

Optical Network Design and Planning

Optical Networks

Series Editor: Biswanath Mukherjee
 University of California, Davis
 Davis, CA

Optical Network Design and Planning

Jane M. Simmons

ISBN 978-0-387-76475-7

Quality of Service in Optical Burst Switched Networks

Kee Chaing Chua, Mohan Gurusamy, Yong Liu, and Minh Hoang Phung

ISBN 978-0-387-34160-6

Optical WDM Networks

Biswanath Mukherjee

ISBN 978-0-387-29055-3

Traffic Grooming in Optical WDM Mesh Networks

Keyao Zhu, Hongyue Zhu, and Biswanath Mukherjee

ISBN 978-0-387-25432-6

Survivable Optical WDM Networks

Canhui (Sam) Ou and Biswanath Mukherjee

ISBN 978-0-387-24498-3

Optical Burst Switched Networks

Jason P. Jue and Vinod M. Vokkarane

ISBN 978-0-387-23756-5

Jane M. Simmons

Optical Network Design and Planning

 Springer

Jane M. Simmons
Monarch Network Architects
Holmdel, NJ
USA

Series Editor
Biswanath Mukherjee
University of California, Davis
Davis, CA
USA

ISBN: 978-0-387-76475-7 e-ISBN: 978-0-387-76476-4
DOI: 10.1007/978-0-387-76476-4

Library of Congress Control Number: 2008920263

© 2008 Springer Science+Business Media, LLC

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer Science+Business Media, LLC, 233 Spring Street, New York, NY 10013, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.

The use in this publication of trade names, trademarks, service marks and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

Printed on acid-free paper

9 8 7 6 5 4 3 2 1

springer.com

To my beautiful mother, Marie

Foreword

The huge bandwidth demands predicted at the start of the millennium have finally been realized. This has been sparked by the steady growth of a variety of new broadband services such as high-speed Internet applications, residential video-on-demand services, and business virtual private networks with remote access to huge databases. In response, carriers are undergoing widespread upgrades to their metro and backbone networks to greatly enhance their capacity. Carriers are demanding WDM optical networking technologies that provide both low capital expenses and low operational expenses. This need has been satisfied by automatically reconfigurable optical networks that support optical bypass. Automatic reconfigurability enables the carriers, or their customers, to bring up new connections and take down existing ones to meet fluctuating bandwidth requirements in near real-time. It also enables rapid automatic restoration from network failures. The optical-bypass property of the network, coupled with long-reach WDM optics, greatly reduces the need for optical-electrical-optical conversion, thus resulting in huge savings in capital and operational expenses.

This book provides a timely and thorough coverage of the various aspects of the design and planning of automatically reconfigurable optical networks in general, with special emphasis on optical-bypass-enabled networks. While the reality of such networks today is somewhat different from the earlier research visions of a purely all-optical network that is transparent to signal format and protocol, the goals of greatly improved economics, flexibility, and scalability have been realized. The optical-bypass networking paradigm has been adopted by many of the major carriers around the world, in both metro and backbone networks. Moreover, efficient optical networking algorithms have emerged as one of the critical components that have enabled this technology to work in practice.

This book provides broad coverage of the architecture, algorithms, and economics of optical networks. It differs from other books on this general subject in that it focuses on real-world networks and it provides good perspective on the practical aspects of the design and planning process. The book serves as a valuable guide to carriers, vendors, and customers to help them better understand the intricacies of the design, planning, deployment, and economics of optical networks. The book also provides practitioners, researchers, and academicians with a wealth of knowledge and ideas on efficient and scalable optical networking algorithms that are suitable for a broad range of optical networking architectures and technologies.

Jane Simmons has been actively working in this area since the mid 1990s. In this time-frame, there was much activity covering all aspect of optical networking - technology, architecture, algorithms, control, and applications. A particularly influential research effort that started in the United States around this time, in which Jane participated, was the government-supported Multiwavelength Optical Networking (MONET) consortium among the telecommunications giants AT&T, Lucent, Verizon and SBC. Just a short time later, much of the vision generated by this research was turned into reality. In the 2000 time-frame, Corvis Corporation became the first company to commercialize the ‘all-optical’ backbone-network vision when it introduced a product with 3,200-km optical reach and the associated optical switching equipment. Jane played a key role at Corvis, where she developed efficient and scalable networking algorithms to support and exploit this technology. This culminated in the first commercial deployment of the ‘all-optical’ vision with Broadwing’s backbone network, in 2001. Jane performed the network design for the Broadwing network, from link engineering to network architecture. Jane also performed network designs for a broad array of North American and European carriers. She successfully showed in these diverse and real environments that ‘all-optical’, or more accurately, optical-bypass-enabled networks are architecturally viable in terms of achieving high network efficiency.

She has continued to work on optical network architecture and algorithms, as a founding partner of Monarch Network Architects, which provides architectural services and networking tools to carriers and system vendors. With this vast experience, and being in the right place at the right time, Jane has developed a unique perspective in the field of optical networking, which she brings forward in this book. I thoroughly enjoyed reading it and I learned a lot from it. I am sure that the reader is in for a real treat!

Dr. Adel A. M. Saleh
Program Manager
DARPA Strategic Technology Office

Preface

I have been involved with the research and development of optical networks for the past 15 years. More specifically, I have worked on the architecture and algorithms of networks with optical bypass, where much of the electronic regeneration is removed from the network. These networks are referred to in this book as optical-bypass-enabled networks.

Optical bypass has progressed from a research topic to a commercial offering in a relatively short period of time. I was fortunate to be in the midst of the activity, as a member of Bell Labs/AT&T Labs Research and Corvis Corporation. There are a few key lessons learned along the way, which I hope have been successfully captured in this book.

First, algorithms are a key component of optical networks. It is not hard to produce studies where poor algorithms lead to inefficient network utilization. Conversely, armed with a good set of algorithms, one can generate efficient designs across a range of network topologies, network tiers, and traffic distributions. It is also important to stress that while replacing electronics with optics in the network poses unique challenges that require algorithms, which is often cited as a concern by the opponents of such networking technology, the design of electronic-based networks requires algorithms as well. Processes such as shared protection or subrate-traffic-grooming are complex enough that algorithms are needed regardless of the nature of the underlying technology.

Second, there should be a tight development relationship between the system engineers, hardware designers, and the network architects of any system vendor developing optical networking equipment. The mantra of many a hardware developer when dealing with the potentially messy consequences of a design decision is often 'the algorithms will take care of it'. While their confidence in the algorithms may be flattering, this is not always the wisest course of action. It is the responsibility of the network architects to push back when appropriate to ensure that the overall system complexity does not grow unwieldy. Based on experience, when challenged, much more elegant solutions were forthcoming. Of course, there are times when the physics of the problem, as opposed to expediency, dictates a solution; it is important to recognize the difference.

This leads to the last point in that the algorithms in a well-designed system do not need to be overly complex. Much effort has been put into algorithm development, which has been successful in producing efficient and scalable algorithms.

Furthermore, it is not necessary that the algorithms take many hours or days to run. With well-honed heuristics, a design that is very close to optimal can often be produced in seconds to minutes.

The primary goal of this book is to cover the aspects of optical network design and planning that are relevant in a practical environment. The emphasis is on planning techniques that have proved to be successful in actual networks, as well as on potential trouble areas that must be considered. While the algorithms and architecture are the core of the content, the various enabling optical network elements and the economics of optical networking are covered as well. The book is intended for both practitioners and researchers in the field of optical networking.

The first two chapters should be read in order. Chapter 1 puts the book in perspective and reviews the terminology that is used throughout the book. Chapter 2 covers the various optical network elements; it is important to understand the functionality of the elements as it motivates much of the remainder of the book. If desired, Section 2.7 and Sections 2.10 through 2.12 can be skipped without affecting the readability of the subsequent chapters.

Chapters 3, 4, and 5 cover routing, regeneration, and wavelength assignment algorithms, respectively. Chapter 3 is equally applicable to O-E-O networks and optical-bypass-enabled networks; Chapters 4 and 5 are relevant only to the latter. The first three sections of Chapter 4 are more focused on physical-layer issues and can be skipped if desired.

Chapters 6 and 7 are standalone chapters on grooming and protection, respectively. Much of these chapters apply to both O-E-O networks and optical-bypass-enabled networks, with an emphasis on the latter. Finally, Chapter 8 presents numerous economic studies.

Acknowledgements

The nucleus of this book began as a Short Course taught at the Optical Fiber Communication (OFC) conference. I would like to thank the students for their suggestions and comments over the past five years that the course has been taught.

I am indebted to Dr. Adel Saleh, with respect to both my career and this book. As a leader of MONET, AT&T optical networking research, and Corvis, he is recognized as one of the foremost pioneers of optical networking. I appreciate the time he put into reading this book and his numerous helpful suggestions and encouragement.

