.4(A)), and, potentially, even extend over non-contiguous spatial dimensions. For fibers exhibiting low (but non-negligible) or localized strong cou-

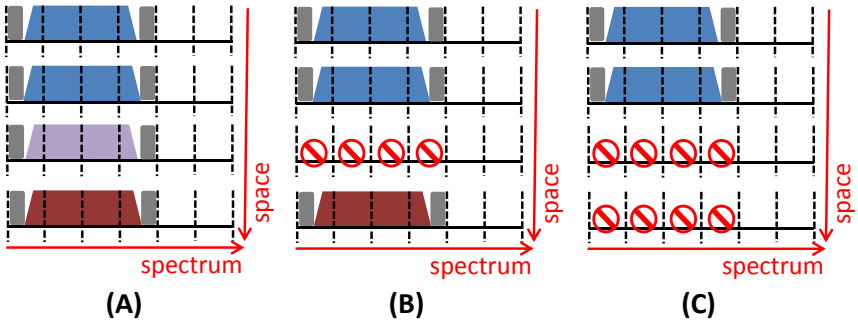

FIGURE 4.4: Feasible spatial allocations with no coupling/crosstalk (A, three super-channels), low or localized (e.g. inter-core) coupling/crosstalk (B, two super-channels) and high coupling (C, one super-channel).

pling/crosstalk (such as certain geometric arrangements of MCFs), a sort of "spatial guard-band" may be needed to isolate different services co-propagating on different spatial dimensions at the same frequency (Figure 4.4(B)). 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 (Figure 4.4(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.

4.3 The Routing, Space and Spectrum Allocation Problem

Given the additional degree of freedom offered by SDM, which transforms the unidimensional spectral resource array into a bidimensional matrix, and may present a number of additional constraints due to restrictions in transceivers, coupling characteristics of fibers and switching paradigms in ROADMs (as modeled, for example, in Section 4.2), in this context the resource allocation problem becomes the "**Routing, Space and Spectrum Allocation problem**" (RSSA), sometimes also called the **Routing, Modulation Level, Space and Spectrum Allocation** (RMLSSA) problem.

Like for its predecessors RWA and R(ML)SA, the final objective of R(ML)SSA is to produce, given a network state and a set containing one or more new connections, the configuration choices needed to place the connection(s) in the network. More formally, given a representation of the SDM network, such as the one described in

Section 4.2, and a set of one or more demands $d = (s, d, b)$, where s and d are the source and destination node (transceiver, or ROADM if transceivers are not modeled), respectively, and b is a measure of the required capacity, often improperly called "bandwidth", of the connection (or, alternatively, a structure describing the optical properties of the desired lightpath, e.g. for handling alien wavelengths originating and possibly terminating outside the network).

Solving the R(ML)SSA problem thus involves solving, individually for each new connection (or jointly for all of them), several sub-problems:

1. **Routing**: finding a route between s and d; by using the proposed model, this is no more difficult than for WDM and flexi-grid network, since the complex SDM links are collapsed to form a "traditional" graph perfectly compatible with well-known routing algorithms, such as Dijkstra's [6] and Yen's [7]. Note that this does not yet specify which spatial dimensions to use (see step 3), but does correctly force all of them to be co-routed (thus minimizing DGD issues).

2. (Optionally) **Modulation Level, Symbol Rate and FEC selection**: once the overall route is known, several transceiver parameters can be optimized (for transceivers supporting this function), which in turn determine the optical reach and capacity of its output(s). Generally speaking, the combination of parameters offering the highest data rate and enough OSNR to traverse the chosen route is to be preferred, but multiple choices may need to be tested, especially for networks where PLI strongly depend on which spatial dimensions are used (e.g. employing FMFs or weakly-isolated MCFs). In addition, since this step determines the capacity that can be squeezed from an optical transceiver, it enables the computation of the number of optical transceivers needed to support the demand (which may, in some cases, be > 1).

3. **Super-Channel shape, Space and Spectrum allocation**: depending on the capabilities of the selected components and characteristics of the underlying network components implementing the selected route, as well as operator policies, several choices are possible with regards to the shape of the super-channel for a new connection. Once a shape is chosen, the spectral resources needed to implement it must be reserved (and, hence, checked for availability) on an adequate number of spatial dimensions. Furthermore, in the case of (F)JoS resources, any additional spectral resource that needs to be co-routed with the chosen ones must be included in the super-channel.

   Also note that, while the choice if spectral resources is fixed for all links in the path (barring the presence of OEO signal regenerators), if lane changes are admitted, then the assignment of spatial resources may, in principle, change on a per-link basis.

4. (Optionally) **Feasibility evaluation**: if a detailed enough model of how optical impairments affect the OSNR at the receiver is available (such as the Gaussian Noise Model for WDM/Flexi-grid [8], [88]), and the related information concerning PLIs encoded in the model e.g. using [89], then it is possible to verify whether the chosen combination of parameters would enable meaningful communication without wasting time on unfruitful connection attempts. Observe that, unlike for RWA/RSA where there is very little correlation between experienced PLIs and the choice of frequency (what is there is mainly due to other active connections, and is accounted for in [8]) and where this step could be performed right after step 2, in the context of SDM this process has to be

performed last, since there may be a strong dependency of PLIs on the choice of spatial dimension.

Given the additional complexity of RSSA compared to RWA/RSA, and given the fact that formal formulations of these problems, generally as ILPs, are of largely academic interest, being of little relevance for real-world sized networks, a choice was made to directly pursue heuristic algorithms to solve RSSA, as discussed in the rest of this chapter.

4.4 Basic RSSA Heuristics

Real optical networks rarely employ the complex RWA/RSA algorithms that can be found in literature. In fact, they generally solve routing using some variation of Dijkstra's [6] Shortest Path (SP) algorithm, or Yen's [7] K-Shortest Paths (KSP) algorithm, and Wavelength Assignment/Spectrum Allocation using the very simple (yet effective) **First-Fit** (**FF**) heuristic, which takes as input the chosen path(s), network model, and, for Flexi-grid, size (in slices) of the frequency slot, and iterates on the chosen paths, selecting the first (i.e., with lowest index) WDM channel or sufficiently large range of adjacent frequency slices (i.e., frequency slot) that is free on all links of that path. The pseudo-code for FF is reported, for completeness, in Algorithm 1 (for both WDM and Flexi-grid).

Algorithm 1 First-Fit (FF) for RWA/RSA.

```
1: for each path (in the set of K shortest path from source to destination) do
2:     Compute the ⋂ of available WDM channels OR spectral slices in the path
3:     for each WDM channel OR spectral slice in the path do
4:         if channel is free or sufficient contiguous slices are free then
5:             return channel or first slice index
6:         end if
7:     end for
8: end for
9: return failure
```

Note that returning a failure state equates to blocking that particular connection demand, while returning a valid channel or frequency slot would result in those resources being reserved to the demand for its lifetime.

In light of the effectiveness of FF, it is only natural to think to expand it to support SDM networks, which was done in a work presented at the 2015 International Conference on Optical Network Design and Modeling [79] (although note that the analysis in that work does not take the existence of optical combs into account), and presented in this section. In this setting FF can take multiple flavors, depending on the shape of super-channel it is trying to allocate, on the coupling constraints of the underlying fibers, and on the switching constraints of the ROADMs.

Spectrum-First Allocation

The **Spectrum-First** (**SpeF**) heuristic extends FF by implementing 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 slot (i.e., a continuous range of spectral slices large enough to accommodate all needed carriers) on the spectrum of the first spatial dimension, from left to right (i.e., lowest to highest index), moving to higher spatial dimensions if no suitable void can be found on the current one.

The pseudo-code for SpeF is reported in Algorithm 2. Observe that, contrary to the name, to achieve preference for spectral super-channels the outer (i.e., first) loop must iterate over the space dimension.

Algorithm 2 Spectrum-First (SpeF) first fit allocation for independent spatial dimensions without lane changes.

```
1: for each path (in the set of K shortest path from source to destination) do
2:     Compute the ⋂ of available spectral slices for each spatial dimension in the path
3:     for each spatial dimension in the path do
4:         for each spectral slice in the dimension do
5:             if sufficient contiguous slices are free then
6:                 return return current spatial dimension and first slice index
7:             end if
8:         end for
9:     end for
10: end for
11: return failure
```

Also note that Algorithm 2 is only suitable for networks employing independent spatial dimensions (bundles of SMFs and well-isolated MCFs) and InS switches, and ignores the characteristics of transceivers.

Space-First Allocation

The **Space-First** (**SpaF**) heuristic extends FF to implement 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 carrier.

In practice, it iterates over all spatial dimensions in ascending order, looking for an available frequency slot large enough to host one carrier. When one such slot 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, i.e., a spatial super-channel may span spatial dimensions 0, 2 and 3, skipping over 2), as show in Algorithm 3.

Note that Algorithm 3 assumes SDM networks with negligible coupling, and no switching constraints (i.e., InS). 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 an electronic grooming mechanism like the one discussed in [78] is employed).

Degenerate-Space-First Allocation

The **Degenerate-Space-First** (**DSpaF**) is a variant of SpaF suited to networks with strongly coupled links (typically based on FMFs) and/or JoS switches, where space

Algorithm 3 Space-First (SpaF) first fit allocation for independent spatial dimensions without lane changes.

```
for each path (in the set of K shortest path from source to destination) do
    D ← ∅
    c ← 0
    Compute the ⋂ of available spectral slices for each spat. dim. in the path
    for each spectral slice in the first spatial dimension of the path do
        for each spatial dimension in the path do
            if sufficient contiguous slices for one carrier are free on the spatial dimension then
                D ← D ∪ {current spatial dimension}
                c ← c + 1
                if c = number of required carriers then
                    return D and first slice index
                end if
            end if
        end for
        D ← ∅
        c ← 0
    end for
end for
```

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.

Algorithm 4 Degenerate Space-First (DSpaF) first fit allocation for strongly coupled spatial dimensions.

```
for each path (in the set of K shortest path from source to destination) do
    Compute the ⋂ of available spectral slices for each spat. dim. in the path
    for each spectral slice in the first dimension do
        if sufficient contiguous slices for one carrier are free then
            return all spatial dimensions, and first slice index
        end if
    end for
end for
return failure
```

In practice it works as SpaF does, but only checks the availability of the initial frequency range on the first spatial dimension, as shown in Algorithm 4

Note that in this degenerate case the presence or absence of lane changes is largely irrelevant without a model that relates individual spatial dimensions to PLIs. In fact, the actual logic behind DSpaF is really not dissimilar to that of FF, except that, in a sense, the "capacity" of each carrier is multiplied by the number of spatial dimensions.

Align-Strict Allocation

In a bid to facilitate the use of splitter-based transceivers, which are only suited to generate spatial super-channels, with the more spectrally efficient spectral super-channels, the **Align-Strict** (**AS**) heuristic assumes foreknowledge of the possible capacities of incoming demands (in order to pre-compute possible super-channels widths), and at least an approximate knowledge of their relative arrival distribution and organizes the spectral-spatial resources of the network to facilitate the exact spectral alignment (i.e., overlap) of super-channels on different spatial dimensions (so that individual carriers in different spatial dimensions may, in principle, be generated using the same splitter-based transceiver). Note that, despite being developed independently, this approach is very similar to the one discussed in [43].

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 splitter-based transceiver sharing than naïve SpeF. In practice, this re-use is further encouraged by placing the second, third and so forth 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.

The pseudo-code for AS (sans the pre-partitioning) is reported in Algorithm 5.

Algorithm 5 Align-Strict (AS) first fit allocation for for independent spatial dimensions without lane changes.

```
1: for each path (in the set of K shortest path from source to destination) do
2:     Compute the ⋂ of available spectral slices for each spatial dimension in the path
3:     for each super-channel-size region of the partition for demands of this size do
4:         for each spatial dimension in the path do
5:             if first spectral slice of the region is free then
6:                 return return current spatial dimension and slice index
7:             end if
8:         end for
9:     end for
10: end for
11: return failure
```

As for SpeF and Spaf, the pseudo-code presented does not take the possibility of lane changes into account (although it could be extended to do so). Another (tested) extension involves relaxing the alignment constraint to allow regions dedicated to large super-channels to host smaller ones under severe load, but this undermines the benefits of alignment and reduces the overall throughput (as larger spectral super-channels are inherently more spectrally efficient than smaller ones).

4.4.1 Simulative Comparison of basic RSSA heuristics

The performance of an hypothetical SDM network employing these basic RSSA heuristic with respect the ideal expected capacity increase of SDN (equal to the number of spatial dimensions) was evaluated, following common practices in the field, in dynamic traffic scenario using a purpose-built simulation tool.

Simulation Parameters

The simulation used 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 slices over the total), computed on the assumption that all connections are accepted on their shortest path.

The simulations were run using the well known Telefónica Spain's national network topology (Figure 4.5, [90]), comprising 30 nodes (average nodal degree: 3.7, maximum nodal degree: 5), of which only 14 are Add/Drop nodes, and 56 links (average length: 148 Km).

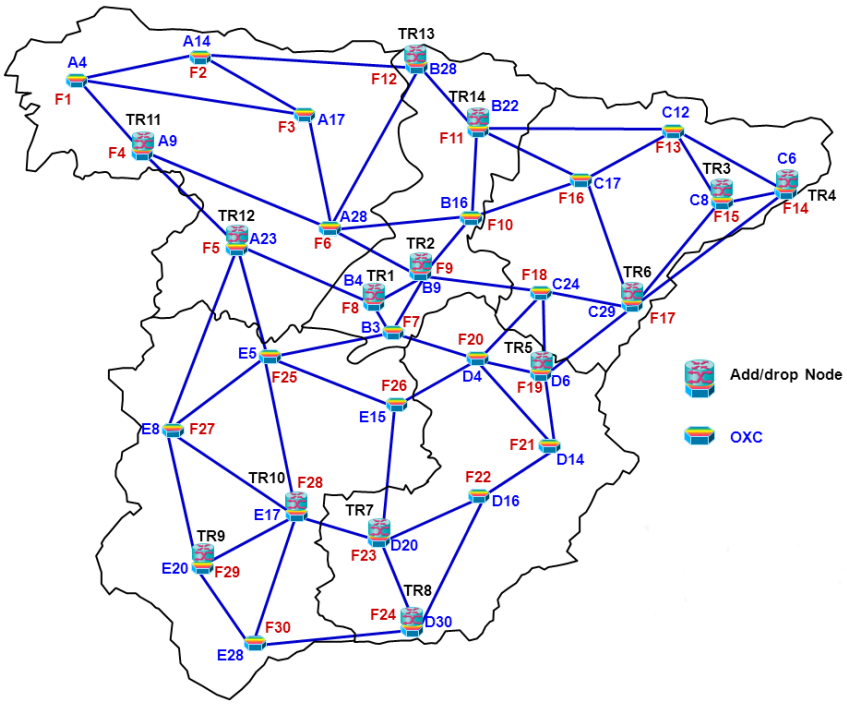

FIGURE 4.5: Telefónica Spain's national network topology.

The simulations used the following parameters:

- 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);
- 384 12.5 GHz spectral slices per spatial dimension (equivalent to a standard full 4.8 THz C-band fiber with 96 50 GHz WDM channels);
- 4 independent spatial dimensions;
- splitter-based transceivers with up to 4 outputs, each a 100 Gb/s carrier modulated using DP-QPSK, requiring 32 GHz of spectrum, and 9 GHz of switching guard-band on each side of each spectral super-channel; the transceivers

are assumed to be able to serve multiple connections (i.e., multiple lighpaths to potentially different destinations) at the same time (similarly to so-called "sliceable" transceivers, see e.g. [91]);

- considered all 30 nodes as potential ROADMs.

Bandwidth demands were randomly sampled, using an uniform distribution, from the set of $\{100, 200, 300, 400\}$ Gb/s, which, using the modulation format described above and standard baud rates and FEC, 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. No restriction was put on the number of total transceivers, nor on the number of transceivers per node. Each experiment simulates 10^6 bidirectional connection requests (plus an initial 10^4 requests not accounted for in the collected statistics, but used to reach an initial network steady state).

The proposed heuristics 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 average **Network Throughput**, and the average **Number of Transceivers** needed to support them.

As a benchmark a fictional heuristic was used, based on a 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 slices equal to the sum of the those available to the other RSSA policies (but spread on all spatial dimensions). This was done to highlight the effect that the limitations posed by SDM have on the resulting capacity increase.

Simulation Results

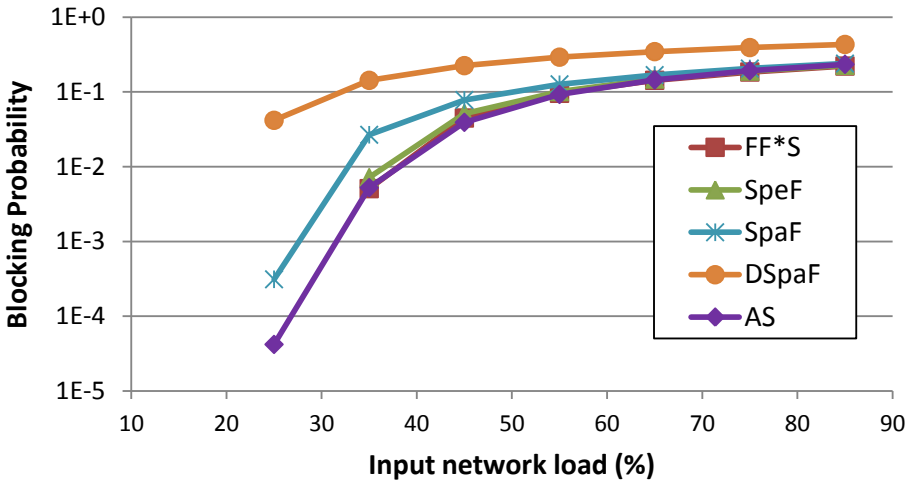

FIGURE 4.6: Blocking Probability (BP) vs. Input Network Load for the basic RSSA heuristics; no blocking occurs below 25% load. Spectral super-channel based heuristics (SpeF, AS) exhibit the lowest BP under load, since they are more spectrally efficient.

The results for measured BP (in log scale) vs. Input Load are depicted in Figure 4.6; a similar result was computed using the Bandwidth Blocking Ratio (BBR) metric, i.e., the ratio between the total bandwidth (throughput) refused over that requested, but due to the log scale is virtually indistinguishable from that of Figure 4.6.

As can be expected, none of the proposed heuristics 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, on average, wasted. As it will be shown later, under more favorable conditions, the performances of DSpaFand, in general, of JoS-based SDM networks can in fact approach those of InS-based networks.

The performance difference between SpaF and SpeF can be explained taking into account that (*a*) SpaF, being space-oriented, is less spectrally efficient than SpeF and AS, and (*b*) since the simulations consider only 4 spatial dimensions, all of which are needed for the largest spatial super-channels, a larger proportion of large, more efficient requests is denied.

