



Dynamic Protection of Optical Networks

Ruepp, Sarah Renée

Publication date:
2008

Document Version
Publisher's PDF, also known as Version of record

[Link back to DTU Orbit](#)

Citation (APA):
Ruepp, S. R. (2008). *Dynamic Protection of Optical Networks*.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Dynamic Protection of Optical Networks

Sarah Renée Ruepp

April 2008



Networks Technologies and Service Platforms
DTU Fotonik
Technical University of Denmark
Building 343
2800 Kgs. Lyngby
DENMARK

Preface

This thesis presents a selection of the research carried out during my Ph.D. study at the Networks Competence Area, DTU Fotonik, Technical University of Denmark, in the period from January 2005 to April 2008.

This Ph.D. study is carried out as a so-called Innovations Ph.D., where potential commercialization aspects of the concepts worked with during the Ph.D. study are reviewed in addition to technical research work. The innovative part of the Ph.D. study was documented in two confidential reports, and approved by the DTU Innovation Committee in 2006 and 2007, respectively. According to the study directives for Innovation Ph.D. studies, the Ph.D. degree shall be awarded based on technical research work, as presented in this thesis.

The thesis is entitled "Dynamic Protection of Optical Networks" and the presented methods aim at increasing network survivability in a dynamic network environment.

Abstract

This thesis deals with making optical networks resilient to failures. The recovery performance of path, segment and span restoration is evaluated in a network with limited wavelength conversion capability using both standard and enhanced wavelength assignment schemes. The enhanced wavelength assignment scheme is based on the Suggested Vector (SV), which is a Generalized Multi-Protocol Label Switching (GMPLS) compliant signalling extension aiming at wavelength conversion minimization.

To increase the recovery percentage, two modifications of the signalling session are proposed and evaluated through simulation. By resolving wavelength contention, the blocking reduction scheme reduces the number of necessary recovery retries and thereby the restoration time and control plane load. The stub-awareness schemes avoids wavelength conversions when merging the restoration segment to the connection stubs at the failure adjacent nodes. Both modifications have a positive influence on the recovery percentage. The recovery enhancements are applicable in both single and multi-domain network environments.

Stub release, where the still working parts of a failure affected connection are released prior to recovery initiation, has shown to increase the recovery percentage. If the lack of span resources limits successful connection recovery, both span and node resources (i.e., Wavelength Converters (WCs)) should be released, while it is sufficient to release the WCs if they are the limiting factor.

The effect of the modularity of capacity units is investigated for resilient network design. Different span upgrading strategies and algorithms for finding restoration paths are evaluated. Furthermore, the capacity efficiency of constraining restoration requests for the same destination node to the same restoration path is evaluated.

Resumé

Denne afhandling behandler metoder, hvormed optiske netværk kan gøres modstandsdygtige over for fejl. Betydningen af "path", "segment" og "span" restoration evalueres i et netværk med begrænsede muligheder for bølgelængdekonvertering, ved brug af en almindelig såvel som en forbedret metode for tildeling af bølgelængder. Metoden for forbedret tildeling af bølgelængder er baseret på "Suggested Vector" som er et GMPLS-kompatibel signaleringsobjekt, som kan minimere brugen af bølgelængdekonvertering.

For at øge fejlretningseffektiviteten opstilles to modificerede signaleringsmodeller, som dernæst evalueres ved simulering. "Blocking reduction" metoden benyttes som en løsning på problemet med konkurrerende bølgelængder og reducerer dermed antallet af genopretningsforsøg, hvilket medfører hurtigere genoprettelse og mindre belastning af kontrolsystemet. "Stub-awareness" metoden tilrettelægger brugen af bølgelængder på backup-ruten, således at unødvendig bølgelængdekonvertering i knudepunkterne nærmest fejlen undgås. Begge metoder forbedrer endvidere genoprettelsesprocenten. De beskrevne metoder kan anvendes i såvel enkelt som multi-domæne netværk.

Genoprettelseeffektiviteten kan yderligere forbedres, hvis de fungerende dele af en ellers fejllende forbindelse frigives inden genopretningsforsøget påbegyndes. Denne proces kaldes "stub release". Hvis mangel på båndbredde hæmmer genoprettelsen, bør både båndbredde og tilhørende knudepunktsressourcer (f. eks. bølgelængdekonverterere) frigives. Ligger begrænsningen derimod i knudepunkterne, er det oftest tilstrækkeligt blot at frigive knudepunktsressourcerne.

Effekten af en modulær kapacitetsopbygning undersøges med henblik på at minimere den totale båndbredde i et pålideligt designet netværk.

Forskellige kapacitetsopgraderingsstrategier og algoritmer til fastlæggelse af backup-ruter evalueres. Desuden undersøges kapacitetseffektiviteten når samtlige forbindelser til en given destination fastlæses til samme backup-rute.

Acknowledgements

The work presented in this thesis would not have been possible without the support of a large number of people, whom I would like to express my gratitude to.

I would like to thank my supervisor Prof. Lars Dittmann for giving me the opportunity to carry out a Ph.D. study in the Networks Area at DTU Fotonik. Also, I would like to thank my co-supervisors Lars Ellegaard, Henrik Christiansen and Peter Tøttrup. The supervisor team has given me inspiration and feedback throughout the project period, which I am very grateful for.

My warmest thanks go to Jakob Buron, for the many hours of joint research, student supervision and for the inspiring and fun collaboration throughout the last three years. Furthermore, I would like to thank Nicola Andriolli, whose enthusiasm and hospitality has made our cooperation as well as visiting SSSUP in Pisa a great experience. Thanks to Anna Manolova, Henrik Wessing, Thomas Stidsen and Arie Koster for the fruitful and enjoyable collaboration. Special thanks go to Arie Koster for welcoming me to Warwick University. I would like to acknowledge the help of IT-Support, and Anders Fosgerau in particular, for setting up dedicated OPNET simulation PCs.

I am indebted to Jakob Buron, Anna Manolova, Henrik Wessing, Lars Ellegaard, Henrik Christiansen and Lars Dittmann for proofreading of this thesis.

Thanks to Jose Soler and Anna Manolova for sharing the office, and to all my colleagues in the Networks Area for providing a great working atmosphere - I have really enjoyed being part of the group.

I would like to thank my family, and in particular my parents, for their love and encouragement. Finally, I would like to express my grat-

itude to Kim for his love, support and understanding, and for always being there for me, especially during the writing of this thesis.

Kgs. Lyngby, April 2008

Sarah Renée Ruepp

Ph.D. Publications

This Ph.D. project has resulted in 24 peer-reviewed publications in total, split into 4 journal publications and 20 peer-reviewed conference contributions. The publications are listed below.

- [1] A. Koster and **S. Ruepp** (2008). "Benchmarking RWA Strategies for Dynamically Controlled Optical Networks", 8 pp., in Proceedings of Networks 2008, Budapest, Hungary
- [2] J. Buron, **S. Ruepp**, H. Wessing, N. Andriolli, A.V. Manolova, and L. Dittmann. (2008). "Wavelength Converter Placement in Optical Networks with Dynamic Traffic", 5 pp., in Proceedings of APOC 2008, Hangzhou, China
- [3] N. Andriolli, A. Giorgetti, **S. Ruepp**, J. Buron, L. Valcarenghi and P. Castoldi. (2008). "Bidirectional Lightpath Provisioning in GMPLS-controlled Optical Networks", 2 pp., in Proceedings of Photonics in Switching 2008, Japan
- [4] X. Li, **S. Ruepp**, L. Dittmann and A.V. Manolova. (2008). "Survivability-Enhancing Routing Scheme for Multi-Domain Networks", 5 pp., in Proceedings of Globecom 2008, New Orleans, USA

- [5] X. Li, **S. Ruepp**, A.V. Manolova, and L. Dittmann. (2008). "OPNET Model of Multi-domain Routing with Enhanced Survivability", in Proceedings of OPNETWORK 2008, Washington DC, USA
- [6] **S. Ruepp**, N. Andriolli, J. Buron, L. Dittmann and L. Ellegaard. (2008). "Restoration in All-Optical GMPLS Networks with Limited Wavelength Conversion", 18 pp. in Computer Networks Journal, Special Issue on Opportunities and Challenges in Optical Networks
- [7] C. Raffaelli, K. Vlachos, N. Andriolli, J. Buron, R. van Caenegem, G. Danilewicz, J. M. Finochietto, J. Garcia-Haro, D. Klonidis, M. O'Mahony, G. Maier, A. Pattavina, P. Pavon-Marino, **S. Ruepp**, M. Savi, M. Scaffardi, I. Tomkos, A. Tzanakaki, L. Wosinska, K. Yannopoulos and F. Neri (2008). "Photonics in Switching: architectures, technologies and systems The e-Photon/ONE+ research results - Part I : system aspects". In Computer Networks Special Issue on Opportunities and Challenges in Optical Networks, authors in alphabetical order
- [8] **S. Ruepp**, J. Buron, N. Andriolli and L. Dittmann. (2008). "Nodal Stub-Release in All-Optical Networks", IEEE Communication Letters, vol. 12, no. 1, pp. 47-49, Jan. 2008
- [9] A.V. Manolova, J. Buron, **S. Ruepp**, Lars Dittmann and Lars Ellegaard. (2007). "Modelling Contention-Resolution in Optical Burst Switching" in Proceedings of Opnetwork 2007, Session 1339, 6 pp., Washington DC, USA, August 2007
- [10] **S. Ruepp**, J. Buron, N. Andriolli and H. Wessing (2007). "Span Restoration in Optical Networks with Limited Wavelength Conversion", in Proceedings of Chinacom 2007, 5 pp., Shanghai, China, August 2007

-
- [11] P. Castoldi, N. Andriolli, I. Cerutti, N. Sambo, L. Valcarenghi, A. Giorgetti, **S. Ruepp** and J. Buron (2007). "A Framework for Label Preference in GMPLS Controlled Optical Networks", in Proceedings of 2007 International Conference on Transparent Optical Networks (ICTON 2007), 4 pp., Rome, Italy, July 2007 (invited)
- [12] A.V. Manolova, J. Buron, **S. Ruepp**, Lars Dittmann and Lars Ellegaard. (2007). "Segmentation-based Path Switching Mechanism for Reduced Data Losses in OBS Networks", in Proceedings of Optical Network Design and Modelling Conference (ONDM), Athens, Greece, May 2007, Published in Lecture Notes in Computer Science (LNCS) 4534, pp. 338-347
- [13] A.V. Manolova, **S. Ruepp**, J. Buron, Lars Dittmann and Lars Ellegaard. (2007). "Advantages and Challenges of the GMPLS/OBS Integration", in Proceedings of The VI Workshop in G/MPLS Networks, pp. 133-144, Girona, Spain, April 2007
- [14] N. Andriolli, J. Buron, **S. Ruepp**, F. Cugini, L. Valcarenghi and P. Castoldi (2006). "Label Preference Schemes in GMPLS Controlled Networks", IEEE Communication Letters, vol. 10, no. 12, pp. 849-851, Dec. 2006
- [15] **S. Ruepp**, J. Buron and N. Andriolli. (2006). "Increasing Restorability for Local-to-Egress Restoration in GMPLS Controlled Networks with Limited Wavelength Conversion", in Proceedings of The Tenth IEEE International Conference on Communication Systems (ICCS), 5 pp., Singapore, October/November 2006

- [16] J. Buron, **S. Ruepp** and N. Andriolli. (2006). "Blocking Reduction of Span Restoration Requests in GMPLS Controlled WDM Optical Networks", in Proceedings of European Conference on Optical Communication (ECOC), 2 pp., P.118, Cannes, France, September 2006
- [17] **S. Ruepp**, N. Andriolli and J. Buron. (2006). "Signaling protocol extensions enhancing provisioning and recovery in GMPLS optical networks", in Proceedings of Opnetwork 2006, Session 1339, 7 pp., Washington DC, USA, August 2006
- [18] J. Buron, **S. Ruepp** and A. V. Manolova. (2006). "Teaching cost-effective and resilient network design with OPNET WDM Guru", in Proceedings of Opnetwork 2006, Session 1020, 6 pp., Washington DC, USA, August 2006
- [19] N. Andriolli, J. Buron, **S. Ruepp** and P. Castoldi. (2006). "Modeling GMPLS-controlled Optical Packet Switched (OPS) Networks", in Proceedings of Opnetwork 2006, Session 1339, 8 pp., Washington DC, USA, August 2006
- [20] **S. Ruepp**, J. Buron and N. Andriolli. (2006). "Wavelength-Converter Saving Span Restoration in GMPLS Controlled WDM Networks", in Proceedings of The Joint International Conferences on Optical Internet and Next Generation Networks (COIN-NGN), 3 pp., Jeju Island, South Korea, July 2006
- [21] N. Andriolli, J. Buron, **S. Ruepp**, F. Cugini, L. Valcarengi and P. Castoldi. (2006). "Signaling protocol extensions for converter-saving wavelength assignment in GMPLS optical networks", in Proceedings of IEEE Workshop on High Performance Switching and Routing (HPSR2006), 7 pp., Poznan, June 7-9, 2006.

- [22] N. Andriolli, H. Harai, J. Buron, **S. Ruepp** and P. Castoldi. (2006). "Wavelength Assignment Schemes for Connection-Oriented OPS Networks", in Proceedings of 10th Conference on Optical Network Design and Modelling (ONDM 2006), 8 pp., Copenhagen, May 22-24, 2006.
- [23] T. Stidsen and **S. Ruepp**. (2006). "Shortcut Span Protection", in Proceedings of The 8th Informs Telecom Conference, 3 pp., Dallas, TX, USA, March/April 2006
- [24] **S. Ruepp**, L. Dittmann and L. Ellegaard. (2005). "Simulation and Comparison of Path Restoration Techniques in SDH Mesh Networks", in Proceedings of The 5th International Workshop on Design of Reliable Communications Networks (DRCN), pp. 47-53, Ischia, Italy, October 2005

Contents

Preface	i
Abstract	iii
Resumé	v
Acknowledgements	vii
Ph.D. Publications	ix
1 Introduction	1
1.1 Motivation	1
1.2 Structure of Thesis	2
1.3 Sections Based on Publications	2
2 Optical Networking Background	5
2.1 Introduction	5
2.2 Data Plane Technologies	7
2.2.1 Switching and Wavelength Conversion	7
2.2.2 Network Topologies	11
2.3 Network Control Plane	12
2.3.1 Control Plane Functions	12
2.3.2 Control Plane Standardization	12
2.3.3 ASON	13
2.3.4 GMPLS	14
2.3.5 UNI and NNI	14
2.4 Network Management Plane	15
2.5 Network Survivability	16

2.5.1	Failure Types	16
2.5.2	Basic Requirements	16
2.5.3	Criteria for Recovery Method Evaluation	18
2.5.4	Fault Management	19
2.5.5	Protection vs. Restoration	19
2.5.6	Linear Protection	21
2.5.7	Ring Protection	23
2.5.8	Mesh Protection	23
2.5.9	Combining Ring and Mesh Protection	25
2.5.10	Reversion	25
2.5.11	Stub Release	25
2.5.12	Single vs. Multi-Layer Recovery	27
2.5.13	Multiple Failures	28
2.5.14	Control Plane Failures	29
2.6	Chapter Summary	29
3	Dynamic Recovery with Limited Wavelength Conversion	31
3.1	Introduction	31
3.2	Previous Work	32
3.3	Network Model	34
3.3.1	Node Architecture	34
3.3.2	Network Architecture	34
3.3.3	Route Computation and Signaling	35
3.4	Wavelength Assignment Schemes	37
3.4.1	Standard Wavelength Assignment - The Label Set	37
3.4.2	Enhanced Wavelength Assignment - The Suggested Vector	39
3.5	Simulation Study	40
3.5.1	Case studies	41
3.5.2	Simulation Parameters	41
3.5.3	Simulation Execution	42
3.6	Results	44
3.6.1	Case 1 - Isolated Restoration Phase	44
3.6.2	Case 2 - Correlated Provisioning and Restoration Phase	49
3.7	Summary and Conclusions on Restoration with Limited Conversion	53

4	Signaling-Based Recovery Enhancements	55
4.1	Introduction	55
4.2	Blocking Reduction of Span Restoration Requests	56
4.2.1	Motivation	56
4.2.2	Blocking Reduction Concept	57
4.2.3	Simulation Study	59
4.2.4	Results	60
4.2.5	Conclusions on Blocking Reduction	64
4.3	Stub-Aware Span Restoration	65
4.3.1	Motivation	65
4.3.2	Method and Operation	66
4.3.3	Simulation Study	68
4.3.4	Results	70
4.3.5	Conclusions on Stub Awareness	72
4.4	Recovery in Multi-Domain Networks	73
4.4.1	Background	73
4.4.2	Policies	73
4.4.3	Multi-Domain Recovery	74
4.4.4	Evaluation Concept	75
4.4.5	Conclusions on Multi-Domain Recovery	76
4.5	Chapter Summary	76
5	Stub Release	79
5.1	Introduction	79
5.2	Full Stub Release	80
5.2.1	Motivation	80
5.2.2	Releasing Span and Node Resources	80
5.2.3	Simulation Study	80
5.2.4	Results	84
5.2.5	Conclusions on Full Stub Release	87
5.3	Nodal Stub Release	88
5.3.1	Motivation	88
5.3.2	Concept and Operation	88
5.3.3	Simulation Study	92
5.3.4	Results	94
5.3.5	Conclusions on Nodal Stub Release	97
5.4	Chapter Summary	97

6	Resilient Network Design	99
6.1	Introduction	99
6.2	Design of Protected Networks with Modular Capacity . .	100
6.2.1	Background	100
6.2.2	Restoration Path Computation	101
6.2.3	Span Upgrading Strategies	104
6.2.4	Path Restoration Simulation	106
6.2.5	Simulation Results	113
6.2.6	Conclusions on Networks with Modular Capacity .	116
6.3	Shortcut Span Protection	118
6.3.1	Motivation	118
6.3.2	Simulation Study	118
6.3.3	Results	120
6.3.4	Conclusions on Shortcut Span Protection	121
6.4	Chapter Summary	122
7	Conclusion and Outlook	123
Appendix A	OPNET Model Details	129
A.1	Introduction	129
A.2	GMPLS Signaling OPNET Model	129
A.2.1	Simulation Parameters	129
A.2.2	Network Model	130
A.2.3	Node Model	131
A.3	Connection Provisioning	132
A.3.1	Request Generation	132
A.3.2	Handling RSVP-TE Messages	132
A.4	Connection Recovery	134
A.4.1	Failure Notification	136
A.4.2	Recovery Operation	136
A.5	Result Presentation	136
A.5.1	Metrics	136
A.5.2	Confidence Intervals	138
A.6	Appendix Summary	139

Appendix B Connection Provisioning	141
B.1 Introduction	141
B.2 Label Preference Schemes in GMPLS Controlled Networks	141
B.2.1 Introduction	141
B.2.2 Connection Setup in GMPLS	142
B.2.3 Tie-Breaking Policies	143
B.2.4 Wavelength/Label Assignment Schemes	143
B.2.5 Summary of Concept for Label Preference Schemes	148
B.3 Wavelength Converter Placement For Dynamic Traffic . .	150
B.3.1 Introduction	150
B.3.2 Converter Placement Heuristics	150
B.3.3 Summary of Concept for Converter Placement Heuris-	
tics	155
B.4 Bidirectional Lightpath Provisioning	156
B.4.1 Introduction	156
B.4.2 Bidirectional GMPLS Lightpath Setup	156
B.4.3 Summary of Concept on Bidirectional Lightpath	
Provisioning	158
B.5 Benchmarking RWA Schemes	160
B.5.1 Introduction	160
B.5.2 Summary of Benchmarking RWA Schemes	161
B.6 Appendix Summary	162
Appendix C Contributions to Papers	163
Bibliography	167
List of Acronyms	183

Chapter 1

Introduction

1.1 Motivation

Our society is increasingly relying on functional communication services. A wide range of services that we use in our everyday life, such as telephone services, Internet browsing, emails, emergency services, business-critical applications, web-banking and lately Voice-over-IP and IP-TV, require functional communication networks. Hence, we require that these services are always available when we want to use them.

Unfortunately, communication networks are affected by many types of errors, both errors occurring within the network and through external impacts. Failures are typically caused by construction workers digging up cables, software errors and misconfigurations, hardware faults, power losses or natural disasters. Even with greatest precaution failures cannot be avoided. Failure situations should therefore be anticipated already when deciding how the network is designed and controlled. The goal of network survivability is to minimize impact of the failures, aiming at the customers not even noticing that a failure occurred.

Despite recent progression within the fields of network control and data transmission, there are still several open questions on how the latest technologies can be applied for the benefit of network survivability. This Ph.D. thesis addresses the challenge of making optical communication networks resilient to failures. The main focus is on dynamic protection, also called restoration, which ensures flexibility and adaptability to a variety of failure scenarios. In the thesis, several methods for enhancing

network survivability are presented, covering network operation, design and efficient span and node resource usage. The presented methods aim at providing high network survivability while keeping the complexity low.

1.2 Structure of Thesis

This Ph.D. study resulted in 24 peer-reviewed journal and conference contributions [1–24]. The publications are within the fields of network survivability, automated connection provisioning, inter-domain routing issues, optical burst switching and optical packet switching. The main focus of this Ph.D. project is on network survivability and therefore, only publications relevant to this field are presented in the main part of this thesis. Some of the network survivability studies stem upon network provisioning concepts, and therefore a summary of these studies is given in the appendix.

The remainder of this thesis is organized as follows: chapter 2 gives an introduction to optical networking concepts, connection switching and control, and network survivability. In chapter 3, three restoration methods are evaluated in a network with limited wavelength conversion capability. Chapter 4 presents signaling-based recovery enhancements in both single and multi-domain environments. The concept of stub release is evaluated in chapter 5, where the cases of releasing both node and span resources, and exclusively releasing node resources are investigated. Chapter 6 presents network design with modular capacity and also evaluates the capacity usage of bundled recovery requests. Conclusions and outlooks for the work presented in this thesis are given in Chapter 7.

In Appendix A, the OPNET model used for the connection provisioning and recovery studies is shown. Appendix B summarizes work on network provisioning, such as different wavelength assignment schemes, converter placement heuristics, bidirectional lightpath provisioning and benchmarking of Routing and Wavelength Assignment (RWA) schemes.

1.3 Sections Based on Publications

Some of the text presented here was written specifically for this thesis, while the research part is mainly based on publications. This section

describes which parts of this thesis are based on research papers and table 1.1 maps the thesis sections and appendices to the respective research papers. A detailed description of my contribution to the individual papers can be found in Appendix C.

Section	Paper	Author number
3	Based on [6]	First
4.2	Based on [16]	Second
4.3	Combination of [10] and [20]	First, First
4.4	Based on [10]	First
5.2	Based on [15]	First
5.3	Based on [8]	First
6.2	Based on [24]	First
6.3	Based on [23]	Second
A	Based on [17]	First
B.2	Summary of [14] and [21]	Third, Third
B.3	Summary of [2]	Second
B.4	Summary of [3]	Third
B.5	Summary of [1]	Second

Table 1.1: Sections of thesis based on research papers.

Chapter 2

Optical Networking Background

2.1 Introduction

Optical networks have been developed as the solution to accommodate the high bandwidth demands of our information society. Traditionally, optical networks have been controlled in a centralized structure with manual interaction. This approach is however inefficient with a dynamic traffic pattern, which creates the need for distributed and highly flexible control mechanisms, where the network can be self-controlling by automatically setting up, tearing down and recovering connections when required. According to the International Telecommunications Union (ITU) an optical network can be modelled as a *transport* plane¹, a *control* plane and a *management* plane [25], shown in figure 2.1. The transport plane refers to the logic and hardware responsible for the physical transfer of data. The control plane covers the infrastructure and distributed intelligence that controls the establishment and recovery of connections in the network. The management plane encompasses systems, interfaces and protocols used to manage the network and its services [26].

During the past years, significant development of both the control plane and data plane has occurred. Generalized Multi-Protocol Label Switching (GMPLS) [27], Automatically Switched Optical Network

¹The transport plane is also called the data plane.




Figure 2.1: Data, control and management plane.

(ASON) [28], and the User-Network Interface (UNI) [29]/Network-to-Network Interface (NNI) [30] specifications are emerging as promising candidates to dynamically control optical networks. Considerable development has also occurred within the data plane. Reconfigurable Optical Add-Drop Multiplexers (ROADMs) and Optical Cross-Connects (OXC), having true optical cores allowing rearrangeability between ports [31] and possibly change of wavelength, have recently been accepted as viable solutions by industry. These components represent a new generation of optical network elements seeking to gradually realize the all-optical network vision hitherto mostly considered in academia.

The interaction of the control and data plane is discussed in this thesis with the main focus of increasing network survivability. This chapter provides some background information on the technologies and mechanisms that are used in the remaining chapters. Section 2.2 introduces data plane technologies. Section 2.3 describes how an optical network can be controlled and discusses standardization efforts. In section 2.4, a brief introduction to network management is given. Section 2.5 presents the requirements for successful connection recovery and gives an overview of survivability methods for various network topologies. Section 2.6 summarizes the concepts presented in this chapter.

2.2 Data Plane Technologies

In this section, connection switching is described with particular focus on how to convert between different wavelengths. Furthermore, different network topologies are presented.

2.2.1 Switching and Wavelength Conversion

In an optical network, each fiber can carry vast amounts of data. It is however very unlikely that a single connection will fill up the fiber. Therefore, the transmission window of a fiber can be divided into different wavelengths (i.e., frequencies)². These systems are called Wavelength Division Multiplexing (WDM) systems. Each wavelength channel can now be used to carry data independently.

At the network nodes two options exist for switching the optical signal. Either, the signal is first converted to the electronic domain, where it is processed, before it is converted back to the optical domain and sent on. This type of switching is called opaque. The advantage of this approach is that the signal can be monitored in the electrical domain and it can be converted to a different wavelength. The conversion to the electrical domain however introduces additional delay and cost [33]. To overcome these drawbacks and to simplify pass-through switching of all wavelengths in a fiber, it is desirable to keep the signal within the optical domain. This type of switching is called transparent. The fact that conversion to the electrical domain is avoided however complicates traffic monitoring, signal regeneration and wavelength conversion [34].

Wavelength conversion is an important capability of the switching nodes, because being able to change wavelengths along the path significantly reduces the blocking probability [35,36]. Setting up a connection in a WDM network traditionally required that the connection is allocated to the same wavelength on its entire route (the so-called wavelength continuity constraint). Emerging technologies, i.e., Wavelength Converters (WCs), allow for the conversion between wavelengths, and thereby allow setting up a lightpath even if it is impossible to reserve the same wavelength on all hops of an envisaged path. An example of a lightpath that would be blocked without wavelength conversion is shown

²Recent studies report 164 wavelength channels at 100 Gb/s over 2550 km [32].




Figure 2.2: Wavelength conversion used in lightpath setup. Free wavelengths are denoted at each span. Color transitions in nodes indicate wavelength conversion.

in figure 2.2.

Wavelength Conversion Techniques

An ideal WC should be able to convert between the entire available wavelength range, ensure fast setup to the chosen output wavelength, should not impair the signal, and be simple and cheap to implement [35]. Wavelength conversion techniques can be classified into two types [33]:

- Opto-electronic wavelength conversion: the optical signal is first converted into an electronic signal, buffered, and then used to drive the input of a tunable laser.
- All-optical wavelength conversion: the signal remains in the optical domain. The implementation of WCs can be fiber-based or semiconductor-based.

Implementation examples for both opto-electronic and all-optical wavelength converters and a comparison of the different methods can be found in [33,35,37,38]. Some WCs may only allow conversion between a limited range of wavelengths [39,40], increasing the complexity of routing and wavelength assignment [41,42]. Despite recent progression, wavelength conversion is still a critical and expensive node capability.

Wavelength Convertible Node Architecture

Equipping switching nodes with WCs has a positive influence on the blocking probability and the network utilization [14, 21, 33, 37]. The question is how the WCs should be integrated into the switching nodes. In [35–37], per-wavelength, per-fiber/port and per-node switch architectures are suggested. Figure 2.3(a) presents an architecture, where a dedicated WC is used for every wavelength at every port. This approach is simple, but also expensive because it is unlikely that every wavelength

requires conversion at the same time. To save some WCs, the architecture in figure 2.3(b) shares the WCs in a per-fiber/port pool. Only the wavelengths that actually need conversion are routed through the WCs. In figure 2.3(c), the WCs are shared in a per-node converter bank. Any of the incoming wavelength from any fiber/port may access the converter bank. Once the wavelengths are converted, the output switch directs them to their respective output fibers/ports.

The per-node converter bank allows wavelengths from all fibers/ports to use the provided WCs, hence giving the greatest flexibility and potentially lowest number of deployed converters. As for all shared architectures, there is a tradeoff between the number of WCs in the pool and the risk of blocking. Considering the current cost of converters, the potential converter saving of a per-node converter pool architecture far outweighs the added cost and complexity in the optical switches [37].

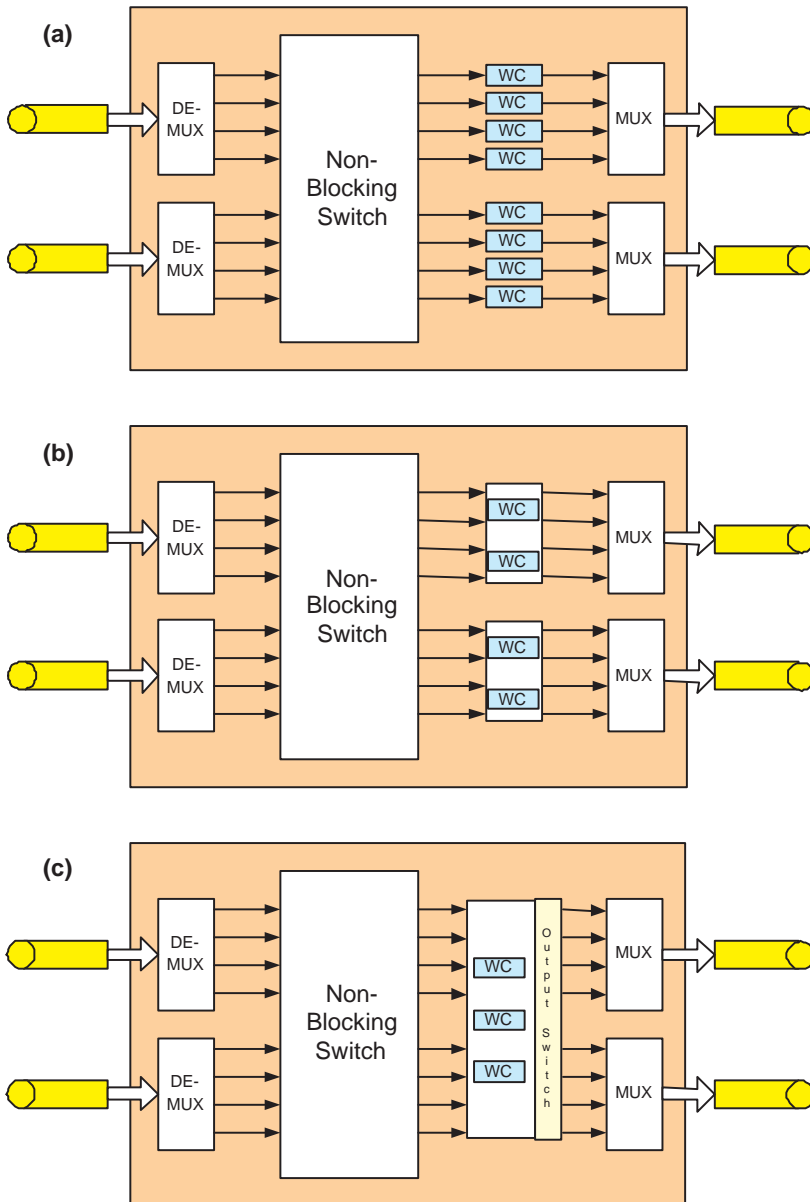


Figure 2.3: Examples of wavelength convertible node architectures. (a) per-wavelength, (b) per-fiber/port, (c) per-node. Adapted from [35–37]

2.2.2 Network Topologies

The structure in which the cables and nodes within an optical network are interconnected is referred to as a network topology. The most common topologies [33] are illustrated in figure 2.4 and described below:




Figure 2.4: Linear/point-to-point, ring and mesh topology.

- Linear/point-to-point: the simplest topology to provide connectivity between two nodes. Currently point-to-point is still the most used topology for the application of WDM networks³. This simple topology is the building block of more advanced topologies.
- Ring: an efficient way of interconnecting a group of nodes, because it reduces the fiber usage compared to point-to-point.
- Mesh: can be seen as a combination of the ring and the point-to-point topologies. The greatest advantage of the mesh topology is its flexibility, both when provisioning and protecting connections. Core networks are moving from ring to mesh topologies, because a mesh network is cheaper in terms of fibers, network cards, switches, etc., than if advanced services and protection should be provided in simpler topologies [45].

How to ensure survivability in different network topologies is explained in sections 2.5.6 - 2.5.8.

³This knowledge is obtained by an interview series related to network design, architecture and survivability covering Danish, European and US network operators, which was carried out in the innovation part of my Ph.D. study [43, 44].

2.3 Network Control Plane

2.3.1 Control Plane Functions

The control plane is responsible for the following functions [26], which are illustrated in figure 2.5:

- Neighbor discovery: allows a network element to automatically discover its neighbors.
- Routing: covers (a) automatic topology and resource discovery (i.e., propagation of connectivity and resource information); and (b) path computation (i.e., identifying a suitable path using available topology and resource information).
- Signaling: specifies the communication between control entities used to establish and maintain connections.
- Local resource management: takes care of book-keeping and advertising of locally available resources.

All of these control plane functionalities must be operational to automatically provision and recover a connection.

It is advantageous to separate the control plane from the data plane to prevent that failures in one plane will not affect the other. The control plane and the data plane can be separated logically or physically (i.e., using separate networks).