I thank the editor of the Springer Optical Networks Series, Prof. Biswanath Mukherjee, for providing guidance and enabling a very smooth publication process. He provided many useful comments that improved the readability and utility of the book.

The team from Springer, Alex Greene and Katie Stanne, has been very professional and a pleasure to work with. Their promptness in responding to all my questions expedited the book.

Table of Contents

Chapter 1. Introduction to Optical Networks.....	1
1.1 Brief Evolution of Optical Networks.....	1
1.2 Geographic Hierarchy of Optical Networks	3
1.3 Layered Architectural Model	5
1.4 Interface to the Optical Layer	7
1.5 Configurable Optical Networks	8
1.6 Terminology	10
1.7 Network Design and Network Planning	14
1.8 Focus on Practical Optical Networks.....	15
 Chapter 2. Optical Network Elements.....	 17
2.1 Basic Optical Components.....	18
2.2 Optical Terminal.....	19
2.2.1 Slot Flexibility	20
2.3 Optical-Electrical-Optical (O-E-O) Architecture	22
2.3.1 O-E-O Architecture at Nodes of Degree-Two	22
2.3.2 O-E-O Architecture at Nodes of Degree-Three or Higher	23
2.3.3 Advantages of the O-E-O Architecture	25
2.3.4 Disadvantages of the O-E-O Architecture	26
2.4 OADMs (ROADMs).....	27
2.4.1 Configurability.....	29
2.4.2 Wavelength vs. Waveband Granularity	30
2.4.3 Wavelength Reuse	30
2.4.4 Automatic Power Equalization.....	32
2.4.5 Edge Configurability	32
2.4.6 Multicast	33
2.4.7 Slot Flexibility	33
2.4.8 East/West Separability.....	34
2.4.9 Broadcast-and-Select and Wavelength-Selective Architectures	34
2.5 Multi-Degree OADMs	36
2.5.1 Optical Terminal to OADM to OADM-MD Upgrade Path.....	40
2.5.2 First-Generation OADM-MD Technology	41
2.5.3 Second-Generation OADM-MD Technology.....	42
2.6 Optical Switches.....	44

2.6.1 O-E-O Optical Switch.....	44
2.6.2 Photonic Switch	46
2.6.3 All-Optical Switch	46
2.7 Hierarchical Switches.....	49
2.8 Adding Edge Configurability to a Node	50
2.8.1 Adjunct Edge Switch	51
2.8.2 Flexible Transponders.....	52
2.9 Optical Reach.....	53
2.10 Integrating WDM Transceivers in the Client Layer	55
2.11 Photonic Integrated Circuits.....	56
2.12 Multi-Fiber-Pair Systems	57
Chapter 3. Routing Algorithms	61
3.1 Shortest Path Algorithms.....	63
3.2 Routing Metrics	65
3.2.1 Fewest-Hops Path vs. Shortest-Distance Path	65
3.2.2 Shortest-Distance Path vs. Minimum-Regeneration Path.....	67
3.3 Generating a Set of Candidate Paths.....	68
3.3.1 K-Shortest Paths Strategy	68
3.3.2 Bottleneck-Avoidance Strategy	69
3.4 Routing Strategies	71
3.4.1 Fixed-Path Routing	71
3.4.2 Alternative-Path Routing.....	71
3.4.3 Dynamic-Path Routing	73
3.5 Avoiding Infeasible Paths	75
3.5.1 Capturing the Available Equipment in the Network Model.....	75
3.5.2 Predeployment of Equipment	78
3.6 Diverse Routing for Protection	79
3.6.1 Shortest Pair of Disjoint Paths.....	82
3.6.2 Shortest Pair of Disjoint Paths: Dual-Sources/Dual-Destinations.....	84
3.6.3 Shared Risk Link Groups (SRLGs).....	87
3.6.4 Routing Strategies With Protected Demands.....	91
3.7 Routing Order	92
3.8 Multicast Routing	93
3.8.1 Multicast Protection.....	97
3.9 Routing with Inaccurate Information.....	97
Chapter 4. Regeneration.....	101
4.1 Factors That Affect Regeneration	102
4.1.1 Optical Impairments	102
4.1.2 Mitigation of Optical Impairments	103
4.1.3 Network Element Effects.....	104
4.1.4 Transmission System Design.....	105
4.1.5 Fiber Plant Specifications	107
4.1.6 System Regeneration Rules	107

4.2 Routing with Noise Figure as the Link Metric	108
4.2.1 Network Element Noise Figure	110
4.2.2 Impact of the OADM without Wavelength Reuse.....	111
4.3 Link Engineering.....	113
4.3.1 Cohesive System Design	113
4.4 Regeneration Strategies.....	114
4.4.1 Islands of Transparency.....	114
4.4.2 Designated Regeneration Sites	116
4.4.3 Selective Regeneration	118
4.5 Regeneration Architectures	119
4.5.1 Back-to-Back WDM Transponders.....	120
4.5.2 Regenerator Cards	122
4.5.3 All-Optical Regenerators.....	125
Chapter 5. Wavelength Assignment.....	127
5.1 Role of Regeneration in Wavelength Assignment	128
5.2 Multi-Step RWA	130
5.2.1 Alleviating Wavelength Contention.....	131
5.3 One-Step RWA.....	132
5.4 Wavelength Assignment Strategies	136
5.4.1 First-Fit.....	137
5.4.2 Most-Used.....	139
5.4.3 Relative Capacity Loss	139
5.4.4 Qualitative Comparison.....	141
5.5 Subconnection Ordering.....	141
5.6 Bi-directional Wavelength Assignment.....	143
5.7 Wavelengths of Different Optical Reach.....	144
5.8 Wavelength Contention and Network Efficiency.....	146
5.8.1 Backbone Network Study	147
5.8.2 Metro Network Study	149
5.8.3 Study Conclusions	151
Chapter 6. Grooming.....	153
6.1 End-to-End Multiplexing	155
6.2 Grooming.....	157
6.3 Grooming Node Architecture.....	159
6.3.1 Grooming Switch at the Nodal Core	159
6.3.2 Grooming Switch at the Nodal Edge.....	160
6.4 Selection of Grooming Sites	163
6.5 Backhaul Strategies	164
6.6 Grooming Tradeoffs.....	166
6.6.1 Cost vs. Path Distance	166
6.6.2 Cost vs. Capacity	168
6.6.3 Grooming Design Guidelines	169
6.7 Grooming Strategies.....	169

6.7.1 Initial Bundling and Routing	169
6.7.2 Grooming Operations	170
6.8 Grooming Network Study	176
6.8.1 Grooming Switches at All Nodes	176
6.8.2 Grooming Switches at a Subset of the Nodes	178
6.9 Evolving Techniques for Grooming Bursty Traffic	179
6.9.1 Selective Randomized Load Balancing	180
6.9.2 Optical Flow Switching (OFS)	180
6.9.3 Optical Burst Switching (OBS)	181
6.9.4 TWIN	181
6.9.5 Lightrail	182
6.9.6 Optical Packet Switching (OPS)	182
Chapter 7. Optical Protection	183
7.1 Dedicated vs. Shared Protection	185
7.1.1 Dedicated Protection	186
7.1.2 Shared Protection	186
7.1.3 Comparison of Dedicated and Shared Protection	187
7.2 Client-Side vs. Network-Side Protection	190
7.3 Ring Protection vs. Mesh Protection	194
7.3.1 Ring Protection	194
7.3.2 Mesh Protection	198
7.4 Fault-Dependent vs. Fault-Independent Protection	200
7.5 Protection Against Multiple Concurrent Failures	205
7.6 Effect of Optical Amplifier Transients on Protection	207
7.7 Shared Protection Based on Predeployed Subconnections	209
7.7.1 Cost vs. Spare Capacity Tradeoff	211
7.8 Shared Protection Based on Pre-cross-connected Bandwidth	213
7.9 Protection Planning Algorithms	217
7.9.1 Algorithms for Dedicated Protection	217
7.9.2 Algorithms for Shared Protection	218
7.10 Protection of Substrate Demands	224
7.10.1 Wavelength-Level Protection	224
7.10.2 Substrate-Level Protection	226
7.10.3 Wavelength-Level vs. Substrate-Level Protection	228
7.10.4 Multilayer Protection	229
7.11 Fault Isolation	231
Chapter 8. Economic Studies	233
8.1 Assumptions	234
8.1.1 Reference Network Topology	234
8.1.2 Reference Traffic Set	235
8.1.3 Cost Assumptions	236
8.2 Prove-In Point for Optical-Bypass Technology	238
8.2.1 Comments on Comparing Costs	240

8.2.2 O-E-O Technology with Extended Optical Reach.....	240
8.3 IP Transport Architectures	242
8.4 Optimal Optical Reach	246
8.4.1 Add/Drop Percentage as a Function of Optical Reach	249
8.5 Architecture of Higher-Degree Nodes	251
8.6 Reduced-Reach Transponders.....	253
8.7 Optimal Topology from a Cost Perspective	256
8.8 Optimal Line-Rate.....	259
8.8.1 Study Assumptions	260
8.8.2 Study Results	262
8.8.3 Less Aggressive Cost Assumptions	264
8.9 Optical Grooming in the Edge Networks	265
8.10 General Conclusions	268
Appendix A. Suggestions for RFI/RFP Network Design Exercises.....	269
Appendix B. C-Code For Routing Routines.....	271
Abbreviations.....	293
Bibliography.....	297
Index	309

Chapter 1

Introduction to Optical Networks

1.1 Brief Evolution of Optical Networks

While the basic function of a network is quite simple – enabling communications between the desired endpoints – the underlying properties of a network can greatly affect its value. Network capacity, reliability, cost, scalability, and operational simplicity are some of the key benchmarks on which a network is evaluated. Network designers are often faced with tradeoffs among these factors, and are continually looking for technological advances that have the power to improve networking on a multitude of fronts.