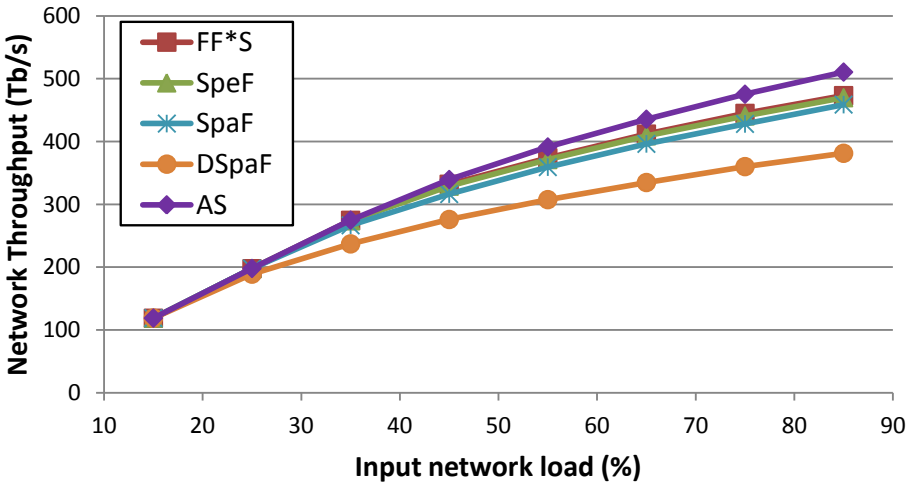

FIGURE 4.7: Total Network Throughput (TP) vs. Input Network Load for the proposed RSSA heuristics. AS achieves the highest total network throughput, thanks to a distribution of accepted demand sizes less biased against larger, more efficient super-channels than the other heuristics.

Figure 4.7 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 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, i.e., 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, tend to exhibit a significantly higher blocking rate for larger, more efficient super-channels than for smaller ones, thus leading to lower spectral efficiency. Note that this behavior would apply even in the context of flexi-grid networks, but is amplified in SDM ones.

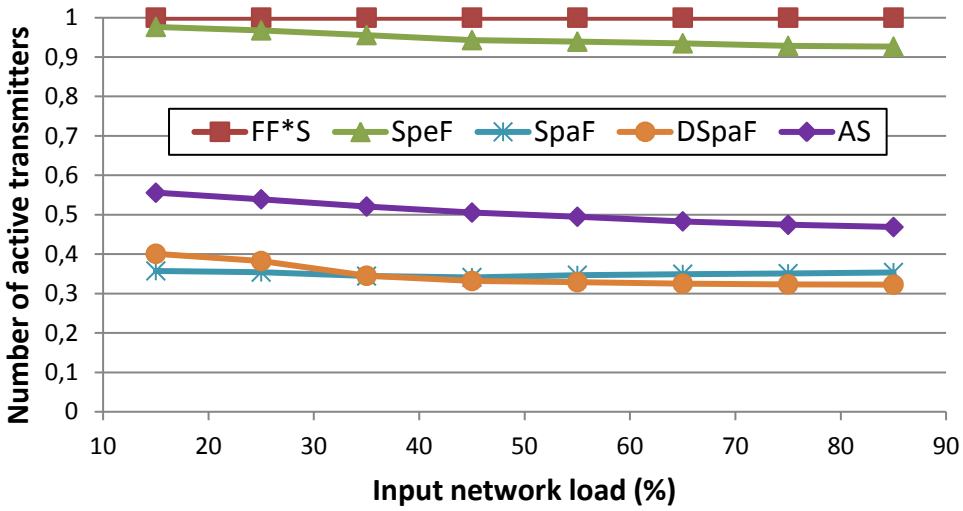

FIGURE 4.8: Average number of active splitter-based transceivers (normalized w.r.t. FF*S) vs. Input Network Load for the proposed RSSA heuristics. Spatial super-channel based heuristics rely on significantly less transceivers than those based on spectral super-channel, except for AS which allows significant transceiver sharing.

Finally, Figure 4.8 shows the average number of active transceivers with respect to the input network load for all policies. All measures are normalized with respect to FF*S (the worst case); for reference, the number of transceivers used by FF*S at 15% load is more than 2300. While such a number is far higher than those of current deployments, it is somewhat misleading, since it is computed on the assumption of using single-output transceivers (since there is only one spatial dimension) with a vastly more capacious fictitious network, where the available spectrum is S (4) times the real one (observe that, in Figure 4.7, at 15% load the simulated network already carries more than 100 Tb/s). Therefore it is more interesting to focus on the ranking between the heuristics rather than the absolute numbers of active transceivers.

Two clusters clearly emerge: the first one comprises FF*S and SpeF, both of which use a much larger number of transceivers 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 source but on different spatial dimensions randomly happen to share the same central frequency, then they can be generated by the same transceiver.

The second cluster, comprising SpaF, DSpaF and AS, uses only a fraction of the transceivers 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 transceivers 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 transceivers needed by SpeF, while using less than 55% more than those needed by SpaF, achieving a sort of balance between hardware and spectral efficiency.

4.5 RMLSSA under InS, JoS and FJoS Paradigms

As outlined in the previous section, different approaches to solving the R(ML)SSA problem, i.e., favoring spatial or spectral super-channels, can have significant impact in the resulting performance (total throughput) and cost of a SDM network.

Furthermore, being able to use more aggressive modulation formats or, in the case of spectral super-channels, sub-channel spacings, which sacrifice optical reach in exchange for higher spectral efficiency (bit/s/Hz) can bring significant benefits.

To incorporate the Modulation Selection step for SpeF and SpaF (Algorithms 2 and 3, turning them into RMLSSA heuristics, they can be modified by evaluating the total length of a candidate path and using a PLI model, such as the GN model of nonlinear interference in coherent (Nyquist) WDM systems [8] used to estimate the receiver OSNR on that path, and selecting the most efficient modulation/spacing that produces a sufficient OSNR for correct reception (fixing some other parameter, such as baud rate).

Alternatively, these computations may be performed offline, and, assuming uniform fiber-induced PLI over the whole network, used to produce a simple table relating modulation or sub-channel spacing with the maximum distance optically traversable by that signal, such as Table 4.1.

TABLE 4.1: Optical reach achievable with DP-8QAM modulated spectral super-channels at 32 GBaud using different sub-channel spacings with a granularity of 3.125 GHz [92], as per the GN model of nonlinear interference [8].

Sub-Channel Spacing [GHz]	Reach [Km]
50	1700
37.5	1200
34.375	800

In addition, SpaF can be modified to consider the size of the underlying spatial groups, as dictated by either the links or ports (as detailed in the model, explained in Section 4.2), by explicitly iterating over the spatial groups, as shown in Algorithm 6. Note that in this manner JoS (i.e., DSpaF) is actually a degenerate case of SpaF where the spatial group size is equal to the number of spatial dimensions, whereas InS is the case where the groups are of size 1. FJoS covers all intermediate cases, including those where groups are of unequal size.

While JoS and FJoS are quite obviously more constraining, and thus potentially less efficient, than InS, their use is justified by the significant savings achievable in the construction of JoS and FJoS based ROADMs compared to InS ones. In general, for a Colorless-Directionless ROADM with $N_{A/D}$ Add/Drop degrees and N_L line degrees, and S spatial dimensions, possibly grouped in G-sized (on average) independent spatial groups, the number of required WSSs would follow the formula reported in Table 4.2.

Algorithm 6 Generalized Space-First (SpaF) first fit allocation for InS/JoS/FJoS SDM networks without lane changes.

```
for each path (in the set of K shortest path from source to destination) do
    D ← ∅
    c ← 0
    Compute the ⋂ of available spectral slices for each spat. dim. in the path
    for each spectral slice in the first spatial dimension of the path do
        for each spatial group G in the path do
            if sufficient contiguous slices for one carrier are free on the first spatial dimension of the group then
                D ← D ∪ G
                c ← c + |G|
                if c = number of required carriers then
                    return D and first slice index
                end if
            end if
        end for
        D ← ∅
        c ← 0
    end for
end for
return failure
```

TABLE 4.2: General formulas of WSS requirements for InS, JoS and FJoS CD ROADMs.

Switching Paradigm	General # of WSSs
InS	$(2 \cdot N_L + 4 \cdot N_{A/D}) \cdot S$
FJoS	$(2 \cdot N_L + 4 \cdot N_{A/D}) \cdot \lceil S/G \rceil$
JoS	$2 \cdot N_L + 4 \cdot N_{A/D}$

In essence, an additional number of WSS proportional to the number of spatial dimensions would be needed for InS, while for JoS the number would not change, and for FJoS the increase depends on the size (or, equivalently, number) of independent spatial groups.

The performances of SpeF and SpaF (as generalized by Algorithm 6) with respect to multiple sub-channel spacings (for SpeF) and switching constraints (for SpaF) where analyzed in a work presented at the 2015 European Conference on Optical Communication (ECOC) [65] and later expanded upon in an article published in the IEEE/OSA Journal of Lightwave Technology [81], and reported in the rest of this section.

4.5.1 Simulative Analysis of RMLSSA under InS, JoS and FJoS

Simulation Parameters

Parameters for this simulation used the same (modified) simulation tool and underlying topology as in Section 4.4, except for these differences:

- only the proper subset of nodes that have ROADM capabilities was used to originate and terminate demands;

- links were assumed to comprise a single bundle of 9 independent SMFs, i.e., no additional PLIs from multiple spatial dimensions were considered to influence the total reach of optical signals;
- demands were assumed to be uniformly distributed over a range of [1,9] sub-channels, with average capacity demand per connection equal to 960 Gbps (192 Gbps per sub-channel using the parameters from Table 4.1);
- each data point was obtained by simulating 10^5 bidirectional connection requests;
- unless otherwise stated, to minimize the impact of intermediate-node WSS filters, a 12.5-GHz guard band (GB) on both sides of spectral super-channels and of each sub-channel in spatial super-channels was assumed.

Simulation Results

Independent Switching

While the simulations described in Section 4.4.1 already quantify the performance differences between InS and JoS in the case of independent Switching, those results, although likely more general, are limited to a single network scenario. A first investigation involved studying the BP performance under InS in order to highlight the influence of the sub-channel spacing. To do so, several scenarios were investigated:

- SpeF using variable spacing adapted to the path length (SpeF-Var);
- SpeF using fixed spacing (34.375, 37.5 GHz and 50 GHz) with 12.5 GHz GBs on both sides of each spectral super-channel;
- both SpeF and SpaF under fixed-grid WDM conditions, with 50 GHz channel spacing including GBs (SpeF-WDM and SpaF-WDM).

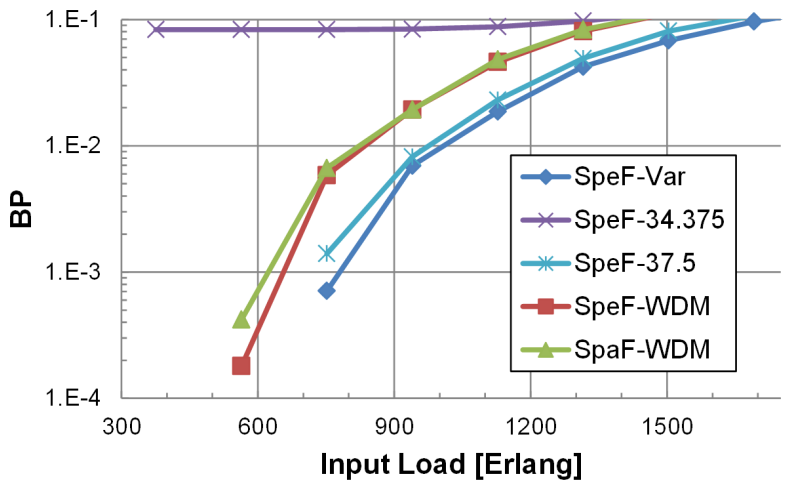

FIGURE 4.9: Blocking Probability vs. Input load under Independent Switching for several sub-channel spacings.

The measured BP vs. the input load (expressed in terms of average active demands, or Erlangs) is depicted in Figure 4.9. Thanks to its superior spectral efficiency in flexi-grid conditions, SpeF is still, in general, superior to SpaF, blocking up

to an order of magnitude fewer connections than it when it can utilize its superior efficiency. That is, however, not the case for SpeF with a 34.375 GHz sub-channel spacing, which presents a horizontal asymptote at 10% BP, because the maximum achievable reach with that spacing is not sufficient to support the establishment of all possible connections.

In the case of WDM-spaced sub-channels, where the efficiency of SpeF and SpaF coincides, both exhibit very similar BP, with a slight edge for SpeF due to the fact that the size (in terms of units) of the spectral dimension is much greater that that of the spatial one (i.e., there are far more spectral slices than spatia dimensions, 384 to 9), therefore large demands start to be starved of resources much sooner using SpaF than SpeF.

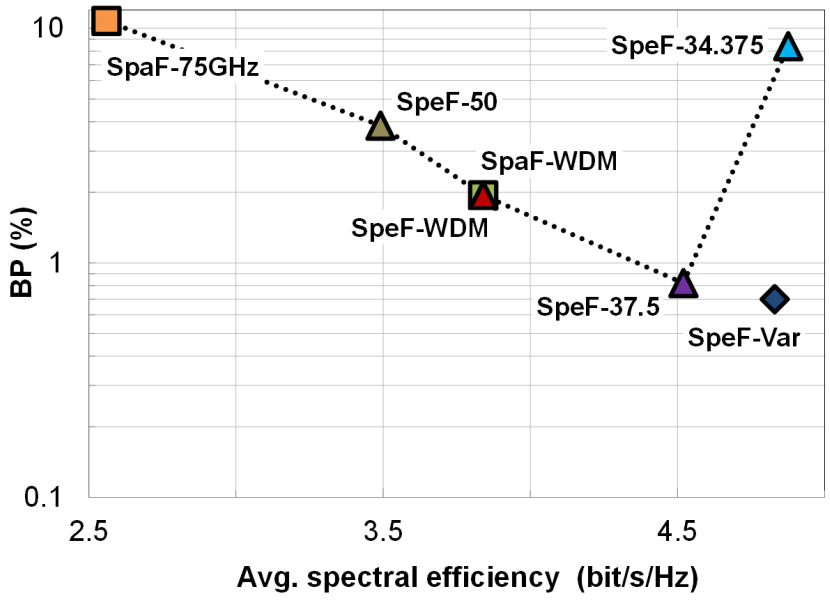

FIGURE 4.10: Blocking Probability vs. Spectral Efficiency of several sub-channel spacings at 1100 Erlang.

Another way to look at this is to observe the percentage of blocked connection requests as a function of the average spectral efficiency for an input load of 1100 Erlang, as shown in Figure 4.10. With respect to static sub-channel spacings a minimum is observed for SpeF-37.5 at a spectral efficiency of about 4.5 bit/s/Hz. As can be expected, lower efficiencies exhibit a gradual increase in experienced BP, owing to their less efficient use of spectral resources. A rapid increase in BP is also experienced for larger efficiencies, this due to steep reduction in optical reach, which prevents the establishment of long connections. As before, in the WDM scenario the spectral efficiency of SpaF and SpeF is the same, and therefore the choice of SpeF/SpaF is largely inconsequential. By optimizing the sub-channel spacing (SpeF-Var) both efficiency and BP can be further optimized, achieving, in the scenario under study, an efficiency of about 4.8 bit/s/Hz.

Joint and Fractional Joint Switching

With the additional support for FJoS it is interesting to revisit the earlier analysis of the performance of jointly switched SDM networks (DSpaF) in a slightly less disadvantageous scenario (9 instead of 4 spatial dimensions).

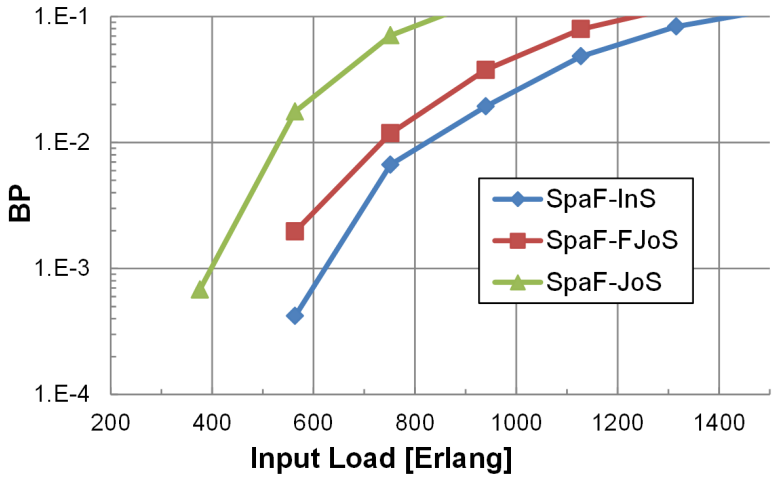

FIGURE 4.11: Blocking Probability vs. Input load for SpaF under InS, JoS and FJoS.

Figure 4.11 depicts the measured BP for InS, JoS and FJoS (with 3 groups of 3 spatial dimensions each) using SpaF-WDM. As observed earlier, JoS imposes a BP penalty compared to the InS case, of about one order of magnitude in the low to medium loads. The reason behind it is that the uniform load profile considered in this study is not well matched to maximum capacity of spatial super-channels. As can be expected, this BP penalty reduces significantly by employing FJoS, Consider, however, that the difference in performance strongly (almost exclusively) depends on the mismatch between the traffic demand sizes and what the JoS network supports.

4.6 Impact of Traffic Scaling on SDM Performance

Traditionally, dynamic traffic scenarios in optical networks have been modeled using a a Poisson process for connection arrivals, and some distribution (often uniform) for selecting the add and drop nodes of each demand. This model offers two "knobs" to affect the load of the network: average Holding Time and average Inter-Arrival Time, which together determine how many connections are, on average, demanded of the network at any one time.

This is perfectly suitable for contexts such as the Plain Old Telephone Service or Wavelength Division Multiplexing (WDM), where all connections require the same bandwidth. However, with the advent of flexi-grid and SDM, a third "knob" must be considered: that of average connection size (in the spectral domain), resulting from the choice of modulation format and super-channel size, spacing and (for SDM) shape. Using these technologies overall network traffic can be scaled by changing either the number of connections or their average capacity (and associated spectral/spatial footprint), or both, as the latter is no longer fixed. Technically, connection capacity can also be tuned in the context of WDM, however there it is done by changing the modulation format while keeping a fixed channel size; generally speaking there is little incentive to use anything but the most aggressive modulation that fits in a channel and ensures enough optical reach for any given WDM lightpath, beyond cost or protection scheme considerations.