2.3.2 Control Plane Standardization

To ensure that multi-vendor networks can operate together, the control plane must be standardized. Different organizations contribute to control plane standardization:

- The ITU, which develops a framework for ASON.
- The Internet Engineering Task Force (IETF), which develops a framework for GMPLS.
- The Optical Internet Forum (OIF), which develops implementation agreements such as the UNI and the NNI.

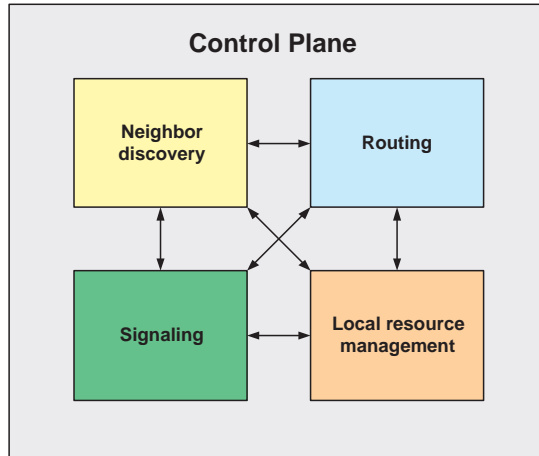


Figure 2.5: Control plane functions.

Several studies, such as [26, 45–47], focus on the difference between these control approaches, but it should be noted that they complement each other [48], since ASON can reference the protocol specifications of the GMPLS protocol suite [47]. The three standardization approaches are described in the following sections.

2.3.3 ASON

ASON [28] is developed by the ITU and is therefore inspired by concepts used in telecommunication transport networks, with well defined interfaces between clients and servers [26]. The goal of ASON is to improve the complex process of provisioning end-to-end transport services [46] by creating a complete definition of the dataplane, operation and management for automatically switched transport networks [47]. ASON specifies a separate Data Communications Network (DCN) [49], which is used as the control communication infrastructure between optical network elements. It also specifies an architecture and requirements for routing [50], a Distributed Connection Management (DCM) [51] using specific protocol mechanisms, and neighbor discovery [52]. ASON is not a protocol or a protocol suite, but a reference architecture that defines the control plane components and how they interact with each other. ASON is protocol

neutral but requires standardized protocols because it defines communication across multi-vendor networks, and any protocol that satisfies its functionality requirements can be included into ASON.

2.3.4 GMPLS

GMPLS [27] is developed by the IETF and uses an IP-based control plane. GMPLS is an extension of the Multi-Protocol Label Switching (MPLS) concept, which was developed to apply traffic engineering to Internet Protocol (IP) networks [45]. In GMPLS, the label is generalized to signify a timeslot, a wavelength or a fiber. GMPLS defines a set of protocols and a framework covering how these protocols should be applied together. The most used protocols within the protocol suite are:

- Open Shortest Path First (OSPF) with Traffic Engineering (TE) extensions [53]: used for routing (i.e., resource information dissemination, path computation).
- Resource ReserVation Protocol (RSVP) with TE extensions [54]: used for signaling (i.e., connection provisioning and maintenance, label assignment).
- Link Management Protocol (LMP) [55]: used for link management (i.e., neighbor discovery, fault localization).

The recovery mechanisms presented in this thesis are based on GMPLS-based recovery and hence more information on the protocols is given in the relevant chapters.

2.3.5 UNI and NNI

The OIF specifies the UNI [29] and NNI [30] based on needs of service providers and equipment vendors. They have the following characteristics [26]:

- UNI: defines the interface between the client and the network, and is used by the client to request a service from the network.
- NNI: defines the interface between different network domains. The NNI is further split into an External Network-to-Network Interface

(E-NNI) and an Internal Network-to-Network Interface (I-NNI) depending on whether the NNI is located between or within administrative domains, respectively.

The OIF also aims at bringing the concepts of ASON and GMPLS together and mediates between the respective standardization bodies [26, 56].

2.4 Network Management Plane

The network management plane performs management functions for a network and coordinates functions of the data plane, control and itself [25]. The management plane is furthermore used to manage individual network devices [57].

Examples of protocols processed in the management plane are Simple Network Management Protocol (SNMP) and Telnet [57].

Since this thesis focusses on the dynamic processes in the control plane to ensure survivability in the data plane, the reader is referred to [58] for further information on the management plane.

2.5 Network Survivability

2.5.1 Failure Types

In an ideal world, communication services that have once been setup would continue running until they are no longer needed. Unfortunately, in real life networks failures do occur, and measures must be taken to ensure the continuation of communication services even when the network experiences failures. The traffic in a network can be affected by many kinds of failures, such as:

- Cable cuts
- Power failures
- Software bugs
- Hardware faults
- Fire
- Natural disasters
- Human errors

Cable cuts due to construction work cause the most significant number of network failures [33, 59]. But no matter the cause of the failure, the goal is the same: Communication networks must be able to survive faults - they must be resilient. A multitude of resilience mechanisms have been developed, and a representative selection of them is described in this chapter.

2.5.2 Basic Requirements

Some general requirements must be fulfilled to make a network resilient. The first requirement is dual homing, which requires that every node in the network must be connected to the rest of the network by at least two spans, else a single cable cut would separate the node from the rest of the network. The dual homing concept is illustrated in figure 2.6. There, if the span connecting the left-most node fails, the network is partitioned, whereas another span can be used in the dual homing case.




Figure 2.6: Dual homing example.

The second requirement does not allow the working⁴ and backup paths to share the same physical route, i.e., they must be disjoint, else a single failure could take both the working path and the backup path out of operation. The two paths can either be span disjoint, meaning that they may not traverse any common spans; or they can be node disjoint where they may not visit the same nodes. An example of span and node disjointness is illustrated in figure 2.7. Disjointness can be ensured by categorizing network elements into Shared Risk Link Groups (SRLGs), which define fate sharing, which means that a single failure can result in the failure of all elements within the same SRLG. By taking SRLGs into account in the route calculation, a working and a backup path that cannot be affected by the same single failure, can be chosen.




Figure 2.7: Span and node disjointness.

The availability of a network element is used to describe the probability that it is operational at some point in time [45]. The availability A

⁴The working path is also called the primary path.

is calculated according to equation 2.1, where $MTTR$ denotes the mean time to repair and $MTBF$ the mean time between failures of a network element.

$$A = 1 - \frac{MTTR}{MTBF} \quad (2.1)$$

The availability of a network can be calculated according to equation 2.2, provided that independence of failures is assumed [45]. $A_{Network}$ is the network availability and $A_1 \dots A_n$ are the availabilities of each of the network elements.

$$A_{Network} = A_1 \cdot A_2 \cdot \dots \cdot A_n \quad (2.2)$$

Decreasing the failure occurrence (i.e., $MTBF$) has a direct effect on the availability of a network and it is therefore desirable to minimize it. This can be done by making some of the node equipment redundant, e.g., the power supply, or by digging the cable deeper into the ground. Despite these measures failures cannot be avoided completely and therefore it is necessary to anticipate them and provide a plan of action for traffic recovery. This is achieved by so-called network recovery or resilience schemes, which operate at a network scale level. The basic idea of network recovery is to divert traffic onto functioning backup paths in case of a failure. As soon as a failure is detected, the recovery mechanism automatically diverts the traffic that is affected by the failure from the working path to the backup path [60].

2.5.3 Criteria for Recovery Method Evaluation

With many recovery methods to choose from, the question of which one is better naturally emerges. To compare the methods, it is necessary to have some common criteria. These criteria include:

- Recovery percentage
- Operational complexity
- Recovery speed
- Node resource usage (e.g., wavelength conversion)

- Capacity usage

The most common approach is to evaluate recovery methods under a single span failure assumption. If a single node failure is assumed, it can be treated as the failure of all spans connected to that node.

2.5.4 Fault Management

The first step in successful connection recovery is to detect that there has been a failure. It takes some time for the failure adjacent nodes to notice a fault, identify exactly which component has failed, and notify the entity responsible for fault management (e.g., the GMPLS control plane). When a failure is suspected, the system moves to a hold-off state to make sure that recovery is only initiated at persistent fault conditions, and also to give techniques of other layers a chance to recover from the failure (see section 2.5.12 for details). If the hold-off period expires and the failure persists, its exact location must be determined to initiate proper recovery actions, circumventing the failed element(s). It is possible that a single failure causes several fault indications. Therefore, to decrease the signaling overhead, the fault indications are correlated before they are dispatched to the node that initiates the connection recovery. These fault notification messages can either be transmitted in the data plane (e.g., as in Synchronous Digital Hierarchy (SDH) using the Remote Defect Indication (RDI) signal), or through the control plane (e.g., using GMPLS RSVP-TE Notify messages) [45]. Once the fault notification message is received, the actual connection recovery is initiated.

2.5.5 Protection vs. Restoration

In the field of network survivability the two terms protection and restoration are widely used. Although both methods are used to ensure that the traffic is switched to a backup path in case of a failure, there is a subtle difference between them, which is related to when the backup resources are provisioned.

Protection refers to the situation where the backup path is provisioned when the working path is set up. Everything is calculated before the failure occurs: the route, which resources to use, and the switches are configured. This means that when a failure occurs, the backup path is

ready for the traffic, ensuring fast and guaranteed traffic recovery. The main drawback of protection is its static nature due to the fact that recovery paths are pre-planned. Therefore, protection can only recover from failures that were anticipated (e.g., a single span failure); if something unforeseen happens the traffic cannot be recovered.

*Restoration*⁵ describes the situation where a spare capacity pool is available in the network instead of pre-assigned backup capacity. In contrast to protection, a connection does not reserve spare capacity when the primary path is provisioned. Only when a failure occurs, a suitable backup path is identified and the necessary resources are reserved. Hence, a restoration scheme allows for more flexibility than a protection scheme when dealing with unexpected failures, because there is often more than one possible restoration path. But since the search for a restoration path only starts when a failure occurs, the recovery time is longer for restoration than for protection. There exist several sub-categories of restoration, depending on whether they contain an element of pre-planning. In so-called pre-planned restoration [34] a set of recovery paths are pre-calculated, but not provisioned until after the failure. This is beneficial for the recovery time if an anticipated failure occurs; else path re-computation is required as well.⁶

Restoration is applied in mesh networks, where the high density allows for several recovery path alternatives. Therefore, span, path and segment protection (which are described in the following sections) also exist in their respective restoration versions.

Protection and restoration can also be combined within the same network, e.g., protection for the most important traffic and restoration for other traffic. This can be used to provide differentiated resilience [61]. Furthermore, restoration can be applied in case protection cannot resolve a given failure situation, e.g., a dual failure affecting both the working and backup path of a protected connection [60].

An overview of the most well-known protection methods is given in the following sections. Further details on the recovery methods can be found in [26, 33, 34, 45, 60].

⁵Restoration is also called dynamic protection.

⁶In this thesis, the term restoration only covers the truly dynamic recovery process where both path computation and connection setup are carried out after a failure occurs.

2.5.6 Linear Protection

Linear protection refers to the situation where one span between two nodes is affected by a failure, the traffic is switched to another span. Several sub-categories of linear protection exist:

1+1 Protection

In 1+1 protection, the traffic is sent on both the working and the protection path, as shown in figure 2.8. This means that the protection path is always ready, and the method is therefore often referred to as hot-standby. The receiving node decides which of the two signals to select based on the signal quality or fault information. The advantage of the 1+1 scheme is that it is simple and has the fastest recovery times, but having to carry a signal copy renders the method capacity inefficient.

1:1 Protection

In 1:1 protection, the traffic is only sent on one of the paths at a time, i.e., on the working path before a failure, as illustrated in figure 2.9. In case of a failure, the traffic is switched from the working path to the backup path at both the sending and the receiving node. During failure free operation, the backup path may carry additional traffic, which is preempted in case the capacity is needed for connection recovery. This increases the capacity efficiency.

m:n Protection

To increase the capacity efficiency even further, a sharing scheme between the working paths and backup paths can be applied. $m:n$ indicates that n protection paths are shared by m working paths. $m:1$ protection is a special case of $m:n$ protection, where one backup path is shared by several working paths, which is illustrated in figure 2.10. The increased capacity efficiency comes with a higher risk, because only n out of m possible failures are repairable. Therefore, finding a suitable ratio between m and n is essential.




Figure 2.8: 1+1 protection.




Figure 2.9: 1:1 protection.




Figure 2.10: m:1 protection.

2.5.7 Ring Protection

The ring protection scheme has been designed for SDH networks but can be extended to WDM networks as well [33]. Many different implementations exist, amongst them Unidirectional Path Switched Ring (UPSR), 2-fiber Bidirectional Line Switched Ring (BLSR) and 4-fiber BLSR. An example of the 4-fiber BLSR is shown in figure 2.11, where a bidirectional connection is in operation between node A and node B. In figure 2.11(a) only the working fibers between node A and node B are affected by a failure. Since the protection fibers are intact between the two nodes they can be used as backup. This concept is called span switching. In figure 2.11(b), the failure affects both the working and the protection fibers between the two nodes. In this case span switching is no longer feasible and the traffic must be routed the other way around the ring to get from A to B and vice-versa. This concept is called ring switching. Ring protection schemes ensure fast and simple switching, but they are not capacity efficient [34, 60].

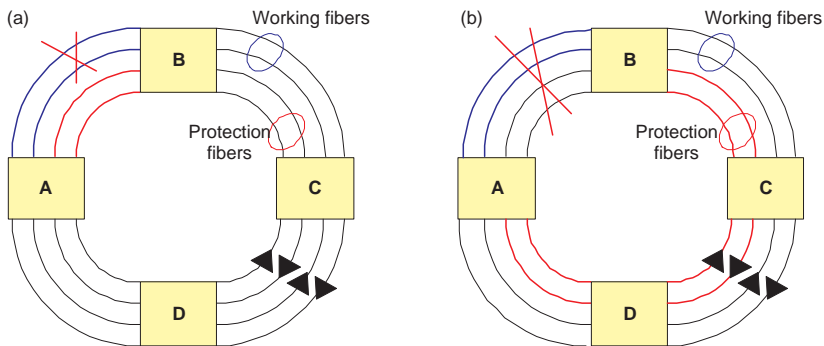


Figure 2.11: 4-fiber BLSR with (a) Span switching and (b) Ring switching, based on [33].

2.5.8 Mesh Protection

Recovery in mesh networks is the main focus of this thesis, therefore special emphasis is given on the description of mesh protection methods. Figure 2.12 illustrates three different recovery methods: span, path and local-to-egress, which are described in this section. The pre-failure connection paths are shown in figure 2.12(a).

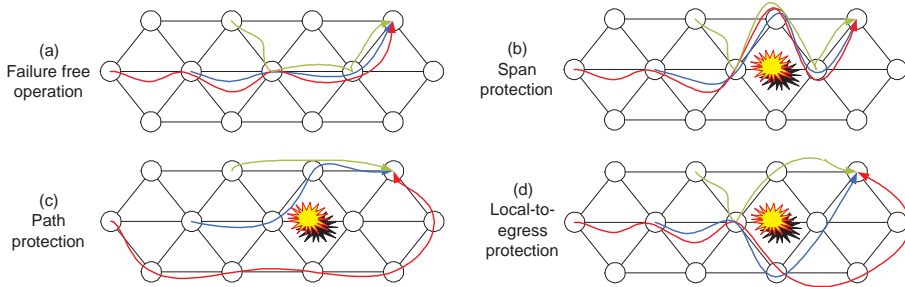


Figure 2.12: Recovery of a failure free connection (a), with (b) Span protection, (c) Path protection and (d) Local-to-egress protection.

Span Protection

In span protection, the connection recovery is carried out between the two failure adjacent nodes, as shown in figure 2.12(b). This means that the failure is dealt with locally, without involving the end nodes of a connection. The advantage of span protection is a short notification time, since the recovery is initiated by the same node that detects the failure. The drawback of span protection is a longer backup path and a large capacity requirement on the failure adjacent spans (i.e., the spans connected to the failure adjacent nodes).

Path / End-to-End Protection

In path protection, each connection is recovered between its end-nodes, as depicted in figure 2.12(c). In case of a failure, the end nodes of each individual connection must be notified of the fault, causing a delay before the traffic is switched to the backup path. The advantage of path protection is its freedom in route choice since there is no requirement of resuming the pre-failure path.

Local-to-egress Protection

Local-to-egress protection [62] (also called local backup dynamic protection [63]) is a special case of so-called segment protection. In segment protection the recovery can be carried out between any two intermediate nodes. Local-to-egress protection is illustrated in figure 2.12(d), where the recovery is carried out between the upstream failure adjacent node

and the destination node. Local-to-egress protection combines the advantages of span and path protection, such as short notification time and capacity efficiency due to freedom in the recovery route selection.

2.5.9 Combining Ring and Mesh Protection

In [34,64], a new protection method with the aim of combining the advantages of ring and mesh protection is described. The method is called *p*-cycles⁷, and promises to achieve ring-like speed with mesh-like efficiency. The main idea of a *p*-cycle is to pre-configure a ring (i.e., a cycle) in a mesh network, where protection traffic is routed after a failure. The advantage of a *p*-cycle compared to a normal ring is that a *p*-cycle not only protects traffic that is on the ring, but also on spans that partition the *p*-cycle. These spans are called straddling spans. Within the field of *p*-cycles, several studies for evaluating *p*-cycle efficiency and finding suitable *p*-cycles exist. A good introduction to the *p*-cycle concept and its different research areas can be found in [34].

2.5.10 Reversion

Once the failure is repaired, the pre-failure routes are in principle operational again. Therefore it must be decided whether the traffic should continue to run on the backup path, or whether it should be reverted to the pre-failure paths. Often, the backup paths are less ideal than the pre-failure paths (e.g., longer, more expensive, less reliable, etc.). Furthermore, if shared protection schemes are used, the current traffic configuration may not be fully protected. Hence, reversion of the traffic can be executed to put the network back into its pre-failure state. Reversion is closely related to protection and restoration operation, but in contrast to connection recovery, which is a reaction to a failure, the reversion operation can be planned and should cause minimal service disruption [60].

2.5.11 Stub Release

When a single span failure occurs, only one span is taken out of operation, but several connections traversing that span are affected. The fate

⁷The *p* stands for protection.




Figure 2.13: Concept of stub release. Green lines represent original connection and dashed lines represent released capacity.

of the surviving parts of the affected connections, i.e., the connection stubs, must be decided before initiating recovery. From a control point of view, the simplest option is to leave the connection stubs in place as unused working capacity. This also makes reversion simple, because the capacity is already in place, but it also gives a poor resource usage, because potentially available capacity is excluded from the spare capacity pool. From a capacity point of view, it makes sense to release the surviving upstream and downstream portions of a path, and use the freed capacity as extra spare capacity. This concept is called stub release and is illustrated in figure 2.13. Note that each failure allows for a release of different resources, and that stub release can be used with any restoration method except span restoration, where the stubs are reused per definition.

Even though the increased spare capacity pool obtained through stub release may seem very tempting at a first glance, there are several operational issues that must be solved. The first challenge is related to the fact that the exact location of the failure, and thus implicitly the surviving parts of the connection, must be known in order to notify the nodes responsible for initiating stub release. The implementation is dependent on how much intelligence is distributed at the individual nodes. If we assume that the failure adjacent nodes see the failure quickly, for example by monitoring each span using LMP, they could initiate stub release by issuing an Alarm Indication Signal (AIS), which would indicate that the capacity is available to the spare capacity pool [34]. The failure detection, and hence appropriate stub release actions, becomes more complicated

when considering a group of transparent nodes, which are only responsible for pass-through switching and therefore without advanced control functions. This causes difficulties in pinpointing the exact location of the failure, but proposals such as Optical-Electro-Optical (OEO)-converting one wavelength at each node to be used as a supervisory channel or using out-of-band pilot tones that provide signal identification and power level information for fault management [65,66], exist. In the studies presented here, exact knowledge of the failure location is assumed.

If we assume that a connection traversing a network experiences a failure, and that it should be restored by using the path restoration technique (i.e., recovery between end nodes), several constraints can be considered for the usage of the spare capacity obtained through stub release. From a capacity point of view, it would be sensible to allow for the use of spare capacity (including the spare capacity gained from stub release) on any span in the network. This unconditional reuse could however result in a situation, where the restoration path of one connection occupies the stub release capacity of another connection, resulting in a cross-entanglement [34]. Such a situation is illustrated in figure 2.14, where the connection stubs are released and are now considered spare capacity. The released capacity of the green connection is used for the recovery (shown in dashed lines) of the blue connection; and vice versa where the blue connection is recovered using the released capacity from the green connection's working path. The resulting cross-entanglement means that none of the connections can be reverted to their pre-failure path without interrupting the other connection's restoration path.

Since cross-entanglement severely complicates the reversion process, a variation termed stub reuse is described in [34], which allows a restoration path to reuse parts of its own pre-failure path as a failure specific response, but the capacity is not released into the spare capacity pool as in stub release. At reversion, the reused path portions are left in place while only parts that departed from the pre-failure path must be switched back.

2.5.12 Single vs. Multi-Layer Recovery

A communication network typically consists of multiple layers, e.g., IP over WDM [67]. In a multi-layer scenario, a failure in the lower layer may




Figure 2.14: Cross-entanglement example.

cause several secondary failures in the higher layer. The simplest implementation is that each layer carries out its own independent recovery actions, but the drawback of this approach is that the recovery actions may obstruct each other. Therefore, the recovery between the different layers should be coordinated, e.g., through timers or token-passing. The subject of multi-layer recovery is covered in detail in [60].

2.5.13 Multiple Failures

The majority of studies within the field of network survivability are carried out under the assumption of a single span failure. Even though they are the most common fault [33], other types of faults, such as node failures affecting multiple spans, dual span failures, multiple failures (more than two spans, combination of node and span failures, etc.) may occur as well. Several studies [34, 68], have shown that the more failure scenarios a network is protected against, the higher the total capacity usage. The probability of multiple failures also increases with the size of the network. This is both in terms of the amount of equipment, each with a given failure probability, and in terms of the geographical extent. Another important assumption is the independence of failures. While two construction-related cable cuts in opposite parts of the network are usually independent, a power-failure may cause malfunctioning of several network components, or damage to a cable duct may cause the failure of all connections associated to that SRLG. Generally speaking,

the more fault possibilities taken into account when designing survivable networks, the more expensive the design will turn out. To reduce the cost, it has been proposed to divide the network traffic into several survivability classes, which differ on how many failures a connection can survive. In [34], the traffic is categorized into several classes depending which protection services they receive:

- Platinum: protected against dual span failures.
- Gold: protected against any single span failure.
- Silver: best effort restoration, using remaining spare capacity after gold services have recovered.
- Bronze: unprotected traffic, but may not be preempted.
- Best effort: may be preempted by higher classes.

This categorization opens the possibility for differentiated cost models, but there is still some doubt within the industry whether customers want to pay for additional protection, and if it really can be guaranteed [43].

2.5.14 Control Plane Failures

While the previous sections have dealt with data plane failures, it should be noted that the control plane can fail as well. Since the control plane may be separated from the data plane, it may also fail independently. The separation of the two planes ensures that already provisioned connections stay operational even though the control plane fails, but they cannot be managed anymore (e.g., torn down or recovered) - they become control plane partitioned. GMPLS contemplates resolving this issue through control plane recovery. Synchronization between control and data plane state information can be reestablished either from local databases or through signaling with modified RSVP Notify messages, which are described in [45].

2.6 Chapter Summary

This chapter has given an introduction into the field of optical networks. The focus has been on data plane technologies, and in particular how

optical networks can be controlled, both for connection provisioning and recovery. Several standardization efforts have been discussed, such as ASON, GMPLS and UNI/NNI. Furthermore, network protection and restoration, and several methods for linear, ring, and mesh protection have been reviewed. The treatment of optical networks given in this chapter is not exhaustive, but it provides an overview of the concepts that are applied in the remaining chapters of this thesis.

Chapter 3

Dynamic Recovery with Limited Wavelength Conversion

3.1 Introduction

In Wavelength-Routed Optical Networks (WRONs), lightpaths, i.e., all-optical end-to-end switched connections between node pairs, are typically established at wavelength granularity [33]. WRONs can be affected by failures. Since these networks carry large amounts of data, survivability is of paramount importance to cause as little disturbance to the traffic as possible. The adaptation to a dynamic network environment makes network restoration a very suitable candidate to provide resilience towards failures. To overcome a failure in a network, different restoration methods are distinguished, mainly characterized by the end nodes of their recovery routes. In literature, three main restoration methods are described, namely span, segment (local-to-egress) and end-to-end/path restoration [34, 60, 69–71], which have been described in section 2.5.8.

Dynamic WRONs leverage a distributed control plane to automatically perform the required network functions, such as lightpath provisioning, fault localization and recovery. A promising control plane standard is the Generalized Multi-Protocol Label Switching (GMPLS) [27]) protocol suite comprising extensions to routing (e.g., Open Shortest Path

First (OSPF)-Traffic Engineering (TE)) and signaling protocols (e.g., Resource ReserVation Protocol (RSVP)-TE), and a Link Management Protocol (LMP). For scalability reasons, the routing protocol typically advertises only summarized link information, which is used during lightpath provisioning and restoration to find a feasible route for the incoming request. Then a signaling session is triggered along the found path in order to perform the wavelength assignment.

Wavelength assignment efficiency, and by extension provisioning and restoration efficiency, are significantly influenced by the wavelength conversion capability [35, 36]. Wavelength Converters (WCs) overcome the wavelength continuity constraint, allowing a lightpath to be setup even when no continuous wavelength is available on each hop of the route. The conversion between wavelengths can be achieved in different ways [37, 38], but all these implementations are very costly. Due to the high price of WCs, several studies focus on minimizing their deployment in a network, either through advanced routing and wavelength assignment heuristics or sharing schemes [14, 21, 36, 42, 72–75]. As a result of the high cost and the WC minimization achieved by design studies, typically only a limited number of WCs are deployed throughout networks.

In this study, three emerging issues in optical networks are combined. We combine (i) dynamic connection provisioning and restoration with (ii) the scenario where there are few available WCs in each node in (iii) a GMPLS controlled all-optical network, and evaluate the effect that limited wavelength conversion has on the restoration performance of span, local-to-egress and end-to-end restoration.

The remainder of this chapter is organized as follows: in section 3.2 the previous work in the field is reviewed and the novelties of this study are highlighted. Section 3.3 presents distributed wavelength assignment schemes leveraging GMPLS signaling, while section 3.4 describes the network model used in this study. In section 3.5 the simulation scenario is detailed, and results are shown and commented in section 3.6. Section 3.7 summarizes and concludes on the presented results.

3.2 Previous Work

Network restoration techniques have been widely studied in the literature, highlighting their performance/complexity trade-offs in terms of

resource efficiency, recovery time, and success rate [69, 70, 76–78]. In particular in [69], path (i.e., end-to-end) restoration is compared to span and segment restoration in a network with full wavelength conversion capability. Furthermore, end-to-end restoration with additional recovery retrials is shown to achieve high network availability at the expense of a restoration time slightly higher than span and local-to-egress restoration, which in turn are characterized by a lower success rate.

Restoration performance can be significantly improved by exploiting WCs, as shown in [79–82]. In [83, 84], an analytical model of a preplanned segment restoration method is developed, showing that the benefits of full wavelength conversion dramatically improve for increasing path length, but become less significant for increasing wavelengths per fiber. An online wavelength assignment algorithm is proposed in [85], where the WC usage and the connection blocking are evaluated, exploiting a multicast conversion model where the output of the WC is split. However only few protection and restoration papers considered limitations to wavelength-conversion capability. In [86], a protection method is presented where limited-range wavelength-converters are considered.

As mentioned in section 2.3, different control and management approaches can be used to control provisioning and restoration actions in an optical network. The desire to provide a unified control framework leads to the Generalized Multi-Protocol Label Switching (GMPLS) protocol suite [27, 87, 88], which has undergone standardization efforts during the past few years, and is emerging as a control plane for next generation networks. Despite this development, only few studies [89, 90] consider the effect of GMPLS signaling when dealing with network restoration. However, these studies do not take the effect of limited wavelength conversion availability into account.

This study leverages the studies in [14, 21], where wavelength assignment schemes were investigated in the provisioning phase. In particular, an enhanced signaling object called Suggested Vector (SV) was introduced, which enables label ranking aimed at WC usage reduction.

The focus is on the restoration performance of all-optical GMPLS controlled networks with limited WC placement. Given a fixed number of WCs deployed in each node, a comparative study on the performance of span, local-to-egress and end-to-end restoration methods is performed, with particular focus on how these restoration methods are affected by

the limitation of WCs and whether using standard or enhanced GMPLS signaling influences the results.

3.3 Network Model

In order to evaluate the influence of WCs on the performance of span, local-to-egress and end-to-end restoration, a model of the network which serves as a basis for the conducted simulation study is created. The model incorporates realistic assumptions about the node and network architectures.

3.3.1 Node Architecture

The employed logical node architecture is illustrated in figure 3.1. Each node is connected to F_i input fibers and F_o output fibers, and each fiber span carries W wavelengths in each direction. Due to distributed control, each node handles both control and data plane functions.

The control plane is responsible for route computation and resource reservation, and accomplishes these functions by exchanging GMPLS control messages. The data plane is responsible for switching data belonging to different connections, based on control plane instructions.

The switch architecture is assumed to be non-blocking, and to have enough add/drop interfaces to establish all locally terminated connections. With reference to pass-through connections, a WC is utilized if the same wavelength is not available at both the incoming and outgoing port. The WCs are assumed to be tunable on the whole wavelength range and shared by all ports as a per-node converter bank [35].

3.3.2 Network Architecture

In order for the network to be independent of a given traffic matrix (i.e., demands between source/destination pairs), all the spans in the network comprise the same number of wavelengths, and all the nodes have the same number of WCs in their converter bank. All the nodes are interconnected through bidirectional fibers spans. Each of these fiber spans comprises unidirectional fibers in each direction. This means that given S bidirectional spans, the number of wavelengths is $2 \cdot S$. Since the investigated network is wavelength routed, the requested connections are

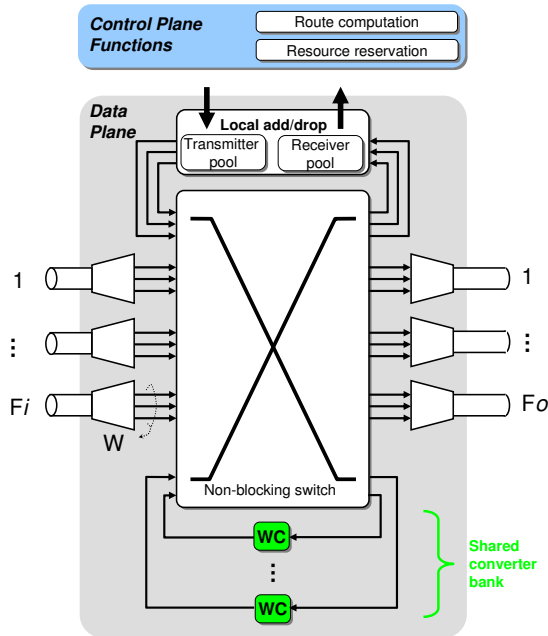


Figure 3.1: Node architecture with shared converter bank.

setup as unidirectional connections at the granularity of one unidirectional wavelength channel¹. The control messages are transmitted along the same route as the actual data. Each node can be the source or the destination of a connection request.

3.3.3 Route Computation and Signaling

In optical networks, a route and a wavelength must be found for each connection demand, which is referred to as the Routing and Wavelength Assignment (RWA) problem. Joint RWA in dynamic wavelength routed networks requires an amount of information which easily exceeds the capability of the routing protocol to maintain an up-to-date network status in each node. Therefore, the route and wavelength of a new connection

¹Only unidirectional connections are considered here, since issues such as outdated wavelength status and forward reservation, which are out of the scope of this study, impact the performance of bidirectional lightpath provisioning. A description of the problems related to bidirectional lightpath provisioning can be found in Appendix B.4.

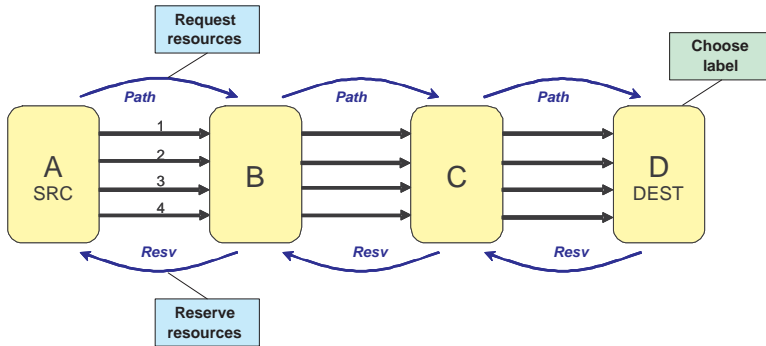


Figure 3.2: Exchange of GMPLS signaling messages.

are often found in decoupled steps [91]. To emulate the behavior of OSPF-TE in the model, the network topology is discovered at the beginning of a simulation run and the route for a connection is computed by running the Dijkstra shortest path algorithm. The simplified modelling of OSPF-TE is justified since routing is not the scope of this study.

The RSVP-TE protocol, which is used to reserve the necessary resources on the identified route, is modelled in detail. RSVP-TE uses a two-way reservation mechanism, where a resource is requested by propagating so-called *Path* messages between the source and destination node. Upon reception of a *Path* message, the destination node chooses a wavelength (i.e., a label) and propagates a *Resv* message upstream towards the source. The resources are only reserved upon reception of such a *Resv* message. The concept of exchanging *Path* and *Resv* messages is illustrated in figure 3.2. The actual wavelength can be chosen based on different wavelength assignment schemes, which are further discussed in section 3.4.

When a span fails, the LMP is modelled in such a way that the upstream failure adjacent node is notified, which then checks which connections are affected by the failure. It initiates recovery actions, either through notification of the end nodes or by itself, depending on the chosen restoration method.