One such watershed development came in the 1980s as the telecommunications industry began migrating much of the physical layer of their inter-city networks to fiber-optic cable. Optical fiber is a lightweight cable that provides low-loss transmission; but clearly its most significant benefit is its tremendous potential networking capacity. Not only did fiber optics open the possibilities of a huge vista for transmission, it also gave rise to optical networks and the field of optical networking.

An optical network is composed of the fiber-optic cables that carry channels of light, combined with the equipment deployed along the fiber to process the light. The capabilities of an optical network are necessarily tied to the physics of light and the technologies for manipulating lightstreams. As such, the evolution of optical networks has been marked with major paradigm shifts as exciting breakthrough technologies are developed.

One of the earliest technological advances was the ability to carry multiple channels of light on a single fiber-optic cable. Each lightstream, or wavelength¹,

¹ The term ‘wavelength’ is commonly used in two different contexts: first, it refers to a channel of light; second, it refers to the specific point in the spectrum of light where the channel is centered (e.g., 1550 nanometers). The context should be clear from its usage; however, when necessary, clarifying text is provided.

is carried at a different optical frequency and multiplexed (i.e., combined) onto a single fiber, giving rise to Wavelength Division Multiplexing (WDM). The earliest WDM systems supported fewer than ten wavelengths on a single fiber. Since 2000, this number has rapidly grown to over one hundred wavelengths per fiber, providing a tremendous growth in network capacity.

A key enabler of cost-effective WDM systems was the development of the Erbium Doped Fiber Amplifier (EDFA). Prior to the deployment of EDFAs, each wavelength on the fiber had to be individually regenerated at roughly 40-km intervals, using costly electronic equipment. The EDFA optically amplifies all of the wavelengths on a fiber at once, allowing optical signals to be transmitted on the order of 500 km before needing to be regenerated.

A more subtle innovation was the migration from an architecture where the optical network served simply as a collection of static pipes to one where it was viewed as another networking layer. In this optical networking paradigm, network functions such as routing and protection are supported at the granularity of a wavelength, which can be operationally very advantageous. A single wavelength may carry hundreds of circuits. If a failure occurs in a fiber cable, restoring service by processing individual wavelengths is operationally simpler than rerouting each circuit individually.

The benefits of scale provided by optical networking have been further accelerated by the increasing capacity of a single wavelength. In the mid 1990s, the maximum capacity of a wavelength was roughly 2.5 Gb/s (Gb/s is 10^9 bits/sec). This has ramped up to 10 Gb/s and 40 Gb/s, with much discussion regarding evolution to 100 Gb/s per wavelength, or higher.

Increased wavelength rate combined with a greater number of wavelengths per fiber has expanded the capacity of optical networks by several orders of magnitude over a period of 25 years. However, transmission capacity is only one important factor. Historically, the contents of each wavelength have undergone electronic processing at numerous points in the network. As networks exploded in size, this necessitated the use of a tremendous amount of electronic terminating and switching equipment, which presented challenges in cost, power consumption, heat dissipation, physical space, and maintenance.

This bottleneck was greatly reduced by the development of *optical-bypass* technology. This technology eliminates much of the required electronic processing and allows a signal to remain in the optical domain for all, or much, of its path from source to destination. Because optical technology can operate on a spectrum of wavelengths at once, and can operate on wavelengths largely independently of their data-rate, maintaining signals in the optical domain allows a significant amount of equipment to be removed from the network and provides a scalable trajectory for network growth.

Achieving optical bypass required advancements in areas such as optical amplification, optical switching, transmission formats, and techniques to counteract optical impairments. Commercialization of optical-bypass technology began in the mid-1990s, leading to its deployment in the networks of several major telecommunications carriers. While reducing the amount of electronic processing ad-

dressed many of the impediments to continued network growth, it also brought new challenges. Most notably, it required the development of new algorithms to assist in operating the network so that the full benefits of the technology could be attained. Overall, the advent of optical-bypass technology has transformed the architecture, operation, and economics of optical networks, all of which is covered in this book.

1.2 Geographic Hierarchy of Optical Networks

When considering the introduction of new networking technology, it can be useful to segment the network into multiple geographic tiers, with key differentiators among the tiers being the number of customers served, the required capacity, and the geographic extent. One such partitioning is shown in Fig. 1.1. (In this section, the standalone term ‘network’ refers to the network as a whole; when ‘network’ is used in combination with one of the tiers, e.g., ‘backbone network’, it refers to the portion of the overall network in that particular tier.)

At the edge of the network, closest to the end-users, is the *access* tier, which distributes/collects traffic to/from the customers of the network. Access networks generally serve tens to hundreds of customers and span a few kilometers. (One can further subdivide the access tier into business-access and residential-access, or into metro-access and rural-access.) The *metro-core* tier is responsible for aggregating the traffic from the access networks, and generally interconnects a number of telecommunications central offices or cable distribution head-end offices. A metro-core network aggregates the traffic of thousands of customers and spans tens to hundreds of kilometers.




Fig. 1.1 Networking hierarchy based on geography.

Moving up the hierarchy, multiple metro-core networks are interconnected via *regional* networks. A regional network carries the portion of the traffic that spans multiple metro-core areas, and is shared among hundreds-of-thousands of customers, with a geographic extent of several hundred to a thousand kilometers. Inter-regional traffic is carried by the *backbone* network². Backbone networks may be shared among millions of customers and typically span thousands of kilometers.

While other taxonomies may be used, the main point to be made is that the characteristics of a tier are important in selecting an appropriate technology. For example, whereas the backbone network requires optical transport systems with very large capacity over long distances, that same technology would not be appropriate for, nor would it be cost effective in, an access network.

As one moves closer to the network edge, the cost of a network in a particular tier is amortized over fewer end users, and is thus a more critical concern. Because of this difference in price sensitivity among the tiers, there is often a trend to deploy new technologies in the backbone network first. As the technology matures and achieves a lower price point, it gradually extends closer towards the edge. A good example of this trend is the deployment of WDM technology, as represented by the timeline in Fig. 1.2 (the costs and dates in this figure are only approximate).

Even as a technology permeates a network, the particular implementation may differ across tiers. For example, with respect to WDM technology, backbone networks generally have 80 to 160 wavelengths per fiber, regional networks have roughly 40 to 80 wavelengths per fiber, metro-core WDM networks have anywhere from 8 to 40 wavelengths per fiber, and access networks typically have no more than 8 wavelengths.




Fig. 1.2 As the cost of WDM infrastructure decreases over time, it is introduced closer to the network edge. (Adapted from [Sale98b].)

² Other common names for this tier are the long-haul network or the core network. These terms are used interchangeably throughout the book.

A similar pattern is emerging with the introduction of optical-bypass technology. Commercial deployment began in backbone networks in the 2000 time-frame, and has gradually spread closer to the network edge. The capabilities of optical-bypass-based systems are tailored to the particular network tier. For example, the distance a signal can be transmitted before it suffers severe degradation is a fundamental attribute of such systems. In backbone networks, technology is deployed where this distance is a few thousand kilometers; in metro-core networks, it is several hundred kilometers.

While optical networking is supported to varying degrees in the different tiers of the network, the architecture of access networks (especially residential access) is very distinct from that of the other portions of the network. For example, one type of access network is based on passive devices (i.e., the devices in the field do not require power); these systems, aptly named Passive Optical Networks (PONs), would not be appropriate for larger-scale networks. Because the topological characteristics, cost targets, and architectures of access networks are so different from the rest of the network, they are worthy of a book on their own; hence, access networks are not covered here. Detailed treatment of access technologies can be found in [Lin06]. Suffice it to say that as optics enters the access network, enabling the proliferation of high-bandwidth end-user applications, there will be increased pressure on the remainder of the network to scale accordingly.

It should be noted that there is a recent trend in the telecommunications industry to ‘blur the boundaries’ between the tiers. Carriers are looking for technology platforms that are flexible enough to be deployed in multiple tiers of the network, with unified network management and provisioning systems to simplify operations [ChSc07].