The choice of traffic profile is significant for works studying SDM networks, as it can significantly affect their outcomes. Specifically, if traffic is scaled by increasing the average number of demands, higher loads result in large numbers of relatively small connections. This has obvious significance for JoS and FJoS architectures, which do not work well in such situations, as we proved in [79], where their performances (in terms of BP) were significantly worse than InS even at high network loads.

Conversely, if traffic is scaled by increasing the average connections size (such as in e.g. [93], [94]), at higher loads almost all demands fill or almost fill an entire spatial group, therefore limiting the penalties imposed by JoS and FJoS. In such situations, their performance can be expected to converge to that of InS at higher loads.

The results discussed in Sections 4.4 and 4.5 suggest that the performance of JoS and FJoS strongly hinges on the characteristics of traffic demands, specifically on them matching (or not) the coarser granularity of JoS and FJoS compared to InS.

In fact, in a planning scenario [95] the performance of JoS and FJoS were shown to converge to those of InS as traffic increases. The main difference between that work and those described in Sections 4.4 and 4.5 is that the traffic was scaled in terms of the size of demands rather than their number, which was fixed.

Realistically, the traffic a network experiences may include:

- a large number of small demands, typically seen in regional part of networks;
- a small number of large demands, common in e.g. inter-datacenter communications;
- a combination of the two, found in e.g. national scale networks serving heterogeneous traffic;

It is therefore reasonable to expect traffic growth due to both increases in the number of demands, and increases of their size over time, and evidence suggests that this has a disproportionate impact on SDM networks employing JoS and FJoS. A study of the sensitivity of InS, JoS and FJoS performance (in terms of induced BP) is the subject of the next batch of simulations, whose results were presented at the 2016 European Conference on Optical Communication (ECOC) [83].

4.6.1 Sensitivity Analysis of InS, JoS and FJoS Performance

Simulation Parameters

As for earlier studies, the Spanish National topology was used in conjunction with a custom discrete event simulation tool different from the one used for all other experiments in the thesis, this one developed by a colleague working on these topics as part of the same project, Behnam Shariati, who is the main author of [83] and who performed the simulations leading to these results, using the following agreed upon parameters:

- 12 spatial dimensions (SMFs in a bundle) for all links, a 2 groups of 6 spatial dimensions for FJoS;
- DP-8QAM modulated carriers at 32 Gbaud, spaced at 50 GHz (WDM grid) except for the last test (Figure 4.18), for a capacity of about 192 Gbps per carrier;
- KSP (K = 3) with SpaF (Algorithm 6) for RSSA.

Traffic demands were generated according to a Poisson process with mean μ and standard deviation σ over the range $[50, 2250]$ Tbps, measuring the resulting BP as performance metric.

Simulation Results

Firstly, in the interest of clarity, consider the Cumulative Distribution Functions (CDFs) for normally distributed demand sizes with fixed σ=200 Gbps and μ=700, 1150, and 1600 Gbps depicted in Figure 4.12.

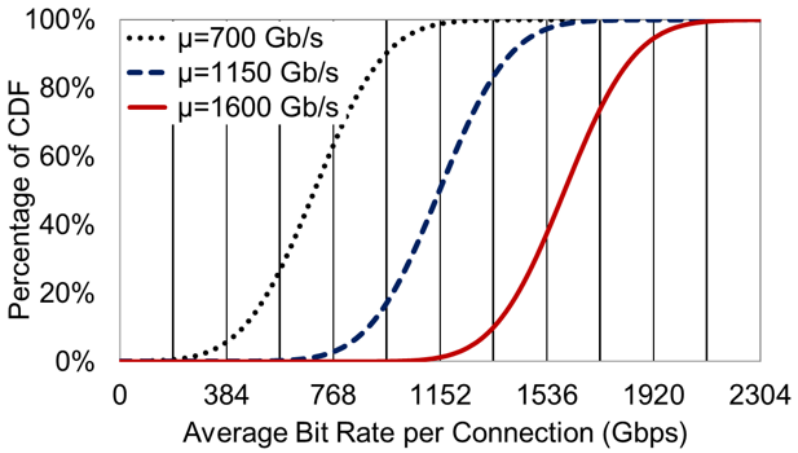

FIGURE 4.12: CDF of normally distributed demand sizes with fixed σ=200 Gbps and μ=700, 1150, and 1600 Gbps.

For μ=700 Gbps most demands are relatively small, resulting in more than 98% of them requiring less than half of the 12 spatial dimensions. For μ=1600 Gbps the opposite is true, and more than 98% of the demands require spectral resources on more than half of the 12 spatial dimensions to be allocated. For μ=1150 Gbps,i.e., 6 carriers, demands have a 50% chance of requiring more, or less, than half of the available spatial dimensions.

The effect of changing σ with a fixed μ=1248 Gbps are depicted in Figure 4.13. In short, the larger the variance the more spread out the possible values of traffic demand sizes.

The parameters from Figure 4.12 are used to produce the BP vs. Input Load (in Erlangs) for Figure 4.14, 4.15 and 4.16, respectively employing μ=1600, 1150, and 700 Gbps, by varying the total number of connection demands.

For large demands (Figure 4.14) the performance of InS, JoS and FJoS is virtually identical. This is because most demands in this scenario require spectral resources on more than half of the available dimensions, therefore the spectral resources that InS and sometimes FJoS leave available are almost never sufficient to allocate another demand, thus leading to the same results as JoS.

For diverse and relatively medium-size demands (Figure 4.15) FJoS and InS exhibit a detectable (even significant, in the case of InS) performance advantage, in terms of BP, over JoS, since in this scenario a sizable fraction of incoming demands is small enough to fit into the uninitialized resources on higher spatial dimensions left available using these switching paradigms.

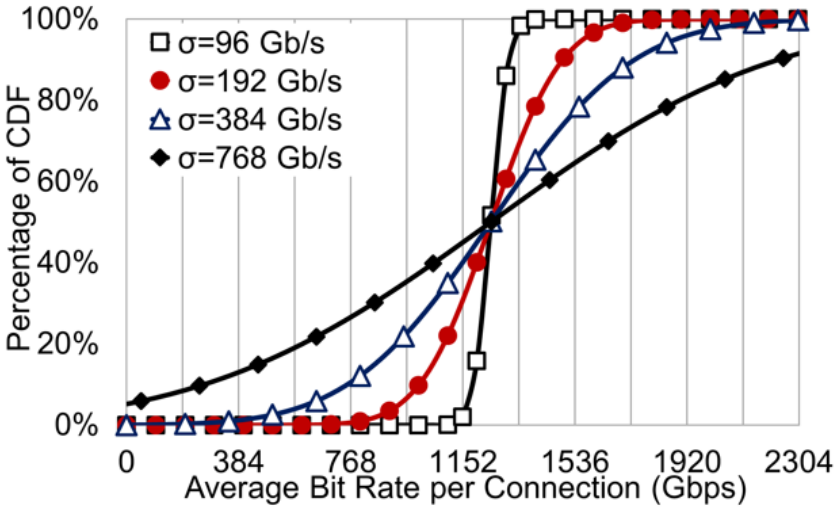

FIGURE 4.13: CDF of nnormally distributed demand sizes with fixed μ=1248 Gbps, and σ=96, 192, 384, 768 Gbps.

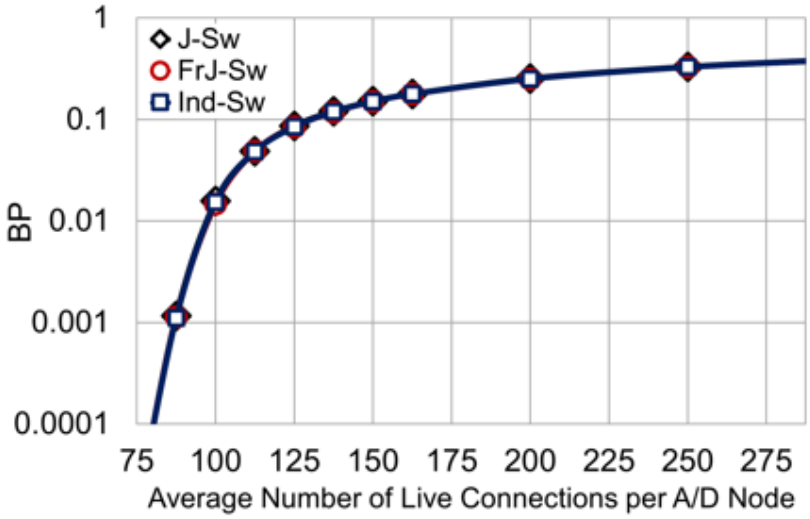

FIGURE 4.14: Blocking Probability vs. Input Load [Erlang] for normally distributed demand sizes with μ=1600 Gbps.

Finally, as the average demand size gets smaller (Figure 4.16) the performance penalty of JoS and FJoS becomes greater, as more and more demands can be placed at the same frequency on different spatial dimensions by InS, and rarely FJoS, but not JoS.

As before, in scenarios with many small connections, despite bringing very significant cost savings [95], JoS and FJoS do not achieve a reasonable level of performance. However, for appropriately-sized demands, which can be achieved by e.g. electronic grooming, this performance penalty is minimized.

To confirm these results, the outcomes of a complementary simulation performed by varying σ and keeping μ fixed at 1248 Gbps (i.e., $>$ 6 spatial dimensions per demand) and the number of live connections per Add/Drop node fixed are depicted in Figure 4.17. Here the average load was kept fixed at $14 * 112 * 1248$ Gbps $\simeq$ 1.95 Pbps. For small deviations the performance of all three switching paradigms

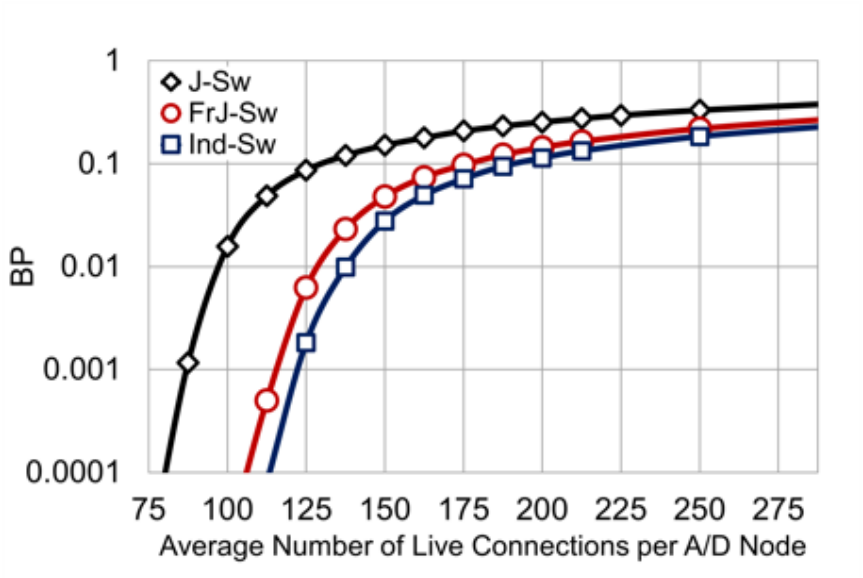

FIGURE 4.15: Blocking Probability vs. Input Load [Erlang] for normally distributed demand sizes with μ=1150 Gbps.

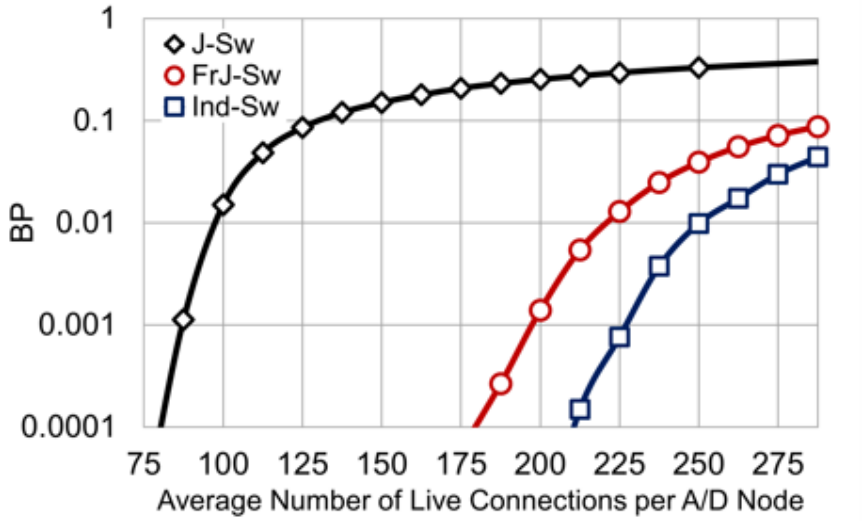

FIGURE 4.16: Blocking Probability vs. Input Load [Erlang] for normally distributed demand sizes with μ=700 Gbps.

coincide, as for the scenario depicted in Figure 4.14. As σ increases, i.e., the diversity of demand sizes becomes larger, the performance of FJoS and, in particular, InS, increases, further confirming that they are the superior choice for networks carrying highly diverse traffic with little to no grooming.

Conversely, by fixing the average number and diversity of connection, and increasing the network's load via larger μ, as shown in Figure 4.18, the performances of InS and FJoS converge to that of InS, offering a strong justification for investing in this cheaper technology.

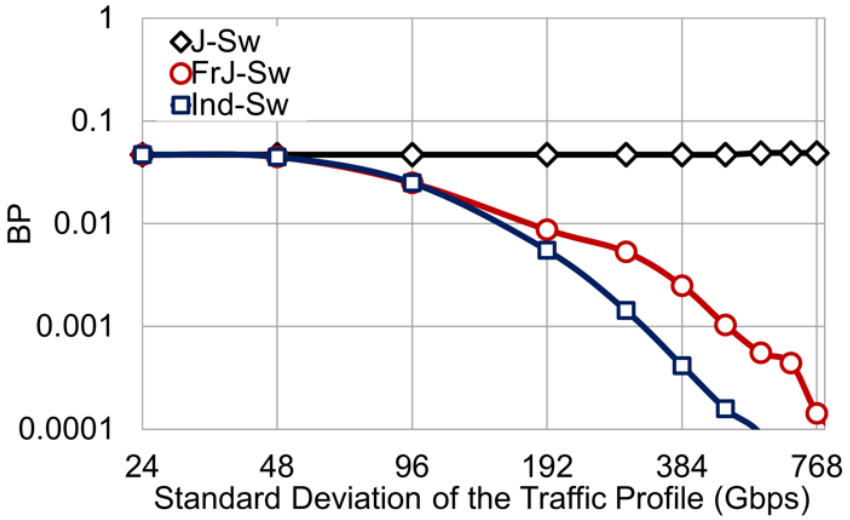

FIGURE 4.17: Blocking Probability vs. σ with μ=1248 Gbps and 112 active demands per Add/Drop node.

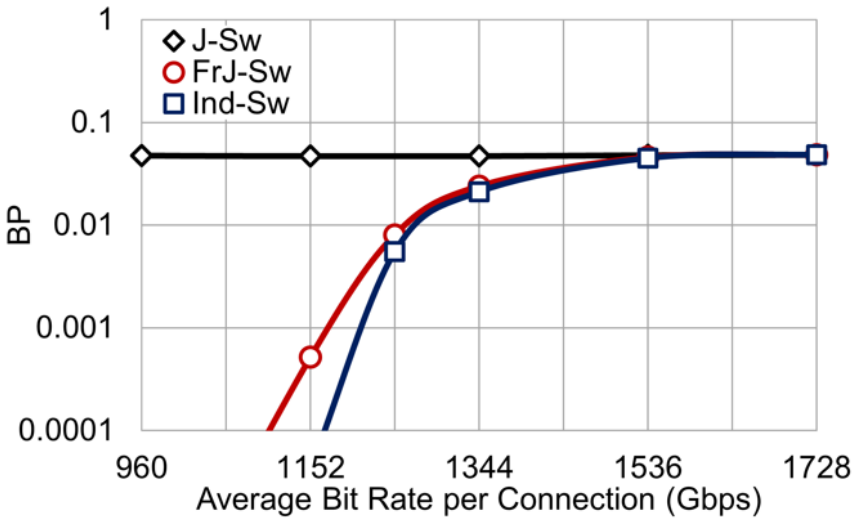

FIGURE 4.18: Blocking Probability vs. Input Load under variable μ and fixed σ.

4.7 Improving Performance of JoS Networks via Spatial Group Sharing

As shown in Section 4.6, the performance of JoS and FJoS strongly depends on the characteristics of the network traffic, but ultimately, if the traffic is scaled appropriately, converges to that of InS, while offering significant cost savings [95].

Is it possible to address the performance penalty of (F)JoS SDM networks and make them more attractive even as short-term solutions? One way to do so is to rely on two slightly tweaked InS ROADM architectures, and properly modified RSSA heuristic to exploit them.

The performance shortfall of both JoS and FJoS with respect to InS in uncoupled

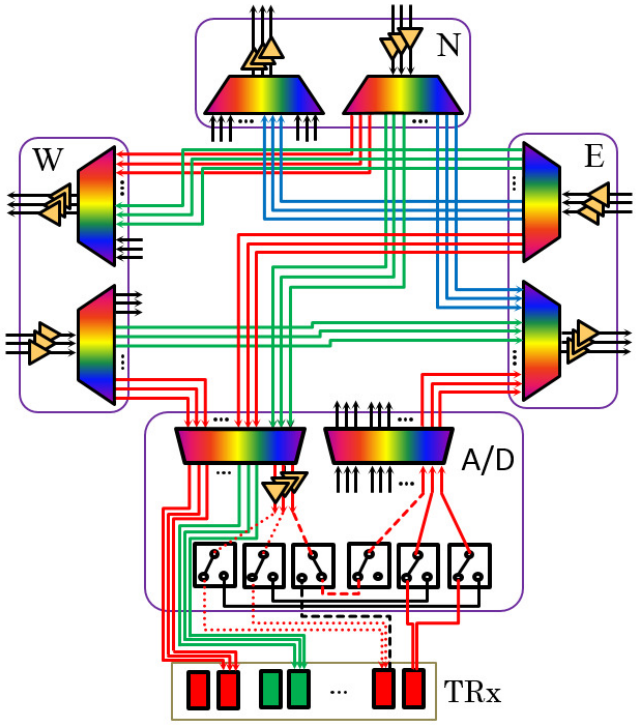

FIGURE 4.19: SDM InS ROADM architecture to enable spatial group sharing in intermediate nodes. Optical Multicast switches are used to recirculate optical connections from locally dropped groups.