3.4 Wavelength Assignment Schemes

WCs play an important role in simplifying the wavelength assignment in all-optical networks, both during the provisioning and the restoration phase. Their presence circumvents the wavelength continuity constraint and consequently increases the probability of successful connection provisioning and recovery. Moreover, WC availability avoids reservation collisions when two or more concurrent signaling sessions contend for the same wavelength channels. This benefit is particularly important in the restoration phase, when many recovery attempts are almost simultaneously triggered and must share a limited amount of available network resources. Additionally, if span or local-to-egress restoration is exploited, WCs are very useful to match the recovery path wavelength with the stub wavelength.

As WCs are a scarce resource, their usage must be rigorously controlled both during the provisioning and the restoration phase. However the standard wavelength assignment performed with GMPLS signaling is not suited to achieve this target. To perform a WC-saving wavelength assignment, an additional signaling object, called Suggested Vector, has been introduced in [14,21]. The following sections describe the standard GMPLS wavelength assignment, exploiting the Label Set object alone, and the enhanced wavelength assignment, exploiting the Suggested Vector object together with the Label Set².

3.4.1 Standard Wavelength Assignment - The Label Set

The Label Set (LS) is a standard protocol object defined in [87] with the purpose of simplifying the wavelength assignment in dynamic WRONs. The resulting scheme enforces the best possible wavelength assignment with current protocols and is considered as a benchmark in this study.

The LS object is propagated within the *Path* message and consists of an array of labels, containing only the wavelengths acceptable to the upstream node. The LS is typically used by an upstream node to control the selection of labels by downstream nodes [87]. If no WCs are available, the LS is restricted to contain those labels that ensure a wavelength-

²Further details on these wavelengths assignment schemes, the algorithm details and their behavior during connection provisioning can be found in Appendix B.2.




Figure 3.3: Standard wavelength assignment using the Label Set.

continuous path. If WCs are available, the LS contains all the labels that are available on a given span. Once the *Path* message arrives at the destination node, a label (i.e., a wavelength) within the LS is reserved according to a tie-breaking policy (e.g., first-fit or random)³. The chosen label is propagated as far as possible along the reverse path. If a node cannot reserve the current wavelength and a WC is available, a new wavelength is chosen within the node's received LS; otherwise the connection is blocked.

In the example illustrated in figure 3.3, the LS is propagated within the *Path* message from source to destination. At the destination, first-fit tie-breaking policy is enforced, meaning that the lowest-indexed label in the received LS (i.e., λ_1) is chosen and propagated upstream within the *Resv* message. However, wavelength conversion becomes necessary at the second node (highlighted in figure 3.3) because λ_1 is not available on the first hop. In contrast, no WC would be needed if wavelength λ_3 had been chosen at destination. From this example it is clear that, using only the Label Set, the destination node lacks information to choose the wavelength that minimizes WC usage.

³A description of the tie-breaking policies is given in Appendix B.2.3.




Figure 3.4: Enhanced wavelength assignment using the Suggested Vector.

3.4.2 Enhanced Wavelength Assignment - The Suggested Vector

A WC-saving wavelength assignment can be ensured by using a novel signaling protocol extension, called Suggested Vector (SV) [14,21]. The SV is used in the *Path* message together with the LS. The SV is an array of the same size as the LS, containing information allowing to rank the labels within the LS.

In this study, the labels are ranked according to the number of WCs each label requires to setup a connection, but other rankings are possible as well. The source node fills the SV with zeros, because all available labels can be reached without WCs. Then the SV propagates along the route. If a label was available on the previous hop, its SV values are propagated without modification. If a label was not available on the previous hop, its SV value is calculated by adding one to the minimum SV value of the previous hop (since the choice of this label implies one conversion at the current node). When the destination node is reached, the label with the lowest SV value (i.e., the one requiring fewest WCs) is chosen and propagated towards the source. If two labels have equal SV values, a tie-breaking policy needs to be applied. If a node cannot further

use the current wavelength and a WC is available, a new wavelength with minimum value of the previous hop SV is chosen; otherwise the connection is blocked.

The SV is updated at each node during *Path* message propagation, as illustrated in the example of the SV scheme operation in figure 3.4. At the destination the SV contains, for each label in the LS, the minimum number of WCs necessary to set up the request using that label: λ_3 is chosen because it requires zero WCs.

3.5 Simulation Study




Figure 3.5: Pan-European network topology.

In this study, the recovery performance of span, local-to-egress and end-to-end restoration is evaluated when the number of available WCs per node is varied, using a standard (LS) and an enhanced (SV) label assignment scheme. The simulations are carried out in OPNET Modeler [92]⁴, and evaluated in the Pan-European triangular topology network [93,94], consisting of 28 nodes and 61 spans, yielding a nodal degree of 4.36, as illustrated in figure 3.5.

⁴Further details of the OPNET model can be found in Appendix A.

3.5.1 Case studies

In order to evaluate the performance of the three restoration methods with standard and enhanced GMPLS signaling, two fundamentally different cases are considered:

1. Comparison of the standard and enhanced scheme exclusively in the restoration phase.
2. Comparison of the two schemes in both the provisioning and restoration phase (real network scenario).

Case 1 - Isolated restoration phase: The connections in the network are provisioned using unlimited conversion capability, and the wavelength assignment scheme in the provisioning phase is fixed to the standard scheme to ensure the same connection setup prior to the restoration phase. Once the desired load is reached, the number of WCs available for the restoration phase becomes limited and the recovery of the span failures is attempted using either wavelength assignment scheme.

Case 2 - Correlated provisioning and restoration phase: The restoration performance of a network which uses the same wavelength assignment scheme in both the provisioning and the restoration phase is investigated. This scenario is also the most realistic one from a real-life network operation point of view. All nodes are given a specific number of WCs at the beginning of the provisioning phase, and no separate WCs are assigned for the restoration phase. This means that the converter-saving property of the enhanced scheme can potentially help to save WCs already in the provisioning phase, which will keep the WCs available for the restoration phase.

3.5.2 Simulation Parameters

We evaluate the three investigated restoration methods (span, local-to-egress and end-to-end), using either standard or enhanced wavelength assignment. The number of WCs available per node is varied and the average network load ANL is kept constant at 0.5. The average network load is defined as the number of occupied unidirectional wavelengths divided by the total number of wavelengths in the network. That is

$$ANL = \frac{N_{Occupied}}{2 \cdot S \cdot W} \quad (3.1)$$

where $N_{Occupied}$ is the number of occupied unidirectional wavelengths in the network, S is the number of spans in the network and W the number of wavelengths on each span. Table 3.1 contains an overview of the parameter values used in both simulation cases.

Parameter	Value
Network topology	pan-european network
Wavelengths per span	10
Wavelength converters per node	0 - 20
Connection interarrival time	10 seconds
Connection end nodes distribution	uniform
Wavelength assignment scheme	standard (LS) or enhanced (SV)
Setup reattempts provisioning	none
Setup reattempts restoration	k-shortest paths with crankback until no route can be found
Connection holding time	infinite
Tie breaking policy	first-fit
Average network load ANL	0.5
Span failures	all spans consecutively failed
Restoration method	span, local-to-egress, end-to-end

Table 3.1: Simulation parameters

3.5.3 Simulation Execution

In the provisioning phase of the simulation, the network is populated with connections of indefinite holding time, meaning that once provisioned, the connections are not torn down or changed. Each connection is unidirectional, occupying one wavelength/label. Labels are assigned using either the standard LS or enhanced SV wavelength assignment scheme. Ties are broken using the first-fit policy [91]. The source and destination pairs are uniformly distributed over the entire network. Connections are routed along the shortest path (hop count) [34, 95]. If a connection request experiences blocking in the setup phase, it is dropped and a new source-destination pair is chosen. The provisioning phase continues until the desired average network load is reached.

At this point, the restoration phase begins and one span failure is si-

mulated at a time. No connection requests are assumed to arrive during failure recovery, which allows a clear differentiation between the provisioning and restoration phases of the simulation. The stub resources are kept occupied (i.e., no stub release is executed) and recovery with span, local-to-egress or end-to-end restoration is attempted using either the standard or enhanced wavelength assignment scheme.

A restoration path can be blocked due to the lack of free WCs or wavelength channels on the route. To increase the restoration percentage [69] the restoration is retried, temporarily excluding the blocking-causing span from the route computation. Information about the blocking span is sent back to the node performing the route calculation using the crankback mechanism described in [96], where the RSVP-TE *PathErr* message carries the information. The source node caches all blocking locations for a particular connection until it has been successfully restored or has been deemed unrecoverable. When all affected connections are either restored or found to be unrecoverable, the network is reverted to its pre-failure state before the next span failure is simulated. This procedure is repeated for all spans in the network. The following performance metrics are adopted:

- Recovery percentage:

$$RP = \frac{Conn_{Recovered}}{Conn_{Failed}} \cdot \frac{100}{1} \quad (3.2)$$

- WC usage per recovered connection:

$$WCRC = \frac{WC_{Restoration\ phase}}{Conn_{Recovered}} \quad (3.3)$$

- Wavelength channel usage in the restoration phase:

$$WLR = \sum WL_{Restoration\ phase} \quad (3.4)$$

- Hop count of recovery segment:

$$HCR = \frac{WLR}{Conn_{Recovered}} \quad (3.5)$$




Figure 3.6: Recovery percentages for Case study 1.

3.6 Results

All results presented are averaged over a number of simulation runs according to table 3.1 with 40 random seeds for each data point, and confidence intervals at 95% confidence level are given⁵.

3.6.1 Case 1 - Isolated Restoration Phase

In this section, results from Case study 1 are presented. As described in section 3.5.1, an identical set of connections is created using the standard signaling scheme by providing an unlimited number of WCs in the provisioning phase. In this way, the restoration methods can be compared across different signaling schemes given the same starting point. In the restoration phase, all restoration methods are evaluated using both the standard and enhanced wavelength assignment schemes. Note that the the number of WCs is limited in the restoration phase.

⁵Details on confidence interval calculation can be found in Appendix A.5.2.

Figure 3.6 shows that end-to-end restoration achieves the best recovery percentage followed by local-to-egress and span restoration in that order. This is explained by the fact that end-to-end restoration has a higher degree of freedom when computing restoration paths on a network-wide basis, whereas the recovery paths of span and local-to-egress restoration emerge from the upstream failure adjacent node, and hence must be merged to their stubs. Above 5 WCs per node available in Restoration Phase (WCiRP), this order is interestingly altered with local-to-egress slightly outperforming end-to-end restoration. At this point, the network moves away from being WC-limited to being span-resource-limited (i.e., available wavelengths on the spans are the scarcest resource). Hence, the potentially shorter restoration segment of local-to-egress restoration is more useful than the freedom of route selection offered by end-to-end restoration. The span restoration method achieves the highest profit from a higher number of WCiRP, since the recovery segment of the connection has to be merged to the stubs of the pre-failure path at the upstream and downstream failure adjacent nodes, thus likely requiring a WC. The enhanced signalling scheme gives a better recovery percentage in all cases except the degenerate case of 0 WCiRP, where both schemes achieve the same performance, as expected. The largest gain of using enhanced signalling is obtained at 1 and 2 WCiRP, resulting in an increase of approximately 5% in Recovery Percentage. The advantage of enhanced signalling diminishes with increasing WCiRP (since the network is no longer WC-limited), but is still visible at 5 WCiRP.

The WC usage per recovered connection (WCRC) is shown in figure 3.7. A general tendency to observe is that the usage of WCs increases when they are widely available throughout the network. More WCs allow more connections to be restored, as can be seen in figure 3.6. It is clearly shown how the enhanced signalling schemes use significantly fewer WCs than the standard scheme regardless of restoration method. The curves flatten from 5 to 10 WCiRP, indicating that the network is saturated with WCs and is no longer WC-limited. Using the enhanced scheme, the WC usage of the restoration methods can be ranked as end-to-end (fewest WCs), local-to-egress and span (most WCs). Using standard signalling schemes, the same ranking is seen above 10 WCiRP. For fewer WCiRP span restoration behaves differently, performing between that of

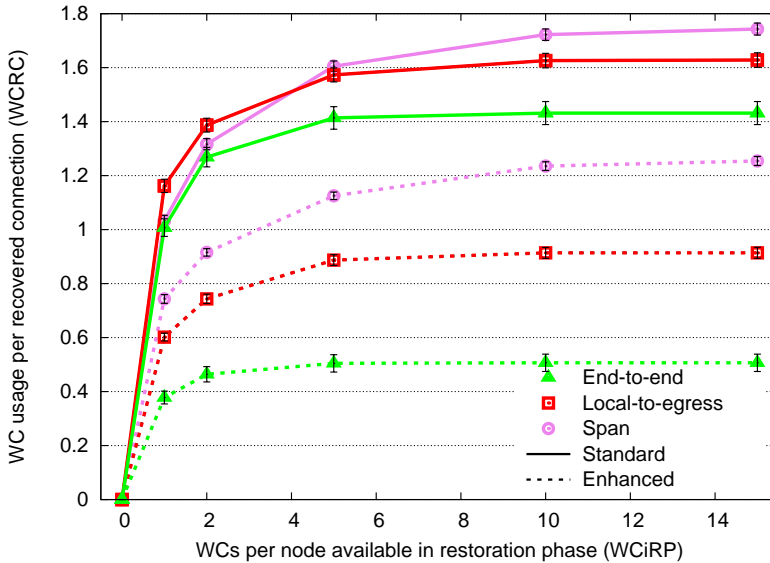


Figure 3.7: WC usage per recovered connection (WCRC) for Case study 1.

end-to-end and local-to-egress restoration. This is due to the fact that the recovery percentage for span restoration is low at this point, because even though converters may be available in the network, they are not available at the needed locations (often the failure adjacent nodes). This is also illustrated in figure 3.6 by the fact that span restoration achieves the highest benefit from more WCiRP, since the recovery percentage steeply increases when they become available.

Figure 3.8 illustrates the hop count of the recovery segment (HCR) for the different restoration methods. End-to-end restoration has the longest restoration segment, since it is restored between the connection's end nodes. local-to-egress restoration has medium length, as paths are restored between the upstream failure adjacent node and the destination node. Span restoration, where connections are restored between the failure adjacent nodes, has the shortest restoration segment. For all methods, the hop count increases at 1 WCiRP. At this point, we also observe a large increase in recovery percentage, which means that also longer paths are recovered. When more WCiRP become available, the hop count decreases because WCs allow the use of shorter routes other-




Figure 3.8: Hop count of recovery segment (HCR) for Case study 1.

wise unavailable. The availability of wavelength converters also has an influence on the hop count for the enhanced and the standard scheme. At few WCiRP, the enhanced scheme achieves a lower hop count, leveraging the saved converters, while the difference diminishes when WCs no longer are the limiting factor.

Figure 3.9 shows the wavelength usage in the restoration phase (WLR) as a function of the WCs per node. Even though the enhanced scheme achieves a higher recovery percentage than the standard scheme, the WLR increase of the enhanced scheme compared to the standard scheme is very limited, which highlights the fact that the difference in recovery percentage is achieved by more efficient WC usage and not longer restoration paths. Comparing the different restoration methods, end-to-end restoration obtains the highest WLR, because it has the longest restoration paths and the highest recovery percentage, while in span restoration WLR is the lowest because the recovery paths just circumvent the failed span and the recovery percentage is the lowest. As expected, local-to-egress restoration lies in between. All restoration methods show a saturation behavior for increasing WCiRP, meaning that WCs are no




Figure 3.9: Wavelength usage in the recovery phase for Case study 1.

more a critical resource, but the network becomes span capacity limited.

3.6.2 Case 2 - Correlated Provisioning and Restoration Phase

For the results presented in this section, the same wavelength assignment scheme is used both in the provisioning and restoration phase, as would be the most realistic choice for real-life network operation. The influence of the different schemes on the RP, WCRC, HCR and WLR is presented.

Figure 3.10 illustrates the recovery percentage obtained for the different restoration methods when the number of WCs per node given at the start of the provisioning phase is varied. A general observation is that the enhanced wavelength assignment scheme outperforms the standard wavelength assignment scheme for all restoration methods, except at 0 converters where both schemes perform the same as expected. Furthermore, we also observe a general tendency that the recovery percentage increases if more WCs are added. However, at 2 WCs using span restoration with the standard scheme, this trend is interestingly broken. This is due to the fact that if no WCs are available, only wavelength continuous connections can be allocated, which also ensures order in the wavelength assignment on the different spans. With very few WCs however, the first few requests may use WCs and hence create disorder in the wavelength assignment on the different spans, without there being enough WCs to use the wavelengths that are available between groups of used wavelengths later on, which results in a lower recovery percentage. Span restoration is particularly vulnerable to disorder in the wavelength assignment, since the restoration path must be merged to the remaining parts (i.e., stubs) of the pre-failure path at the two failure adjacent nodes. The reason why the difference in recovery percentage between the standard and enhanced schemes is bigger compared to Case 1 is that the enhanced wavelength assignment scheme can save WCs already in the provisioning phase, which in turn will be available in the restoration phase. Span restoration achieves the largest performance increase, going from an average restoration percentage of 63% to 88% at 5 WCs per node. For local-to-egress and end-to-end restoration, the increase in recovery percentage is smaller, but still significant. When 20 WCs are available per node, the difference diminishes not only between the different wavelength assignment schemes, but also between the different restoration methods.

The WC usage per recovered connection (WCRC) is shown in fig-




Figure 3.10: Recovery percentages for Case study 2.

ure 3.11. A general trend is that the usage of WCs increases when they are widely available throughout the network. The higher WC usage of the enhanced scheme at few WCs allows for more connections to be restored, as can be seen in figure 3.10. Furthermore, we observe a saturation behavior above 10 WCs, where neither the WCRC nor the RP increases. This is due to the fact that the network is becoming capacity limited and adding more WCs does not remedy this. For the standard scheme, which has a more wasteful wavelength assignment, more WCs are needed to achieve the saturation behavior. Using the enhanced scheme, the WC usage of the restoration methods can be ranked as end-to-end (fewest WCs), local-to-egress and span (most WCs). Using standard signalling schemes, the same ranking appears at high WC numbers. The reason for this ranking is that span restoration must merge the restoration path to the stubs at the two failure adjacent nodes, which is potentially WC consuming. local-to-egress restoration only has one merger node, while end-to-end restoration entirely avoids merging.

Figure 3.12 shows the hop count of the recovery segment (HCR) for the different restoration methods. Span restoration, where connections




Figure 3.11: WC usage per recovered connection (WCRC) for Case study 2.

are restored between the failure adjacent nodes, has the shortest restoration segment. local-to-egress restoration has a medium path length, since paths are restored between the upstream failure adjacent node and the destination node. End-to-end restoration has the longest restoration path, since the recovery path is found between the end nodes of a connection. The HCR has a similar performance for the two schemes, with slight variation at few WCs, which is due to the fact that a varying WC count allows for different connections to be restored. When WCs are abundant, both schemes achieve the same HCR.

Figure 3.13 shows the wavelength usage in the restoration phase (WLR) as a function of the WCs per node. End-to-end restoration used most wavelength resources for its restoration paths, followed by local-to-egress and span restoration, which is due to the choice of merging nodes. The enhanced scheme uses more wavelengths than the standard scheme, which is due to the fact that leveraging the available WCs in the restoration phase results in a higher recovery percentage, as shown in figure 3.10. Once WCs are no longer the limiting resource and the network becomes span capacity constrained, both signaling schemes show

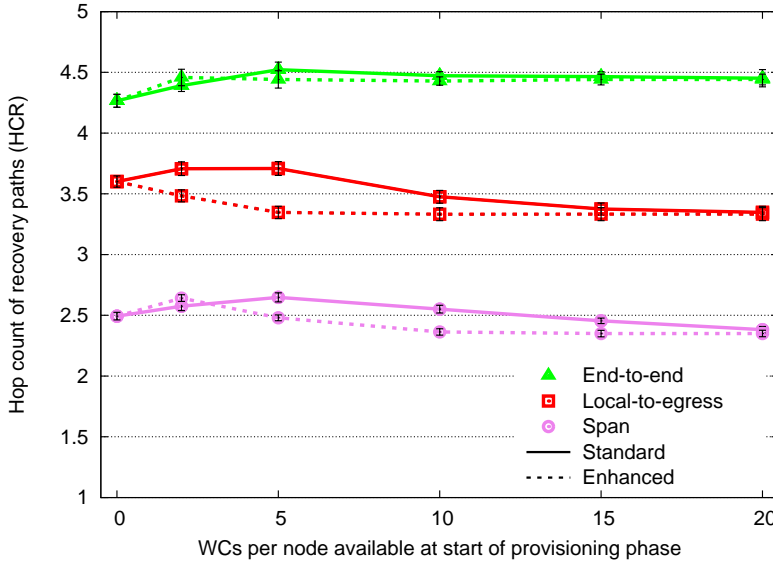


Figure 3.12: Hop count of recovery segment (HCR) for Case study 2.

equal wavelength usage in each restoration method.

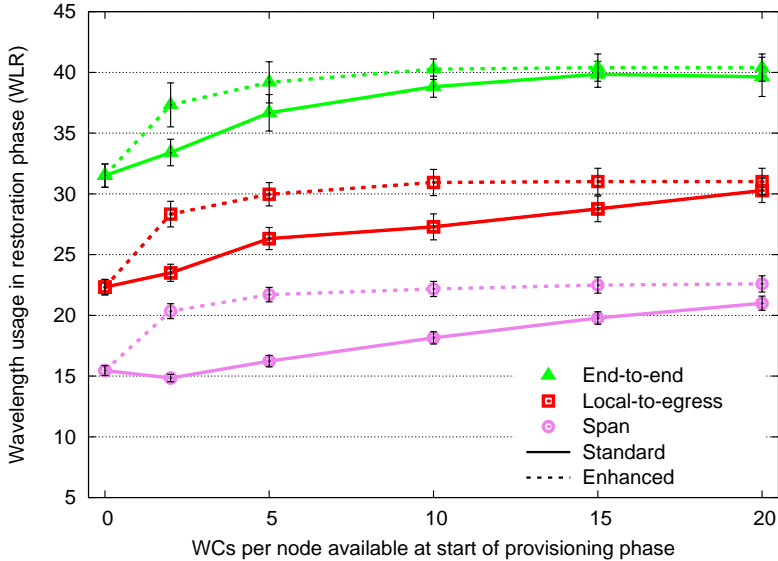


Figure 3.13: Wavelength usage in the recovery phase for Case study 2.

3.7 Summary and Conclusions on Restoration with Limited Conversion

In this chapter, the effect of limited wavelength converter availability on the span, local-to-egress and end-to-end restoration methods are evaluated. The performance of an all-optical network in a GMPLS controlled environment is simulated, thereby evaluating the effect of both standard and enhanced (aiming at minimizing converter usage) wavelength assignment schemes for the different restoration methods. In two case studies, the effect of limited wavelength conversion was investigated: 1) exclusively in the restoration phase, and 2) in a network operation scenario with correlated provisioning and restoration phase.

The study shows that limited wavelength conversion availability has a different effect on the three restoration methods. The recovery percentage of span restoration is most severely affected, due to its nature of having to merge the restoration segment to the pre-recovery path at the failure adjacent nodes, which often requires a conversion between wavelengths. This is consistent with the number of converters that are used

per recovered connection. End-to-end restoration is the least affected, because it restores affected connection between the connection's source and destination nodes, and hence profits from a high degree of freedom and the ability to avoid potential bottlenecks when setting up restoration paths. The performance of local-to-egress restoration lies between the previous two methods.

Another interesting observation is that if abundant WCs are available, all three restoration methods achieve very similar restoration performance. Since the choice of restoration method becomes less relevant for the actual recovery percentage, an operator can choose a restoration method based on other criteria, such as notification time or manageability, given that enough WCs are available.

The benefit of the standard or enhanced wavelength assignment scheme highly depends on the number of WCs that is provided. If WCs are limited, all restoration methods can benefit from the converter saving properties of the enhanced scheme. Especially span and local-to-egress restoration can significantly increase their performance with the enhanced scheme. If span or local-to-egress restoration is the preferred restoration method and WCs are very limited, it is advisable to use the enhanced scheme to minimize performance penalty compared to end-to-end restoration.

Chapter 4

Signaling-Based Recovery Enhancements

4.1 Introduction

In the field of network recovery, much research is devoted to finding suitable restoration paths. This is generally achieved by off-line optimization schemes which may take a large variety of constraints into account. In a dynamic network environment the routing and wavelength assignment is normally treated as two separate sub-problems due to scalability issues [91], in such a way that a route is identified first and subsequently resource reservation is attempted on that given route. This chapter employs this separated approach. A simple shortest path based routing scheme is used to find a path, and the subsequent wavelength assignment is tailored for recovery enhancement. Two methods for improving restoration performance through sophisticated wavelength assignment in a signaling session for resource reservation in dynamic connection recovery are presented. Section 4.2 present a scheme that reduces blocking of restoration requests. In section 4.3, a stub-awareness scheme aiming at wavelength conversion reduction is described. Section 4.4.3 extends the restoration concept to a multi-domain environment. Section 4.5 summarizes the chapter.

4.2 Blocking Reduction of Span Restoration Requests

4.2.1 Motivation

This section presents a modification of the Suggested Vector (SV) [14,21], which aims at reducing blocking of recovery requests in addition to minimization of Wavelength Converter (WC) usage. The modification allows spreading of the wavelength requests throughout the range of available wavelengths and thereby reduces the number of blocked restoration attempts. As a consequence, recovery retries are reduced, thus lessening control plane load, and the restoration time is reduced.

In span restoration the upstream failure adjacent node initiates the recovery for all affected connections. To setup a recovery path in a Generalized Multi-Protocol Label Switching (GMPLS) environment using Resource ReserVation Protocol (RSVP)-Traffic Engineering (TE) signaling, the recovery initiating node finds a route around the failure and then sends a *Path* message towards the downstream failure adjacent node, thereby collecting information on available wavelengths along the envisaged route. Once the end node is reached, it chooses a wavelength and sends a *Resv* message towards the initiating node. Several recovery requests may want to use the same wavelength, leading to resource contention. This contention can be resolved through wavelength conversion, if WCs are available at the appropriate nodes.

Successful connection setup and the concepts of forward and backward blocking due to WC unavailability are illustrated in figure 4.1. Forward blocking occurs if a suitable route has been identified, but the desired wavelength is no longer available during *Path* message propagation, and no converter is available to divert to another wavelength. A connection experiences backward blocking if a given wavelength was available during *Path* message propagation, but has been occupied by another connection when the *Resv* message is propagated. If a span becomes completely occupied in the timeframe between route calculation or *Path* or *Resv* message propagation, blocking occurs despite WC availability.

Blocking is likely to arise in span restoration, where the recovery for all affected connections are initiated between the same two nodes during

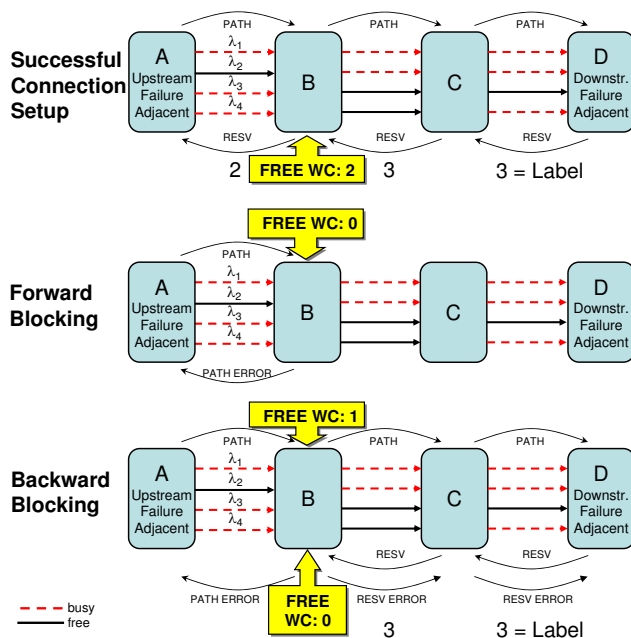


Figure 4.1: Successful connection setup compared to forward and backward blocking.

a short timeframe. In case of blocking, the recovery is re-attempted over a different route to increase the recovery percentage [69], but each retry causes restoration delay and the exchange of additional control messages.

4.2.2 Blocking Reduction Concept

The exhaustion of WCs is the cause of many blocking incidents, since the flexibility of diverting to another wavelength during the reservation process is lost. To minimize the blocking probability and thereby the number of signaling messages and recovery retries, the wavelength assignment of the recovery routes should be administered in a WC-saving manner. The signaling protocol extension called SV [14,21] is modified to achieve this, resulting in the Blocking Reduction (BR) scheme described here. The SV is used in combination with the Label Set, which is a standard protocol extension allowing an upstream node to control the wavelength/label selection of a downstream node [87].

By means of the SV, a label preference level is recorded at each node, enabling the downstream node to choose the most preferable label. As the SV scheme was invented with connection provisioning in mind, all labels are assigned identical initial SV values, so the downstream failure adjacent node has a large selection of equally preferable wavelengths to choose from. In a span restoration scenario this behavior has two drawbacks: first, the probability of collision with simultaneous recovery attempts increases; second, on a previously wavelength continuous lightpath a WC is required at failure adjacent nodes, if a wavelength different from the one used on the failed span (pre-failure wavelength) is chosen. The BR scheme presented here aims at solving both issues. It is based on the idea that all connections were assigned a particular wavelength on the failed span, and by letting the restoration path prefer that wavelength, the recovery attempts now prefer wavelengths which are distributed over the spectrum of available wavelengths. To get the SV to reflect this, the pre-failure wavelength is assigned an SV value of 0. The pre-failure wavelengths of other affected connections are penalized by giving them a SV value of 2^1 , reducing their chance of being chosen. Unused labels are assigned a SV value of 1. At each WC en-route, the SV value is increased by 1, as in normal SV operation. The principle of the BR scheme is illustrated in figure 4.2.

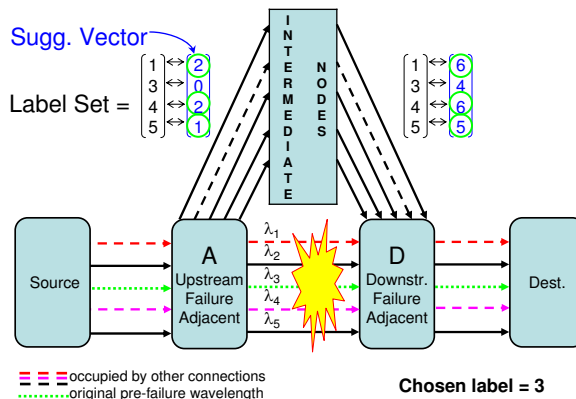


Figure 4.2: SV modified for blocking reduction.

¹The value of 2 was found optimal after investigating several other values.

4.2.3 Simulation Study

The BR scheme's performance is investigated in terms of recovery retries due to blocking, recovery percentage, used WCs and length of the restoration path. The performance of the BR scheme is compared to the aforementioned SV scheme and a No Preference (NP) scheme. The NP scheme blindly selects a wavelength amongst the wavelengths that are free on the last hop².

Topology

The schemes are simulated in OPNET Modeler [92] and evaluated on the Pan-European triangular topology network [93,94], which is shown in figure 4.3. The network consists of 28 nodes and 61 spans, resulting in a nodal degree of 4.36.

Simulation Execution

Each span has a capacity of 10 wavelengths in each direction. The network is populated with unidirectional connections up to a desired average network load using the SV scheme. The same scheme is used in all cases to provide equal starting points for the recovery study. Each connection has the size of one full wavelength. Both the number of available WCs per node and the load is varied. After reaching the desired load, the spans are consecutively failed. The restoration paths are found using the Dijkstra k-shortest path algorithm [34,95]. Recovery is attempted using appropriate RSVP-TE signaling, using either NP, SV or BR in the wavelength selection process. First-fit is used as a tie-breaking policy, meaning that the wavelength with the lowest index is preferred amongst equally good wavelengths [91]. If the restoration request experiences blocking, the causing span is temporarily removed from the path computation using information from the GMPLS experimental crankback feature [96]. Recovery is then re-attempted over the k-shortest path.

²Further details on the NP scheme can be found in Appendix B.2.




Figure 4.3: Pan-European network used in the simulation study.

Scenarios

The scenarios shown in table 4.1 are evaluated. Note that all scenarios use the same scheme for provisioning to achieve equal starting points, and that it is only relevant to use the BR scheme in the recovery phase.

Scenario name	Provisioning scheme	Recovery scheme
SV/NP	SV	NP
SV/SV	SV	SV
SV/BR	SV	BR

Table 4.1: Overview of simulation scenarios.

4.2.4 Results

All results presented in this section are obtained by averaging 40 repetitions with different random seeds. Each data point shows 95% confidence

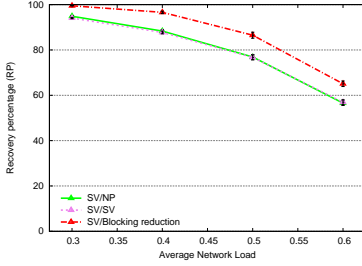
interval.

Figure 4.4 shows the recovery percentage under varying (a) network load and (b) number of available WCs per node. The recovery percentage is calculated as the ratio between the number of recovered connections and the number of affected connections for each span failure, averaged over all span failures. For all schemes, the recovery percentage decreases when the network load increases, which is due to the lack of spare resources. The BR scheme achieves the highest recovery percentage. This is also the case under varying WCs per node but the advantage diminishes when WCs are plentiful.