1.3 Layered Architectural Model

Another useful network stratification is illustrated by the three-layered architectural model of Fig. 1.3. At the top of this model is the applications layer, which includes all types of services, such as voice, video, and data. The intermediate layer encompasses multiplexing, transport, and switching based on electronic technology. For example, this layer includes Internet Protocol (IP) routers, Ethernet switches, Asynchronous Transfer Mode (ATM) switches, and Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH) switches. Each of these protocols has a particular method for partitioning data and moving the data from source to destination.

The payloads of the electronic layer are passed to the optical layer, where they are packed into wavelengths. In the model of interest, the optical layer is based on WDM technology and utilizes optical switches that are capable of dynamically routing wavelengths. Thus, the bottom tier of this particular model can also be referred to as the ‘configurable WDM layer’.




Fig. 1.3 Three-layered architectural model. In the systems of interest, the optical layer is based on WDM technology with configurable optical switches. (Adapted from [WASG96]. © 1996 IEEE)

From the viewpoint of the electronic layer, the wavelengths form a *virtual topology*. This concept is illustrated in Fig. 1.4 by a small network interconnecting five points. In Fig. 1.4(a), the solid lines represent fiber optic cables, or the physical topology, and the dotted lines represent the paths followed by two of the wavelengths. This arrangement of wavelengths produces the virtual topology shown in Fig. 1.4(b); i.e., this is the network topology as seen by the electrical layer. In contrast to the fixed physical topology, the virtual topology can be readily modified by reconfiguring the paths of the wavelengths.

Note that it is possible for the application layer to directly access the optical layer, as represented in Fig. 1.3 by the optical access services. This capability could be desirable, for example, to transfer very large streams of protocol-and-format-independent data. Because the electronic layers are bypassed, no particular protocol is imposed on the data. By transporting the service completely in the optical domain, the optical layer potentially provides what is known as *protocol and format transparency*. While such transparency has often been touted as another benefit of optical networking, thus far these services have not materialized in a major way in practical networks.




Fig. 1.4 In (a), the solid lines represent the physical fiber optic links and the dotted lines represent the paths of two routed wavelengths. The two wavelength paths create the virtual topology shown in (b), where the solid lines represent virtual links. The virtual topology can be modified by setting up different wavelength paths.

1.4 Interface to the Optical Layer

One difficulty with carrying services directly in wavelengths is that the network can be difficult to manage. Network operations can be simplified by using standard framing that adds overhead for management. For example, the SONET and SDH specifications define a standard framing format for optical transmission, where the frame includes overhead bytes for functionality such as performance monitoring, path trace, and Operations, Administration, and Maintenance (OAM) communication. SONET/SDH is commonly used as the interface to the optical layer; standards exist to map services such as ATM and IP into SONET/SDH frames. (In addition to using SONET/SDH for framing, it is often used for switching and multiplexing in the electronic domain, as was shown in Fig. 1.3.)

The SONET and SDH specifications are very closely related: SONET is the American National Standards Institute (ANSI) standard, whereas SDH is the International Telecommunication Union (ITU) standard. SONET defines a base signal with a rate of 51.84 Mb/s, called Synchronous Transport Signal level-1 or STS-1 (Mb/s is 10^6 bits/sec). Multiple STS-1 signals are multiplexed together to form higher rate signals, giving rise to the SONET rate hierarchy. For example, three STS-1 signals are multiplexed to form an STS-3 signal. The optical instantiation of a general STS-N signal is called Optical Carrier level-N, or OC-N. SDH is similar to SONET, although the framing format is somewhat different. The SDH base signal is defined as Synchronous Transport Module level-1, or STM-1, which has a rate equivalent to an STS-3. Some of the most commonly used SONET and SDH rates are shown in Table 1.1. The bit-rates shown in parentheses for some of the signals are the nominal rates commonly used in reference to these signals. For more details on SONET/SDH technology, see [Tek01, Gora02, Telc05].

Table 1.1 Commonly Used SONET/SDH Signal Rates

SONET Signal	SDH Signal	Bit-Rate
STS-1, OC-1	-	51.84 Mb/s
STS-3, OC-3	STM-1	155.52 Mb/s
STS-12, OC-12	STM-4	622.08 Mb/s
STS-48, OC-48	STM-16	2.49 Gb/s (2.5 Gb/s)
STS-192, OC-192	STM-64	9.95 Gb/s (10 Gb/s)
STS-768, OC-768	STM-256	39.81 Gb/s (40 Gb/s)

The SONET/SDH standards were initially developed in the 1980s with a focus on voice traffic, although features have been added to make them more suitable for data traffic. More recently, the ITU has developed a new architectural paradigm to better address the needs of optical networking, called the Optical Transport Network (OTN). The associated transport hierarchy and formats are defined in ITU standard G.709, with the basic frame called an Optical channel Transport Unit (OTU). The bit-rate of the OTU hierarchy is slightly higher than the SONET/SDH rates, with an OTU1 having a rate of 2.67 Gb/s, OTU2 at 10.71 Gb/s, and OTU3 at 43.02 Gb/s. Some of the relevant standards documents are [ITU01, ITU03].

Compared with SONET/SDH, OTN provides benefits such as more efficient multiplexing and switching of high-bandwidth services, enhanced monitoring capabilities, and stronger forward error correction (FEC). FEC allows bit errors picked up during signal transmission to be corrected when the signal is decoded. Enhanced FEC can be used to compensate for more severe transmission conditions. For example, it potentially allows more wavelengths to be multiplexed onto a single fiber, or allows a signal to remain in the optical domain for longer distances, which is important for optical-bypass systems.

OTN provides a uniform method of multiplexing a range of protocol types, essentially by providing a generic ‘digital wrapper’ for the payload. It is envisioned as a step towards network convergence, where carriers can support multiple services with a single network rather than deploying parallel networks. One of its main drivers for acceptance is it potentially offers carriers a managed, cost-effective means of adapting their networks to the growing demand for Ethernet services. OTN has gradually entered commercial applications; however, there is still a great deal of deployed legacy SONET/SDH-based equipment.

1.5 Configurable Optical Networks

Networks undergo continuous cycles of evolution where the requirements of the applications drive the development of innovative technology, and where improved technology encourages the development of more advanced applications. One such example is the automated configurability of an optical network.

The initial driver for automated configurability was the normal forecast uncertainty and churn that occur in a network (churn is the process of connections being established and then later taken down as the demand patterns change). It is difficult to forecast the precise endpoints and bandwidth of the traffic that will be carried in a network. Furthermore, while most traffic has historically been fairly static, with connection holding times on the order of months or longer, there is a subset of the traffic that has a much shorter lifetime, leading to network churn. In addition, as carriers move away from protection schemes where the backup paths are pre-established, reconfigurability³ is needed to dynamically create a new path for failure recovery.

Thus, it is necessary that the network be able to adapt to inaccurate forecasts, changing demand patterns, and network failures; moreover, it is desirable that the process be automated to eliminate the labor cost and potential errors involved with manual configuration. The infrastructure and distributed intelligence to enable automated reconfigurability are collectively known as the *control plane*. (This is in contrast to the typically centralized *management plane* that has historically been responsible for network operations such as fault management and security.)

Various organizations have developed standards in support of the control plane. For example, the ITU has developed the Automatically Switched Optical Networks (ASON) architecture, and the Internet Engineering Task Force (IETF) has developed the Generalized Multi-Protocol Label Switching (GMPLS) paradigm. These specifications include signaling protocols to automate control of the optical network and enable features such as discovery of the network topology and network resources, and connection establishment. Some of the relevant standards and specifications can be found in [SDIR04, Mann04, ITU06]. A more detailed discussion of the topic is provided in [BeRS03].

GMPLS includes three models for interacting with the optical layer: peer, overlay, and augmented. For concreteness, the discussion will focus on the interaction of the IP and optical layers, but the principles apply to other electronic layers. In the peer (or integrated) model, the IP and optical layers are treated as a single administrative domain, with IP routers having full knowledge of the optical topology. The IP routers can determine the entire end-to-end path of a connection including how it should be routed through the optical layer. In the overlay model, the IP and optical layers are treated as distinct domains, with no exchange of routing and topology information between them. The IP layer is essentially a *client* of the optical layer and requests bandwidth from the optical layer as needed. The augmented model is a hybrid approach where a limited amount of information is exchanged between layers.

Given the amount of information that needs to be shared in the peer model, and the potential trust issues between the layers (e.g., the IP and optical layers may be operated by different organizations), the overlay and augmented models are generally more favored by carriers. In the overlay model, which is more established, the boundary between the client and the optical layers is called the

³ The terms reconfigurability and configurability are used interchangeably in this book.

User-Network Interface (UNI). Signaling specifications for the UNI have been developed by the IETF as well as the Optical Internetworking Forum (OIF) [SDIR04, OIF04].