SDM networks exhibited in Section 4.5 can be traced to the fact that when several demands smaller than the capacity of a spatial group (which defines the minimum service granularity of the network) are requested, they do not fill its whole capacity at a given frequency, thus resulting in "busy" (i.e., already switched on an end-to-end path) yet unused spectral resources. In other words, additional capacity could be extracted if partially-occupied spatial groups could be shared between multiple demands, even without the aid of electronic grooming. This could be achieved under two main scenarios:

1. *Demands that share the same path*, and therefore only require additional flexibility in the source and destination ROADMs, and

2. *Demands that share at least one link on their path*, which would require additionally flexibility in most intermediate nodes.

Scenario 1 could be realized in at least two ways: either via electronic grooming (as proposed in [78] and other works), i.e., merging multiple incoming data streams into a larger one that more closely matches the increased granularity of (F)JoS SDM networks, or by utilizing e.g. (F)JoS Colorless, Directionless, Contentionless (CDC) ROADMs with and InS Add/Drop degree. In the latter case, multiple connections from the same source node may be multiplexed within the same spatial group and then de-multiplexed and the destination node directly in the optical domain, albeit at a slight increase in the ROADM cost (due to the additional WSSs needed in the Add/Drop degree).

Scenario 2, could instead be implemented with e.g. the three-degree CDC route-and-select SDM ROADM architecture for the joint-switching of three single-mode fibers depicted in Figure 4.19, thus enabling the establishment of one, two or three

connections within a spatial group. WSSs with a minimum port count of $3 \times (1 \times 3)$ are required, and contentionless operation is achieved through the use of one optical Multicast Switch per each spatial dimension or one $P \times Q$ WSS [96], both of which may lead to the requirement of further amplification, not shown in the figure, if not already present. Spatial groups can be routed toward the other two directions or locally added/dropped via the Add/Drop stage, which leads to two arrays of three 1×2 opto-mechanical fiber optic switches (whose insertion loss, for commercially available realizations, is $<$1 dB) hooked up to selected input/output ports of the MCSs or $P \times Q$ WSSs to allow the selective extraction of channels from a spatial group (red dotted line in Figure 4.19), as well as the integration of pass-through (red dashed line) and newly added (red solid line) channels into the same spatial group.

Such an architecture adds a small cost to traditional (F)JoS solutions, yet the overall cost would still be well below that of InS designs. Technically, such an architecture could also be used to support scenario 1, but at a greater cost than strictly necessary.

4.7.1 RSSA Algorithms for Spatial Group Sharing

Back in Section 4.5 an RSSA heuristic that instantiates spatial super-channels by assigning resources according to a Space-First (SpaF) First Fit policy (using K-Shortest Paths for routing) was described, working by combing the spectrum matrix (where each row represents the WDM channels or spectrum slices of a single spatial dimension) column by column, trying to place all carriers of a demand at the same frequency on different spatial dimensions. A simple yet effective modification to SpaF, described in Algorithm 7, is to relax its constraint that all carriers of a single demand must utilize the same frequency. In this manner, demands for large spatial super-channels can be more easily accepted, by splitting them into smaller chunks that can be placed at different frequencies, in a sort of mixed or even "distributed" super-channel. This does not impact the overall efficiency, as the same amount of spectral guard-bands are used. The algorithms has a complexity of $\mathcal{O}(K \cdot E \cdot S \cdot W)$, where K is the number of paths considered, E is the number of (aggregated) edges in the network graph (in the worst case, e.g. a linear topology, a path may traverse all of them), S is the number of spatial dimensions and W is the number of WDM channels or spectrum slices. In other words, this simple heuristic has, at worst, linear complexity with respect to all parameters. Note that entire spatial groups are allocated at a time, therefore ensuring that the granularity of resources assigned to a demand matches that of a spatial group, whether an actual carrier is or is not carried in that position. We called this heuristic **SpaF with no-reuse** (**SpaF-N**), and we used it as a benchmark for measuring the performance of the other proposed spatial group reuse techniques.

Scenario 1 can be supported, assuming a network employing uncoupled spatial dimensions (e.g. bundles of single-core, single mode fibers), using an RSSA algorithm that computes as follows: upon receiving a new demand between a certain source and destination, the algorithm operates like SpaF, by firstly iterating over the K-shortest paths computed between those nodes. After selecting a path, it searches, among the active demands, for those that share the exact same path of the new one, and re-uses any spectral resources that are free and already switched on that path (because they belong to spatial groups that are partly occupied by the co-routed demands). Finally, if there are yet more carriers to place, the algorithm uses new (i.e., free and currently not switched) spatial groups, looking for the left- and top-most first, as per the basic SpaF heuristic, as shown in Algorithm 8. If the process is unsuccessful, it is repeated on any computed secondary paths, until either a valid

Algorithm 7 Generalized Space-First with no resource re-use (SpaF-N) first fit allocation using disjoint and mixed spatial super-channels for InS/JoS/FJoS SDM networks without lane changes.

```
1: for each path (in the set of K shortest path from source to destination) do
2:     D ← Vector()
3:     F ← Vector()
4:     c ← 0
5:     a ← 0
6:     Compute the ⋂ of available spectral slices for each spat. dim. in the path
7:     for each spectral slice s in the first spatial dimension of the path do
8:         for each spatial group G in the path do
9:             if sufficient contiguous slices for one carrier are free on the first spatial
   dimension of the group then
10:                for each spatial dimension g in the group do
11:                    D ← D ∪ g
12:                    F ← F ∪ s
13:                end for
14:                c ← c + |G|
15:                a ← 1
16:                if c = number of required carriers then
17:                    return D and F
18:                end if
19:            end if
20:        end for
21:        if a = 1 then
22:            Skip to the next spectral slice
23:        end if
24:    end for
25:    D ← ∅
26:    c ← 0
27:    a ← 0
28: end for
29: return failure
```

solution is found or K is reached, at which point the demand is blocked. The complexity of this algorithm is $\mathcal{O}(K \cdot D \cdot E \cdot S \cdot W)$, where D is the number of still-active past demands. We called this **SpaF with Full-match resource re-use** (**SpaF-F**).

Likewise, scenario 2 can be supported with a largely similar heuristic, where, however, the algorithm also attempts to re-use resources from past demands whose paths only partially match the current one, but only after exploiting any full-matches it finds, as shown in Algorithm 9. The worst-case complexity of this algorithm is the same as that of the previous one, albeit the number of past demands (D) is potentially twice as impacting (a fact that is hidden in Big-O notation). Nonetheless, in the tests reported in Section 4.7.2, the execution time per each incoming demand was a small fraction of a second (on a modern laptop, and with the number of previously instantiated demands rarely going above a few thousands). Observe that D is not a static network parameter like W or S, but will tend to grow as the network ages (at least initially). We call this heuristic **SpaF with Partial-match resource re-use** (**SpaF-P**).

Algorithm 8 Generalized Space-First with Full-path match resource re-use (SpaF-F) first fit allocation using disjoint and mixed spatial super-channels for InS/JoS/FJoS SDM networks without lane changes.

```
for each K shortest path from source to destination do
    Collect the spectral availability mask of the path
    for each active past dmnd do
        if past dmnd path == incoming dmnd path then
            for each resource ∈ past dmnd do
                if resource is free then
                    Allocate carrier here if any remain
                end if
            end for
        end if
    end for
    allocate remaining carriers like in SpaF-N
    if all carriers have been allocated then
        return success
    else
        remove allocations and mark new resources as free
    end if
end for
return failure
```

Please note that the implementations of these algorithms use a few extra optimizations, such as conditions within the loops to ensure that execution exits them at the earliest opportunity once all carriers of a demand are in place. An additional optimization to remove the dependency on E, by e.g. pre-hashing all paths, and using the hashes (or some IDs) to compare them for equality (rather than comparing them link by link), was not implemented in this work.

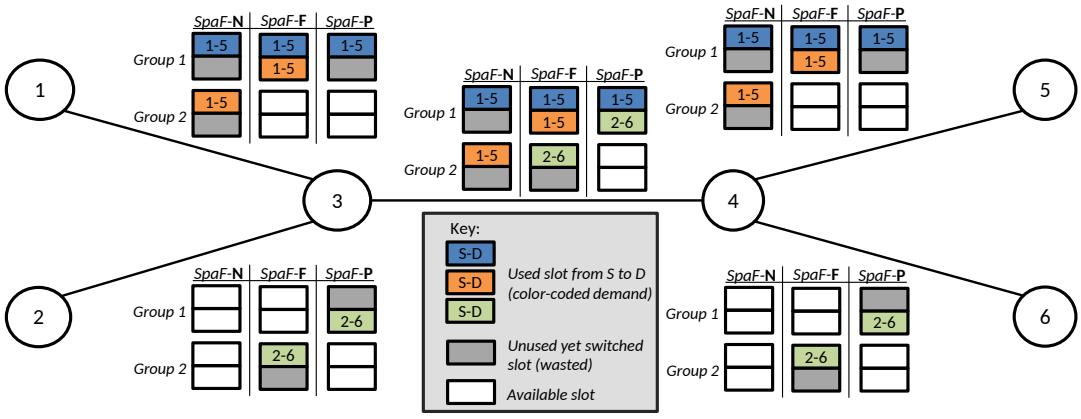

FIGURE 4.20: Visualization of the differences between SpaF-N, SpaF-F and SpaF-P: SpaF-F allows re-using resources on path-matching demands which would be wasted using SpaF-N, while SpaF-P works even with partial path matches.

The operation of the three heuristics is visually outlined in Figure 4.20, where a simple network composed of 4 spatial dimensions divided into 2 spatial groups is shown. Given two demands between nodes 1-5 of 1 carrier each, SpaF-N places them in different spatial groups (wasting the dark gray resources in the figure), switching

Algorithm 9 Generalized Space-First with Partial-path match resource re-use (SpaF-P) first fit allocation using disjoint and mixed spatial super-channels for In-S/JoS/FJoS SDM networks without lane changes.

```
for each K shortest path from source to destination do
    Collect the spectral availability mask of the path
    for each active past dmnd do
        if past dmnd path == incoming dmnd path then
            for each resource ∈ past dmnd do
                if resource is free then
                    Allocate carrier here if any remain
                end if
            end for
        end if
    end for
    for each active past dmnd do
        if past dmnd.path ⊇ incoming dmnd.path then
            for each resource ∈ past dmnd do
                if resource is free then
                    Allocate carrier here if any remain
                end if
            end for
        end if
    end for
    allocate remaining carriers like in SpaF-N
    if all carriers have been allocated then
        return success
    else
        remove allocations and mark new resources as free
    end if
end for
return failure
```

but not utilizing the remaining spectral resources along the path. SpaF-F is instead able to share a single spatial group for these path-matching demands (thus saving the resources belonging to the second spatial groups for other connections). In the same conditions, SpaF-P would perform in the same way as SpaF-F. However, if subsequent requests for services between nodes 1-5 and 2-6 were received, while SpaF-F would require the use of two different spatial groups on link 3-4, SpaF-P can re-use spatial groups for partial path matches and make do with a single spatial group.

4.7.2 Simulative Analysis of Spatial Group Reuse Heuristics

Simulation Parameters

Like for the simulations described earlier, a custom tool was used to simulate several dynamic scenarios (i.e., assuming that connections arrive one by one, have a finite duration and are eventually removed from the network) to measure the performance gains of the proposed ROADM architecture and heuristics.

As before, the Telefónica Spain's national network topology was used. In order to have a fair comparison among the three SDM switching paradigms, regardless of any transmission medium related performance constraints, bundles of 9 or 12 (depending on the experiment) independent single core, single mode fibers, each with an available spectrum of 4.8 THz (i.e., the entire C-band), equivalent to 96 DWDM channels or 384 12.5 GHz frequency slots, were considered for all links in the network. As an additional justification for the focus on SMF bundles, rather than more exotic SDM fibers, observe that utilizing already deployed fiber plants (that is, bundles of fibers) is the most cost-effective option for near-term realizations of SDM networks. Routing was performed using the well-known Yen's K-Shortest Paths algorithm [7] based on Dijkstra's algorithm [6] to compute a single Shortest Path, using a value of $K = 3$ for all experiments. Unless otherwise stated, the spatial group size was set to 3, although, obviously, InS scenarios are implemented with a group size of 1, and JoS with a size equal to the number of spatial dimensions. DP-8QAM was assumed as the modulation format, which, under the conditions of the simulations, should give a reach of about 1700 Km [8], which is enough to support even the longest demand between any of the Add/Drop nodes on the third-shortest path between them (i.e., no demands are rejected due to routing and optical reach issues). Such a modulation, assuming a 32 Gbaud baudrate, allows the transmission of up to 192 Gb/s on each carrier, while fitting in the standard DWDM 50 GHz grid (which was used in all experiments). Each data point is obtained by simulating 10^5 bidirectional demands.

As for earlier experiments, two network performance metrics were measured: Blocking Probability (BP), i.e., the ratio between the number of refused and total demands over the whole network, and network Throughput, i.e., the average of the sum of the capacities (in Tb/s) of all active optical connections in the network, against network load or the characteristics of the simulated traffic demands, under a number of conditions discussed in the rest of this section.

Both traffic scaling models discussed in Section 4.6, namely scaling the number and size of connections, were used.

In all experiments source and destination nodes were uniformly distributed between all Add/Drop nodes in the network.

Simulation Results: Scaling Average Number of Demands

A first batch of simulations involved scaling the load of the network by increasing the average number of connected demands, modeled as uniformly distributed over 1-9 carriers (out of 9 spatial dimensions), for an average requested throughput of 960 Gb/s (an average of 5 carriers per connection request), selecting a fixed inter-arrival time and variable holding time so as to achieve a desired average input load. This means that the average size of the demanded connections connections does not change, while their number does as the load on the network increases.

Figure 4.21 depicts the measured BP against the average input network load, expressed in Erlang (i.e., requested average number of active connection in the whole network at a time), for a network employing either SpaF-F, SpaF-N or SpaF-P, for each of the InS, FJoS or JoS switching paradigms. As expected, there is no difference between any of the algorithms when InS is used, since there are no spectral resources that are already switched but not used in that case, therefore InS-F, InS-N and InS-P were collapsed into a single curve (InS), which is the best performing. Consistently with previous results (Section 4.5, the JoS solution perform significantly worse than FJoS, which in turn is worse than InS in terms of BP in the context of dynamic small

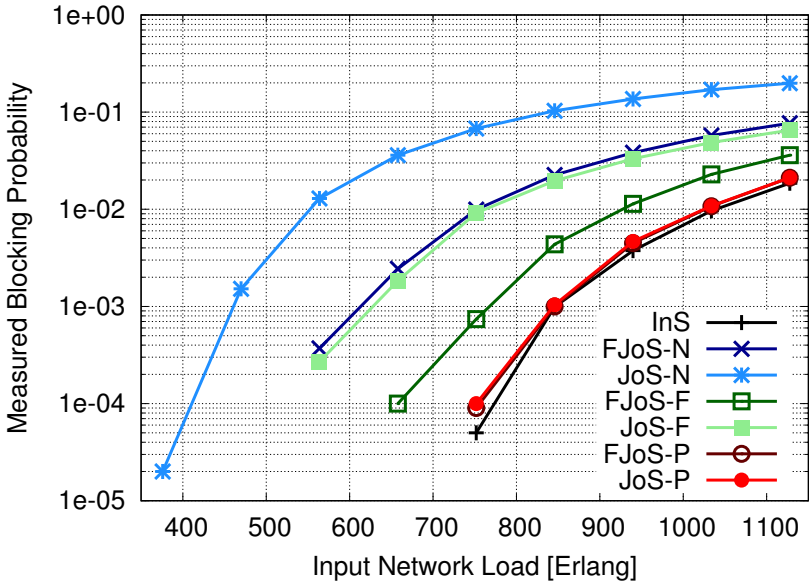

FIGURE 4.21: Measured Blocking Probability vs. Input Load for No, Full and Partial path matching for resource re-use.

connections. However, by re-using resources on full path matches (SpaF-F), significant reductions in BP, of more than an order of magnitude, can be achieved for both JoS and FJoS (with respect to SpaF-N). Furthermore, SpaF-P achieves a similar gain to reach nearly the same performance of InS even when using FJoS and JoS architectures, thus justifying the required slight increase in node complexity (and thus cost), despite requiring a slightly more complex switching node architecture.

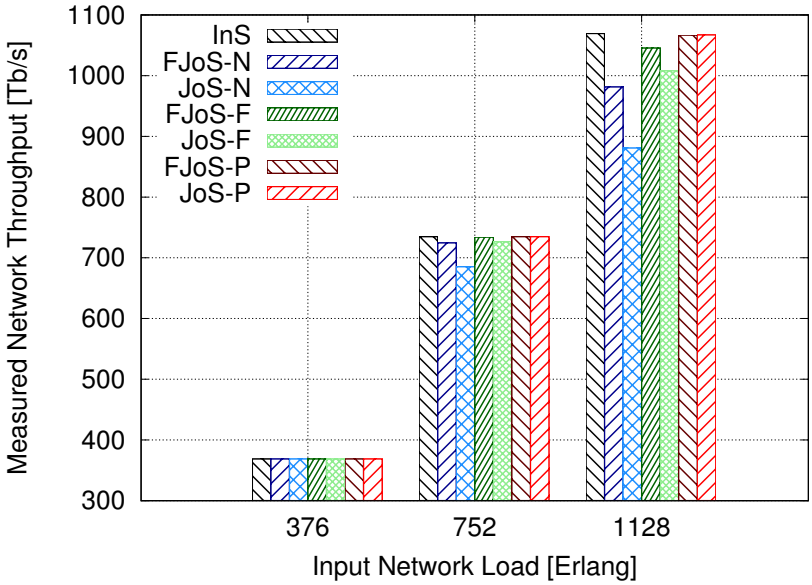

FIGURE 4.22: Measured Network Throughput vs. Input Load for No, Full and Partial path matching for resource re-use.

Accordingly, Figure 4.22 shows the measured throughput (in Tb/s) against the average input load for all algorithms and switching paradigms in the first batch of simulations. As expected from Figure 4.21, at very low loads all algorithms exhibit

the same performance, as there is no significant blocking to prevent connections from being established. At higher loads, however, first SpaF-N and then SpaF-F start to fall significantly behind the performance exhibited by InS when JoS and FJoS are used. Again, SpaF-P instead maintains comparable throughput (within a 0.3% difference) even when FJoS or JoS are used.

Simulation Results: Scaling Average Size of Demands

A second batch of simulations involved scaling the load of the network by instead increasing the average size of demands, by modeling the connection arrival as a (truncated and discretized) normal distribution, fixing both the inter-arrival and holding times (for a demanded load of 940 Erlang), and scaling the load of the demands by changing the average of the distribution (from 1 to 9 carriers) with a fixed standard deviation, i.e., we scaling the average size of connections while keeping their number fixed.