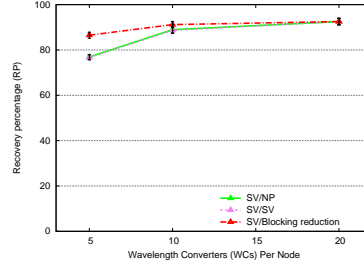
In figure 4.5 the number of recovery retries due to forward and backward blocking is illustrated. Only recovery requests that are successfully restored are counted, i.e., recovery attempts that fail even after several recovery retries are not included in the statistics. The number of recovery retries is divided by the number of recovered connections to provide a fair comparison of scenarios under different recovery percentages. In figure 4.5(a) under varying load conditions, the BR scheme causes fewer blockings and thereby fewer recovery retries than the NP and SV schemes. At high loads, all schemes achieve similar performance. In figure 4.5(b) the BR scheme again outperforms the others at few WCs per node whereas the number of recovery retries is similar when many WCs are available.

The average WC usage per recovered connection is depicted in figure 4.6. The converter usage increases with increased network load, because the chance for getting a wavelength continuous path is reduced at high loads, since connections must be placed between already provisioned ones using whatever wavelength is free. The converter usage, in increasing order, is $SV/BR \rightarrow SV/SV \rightarrow SV/NP$, for both varying load and WCs per node. The strength of the BR scheme is a potentially lower converter usage at the failure adjacent nodes which can result in an increased recovery percentage if WCs are the limiting factor for successful recovery.

Figure 4.7 shows the average hop count of the recovered connections. Under varying load the BR scheme has a lower hop count compared to the other schemes. This is explained by the fact that it experiences fewer blockings. Since each blocking results in the removal of the span causing the blocking from the path computation, a recovery retry will



(a) Recovery Percentage with varying load. 5 WCs per node.




(b) Recovery Percentage with varying number of WCs. Network load fixed at 0.5.


Figure 4.4: Recovery Percentages for varying (a) load and (b) number of WCs per node.

generally have to be attempted over a longer path. An increasing number of available converters per node significantly reduces the hop count for all schemes.

As a general comment on the results it should be noted that the performance of the SV/NP and SV/SV are very similar in terms of recovery percentage, recovery retries and hop count. This is likely caused by the fact that the SV scheme is used during the provisioning phase for all scenarios. An economic converter usage results in many WCs being available in the recovery phase. Converter minimization schemes show their strength when WCs are deployed very sparsely (i.e., even less than 5 WCs per node) and it is therefore future work to extend the simulation studies to cover severely limited converter placement allowing a fuller evaluation of the benefits of the proposed schemes.




(a) Recovery Retries with varying load. 5 WCs per node.




(b) Recovery Retries with varying number of WCs. Network load fixed at 0.5.

Figure 4.5: Recovery Retries for varying (a) load and (b) number of WCs per node.




(a) WC Usage per recovered connection with varying load. 5 WCs per node.




(b) WC Usage per recovered connection with varying number of WCs. Network load fixed at 0.5.

Figure 4.6: WC Usage per recovered connection for varying (a) load and (b) number of WCs per node.



(a) Hop Count of recovered connection with varying load. 5 WCs per node.



(b) Hop Count of recovered connection with varying number of WCs. Network load fixed at 0.5.

Figure 4.7: Hop Count of recovered connection for varying (a) load and (b) number of WCs per node.

4.2.5 Conclusions on Blocking Reduction

This study shows that the presented BR scheme for blocking reduction in span restoration causes less blocking in comparison to the NP and the SV schemes. Blocking minimization increases the recovery speed since fewer recovery retries are necessary, which eases the load on the control plane. The study has also shown that the BR scheme achieves a higher recovery percentage. The advantage of the BR scheme is predominant if WCs are limited and at low-medium load ranges.

It is expected that the results are dependent not only on the recovery phase, but also on how the WCs are administrated in the provisioning phase. The results have shown that the BR scheme outperforms the other schemes especially if few WCs are available. Therefore, it is expected that the performance gap between the schemes increases under the following two conditions: first, if a more wasteful scheme such as NP is used in the provisioning phase, leaving less WCs available for the recovery; and second, if WCs are deployed extremely sparse (i.e., less than 5 per node). The BR scheme is expected to outperform the other schemes under these conditions, since it uses fewest conversions. This reasoning is supported by the fact that the BR scheme achieves the best recovery performance especially if few WCs are available per node. Evaluation of these two effects is left for further study.

In this study, the BR scheme has been used in conjunction with the converter-saving feature of the SV. The reduction of blocking of recovery requests has shown that SV based schemes can be used for traffic-engineering purposes in addition to converter minimization. The scope of SV-based schemes can be extended to allow for other wavelength specific traffic engineering metrics. It should also be noted that BR is a simple scheme which only requires modifications at the failure adjacent nodes, while the intermediate nodes do not need to take any actions other than the usual SV operation.

4.3 Stub-Aware Span Restoration

4.3.1 Motivation

Span restoration has the advantage of a short notification time and few signaling actions due to the proximity of the failure location and the place where the recovery is initiated. The connection stubs are reused as a characteristic of the method, so the stubs must be merged with the restoration path at the failure adjacent nodes, which is illustrated in figure 4.8. This implies that if the stub wavelength is different from the restoration path wavelength, a WC must be utilized at the failure adjacent nodes. In converter-limited networks this merging may be complicated by the lack of needed WCs at the failure adjacent nodes. As described in chapter 3, WC availability heavily influences the restoration percentage. This section presents an extension of the SV, which aims at reducing wavelength conversion at the failure adjacent nodes. It should be noted that the stub-aware scheme can be used in any recovery context except where connections are restored on an end-to-end basis, meaning that the scheme is relevant for segment restoration as well. The following explanation of the scheme is based on span restoration since it has to merge the recovery path to the connection stubs at two nodes.




Figure 4.8: Span restoration concept.

4.3.2 Method and Operation

When a route has been found using the Open Shortest Path First (OSPF)-TE protocol, the RSVP-TE protocol is used to reserve the necessary resources on that route. This is done by propagating *Path* and *Resv* messages between the end points of the connection. The *Path* message is used to request a resource, and it may contain optional objects such as the Label Set (LS) [87] or the SV [14,21]. Two basic decisions must be made before selecting a wavelength (i.e., a label) for a given connection:

- Which labels are included in the Label Set, i.e., Label Set composition.
- Which label is preferred within the Label Set, i.e., label selection.

Stub-awareness is designed with the main goal of influencing the label selection and is therefore based on an extension of the SV, but it can also influence the Label Set composition.

When setting up the restoration path, the wavelength assignment schemes, i.e., LS and SV, can either disregard (stub-unaware) or use (stub-aware) the information about the stubs of the failed connection. In the stub-unaware scheme, no stub information is used and the best possible wavelength for the recovery path is found without taking potential conversions at the failure adjacent nodes into account. The stub-aware scheme takes stub information and potential wavelength conversion at the failure adjacent nodes into account during the wavelength selection process.

When executing span restoration with the stub-unaware scheme, the label selection works as follows: all wavelengths are given equal preference at the node initiating the restoration path. Similarly, the destination node does not consider the stubs when choosing the wavelength. Hence, the SV-value only contains the number of WCs used on the bypass route. This means that the potential conversions to the stubs at the failure adjacent nodes are not counted, which may lead to a waste of WCs and hence lower restorability. When using both SV and LS in stub-unaware mode, all free wavelengths are included in the LS. In the special case of 0 available WCs in the upstream failure adjacent node, the stub-unaware scheme has the unwanted behavior of adding all available wavelengths on the next span when composing the LS, although

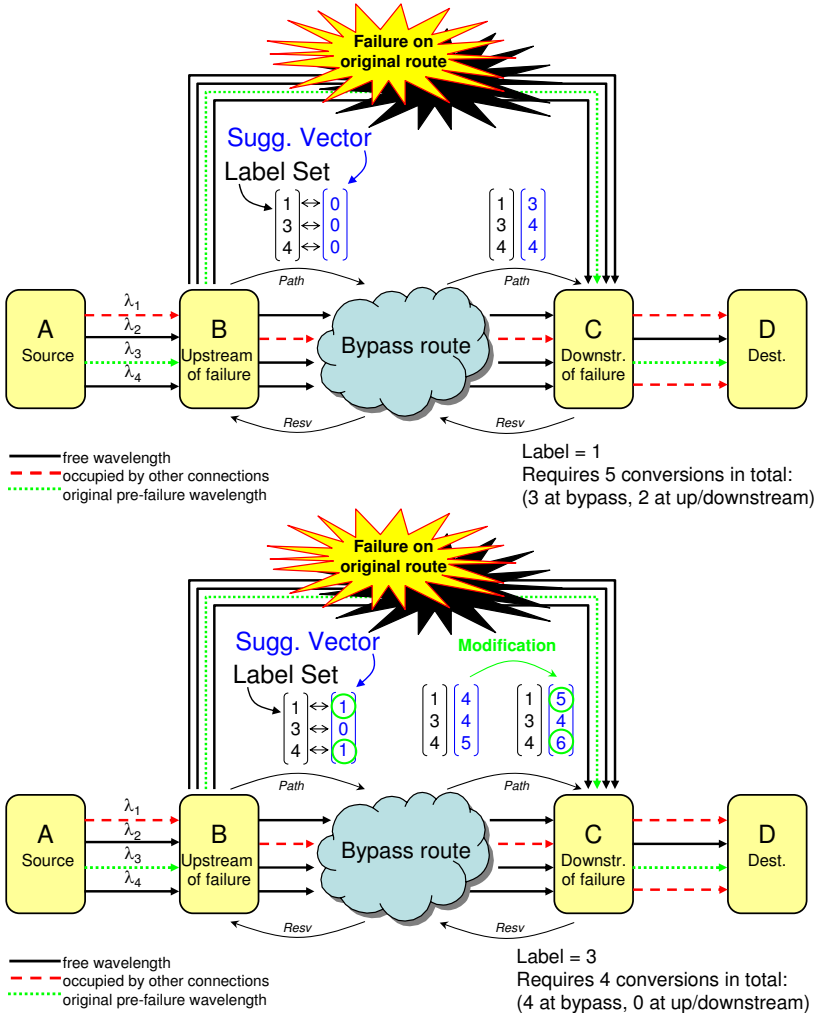


Figure 4.9: Suggested Vector application in span restoration. Top: stub-unaware scheme, Bottom: stub-aware scheme with modifications circled.

only the stub wavelength can be successfully merged to the restoration path. With the stub-aware modification the LS scheme disregards unusable labels if there are no free WCs at the upstream failure adjacent node.

Four different combinations of Label Set composition and label selection are possible, which are illustrated in table 4.2.

Scheme	Label Set Composition	Label Selection
LS	Unaware	Not based on WC info
Stub-aware LS	Stub-aware	Not based on WC info
SV	Unaware	Using WC info of bypass route
Stub-awareSV	Stub-aware	Using WC info of bypass route and failure adjacent nodes

Table 4.2: Label Set composition and label selection for stub-awareness

The aim of the stub-aware extension is to avoid unnecessary conversion at the failure adjacent nodes. The SV scheme is therefore modified by giving the non-stub wavelengths a value of 1, as their choice requires a WC at the upstream failure adjacent node. The stub wavelength keeps its SV-value of 0, since choosing it requires no conversion when merging the recovery path to the original path. At the downstream failure adjacent node accordingly, the SV is examined and modified by adding 1 to the SV-values for all wavelengths except the stub wavelength, since their choice induces a WC. The SV-values now represent the necessary WCs at the upstream failure adjacent node, the bypass route and the downstream failure adjacent node. Afterwards, the wavelength with the lowest SV-value is chosen and propagated in the *Resv* message. The concept of the SV and the stub-aware SV is illustrated in figure 4.9. The stub-aware scheme has the advantage that only the failure adjacent nodes have to modify their behavior, whereas no change is required at intermediate nodes on the restoration path.

4.3.3 Simulation Study

In this study, the effect on the recovery percentage of varying the number of WCs per node and the network load, when using the standard LS scheme and the enhanced SV scheme, both with and without stub-awareness in span restoration, is investigated.




Figure 4.10: NSFNET used in simulation study.

Topology

We evaluate the schemes' performance in the NSFNET [97], illustrated in figure 4.10, using OPNET Modeler [92]. The network consists of 14 nodes and 21 spans, resulting in a nodal degree of 3.14. There are 10 wavelengths per link per direction.

Simulation Execution

The network is populated with unidirectional connections with infinite holding time based on a uniform distribution of the source/destination pairs until a specified average network load is reached. The chosen wavelength assignment scheme is used in both the provisioning and restoration phase. Once the desired load is reached, a span is failed and recovery is attempted. If the recovery request experiences blocking, the blocking-causing span is temporarily removed from the topology database using the experimental crankback feature [96], before restoration is reattempted. When all connections are either restored or found unrecoverable, the network is reversed to its pre-failure state before the next span failure is evaluated. This procedure continues until all possible single span failures have been simulated.

Scenarios

The LS-based scenarios represent the case where the network nodes support only standard RSVP-TE protocol objects. In the SV-enabled sce-

nario, it is assumed that all nodes support the SV optional object, and the SV is used in both the provisioning and restoration phase. In the stub-aware scenarios, plain LS or SV is used in the provisioning phase, since stub-awareness is only relevant during the recovery phase. If several equally good wavelength choices exist a first-fit tie breaking policy is applied [91]. An overview of the different simulation scenarios is given in table 4.3. Details of the individual schemes' Label Set composition and label selection were described in table 4.2.

Table 4.3: Simulation scenarios

Simulation scenario	Provisioning scheme	Restoration scheme
LS	LS	LS
Stub-aware LS	LS	Stub-aware LS
SV	SV	SV
Stub-aware SV	SV	Stub-aware SV

4.3.4 Results

All results are obtained by averaging 20 repetitions with different random seeds. The recovery percentage is calculated as the ratio between the number of recovered connections and the number of affected connection for each span failure, averaged over all span failures. The error bars show the confidence intervals at 95% confidence level.

Figure 4.11 illustrates the recovery percentage obtained with a varying number of WCs per node at a constant average network load of 0.5. If no WCs are available, the LS and the SV perform the same, whereas the stub-aware extension for both schemes achieves a better performance. This is due to the fact that the LS only contains the wavelength that ensures a wavelength-continuous path on the bypass route and the connection stubs, hence avoiding potential merging issues at the failure adjacent nodes. When more WCs are available the recovery percentage increases. Also, the SV outperforms the LS in terms of recovery percentage. The advantage of the stub-aware modification for both the LS and the SV is especially predominant if WCs are allocated very sparsely, i.e., 5 or fewer WCs per node. If many WCs are available, they are no longer the limiting factor, and the difference between the schemes diminishes.

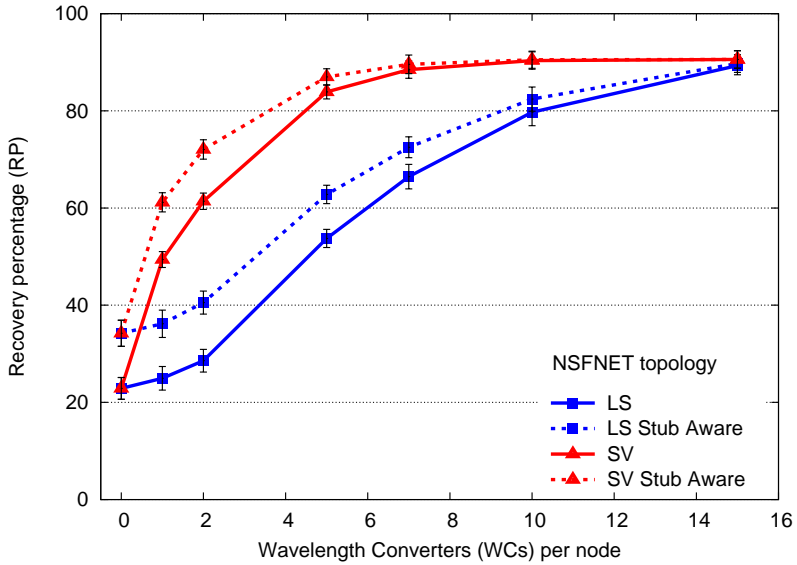


Figure 4.11: Recovery percentage with varying number of WCs per node. Average network load is fixed at 0.5.

Figure 4.12 shows how the recovery percentage for the different schemes changes with a varying average network load. There are 5 WCs available per node. As expected, the recovery percentage decreases when the average network load is increasing, because more connections are affected and fewer spare resources are available for recovery purposes. The stub-aware modification gives the highest performance improvement (compared to the stub-unaware case) at medium load ranges, where the WCs are the limiting resource for restoration.

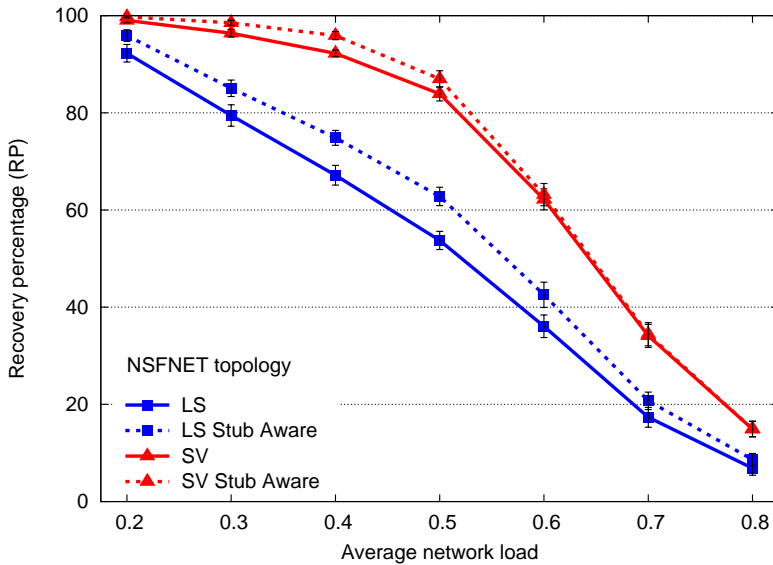


Figure 4.12: Recovery percentage with varying average network load. 5 WCs per node.

4.3.5 Conclusions on Stub Awareness

The recovery performance of a standard (LS) and WC-saving (SV) wavelength assignment scheme with and without stub-awareness is evaluated in a GMPLS controlled optical network with limited wavelength-conversion capability. The recovery percentage is investigated in the scenario of a single span failure using span restoration. The performance is compared under varying network load and WC availability. Simulation results show that both the LS and the SV scheme benefit from the stub-aware modification. The modification is especially useful if a limited number of WCs are available and at medium load ranges. The SV scheme outperforms the LS scheme. The stub-aware modification of the SV only has to be executed at the failure adjacent nodes and does not affect the operation at the intermediate nodes, so the added complexity is minimal.

4.4 Recovery in Multi-Domain Networks

4.4.1 Background

In this section, some consideration is given to extending the SV scheme to a multi-domain environment with heterogeneous technologies and administrative policies. Automatically Switched Optical Network (ASON) [28] provides a framework to interconnect several GMPLS domains. Such an environment is shown in figure 4.13, where the interfaces between GMPLS domains are specified by External Network-to-Network Interface (NNI). Topology information is normally not disclosed through the E-NNI [98]. The reason for this is twofold: first, disclosing the routing information would limit the scalability of the interconnected network by significantly increasing the routing overhead. Second, a network operator would not disclose his topology information for competitive reasons and to avoid exposing weaknesses in the network design. Hence scalability and confidentiality are the key arguments to hide the topology information. The SV does not breach these confidentiality requirements since it only contains summarized information. Therefore the information contained in the SV cannot be used to derive the other domains' topologies.

4.4.2 Policies

An issue that arises with multi-domain networks and the SV scheme is possible unfair WC distribution between the domains. The SV scheme provides information on the required number of WCs in the available optical paths. The overall number of WCs is useful for the destination node to choose the wavelength causing fewest conversions overall. An inherent assumption in the SV scheme is that all nodes agree on the optimal choice of wavelength (i.e., the one leading to minimum wavelength conversion for the entire path). In a multi-domain environment, the opinion on the optimal label may differ between domains, since the choice of one label may minimize the WC usage in one domain at the expense of more conversions in the other. A malicious neighboring network might choose a path with unfair distribution of the WCs and a cost model should thus be agreed between the networks participating in the E-NNI.




Figure 4.13: Multi-domain network environment.

4.4.3 Multi-Domain Recovery

The lack of routing information flow across E-NNI interfaces poses a challenge in the backup route computation of span restoration. Normally, when the working path traverses several domains, the backup path is computed separately for each domain [45]. For span restoration, this will obviously cause problems when the failed span connects two domains. Even if the failed span is entirely within one domain, the best backup route may go through a second domain. Such a situation is depicted in figure 4.14, where the neighboring domain provides a lower hop count, which makes this path preferable (any other metric could be used to decide the most feasible path). However, finding this path requires modifications to the traditional domain-by-domain route computation. Today the use of a neighboring domain for failure recovery is not used, however, new approaches for multi-domain inter-operability, as described in [99], are emerging. Increased cost efficiency obtained by using another network are the main driver for such recovery actions.




Figure 4.14: Span restoration in single-domain and multi-domain environments. Restoration through a neighboring network may provide a more cost-efficient solution in terms of hop count or any other measure.

4.4.4 Evaluation Concept

The impact of the limitation imposed by the E-NNI is in the following used to define three simulation scenarios. The simulations have not been performed yet, but the discussion is presented here to guide and illustrate future work. The scenarios differ on the amount of information that is disclosed from each network through the domain interface. In the first scenario, the reference scenario, both routing and SV information is disclosed. As such, this scenario is similar to the situation where all the nodes in question belong to the same administrative domain. The second scenario is restricted to the current implementations of the E-NNI [28]. Hence, neither routing nor SV information is passed. Finally, in the third scenario only the SV information is passed. An overview of the simulation scenarios is provided in table 4.4.

Scenario 1 should be used for reference only as the scalability and the

Table 4.4: Multi-domain information disclosure scenarios

Scenario	Disclosed information between domains
Scenario 1	Routing and SV information
Scenario 2	None
Scenario 3	SV information

confidentiality issues related to visible topology information cancel the interest for further implementation. Scenario 2 should be evaluated to determine the impact of restricting routing and SV information. Finally, for scenario 3, the performance compared to scenario 2 should be clarified and the inclusion of the SV information in the signalling protocols should be considered. In addition, measures to ensure a fair distribution of WC resources between the involved domains must be agreed upon, either by strict policies or by cost models.

4.4.5 Conclusions on Multi-Domain Recovery

The previous sections discussed how the SV scheme can be applied to multi-domain networks with the main goal of multi-domain recovery. The advantage of using the SV scheme is scalability and non-disclosure of topological information. When passing SV information between domains, the involved domains must agree on a policy and cost model in terms of how the obtained information should be used to ensure a fair distribution of resources (e.g., WCs) throughout the different domains.

One could argue that even future optical networks would employ Optical-Electro-Optical (OEO) conversion at domain boundaries. From a converter point of view this would render SV passing between domains obsolete, as the border node could use the OEO conversion to choose the best wavelength for its domain. However, the concept of the SV can be extended from carrying WC information to containing traffic engineering information that can be accumulated throughout different domains.

4.5 Chapter Summary

This chapter presents several methods that improve the recovery performance through signaling enhancements. Two simulation studies are

presented. The first one allows for a blocking reduction (BR) of restoration requests through spreading of the requests based on their original wavelength on the failed span. The BR scheme improves the recovery percentage as well. The second approach avoids WC usage at the failure adjacent nodes through stub-awareness and thereby solves path merging issues in span and segment restoration, which has a positive influence on the recovery percentage if WCs are limited. Furthermore, the concept of extending the SV to a multi-domain environment is considered.

Chapter 5

Stub Release

5.1 Introduction

Stub release refers to freeing the still functioning parts of a connection that is affected by a failure. In several studies stub release has been treated within an optimization problem [34, 68, 100, 101]. A given set of working and backup connections must be routed in the network in such a way that the total used capacity is minimized, depending on whether stub release is used or not. An overview of different approaches for capacity usage optimization for path protection is given in [102]. These studies show that stub release can increase the capacity efficiency but that the actual benefit is dependent on the studied network scenario. The main reason why the effect of stub release is small in these studies is that all connections are known beforehand and can be routed off-line aiming at capacity minimization.

This situation is different from the scenario investigated here, where predefined network resources and dynamic connection arrival and recovery are considered. A tradeoff between the potentially higher recovery percentages and the operational complexity of failure location detection, signalling and reversion for the use of stub release emerges. Ideally, the drawbacks of stub release could be avoided without penalizing restorability. In this chapter, two methods are presented that seek this goal. The first one substitutes the need for stub release by using an enhanced wavelength assignment scheme, while the second method modifies traditional stub release to exclusively release node resources.

5.2 Full Stub Release

5.2.1 Motivation

In this section, the recovery percentage for a converter-saving wavelength assignment scheme (i.e., Suggested Vector (SV) [14,21]) is compared to a simple wavelength assignment scheme with and without stub release, in order to investigate if the drawbacks of stub release can be avoided by using intelligent wavelength assignment.

5.2.2 Releasing Span and Node Resources

In this study, local-to-egress restoration is used as recovery method. This means that stub release is performed by letting the downstream failure adjacent node send teardown messages towards each affected connection's destination node. Upon reception of such a message each node frees the stub's labels and Wavelength Converters (WCs), if used. These resources are now available for establishing restoration paths. The principle of stub release in a network equipped with WCs is illustrated in figure 5.1. Due to the freeing of resources, a better restoration success rate is expected when stub release is used, especially when available labels and WCs are limited. The question remains whether a more sophisticated label assignment scheme (i.e., the SV) can achieve the same performance without using stub release. If yes, then the drawbacks of stub release can be avoided without sacrificing restorability.

5.2.3 Simulation Study

The restoration efficiency of the SV scheme compared to the simple No Preference (NP)¹ scheme with and without stub release applied to local-to-egress restoration is evaluated. The SV scheme for local-to-egress restoration is shown in figure 5.2. From the upstream failure adjacent node, the Label Set (LS) and the SV are propagated towards the connection's destination. Along the restoration path, the SV is updated when WCs are required and the destination node can then choose the

¹Further details on these wavelengths assignment schemes can be found in Appendix B.2.




Figure 5.1: Restoration of the connection shown at the top with and without stub release. Bold spans are fully occupied. Solid line denotes working path. Dashed lines denote restoration path. The numbers in some nodes state the number of free WC's. A WC is required at the nodes with yellow background.

label requiring fewest WCs along the route and signal this in its *Resv* messages.

Scenarios

Four different scenarios are evaluated:

1. **NP:** Connections are setup and restored with the NP label assignment scheme. Stub labels and WCs are not released.
2. **NP Stub Release:** Connections are setup and restored with the NP scheme. When a failure occurs and before initiating restoration requests, stub release is performed (i.e., freeing the connections labels and WC on the non-failed spans).
3. **SV:** Connections are setup and restored with the SV label assignment scheme. Stub labels and WCs are not released.

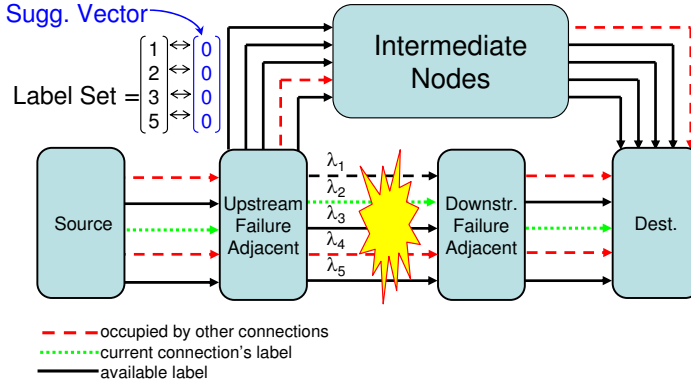


Figure 5.2: Suggested Vector used in local-to-egress restoration.

4. **SV Stub Release:** Connections are setup and restored with the SV label assignment scheme. Stub labels and WCs are released.

Evaluated Topologies

The simulations are carried out in OPNET Modeler [92] and evaluated in the NSFNET [97] and the Pan-European triangular topology network [93,94]. The network topologies are illustrated in figure 5.3 and 5.4, respectively. The number of nodes, spans and the nodal degree for both network topologies are illustrated in table 5.1.

Topology	Nodes	Spans	Nodal Degree
NSFNET	14	22	3.14
Pan-European	28	61	4.36

Table 5.1: Evaluated network topologies.

The simulation process is the same as described in chapter 3 and chapter 4, except for the fact that stub release actions may be performed before restoration is attempted. Each span has a capacity of 10 wavelengths (i.e., labels) in each direction. The network is incrementally populated with unidirectional connections of infinite holding time up to a given average network load (0.4 to 0.6). Each connection occupies one wavelength. The connections are uniformly distributed over all source-




Figure 5.3: NSFNET network topology.




Figure 5.4: Pan-European triangular topology network.

destination pairs in the network. When the desired load is reached, stub release may be executed. The spans are failed consecutively, meaning that the network is reverted to its pre-failure state before the next span failure is simulated.

5.2.4 Results

The results in figure 5.5 and figure 5.6 show the recovery percentages for the NSFNET and the Pan-European network, respectively. Results calculated with a confidence interval of 95% confidence level over 40 random seeds. The recovery percentage decreases at increasing network load, since more connections are affected by the failure and they must share fewer residual resources. The NP scheme allows for a higher restoration percentage when stub release is enabled, compared to when no stub resources are freed. The NP scheme may assign labels in an inefficient manner, resulting in a high WC usage per connection. This leaves few WCs for the restoration process, meaning that the possibility for successful connection restoration is diminished especially when WCs are sparse. If the stubs are released, these resources can be used during the restoration process and more connections can be restored.

More importantly, we see that the SV scheme without using stub release gives a better restoration success rate than both NP and NP with stub release. This is due to the fact that the SV scheme assigns labels in a converter-saving manner, providing more WCs for the restoration process, hence allowing more connections to be restored. The influence of the network topology on the recovery percentage can be seen by comparing figure 5.5 and figure 5.6. For the NSFNET topology, the recovery percentage for SV compared to NP with stub release is relatively increased by approximately 15%.² For the Pan-European topology, the increase is in the order of 25%. The combination of SV with stub release has also been investigated. In this case however, stub release only provides a small performance increase, which does not really justify adding the additional complexity to the restoration process.

²calculated as:

$$\frac{\text{mean rec. \% (SV)} - \text{mean rec. \% (NPwithStubRel)}}{\text{mean rec. \% (NPwithStubRel)}}, \text{ averaged over all loads}$$




Figure 5.5: NSFNET recovery percentages for varying loads. 5 WCs per node.

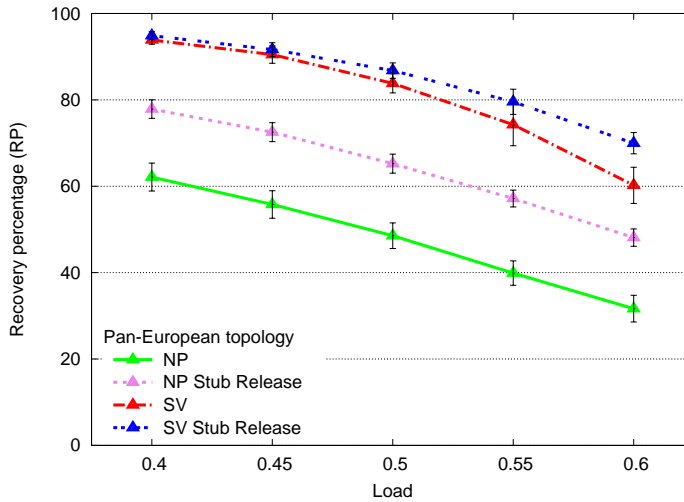


Figure 5.6: Pan-European network recovery percentages for varying loads. 5 WCs per node.

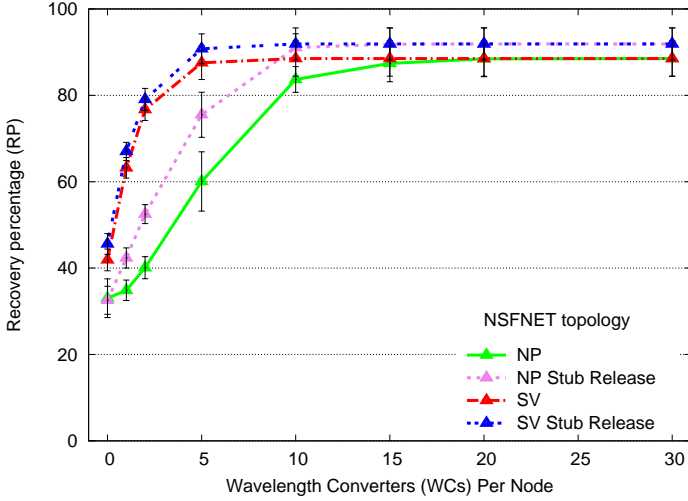


Figure 5.7: NSFNET recovery percentages for varying WCs per node. The network load is fixed at 0.5.

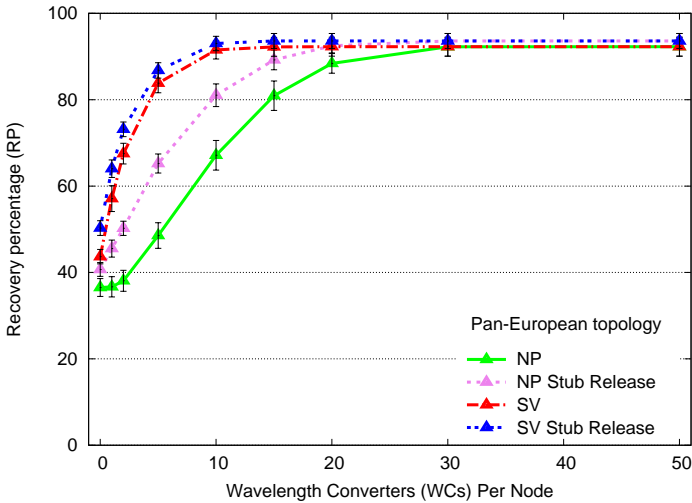


Figure 5.8: Pan-European triangular topology network recovery percentages for varying WCs per node. The network load is fixed at 0.5.