As these protocols for automated configurability have begun to make their way into carrier networks, the need to support more advanced dynamic services has emerged. In one flavor of dynamic service, the application requests a connection and requires that it be established very rapidly (e.g., in less than a second). For example, in large-scale distributed computing, there may be hundreds of computers that continually need to change their interconnection pattern as the computation evolves. In a second type of dynamic application, very-high-bandwidth transmission is periodically required but only for a short time. The need for the bandwidth is often known in advance, providing the opportunity to schedule the network resources as needed. One example of this is grid computing, which is a means of sharing distributed processing and data resources in order to achieve very high performance. This may require that huge datasets be disseminated to multiple locations in a very short period of time.

The stringent requirements of these applications will require the development of more advanced cross-layer bandwidth optimization, where the bandwidth allocation is dynamically optimized across multiple layers [EIMW06]. For example, the IP layer may automatically initiate a request for more bandwidth from the optical layer via the control plane. Additionally, more sophisticated provisioning protocols that can establish connections across multiple domains are also needed. (A domain is defined as an area of the network under the control of a single entity. The interface between domains is known as the External Network-Network Interface (E-NNI), whereas the interface between networks within a domain is the Internal NNI (I-NNI) [ITU06].)

This book focuses on the optical layer and does not consider topics such as cross-layer bandwidth management. While this approach is more in-line with the overlay model, the general network design principles discussed would need to be incorporated in any of the models.

1.6 Terminology

This section introduces some of the terminology that is used throughout the book. Refer to the small network shown in Fig. 1.5. The circles represent the network *nodes*. These are the points in the network that source/terminate and switch traffic. The lines interconnecting the nodes are referred to as *links*. While the links are depicted with just a single line, they typically are populated by one or more fiber-pairs, where each fiber in a pair carries traffic in just one direction. (It is possible to carry bi-directional traffic on a single fiber, but not common.) Optical amplifiers may be periodically located along each fiber, especially in regional and backbone networks. Sites that solely perform amplification are not considered nodes. The portion of a link that runs between two amplifier sites, or between a node and an amplifier site, is called a *span*.




Fig. 1.5 Nodes are represented by circles and links are represented by solid lines. Nodes A and B have a degree of two, Node C has a degree of three, and Node D has a degree of one.

A very important concept is that of *nodal degree*. The degree of a node is the number of links incident on that node. Thus, in the figure, Nodes A and B have a degree of two, Node C has a degree of three, and Node D has a degree of one. Nodal degree is very important in determining the type of equipment appropriate for a node.

The specific arrangement of nodes and links constitutes the network topology. Early networks were almost always based on ring topologies due to the simple restoration properties of rings. More recently, networks, especially those in the backbone, have migrated to more flexible *mesh* topologies. In mesh networks, the nodes are arbitrarily interconnected, with no specific routing pattern imposed on the traffic. In Fig. 1.1, the topologies in the metro-core tier are shown as rings, whereas the regional and backbone topologies are mesh. While it is possible to develop network design techniques that are specifically optimized for rings, the approach of this book is to present algorithms and design methodologies that are general enough to be used in any topology.

The *traffic* in the network is the collection of services that must be carried. The term *demand* is used to represent an individual traffic request. For the most part, demands are between two nodes and are bi-directionally symmetric. That is, if there is a traffic request from Node A to Node B, then there is equivalent traffic from Node B to Node A. In any one direction, the originating node is called the source and the terminating node the destination. In multicast applications, the demands have one source and multiple destinations; such demands are typically one-way only. (It is also possible to have demands with multiple sources and one or more destinations, but not common.)

The term *connection* is used to represent the path allocated through the network for carrying a demand. The process of deploying and configuring the equipment to support a demand is called *provisioning*, or *turning up*, the connection. The rate of a demand or a connection will be referred to in either absolute terms (e.g., 10 Gb/s) or using SONET terminology (e.g., OC-192), depending on the context.

The optical networks of interest in this book are based on WDM technology. Figure 1.6 shows the portion of the light spectrum where WDM systems are generally based, so chosen because of the relatively low fiber attenuation in this re-

gion (as shown in the figure, the loss is typically between 0.20 and 0.25 dB/km). This spectrum is broken into three regions: the conventional band or C-band; the long wavelength band or L-band; and the short wavelength band or S-band. Most WDM systems make use of the C-band, however, there has been expansion into the L and S bands to increase system capacity.

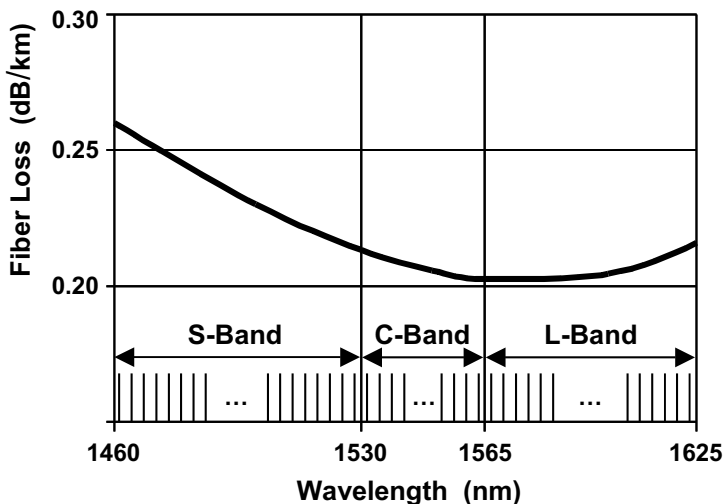


Fig. 1.6 Approximate S, C, and L wavelength bands, and the corresponding typical fiber loss.

An optical channel can be referred to as operating at a particular wavelength, in units of nanometers (nm), or equivalently at a particular optical frequency, in units of Terahertz (THz). The term *lambda* is frequently used to refer to the particular wavelength on which an optical channel is carried; *lambda* *i*, or λ_i , is used to represent the *i*th wavelength in the WDM system. The distance between adjacent channels in the spectrum is generally noted in frequency terms, in units of Gigahertz (GHz). For example, a 40-channel C-band system is achieved with 100-GHz spacing between channels, whereas an 80-channel C-band system is obtained using 50-GHz spacing.

An important piece of equipment is the WDM *transponder*, which is illustrated in Fig. 1.7(a). One side of the transponder is termed the *client side*, which takes a signal from the client of the optical network, e.g., an IP router. The client optical signal is generally carried on a 1310-nm wavelength. (1310-nm is outside the WDM region; WDM is usually not used for intra-office⁴ communication.) Various interfaces can be used on the client side of the transponder, depending on how much optical loss is encountered by the client signal. For example, *short-reach*

⁴ Office refers to a building that houses major pieces of telecommunications equipment, such as switches and client equipment.

interfaces tolerate up to 4 dB or 7 dB of loss depending on the signal rate, whereas *intermediate-reach interfaces* tolerate up to 11 dB or 12 dB of loss⁵. The interface converts the client optical signal to the electronic domain; the electronic signal modulates (i.e., drives) a WDM-compatible laser such that the client signal is converted to a particular wavelength (i.e., optical frequency) in the WDM region. The WDM side of the transponder is also called the *network side*. In the reverse direction, the WDM-compatible signal enters from the network side and is converted to a 1310-nm signal on the client side.

A single WDM transponder is shown in more detail in Fig. 1.7(b), to emphasize that there is a client-side receiver and a network-side transmitter in one direction and a network-side receiver and a client-side transmitter in the other direction. For simplicity, the transponder representation in Fig. 1.7(a) is used in the remainder of the book; however, it is important to keep in mind that a transponder encompasses separate devices in the two signal directions.




Fig. 1.7 (a) A simplified depiction of a WDM transponder that converts between a 1310-nm signal and a WDM-compatible signal. (b) A more detailed depiction of the WDM transponder that emphasizes its bi-directional composition. There is both a 1310-nm transmitter/receiver and a WDM-compatible transmitter/receiver.

In fixed-tuned transponders, the client signal can be converted to just one particular optical frequency. In transponders equipped with tunable lasers, the client signal can be converted to any one of a range of optical frequencies. Some architectures require that the transponder have an optical filter on the network side to receive a particular frequency. Tunable filters allow any one of a range of optical frequencies to be received. Since the early 2000s, most networks have been equipped with transponders with tunable lasers; transponders with both tunable lasers and filters were commercially available some time later. While they cost somewhat more, tunable transponders greatly improve the flexibility of the network as well as simplify the process of maintaining spare equipment for failure events.