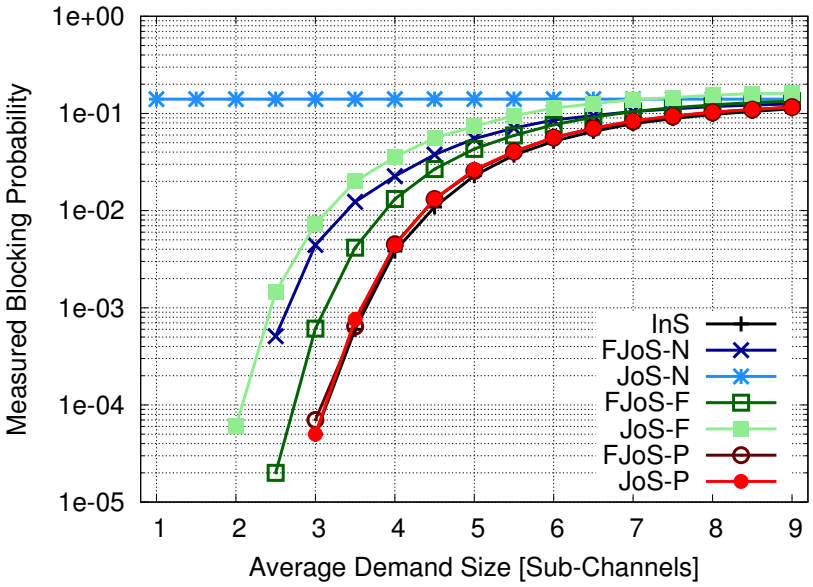

FIGURE 4.23: Measured Blocking Probability vs. Average Connection Size at 940 Erlang for No, Full and Partial path matching for resource re-use.

Figure 4.23 depicts the measured BP against increasing average connection size for the various algorithms and switching paradigms under study. Once again, the performance when InS is used is collapsed into a single curve. using this model the ranking of the various algorithms is largely the same, but with some notable differences. First of all, the performance of JoS-N is static. This is because this algorithm uses full Joint Switching of all spatial dimensions, without attempting any re-use of already switched resources; therefore, irrespective of the size of a demand, the same amount of resources (a spectral super-channel of 9 carriers) is allocated for all points in the X axis, which, at the chosen load exhibits significant blocking (consistently with Figure 4.21). Secondly, unlike Figure 4.21, here all algorithms and switching paradigms tend to converge to the same performance at higher loads (at least in log scale), and as connections mostly fill the available super-channel the penalty from using JoS and FJoS becomes relatively small. Again, both JoS-P and FJoS-P perform

admirably close to InS, still providing an advantage over the other approaches at lower loads.

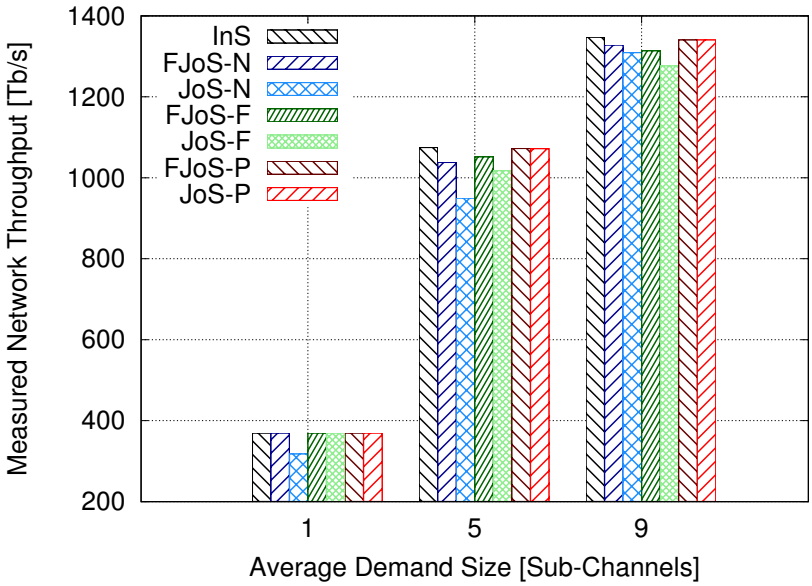

FIGURE 4.24: Measured Network Throughput vs. Average Connection Size at 940 Erlang for No, Full and Partial path matching for resource re-use.

Lastly, Figure 4.24 depicts the measured throughput (in Tb/s) against increasing average connection size for all algorithms and switching paradigms. Here we can observe a few differences compared to Figure 4.22 (where traffic was scaled increasing the number of connections). First of all, at low loads (with small connections) JoS-N performs significantly worse than the other algorithms, owing to it's high constant BP exhibited in Figure 4.23. As the average connection size increased, we observe the performance of the various algorithms first approaches that of Figure 4.22; indeed, with an average connection size of 5 (the same as in the number of connections experiment) the performance of the algorithms is the same. However, as the average connection size approaches that of the spatial super-channel used by JoS, the performance of all algorithms tends to converge. Once again, the performance of JoS-P and FJoS-P is largely the same of that of InS (as expected from Figure 4.23), therefore potentially justifying the extra investment compared to JoS-N and FJoS-N even in this scenario.

Simulation Results: Sensitivity Analysis to Traffic Parameters

A third batch of simulations, borrowing from the experience discussed in Section 4.6, studied the impact of different parameters of the process generating the input traffic to the performance of the proposed algorithms, focusing on the behavior of SpaF-F and SpaF-P since the behavior of simple SpaF was analyzed previously.

A normal distribution was used to draw the incoming demand size. Links were configured to use 12 spatial dimensions, which allows to simulate two levels of symmetric FJoS, with groups of 3 and 6 spatial dimensions, respectively.

First, the average input demand size and standard deviation (200 Gbps) were fixed, studying the BP resulting from changing the input load at several average

sizes. Then, the impact of changing the standard deviation at a fixed network load (1000 Erlang), against different average connection sizes was investigated.

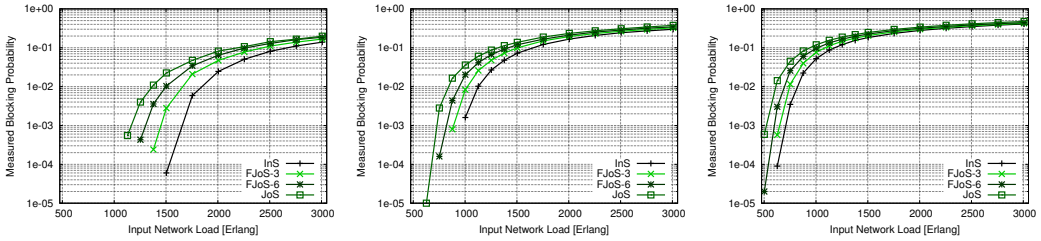

(A) (B) (C)

FIGURE 4.25: Measured Blocking probability vs. Input Network Load for SpaF-F with average connection sizes of 700 (a), 1150 (b) and 1600 (c) Gbps.

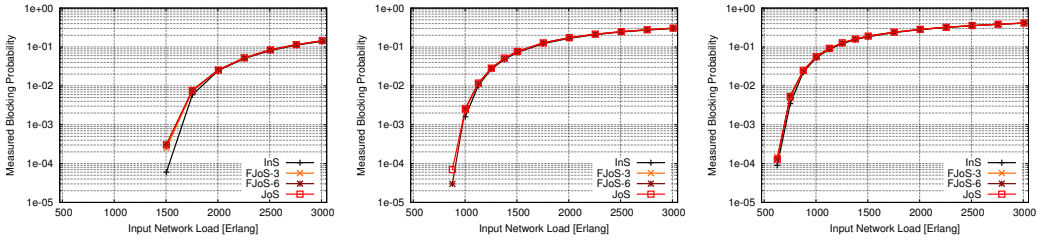

(A) (B) (C)

FIGURE 4.26: Measured Blocking probability vs. Input Network Load for SpaF-P with average connection sizes of 700 (a), 1150 (b) and 1600 (c) Gbps.

The sensitivity of SpaF-F and SpaF-P to different average connections sizes (700, 1150, and 1600 Gbps) was measured by studying the resulting BP as a function of network load, as shown in Figure 4.25 and Figure 4.26 for SpaF-F and SpaF-P, respectively. These results are largely comparable to those described in Section 4.6, in the sense that, broadly speaking, larger connections ameliorate the BP penalty suffered by FJoS and JoS designs compared to InS. Furthermore, the additional penalty resulting from increasing the size of the atomically switched spatial group in FJoS is clearly visible.

Observing the results for SpaF-F (Figure 4.25), with a majority of small connections (Figure 4.25a) the performance convergence of the four plots is slow and imperfect (note the log scale on the Y axis), while for larger ones (Figure 4.25b and Figure 4.25c) it happens much more quickly. One notable difference is that, due to the ability of re-using partially occupied spectral resources enabled by SpaF-F and SpaF-P, even with large connections (Figure 4.25c) there is a detectable difference in the performance of different switching algorithms (unlike in Section 4.6). This is because SpaF-F allows the splitting of a single large spatial super-channel into multiple smaller ones, inversely multiplexed into the available spectral/spatial resources. For the same reason, the overall performance of JoS exhibits a dependence on the average connection size; this was not the case for the original SpaF implementation (Section 4.6). Note that the first data point for JoS in Figure 4.25b has a very low statistical confidence (exactly one occurrence out of 10^5), and can thus be disregarded.

The results for SpaF-P (Figure 4.26) exhibit a behavior similar to that of Figure 4.21, i.e., JoS and FJoS block about as much as InS, with the latter performing

exactly as in Figure 4.25. Again, the performance of JoS is not constant, as in in Section 4.6, but closely tracks that of InS, as does FJoS, for the same reasons as before, magnified by the greater likelihood of resource re-use of SpaF-P (which allows link-matches in addition to full-path matches).

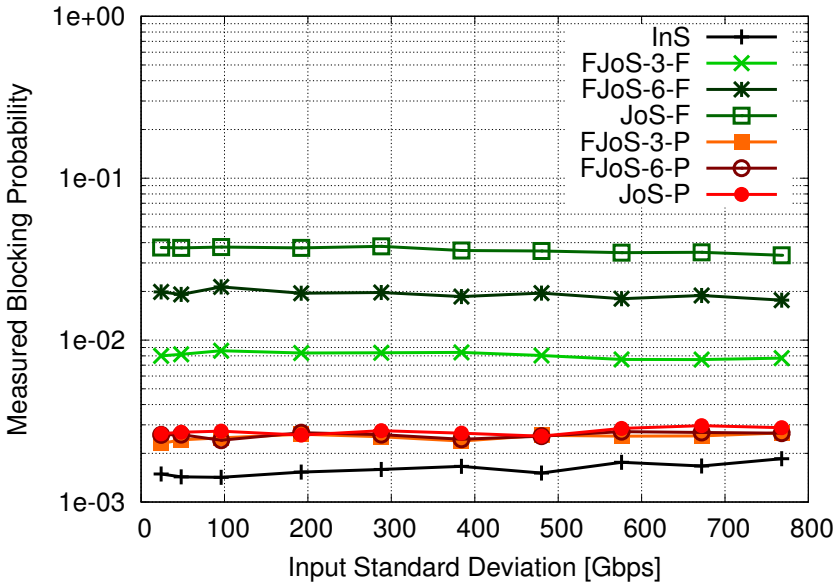

FIGURE 4.27: Measured Blocking probability vs. Input Traffic Standard Deviation for SpaF-N and SpaF-P, with average connection sizes of 1150 Gbps at 1000 Erlang.

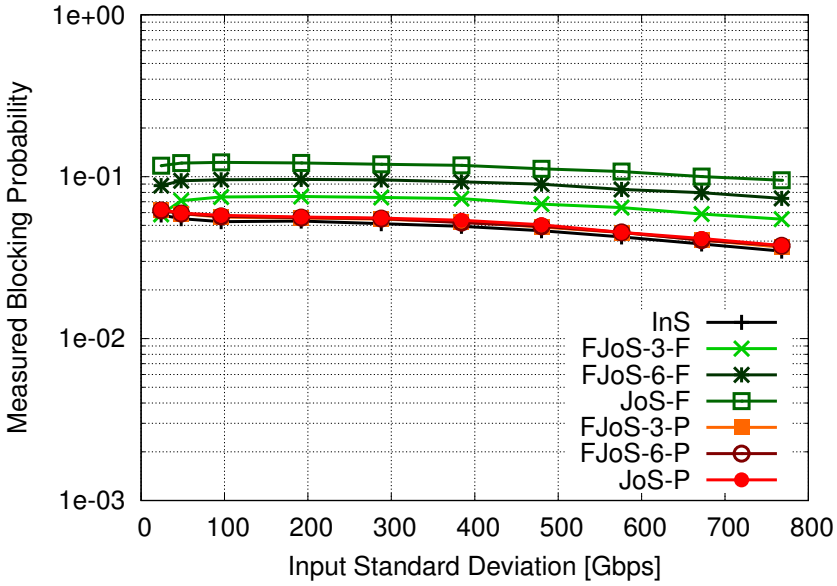

FIGURE 4.28: Measured Blocking probability vs. Input Traffic Standard Deviation for SpaF-N and SpaF-P, with average connection sizes of 1600 Gbps at 1000 Erlang.

Lastly, a study of the sensitivity of SpaF-F and SpaF-P to the diversity of the input traffic was performed. This was measured through its standard deviation: a low value indicates fairly uniform connection sizes (close to the chosen average), while a large one results in connection requests ranging over the whole space of

possible sizes. The resulting BP for both SpaF-F and SpaF-P with respect to the standard deviation of the input traffic model, at a load of 1000 Erlang, is depicted in Figure 4.27 and Figure 4.28 for average connection sizes of 1150 and 1600 Gbps, respectively. No result is shown for 700 Gbps, since, as shown in Figure 4.25a and Figure 4.26a, neither SpaF-F nor SpaF-P exhibit any blocking at 1000 Erlang with this average connection size.

In both Figure 4.27 and Figure 4.28 there is no difference between SpaF-F and SpaF-P under InS conditions, and the exhibited BP increases steadily with the size of the switching group. Also in both cases the resulting BP appears to be far less sensitive to the value of the standard deviation compared to what was observed in Section 4.6, where InS and FJoS performed significantly better than JoS with higher levels of standard deviation. This was assumed to be because FJoS and especially InS are better suited to accommodate a mixture of big and small demands, due to their smaller granularity. This effect is still somewhat visible in Figure 4.27 and especially Figure 4.28, however is is far less marked and also applies to JoS. This is because SpaF-F and SpaF-P allow both the re-use of partially occupied spatial groups, and the inverse multiplexing of large connections into a number of spatial groups, therefore suffering far less from the increased granularity issue. Very similar results were obtained for higher input loads, consistently with Figure 4.25 and Figure 4.26.

4.8 YAMATO: an Experimental SDN Control Plane for SDM Networks

In addition to devising a model to represent SDM networks and resource allocation heuristics to install lightpahts over them, as part of the INSPACE project [97] an experimental control plane for SDM networks was designed and built (although the author of this thesis mostly only dealt with the former), dubbed YAMATO.

Before delving into a description of the control plane design and an analysis of its performance, this section describes how to deal with the general case of SDM networks employing lane changes and/or presenting cross-wired spatial dimensions, and the additional challenges facing reactive reliability in such networks.

YAMATO: Origins of a Name

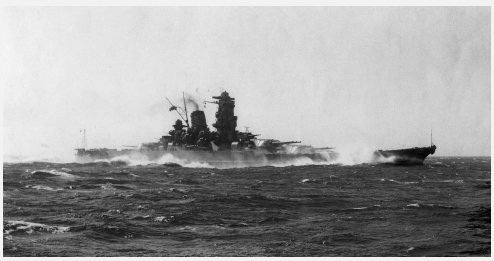

The name YAMATO comes, funnily enough, from a fictional space-worthy battleship (see the joke? *Space*-ship…), depicted above (left), itself inspired by a real world WW2 Japanese battleship (right). Many thanks to Antonio Francescon for coming up with the name and reference, and to my supervisor, Dr. Domenico Siracusa, for the original joke.
All images © of their original owners.

4.8.1 Resource Allocation with Spatial Lane Changes and Cross-wired Spatial Dimensions

All the heuristics presented in this chapter produce a single assignment of spatial resources, which is fine for networks where spatial dimensions cannot change at intermediate nodes, but otherwise excessively limiting.

Algorithm 10 Generalized Space-First first fit allocation using disjoint and mixed spatial super-channels for InS/JoS/FJoS SDM networks with lane changes.

```
1: for each path (in the set of K shortest path from source to destination) do
2:     D ← Vector(i)
3:     for each link i in the path do
4:         D[i] ← Vector()
5:     end for
6:     F ← Vector()
7:     a ← 0
8:     b ← 0
9:     c ← 0
10:    Compute the ⋂ of available spectral slices for each spat. dim. in the path
11:    for each spectral slice s in the first spatial dimension of the path do
12:        for each spatial group G in the path do
13:            for each link i in the path do
14:                if sufficient contiguous slices for one carrier are free on the first
   spatial dimension of the group in this link then
15:                    for each spatial dimension g in the group do
16:                        D[i] ← D[i] ∪ g
17:                        F ← F ∪ s
18:                    end for
19:                    c ← c + |G|
20:                    a ← 1
21:                    if c = number of required carriers then
22:                        b ← b + 1
23:                        if then b = number of links in the path)
24:                            return D and F
25:                        end if
26:                        Skip to the next link in the path
27:                    end if
28:                end if
29:            end for
30:        end for
31:        if a = 1 then
32:            Skip to the next spectral slice
33:        end if
34:    end for
35:    D ← ∅
36:    c ← 0
37:    a ← 0
38: end for
39: return failure
```

One way to deal with this problem is to turn the D parameter, i.e., the set of

selected spatial dimensions returned by the algorithm, into a vector of such sets, each representing the spatial dimensions selected for the i^{th} link in the path. An example of such an heuristic, which is an extended version of Algorithm 7, is reported in Algorithm 10.

Such an heuristic can deal with the general case of lane changes, and can also work on networks employing heterogeneous SDM links, i.e., links with different numbers of arrangement of spatial dimensions. The latter case is not a complete fancy, as the piecemeal upgrade of SMF-based SDM networks, starting from the busiest links, is one of the attractive options of this paradigm.

Another mean to achieve compatibility with such heterogeneous networks, even if they do not support lane changes, is to introduce a fifth sub-step in the RSSA problem (described back in Section 4.3) to produce the per-link spatial dimension assignment and verify its feasibility. This can be easily done using the well-known Max-Flow algorithm [98], on a network graph restricted to the representation of the intra- and inter-node links of the chosen path (to "herd" spatial dimensions on that path). The algorithm returns the maximum number of spatial dimensions occupied (which must match with the number of expected carriers in an exclusively spatial super-channel) and the set of links (switching ports, in Section 4.2 parlance) used to derive that number.