Figure 5.7 and figure 5.8 illustrate the recovery percentages of the different schemes when varying the number of WCs that are available in the nodes, while the load is kept constant at 0.5. It can be seen that the SV scheme is especially useful in a situation where WCs are limited. If numerous WCs are available the advantage diminishes, which is due to the fact that the WCs no longer are the limiting factor. In turn, the span resources are now the limiting factor for successful recovery. Hence, all four schemes converge to an upper limit where adding more WCs does not improve the restoration behavior. At this saturated state, the schemes with stub release achieve a higher restoration percentage than the ones without, because their restoration process can make use of the freed span resources. The effect of stub release in a span limited network is predominant in the NSFNET topology, while the Pan-European benefits from its dense topology to provide more path alternatives. Also, it can be noted that the span limitation occurs at fewer WCs per node in the NSFNET topology compared to the Pan-European topology, which is again due to their topological characteristics.

5.2.5 Conclusions on Full Stub Release

The simulation results presented in this section show that the SV scheme without stub release outperforms the NP scheme with and without stub release. This means that the SV scheme is a good candidate for provisioning and local-to-egress restoration when the drawbacks of stub release must be avoided without sacrificing restorability. The best performance is expected when combining the SV with stub release, since the SV is advantageous at few WCs while stub release shows its benefits at few span resources. The simulation results have however shown that the recovery gain of adding stub release to the SV scheme is only in the order of a few percent in sparsely connected networks and negligible in dense topologies. When considering the complexity stub release adds to the recovery and reversion cycle, the benefit of adding stub release to the SV scheme is questionable. Summing up, a more efficient label assignment scheme (performed through proper control plane enhancements, i.e., adding the optional SV object) can significantly increase restoration percentages, even without resorting to stub release actions.

5.3 Nodal Stub Release

5.3.1 Motivation

To setup a lightpath in a wavelength convertible network, two resources must be considered: the span resources (i.e., the wavelength channels) and the node resources (i.e., wavelength-converters). The configuration of wavelength channels in a dynamic network environment requires signalling between the end nodes. In contrast, WCs are managed locally in each node. WCs are costly components and therefore several studies deal with WC placement minimization. In [73] and [74], the Shared Risk Link Group (SRLG) concept was expanded to cover node resources as well, meaning that several backup paths could share for example an Optical-Electro-Optical (OEO) device, and not just span capacity. The work in [83] and the results presented in section 5.2 show that WCs play an important role in the restoration process, as their exhaustion severely diminishes the chance of successful recovery. As stated in section 5.1, the effect of stub release on the total capacity usage for off-line provisioning has been investigated, but conversion limitations are not considered in these studies. This section addresses the challenge of making highly scarce and needed nodal resources, such as WCs, available for the restoration process.

5.3.2 Concept and Operation

Due to the high price of WCs, it is a realistic assumption that only a limited number of these components is provided in the network nodes. Few WCs create a bottleneck situation for conversion. A nodal stub release method where only nodal resources (i.e., WCs) are released during the stub release operation is described in this section. The basic concept of nodal stub release is that during the stub release phase only node resources are released into the spare capacity pool, while span resources are not released. The nodal stub release method provides access to highly demanded WCs during the recovery phase and simplifies the reversion process as re-negotiation of span resources is avoided. A node architecture where the WCs are shared in a per-node converter bank [36] provides high flexibility in terms of accessing WCs through different input ports. The node architecture is illustrated in figure 5.9, and used

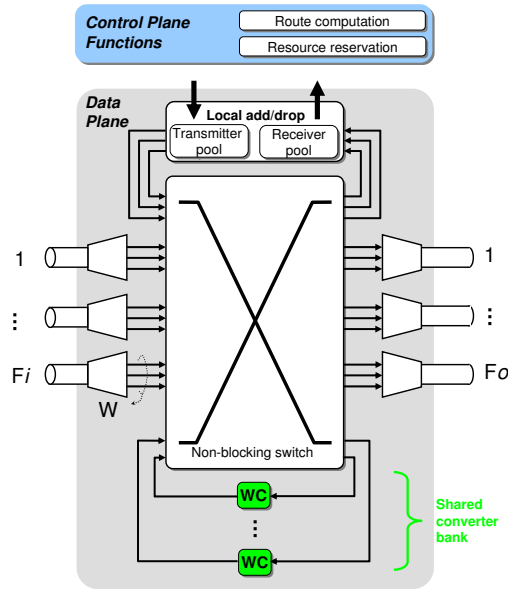


Figure 5.9: Node architecture with shared converter bank.

in the presented study. The closest related work found is the study in [73, 74], where node resources are shared as an extension of the SRLG concept for off-line optimization. This differs from the nodal stub release approach, where the node resources are only released (or shared in other words) after the failure has occurred. Furthermore, the nodal stub release concept is evaluated in a dynamic network scenario without preplanned off-line optimization.

Complication of the reversion phase due to cross-entanglement is one of the major challenges faced by stub release and leads to the development of stub reuse, as described in section 5.1. Stub reuse locks the released resource configuration to a particular connection in such a way that a given connection may only reuse its own released resources. In this way, cross-entanglement is avoided. These restrictions can however lead to a situation where the locking of resources impairs the restoration process. An example is illustrated in the upper part of figure 5.10, where the last available WC is locked to a stub reuse configuration, which hinders the recovery of the connection, as wavelengths which match the WC




Figure 5.10: Restoration example using Stub Reuse (top) and Nodal Stub Release (bottom). Nodes with color transitions denote the WC state. Solid lines mark used capacity, dashed lines is released capacity. Black solid lines denote fully occupied spans. Free wavelengths and WC are denoted at spans and nodes where relevant.

state are not free.

Nodal stub release intrinsically avoids cross-entanglement of span resources since they are not released, but the method does not avoid cross-entanglement of WCs. If the use of released WCs is constrained to a connection's own pre-failure WCs, the cross-entanglement situation is avoided. In that case, the nodal stub release concept is similar to the one of stub reuse, but has greater flexibility, since the node resources can be accessed from many spans. Several options exist on how tight the released WCs should be locked to their pre-failure situation. In the freest condition, the WC may convert between all possible wavelengths at all ports, as long as it is used for the recovery of the particular connection that released it. Under stricter conditions, the WC must still convert between the same wavelengths, but from different input ports or fibers. An example where nodal stub release leads to a successful recovery in contrast to stub reuse is shown at the bottom of figure 5.10.

The concept of nodal stub release and reversion is shown in figure 5.11. To perform nodal stub release, a notification message is propagated from the failure adjacent node (towards the destination node, and possibly source node, depending on the used restoration method), which

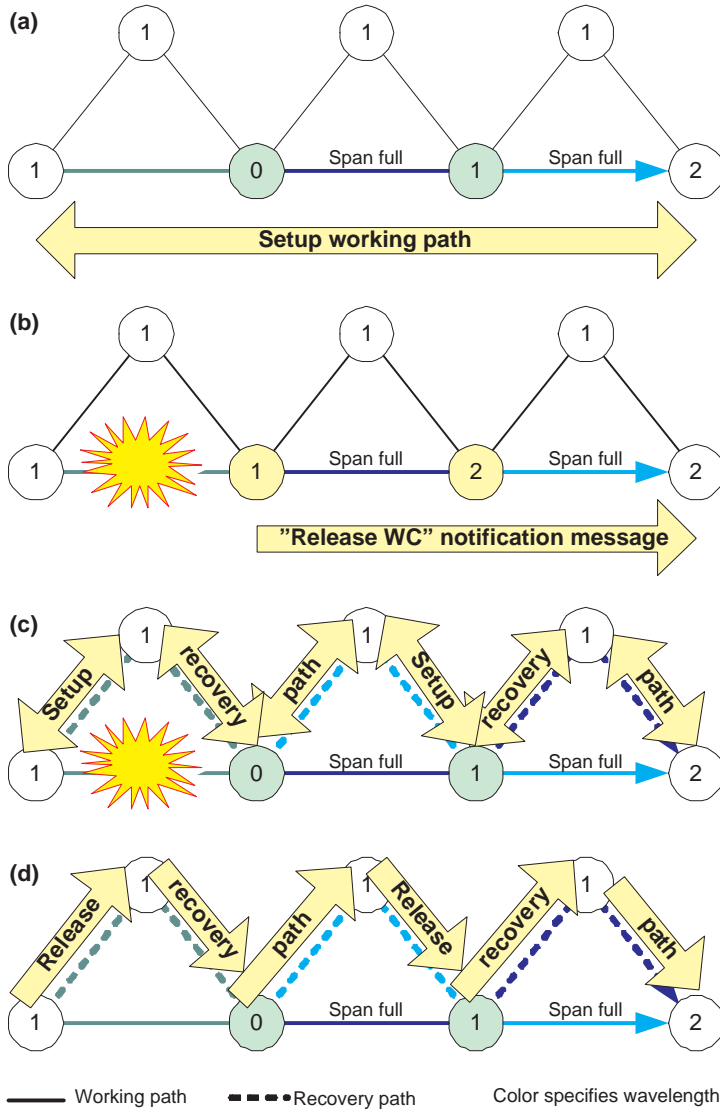


Figure 5.11: Recovery and reversion of one lightpath with nodal stub release. Numbers in nodes indicate available WCs. (a) failure-free lightpath, (b) nodal stub release (WC release), (c) restored lightpath, (d) reversion process.

instructs the nodes along the route to release the WCs of the failure affected lightpaths, while the span resources remain occupied. The released WCs can then be used by restoration requests to restore lightpaths.

Once the failed span is repaired, the traffic is reverted to their original paths. While full stub release requires re-provisioning and re-signaling of the working path and tear-down of the recovery path, nodal stub release only entails tear-down of the recovery path, since the span resources have not been released and the originally used WCs become available by releasing the recovery path. Implementation details depend on the network architecture and the used protocol, and alarm suppression mechanisms must be used during the reversion phase to prevent the temporary connection interruption being treated as a failure.

5.3.3 Simulation Study

This study investigates how the nodal stub release method performs in a dynamic restoration study in three networks with limited WC availability. Generalized Multi-Protocol Label Switching (GMPLS) control plane signaling is used with standard protocol extensions (i.e., the LS) [87]. Local-to-egress is used as restoration method, which combines the advantages of the well-known span and path restoration methods, such as short notification time and high resource efficiency [62]. The simulation procedure is the same as in section 5.2, with the addition of the nodal stub release method.

Scenarios

The following stub release methods are compared:

1. **No stub release:** The stubs' span and node resources (i.e., WCs) are kept occupied during restoration.
2. **Full stub release:** Both the stubs' span resources and the corresponding WCs are released.
3. **Nodal stub release:** Only the stubs' WCs are released, the span resources are kept occupied.




Figure 5.12: German network topology

Evaluated Topologies

The recovery performance is evaluated in three network topologies using OPNET Modeler [92]: the NSFNET [97], the Pan-European triangular topology network [93,94] and a German network [103]. The NSFNET and Pan-European topologies were illustrated in figure 5.3 and figure 5.4 on page 83, while the German network topology is depicted in fig 5.12. The network properties are shown in table 5.2.

Table 5.2: Network topologies.

Network	Nodes	Spans	Nodal degree
German	11	33	6.00
NSFNET	14	22	3.14
Pan-EU	28	61	4.36

5.3.4 Results

In figure 5.13 - 5.15, the recovery percentage is illustrated for a varying number of WCs per node, while the average network load is kept constant at 0.5. These results focus on the performance of nodal stub release relative to full and no stub release, rather than the absolute recovery percentages, which are dependent on external factors such as span dimensioning and WC distribution. For all three networks, if 0 WCs are available in the nodes, the nodal stub release scheme and the no stub release scheme have the same performance, which is expected. In the Pan-European and the NSFNET topology the full stub release scheme, where also span resources are released, slightly outperforms the other schemes, while the improvement is negligible for the German network.

The highest recovery percentages are achieved in the German network (tested topology with highest nodal degree), as shown in figure 5.13. The nodal stub release scheme equals the full stub release scheme, while the no stub release scheme achieves lower performance. If more than 10 WCs are available per node, all three schemes achieve the same performance, since the abundance of WCs and multiple recovery path options due to a high nodal degree compensate for the occupied resources.

Results for the NSFNET (tested topology with lowest nodal degree) are shown in figure 5.14. If few WCs are available, the nodal stub release method achieves a recovery percentage close to the one of full stub release, while the no stub release method has considerably lower performance. This is due to the fact that WCs are a limiting factor for successful recovery, and the nodal stub release scheme makes them available. When many WCs per node are provided, they are no longer the limiting factor, and full stub release, which can also release span resources, obtains better performance.

In figure 5.15 Pan-European network, results are similar to figure 5.14, but the network gets span-limited at more WCs per node due to its larger size. The strength of the nodal stub release scheme is especially dominant when WCs are sparse: nodal stub release can significantly increase the restoration percentage compared to the case of no stub release, while only a minor performance penalty compared to full stub release is observable.

Figure 5.16 depicts the recovery percentage as a function of the network load for the NSFNET. The number of WCs per node is fixed to

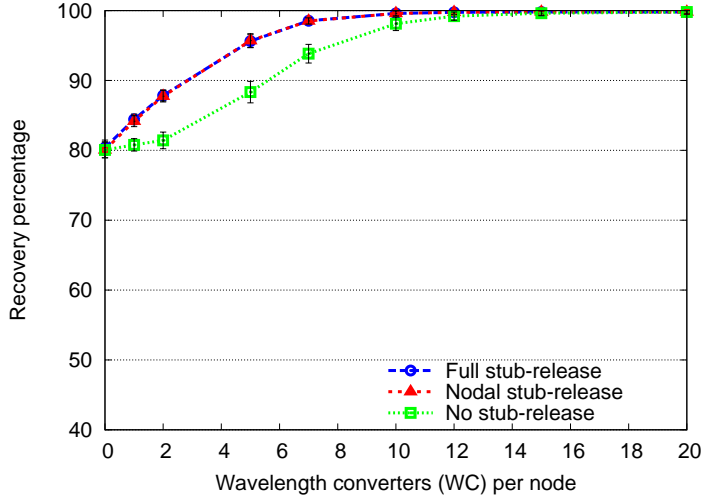


Figure 5.13: German network with varying number of WCs. Average network load is 0.5.




Figure 5.14: NSFNET with varying number of WCs. Average network load is 0.5.




Figure 5.15: Pan-European network with varying number of WCs. Average network load is 0.5.

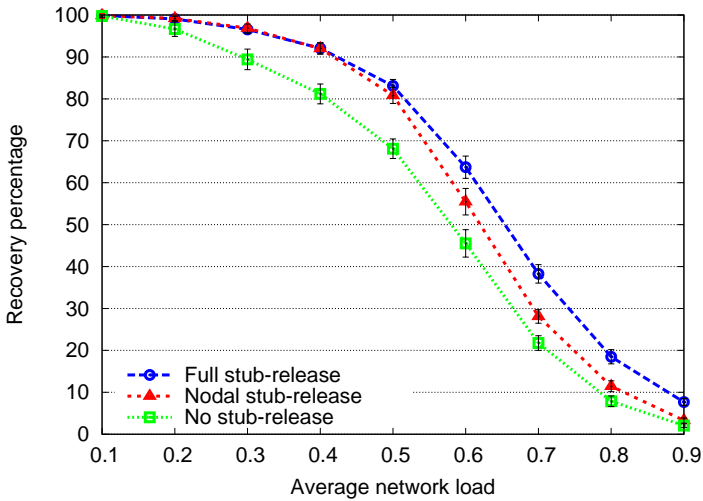


Figure 5.16: NSFNET with varying network load. 5 WC deployed at each node.

5. At all loads, nodal and full stub release achieve better recovery percentages than no stub release. At low loads, the nodal and full stub release methods achieve the same performance, because span resources are widely available. As the network load increases, the performance of the nodal stub release method approaches the no stub release method's performance, as restoration success now mainly depends on the availability of span resources.

5.3.5 Conclusions on Nodal Stub Release

The simulation results show that the nodal stub release method performs well in both highly and sparsely connected topologies and is especially useful when WCs are sparse in low-medium loaded networks. Under these conditions, the scheme matches the performance of full stub release, and it always matches or outperforms the no stub release scheme. Furthermore, nodal stub release avoids re-signaling of span resources and hence facilitates the control plane load.

5.4 Chapter Summary

In this chapter, the benefits and drawbacks of stub release have been discussed, which illustrates a tradeoff between capacity efficiency and complexity. Two simulation studies have been conducted in converter-limited networks. The first one avoids the need for stub release by saving WCs through a sophisticated wavelength assignment scheme without restoration percentage reduction. The second study is based on the idea of making the scarce WC resources available in the recovery phase while still allowing fast and simple reversion to the pre-failure network configuration. Both methods have shown to give good performance from a restoration point of view, and they also have the potential to ease the load on the control plane.

Chapter 6

Resilient Network Design

6.1 Introduction

Network design with the goal of capacity deployment minimization has been the focus of much research in the area of network survivability. Some of the early studies are based on the max-flow min-cut approach [104], where Linear Program (LP) formulations were presented for spare capacity placement [105]. A basic idea of assigning spare capacity based on a set of restoration trails of all spans one at a time was presented in [106].

Based on operations research two main optimization strategies exist: either the working paths for all connections are allocated before the spare capacity is considered (Spare Capacity Allocation (SCA)); or a joint optimization of working and spare capacity is executed (Joint Capacity Allocation (JCA)) [34, 100, 107]. Joint optimization can provide total capacity improvements for most types of networks [107], while for others it offers little improvement over separate optimization [101].

This chapter presents two SCA studies. In section 6.2, an iterative span capacity upgrading study is presented. Section 6.3 describes a recovery method based on local-to-egress protection, where all recovery requests towards the same destination node are bundled to use the same route. Section 6.4 summarizes the findings of this chapter.

6.2 Design of Protected Networks with Modular Capacity

6.2.1 Background

Identifying the best possible resource allocation is complicated by the fact that some capacity units are modular, meaning that they can only be assigned as chunks of a certain size, such as Synchronous Digital Hierarchy (SDH) modules [26] or Wavelength Division Multiplexing (WDM) systems with a certain number of wavelengths. The modularity complicates network optimization problems and increases the network cost due to modularity-related over-provisioning. An example is shown in [108], where modelling with modularity constraints results in less overhead (i.e., assigned, but not used) capacity with increments of 1 than when larger modularity increments are used. Previously, it was a common approach to compute the needed spare capacity for all spans, and then round up to the next modular value [34]. But in [107, 109], modularity is taken into consideration already in the design formulation. In [107], two approaches are presented. The first one is a modularity aware spare capacity placement, where the working path is routed first and then the sum of the working and spare capacities are modularized. The second one employs a joint approach, where the working path may deviate from its shortest path to coordinate with other paths in terms of filling up modules, which gives overall capacity savings.

The work presented in this section deals with iterative span upgrading steps for allocating modular spare capacity for restoration, i.e., SCA. The goal is not capacity optimization as such, but finding a simple comparison of the influence of the upgrading step size and the path selection algorithms. Iterative upgrading based on path selection of a maximum flow minimum cut approach [104, 110] is compared to a shortest path approach [95, 111]. The upgrading modularities used in this work are based on SDH module size because the work was done in collaboration with Tellabs Denmark A/S, who has a strong portfolio of SDH equipment. The results presented here can however be generalized to any modular capacity network, such as a WDM network. SCA is executed, meaning that all working connections are routed before spare capacity is allocated. The modularity requirement is fulfilled by rounding up.

The spare capacity is allocated based on the restoration requirements of all span failures in the network. The failure of the individual spans is ordered, meaning that the network is sufficiently upgraded to recover all traffic demands before the next span failure is simulated. The study is carried out for two path selection algorithms, two upgrade sizes and the case where additional spans are added to the network for restoration purposes.

6.2.2 Restoration Path Computation

Mesh restoration has been shown to require considerably less redundant capacity than rings in order to provide full restoration against any single span failure [112], and is hence chosen as the recovery method to be investigated in this study. Path restoration is chosen for this study, because it operates at the granularity of the individual connection and therefore provides high bandwidth efficiency [26]. The scenarios investigated in this study are based on the assumption that a given set of connections are routed in a network in the best possible way, but without survivability in mind. The case of a single span failure is considered, where the restoration path is computed after the failure has occurred. Restoration paths are chosen based on the Ford-Fulkerson algorithm [110, 113] and the Dijkstra algorithm [95, 111], which are described in the following sections.

Ford-Fulkerson Method

The first restoration method is based on the Ford-Fulkerson algorithm that finds flow augmenting paths through a network [104, 110, 113]. A flow augmenting path is a path that has enough free capacity to accommodate another flow thereby increasing the total flow between a source and destination pair. It is based on the max-flow min-cut theorem which is used to find the maximum flow over a minimum capacity cut, identifying a bottleneck in a network. The graph can be thought of as a network of pipes, where as much material should be pumped from the source to the destination without exceeding the capacity of any pipe. A flow from source to destination assigns each span a value, which is larger than zero, but smaller than or equal to the capacity of the span. At each intermediate node, the flow entering must be equal to the flow leaving

the node. A cut on a graph is composed of a set of spans, such that every path from a source to a destination passes through at least one span of the cut. If the spans of the cut are taken out of the graph, the source and the destination nodes are no longer connected. The capacity of the cut is the sum of the capacities of all spans on the cut. The max flow-min cut theorem states that the maximum amount of a flow is equal to the capacity of a minimal cut. There may be several flows which have the maximum amount of flow, and there may be several cuts with minimal value [104].

The Ford-Fulkerson algorithm can be used to find the max flow and min cut from source to destination in a transport network. The value of the max flow from source to destination is equal to the capacity of a source-destination cut with minimal capacity. The Ford-Fulkerson method consists of the following operations:

1. Set the flow in each span to zero.
2. Find a flow augmenting path from source to destination through the network (a path with free capacity available).
3. Increase the flow along the flow augmenting path with the maximum value that can be carried along the different spans.
4. Go back to step 2.

A special version of the Ford-Fulkerson algorithm only traverses a span if it has free capacity on the span leaving the node, meaning that negative flow is not considered [110]. In figure 6.1, an example network is shown, where the numbers on the span denote "Capacity/Flow". The goal is to find flow augmenting paths between node 1 and node 8, and to identify the cut set. In figure 6.2, all these paths are shown, and the cut is identified at node 1 (marked with a red star). The accumulated maximum flow (15 in the example) is equal to the capacity of the cut. The flow augmenting paths and the cut are shown in the network in figure 6.3. The cut runs between nodes out of which no more capacity is available and therefore one of the spans on the cut must be upgraded if more traffic should be accommodated between node 1 and node 8. In this example the network is fully occupied and therefore the cut is located around the source node.

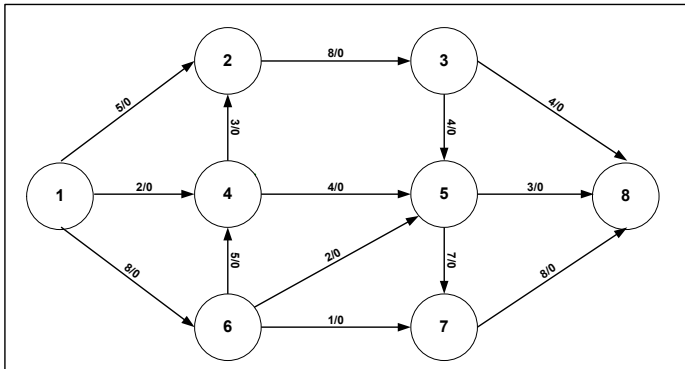


Figure 6.1: Example network. Notation on spans is "Capacity/Flow".

Iteration	Flow augmenting path	Flow along path
1	1 6 4 5 8 0 0 0	3
2	1 6 4 2 3 5 7 8	2
3	1 6 5 7 8 0 0 0	2
4	1 6 7 8 0 0 0 0	1
5	1 2 3 8 0 0 0 0	4
6	1 2 3 5 7 8 0 0	1
7	1 4 2 3 5 7 8 0	1
8	1 4 5 7 8 0 0 0	1
9	problem = dead end ans = here we have a min cut at source ans = Marked nodes with spare cap before cut: ans = 1 0 0 0 0 0 0 0 accmaxflow = 15	

Figure 6.2: Flow augmenting paths and cut identification. The column "Flow augmenting path" specifies the route, while "Flow along path" specifies the size of the flow.

For this study, the original Ford-Fulkerson algorithm is modified in such a way that only flows which have at least the size of the traffic demand for each connection are considered. The flow augmenting paths are therefore the paths where a demand can be routed from the source to the destination. The first identified flow augmenting path is used as a restoration path.

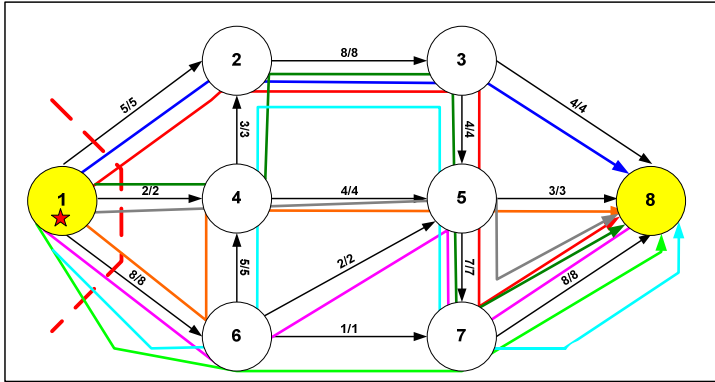


Figure 6.3: Illustration of flows and cut. Red star denotes last marked node. Cut is illustrated by red dashed line.

Dijkstra Method

The second method is based on the Dijkstra algorithm for finding the shortest path between the source node and the other nodes in a network based on a weighted graph. The main idea of the Dijkstra algorithm is to keep identifying the closest nodes from the source node in order of increasing path cost. At the first iteration, the algorithm finds the closest node from the source node. At the second iteration, the algorithm finds the second-closest node from the source node. This node must either be a neighbor to the source node or to the node that was identified as the closest node in the first iteration; else there is a closer node. At the third iteration, the closest node has to be a neighbor of the first two closest nodes, etc. If two nodes can be the closest node because they have the same value, the closest node can be chosen randomly. Additional information about the Dijkstra algorithm can be found in [34, 110, 111].

6.2.3 Span Upgrading Strategies

Upgrading Steps

Two different approaches for upgrading a span are investigated and simulated with the Ford-Fulkerson algorithm and the Dijkstra algorithm for finding end-to-end restoration paths. The first method consists of

upgrading the span capacity by inserting a larger SDH module at the adjacent nodes, which upgrades each span to four times its capacity (Synchronous Transport Module (STM)-16 \rightarrow STM-64 \rightarrow STM-256). This method will be referred to as "x4 capacity" upgrade in the following sections. The second method is applicable when WDM systems are used in the underlying layer. The upgrades are made by taking up an extra wavelength in a WDM system. If the original capacity of the span is e.g., STM-64, the next steps are 128 \rightarrow 192 \rightarrow 256. This method will be referred to as "+own capacity" upgrade in the remaining sections. Furthermore, a simulation where additional spans are added to the network in the beginning of the upgrading procedure is carried out. It is evaluated with both the "x4 capacity" and the "+own capacity" span upgrading strategies. Each upgrading step counts as a span upgrade.

Span Upgrades with Ford-Fulkerson Method

When a connection is affected by a span failure, the algorithm searches the network for flow augmenting paths between the connections' end nodes. The aim is to use the flow augmenting path as a restoration path. If no such paths can be found (i.e., the connection is not restorable), the algorithm searches until it reaches a node which does not have enough free capacity available on its outgoing span(s) and thereby is a dead end. If several spans leave the node, the decision of which one of them should be upgraded first is made according to the following decision steps:

1. Upgrade the span leading to the node which has the shortest distance (hops) to the destination.
2. If there are several possibilities, upgrade the span with the lowest capacity.
3. If there are several possibilities, upgrade the span with the highest utilization (least amount of free capacity).

With this method the spans are upgraded in an iterative manner, meaning that only one span is upgraded at a time and then restoration is retried.

Span Upgrades with Dijkstra Method

To restore a connection, the Dijkstra algorithm finds the shortest path with available capacity from source to destination. If a connection is not restorable, the necessary spans to upgrade need to be identified. With the Dijkstra method, the algorithm calculates a shortest path that only depends on the hop count; whether free capacity actually is available on the used links is not considered at first. When the shortest path is found, the algorithm identifies the spans whose capacity needs to be upgraded to accommodate the connection. Thus several spans that need upgrading can be found in one program iteration with the Dijkstra method, whereas the Ford-Fulkerson method only allows upgrading of one span per iteration.

6.2.4 Path Restoration Simulation

Provisioning Phase

The first step in a network restoration simulation is to decide on a reference network and to provision the working connections. Figure 6.4 illustrates the examined network, which is based on information obtained from Tellabs. It consists of 17 nodes and 27 links and it has a nodal degree of 3.2. A total of 24 bidirectional connections are setup in the network at Virtual Circuit (VC)-4 level [60]. The traffic demands can be seen in figure 6.5. The primary connections are routed by using the built-in OSPF [114] protocol in OPNET [92]. It is assumed that the bidirectional traffic is routed along the same path for both directions. Consequently, only one direction must be calculated. The connection end node with the lower id is chosen as the starting point of the route calculation and is hence called the source. A sample of these paths is shown in figure 6.6.

All spans are assigned equal cost, meaning that the shortest path is equal to the path with the lowest hop count. The average hop count for the primary connections is 2.6 hops. The total traffic on all spans in the network sums up to 1202 VC-4s. For all connections, statistics are kept on how much traffic is assigned to each span. Based on these traffic assignments, the individual spans can be dimensioned by rounding the capacity to the smallest possible module that fulfils the capacity




Figure 6.4: Reference network.

requirements [107] according to the following criteria:

- $0 < \text{VC-4 traffic} \leq 16$: results in an STM-16
- $16 < \text{VC-4 traffic} \leq 64$: results in an STM-64
- $64 < \text{VC-4 traffic} \leq 256$: results in an STM-256

This dimensioning results in a network capacity of 2352 STM units, which is needed to accommodate the total of 1202 VC-4 stated in the traffic demands, which can be seen in figure 6.7. This means that only 51% of the total capacity allocated in the network is actually used for traffic when the network operates in a failure free condition. So in case of a failure, the overhead capacity is available for restoration purposes. It is therefore expected that the overhead obtained from dimensioning the original network could be used to provide restoration paths for a large number of the connections, but preliminary simulations [115] have shown that this is not enough due to the following two reasons. First, even though a large amount of free capacity is available in the network, it is not necessarily located at the appropriate location from a restoration point of

FROM /TO	(8)	(7)	(4)	(3)	(1)	(12)	(15)	(6)	(17)	(2)	(9)	(10)	(11)	(13)	(14)	(16)	(5)
(8)	X	32						18						20			
(7)	32	X				12											
(4)			X		20						28					18	
(3)				X			20		12								20
(1)			20		X			24					12				
(12)		12				X					18		24				
(15)				20			X								20		
(6)	18				24			X		18							16
(17)				12					X				20				
(2)								18		X		24					
(9)			28			18					X			20			28
(10)										24		X			12		
(11)					12	24			20				X			18	
(13)	20										20			X			20
(14)							20					12			X		
(16)			18										18			X	
(5)				20				16			28				20		X

Figure 6.5: Bidirectional traffic demands.

view. Second, six of the nodes only have a connectivity of two. If one of these two links fails, only one span can be used for restoration purposes, which causes bottlenecks if the remaining span is under-dimensioned. Therefore, span upgrades are required to provide full restoration.

Selection of Span Failure Scenario

The connections in the network can be affected by different span failures and each failure results in another group of spans requiring upgrades. At the end of the simulation, all connections should be restorable in any failure scenario. Since this simulation is based on iterative failure evaluation and upgrading, a decision regarding the order of the failure simulations must be made. In this study, the connection carrying the largest amount of traffic is restored first, based on the assumption that

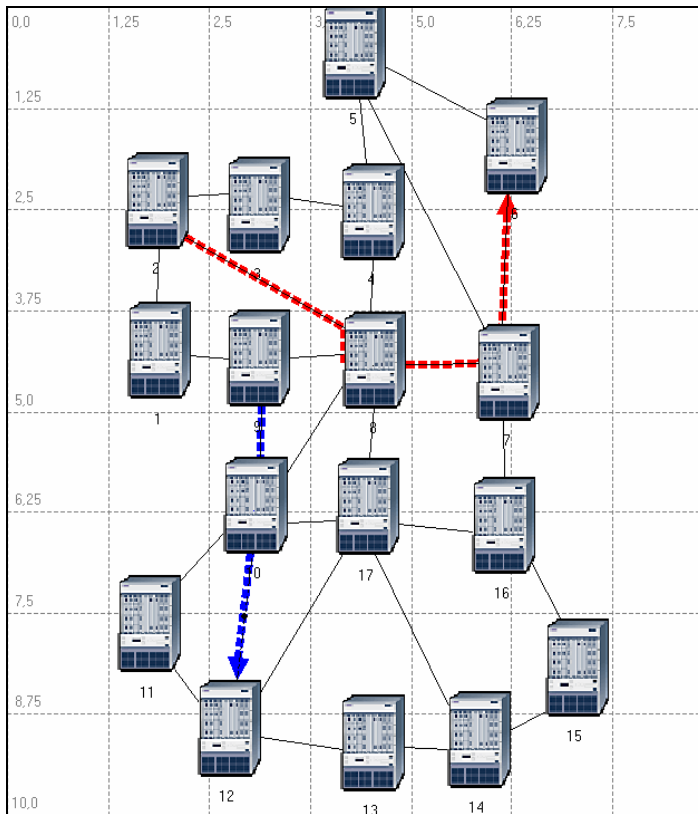


Figure 6.6: Example of OSPF simulation in OPNET. Blue arrow is the shortest path between 9-12 and red arrow between 2-6.

a smaller connection might be restorable from the overhead caused by a span upgrade originating from the restoration of a larger connection. Hence, the span failure scenario is chosen in the following order¹:

1. Choose the largest connection. Choose a span failure that affects it.
2. If more than one connection contains the same amount of traffic, choose the working connection with the lowest hop count. Choose a span failure that affects it.