The signal rate carried by a wavelength is called the *line-rate*. It is often the case that the clients of the optical network generate traffic that has a lower rate than the wavelength line-rate. This is referred to as *subrate* traffic. For example,

⁵ The loss increases with fiber distance and the number of fiber connectors; thus, these various types of interfaces determine the allowable interconnection arrangements within an office.

an IP router may generate 10 Gb/s signals but the line-rate may be 40 Gb/s. This mismatch gives rise to the need to *multiplex* or *groom* traffic, where multiple client signals are carried on a wavelength in order to improve the network efficiency. (End-to-end multiplexing bundles together subrate traffic with the same endpoints; grooming uses more complex aggregation than multiplexing and is thus more efficient, though more costly.) It is also possible, though less common, for the client signal rate to be higher than the wavelength line-rate. In this scenario, *inverse multiplexing* is used, where the client signal is carried over multiple wavelengths.

1.7 Network Design and Network Planning

As indicated by the title of the book, both network design and network planning are covered. Network design encompasses much of the up-front work such as selecting which nodes to include in the network, laying out the topology to interconnect the nodes, selecting what type of transmission and switching systems to deploy (e.g., selecting the line-rate and whether to use optical bypass), and what equipment to deploy at a particular node. Network planning is more focused on the details of how to accommodate the traffic that will be carried by the network. For example, network planning includes selecting how a particular demand should be routed, protected, and groomed, and what wavelength(s) in the system spectrum should be assigned to carry it.

Network planning is carried out on two time scales, both of which are covered in this book. In *long-term network planning*, there is sufficient time between the planning and provisioning processes such that any additional equipment required by the plan can be deployed. In the long-term planning that typically occurs before a network is deployed, there is generally a large set of demands to be processed at one time. In this context, the planning emphasis is on determining the optimal strategy for accommodating the whole traffic set. After the network is operational, long-term planning is performed for the traffic that does not need to be provisioned immediately; however, typically it is performed for a smaller number of demands at one time. Again, the focus is on determining optimal strategies, as there is enough time to deploy equipment to accommodate the design.

In *real-time network planning*, there is little time between planning and provisioning, and demands are generally processed one at a time. It is assumed that the traffic must be accommodated using whatever equipment is already deployed in the network. Thus, the planning process must take into account any constraints posed by the current state of deployed equipment, which, for example, may force a demand to be routed over a sub-optimal path. (A related topic is *traffic engineering*, which in this context is a process where traffic is controlled to meet specific performance objectives; e.g., a demand may be routed over a specific path to meet a particular availability metric, or a demand may be routed such that it avoids a heavily utilized region of the network. Traffic-engineering support for real-time routing has been incorporated in several protocols; e.g., see [ABGL01, KaKY03].)

1.8 Focus on Practical Optical Networks

This book examines the design and planning of state-of-the-art optical networks, with an emphasis on the ramifications of optical-bypass technology. It expands on the aspects of optical network design and planning that are relevant in a practical environment, as opposed to taking a more theoretical approach. Much research has focused on idealized optical-bypass systems where all intermediate electronic processing is removed; such networks are often referred to as ‘all-optical’. However, in reality, a small amount of intermediate electronic processing may still be required, for example, to improve the quality of the signal or to more efficiently pack data onto a wavelength. This small deviation from the idealized ‘all-optical’ network can have a significant impact on the network design, as is covered in later chapters. Thus, rather than use the term ‘all-optical network’, this book uses the term *‘optical-bypass-enabled network’*.

Many of the principles covered in the book are equally applicable in metro-core, regional, and backbone networks. However, it will be noted when there are significant differences in the application of the technology to a particular tier.

The foundation of today’s optical networks is the network elements; i.e., the major pieces of equipment deployed at a node. Chapter 2 discusses the various network elements in detail, with a focus on functionality and architectural implications. The underlying technology will be touched on only to the level that it affects the network architecture. Traditional network elements are covered as well as the new elements that enable optical bypass. From the discussion of the network elements, it will be apparent why algorithms play an important role in optical-bypass-enabled networks. Chapters 3 through 5 focus on the algorithms that are an integral part of operating an efficient and cost-effective optical network. The goal is not to cover all possible optical networking algorithms, but to focus on techniques that have proved useful in practice. Chapter 3, on routing algorithms, is equally applicable to optical-bypass-enabled networks as well as more traditional networks. Chapters 4 and 5, on regeneration and wavelength assignment, respectively, are relevant just to optical-bypass-enabled networks.

As mentioned earlier, treating the optical network as another networking layer can be very advantageous. However, networking at the wavelength level can potentially be at odds with operating an efficient network if the wavelengths are not well packed. Chapter 6 looks at efficient grooming of substrate demands, with an emphasis on various grooming architectures and methodologies that are compatible with optical bypass. It also considers strategies for grooming in the optical domain. Chapter 7 discusses protection in the optical layer. Rather than covering the myriad variations of optical protection, the discussion is centered on how protection in the optical layer is best implemented in a network with optical-bypass technology.

From the viewpoint of the network operator, perhaps the single most important characteristic of a network is its cost, both capital cost (i.e., equipment cost) and operating cost. Chapter 8 includes a range of economic studies that probe how

and when optical networking can improve the economics of a network. These studies can serve as a guideline for network architects planning a network evolution strategy, as well as equipment vendors analyzing the potential benefits of a new technology. The emphasis of the studies in this chapter, as well as the book as a whole, is on real-world networks. For general books on optical networking, see reference texts such as [StBa99, RaSi01, Mukh06].

Chapter 2

Optical Network Elements

The dramatic shift in the architecture of optical networks is chiefly due to the development of advanced optical network elements. These elements are based on the premise that the majority of the traffic that enters a node is being routed *through* the node en-route to its final destination as opposed to being *destined for* the node. This transiting traffic can potentially remain in the optical domain as it traverses the node rather than be electronically processed. By deploying technology that enables this so-called *optical bypass*, a significant reduction in the amount of required nodal electronic equipment can be realized.

After briefly discussing some basic optical components and switch terminology, Sections 2.2 and 2.3 review the traditional network architecture where all traffic entering a node is electronically processed. This architecture is based on a simple optical network element, the *Optical Terminal*. The economic and operational challenges of these legacy networks motivated the development of optical-bypass technology. The three major network elements that are capable of optical bypass are the *Optical Add/Drop Multiplexer* (OADM), the *Multi-Degree OADM* (OADM-MD), and the *All-Optical Switch*, all of which are discussed in Sections 2.4 through 2.7. These elements come in many flavors; the chief attributes that affect their efficiency, cost, and flexibility are covered in these sections.

The configurability of the network elements in response to changing traffic is one of the most important attributes. In some implementations, the network element does not provide sufficient configurability at the edge of the optical network (i.e., the interface between the client network and the optical network). Alternative methods of achieving the desired edge flexibility are presented in Section 2.8. Note that the terms ‘configurability’ and ‘reconfigurability’ are used interchangeably throughout the book.

While the network elements are a necessary ingredient for enabling optical bypass, they are only half of the story. Another key requirement is extended *optical reach*, which is the distance an optical signal can travel before it degrades to a level that necessitates it be ‘cleaned up’, or regenerated. The interplay of optical reach and optical-bypass-enabled elements is presented in Section 2.9.

Integration of elements or components within a node is a more recent development, motivated by the desire to eliminate individual components and reduce cost. There can be a range of levels of integration as illustrated by the discussions of Sections 2.10 and 2.11.

Throughout this chapter, it is implicitly assumed that there is one fiber-pair per link; e.g., a degree-two node has two incoming and two outgoing fibers. Due to the large capacity of current transmission systems, single-fiber-pair deployments are common. However, the last section of the chapter briefly addresses multi-fiber-pair scenarios.

Throughout this chapter, the focus is on the functionality of the network elements. The underlying technology of the elements is discussed only when it has a major impact on how a particular element is used.

2.1 Basic Optical Components

Some of the optical components that come into play throughout this chapter are discussed here. One very simple component is the *wavelength-independent optical splitter*, which is typically referred to as a *passive splitter*. A splitter has one input port and N output ports, where the input optical signal is sent to all of the output ports. Note that if the input is a WDM signal, then the output signals are WDM as well. In many splitter implementations, the input power level is split equally across the N output ports, such that each port receives $1/N$ of the original signal power level. This corresponds to a nominal optical loss of $10 \cdot \log_{10} N$, in units of decibels (dB). Roughly speaking, for every doubling of N , the optical loss increases by another 3 dB. It is also possible to design optical splitters where the power is split non-uniformly across the output ports so that some ports suffer lower loss than others.

The inverse device is called a *passive optical coupler* or *combiner*. This has N input ports and one output port, such that all of the inputs are combined into a single output signal. The input signals are usually at different optical frequencies to avoid interference when they are combined. The coupler losses are the same as for the splitter.