4.8.2 Optical Restoration Challenges in Heterogeneous SDM Networks

Optical restoration [12] is a reliability mechanism through which a transparent optical connection (i.e., a lightpath), which fails because either a link or a node in its path suffers a catastrophic issue (e.g. a fiber cut), is re-routed on an alternative path (typically computed a posteriori, i.e., dynamically after the failure occurs) in order to resume service.

This process is subject to additional complexity in the context of SDM networks whenever these are composed of heterogeneous links and/or nodes; this can be expected to be a common occurrence, both during a transition phase towards SDM (e.g. exploiting what bundles of fibers are already deployed), and to better cater to a non-uniform traffic matrix (adding more spatial dimensions to the more heavily utilized links).

For example, imagine that a connection is transmitted as a spatial super-channel of two signals (transmitted at the same frequency), over the path A-B on the simplified network depicted in Figure 4.29. Upon a failure on link A-B, a restoration scheme would then choose to re-route on the A-C-B path (the only alternative). However, the C-B link has only one spatial dimension. Therefore, to correctly restore this service on the chosen path, it must be split and sent over a single spatial dimension, if feasible. Note that if e.g. the super-channel is generated by a single SDM transmitter, which can only output a single frequency at a time, this would not be an acceptable solution. A similar reasoning could be needed even in the case of paths with homogeneous links: it is sufficient that no two same channels or spectrum slices be available at the same frequency on both spatial dimensions.

In general, if the resources on, number of or relationship between spatial dimensions differs from primary and backup path(s), a restoration scheme for SDM networks must be smart enough to recognize cases where it is necessary, and feasible, to alter the shape of a super-channel. The same requirements also apply to any pre-planned protection scheme.

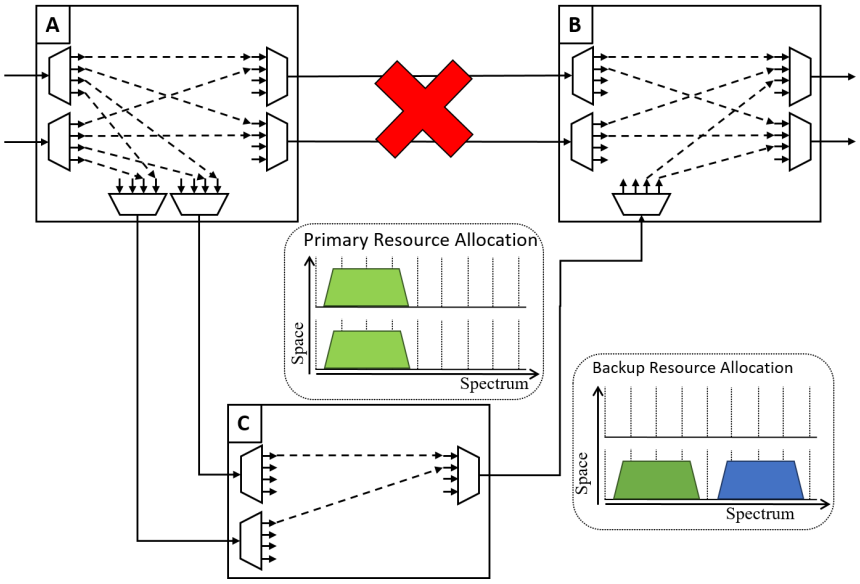

FIGURE 4.29: Restoration in heterogeneous SDM networks: alternate paths may have different spatial configurations; only relevant WSSs are depicted.

4.8.3 YAMATO Control Plane Architecture and Operation

The main component of YAMATO is a spatially-spectrally flexible SDN optical network controller developed on top of the OpenDaylight SDN framework (ODL) [99], capable of managing heterogeneous SDM networks employing InS, JoS and FJoS ROADMs, and any mixture of SDM links. The high-level architecture of YAMATO is depicted in Figure 4.30. The main controller functions are implemented by several sub-modules, most of which where developed from scratch.

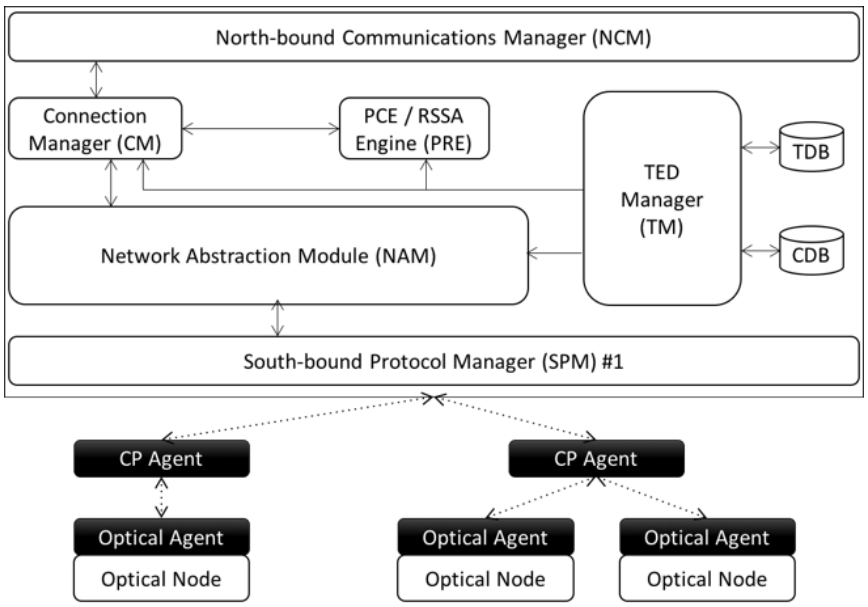

FIGURE 4.30: YAMATO architecture schema.

The North-bound Communications Manager (NCM) module exposes an HTTP REST JSON interface to external applications, through which they can retrieve network and connection information and manage optical services. Furthermore, the

NCM has been designed to be easily extended with a Graphical User Interface without changing the internal functionalities and object representations.

The Connection Manager (CM) handles connection requests coming from the NCM, enforces their serialization, and manages the life-cycle of optical services. In particular, when a new connection service request comes from the NCM module, the CM validates it, asks for a path and spatial-spectrum allocation to the PRE, and sends the request, with full path and spectrum information, to the NAM. Finally, after receiving the required service updates from the lower modules, it notifies the NCM with the updated status (installed/failed) of the request.

The PCE/RSSA Engine (PRE) performs Routing, Space and Spectrum Allocation (RSSA) using information retrieved from the TED manager. The well-known Yen's K-Shortest Paths algorithm [7] is used to compute the end-to-end, Inter-node path, while Intra-node routing (i.e., selecting which intra-node cables to cross for a given service, see e.g. Figure 4.2) is computed after RSSA using the Max-Flow algorithm [98]. This is done to constrain the algorithm to use the same inter-node links for all sub-channels of a service, which is paramount in the case of coupled SDM links, as MIMO DSP techniques are used to compensate from crosstalk and core coupling. Note that, if all ROADMs are symmetrically wired, Intra-node routing is a fairly straightforward process, which could be computed using simple heuristics rather than the more computationally expensive Max-Flow algorithm, which, however, is general enough to be able to handle even ROADMs with asymmetric patch cable configurations. As space-spectrum resource allocation strategy, PRE embeds a Space-First First-Fit algorithm (specifically, Algorithm 6) for spectrum and space allocation. Optical feasibility estimation is currently carried out by means of a table that provides a pre-computed relation between transmission parameters and reach, similarly to Table 4.1.

The TED Manager (TM) simply provides a relational Database as a Service (DaaS) to other modules, ensuring data persistence and consistency.

The Network Abstraction Module (NAM), the only module used largely unmodified from the stock ODL sources, collects topology information and changes, and provides an abstract interface towards the underlying network to the upper modules.

The South-bound Protocol Manager (SPM) implements a protocol specific driver to communicate to the remote CP Agents, based on a RESTCONF [100] XML representation of a node in our network model. RESTCONF is an HTTP-based protocol that provides a programmatic interface for accessing data defined in YANG [101].

Finally, YAMATO leverages two different external software artifacts, running remotely: (*i*) the Control Plane Agents (CP Agents) and (*ii*) the Optical Agents (OAs), interconnected through another custom HTTP REST JSON API. Each CP Agent represents a different Optical Node. It interacts with the YAMATO controller on one side, while connecting to multiple Optical Drivers on the other. Its main functionalities include:

- Translating between the OAs and YAMATO. The data model used by YAMATO and by the Optical Agents is different (network-centric vs. device-centric). Moreover, each CP Agent can control several devices and thus needs to correlate all their configurations to provide an integrated node view to YAMATO.
- Collecting information local to a node (WSSs, patch cable configurations, transceivers). The CP Agent is responsible for maintaining an up-to-date view of the status of all the inventory items belonging to a node, and update YAMATO of any changes.

- Configuring the optical devices according to incoming requests. When a configuration request is received from the HTTP RESTCONF interface with YAMATO, the CP Agent contacts the Optical Drivers, sends them the updates that pertain them, and waits for installation notifications, that are then forwarded back to the controller.
- Reacting to network failures and escalating them up to YAMATO. If a network event occurs, such as WSS-, link- or port-down event, the CP Agents are in charge of notifying YAMATO that will then take call the appropriate logic to overcome or mitigate the failure.

Each Optical Agent (OA) abstracts device-specific commands and configuration (generally vendor-dependent) and stores their up-to-date configuration in a database. The Optical driver has two interfaces: a JSON-REST Northbound AP, used to communicate with the CP Agent, that allows the Control Plane to configure each WSS and transceiver on the node, and Southbound interface (e.g. RS422 Serial) to communicate with each WSS and transceiver and configure them.

YAMATO Operation

This section outlines the procedures carried out by YAMATO in case of: (*i*) a new connection service request and (*ii*) service restoration after a network failure.

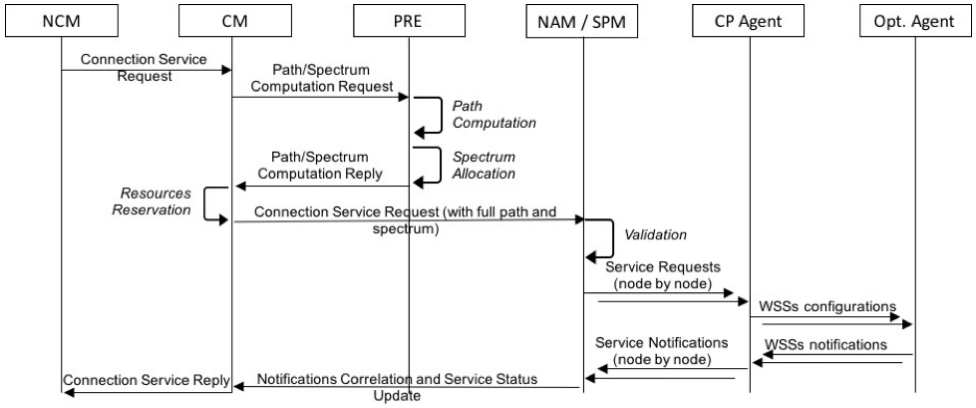

FIGURE 4.31: Connection Service work-flow.

Figure 4.31 describes the work-flow in the case of a new connection service request. The request arrives to YAMATO through the NCM module, that parses it, converts the JSON representation into a Connection Service Request (a Java object) and sends it to the CM module. The Connection Manager validates the requests (verifying e.g. the presence of required fields, the correctness of IDs with respect to those stored in the internal database, etc.) and then requests a path and spectrum computation to the PRE. The PCE/RSSA Engine performs path computation and space-spectrum allocation in series, and replies with a full path and selected spatial and spectral resources (if at least one solution was found) or with a negative feedback. The CM then pre-reserves the requested resources on the TM and sends the fully populated request (with full path and spectrum information) to the NAM, which again validates it with the updated status of the physical resources and sends it to the SPM driver(s) in charge of the optical nodes on the selected path. These send a set of partial connection service requests to the remote Control Plane Agents responsible for the optical nodes affected by the new service. Each CP Agent unpacks its local service request into multiple WSSs configurations that are then sent

to the Optical Agents, which configure the physical elements and send back to the CP Agent a notification when the process completes. The CP Agent then forwards the service status updates back up the NAM, which collects and correlates them to provide an up-to-date service status to the higher modules.

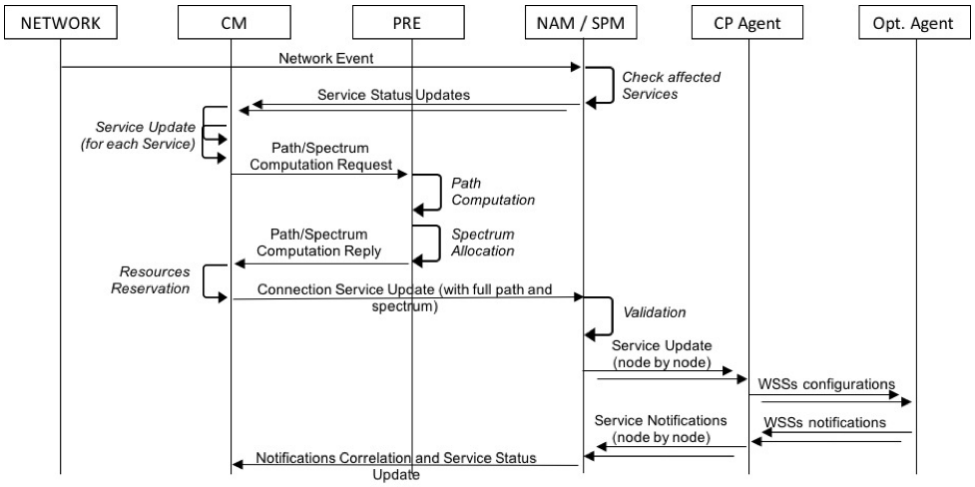

FIGURE 4.32: Connection Restoration (after Network event) work-flow.

The work-flow for the restoration of affected connections after a network event (i.e., node, port or link down/up) is depicted in Figure 4.32. In the case of such a network event, the NAM manager is responsible for identifying the affected connection services and for initiating the restoration procedure, updating their status and notifying both the TED (that will update the status of their resources in the database) and the CM about these changes. For each connection service, the CM first stores the outdated service properties (such as primary failed path and related spatial and spectral resources) to later re-use them (e.g. in the case of revertible services, that is, services that can be moved back to the original path after a failure has been repaired). Then, it starts a new service request process, similar to the one detailed previously. The main difference is that the service request(s) generated by the NAM to the Control Plane Agents also include request(s) to release the resources associated to the failed services before reserving the new ones. This does not lead to resource preemption, since these resources are still known to belong to the original primary path(s) in the SDN controller.

4.8.4 Experimental Setup & Performance Evaluation of YAMATO

As for earlier experiments, the emulated Spanish national backbone topology [90], depicted in Figure 4.5 back in Section 4.4.1, was used, configured as comprising 30 FJoS nodes (switching two symmetrical spatial groups) and 56 links, each made of 4 uncoupled spatial dimensions. The work "emulation" is used to describe the fact that CP agent for each node in the topology was instantiated, and furnished with a number of OAs emulating the respective WSSs by accepting configuration commands and replying successfully (as long as they are well-formatted).

As a performance metric, the time (in ms) needed to set up a new optical connection (from the point of view of the CP, i.e., ignoring amplifier setup times) was used, as a function of the length (number of hops) of a connection, and as a function of the number of previously instantiated (and still active) connections in the

network. These experiments outline the scalability of the system, and show which sub-processes need further refinements. In addition, the time needed to compute and install restoration services after a failure was also measured, as a function of the number of failed connections. For all experiments, each data point shows the average of 10^3 measurements, and is divided into several factors contributing to the overall setup time.

Scalability vs. Connection Size

In a first experiment, whose results are depicted in Figure 4.33, the average amount of time needed to instantiate a new connection, as a function of the number of links it traverses, was measured. This value is broken down into several sub-steps: Routing, i.e., the selection of the inter-node path to be traversed, then Modulation Level, Spectrum and Space allocation (MLSSA), i.e., the selection of the spatial and spectral resources to be used, and intra-node routing, then generation of the required YANG configuration, i.e., computation of the difference between the existing configuration of the nodes and that needed to implement the new service, and finally time needed by the remote agents to configure the emulated nodes and report back.

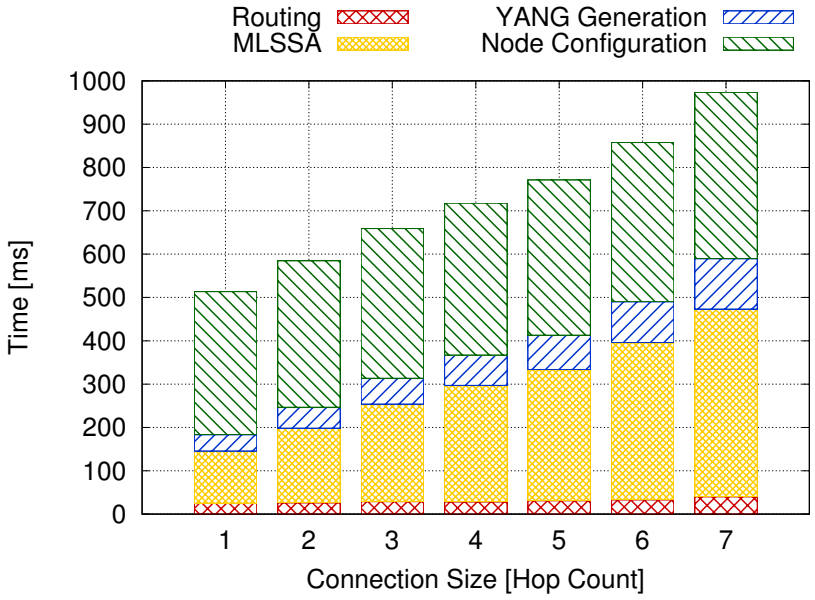

FIGURE 4.33: Setup time vs. demand size (hops).

Figure 4.33 shows that for all hop counts the process terminates in less than a second, and all sub-steps scale more or less linearly with respect to the hop count, but at significantly different rates. The routing sub-process scales very well (the worst-case complexity of our implementation is $\mathcal{O}(K \cdot N^3)$, where N is the number of nodes in the network and K the parameter of KSP). Likewise, the MLSSA and the YANG generation steps clearly scale linearly (at best) with the hop count, which is reasonable considering that MLSSA must check the availability of resources in all the hops of a path, and the length of the generated YANG clearly depends on the amount of hops to describe. Lastly, the node configuration step scales very well (it is nearly constant), due to the fact that YAMATO configures the nodes in parallel through the CP Agents, rather than waiting for a node to finish its configuration before moving to the next one, as done in the Resource Reservation Protocol traditionally used for this purpose.

A first lesson that can be gleamed from this result, and will be further reinforced by the others, is that while translating between several representations (JSON request, Java Object, YANG model(s), OA JSON model(s) can be necessary to tying together fairly disparate technologies, it has a non-negligible, and in fact rather significant, cost in terms of setup time.