¹Other orderings are also possible, but have not been evaluated here.

Total VC-4 Traffic	140	28	40	52	44	12	56	98	38	66	16	60	30	20
STM-Unit	256	64	64	64	64	16	64	256	64	256	16	64	64	64
For Span	7-8	4-8	3-4	2-3	1-2	1-9	8-9	2-8	4-5	5-7	5-6	6-7	7-16	15-16
Total VC-4 Traffic	20	52	50	72	24	50	50	24	38	50	12	20	40	1202
STM-Unit	64	64	64	256	64	64	64	64	64	64	16	64	64	2352
For Span	14-15	14-17	16-17	8-17	8-10	9-10	10-11	11-12	10-12	10-17	12-17	12-13	13-14	All Spans

Figure 6.7: Rounding of traffic demands.

3. If several spans have the same traffic and the same hop count, fail the span that carries most connections.
4. If several spans have the same amount of connections, fail the span with the highest capacity value.

By following these steps, a prioritized list of span failure scenarios is obtained to enter into the simulation.

Restoration Simulation

The restoration simulation is carried out in MATLAB [116]. The basic procedure consists of the steps illustrated in figure 6.8, independent of which restoration method and span upgrading strategy is used. The steps with yellow background are specific to the restoration method, i.e., Ford-Fulkerson or Dijkstra. Green steps are related to the upgrading method, i.e., "x4 capacity", "+own capacity". The input to the simulation is the number of the failed span. The first step is to set the capacity of the failed span to zero. Then, the affected connections are identified. Stub release² is executed for these connections, which means that the capacity they occupy on the span they traverse is made available. Connections that do not experience failures may not be re-routed during the recovery phase. This gives less freedom and therefore higher capacity usage than in a complete re-routing case, but makes the operation simpler [23]. According to the procedure described in section 6.2.4 the largest affected connection is identified and the restoration method (Ford-Fulkerson or Dijkstra) is called to find a restoration path. The upstream and the downstream directions of a connection are routed along

²Details on stub release can be found in chapter 5.

the same path, so only one direction needs to be calculated. If the connection is restorable, it is setup via the identified backup route, which means that capacity is occupied at the relevant spans. If the connection is not restorable, the appropriate spans are upgraded according to the description section 6.2.3. For one connection, several iterations are normally necessary in the Ford-Fulkerson method to identify the spans requiring upgrades, whereas all spans that need upgrading are found in one iteration with the Dijkstra method. It should be noted that the location of the upgrade capacity differs between the Ford-Fulkerson method and the Dijkstra method due to their different choice of restoration path. If there are more affected connections, they are restored after the same principle. If the current connection is the last one affected, the simulation of that particular span failure is finished. The starting point for the next failure simulation is the reference network with its original traffic connections in place, i.e., the network is reverted to its pre-failure state.

Adding Additional Spans

In the basic span upgrading strategies "x4 capacity" and "+own capacity" only existing links are upgraded and no additional spans are added once the restoration process has begun. The question is whether providing extra spans to the network results in the restoration possibility of more connections, which means that fewer upgrades are required, resulting in less overhead. When looking at the reference network shown in figure 6.4, it can be seen that six nodes only are connected via two links. This makes them very vulnerable to a span failure, because if one of the adjacent spans fails, only one span remains as a choice for the restoration path. Therefore, the nodes with the lowest connectivity (here, only two adjacent spans) are supplied with an extra span for restoration purposes. These spans are only available in the restoration phase, meaning that the working connections are routed in the original network to allow comparison with the previous simulations. The new topology is illustrated in figure 6.9, where the four additional links are shown as dashed lines³. As a starting point, the additional spans are dimensioned

³It should be noted that in a real network environment other factors such as geographical location, network policies and cost must be included into the decision process on where to add spans.

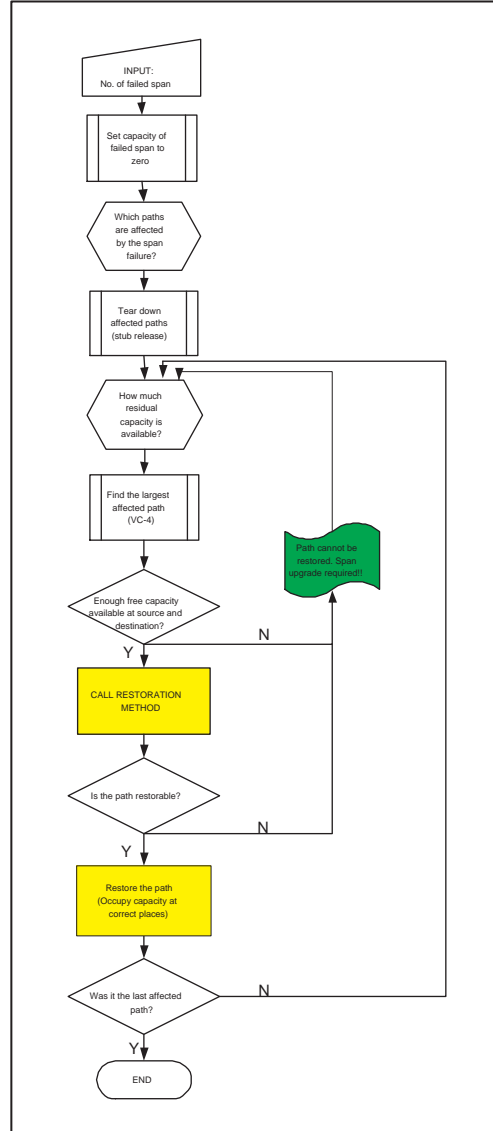


Figure 6.8: Restoration simulation procedure. Steps with yellow background are restoration method specific. Steps with green background are upgrading method specific.

to STM-16, which is the smallest value used in the original network. The restoration procedure is the same as described in section 6.2.4.




Figure 6.9: Reference network with additional links (dashed).

6.2.5 Simulation Results

After all upgrades are performed, all connections in the network can be restored after any single span failure. Three different parameters are investigated for all upgrading schemes: required capacity, average length of restoration path (hop count) and number of necessary upgrades. The total capacity for the protected network is shown in figure 6.10. The general tendency to observe is that the "x4 capacity upgrades" use more capacity for the Ford-Fulkerson, Dijkstra and adding of additional spans⁴, compared to their respective "+own capacity" methods. When comparing the Ford-Fulkerson and the Dijkstra method, the Dijkstra method uses less capacity; and least capacity is used when additional spans are added based on the Dijkstra method. This can be explained by looking at the number of span upgrades required by the different methods,

⁴The adding of additional spans is only executed for the Dijkstra method.

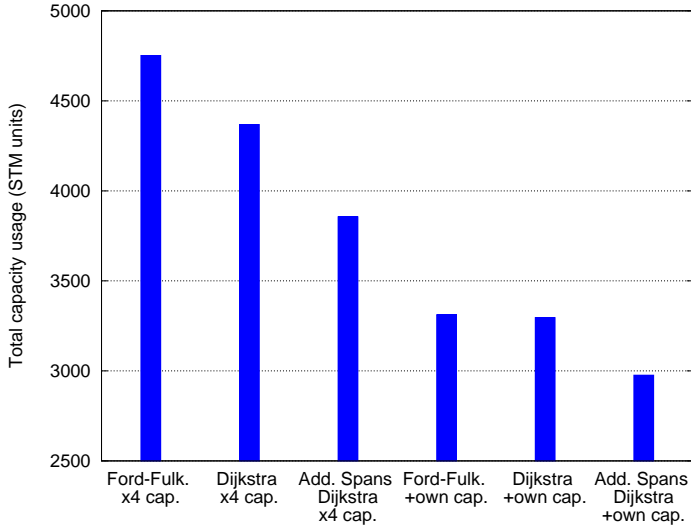


Figure 6.10: Total restoration capacity.




Figure 6.11: Number of necessary span upgrades.




Figure 6.12: Average restored connection hop count.

which is illustrated in figure 6.11. The "x4 capacity" methods require fewer upgrades than the "+own capacity" upgrading methods. This is explained by the granularity of the upgrades, since the "x4 capacity" method quadruples the capacity for each step, while the "+own capacity" method only doubles it. This means that some connections can be recovered by using the overhead that is created by other upgrades, which is higher for the "x4 capacity" method. Also due to its coarser granularity, fewer upgrades are required. In figure 6.12 the average hop count for the restoration paths is illustrated. The value for the Ford-Fulkerson method is higher compared to the Dijkstra method, since the Ford-Fulkerson method finds a flow augmenting path through the network, but this is not necessarily the shortest path. The high hop count is likely to result in a long restoration time [117]. The Dijkstra method finds the shortest path through the network and therefore has a lower hop count. The lowest hop count is obtained in the case where additional spans are added. This is because these spans allow for the possibility of taking "shortcuts" for the recovery paths. The hop count is also related to the capacity usage, since longer paths potentially result in high capacity usage unless sophisticated sharing schemes are used.

6.2.6 Conclusions on Networks with Modular Capacity

In this study, the total capacity usage, the number of span upgrades and the average hop count of the recovery paths are investigated. The restoration paths are chosen based on either flow augmenting paths (i.e., Ford-Fulkerson method), shortest paths (i.e., Dijkstra method), and also when additional spans are added to the network. Two different upgrading schemes are compared, based on whether the next larger size STM module was deployed (called "x4 capacity"), or if an additional wavelength is taken up (called "+own capacity"). The simulation results show that the Ford-Fulkerson method with the "x4 capacity" upgrades uses most capacity for restoration which makes it an expensive choice. Furthermore, the high hop count will result in a long restoration time, which contributes to its unsuitability. The methods using the "x4 capacity" upgrading scheme create a large amount of overhead per upgrade, which means that the total capacity increases significantly with only a few span upgrades. In some cases, the overhead capacity can be used for the restoration of other connections instead of having to upgrade an additional span, but this depends on the location of the extra capacity. The methods using the "+own capacity" upgrading scheme have a finer upgrading granularity, which means that less overhead is created per upgrade. But since a "+own capacity" upgrade creates less additional capacity to accommodate further connections, more upgrades are required.

The hop count for the Ford-Fulkerson method is higher compared to the Dijkstra method, since the Ford-Fulkerson method finds a flow augmenting path through the network, but this is not necessarily the shortest path. The Dijkstra method finds the shortest path through the network and therefore has a lower hop count. The alternative restoration approach, where additional links are added to the network and the necessary upgrades are made with the "x4 capacity" scheme, has resulted in a lower capacity usage compared to when no extra spans added. When the upgrades were made with the "+own capacity" scheme, the capacity usage was lower, but the number of necessary span upgrades was higher. In this simulation, the extra spans are assigned to the nodes with the lowest connectivity, but several other criteria could be used (e.g., nodes with lowest total capacity on adjacent spans, nodes that are often traversed, etc.). The placement of extra spans in the network is very important,

since the correct location of spare capacity is important for the success of network restoration. The geographical characteristics, network policies and economical issues must also be taken into account when deciding on the exact location of additional spans.

Based on the results obtained from the simulations, a restoration method based on the Dijkstra approach can be recommended over a restoration based on the Ford-Fulkerson method, since it performs better in capacity usage, number of required span upgrades and average hop in both the "x4 capacity" and "+own capacity" upgrading strategy respectively. A fine granularity in the upgrading steps results in a low total capacity usage but requires more span upgrades. The alternative restoration approach, where additional links are supplied to the network, has shown to have a potential for saving capacity and results in a lower hop count compared to when only existing links are upgraded. The results are however dependent on the traffic distribution and network topology, which should be varied to generalize the observed tendencies.

This study is mainly a graph theoretical experiment with iterative upgrading strategies. Off-line optimization SCA or JCA are able to significantly reduce the total capacity usage, especially if modularity is taken into account in the beginning of the design problem [107].

6.3 Shortcut Span Protection

6.3.1 Motivation

Designing resilient networks has been in focus of research and industry for some years now, resulting in a large variety of both well-known and experimental resilience provisioning mechanisms. Many of these studies focus on providing networks with the least amount of capacity possible to fulfil the required level of protection.

As mentioned in section 2.5 the local-to-egress [118] method, where the traffic is re-routed between the upstream failure adjacent node and the destination node, has a good performance trade-off in terms of notification time and capacity efficiency. The local-to-egress method was further studied in [62], where a LP to obtain optimal working and backup paths in mesh networks was provided. In that model, each affected connection was re-routed individually between the local failure adjacent node and the destination node.

However, there is also a tradeoff between capacity efficiency and the complexity of fully individual connection re-routing in the local-to-egress protection method. Furthermore, with the prices for fiber (i.e., capacity) dropping [119], capacity usage probably becomes a less important factor for deciding which protection method should be employed, compared to complexity, manageability and speed.

With the aim of reducing the complexity of the protection method without sacrificing restorability, a variation of the local-to-egress protection method called Shortcut Span Protection (SSP) is presented here. In SSP, the traffic is bundled between the failure adjacent node and the egress node if several affected connections have the same destination node. The method is illustrated in Figure 6.13. The advantage of the proposed method lies in a less complex route calculation and signaling process, since all connections going to the same destination node must follow the same route from the upstream failure adjacent node.

6.3.2 Simulation Study

The capacity efficiency of the SSP method depends on the routing of primary paths and the backup paths. Hence, to evaluate the efficiency the optimization problem must be solved: Route the primary and backup




Figure 6.13: Shortcut Span Protection.

paths, using as little network capacity as possible. To do this, a Linear Program is used to model the capacity usage and route the primary paths and backup paths optimally. The LP model is based on a graph representing the network such that the graph nodes correspond to network switches and the graph edges corresponds to the network cables. The objective is to find the total required capacity for the network. The following constraints are setup:

- Demand constraint: requires that for each demand enough capacity is assigned to the primary path.
- Backup constraint: ensures that if a span fails and a number of primary paths which use that span and which end at a given node, enough capacity is assigned on the path to that node.
- Capacity constraint: calculates the necessary capacity on each span for all failure situations.

In order to evaluate the effectiveness of SSP it is tested by optimizing over a set of 5 networks for a demand volume of 1 between all pairs of nodes in the network. In Table 6.1 the network properties are summarized. The contents of each column in the table is given as:

1. **Network:** Network name
2. **Nodes:** Number of nodes in the network
3. **Spans:** Number of spans in the network

Network	Nodes	Spans	Avg. N.D.	N.-F. Capacity
Cost239 [120]	11	26	4.73	86
PanEuropean	13	21	3.23	158
USANetwork [121]	28	45	3.21	1273
Italy [122]	33	68	4.12	1718
France [121]	43	71	3.30	3473

Table 6.1: Network properties of test networks.

4. **Avg. N.D.:** Average nodal degree in the network
5. **N.-F. Capacity:** Non-Failure network capacity, i.e., the summed capacity necessary for non-failure routing.

The following protection methods are evaluated:

- **Complete rerouting (CR):** allows the re-routing of non-failed connections as well [123]. Not practical for real networks, used for benchmarking only.
- **Path protection (PP):** connections are recovered on an end-to-end basis.
- **Local-to-egress protection (LtE):** recovery between the upstream failure adjacent node and the destination node.
- **Shortcut span protection (SSP):** as local-to-egress, but all protection paths for a given destination must follow the same route [23].
- **Span protection (SP):** the failure is restored between the failure adjacent nodes.

6.3.3 Results

This section presents the capacity usage of SSP in comparison to CR, PP, LtE and SP. In figure 6.14, the capacity usage is depicted as relative protection overbuild, defined as: necessary extra network capacity relative to the non-failure network capacity.

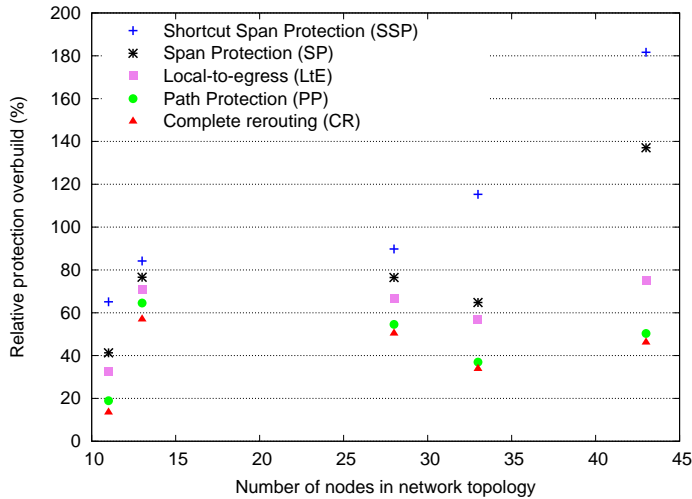


Figure 6.14: Relative protection overbuild.

The results show the following ordering of protection capacity, starting from the lowest usage: Complete re-routing \rightarrow path protection \rightarrow local-to-egress protection \rightarrow span protection \rightarrow SSP. This tendency shows that more freedom in terms of route choice results in a better capacity efficiency. SSP uses most resources. For local-to-egress protection the requirement of bundling results in a higher capacity usage as can be seen in SSP. The pattern of the individual curves follow each other, except in the case of SSP in the Italian network, where the capacity usage increase of SSP is larger than in other network instances. This is likely due to the topological characteristics of the network.

6.3.4 Conclusions on Shortcut Span Protection

In this study, the capacity usage of SSP has been evaluated. In comparison to local-to-egress protection, which is the closest related protection method in terms of involved nodes, the capacity usage of SSP is higher. The increase in capacity can however be compensated by an easier control since all recovery requests are constrained to the same route to a particular destination. The quantification of the control plane relief depends on the specific protocol used and is left for further study.

6.4 Chapter Summary

This chapter deals with design of resilient optical networks. It describes an iterative capacity upgrading study for networks using modular capacity; and illustrates the effect of bundling recovery paths to the same route towards a destination.

In the iterative upgrading study, two restoration methods were simulated, one based on the Ford-Fulkerson maximum flow - minimum cut theorem and the other one based on the Dijkstra shortest path algorithm. Both methods were tested with two different span upgrading strategies: one where the capacity of each span was upgraded corresponding to upgrading an SDH module in a node; and one where the capacity of the span to be upgraded was increased corresponding to taking up an extra wavelength in an underlying WDM system. Additionally, the effect of supplying additional links to the network was investigated with the Dijkstra algorithm and both span upgrading strategies. Results have shown that using small upgrading modularities over short paths give the most capacity efficient result. These results can be generalized to any network technology that uses modular capacity.

The second study is motivated by the trade-off between capacity and control efficiency. It illustrates the effect of bundling recovery paths in such a way that the same route must be chosen for all affected connections between the recovery end points. The method investigated in this study is SSP, where all failed connections are bundled "per-destination" meaning that they must use the same route between the upstream failure adjacent node and the destination node. Results have shown that the bundling increases the capacity usage, but this may be compensated by simpler control of the protection requests.

Chapter 7

Conclusion and Outlook

Functioning communication networks are necessary for the growth and prosperity of our society. Unfortunately, these networks are regularly affected by failures. If these failures remain unresolved, they do not only cause loss of revenue, but may also lead to serious consequences for the people affected by the failure.

The work presented in this thesis aimed at increasing the survivability of optical networks. The main focus has been on restoring traffic in a dynamic network environment. Throughout this thesis, it has been highlighted that both sophisticated route selection and enhancing the signaling session have a positive influence on the restoration performance.

The presented studies show that limited wavelength conversion capability has a different effect on the span, local-to-egress and path restoration methods. The recovery percentage of span restoration is most severely affected, because it has to merge the restoration segment to the pre-recovery path at the failure adjacent nodes, which often requires a conversion between wavelengths. Path restoration is the least affected, because it restores affected connection between the connection's end nodes, and is therefore not bound by merging issues. The performance of local-to-egress restoration lies between the other two methods.

In a wavelength convertible network, it is crucial to use the few available Wavelength Converters (WCs) wisely. The Suggested Vector (SV) has proven to be beneficial for saving WCs during connection provisioning. This thesis proposes modifications to the SV aiming at increasing the recovery performance as well. The SV-based stub awareness

scheme avoids unnecessary wavelength conversions at failure adjacent nodes. This is valuable for span and segment restoration, where the wavelength of the bypass path must be merged to the wavelengths of the stubs at the failure adjacent nodes. The versatility of the SV is shown in the blocking reduction scheme, which can be used for traffic engineering purposes. The blocking reduction scheme both reduces the number of recovery retries and increase the recovery percentage. Both schemes adhere to the Generalized Multi-Protocol Label Switching (GMPLS) protocol suite and their complexity increase is minimal.

The SV only contains summarized information. It does not disclose details of a network domain and it can therefore be passed between domains in a multi-domain scenario without violating confidentiality requirements. Multi-domain survivability with the SV requires that all participating domains agree on how the information within the SV should be interpreted. The more information is exchanged between domains, the higher the probability for finding a route and a wavelength that can satisfy all participating domains.

It has been demonstrated that the restoration performance benefits from stub release actions. Full stub release is beneficial if capacity availability is the limitation of connection recovery. If WCs are the limiting factor, the benefit of stub release depends on the used wavelength assignment scheme. A simple wavelength assignment scheme, such as e.g., No Preference (NP), does not use WCs efficiently. Making them available for the recovery phase through stub release can therefore boost the recovery percentage significantly. An advanced scheme, such as e.g., the SV, ensures an efficient WC usage and is therefore less dependent on freed resources from stub release. It can therefore be concluded that in a wavelength convertible network, either an advanced wavelength assignment scheme or stub release can be used to achieve good restoration performance, depending on whether the additional complexity should be allocated in the signalling protocol or the recovery operation. Unfortunately, stub release complicates reversion to the primary path once the failure is repaired. This is due to the fact that the released portions of the path must be re-provisioned. With the nodal stub release method, it has been shown that if WCs are the limiting components, good recovery performance can be achieved by exclusively releasing the WCs during the stub release phase.

A network design study has illustrated that modularity of capacity units complicates efficient capacity assignment. Larger capacity modules result in more provisioned but unused capacity, leading to high network cost. The study has also revealed that short recovery paths and the addition of spans exclusively for the restoration phase decreases the overall capacity needed.

In terms of the granularity of the restoration paths, a trade-off between capacity efficiency and operational complexity has been identified. From a capacity point of view, it is beneficial to restore each connection individually, but individual connection restoration requires a large amount of route computation and signalling. Shortcut Span Protection, which bundles the restoration path of all affected connections going to the same node reduces the control plane load at the expense of a higher capacity usage. While previous studies focussed on minimizing capacity, Shortcut Span Protection (SSP) can be used by an operator who likes to minimize complexity and still have a reasonable capacity usage.

This work presented in this thesis has contributed to the progress of dynamic network survivability by analyzing existing solutions, extending recovery methods tailored to wavelength convertible networks and proposing new recovery methods. The results have been published in international journals and conferences. Gaining full benefits of the presented dynamic recovery methods however requires that network operators allow for automated network control processes to occur in their network. In the past, giving up control and allowing the network to make automated decisions has been dreaded by operators, but recently several operators have started to gradually introduce GMPLS-based dynamic processes into their networks. This tendency opens the possibility for commercial deployment of advanced restoration methods, with the long term goal of introducing fully dynamic multi-layer recovery. In the near future, pre-configured protection will most likely still be the preferred recovery method for single span failures in optical networks, at least until sufficient trust in fully dynamic network control is established. During this phase, network restoration can be used to recover from unanticipated failures, such as dual and multiple span failures, node failures or combinations thereof.

In addition to applying network restoration methods to telecommunication networks, both electric power companies and companies dealing with water supply have shown interest in the dynamic recovery methods proposed in this thesis. This shows that reliable services are not only important for communication networks, but also for other basic services relevant for our society. The possibility of extending and adapting the presented restoration methods for fault recovery in electrical power and water networks opens promising possibilities for combining telecommunication network recovery concepts with other industries for the benefit of our society as a whole.

Appendices

Appendix A

OPNET Model Details

A.1 Introduction

This appendix presents details of the OPNET simulator that was used in this thesis. The simulation framework is developed in OPNET Modeler [92] for evaluating the performance of Generalized Multi-Protocol Label Switching (GMPLS) signaling protocol extensions in dynamic Wavelength Routed Optical Networks (WRONs). The separation between the message handling and the functional operation on the protocol objects allows to easily adapt the simulator to the studied scenarios. Both network provisioning and recovery (span, local-to-egress and path) can be evaluated.

A.2 GMPLS Signaling OPNET Model

In the section, the development of the provisioning and recovery OPNET model is described in a top-down approach.

A.2.1 Simulation Parameters

The simulation parameters shown in table A.1 can be chosen by the user via the simulation configuration dialog box:

Parameter	Range
Wavelengths per span	1 - unlimited
WCs per node	0 - unlimited
Connection duration	Seconds (exponential distribution)
Mean inter-arrival time	Seconds (exponential distribution)
Provisioning label preference scheme	NP, LS, SL, SV
Recovery label preference scheme	NP, LS, SL, SV
Tie-breaking policy	Fist-fit, random
Span to fail	0 - number of spans in network
Recovery method	Span, Local-to-egress, Path
Stub release	Full and nodal
Recovery route calculation	Single try or k-shortest
Pre-processing	Match original span (stub awareness) or match stubs (blocking reduction)
Directionality of connection	Unidirectional or bidirectional

Table A.1: Simulation parameters for OPNET model.

A.2.2 Network Model

The network model shown here is based on the NSFNET topology [97], shown in figure A.1. The network consists of 14 nodes. Two additional nodes are shown on the figure: netinit and failure_control. The netinit node initializes all data structures and discovers the network topology by stepping through all the nodes in the network, thereby investigating which other nodes are connected in the other end of each span. Because the dynamics of Open Shortest Path First (OSPF) are outside




Figure A.1: Network model.

the scope of this work, this simple centralized model of the topology discovery is deployed instead of modelling distributed OSPF. The failure_control node takes care of initiating span failures and control which nodes should be invoked for proper recovery actions. More information on network recovery is given in section A.4.

A.2.3 Node Model




Figure A.2: Node model of network node.

The node model of a network node is shown in figure A.2. Seven transceivers are deployed in each network node, which should fit the nodal degree of most networks¹. Connections can originate and terminate at any network node with equal probability; they are initiated from the ReqGen processor. The processor named RSVP-TE takes care of

¹More transceivers can be added easily, since the model discovers the number of them automatically at simulation start.

handling RSVP-TE messages, which is described in the following sections.

A.3 Connection Provisioning

A.3.1 Request Generation

Each network node has a request generator (ReqGen in figure A.2), which initiates the setup of connections. The process model of the request generator is illustrated in figure A.3. In the Init state, simulation attributes entered by the user are collected. In the Idle state, the next request is scheduled as a self-interrupt based on the connection inter-arrival time. When this interrupt arrives, the Gen_req state schedules a remote interrupt to the RSVP-TE module in the node, which initiates connection setup and schedules another self-interrupt.

A.3.2 Handling RSVP-TE Messages

When an interrupt is received by the RSVP-TE module, the interrupt is processed by the process model illustrated in figure A.4. This is a root process that controls interrupts destined for child processes (explained below) that take care of handling the individual connections. A REQ_ARR_INTRPT means that a new connection request arrives and the Req_handle state finds a route based on the Dijkstra shortest path algorithm and creates a new child process for the specific connection.




Figure A.3: Proces model of request generator (ReqGen).




Figure A.4: Process model of RSVP-TE (root).

The child process then takes care of the exchange of the proper RSVP-TE messages. When such an RSVP-TE message arrives (`MSG_ARR_INTRPT`), the `Pk_handle` state identifies the type of packet (`Path`, `Resv`, `PathErr`, `ResvErr`, `PathTear`) and for which connection (i.e., child process) it is intended. When the connection expires after the specified connection duration (`EXP_NOTIFY`), it is removed in the `Rem_conn` state. The remaining three states (`RecoverySource`, `RecoveryDest` and `Failure`) are related to the handling of failures and are used for notification purposes.

The child processes take care of handling RSVP messages for the individual connections. Depending on the type of message that arrives at the child process, different actions need to be taken. The complete state-transition diagram of the RSVP-TE child process has been replaced by the simplified figure A.5 for convenience.

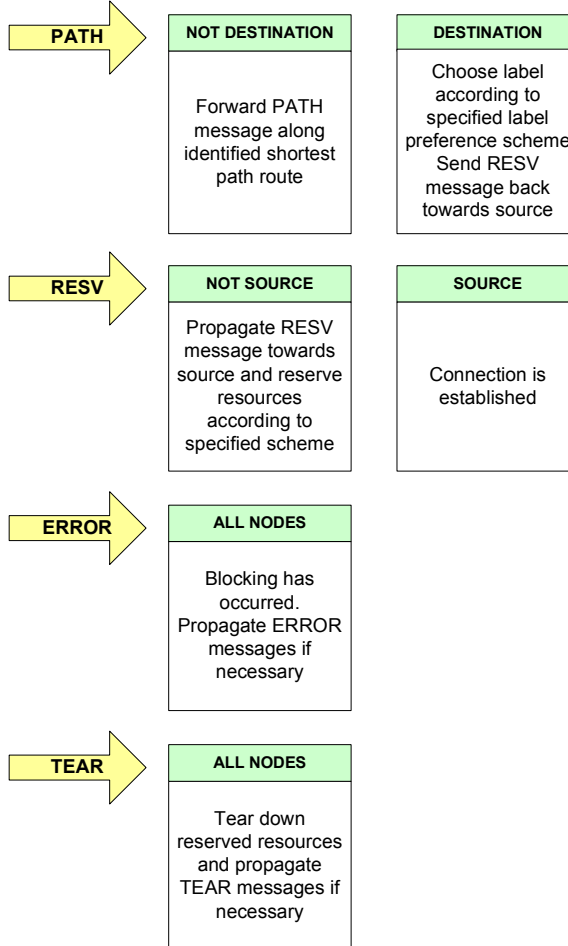
Message Type:

Figure A.5: Simplified process model of RSVP-TE (child).

A.4 Connection Recovery

Each fault recovery simulation run covers network initialization and one failure and terminates when all affected connections are either restored or found to be non-recoverable. Before any failure occurs, the network is

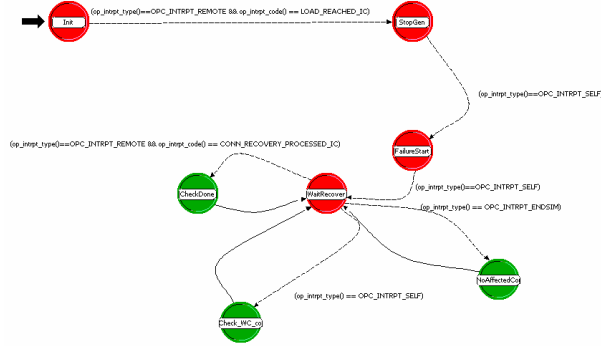


Figure A.6: Process model of failure_control.

populated with connections, until the average network load has reached a user-specified value. The average network load is calculated as in equation A.1, where S denotes the number of spans in the network and W specifies how many wavelengths are available per span. The 2 is introduced because the wavelengths are available in both directions.

$$AverageNetworkLoad = \frac{Activewavelengthchannels}{2 \cdot S \cdot W} \quad (A.1)$$

In contrary to the provisioning study, no connections are torn down once they are established. This is due to the fact, that the goal is to investigate the performance of the different label assignment schemes in the recovery phase under equal conditions (i.e., routes of provisioned connections and residual capacity on the different spans). Once the desired average network load is reached, the failure phase begins, where connection provisioning is stopped and a span failure is simulated on a user-selectable link. The failure_control node (shown on the network overview in figure A.1) initiates the chosen failures. The process model of the failure_control is shown in figure A.6. The connection generators are stopped when the failure_control receives an interrupt signaling that the desired load has been reached. After a settling period, allowing ongoing connection provisioning in the network to finish, a stable network state is reached and the failure phase can commence.

A.4.1 Failure Notification

In the FailureStart state, the nodes attached to the failed span are notified (with a remote interrupt to the relevant nodes' RSVP-TE module's root process shown in figure A.4) and the state machine goes to the Wait-Recover state, where it is notified of all successful and permanently failed recovery attempts. Failures are detected by the adjacent nodes through Link Management Protocol (LMP) signaling [45,55], but in this model explicit LMP signaling has been replaced by direct notification to the failure adjacent nodes. This simplification is justified, since quantifying the signaling delay during recovery is not the goal of this simulation.

A.4.2 Recovery Operation

Similar to the connection provisioning, the child processes take care of exchanging the appropriate RSVP-TE messages from the failure notification until either a successful recovery or a permanent failure (i.e., connection is non-recoverable) is reached.

A flowchart of the recovery operation performed in the upstream failure adjacent child process is shown in figure A.7. Optionally, stub release is performed first. After a settling period, a shortest path is calculated. If no route can be found, permanent failure is reported. Otherwise, route reservation is initiated. If successful, the recovery is successful. If the reservation fails, the causing span is removed [96] from the Traffic Engineering Database and the procedure retried from the path calculation. When all affected connections have finished and reported in, failure_control ends the simulation and saves the collected data.

A.5 Result Presentation

The developed simulation model allows extracting a large variety of result metrics, which are described below. Furthermore, an explanation of the confidence interval calculation is given.