Another important component is the $1 \times N$ demultiplexer, which may be built using an *arrayed waveguide grating* (AWG) [Okam98, RaSi01, DoOk06]. (AWGs are also called *wavelength grating routers*.) For large N , the loss through an AWG is on the order of 4 dB to 6 dB. This device can be viewed as having one input port and N output ports. With this configuration, an input WDM signal is internally demultiplexed into its constituent wavelengths, and specific wavelengths are sent to each output port. In a common implementation, the number of output ports and the number of wavelengths in the WDM signal are the same, such that exactly one wavelength is sent to each output port. The inverse device is an $N \times 1$ multiplexer, with N input ports and one output port. Only specific wavelengths from each of the input ports are multiplexed together on the output port. Again, in

a common implementation where the number of input ports equals the number of wavelengths, just one wavelength from each input port is multiplexed in the outgoing WDM signal.

Throughout this chapter, various types of switches are mentioned. For the most part, the relevant details of the switch are presented when the application is discussed; however, it is advantageous to introduce some terminology here. First, there is a broad class of switches known as *optical switches*. Contrary to what the name implies, these switches do not necessarily have a switch fabric that operates on optical signals. (The switch fabric is the ‘guts’ of the switch, where the interconnection between the input and output ports is established.) Rather, the term ‘optical switch’ is used to indicate a switch where the ports operate on the granularity of a wavelength or a group of wavelengths.

Wavelength-selective is a term used to classify devices that are capable of treating each wavelength differently. A wavelength-selective optical switch of size $1 \times N$ is known as a *wavelength-selective switch* (WSS) [MMMT03, Maro05]. A WSS can direct any wavelength on the input port to any of the N output ports. A wavelength-selective optical switch of size $M \times M$ is known as a *wavelength-selective cross-connect* (WSXC) [KYJH06]. This device can direct any wavelength from any of the M input ports to any of the M output ports.

Micro-electro-mechanical-system (MEMS) technology [WuSF06] is often used to fabricate switches with an optical switch fabric. This technology essentially uses tiny movable mirrors to direct light from input ports to output ports. Note that an individual MEMS switching element is not wavelength selective; it simply switches whatever light is on the input port without picking out a particular wavelength. However, when combined with multiplexers and demultiplexers that couple the individual wavelengths of a WDM signal to the ports of the MEMS switch, the combination is wavelength-selective, capable of directing any wavelength to any output port.

Various optical switch configurations are discussed in the chapter.

2.2 Optical Terminal

In traditional optical network architectures, optical terminals are deployed at the endpoints of each link. Figure 2.1 illustrates a single optical terminal equipped with several WDM transponders. An optical terminal is typically depicted in figures as a trapezoid to capture its multiplexing/demultiplexing functionality; i.e., there are individual wavelengths on the client side of the terminal and a WDM signal on the network side. Unfortunately, a trapezoid is often used to specifically represent a $1 \times N$ AWG. While an optical terminal can be based on AWG technology, there are other options as well, some of which are discussed in Section 2.2.1. *Throughout this book, the trapezoid is used to represent a general optical terminal, not necessarily one based on a specific technology.*




Fig. 2.1 An optical terminal equipped with WDM transponders.

Tracing the flow from left to right in the figure, a client signal enters a transponder, which converts the signal to a WDM-compatible optical frequency. (Only one client is shown in the figure, an IP router.) The optical terminal multiplexes the signals from all of the transponders onto a single network fiber. In general, the transponders plugged into an optical terminal generate different optical frequencies; otherwise, the signals would interfere with each other after being multiplexed together by the terminal.

In the reverse direction, a WDM signal is carried by the network fiber into the optical terminal, where it is demultiplexed into its constituent frequencies. Each transponder receives its signal on a particular optical frequency and converts it to a client-compatible signal.

Recall from Fig. 1.7(b) in Section 1.6 that the WDM transponder encompasses both a client-side receiver/network-side transmitter in one direction and a network-side receiver/client-side transmitter in the other direction. Similarly, the optical terminal is composed of both a multiplexer and a demultiplexer. Note that it is possible for the network-side signal transmitted by a transponder to be at a different optical frequency than the network-side signal received by the transponder; however, in most scenarios the frequencies are the same.

2.2.1 Slot Flexibility

An optical terminal is deployed with equipment shelves in which the transponders are inserted. One figure of merit of an optical terminal is the density of the transponders on a shelf, where higher density is preferred as it requires less space. For example, if a shelf holds up to sixteen 10 Gb/s transponders, the shelf density is 160 Gb/s.

The flexibility of the individual slots in the transponder shelves is another important attribute. In the most flexible optical terminal architecture, any slot can

accommodate a transponder of any frequency. This ‘colorless’ slot architecture allows for the most efficient deployment of optical terminals; i.e., the number of slots deployed needs to be only as large as the maximum number of transponders required at the node (subject to the shelf granularity). This architecture also maximizes the benefits of tunable transponders for dynamic networking, as it allows a transponder to tune to a different frequency without needing to be manually moved to a different slot.

One implementation of a colorless optical terminal is based on a passive splitter and coupler. The received WDM signal is passively split and sent to each of the slots. The transponder receiver (on the network-side) is equipped with an optical filter to select the desired optical frequency; for maximum transponder flexibility, this filter should be tunable. In the reverse direction, the signals from the transponders are passively coupled together into a WDM signal; again, for maximum flexibility, the transponders should be equipped with tunable lasers. Because passive splitters and couplers can result in significant optical loss, this architecture often requires optical amplifiers to boost the signal level, especially if the number of supported slots is large.

Another flexible optical-terminal architecture is based on a $1 \times N$ wavelength-selective-switch (WSS), where the switch can be configured to direct any wavelength from the input WDM signal to any of the optical terminal slots. A separate $N \times 1$ WSS is used in the reverse direction to multiplex the signals. In this architecture, the transponder receiver does not need to have an optical filter because the wavelength selection is carried out in the WSS (i.e., the transponder is capable of receiving whatever optical frequency is directed to its slot). One drawback of this approach is the limited size of the WSS, which limits the number of supported slots on the terminal. Commercially available WSSs originally had a maximum size on the order of 1×8 , although they have been continually increasing in size. Another disadvantage is the cost of the WSS relative to passive splitters and couplers.

In contrast to these colorless optical terminals, there are also fixed optical terminals where each slot can accommodate a transponder of only one particular frequency. This type of terminal is often implemented using AWG technology. Again, the transponder receiver does not need an optical filter. Though relatively cost effective and low loss, this architecture can lead to inefficient shelf packing and ultimately higher cost in networks where the choice of optical frequencies is very important. A new shelf may need to be added to accommodate a desired frequency even though the shelves that are already deployed have available slots. This architecture also negates the automated configurability afforded by tunable transponders.

In an intermediary optical-terminal architecture, the WDM spectrum is partitioned into groups, and a particular slot can accommodate transponders only from one group [ChLH06]. This type of terminal can be architected with lower loss or cost than a completely colorless design, but it has only limited configurability.

2.3 Optical-Electrical-Optical (O-E-O) Architecture

2.3.1 O-E-O Architecture at Nodes of Degree-Two

The traditional (non-configurable) optical-terminal-based architecture for a node of degree-two is shown in Fig. 2.2. There are two network links incident on the node, where it is common to refer to the links as the ‘East’ and ‘West’ links (there is not necessarily a correspondence to the actual geography of the node). As shown in the figure, the node is equipped with two optical terminals arranged in a ‘back-to-back’ configuration. The architecture shown does not support automated reconfigurability. Connectivity is provided via a manual patch-panel; i.e., a panel where equipment within an office is connected via fiber cables to one side (typically in the back), and where short patch cables are used on the other side (typically in the front) to manually interconnect the equipment as desired. Providing automated reconfigurability is discussed in the next section in the context of higher-degree nodes.




Fig. 2.2 O-E-O architecture at a degree-two node (without automated reconfigurability). Nodal traffic is characterized as either add/drop traffic or through traffic. All traffic entering and exiting the node is processed by a transponder. Note that the through traffic can undergo wavelength conversion, as indicated by interconnected transponders of different wavelengths (e.g., the bottom pair of interconnected transponders converts the signal from λ_6 to λ_9 in the East-to-West direction).

Tracing the path from right to left, the WDM signal enters the East optical terminal from the East link. This WDM signal is demultiplexed into its constituent wavelengths, each of which is sent to a WDM transponder that converts it to a 1310-nm optical signal. (Recall that 1310-nm is the typical wavelength of the client-side optical signal.) At this point, it is important to distinguish two types of traffic with respect to the node. For one type of traffic, the node serves as the exit

point from the optical layer. This traffic ‘drops’⁶ from the optical layer and is sent to a higher layer (the higher layers, e.g., IP, are the clients of the optical layer). The other type of traffic is transiting the node en-route to its final destination. After this transiting traffic has been converted to a 1310-nm optical signal by its associated transponder, it is sent to a second transponder located on the West optical terminal. This transponder converts it back into a WDM-compatible signal, which is then multiplexed by the West optical terminal and sent out on the West link. There are also transponders on the West terminal for traffic that is being ‘added’ to the optical layer, from higher layers, that needs to be routed on the West link.