Scalability vs. Number of Connections

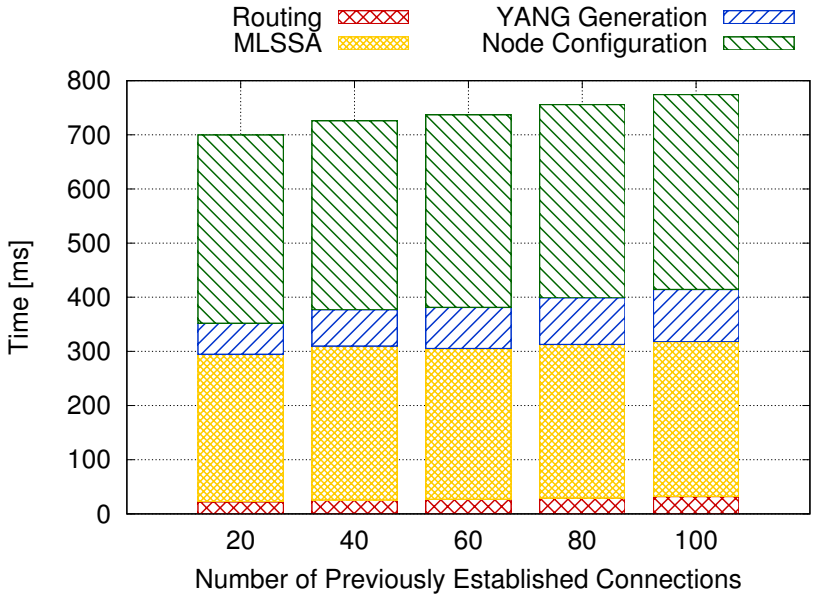

FIGURE 4.34: Setup time vs. number of connections already established.

In a second experiment, whose results are shown in Figure 4.34, the average amount of time needed to instantiate a new connection as a function of the number of random-length previously instantiated connections in the network was measured, again decomposed as before, be first installing the desired amount of connections and then repeatedly requesting and removing the same 4-hops test demand. Observe that in this case the scaling is nearly flat, which suggests that YAMATO can handle a significant number of optical connections. The slight increase in setup times in this case is mostly due to the YANG generation step, which includes several checks about the computed candidate connection against those already established, and therefore requires a longer time the larger the number of existing active connections in the network.

SDM Restoration Delay

The last experiment, whose results are depicted in Figure 4.35, shows instead the time needed to complete the restoration of a given number of affected connections after a simulated failure occurs. Clearly, this process is dominated by the MLSSA computation. This is because we did not implement an algorithm dedicated to this task, so the system simply performs a series of allocations using the specifications (source, destination, bandwidth) of the original requests. In fact, the time required by the MLSSA process in the context of restoration is almost exactly equal to the average time it needed for a single connection (Figure 4.33, times the number of

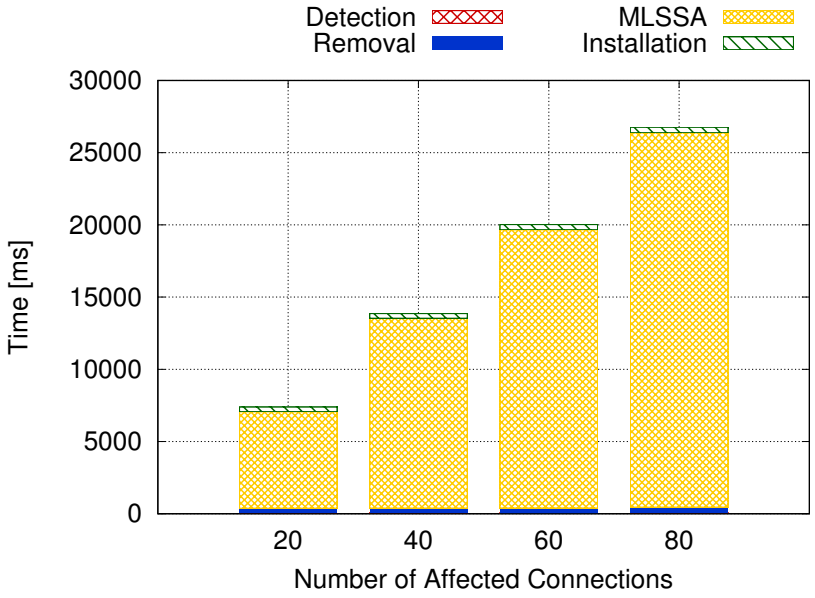

FIGURE 4.35: Time needed to setup a given number of failed optical connections.

affected connections. We expect that by implementing a proper heuristic able to accommodate all affected demands in a single run this time could be significantly reduced.

While the overall performance of YAMATO is far from record-breaking, correctness, rather than speed, was our primary concern during development. This is reflected in the many layer of redundant checks that YAMATO performs at every step of the process. Low-hanging fruits to improve YAMATO's performances include parallelizing the configuration of each device, which is currently done serially, and reducing the current debug-level of internal checks. Yet more performance may be obtained by realizing the aforementioned "planning" MLSSA heuristic, as well as by profiling the code to locate the most critical areas for optimization. Finally, compiler-level optimizations may also allow us to squeeze a bit more performance from our implementation.

Chapter 5

Summary of Contributions and Conclusions

The stated objective of this thesis was to "contribute to the realization of the needed increase in the performance of future transport networks, by means of improved allocation resource algorithms". In order to achieve it, three main contributions were identified.

Firstly, in Chapter 3 we tackled the the problem of **Spectrum Fragmentation in Flexi-grid networks**, solving which would enable network operators to squeeze additional capacity from existing infrastructure before it requiring upgrades. After describing the problem and various approaches to quantify and mitigate in the scientific literature, we presented a new fragmentation metric, Wasted Slices and Usability Factor (WSUF, Section 3.3), which measures the probability for free spectral resources of not being useful to configure a lightpath, and two simple families of RSA heuristics, Minimize-Fragmentation-Routing (M-F-R) and Minimize-Fragmentation-NoRouting (M-F-NR), to minimize fragmentation (Section 3.4). Using simulations, we showed (Section 3.5.2) how the proposed metric outperforms two existing ones (and a baseline fragmentation-unaware RSA, *K*-Shortest Path with First Fit spectrum allocation) in terms of resulting network performance in a dynamic scenario. The results show that WSUF is a more effective metric at capturing the nuances of fragmentation than the other metrics from literature (thus resulting in better network performance in terms of accepted connections), and that both families of algorithms perform better (M-F-R potentially much better) than FF in terms of both blocking probability and network throughput. Furthermore, results show that M-F-R algorithms use consistently more spectral resources than FF and M-F-NR ones, because they load balance busy links as a side effect of minimizing fragmentation, thus utilizing longer-than-shortest paths well before the others. On balance, this behavior is found to be beneficial to the network's performance.

Secondly, in Chapter 4 we identified **Space Division Multiplexing** (SDM) as a promising solution for longer then capacity growth, overcoming the finite capacity of fibers by using multiple ones in parallel, and provided an overview of the technique and its related literature (Section 4.1). Two main weaknesses in the existing scientific literature were identified: **modeling** (Section 4.2)and solving the **resource allocation** problem in (Sections 4.4 to 4.7) SDM networks.

With respect to the former, in Section 4.2 we described a network model for SDM networks, showed how to use it to encode complex SDM configurations, and discussed how to incorporate its novelties in existing protocol models. On a related note, in Section 4.8 we detailed the architecture and work-flow of YAMATO, an experimental SDN control plane for SDM optical networks employing independent, joint and fractional joint switching paradigms and/or coupled SDM links, implementing and exploiting the proposed model capable of abstracting the features of

SDM devices. We also described the issues for the process of optical restoration stemming from the introduction of SDM, and evaluated the performance and scalability of YAMATO using an emulated SDM extension of a European national network, showing that it is capable of handling a significant number of connections with short setup times, and is capable of coordinating optical restoration.

With respect to the latter, a number of heuristics to solve the Routing, Space and Spectrum Allocation (RSSA) problem were presented and iteratively refined. A simulative study of simple extensions to the First Fit heuristic (Section 4.4.1) for solving RSSA showed that the potential capacity increase of SDM is in fact limited by some limitations in the way optical connections can be allocated, and that, while space-oriented heuristics work well with splitter-based SDM transceivers, they are less spectrally efficient (and, thus, exhibit lower overall network throughput) than spectrum-oriented ones, unless care is taken to precisely align spectral super-channels over multiple spatial dimensions. A further set of simulations (Section 4.5.1) reinforced the notion that different approaches to creating SDM super-channels lead to different spectral efficiency, but also highlighted how the less efficient spatial super-channels can benefit from significant reductions in cost due to the use of Joint and Fractional Joint Switching (JoS/FJoS), and quantified the performance loss (in terms of connection blocking probability) of these paradigms. Another study (Section 4.6.1) expanded the set of network traffic conditions studied to prove that, under the right scaling (namely, relatively low traffic diversity), (F)JoS network do not exhibit higher blocking than the much more expensive InS ones. We then showed (Section 4.7) an approach to improve the performance of (F)JoS in uncoupled SDM networks when combined with relatively small connection demands, by means of two RSSA heuristics requiring ROADM designs with slight modifications to the Add/-Drop stage. We used simulations to prove that the proposed algorithms can significantly improve the performance of SDM networks employing Joint and Fractional Joint switching, in some cases achieving performance parity with the much more expensive Independent Switching paradigm, and verified this claim both when network traffic scales due to the number of demands (resulting in many small connections) or their size (resulting in fewer, larger connections), showing that the proposed algorithms are insensitive to average connection size and standard deviation.

Overall, we believe the contributions of the thesis adequately provide at least some answers to the challenges faced by future optical transport networks, both in the short (Flexi-grid) and long (SDM) terms. Especially with respect to the latter, however, the findings outlined in this thesis, while hopefully useful first steps, are far from a complete analysis of the topic of SDM, and further work is needed to unlock the full potential of this technology.

Bibliography

[1] C. J. Hélio, A. L. N. Souza, J. C.S. S. Januário, S. M. Rossi, A. Chiuchiarelli, J. D. Reis, S. Makovejs, and D. A. A. Mello, "Single-carrier 400g unrepeatered wdm transmission using nonlinear compensation and dd-lms with fec feedback", in **2017 SBMO/IEEE MTT-S International Microwave and Optoelectronics Conference (IMOC)**, 2017, pp. 1–5.

[2] O. Gerstel, M. Jinno, A. Lord, and S. Yoo, "Elastic optical networking: A new dawn for the optical layer?", **Communications Magazine, IEEE**, vol. 50, no. 2, s12–s20, 2012.

[3] E. Mannie, "Generalized multi-protocol label switching (gmpls) architecture", RFC Editor, RFC 3945, 2004.

[4] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "Openflow: Enabling innovation in campus networks", **SIGCOMM Comput. Commun. Rev.**, vol. 38, no. 2, pp. 69–74, Mar. 2008.

[5] R. Enns, M. Bjorklund, J. Schoenwaelder, and A. Bierman, "Network configuration protocol (netconf)", RFC Editor, RFC 6241, 2011, http://www.rfc-editor.org/rfc/rfc6241.txt.

[6] E. W. Dijkstra, "A note on two problems in connexion with graphs", **Numer. Math.**, vol. 1, no. 1, pp. 269–271, Dec. 1959.

[7] J. Yen, "Finding the k shortest loopless paths in a network", **management Science**, pp. 712–716, 1971.

[8] A. Carena, V. Curri, G. Bosco, P. Poggiolini, and F. Forghieri, "Modeling of the impact of nonlinear propagation effects in uncompensated optical coherent transmission links", **Journal of Lightwave Technology**, vol. 30, no. 10, pp. 1524–1539, 2012.

[9] I. Chlamtac, A. Ganz, and G. Karmi, "Lightpath communications: An approach to high bandwidth optical wan's", **IEEE Transactions on Communications**, vol. 40, no. 7, pp. 1171–1182, 1992.

[10] H. Zang and J. P. Jue, "A review of routing and wavelength assignment approaches for wavelength-routed optical wdm networks", **Optical Networks Magazine**, vol. 1, pp. 47–60, 2000.

[11] W. Shi, Z. Zhu, M. Zhang, and N. Ansari, "On the effect of bandwidth fragmentation on blocking probability in elastic optical networks", **Communications, IEEE Transactions on**, vol. 61, no. 7, pp. 2970–2978, 2013.

[12] Z. Lu, Y. Jayabal, M. Razo, M. Tacca, A. Fumagalli, G. M. Galimberti, G. Martinelli, and G. Swallow, "Optical layer-driven network restoration and redesign for improved fast reroute reliability", in **2016 18th International Conference on Transparent Optical Networks (ICTON)**, 2016, pp. 1–4.

[13] S. Kohei, I. Inoue, and E. Oki, "Multi-layer network operation and management for future carrier backbone networks", in **Global Telecommunications Conference, 2008. IEEE GLOBECOM 2008. IEEE**, 2008, pp. 1–5.

[14] A. D. Ellis, J. Zhao, and D. Cotter, "Approaching the non-linear shannon limit", **Journal of Lightwave Technology**, vol. 28, no. 4, pp. 423–433, 2010.

[15] R. J. Essiambre and R. W. Tkach, "Capacity trends and limits of optical communication networks", **Proceedings of the IEEE**, vol. 100, no. 5, pp. 1035–1055, 2012.

[16] B. Shariati, P. S. Khodashenas, J. M. Rivas-Moscoso, S. Ben-Ezra, D. Klonidis, F. Jiménez, L. Velasco, and I. Tomkos, "Investigation of mid-term network migration scenarios comparing multi-band and multi-fiber deployments", in **2016 Optical Fiber Communications Conference and Exhibition (OFC)**, 2016, pp. 1–3.

[17] X. white paper, "Nxwdm solution for 64 terabit optical networking", Tech. Rep., 2013.

[18] W. S. Pelouch, "Raman amplification: An enabling technology for long-haul coherent transmission systems", **Journal of Lightwave Technology**, vol. 34, no. 1, pp. 6–19, 2016.

[19] L. Barker, A. E. El-Taher, M. Alcon-Camas, J. D. Ania-Castanon, and P. Harper, "Extended bandwidth for long haul dwdm transmission using ultra-long raman fiber lasers", in **36th European Conference and Exhibition on Optical Communication**, 2010, pp. 1–3.

[20] A. E. Willner, "Communication with a twist", **IEEE Spectrum**, vol. 53, no. 8, pp. 34–39, 2016.

[21] R. M. Nejad, K. Allahverdyan, P. Vaity, S. Amiralizadeh, C. Brunet, Y. Messaddeq, S. LaRochelle, and L. A. Rusch, "Mode division multiplexing using orbital angular momentum modes over 1.4-km ring core fiber", **Journal of Lightwave Technology**, vol. 34, no. 18, pp. 4252–4258, 2016.

[22] J. Wang, "Twisted communications using orbital angular momentum", in **2016 Optical Fiber Communications Conference and Exhibition (OFC)**, 2016, pp. 1–1.

[23] P. Castoldi, T. Foggi, F. Paolucci, and F. Cugini, "Next steps in elasticity: Enabling signal overlap in optical networks", in **2016 18th International Conference on Transparent Optical Networks (ICTON)**, 2016, pp. 1–4.

[24] N. Wada, B. J. Puttnam, R. S. Luís, J. Sakaguchi, W. Klaus, J. M. D. Mendinueta, and Y. Awaji, "Space division multiplexing (sdm) transmission and related technologies", in **2016 18th International Conference on Transparent Optical Networks (ICTON)**, 2016, pp. 1–4.

[25] P. J. Winzer, "Spatial multiplexing: The next frontier in network capacity scaling", in **Optical Communication (ECOC 2013), 39th European Conference and Exhibition on**, 2013, pp. 1–4.

[26] R. Wang and B. Mukherjee, "Spectrum management in heterogeneous bandwidth networks", in **Global Communications Conference (GLOBECOM), 2012 IEEE**, 2012, pp. 2907–2911.

[27] E. Horowitz and S. Sahni, "Computing partitions with applications to the knapsack problem", **J. ACM**, vol. 21, no. 2, pp. 277–292, Apr. 1974.

[28] P. Wright, M. Parker, and A. Lord, "Minimum- and maximum-entropy routing and spectrum assignment for flexgrid elastic optical networking [invited]", **Optical Communications and Networking, IEEE/OSA Journal of**, vol. 7, no. 1, A66–A72, 2015.

[29] ——, "Simulation results of shannon entropy based flexgrid routing and spectrum assignment on a real network topology", in **Optical Communication (ECOC 2013), 39th European Conference and Exhibition on**, 2013, pp. 1–3.

[30] C. E. Shannon, "A mathematical theory of communication", **SIGMOBILE Mob. Comput. Commun. Rev.**, vol. 5, no. 1, pp. 3–55, Jan. 2001.

[31] D. Amar, E. Le Rouzic, N. Brochier, E. Bonetto, and C. Lepers, "Traffic forecast impact on spectrum fragmentation in gridless optical networks", in **Optical Communication (ECOC), 2014 European Conference on**, 2014, pp. 1–3.

[32] Y. Yin, M. Zhang, Z. Zhu, and S. Yoo, "Fragmentation-aware routing, modulation and spectrum assignment algorithms in elastic optical networks", in **Optical Fiber Communication Conference and Exposition and the National Fiber Optic Engineers Conference (OFC/NFOEC), 2013**, 2013, pp. 1–3.

[33] S. Yan, J. Kim, X. Wang, and A. Fumagalli, "Spectrum fragmentation analysis in a two-rate elastic optical link", in **Advanced Photonics 2015**, Optical Society of America, 2015, NeM2F.4.

[34] R. Wang and B. Mukherjee, "Provisioning in elastic optical networks with non-disruptive defragmentation", **Lightwave Technology, Journal of**, vol. 31, no. 15, pp. 2491–2500, 2013.

[35] Z. shu Shen, H. Hasegawa, and K.-I. Sato, "Integrity enhancement of flexible/semi-flexible grid networks that minimizes disruption in spectrum defragmentation and bitrate-dependent blocking", **Optical Communications and Networking, IEEE/OSA Journal of**, vol. 7, no. 4, pp. 235–247, 2015.

[36] F. Cugini, F. Paolucci, G. Meloni, G. Berrettini, M. Secondini, F. Fresi, N. Sambo, L. Poti, and P. Castoldi, "Push-pull defragmentation without traffic disruption in flexible grid optical networks", **Lightwave Technology, Journal of**, vol. 31, no. 1, pp. 125–133, 2013.