A.5.1 Metrics

The following result metrics can be extracted from the simulation results:

- Blocking probability

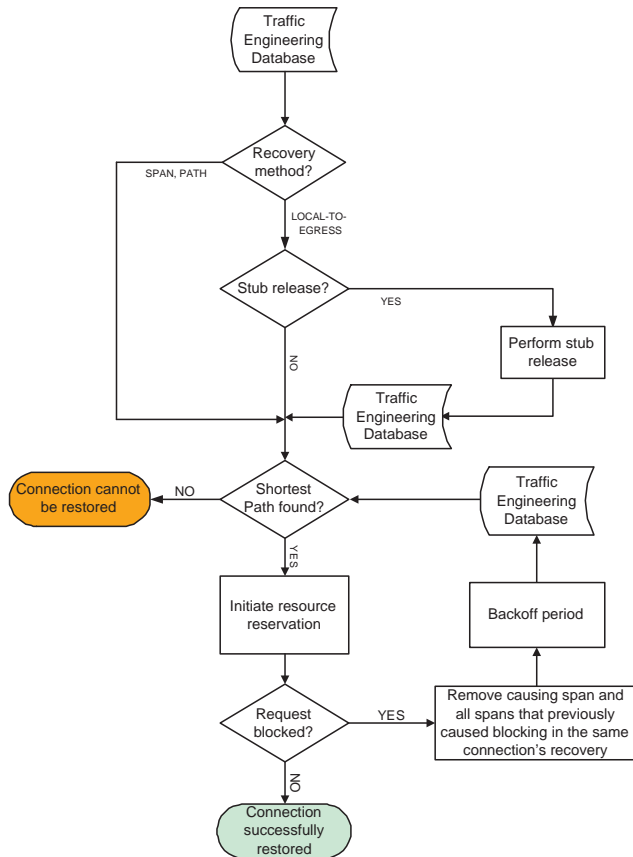


Figure A.7: Flowchart for connection recovery.

- Wavelength converter usage during provisioning or recovery
- Number of provisioned, failure affected, recovered or permanently failed connections
- Routing blocking in provisioning or recovery phase
- Forward blocking in provisioning or recovery phase
- Backward blocking in provisioning or recovery phase
- Number of recovery retries

- Provisioning or recovery hop count
- Provisioning or recovery wavelength usage
- Wavelength converters released during stub release
- Wavelengths released during stub release

A.5.2 Confidence Intervals

OPNET gives the possibility to use several random seeds for evaluating each data point. By using confidence intervals, an interval likely to include a certain value is given in addition to the mean value obtained by averaging the results of all runs with different random seeds. The likelihood that a certain value is included in the confidence interval is specified by the confidence level. 95% confidence level is a typical value.

Confidence interval calculation, based on [124]: We have n results X_1, X_2, \dots, X_n for different random seeds. The mean value is calculated in equation A.2. The variance of the sample is calculated in equation A.3.

$$\tilde{X} = \frac{1}{n} \cdot \sum_{i=1}^n X_i \quad (\text{A.2})$$

$$s^2 = \frac{1}{n-1} \left(\sum_{i=1}^n X_i^2 - n \cdot \tilde{X}^2 \right) \quad (\text{A.3})$$

The confidence interval of the mean value is found in equation A.4, where $t_{n-1, 1-\alpha/2}$ is the upper $(1 - \alpha/2)$ percentile of the t-distribution with $n-1$ (degrees of freedom) [124]. The probability that the confidence interval includes another mean value is equal to $(1 - \alpha)$. This is called the level of confidence. Consequently, confidence intervals at 95% confidence level require an α -value of 0.05.

$$\tilde{X} \pm t_{n-1, 1-\alpha/2} \cdot \sqrt{\frac{s^2}{n}} \quad (\text{A.4})$$

A.6 Appendix Summary

This appendix presented details on the simulation framework used in this thesis to evaluate GMPLS standard and novel signaling extensions. First, the simulator for provisioning studies was presented. The versatility of the framework is demonstrated by the straightforward extension of the simulator from connection provisioning to also cover span, local-to-egress and path restoration. The capacity of the spans (number of labels), the number of Wavelength Converters (WCs), the desired average network load, the label assignment scheme, the tie-breaking policy, the restoration method, etc., can be conveniently chosen by the user through the simulation configuration dialog box, which makes the simulation framework both a convenient and versatile tool to evaluate optical network provisioning and restoration.

Appendix B

Connection Provisioning

B.1 Introduction

This section presents some concepts used in connection provisioning, that I have been involved in during my thesis work. Since they are not concerned with network survivability as such, they are only presented as summaries. Section B.2 presents different label preference schemes. Section B.3 shows evaluates strategies for Wavelength Converter (WC) placement in a dynamic network environment. In section B.4 the challenges of bidirectional lightpath provisioning are discussed. Section B.5 benchmarks dynamic Routing and Wavelength Assignment (RWA) compared to off-line optimization. Section B.6 summarizes the appendix.

B.2 Label Preference Schemes in GMPLS Controlled Networks

B.2.1 Introduction

Generalized Multi-Protocol Label Switching (GMPLS) [27] is emerging as a unified framework to provide next generation network architectures, characterized by multiple switching layers, with a common control plane. In this section, different wavelength preference schemes are presented, which have been compared and the results have been published in [14,21].

B.2.2 Connection Setup in GMPLS

In this section the basic mechanism for connection (i.e., Label Switched Path (LSP)) setup in GMPLS networks is described exploiting Resource ReserVation Protocol (RSVP)-Traffic Engineering (TE) notation. After a unidirectional¹ LSP setup request (either by the network management system, or by a client at the source node), a route from source to destination is either retrieved from a static database, or dynamically computed using the information gathered by the routing protocol. In a typical optical network scenario, the routing protocol is not aware of the occupation of each wavelength, but it advertises the availability of a network link as long as free capacity is available.

Then the signaling session is triggered along the found route. A *Path* message is sent from source to destination (i.e., in the so-called downstream direction), carrying the necessary information for LSP setup. It contains a Generalized Label Request object, which every node issues to the downstream neighbor to notify the request of a label binding. It may include optional objects, such as the Label Set and the Suggested Label, which have been standardized [87], and the Suggested Vector, which is presented in this section. Intermediate nodes along the route process the *Path* message, creating a proper state for the session in a local database to store the received objects, and forward it to the next hop.

When the *Path* message eventually reaches the destination node, it responds (in the upstream direction) with a *Resv* message carrying a Generalized Label object, where every node indicates to the upstream neighbor the label assigned to the incoming request. When the *Resv* message finally reaches the source node, the LSP is established.

Besides *Path* and *Resv*, other signaling messages are defined. Among them we cite *PathErr/ResvErr*, used to notify error conditions (e.g., LSP setup failure) to upstream and downstream nodes respectively, and *PathTear/ResvTear*, used to delete the LSP and the relative session state in all traversed nodes.

After a general overview of LSP setup in GMPLS, the following sections first define the tie-breaking policies enforced at network nodes to select a label if two or more have equal preference. Then the signaling

¹Issues and strategies related to bidirectional LSP setup are discussed in section B.4.

protocol objects are detailed, i.e., the standardized Label Set and Suggested Label, and the proposed Suggested Vector, used by each node to specify the preference for each available label.

B.2.3 Tie-Breaking Policies

In the context of GMPLS-controlled optical networks tie-breaking policies are similar to classical wavelength assignment policies. In particular, tie-breaking policies specify the behavior of network nodes when they need to choose a label within a pool of labels with equal ranking. This is a common situation when the aforementioned preference mechanisms are enforced: for instance, if more than one label can be used to set up a lightpath in an optical network without wavelength converters, then a choice must be performed to select one of them.

It is important to notice that tie-breaking policies are a local matter of each network node, and no standard behavior is needed. They are therefore not part of the protocol specification, instead they are left to the implementation.

Two common policies are used here:

- **First-fit:** with the First-fit policy, the tie is broken by choosing the lowest-index label within the pool.
- **Random:** with the Random policy, the tie is broken by choosing a random label within the pool.

In this network setup, advanced policies, such as Most-used and Least-used [91], selecting the label used on the greatest and least number of links respectively, cannot be used because nodes are not aware of the detailed wavelength occupation state of all network spans.

B.2.4 Wavelength/Label Assignment Schemes

No Preference

This simple scheme does not use any protocol object to specify label preferences. When a request reaches the destination, a wavelength free on the last hop is reserved and propagated as far as possible upstream. If along the reverse path a node cannot further use the current wavelength because it is busy on its previous hop, and a converter is available, a new

wavelength free on the previous hop is chosen, and so on; otherwise the connection is blocked.

Label Set

The Label Set is carried in the *Path* message and is used by an upstream node to control the selection of labels by downstream nodes. The Label Set is an array of labels, containing only the labels acceptable by the upstream node. Therefore the downstream node is forced to select a label from this set. The Label Set is an optional object: if it is not present, all available labels on the previous hop may be chosen.

The Label Set object is particularly useful in optical networks without wavelength converters. The source node indicates in the Label Set the available wavelengths on the first hop. Along the path each intermediate node considers the wavelengths available on its next hop, and potentially decreases the set of usable labels by modifying the Label Set. When the request eventually reaches the destination node, the Label Set contains only the wavelengths available on the entire end-to-end path; therefore the destination node can safely choose (with a given tie-breaking policy) one of them and send back the *Resv* message with the proper Generalized Label object. During its reverse path, the *Resv* message reserves the selected wavelength along each hop.

However, in networks with wavelength converters the benefits of Label Set object are not so relevant. This is explained in figure B.1 (for the moment focus on the Label Set only), where a lightpath from A to D is being established. The Label Set traverses the path and, when it reaches destination node D, it contains both λ_1 and λ_3 . We notice that λ_1 guarantees an end-to-end wavelength-continuous path, while with λ_3 a wavelength conversion is required in node B. It is clear that, using just the Label Set, the destination node lacks information to choose the wavelength that minimizes the number of wavelength conversions. To solve this issue we propose to support the Label Set with the Suggested Label object.

Suggested Label

The Suggested Label is an optional object carried in the *Path* message. It is used to provide a downstream node with the upstream node's la-

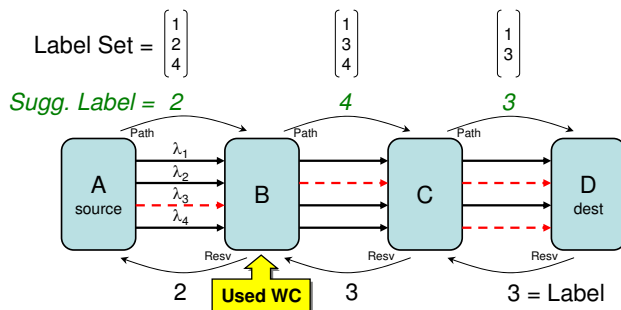


Figure B.1: Example of non-optimal utilization of the Suggested Label in a network with wavelength converters. Busy wavelengths are dashed.

bel preference. As stated in [87], the main purpose of this extension is to reduce the setup time when establishing or recovering LSPs through optical switching fabrics characterized by a non-negligible configuration time, because the upstream node can start configuring its hardware before receiving the actual label from the downstream neighbor. However this mechanism is just an optional optimization: the downstream node still maintains the full control of the label choice. If the downstream node does not accept the suggestion and passes a different label to the upstream node, the latter has to reconfigure itself accordingly.

A novel use of the Suggested Label is proposed here, aiming at reducing the amount of wavelength converters utilized to set up a lightpath. For this purpose an algorithm for computing the Suggested Label for the next hop has been designed. It is presented below, where *ph* and *nh* represent previous hop and next hop, respectively.

```

if (thisnode == source)
  nh_sugglab := one_of(nh_labelset);
else
  if (nh_avail(ph_sugglab) == TRUE)
    nh_sugglab := ph_sugglab;
  else
    if ((ph_avail_set && nh_avail_set) != EMPTY)
      nh_sugglab :=
        one_of(ph_avail_set && nh_avail_set);
    else
      nh_sugglab := one_of(nh_labelset);

```


The source node is oblivious of the optimal wavelength to minimize wavelength conversion, so it simply chooses one (with a given tie-breaking policy) within the Label Set for the next hop. An intermediate node first checks if the suggested wavelength is available on the next hop. If it is free, it continues suggesting the same label. If it is busy, the intermediate node computes the set of wavelengths available on both its previous and next hops, not requiring conversion at the node. If the set is not empty, the node chooses one of them (with a tie-breaking policy); if the set is empty, meaning that in the node a wavelength conversion is needed in any case, a wavelength is chosen from the Label Set for the next hop (using a tie-breaking policy).

When the *Path* message reaches the destination node, a reservation is attempted on the suggested label, backtracking to the source node. If a wavelength conversion becomes necessary because the current label is busy on the upstream hop, the suggested label for that hop is tested first: if it has been used by another connection in the meantime, a new label is chosen among the remaining ones of the Label Set.

However, the constraint that only a single wavelength can be notified with the Suggested Label object, and the consequent choices that must be performed to break ties, make this approach sub-optimal, especially for long paths. The shortcoming is illustrated in figure B.1. Source node A, adopting the Random tie-breaking policy, suggests λ_2 , but node B realizes that it is busy on its next hop. Node B thus computes the set of wavelengths available on both its previous and next hops (i.e., λ_1 and λ_4): if it chooses λ_1 , this suggested label is propagated up to destination, and the connection is smoothly setup on λ_1 . On the contrary, if node B randomly chooses to suggest λ_4 , node C realizes that λ_4 is not available on the next hop, thus computes the set of wavelengths available on both its previous and next hops (i.e., λ_1 and λ_3), and it may suggest λ_3 . Destination node D then follows the suggestion and reserves λ_3 on its previous hop, node C continues on λ_3 , following the rule of not changing wavelength unless necessary, but eventually node B is forced to use a wavelength converter. To overcome this issue, a new protocol object, namely the Suggested Vector, is proposed and described in the following section.

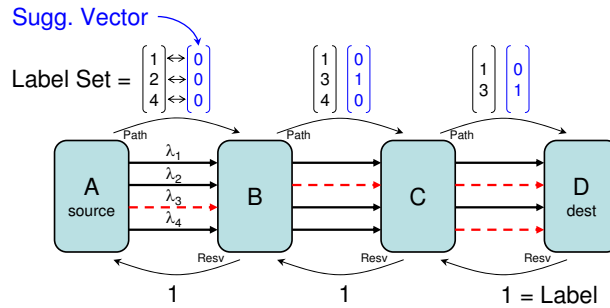


Figure B.2: Example of utilization of the Suggested Vector.

Suggested Vector

The Suggested Vector is proposed as an optional object, carried in the *Path* message and aimed at specifying the preference level for each wavelength. It is meant to be used together with the Label Set object: they have the same structure, being arrays with equal number of elements. For each label present in the Label Set, the Suggested Vector contains the number of wavelength conversions needed to use the specific label on the next hop. An algorithm for computing the Suggested Vector for the next hop has been defined and it is shown below.

```

if (thisnode == source)
  for (each elem in nh_labelset)
    nh_suggvec[elem] := 0;
else
  for (each elem in nh_labelset)
    if (ph_avail(elem) == TRUE)
      nh_suggvec[elem] := ph_suggvec[elem];
    else
      nh_suggvec[elem] := min(ph_suggvec) + 1;

```

The source node can reach every wavelength within the Label Set without wavelength conversion, hence it fills the Suggested Vector with zeros. Intermediate nodes check every wavelength within the Label Set for the next hop to determine if it is available also on the previous hop. If the condition is true, it means that the specific wavelength can traverse

the node without wavelength conversion, so the number of necessary conversions remains the same as on the previous hop. On the contrary, if the label is busy on the previous hop, it means that a wavelength conversion is necessary to use the label on the next hop. In this case the value specified in the Suggested Vector is the minimum amount of wavelength conversions needed to use a label on the previous hop incremented by one².

When the *Path* message reaches the destination node, a reservation is tried on the label with minimum Suggested Vector value (if more than one is available, a tie-breaking policy is used), backtracking to the source node. When a wavelength conversion becomes necessary because the current label is busy on the upstream hop, labels within the previous hop Label Set are tried for increasing values of the previous hop Suggested Vector, thus guaranteeing that the amount of used converters is always minimized.

Figure B.2 exemplifies the concept, referring to the same situation of figure B.1. Node A adds a zero in the Suggested Vector for each element of the Label Set. Node B keeps the values for λ_1 and λ_4 unchanged because they are free on its previous hop, but assigns 1 to λ_3 , since it is busy on its previous hop and the minimum value of the Suggested Vector on its previous hop is zero. Node C simply keeps the values for λ_1 and λ_3 unchanged. Thanks to the information carried in the Suggested Vector, the destination node D is fully aware of the amount of converters needed for each wavelength present in the received Label Set, and can choose λ_1 , which minimizes this value.

B.2.5 Summary of Concept for Label Preference Schemes

The four label preference schemes and the two tie-breaking policies have been evaluated by means of event-driven simulations using OPNET Modeler [92]. The results have been published in [21] and [14], where more details on the simulation study can be found as well. The simulations results show that both Suggested Label (SL) and Suggested Vector (SV)

²For simplicity, we suppose to use full-range wavelength converters, i.e., capable of converting any wavelength to any wavelength of the Wavelength Division Multiplexing (WDM) comb. A limitation on the conversion range would only reduce the set of previous hop labels on which to compute the minimum value of the Suggested Vector.

schemes significantly reduce the wavelength conversion utilization compared to currently exploited No Preference (NP) and Label Set (LS) schemes. When the number of converters is limited, SL and SV also reduce the blocking. For a simple scheme, such as NP, it is beneficial to spread the wavelength usage by using random tie-breaking whereas SV with the First-fit policy outperforms all the other.

B.3 Wavelength Converter Placement For Dynamic Traffic

B.3.1 Introduction

The previous section has shown that adding wavelength preference information to the *Path* messages can significantly reduce the WC usage and hence improve the blocking probability at the expense of a few more bits in the signaling messages. The Suggested Vector object has been proposed in [14, 21] for this purpose: it includes complete information on the number of required WCs for a specific wavelength selection on all traversed links along the chosen path. Using the above approaches it is possible to reduce the number of expensive wavelength converters throughout the network. However, it is not clear whether the wavelength converters should be distributed uniformly or based on specific optimization schemes. Evidently, a uniform distribution of the WC resources would ease the network management in case of node addition and thus increase the scalability. On the contrary, advanced optimization techniques may further reduce the wavelength converter count without degrading the blocking probability. However, the drawback of optimization techniques is the reliance on complete knowledge of a static traffic distribution. This requirement conflicts with the dynamic nature of on-demand traffic requests in next-generation optical networks. In this section, WC placement heuristics are presented, which were evaluated through modelling under dynamic traffic assumptions.

B.3.2 Converter Placement Heuristics

Ever since WCs were first demonstrated, one of the main questions has been where in the network they should be placed. Several studies [36, 125–127] have used mathematical optimization methods to achieve an optimal converter placement based on a static traffic matrix. This problem however was found to be difficult to solve, and therefore several heuristics have been proposed to achieve near-optimal performance [128, 129]. Since WCs are expensive components they are generally placed in such a way that the traffic demands can be satisfied with the least number of converters. The optimal converter placement is therefore highly dependent on the characteristics of the traffic demands. Under dynamic traffic

assumptions the traffic distribution will change over time and therefore the once optimal converter location may now cause bottlenecks at some nodes, whereas WCs are abundant at others. If converters are severely limited, only small fluctuations in the traffic may result in blocking of requests due to WC unavailability. To fulfill the changed traffic demands, it may be necessary to re-deploy the converter pool, or to upgrade the number of converters available in the network. Both scenarios cause significant economic drawbacks. It is therefore questionable whether specific traffic matrices should be considered during WC placement, or if more general rules are better to accommodate variations in the traffic locations caused by shifted traffic demands. Here, the goal is not wavelength converter deployment minimization. Instead, a case study is conducted where a network operator has deployed a given number of wavelength converters based on the well-known results that WCs can reduce the blocking probability, as demonstrated in e.g., [14,21].

The converters should be located in the network so that good blocking performance is achieved and that the placement is future-proof in terms of scalability and varying traffic demands. Three different converter placement heuristics are presented. They are developed based on different reasoning in terms of simplicity, network characteristics and demand generation. A general feature of all methods is the independence of specific traffic demand matrices, which results in high scalability and simple management of the converter deployment. The heuristics are based on the following features:

- Uniform distribution
- Nodal degree proportional
- Geographically dependent

Uniform Converter Distribution

The uniform distribution is the simplest heuristic. As the name implies, each node in the network is assigned the same number of WCs, without considering further measures such as traffic demands or network topology. The strength of this scheme is its simplicity and its flexibility, since it is completely independent of the traffic distribution in the network when the converters are provisioned. The total number of WCs in a




Figure B.3: Uniform converter assignment.

network using uniform placement, $C_{tot,u}$, is calculated in equation B.1, where N is the number of nodes in the network and C_u is the number of converters in each node.

$$C_{tot,u} = N \cdot C_u \quad (\text{B.1})$$

An example of uniform converter assignment is shown in figure B.3, where all nodes contain the same number of WCs.

Nodal Degree Proportional Converter Placement

The nodal degree specifies the number of spans that enter/leave a given node. The motivation behind placing converters related to the nodal degree is that nodes with many connecting spans often carry a large amount of traffic. In [125], it has been shown that nodes with a high nodal degree are more likely to require conversion capability, which is consistent with the approach of nodal degree proportional converter placement. The strength of this scheme is that it reflects the traffic distribution on a general level (i.e., more spans into a node gives potentially higher demand for conversion) while it is independent of specific traffic matrices. Furthermore, the scheme is still relatively simple to manage, as specific topological characteristics decide on the converter placement. To map the total number of available WCs to the individual nodes equation B.2 is applied.




Figure B.4: Converter assignment based on nodal degree.

In equation B.2, the number of converters in a node i with nodal degree proportional placement, $C_n(i)$ is calculated, where $n_d(i)$ is the nodal degree of node i and α is a proportionality constant. α is real-valued and can be fitted to control the total number of converters in the network. The rounding to nearest integer in equation B.2 is necessary since WCs are deployed in whole units. The total number of converters in the network using nodal degree proportional placement, $C_{tot,n}$, is obtained by summing $C_n(i)$ in equation B.2 over all nodes.

$$C_n(i) = \text{round}(n_d(i) \cdot \alpha) \quad (\text{B.2})$$

Figure B.4 shows how converters are distributed based on the nodal degree of the nodes. In the figure, a proportionality constant of 1.0 is applied, i.e., the converter count is equal to the nodal degree.

Geographically Dependent Converter Placement

This approach is motivated by the geographical placement of the individual nodes in a network. Its basic idea is to divide the nodes into two groups since networks are generally composed of a group of core nodes and a group of remote edge nodes. If two edge nodes from opposite sides of the network want to communicate they habitually must pass through the core of the network. Hence nodes located in the center of the network carry great amounts of pass-through traffic in addition to




Figure B.5: Converter assignment based on geographical location.

the locally generated traffic, which increases the conversion probability. The geographical placement is again independent of specific traffic demands, but it acknowledges the fact that much traffic is carried between core nodes. One of the challenges in this placement heuristic is making the actual division of whether a node belongs to the edge or the core group. A current traffic distribution could be used to break ties in case of uncertainties. The total number of converters using geographically dependent heuristic, $C_{tot,g}$, is calculated in equation B.3, where N_c and N_e are the number of core and edge nodes, and C_c and C_e are the number of converters assigned to each core and edge node, respectively.

$$C_{tot,g} = N_c \cdot C_c + N_e \cdot C_e \quad (\text{B.3})$$

Figure B.5 illustrates the concept of geographical converter placement. In this example, a 2:1 ratio of the core:edge converter number has been chosen, but it should be noted that any other ratio can be used. Furthermore, the concept can be extended from a core and edge division to partitioning the network into multiple zones of different converter numbers.

B.3.3 Summary of Concept for Converter Placement Heuristics

This section presented three converter placement heuristics (uniform, nodal degree and geographical). They have been evaluated for the NP, LS, SL and SV, and the results of the study are published in [2]. The results show that WC placement using nodal degree proportional or geographically dependent heuristics perform only marginally better than simple uniform distribution. This conclusion contradicts previous studies, which indicated that WC placement optimization has a large impact on network performance [36, 125–128]. The presented results, however, differ from these previous works on two major assumptions: 1) Dynamic (random) traffic is used and 2) the intelligence of the label preference schemes. Due to the dynamic traffic assumption, the results are valid for the very useful case of provisioning with unknown or uncertain traffic matrices. Furthermore, the label preference schemes may also contribute to the WC placement indifference. In summary, it has been shown that uniform WC placement can match the performance of an advanced WC placement heuristic in optical networks under dynamic traffic assumptions using advanced label preference schemes.

B.4 Bidirectional Lightpath Provisioning

B.4.1 Introduction

In a transport network, it is often desired that connections are bidirectional, with the same data transfer capabilities in each direction. For control simplicity, another common requirement is that both directions share the same spans to provide fate sharing. Moreover in Wavelength-Routed Optical Networks (WRONs) without wavelength conversion capability, both directions have typically to be routed on the same wavelength due to physical constraints of optical devices and aimed at simplifying flexible device operation [130]. In GMPLS controlled networks, these requirements can be satisfied by establishing a pair of unidirectional connections (or LSP) with two separate signaling sessions.

A smarter solution requiring a single signaling session to set up both directions of a bidirectional LSP has been standardized in the GMPLS signaling [87, 131]. This scheme reduces the setup latency, the number of exchanged control plane messages, and the memory requirements in the traversed nodes, since only one control plane state has to be stored. To setup a bidirectional LSP, a forward reservation of the resources belonging to the reverse path is performed by exploiting the Upstream Label object, while the forward path is reserved with the usual backward reservation. However, if the status of the wavelengths on the traversed spans is not known in detail, the source node may select a wavelength not available on the whole path, causing the blocking of the setup attempt.

Recently, an alternative scheme based on the Label Set object has been proposed, where both lightpath directions are established on the same wavelength with the backward reservation after collecting information on the wavelength utilization of both directions [130].

In this section, the Upstream Label and the Label Set schemes are presented together with three typical routing scenarios, where the routing protocol advertises more or less detailed network status information [132]. The results of the study are published in [3].

B.4.2 Bidirectional GMPLS Lightpath Setup

In this section the mechanism for LSP setup in a GMPLS-controlled WRON is described. Upon an LSP setup request, a route from source

to destination is computed by using the information advertised by the routing protocol. The three considered routing scenarios are detailed in section B.4.2. After the path computation, a signaling session is triggered to establish the lightpath: the two signaling schemes, namely the Upstream Label scheme and the Label Set scheme, are detailed in section B.4.2.

Routing Scenarios and Path Computation

Three routing scenarios are envisioned, where an increasing amount of network status information is advertised. In the No Information (NI) scenario, Open Shortest Path First (OSPF)-TE is used to advertise the network topology only. In the Aggregated Information (AI) scenario, OSPF-TE is exploited to flood the number of wavelength channels available on each link. In the Detailed Information (DI) scenario, an extended version of OSPF-TE [133] is implemented to advertise the status of each wavelength on every span. In all the considered routing scenarios, the path between each node pair (s, d) is selected at source within a set of candidate paths $P_{s,d}$. In the NI scenario, a path within $P_{s,d}$ is randomly selected. In the AI scenario, the path with the largest number of available wavelengths on its most congested span is selected, by using a link-state database containing the number of available wavelengths on every span. Finally, in the DI scenario, the path that can accommodate the largest number of wavelength-continuous lightpaths is selected, using a link-state database containing the status (i.e., busy or available) of each wavelength on every span. In AI and DI scenarios, if more than one path satisfies the condition, one of them is randomly selected. In these scenarios a lightpath request may be blocked during path selection if the source node is unable to find a path with available wavelength channels. In all the considered routing scenarios the information is flooded by Link State Advertisementss (LSAs). To limit LSA generation in AI and DI scenarios an LSA update timeout is used. Once an LSA has been generated for a given link, all link-state changes detected on the link before the timeout are not immediately advertised, but delayed after the timeout expiration. OSPF-TE minimum LSA update timeout is 5 s [54].

Bidirectional Wavelength Assignment

After the path selection, a signaling session is triggered, which may be blocked during *Path* or *Resv* message propagation. The two signaling schemes under consideration are detailed in the following

Upstream Label (UL) scheme:

This scheme adheres to [87]. A *Path* message is sent from source to destination, containing an Upstream Label, which reserves a wavelength on the reverse path. It may include optional objects, such as the Label Set. However, since both LSP directions must be routed on the same wavelength, the usefulness of the Label Set is severely limited as the only acceptable label is the one indicated in the Upstream Label³ Assuming that no error occurs, the reverse path is completely set up when the *Path* message reaches the destination. A *Resv* message is sent in the upstream direction to reserve the forward path, carrying a label identical to the one used on the reverse path. When the *Resv* message finally reaches the source node, the LSP is established.

Label Set (LS) scheme:

This scheme has been proposed in [130] to avoid the drawbacks of the forward reservation performed with the Upstream Label. The request for a bidirectional LSP is indicated with a flag in the LSP_ATTRIBUTES object [134] carried by the *Path* message. The Label Set is mandatory and is updated at each intermediate node by checking the availability in both directions jointly. When the LSP setup request arrives at the destination node, it selects an available wavelength within the received Label Set and starts the backward reservation of both LSP directions on the chosen wavelength.

B.4.3 Summary of Concept on Bidirectional Lightpath Provisioning

In this section the Upstream Label scheme and the Label Set scheme for bidirectional lightpath setup in GMPLS WRONs have been compared. The results are submitted to Photonics in Switching 2008. The simulation results show that the Upstream Label (UL) scheme performs poorly unless detailed routing information is frequently advertised. The

³In this case, the Label Set only serves to fill out the Acceptable Label Set object carried in the PathErr message.

LS scheme outperforms the UL scheme in all the considered routing scenarios, especially at low and medium loads.

B.5 Benchmarking RWA Schemes

B.5.1 Introduction

WDM networks have shown to fulfill the increased capacity demands at the core of communication networks. In optical networks, a route and a wavelength must be found for each connection demand, which is referred to as the RWA problem. Obtaining the best possible assignment of routes and wavelengths for a set of connection requests requires joint optimization of the two tasks. Both tasks are however already difficult to solve on their own, not only in theory (NP-hard) [135] but also in practice (e.g., [136, 137]). In particular, the complexity of the wavelength assignment task increases rapidly with increasing number of requests. Consequently, their joint optimization is even more complex, see for example [138].

When looking into RWA problems, it is important to differentiate whether connection demands are known beforehand, or whether they arrive dynamically. Dynamic connection arrival is often more realistic, but the admission of the newest path depends on how the previous paths were allocated, if re-ordering of already provisioned paths is prohibited. If the demands are static, they can be allocated more efficiently in the network. Off-line optimization with a static set of demand requests can therefore serve as a benchmark for dynamic RWA strategies. In [139] such a study in terms of bandwidth efficiency and shared resource usage has been investigated for routing schemes.

During the last few years, Automatically Switched Optical Network (ASON) and GMPLS have been emerging as promising candidates to control wavelength routed networks. In dynamic GMPLS networks, knowledge of each individual wavelength's availability on each span is unavailable during the routing phase due to scalability issues [140], and the wavelength is hence only assigned during the signaling phase. The standard GMPLS signaling procedure does not contain functions that minimize the usage of wavelength converters during the signaling phase. To improve the performance, an additional object called the Suggested Vector [14, 21] was introduced in the signalling message and compared to the WC usage of the Label Set, which is standardized in the GMPLS framework [87], but the scheme has not been compared to the results of an off-line optimization for benchmarking. The concept for the bench-

marking study is presented here, while the results have have been published in [1].

Benchmarking Study

In this study, the performance of routing and wavelength assignment using off-line optimization is compared to a dynamic distributed provisioning scenario in wavelength convertible networks. This is the first time that the SV scheme performance is compared to off-line optimization. Furthermore, the concept of the SV is extended to cover bidirectional connections.

Two cases are considered: in the first case, the wavelength assignment (WA) is considered. Unlimited WCs are available in all nodes, and we compare the performance of the dynamic and off-line WA in terms of WC usage over the same paths. The computational experiments reveal that the number of converters could be significantly reduced if reassignment of wavelengths with an off-line algorithm would be possible.

In the second case, different routes for the dynamic and off-line optimized cases are allowed as well. This time, the number of available WCs per node is limited. In this case, we investigate if the optimized approach allows to accommodate connections that were blocked in the dynamic case, and how many WCs are used per provisioned connection in average. Experiments not only show that, again, there is a significant difference in the number of provisioned connections, but also that there seems to be a natural bound on the number of extra connections that could be established, compared to the dynamic setting. That is, the number of extra connections establishable if rerouting and reassignment of wavelengths would be possible is limited, most likely to a bound depending on the network topology and the demand distribution.

B.5.2 Summary of Benchmarking RWA Schemes

In this section, the concept of benchmarking routing and wavelength assignment schemes for dynamically controlled optical networks was described. Specifically, the performance of an advanced wavelength assignment scheme called SV, which was developed with the aim of ranking wavelengths (e.g., according to wavelength converter usage, as applied here) was evaluated during dynamic connection provisioning. The SV

scheme was compared to off-line routing and wavelength assignment for benchmarking. Two case studies were performed: in the first case, where only the wavelength assignment differs between the dynamic and the off-line approach differs, the wavelength converter usage is compared in a network with unlimited converter pools. In the second case, where both the routing and the wavelength assignment may be changed, the ability to provision connections and the wavelength converter usage when the wavelength conversion capability is severely limited was compared. Both studies show that the off-line algorithm can save considerably on number of wavelength converters and increase the capacity of the network, compared to the Suggested Vector scheme used for dynamic control. Considering that the latter cannot make use of knowledge of future demands, its performance is acceptable. The results of the study together with details of the mathematical model are submitted to Networks 2008.

B.6 Appendix Summary

This appendix summarized the concepts behind four network provisioning studies. Section B.2 presented details on label preference schemes and tie-breaking policies, which also were used in the restoration studies. Section B.3.2 presented three converter placement heuristics for evaluation under dynamic traffic. In section B.4, connection setup was extended to cover bidirectional connections. Finally, the SV as benchmarked against off-line RWA in section B.5.