[37] N. Sambo, F. Paolucci, G. Meloni, F. Fresi, L. Poti, and P. Castoldi, "Control of frequency conversion and defragmentation for super-channels [invited]", **Optical Communications and Networking, IEEE/OSA Journal of**, vol. 7, no. 1, A126–A134, 2015.

[38] D. Siracusa, A. Broglio, A. Zanardi, E. Salvadori, G. Galimberti, and D. La Fauci, "Hitless network re-optimization to reduce spectrum fragmentation in distributed gmpls flexible optical networks", in **Optical Communication (ECOC 2013), 39th European Conference and Exhibition on**, 2013, pp. 1–3.

[39] M. Zhang, W. Shi, L. Gong, W. Lu, and Z. Zhu, "Bandwidth defragmentation in dynamic elastic optical networks with minimum traffic disruptions", in **Communications (ICC), 2013 IEEE International Conference on**, 2013, pp. 3894–3898.

[40] L. Altarawneh and S. Taebi, "Bandwidth granularity adaptation for multipath provisioning in elastic optical ofdm-based networks", in **Electro/Information Technology (EIT), 2015 IEEE International Conference on**, 2015, pp. 236–240.

[41] T. Fukuda, L. Liu, K.-I. Baba, S. Shimojo, and S. Yoo, "Fragmentation-aware spectrum assignment for elastic optical networks with fully-distributed gmpls", in **Optical Fiber Communications Conference and Exhibition (OFC), 2015**, 2015, pp. 1–3.

[42] H.-C. Le, L. Liu, and S. Ben Yoo, "Distributed control plane with spectral fragmentation-aware rmsa and flexible reservation for dynamic multidomain software-defined elastic optical networks", in **Optical Fiber Communications Conference and Exhibition (OFC), 2015**, 2015, pp. 1–3.

[43] J. Comellas, A. Asensio, M. Ruiz, G. Junyent, and L. Velasco, "Using spectrum management techniques to differentiate services in elastic optical networks", in **Transparent Optical Networks (ICTON), 2015 17th International Conference on**, 2015, pp. 1–4.

[44] Y. Hirota, Y. Hatada, T. Watanabe, and H. Tode, "Dynamic spectrum allocation based on connection alignment for elastic optical networks", in **Information and Telecommunication Technologies (APSITT), 2015 10th Asia-Pacific Symposium on**, 2015, pp. 1–3.

[45] D. Borquez, A. Beghelli, and A. Leiva, "Deadlock-avoiding vs. greedy spectrum allocation algorithms in dynamic flexible optical networks", in **Transparent Optical Networks (ICTON), 2015 17th International Conference on**, 2015, pp. 1–4.

[46] M. Zhang, W. Lu, Z. Zhu, Y. Yin, and S. Yoo, "Planning and provisioning of elastic o-ofdm networks with fragmentation-aware routing and spectrum assignment (rsa) algorithms", in **Communications and Photonics Conference (ACP), 2012 Asia**, 2012, pp. 1–3.

[47] M. Zhang, C. You, and Z. Zhu, "On the parallelization of spectrum defragmentation reconfigurations in elastic optical networks", **Networking, IEEE/ACM Transactions on**, vol. PP, no. 99, pp. 1–1, 2015.

[48] N. Wang, J. Jue, X. Wang, Q. Zhang, H. Cankaya, and M. Sekiya, "Holding-time-aware scheduling for immediate and advance reservation in elastic optical networks", in **Communications (ICC), 2015 IEEE International Conference on**, 2015, pp. 5180–5185.

[49] R. Proietti, L. Liu, R. Scott, B. Guan, C. Qin, T. Su, F. Giannone, and S. Yoo, "3d elastic optical networking in the temporal, spectral, and spatial domains", **Communications Magazine, IEEE**, vol. 53, no. 2, pp. 79–87, 2015.

[50] M. Imran, F. Paolucci, F. Cugini, A. D'Errico, L. Giorgi, T. Sasaki, P. Castoldi, and L. Poti, "Quasi-hitless software-defined defragmentation in space division multiplexing (sdm)", in **2015 European Conference on Optical Communication (ECOC)**, 2015, pp. 1–3.

[51] G. Meloni, F. Fresi, M. Imran, F. Paolucci, F. Cugini, A. D'Errico, L. Giorgi, T. Sasaki, P. Castoldi, and L. Pot, "Software-defined defragmentation in space-division multiplexing with quasi-hitless fast core switching", **Journal of Lightwave Technology**, vol. 34, no. 8, pp. 1956–1962, 2016.

[52] F. Paolucci, F. Cugini, A. Giorgetti, N. Sambo, and P. Castoldi, "A survey on the path computation element (pce) architecture", **Communications Surveys Tutorials, IEEE**, vol. 15, no. 4, pp. 1819–1841, 2013.

[53] F. Agraz, S. Azodolmolky, M. Angelou, J. Perello, L. Velasco, S. Spadaro, A. Francescon, C. Saradhi, Y. Pointurier, P. Kokkinos, E. Varvarigos, M. Gunkel, and I. Tomkos, "Experimental demonstration of centralized and distributed impairment-aware control plane schemes for dynamic transparent optical networks", in **Optical Fiber Communication (OFC), collocated National Fiber Optic Engineers Conference, 2010 Conference on (OFC/NFOEC)**, 2010, pp. 1–3.

[54] J. Sakaguchi, B. Puttnam, W. Klaus, Y. Awaji, N. Wada, A. Kanno, T. Kawanishi, K. Imamura, H. Inaba, K. Mukasa, R. Sugizaki, T. Kobayashi, and M. Watanabe, "19-core fiber transmission of 19 172-gb/s sdm-wdm-pdm-qpsk signals at 305tb/s", in **Optical Fiber Communication Conference and Exposition (OFC/NFOEC), 2012 and the National Fiber Optic Engineers Conference**, 2012, pp. 1–3.

[55] M. Cvijetic, I. B. Djordjevic, and N. Cvijetic, "Dynamic multidimensional optical networking based on spatial and spectral processing", **Opt. Express**, vol. 20, no. 8, pp. 9144–9150, 2012.

[56] Y. Sasaki, K. Takenaga, S. Matsuo, K. Aikawa, and K. Saitoh, "Few-mode multicore fibers for long-haul transmission line", **Optical Fiber Technology**, vol. 35, pp. 19 –27, 2017, Next Generation Multiplexing Schemes in Fiber-based Systems.

[57] F. M. Ferreira, N. M. Suibhne, C. S. Costa, M. Sorokina, S. Sygletos, and A. Ellis, "On the group delay statistics of few-mode fibres with intermediate linear mode coupling", in **2016 18th International Conference on Transparent Optical Networks (ICTON)**, 2016, pp. 1–5.

[58] X. Jin, A. Gomez, K. Shi, B. C. Thomsen, F. Feng, G. S. D. Gordon, T. D. Wilkinson, Y. Jung, Q. Kang, P. Barua, J. Sahu, S. u. Alam, D. J. Richardson, D. C. O'Brien, and F. P. Payne, "Mode coupling effects in ring-core fibers for space-division multiplexing systems", **Journal of Lightwave Technology**, vol. 34, no. 14, pp. 3365–3372, 2016.

[59] T. Hayashi, T. Taru, O. Shimakawa, T. Sasaki, and E. Sasaoka, "Characterization of crosstalk in ultra-low-crosstalk multi-core fiber", **Journal of Lightwave Technology**, vol. 30, no. 4, pp. 583–589, 2012.

[60] X. Jiang and B. Zhu, "Experimental evaluation of the crosstalk impacts of cascaded roadms supporting space division multiplexing", in **2016 Optical Fiber Communications Conference and Exhibition (OFC)**, 2016, pp. 1–3.

[61] N. Sambo, A. D'Errico, C. Porzi, V. Vercesi, M. Imran, F. Cugini, A. Bogoni, L. Poti', and P. Castoldi, "Sliceable transponder architecture including multi-wavelength source", **Optical Communications and Networking, IEEE/OSA Journal of**, vol. 6, no. 7, pp. 590–600, 2014.

[62] S. Randel, "Space division multiplexed transmission", in **Optical Fiber Communication Conference/National Fiber Optic Engineers Conference 2013**, Optical Society of America, 2013, OW4F.1.

[63] D. Marom and M. Blau, "Switching solutions for wdm-sdm optical networks", **Communications Magazine, IEEE**, vol. 53, no. 2, pp. 60–68, 2015.

[64] R. Ryf, S. Chandrasekhar, S. Randel, D. T. Neilson, N. K. Fontaine, and M. Feuer, "Physical layer transmission and switching solutions in support of spectrally and spatially flexible optical networks", **IEEE Communications Magazine**, vol. 53, no. 2, pp. 52–59, 2015.

[65] D. Siracusa, F. Pederzolli, P. S. Khodashenas, J. M. Rivas-Moscoso, D. Klonidis, E. Salvadori, and I. Tomkos, "Spectral vs. spatial super-channel allocation in sdm networks under independent and joint switching paradigms", in **2015 European Conference on Optical Communication (ECOC)**, 2015, pp. 1–3.

[66] N. K. Fontaine, T. Haramaty, R. Ryf, H. Chen, L. Miron, L. Pascar, M. Blau, B. Frenkel, L. Wang, Y. Messaddeq, S. LaRochelle, R. J. Essiambre, Y. Jung, Q. Kang, J. K. Sahu, S. U. Alam, D. J. Richardson, and D. M. Marom, "Heterogeneous space-division multiplexing and joint wavelength switching demonstration", in **2015 Optical Fiber Communications Conference and Exhibition (OFC)**, 2015, pp. 1–3.

[67] B. Shariati, P. S. Khodashenas, J. M. Rivas-Moscoso, S. Ben-Ezra, D. Klonidis, F. Jiménez, L. Velasco, and I. Tomkos, "Evaluation of the impact of different sdm switching strategies in a network planning scenario", in **2016 Optical Fiber Communications Conference and Exhibition (OFC)**, 2016, pp. 1–3.

[68] M. D. Feuer, "Optical routing for sdm networks", in **2015 European Conference on Optical Communication (ECOC)**, 2015, pp. 1–3.

[69] A. Muhammad, G. Zervas, and R. Forchheimer, "Resource allocation for space-division multiplexing: Optical white box versus optical black box networking", **Journal of Lightwave Technology**, vol. 33, no. 23, pp. 4928–4941, 2015.

[70] K. S. Abedin, T. F. Taunay, M. Fishteyn, M. F. Yan, B. Zhu, J. M. Fini, E. M. Monberg, F. Dimarcello, and P. Wisk, "Amplification and noise properties of an erbium-doped multicore fiber amplifier", **Opt. Express**, vol. 19, no. 17, pp. 16 715–16 721, 2011.

[71] G. Bernstein, Y. Lee, D. Li, W. Imajuku, and J. Han, "Routing and wavelength assignment information encoding for wavelength switched optical networks", RFC Editor, RFC 7581, 2015.

[72] G. Bernstein, S. Xu, Y. Lee, G. Martinelli, and H. Harai, "Signaling extensions for wavelength switched optical networks", RFC Editor, RFC 7689, 2015.

[73] H. Zhu, H. Zang, K. Zhu, and B. Mukherjee, "A novel generic graph model for traffic grooming in heterogeneous wdm mesh networks", **IEEE/ACM Transactions on Networking**, vol. 11, no. 2, pp. 285–299, 2003.

[74] A. Muhammad, G. Zervas, D. Simeonidou, and R. Forchheimer, "Routing, spectrum and core allocation in flexgrid sdm networks with multi-core fibers", in **2014 International Conference on Optical Network Design and Modeling**, 2014, pp. 192–197.

[75] S. Fujii, Y. Hirota, H. Tode, and K. Murakami, "On-demand spectrum and core allocation for multi-core fibers in elastic optical network", in **2013 Optical Fiber Communication Conference and Exposition and the National Fiber Optic Engineers Conference (OFC/NFOEC)**, 2013, pp. 1–3.

[76] ——, "On-demand spectrum and core allocation for reducing crosstalk in multicore fibers in elastic optical networks", **IEEE/OSA Journal of Optical Communications and Networking**, vol. 6, no. 12, pp. 1059–1071, 2014.

[77] S. Fujii, Y. Hirota, and H. Tode, "Dynamic resource allocation with virtual grid for space division multiplexed elastic optical network", in **39th European Conference and Exhibition on Optical Communication (ECOC 2013)**, 2013, pp. 1–3.

[78] Y. Zhang, L. Yan, H. Wang, and W. Gu, "Routing, wavelength and mode assignment algorithm for space division multiplexing transmission network", in **2012 Second International Conference on Instrumentation, Measurement, Computer, Communication and Control**, 2012, pp. 1383–1385.

[79] D. Siracusa, F. Pederzolli, D. Klonidisz, V. Lopezy, and E. Salvadori, "Resource allocation policies in sdm optical networks (invited paper)", in **2015 International Conference on Optical Network Design and Modeling (ONDM)**, 2015, pp. 168–173.

[80] F. Pederzolli, D. Siracusa, J. M. Rivas-Moscoso, B. Shariati, E. Salvadori, and I. Tomkos, "Spatial group sharing for sdm optical networks with joint switching", in **2016 International Conference on Optical Network Design and Modeling (ONDM)**, 2016, pp. 1–6.

[81] P. S. Khodashenas, J. M. Rivas-Moscoso, D. Siracusa, F. Pederzolli, B. Shariati, D. Klonidis, E. Salvadori, and I. Tomkos, "Comparison of spectral and spatial super-channel allocation schemes for sdm networks", **Journal of Lightwave Technology**, vol. 34, no. 11, pp. 2710–2716, 2016.

[82] F. Pederzolli, M. Gerola, A. Zanardi, X. Forns, J. F. Ferran, and D. Siracusa, "Experimental evaluation of yamato, a sdn control plane for joint and fractional-joint switched sdm optical networks", in **ECOC 2016; 42nd European Conference on Optical Communication**, 2016, pp. 1–3.

[83] B. Shariati, D. Klonidis, D. Siracusa, F. Pederzolli, J. M. Rivas-Moscoso, L. Velasco, and I. Tomkos, "Impact of traffic profile on the performance of spatial superchannel switching in sdm networks", in **ECOC 2016; 42nd European Conference on Optical Communication**, 2016, pp. 1–3.

[84] F. Pederzolli, D. Siracusa, B. Shariati, J. M. Rivas-Moscoso, E. Salvadori, and I. Tomkos, "Improving performance of spatially joint- switched space division multiplexing optical networks via spatial group sharing", **IEEE/OSA Journal of Optical Communications and Networking**, vol. 9, no. 3, B1–B11, 2017.

[85] F. Pederzolli, M. Gerola, A. Zanardi, X. Forns, J. F. Ferran, and D. Siracusa, "Yamato: The first sdn control plane for independent, joint, and fractional-joint switched sdm optical networks", **Journal of Lightwave Technology**, vol. 35, no. 8, pp. 1335–1341, 2017.

[86] M. Klinkowski, P. Lechowicz, and K. Walkowiak, "Survey of resource allocation schemes and algorithms in spectrally-spatially flexible optical networking", **Optical Switching and Networking**, vol. 27, no. Supplement C, pp. 58 –78, 2018.

[87] A. Clemm, J. Medved, R. Varga, N. Bahadur, H. Ananthakrishnan, and X. Liu, "A data model for network topologies", IETF Secretariat, Internet-Draft draft-ietf-i2rs-yang-network-topo-20, 2017, http://www.ietf.org/internet-drafts/draft-ietf-i2rs-yang-network-topo-20.txt.

[88] P. Poggiolini, "The gn model of non-linear propagation in uncompensated coherent optical systems", **Journal of Lightwave Technology**, vol. 30, no. 24, pp. 3857–3879, 2012.

[89] G. Martinelli, X. Zhang, G. Galimberti, D. Siracusa, A. Zanardi, F. Pederzolli, Y. Lee, and F. Zhang, "Information encoding for wson with impairments validation", IETF Secretariat, Internet-Draft draft-martinelli-ccamp-wson-iv-encode-08, 2017, http://www.ietf.org/internet-drafts/draft-martinelli-ccamp-wson-iv-encode-08.txt.

[90] F. Rambach, B. Konrad, L. Dembeck, U. Gebhard, M. Gunkel, M. Quagliotti, L. Serra, and V. Lopez, "A multilayer cost model for metro/core networks", **IEEE/OSA Journal of Optical Communications and Networking**, vol. 5, no. 3, pp. 210–225, 2013.

[91] R. Martinez, F. Cugini, R. Casellas, F. Paolucci, R. Vilalta, P. Castoldi, and R. Munoz, "Control plane solutions for sliceable bandwidth transceiver configuration in flexi-grid optical networks", **IEEE Communications Magazine**, vol. 54, no. 8, pp. 126–135, 2016.

[92] P. S. Khodashenas, J. M. Rivas-Moscoso, D. Klonidis, D. M. Marom, and I. Tomkos, "Evaluating the optimum filter resolution and sub-channel spectrum granularity for flexible super-channels", in **2015 Optical Fiber Communications Conference and Exhibition (OFC)**, 2015, pp. 1–3.

[93] P. S. Khodashenas, J. M. Rivas-Moscoso, B. Shariati, D. M. Marom, D. Klonidis, and I. Tomkos, "Investigation of spectrum granularity for performance optimization of flexible nyquist-wdm-based optical networks", **Journal of Lightwave Technology**, vol. 33, no. 23, pp. 4767–4774, 2015.

[94] L. Velasco, F. Morales, L. Gifre, A. Castro, O. G. de Dios, and M. Ruiz, "On-demand incremental capacity planning in optical transport networks", **IEEE/OSA Journal of Optical Communications and Networking**, vol. 8, no. 1, pp. 11–22, 2016.

[95] B. Shariati, P. S. Khodashenas, J. M. R. Moscoso, S. Ben-Ezra, D. Klonidis, F. Jimenez, L. Velasco, and I. Tomkos, "Evaluation of the impact of different sdm switching strategies in a network planning scenario", in **Optical Fiber Communication Conference**, Optical Society of America, 2016, Tu2H.4.

[96] L. Pascar, R. Karubi, B. Frenkel, and D. M. Marom, "Port-reconfigurable, wavelength-selective switch array for colorless/directionless/contentionless optical add/drop multiplexing", in **2015 International Conference on Photonics in Switching (PS)**, 2015, pp. 16–18.

[97] **Inspace eu project [web]**, http://www.ict-inspace.eu/.

[98] L. R. Ford and D. R. Fulkerson, **Flows in Networks**. Princeton University Press, 1962.

[99] **Opendaylight [web]**, https://www.opendaylight.org/.

[100] A. Bierman, M. Bjorklund, and K. Watsen, "Restconf protocol", RFC Editor, RFC 8040, 2017.

[101] M. Bjorklund, "Yang - a data modeling language for the network configuration protocol (netconf)", RFC Editor, RFC 6020, 2010, http://www.rfc-editor.org/rfc/rfc6020.txt.