Appendix C

Contributions to Papers

My contributions to the papers used in the main part of this thesis are outlined below:

- [6] This work gives a comprehensive overview on how limitation in the wavelength conversion capability affects both the provisioning and recovery phase in optical networks. I developed the framework of the paper and the idea for the two case studies. Implementation of the path restoration model was done together with Jakob Buron (JBU), while Nicola Andriolli (NAN) contributed to the theoretical sections of the SV. I generated the results and was the main responsible for the paper. All authors gave feedback on the paper and provided input and suggestions.
- [16] I came up with the basic idea for the BR scheme, but it was refined into its current form in collaboration with JBU. We were both involved in the modelling of this scheme and wrote sections of the paper. JBU was responsible for preparing the poster for the conference. All authors provided comments to the paper.

- [20] I presented the idea for the stub awareness study. JBU and I developed a survivability model in OPNET based on the provisioning model of NAN. I was the main responsible for the paper and presented it at the conference. All authors provided suggestions and feedback on the paper and presentation.
- [10] This work is an extension of the [20]. I was the main responsible for writing the paper, including result generation. Henrik Wessing (HW) and JBU were mainly responsible for the section of applying the SV in a multi-domain environment. I presented the paper at the conference, and all authors contributed with suggestions, comments and feedback to both the article and presentation.
- [15] The idea of comparing stub release in a simple scheme to a more complex scheme (SV) in local-to-egress restoration was proposed by me. The concept was implemented in OPNET together with JBU. I generated the results and was the main responsible for the paper and presented it at the conference. All authors provided suggestions and feedback on the paper and presentation.
- [8] I came up with the nodal stub release idea and discussed the concept with Lars Dittmann (LD). I implemented the nodal stub release model in OPNET, generated the results and was the main responsible for the paper. JBU and NAN gave feedback and commented on the paper.
- [24] This publication is an extension of the work commenced during my master's thesis. I was responsible for the research work (theory, modelling, writing of article and presentation), while LD and Lars Ellegaard (LE) gave feedback and proofread the article.
- [23] The idea of Shortcut Span Protection was developed jointly by Thomas Stidsen (TS) and myself. I was responsible for explaining the concept of SSP. TS was responsible for the mathematical formulation, the LP modelling and the presentation of the work at the conference.

My contribution to papers whose concept is summarized in the appendices is outlined below:

- [17] This paper describes the OPNET model used throughout this thesis in more detail. All authors have contributed to the code described in the paper. I was the main responsible of the paper, and I designed figures and flowcharts that illustrated the operation of the developed simulation model. I also presented the paper at the conference. All authors provided comments and feedback on the paper and presentation.
- [21] The main work of this paper was carried out by NAN during an exchange visit at DTU-Fotonik. I provided feedback on the paper, but was not involved in the modelling of the proposed SV concept.
- [14] This paper is a generalization of the concept proposed in [21], and NAN was the main responsible. I contributed to the paper through discussion, suggestions and proofreading.
- [2] The idea for the paper was developed jointly by JBU and myself. JBU, HW and myself wrote sections of the paper. The results were generated by JBU. All authors gave feedback on the paper.
- [3] The idea of working with bidirectional lightpath provisioning the GMPLS networks was initiated and discussed by NAN and myself. NAN was the main responsible for the paper. The simulations were carried out by Alessio Giorgetti. I contributed to the paper with content suggestions, feedback and proofreading.

- [1] This work was carried out during a Cost 293 STSM at Warwick University. I proposed the idea of the study, which consists of combining ILP and OPNET modelling. Arie Koster was responsible for the mathematical modelling. I was responsible for the discrete event simulation in OPNET. I extended the simulator to work with bidirectional connections, and wrote code that allows the integration between OPNET and CPLEX modelling. For the result generation and paper writing, we each contributed with our area and gave feedback to the other parts.

Bibliography

- [1] A. Koster and S. Ruepp, "Benchmarking RWA strategies for dynamically controlled optical networks," in *Proc. Networks 2008*.
- [2] J. Buron, S. Ruepp, H. Wessing, N. Andriolli, A. Manolova, and L. Dittmann, "Wavelength converter placement in optical networks with dynamic traffic," in *Proc. APOC 2008*.
- [3] N. Andriolli, A. Giorgetti, S. Ruepp, J. Buron, L. Valcarenghi, and P. Castoldi, "Bidirectional lightpath provisioning in GMPLS-controlled optical networks," in *Proc. Photonics in Switching 2008*.
- [4] X. Li, S. Ruepp, L. Dittmann, and A. Manolova, "Survivability-enhancing routing scheme for multi-domain networks," in *Proc. Globecom 2008*.
- [5] A. M. X. Li, S. Ruepp and L. Dittmann, "Opnet model of multi-domain routing with enhanced survivability," in *Proc. Opnetwork 2008*.
- [6] S. Ruepp, N. Andriolli, J. Buron, L. Dittmann, and L. Ellegard, "Restoration in all-optical GMPLS networks with limited wavelength conversion," *Computer Networks Special Issue on Opportunities and Challenges in Optical Networks*, 2008, accepted for publication.
- [7] C. Raffaelli, K. Vlachos, N. Andriolli, J. Buron, R. van Caenegem, G. Danilewicz, J. M. Finochietto, J. Garcia-Haro, D. Klondis, M. O'Mahony, G. Maier, A. Pattavina, P. Pavon-Marino, S. Ruepp, M. Savi, M. Scaffardi, I. Tomkos, A. Tzanakaki, L. Wosinska,

- K. Yannopoulos, and F. Neri, "Photonics in switching: architectures, technologies and systems the e-photon/one+ research results," *Computer Networks Special Issue on Opportunities and Challenges in Optical Networks*, 2008, accepted for publication.
- [8] S. Ruepp, J. Buron, N. Andriolli, and L. Dittmann, "Nodal stub release in all-optical networks," *IEEE Commun. Lett.*, vol. 12, no. 1, pp. 47–49, Jan. 2008.
- [9] A. Manolova, J. Buron, S. Ruepp, L. Dittmann, and L. Ellegaard, "Modelling contention-resolution in optical burst switching," in *Proc. Opnetwork*, Aug. 2007.
- [10] S. Ruepp, J. Buron, N. Andriolli, and H. Wessing, "Span restoration in optical networks with limited wavelength conversion," in *Proc. Chinacom*, Aug. 2007.
- [11] P. Castoldi, N. Andriolli, I. Cerutti, N. Sambo, L. Valcarenghi, A. Giorgetti, S. Ruepp, and J. Buron, "A framework for label preference in gmpls controlled optical networks," in *Proc. ICTON*, July 2007, invited.
- [12] A. Manolova, J. Buron, S. Ruepp, L. Dittmann, and L. Ellegaard, "Segmentation-based path switching mechanism for reduced data losses in OBS networks," in *Proc. Optical Network Design and Modelling Conference (ONDM)*, May 2007.
- [13] A. Manolova, S. Ruepp, J. Buron, L. Dittmann, and L. Ellegaard, "Advantages and challenges of the GMPLS/OBS integration," in *Proc. The VI Workshop in G/MPLS Networks*, April 2007.
- [14] N. Andriolli, J. Buron, S. Ruepp, F. Cugini, L. Valcarenghi, and P. Castoldi, "Label preference schemes in GMPLS controlled networks," *IEEE Commun. Lett.*, vol. 12, no. 10, p. 1, Dec. 2006.
- [15] S. Ruepp, J. Buron, and N. Andriolli, "Increasing restorability for local-to-egress restoration in GMPLS controlled networks with limited wavelength conversion," in *Proc. ICCS*, Oct. 2006.

-
- [16] J. Buron, S. Ruepp, and N. Andriolli, "Blocking reduction of span restoration requests in GMPLS controlled WDM optical networks," in *Proc. ECOC*, Sept. 2006.
- [17] S. Ruepp, N. Andriolli, and J. Buron, "Signaling protocol extensions enhancing provisioning and recovery in GMPLS optical networks," in *Proc. Opnetwork*, Aug. 2006.
- [18] J. Buron, S. Ruepp, and A. V. Manolova, "Teaching cost-effective and resilient network design with OPNET WDM guru," in *Proc. Opnetwork*, Aug. 2006.
- [19] N. Andriolli, J. Buron, S. Ruepp, and P. Castoldi, "Modeling GMPLS-controlled optical packet switched (OPS) networks," in *Proc. Opnetwork*, Aug. 2006.
- [20] S. Ruepp, J. Buron, and N. Andriolli, "Wavelength-converter saving span restoration in GMPLS controlled WDM optical networks," in *Proc. Conference of Optical Internet, COIN*, July 2006.
- [21] N. Andriolli, J. Buron, S. Ruepp, F. Cugini, L. Valcarengi, and P. Castoldi, "Signaling protocol extensions for converter-saving wavelength assignment in GMPLS optical networks," in *Proc. HPSR*, Jun. 2006.
- [22] N. Andriolli, H. Harai, J. Buron, S. Ruepp, and P. Castoldi, "Wavelength assignment schemes for connection-oriented OPS networks," in *Proc. ONDM*, May. 2006.
- [23] T. Stidsen and S. Ruepp, "Shortcut span protection," in *Extended Abstract for Inform's Telecom*, 2006.
- [24] S. Ruepp, L. Dittmann, and L. Ellegaard, "Simulation and comparison of path restoration techniques in sdh mesh networks," in *Proc. Fifth International Workshop on the Design of Reliable Communication Networks (DRCN)*, Oct. 2005, pp. 47–53.
- [25] International Telecommunication Union (ITU-T), "Recommendation I.321: B-ISDN protocol reference model and its application."

- [26] G. Bernstein, B. Rajagopalan, and D. Saha, *Optical Network Control - Architecture, Protocols and Standards*. Boston, USA: Pearson Education Inc., 2004, ISBN: 0-20-175301-4.
- [27] E. Mannie (ed.), "Generalized multi-protocol label switching (GMPLS) architecture," *RFC 3945*, Oct. 2004.
- [28] International Telecommunication Union (ITU-T), "Recommendation G.8080: Architecture for the automatically switched optical network (ASON)."
- [29] OIF, "User network interface (UNI) 1.0 signaling specification, release 2," Feb. 2004.
- [30] ———, "Intra-carrier E-NNI signaling specification," Feb. 2004.
- [31] P. Blair and J. B. (Ovisor Technologies), "ROADMs - the traffic cops of agile optical networks," *Europhotonics*, European convergence of optics, lasers, imagine, fibre optics, electro-opticas and photonic component manufacturing - Laurin Publications, Feb./Mar. 2007.
- [32] G. Charlet, J. Renaudier, H.Mardoyan, P.Tran, O. Pardo, F. Verluise, M. Achouche, A. Boutin, F. Blache, J.-Y.Dupuy, and S. Bigo, "Transmission of 16.4 Tbit/s capacity over 2,550 km using PDM QPSK modulation format andcoherent receiver," in *Proc. OFC*, 2008, PDP3.
- [33] B. Mukherjee, *Optical WDM Networks*. Springer, 2006, ISBN: 0-387-29055-9.
- [34] W. Grover, *Mesh-Based Survivable Networks - Options and Strategies for Optical, MPLS, SONET, and ATM Networking*. Prentice-Hall, 2004, ISBN: 0-13-494576-X.
- [35] B. Ramamurthy and B. Mukherjee, "Wavelength conversion in WDM networking," *IEEE J. Sel. Areas Commun.*, vol. 16, no. 7, pp. 1061–1073, 1998.
- [36] K.-C. Lee and V. Li, "A wavelength-convertible optical network," *J. Lightw. Technol.*, vol. 11, no. 5/6, pp. 962–970, May/Jun. 1993.

- [37] J. Elmirghani and H. Mouftah, "All-optical wavelength conversion: technologies and applications in DWDM networks," *IEEE Commun. Mag.*, vol. 38, no. 3, pp. 86–92, Mar. 2000.
- [38] S. Yoo, "Wavelength conversion technologies for WDM network applications," *J. Lightw. Technol.*, vol. 14, no. 6, pp. 955–966, Jun. 1996.
- [39] J. Iness and B. Mukherjee, "Sparse wavelength conversion in wavelength-routed WDM optical networks," *Photonic Network Communication*, vol. 1, no. 3, pp. 183–205, 1999.
- [40] G. Shen, S. Bose, T. Cheng, C. Lu, and T. Chai, "Operation of WDM networks with different wavelength conversion capabilities," *IEEE Commun. Lett.*, vol. 4, no. 7, pp. 239–241, July 2000.
- [41] M. Swaminathan and K. Sivarajan, "Practical routing and wavelength assignment algorithms for all optical networks with limited wavelength conversion," in *Proc. of ICC*, 2002, pp. 2750–2755.
- [42] Q.-D. Ho and M.-S. Lee, "Converter-aware wavelength assignment in WDM networks with limited-range conversion capability," *IEICE - Transactions on Communications*, vol. Vol.E89-B, no. 2, pp. 436–445, Feb. 2006.
- [43] S. Ruepp, "Network survivability - market research based on interviews with telecom network operators [confidential]," carried out as part of the requirements for Innovation Ph.D. students, approved by DTU Innovation Committee in 2006.
- [44] —, "Network survivability - business plan [confidential]," carried out as part of the requirements for Innovation Ph.D. students, approved by DTU Innovation Committee in 2007.
- [45] A. Farrel and I. Bryskin, *GMPLS Architecture and Applications*. Morgan Kaufmann, Elsevier, 2006, ISBN: 0-12-088422-4.
- [46] S. Tomic, B. Statovci-Halimi, A. Halimi, W. Müllner, and J. Frühwirth, "ASON and GMPLS - overview and comparison," *Photonic Network Communications*, vol. 7, no. 2, pp. 111–130, 2004.

- [47] N. Larkin, "ASON and GMPLS - the battle of the optical control plane: An overview of ongoing work of the IETF and ITU to standardize optical control plane protocols," Data Connection Whitepaper, Aug. 2002. [Online]. Available: <http://www.dataconnection.com/network/download/whitepapers/asongmpls.pdf>
- [48] D. Papadimitriou, J. Drake, J. Ash, A. Farrel and L. Ong (eds.), "Requirements for generalized MPLS (GMPLS) signaling usage and extensions for automatically switched optical network (ASON)," *RFC 4139*, July 2005.
- [49] International Telecommunication Union (ITU-T), "Recommendation G.7712: Architecture and specification for the data communications network (DCN)."
- [50] —, "Recommendation G.7715: Architecture and requirements for routing."
- [51] —, "Recommendation G.7713: Distributed connection management (DCM)."
- [52] —, "Recommendation G.7714: Neighbor discovery."
- [53] K. Kompella and Y. Rekhter (eds.), "OSPF extensions in support of generalized multi-protocol label switching (GMPLS)," *RFC4203*, Oct. 2005.
- [54] D. Katz, K. Kompella and D. Yeung (eds.), "Traffic engineering (TE) extensions to OSPF version 2," *RFC 3630*, Sept. 2003.
- [55] J. Lang (ed.), "Link management protocol (LMP)," *RFC4204*, Oct. 2005.
- [56] Optical Internet Forum, "Intra-carrier E-NNI signaling specification," *OIF-E-NNI-Sig-01.0*, Feb. 2004.
- [57] Cisco, "Management plane protection," Feb. 2006. [Online]. Available: http://www.cisco.com/en/US/docs/ios/12_4t/12_4t11/htsecmpp.pdf

- [58] ———, “Network management system: Best practises whitepaper,” July 2007. [Online]. Available: http://www.cisco.com/warp/public/126/NMS_bestpractice.pdf
- [59] TDC Pressesekretariat, “Flere kabler graves over,” cited: 3. June 2005. [Online]. Available: <http://www.tdc.dk>
- [60] J.-P. Vasseur, M. Pickavet, and P. Demeester, *Network Recovery, Protection and Restoration of Optical, SONET-SDH, IP, and MPLS*. Morgan-Kaufmann Publishers, Elsevier, 2004, ISBN: 0-12-715051-x.
- [61] S. Dong, C. Phillips, and R. Friskney, “Differentiated resilience provisioning for the wavelengthrouted optical network,” *J. Lightw. Technol.*, vol. 24, no. 2, pp. 667–673, Feb. 2006.
- [62] C. Gruber, “A comparison of bandwidth requirements of path protection mechanisms,” *ICN 2005, LNCS Springer*, vol. 3420, pp. 133–143, 2005.
- [63] E. Calle, J. Marzo, and A. Urra, “Protection performance components in MPLS networks,” *Computer Communications, Elsevier*, vol. 27, pp. 1220–1228, 2002.
- [64] W. Grover and D. Stamatelakis, “Cycle-oriented distributed configuration: ring-like speed with mesh-like capacity for self-planning network restoration,” *International Conference on Communications ICC*, vol. 1, pp. 537–543, 1998.
- [65] C.-S. Li and R. Ramaswami, “Automatic fault detection, isolation and recovery in transparent all-optical networks,” *J. Lightw. Technol.*, vol. 15, pp. 1784–1793, Oct. 1997.
- [66] C. Mas, I. Tomkos, and K. Tonguz, “Failure location algorithm for transparent optical networks,” *IEEE J. Sel. Areas Commun.*, vol. 23, pp. 1508–1519, Aug. 2005.
- [67] R. Gallaher, “Introduction to multi-protocol lambda switching (MPLS) and generalized multi-protocol label switching (GMPLS),” cited: 22. Feb. 2008. [Online]. Available:

[http://www.convergedigest.com/
tutorials/mpls6/page1.asp](http://www.convergedigest.com/tutorials/mpls6/page1.asp)

- [68] D. Schupke and R. Prinz, "Capacity efficiency and restorability of path protection and rerouting in WDM networks subject to dual failures," *Photonic Network Communications*, vol. 8, no. 2, pp. 191–207, 2004.
- [69] J. Wang, L. Sahasrabudde, and B. Mukherjee, "Path vs. subpath vs. link restoration for fault management in IP-over-WDM networks: Performance comparisons using GMPLS control signaling," *IEEE Commun. Mag.*, vol. 40, no. 11, pp. 80–87, Nov. 2002.
- [70] J. Zhang and B. Mukherjee, "A review of fault management in WDM mesh networks: Basic concepts and research challenges," *IEEE Netw.*, vol. 18, no. 2, pp. 41–48, Mar./Apr. 2004.
- [71] E. Mannie and D. Papadimitriou (eds.), "Recovery (protection and restoration) terminology for generalized multi-protocol label switching (GMPLS)," *RFC 4427*, Mar. 2006.
- [72] G. Suixiang, J. Xiaohua, H. Chuanhe, and D. Ding-Zhu, "An optimization model for placement of wavelength converters to minimize blocking probability in WDM networks," *J. Lightw. Technol.*, vol. 21, no. 3, pp. 684–694, Mar. 2003.
- [73] M. Ali, "Shareability in optical networks: beyond bandwidth optimization," *IEEE Commun. Mag.*, vol. 42, no. 2, pp. S11–S15, Feb. 2004.
- [74] —, "Generalized sharing in survivable optical networks," *IEEE/ACM Trans. on Networking*, vol. 14, no. 6, pp. 1388–1399, Dec. 2006.
- [75] A. Koster and A. Zymolka, "Tight LP-based lower bounds for wavelength conversion in optical networks," *Statistica Neerlandica*, vol. 61, no. 1, pp. 115–136, Feb. 2007.
- [76] A. Fumagalli and L. Valcarenghi, "IP restoration vs. WDM protection: Is there an optimal choice?" *IEEE Netw.*, vol. 14, no. 6, pp. 34–41, Nov./Dec. 2000.

-
- [77] S. Ramamurthy, L. Sahasrabudde, and B. Mukherjee, "Survivable WDM mesh networks," *J. Lightw. Technol.*, vol. 21, no. 4, pp. 870–883, Apr. 2003.
- [78] A. V. Sichani and H. T. Mouftah, "Performance evaluation of dynamic restoration techniques for survivable optical networks," in *Proc. ISCC*, 2006, pp. 961–966.
- [79] Y. Luo and N. Ansari, "Restoration with wavelength conversion in WDM networks," *Electronic Letters*, vol. 38, no. 16, pp. 900–901, Aug. 2002.
- [80] —, "Split restoration with wavelength conversion in WDM networks," in *Proc. of ICC*, 2003, pp. 1423–1427.
- [81] G. Conte, M. Listanti, M. Settembre, and R. Sabella, "Strategy for protection and restoration of optical paths in WDM backbone networks for next-generation internet infrastructures," *J. Lightw. Technol.*, vol. 20, no. 8, pp. 1264–1276, 2002.
- [82] M. Keshtgary, F. Al-Zahrani, A. Fahad, M. Jahangir, and A. Jayasumana, "Survivability performance evaluation of WDM networks with wavelength converters," *Photonic Network Communications*, vol. 11, no. 1, pp. 15–27, 2006.
- [83] M. Azim, X. Jiang, P. Ho, and S. Horiguchi, "Models of restoration probability in WDM networks employing active restoration," *Photonic Network Communications*, vol. 10, no. 2, pp. 141–153, 2005.
- [84] —, "Behaviors of active lightpath restoration in all-optical networks with and without wavelength converters," in *Proc. of Asia Pacific Conference*, Oct. 2005, pp. 941–945.
- [85] Z. Zhang and Y. Yang, "On-line optimal wavelength assignment in WDM networks with shared wavelength converter pool," *IEEE/ACM Trans. Netw.*, vol. 15, no. 1, pp. 234–245, 2007.
- [86] P.-H. Ho and H. Mouftah, "A novel survivable routing algorithm for shared segment protection in mesh WDM networks with partial

- wavelength conversion,” *J. Lightw. Technol.*, vol. 22, no. 8, pp. 1548–1560, 2004.
- [87] L. Berger (ed.), “Generalized multi-protocol label switching (GMPLS) signaling functional description,” *RFC 3471*, Jan. 2003.
- [88] D. Papadimitriou and E. Mannie (eds.), “Analysis of generalized multi-protocol label switching (GMPLS)-based recovery mechanisms (including protection and restoration),” *RFC 4428*, Mar. 2006.
- [89] J. Zheng, B. Zhang, and H. Mouftah, “Dynamic path restoration based on multi-initiation for GMPLS-based WDM networks,” in *Proc. of ICC*, 2004, pp. 1639–1643.
- [90] S. Albarrak, “Pre-allocated restoration technique in a distributed GMPLS-based IP/WDM mesh network,” in *Proc. of Information and Computer Technologies (ICTTA)*, 2006, pp. 2236–2241.
- [91] H. Zang, J. P. Jue, and B. Mukherjee, “A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks,” *Optical Netw. Mag.*, Jan. 2000.
- [92] “OPNET modeler, OPNET technologies, inc.” [Online]. Available: <http://www.opnet.com>
- [93] R. Inkret, A. Kuchar, and B. Mikac, Eds., *Extd. Final Report of COST Action 266*. Univ. of Zagreb, 2003.
- [94] S. D. Maesschalck, D. Colle, I. Lievens, M. Pickavet, P. Demeester, C. Mauz, M. Jäger, R. Inkret, B. Mikac, and J. Derkacz, “Pan-european optical transport networks: an availability-based comparison,” *Photonic Network Communications*, vol. 5, no. 3, pp. 203–225, May 2003.
- [95] E. Dijkstra, “A note on two problems in connection with graphs,” *Numerische Matematik*, vol. 1, pp. 269–271, 1959.
- [96] A. Farrel (ed.), “Crankback signaling extensions for MPLS and GMPLS RSVP-TE,” *RFC 4920*, Jan. 2007.

- [97] M. Miyao and H. Saito, "Optimal design and evaluation of survivable WDM transport networks," *IEEE J. Sel. Areas Commun.*, vol. 16, no. 7, pp. 1190–1199, 1998.
- [98] Y. Xue, M. Lazer, A. Nagarajan, H. Ishimatsu, O. Aparicio, and S. Wright, "Carrier optical services requirements," *draft-ietf-ipo-carrier-requirements-03.txt*, June 2002, expires: Dec. 2002.
- [99] H.-M. Foisel, J. Spaeth, C. Cavazzoni, H. Wessing, and M. Popov, "MUPBED - interworking challenges in a multi-domain and multi-technology network environment," *Proc. of Optical Network Design and Modeling (ONDM), Athens, Greece*, pp. 269–278, May 2007.
- [100] R. Iraschko, M. McGregor, and W. Grover, "Optimal capacity placement for path restoration in STM or ATM mesh-survivable networks," *IEEE/ACM Trans. on Networking*, vol. 6, no. 3, pp. 325–336, June 1998.
- [101] M. Lewis, "Joint capacity allocation with modularity and stub release in mesh telecommunication networks: an experimental design approach," *Technical Report, School of Business, University of Mississippi*, 2002, cited: 25. Feb. 2008. [Online]. Available: <http://faculty.bus.olemiss.edu/mlewis/ResearchPapers/JointCapacityAllocationR2.pdf>
- [102] S. Orłowski and R. Wessälly, "Comparing restoration concepts using optimal network configurations with integrated hardware and routing decisions," *Journal of Network and Systems Management*, vol. 13, no. 1, pp. 99–118, March 2005.
- [103] phd-DBENCANN, "German network." [Online]. Available: <http://www.sndlib.zib.de>
- [104] "Max-flow min-cut theorem," cited: 11. Feb. 2008. [Online]. Available: http://www.fact-index.com/m/ma/max_flow_min_cut_theorem.html
- [105] M. Herzberg, S. Bye, and A. Utano, "The hop limit approach for spare capacity assignment in survivable networks," *IEEE/ACM Trans. on Networking*, vol. 13, 1995.

- [106] W. Grover, V. Rawat, and M. McGregor, "A fast heuristic principle for spare capacity placement in mesh-restorable SONET/SDH transport networks," *Electronic Letters*, vol. 33, pp. 195–196, Jan. 1997.
- [107] J. Doucette and W. Grover, "Influence of modularity and economy-of-scale effects on design of mesh-restorable DWDM networks," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 10, pp. 1912–1923, October 2000.
- [108] W. Grover, T. Bilodeau, and B. Venables, "Near optimal synthesis of a mesh restorable network," in *Proc. GLOBECOM*, 1991, pp. 2007–2012.
- [109] J. Kennington and M. Lewis, "The path restoration version of the spare capacity allocation problem with modularity restrictions: models, algorithms, and an empirical analysis," *INFORMS Journal on Computing*, vol. 13, pp. 181–190, 2001.
- [110] O. Holst, *Optimering på netværk - Forlæsningsnote*, 2nd ed. Danmarks Tekniske Universitet, 1980.
- [111] A. Leon-Garcia and I. Widjaja, *Communication Networks - Fundamental Concepts and Key Architectures*. McGraw-Hill, 2003, ISBN: 0-07-123026-2.
- [112] W. Grover and J. Doucette, "Topological design of survivable mesh-based transport networks," *Annals of Operations Research, Kluwer Academic Publishers*, vol. 106, pp. 79–125, 2001.
- [113] F. Nielsen, *Grafteori - Algoritmer og netværk*. Danmarks Tekniske Universitet, 1995, ISBN: 87-88764-49-4.
- [114] J. Moy (ed.), "OSPF version 2," *RFC 2328*, Apr. 1998.
- [115] S. Ruepp, *Network Restoration*. M.Sc. Thesis, Oct. 2004, carried out at COM-DTU, Technical University of Denmark.
- [116] "The MathWorks, inc." [Online]. Available: <http://www.mathworks.com>

- [117] J.-F. Laboudette, "Shared mesh restoration in optical networks," *Optical Fiber Communication Conference (OFC)*, March 2004.
- [118] A. Authenrieth and A. Kirstädter, "Engineering end-to-end IP resilience using resilience-differentiated QoS," *IEEE Commun. Mag.*, pp. 50–57, January 2002.
- [119] "Optical fiber supply meets expanding application," *SK Opto-Electronics Newsletter*, vol. 12, pp. 50–57, February 2003.
- [120] P. Batchelor, B. Daino, P. Heinzmann, D. R. Hjelme, P. Leuthold, R. Inkret, G. D. Marchis, H. A. Jager, F. Matera, M. Join-dot, B. Mikac, A. Kuchar, H.-P. Nolting, E. Coquil, J. Späth, F. Tillerot, B. Caenegem, N. Wauters, and C. Weinert, "Study on the implementation of optical transparent transport networks in the european environment-results of the research project COST 239," *Photonic Network Communications*, vol. 2, no. 1, pp. 15–32, May 2000.
- [121] J. Doucette, D. He, W. Grover, and O. Yang, "Algorithmic approaches for efficient enumeration of candidate p -cycles and capacitated p -cycle network design," in *Proc. Workshop on Design of Reliable Communication Networks (DRCN)*, 2003, pp. 212–220.
- [122] W. Grover, J. Doucette, M. Clouqueur, D. Leung, and D. Stamatakis, "New options and insights for survivable transport networks," *IEEE Commun. Mag.*, vol. 40, no. 1, pp. 15–32, 2002.
- [123] T. Stidsen and P. Kjaerulff, "Complete rerouting protection," *Journal of Optical Networking*, vol. 5, no. 6, pp. 481–492, 2006.
- [124] V. Iversen, *Teletraffic Engineering - Basic Teletraffic Theory*. Course Compendium, Technical University of Denmark, Jan. 2002.
- [125] P. Belotti and T. Stidsen, "Optimal placement of wavelength converting nodes," in *Proc. Third International Workshop on the Design of Reliable Communication Networks (DRCN)*, 2004.
- [126] S. Subramaniam, M. Azizoglu, and A. K. Somani, "On optimal converter placement in wavelength-routed networks," *IEEE/ACM Trans. Networking*, vol. 7, no. 5, pp. 754–766, 1999.

- [127] H.-Y. Jeong and S.-W. Seo, "A binary (0-1) linear program formulation for the placement of limited-range wavelength converters in wavelength-routed WDM networks," *J. Lightw. Technol.*, vol. 23, no. 10, pp. 3076–3091, 2005.
- [128] K. R. Venugopal, M. Shivakumar, and P. S. Kumar, "A heuristic for placement of limited range wavelength converters in all-optical networks," in *Proc. IEEE INFOCOM*, 1999, pp. 908–915.
- [129] S. K. Bose, Y. N. Singh, A. Raju, and B. Papat, "Sparse converter placement in WDM networks and their dynamic operation using path-metric based algorithms," in *Proc. ICC*, 2002, pp. 2855–2859.
- [130] S. Xu, H. Harai, and D. King, "Extensions to GMPLS RSVP-TE for bidirectional lightpath with the same wavelength," *draft-xu-rsvp-te-bidir-wave-01.txt*, Nov. 2007, expires: May 2008.
- [131] E. Oki, N. Matsuura, K. Shiimoto, and N. Yamanaka, "Bidirectional path setup scheme using an upstream label set in optical GMPLS networks," *IEICE - Transactions on Communications*, vol. E87-B, no. 6, pp. 1569–1576, June 2004.
- [132] G. Bernstein, Y. Lee, and W. Imajuku, "Framework for GMPLS and PCE control of wavelength switched optical networks," *draft-bernstein-ccamp-wavelength-switched-03.txt*, Feb. 2008, expires: August 2008.
- [133] D. Li, Y. Lee, and J. Gao, "Evaluation of possible interior gateway protocol extensions for wavelength switching optical networks," *draft-li-ccamp-wson-igp-eval-00.txt*, Nov. 2007, expires: May 2008.
- [134] A. Farrel and D. Papadimitriou and J.P. Vasseur and A. Ayyangar (eds.), "Encoding of attributes for MPLS LSP establishment using RSVP-TE," *RFC 4420*, Feb. 2006.
- [135] A. Ferreira, S. Pérennes, A. W. Richa, H. Rivano, and N. S. Moses, "Models, complexity and algorithms for the design of multifiber WDM networks," *Telecommunication Systems*, vol. 24, pp. 123–138, 2003.

-
- [136] M. Andrews and L. Zhang, "Complexity of wavelength assignment in optical network optimization," in *Proc. IEEE INFOCOM*, 2006.
- [137] A. Zymolka, "Design of survivable optical networks by mathematical optimization," 2007.
- [138] B. Jaumard, C. Meyer, and B. Thiongane, "ILP formulations for the RWA problem: Symmetric systems," pp. 637–677, 2006.
- [139] A. Elwalid, D. Mitra, I. Saniee, and I. Widjaja, "Routing and protection in GMPLS networks: From shortest paths to optimized designs," *J. Lightw. Technol.*, vol. 21, no. 11, pp. 2828–2838, 2003.
- [140] Z. Zhang, J. Fu, D. Guo, and L. Zhang, "Lightpath routing for intelligent optical networks," *IEEE Networks*, pp. 28–35, July/Aug. 2001.

List of Acronyms

AI	Aggregated Information
AIS	Alarm Indication Signal
ASON	Automatically Switched Optical Network
BLSR	Bidirectional Line Switched Ring
BR	Blocking Reduction
DCM	Distributed Connection Management
DCN	Data Communications Network
DI	Detailed Information
E-NNI	External Network-to-Network Interface
GMPLS	Generalized Multi-Protocol Label Switching
IETF	Internet Engineering Task Force
I-NNI	Internal Network-to-Network Interface
IP	Internet Protocol
ITU	International Telecommunications Union
JCA	Joint Capacity Allocation
LMP	Link Management Protocol
LP	Linear Program

LS	Label Set
LSA	Link State Advertisements
LSP	Label Switched Path
MPLS	Multi-Protocol Label Switching
NI	No Information
NNI	Network-to-Network Interface
NP	No Preference
OEO	Optical-Electro-Optical
OIF	Optical Internet Forum
OSPF	Open Shortest Path First
OXC	Optical Cross-Connects
RDI	Remote Defect Indication
ROADM	Reconfigurable Optical Add-Drop Multiplexer
RSVP	Resource Reservation Protocol
RWA	Routing and Wavelength Assignment
SCA	Spare Capacity Allocation
SDH	Synchronous Digital Hierarchy
SL	Suggested Label
SRLG	Shared Risk Link Group
SSP	Shortcut Span Protection
STM	Synchronous Transport Module
SV	Suggested Vector
TE	Traffic Engineering

UNI	User-Network Interface
UL	Upstream Label
UPSR	Unidirectional Path Switched Ring
VC	Virtual Circuit
WC	Wavelength Converter
WDM	Wavelength Division Multiplexing
WRON	Wavelength-Routed Optical Network