In the left to right direction of the figure, the operation is similar. Some of the traffic from the West link drops from the optical layer and some is sent out on the East link. Additionally, there are transponders on the East terminal for traffic that is added to the optical layer at this node that needs to be routed on the East link.

The traffic that is being added to or dropped from the optical layer at this node is termed *add/drop* traffic; the traffic that is transiting the node is called *through* traffic. Regardless of the traffic type, note that all of the traffic entering and exiting the node is processed by a WDM transponder. In the course of converting between a WDM-compatible optical signal and a client optical signal, the transponder processes the signal in the electronic domain. Thus, all traffic enters the node in the optical domain, is converted to the electronic domain, and is returned to the optical domain. This architecture, where all traffic undergoes optical-electrical-optical conversion, is referred to as the *O-E-O architecture*.

2.3.2 O-E-O Architecture at Nodes of Degree-Three or Higher

The O-E-O architecture readily extends to a node of degree greater than two. In general, a degree- N node will have N optical terminals. Figure 2.3 depicts a degree-three node equipped with three optical terminals, with the third link referred to as the ‘South’ link. The particular architecture shown does not support automated reconfigurability.

As with the degree-two node, all of the traffic entering a node, whether add/drop or through traffic, is processed by a transponder. The additional wrinkle with higher-degree nodes is that the through traffic has multiple possible path directions. For example, in the figure, traffic entering from the East could be directed to the West or to the South; the path is set by interconnecting a transponder on the East optical terminal to a transponder on the West or the South optical terminal, respectively. In many real-world implementations, the transponders are interconnected using a manual patch-panel. Modifying the through path of a connection requires that a technician manually rearrange the patch-panel, a process that is not conducive to rapid reconfiguration and is subject to operator error.

⁶ While ‘drop’ often has a negative connotation in telecommunication networks (e.g., dropped packets, dropped calls), its usage here simply means a signal is exiting from the optical layer.

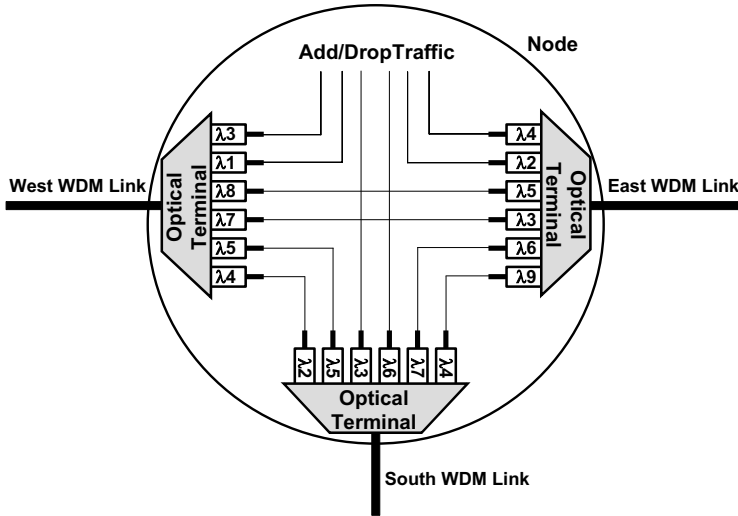


Fig. 2.3 O-E-O architecture at a degree-three node (without automated reconfigurability). There are three possible directions through the node. The path of a transiting connection is set by interconnecting a pair of transponders on the associated optical terminals.

The reconfiguration process can be automated through the addition of an optical switch, as shown in Fig. 2.4. Each transponder at the node feeds into a switch port, and the switch is configured as needed to interconnect two transponders to create a through path. Additionally, the add/drop signals are fed into ports on the switch, so that they can be directed to transponders on any of the optical terminals. Furthermore, the switch allows any transponder to be flexibly used for either add/drop or through traffic, depending on how the switch is configured. Note that the degree-two architecture of Fig. 2.2 could benefit from a switch as well with respect to these latter two applications; i.e., redirecting add/drop traffic to any optical terminal and flexibly using a transponder for either add/drop or through traffic.

While deploying a switch enhances the network flexibility and reduces operational costs due to less required manual intervention, the downside is the additional equipment cost. There are generally two types of optical switches that are used in such applications. Most commonly, a switch with an electronic switching fabric is used, where each port is equipped with a short-reach interface to convert the 1310-nm optical signal from the transponder to an electrical signal (this is the option shown in the figure). A second option is to use a switch with an optical switch fabric such as a MEMS-based switch. This technology can directly switch an optical signal, thereby obviating the need for short-reach interfaces on the switch ports (although it may require that the transponders be equipped with a special interface that is tolerable to the optical loss through the switch). As MEMS technology comes down in price, this type of switch may be the more cost-effective option. (Electronic switches and MEMS switches are revisited in Sections 2.6.1 and 2.6.2, respectively.)




Fig. 2.4 A switch is used to automate node reconfigurability. The particular switch shown has an electronic switch fabric and is equipped with short-reach interfaces on all of its ports.

2.3.3 Advantages of the O-E-O Architecture

The fact that the O-E-O architecture processes all traffic entering the node in the electronic domain does offer some advantages. First, converting the signal to the electronic domain and back to the optical domain ‘cleans up’ the signal. Optical signals undergo degradation as they are transmitted along a fiber. The O-E-O process re-amplifies, reshapes, and retimes the signal, a process that is known as *3R-regeneration*.

Second, it readily allows for performance monitoring of the signal. For example, assuming the SONET framing format is being used, a SONET-compatible transponder can examine the overhead bytes using electronic processing to determine if there are errors in the signal. Because this process can be done at every node, it is fairly straightforward to determine the location of a failure.

Third, the O-E-O architecture is very amenable to a multi-vendor architecture because all communication within a node is via a standard 1310-nm optical signal. Whereas the WDM transmission characteristics on a link may be proprietary to a vendor, the intra-nodal communication adheres to a well-defined standard. This allows the transmission systems on the links entering a node to be supplied by different vendors. Furthermore, the vendors of the switch (if present) and the transmission system can be different as well.

The O-E-O architecture also affords wavelength assignment independence to the traffic passing through a node. Here, the term wavelength is used to indicate a

particular frequency in the WDM spectrum. As noted above, the through traffic enters the node on one transponder and exits the node using a second transponder. These two transponders communicate via the 1310-nm optical signal; there is no requirement that the WDM-compatible wavelengths of these two transponders be the same. This is illustrated in Fig. 2.2 through Fig. 2.4, where the wavelengths of two interconnected transponders are not necessarily the same. This process accomplishes what is known as *wavelength conversion*, where the signal enters and exits the node on two different wavelengths. There is complete freedom in selecting the wavelengths of the two transponders, subject to the constraint that the same wavelength cannot be used more than once on any given fiber. The most important implication is that the choice of wavelength is local to a particular link; the wavelength assignment on one link does not affect that on any other link.

2.3.4 Disadvantages of the O-E-O Architecture

While the O-E-O architecture does have advantages, terminating every wavelength entering every node on a transponder poses many challenges in scaling this architecture to larger networks. For example, as capacities grow to over 100 wavelengths per fiber, this architecture could potentially require several hundred transponders at a node. The first barrier is the cost of all this equipment, although transponder costs do continue to decrease (see Section 2.11). Second, there are concerns regarding the physical space at the sites that house the equipment. More transponders translate to more shelves of equipment, and space is already at a premium in many carrier offices. Furthermore, providing the power for all of this electronics and dissipating the heat created by them is another operational challenge that will only worsen as networks continue to grow.

Another barrier to network evolution is the fact that electronics are often tied to a specific technology. For example, a short-reach interface that supports a 10-Gb/s signal typically does not also support a 40-Gb/s signal. Thus, if a carrier upgrades its network from a 10-Gb/s line-rate to a 40-Gb/s line-rate, a great deal of equipment needs to be replaced.

Provisioning a connection can be cumbersome in the O-E-O architecture. A technician may need to visit every node along the path of a connection to install the required transponders. Even if transponders are predeployed in a node, a visit to the node may be required, for example, to manually interconnect two transponders via a patch-panel. As noted above, manual intervention can be avoided through the use of a switch; however, as the node size increases, the switch size must grow accordingly. A switch, especially one based on electronics, will have its own scalability issues related to cost, size, power, and heat dissipation.

Finally, all of the equipment that must be deployed along a connection is potentially a reliability issue. The connection can fail, for example, if any of the transponders along its path fails.

2.4 OADMs (ROADMs)

The scalability challenge of the O-E-O architecture was a major impetus to develop alternative technology where much of electronics could be eliminated. Given that through traffic is converted from a WDM-compatible signal to a 1310-nm signal only to be immediately converted back again to a WDM-compatible signal, removing the need for transponders for the through traffic was a natural avenue to pursue.

This gave rise to the Optical Add/Drop Multiplexer (OADM) network element for nodes of degree two, as shown in Fig. 2.5. With an OADM, the through traffic remains in the optical domain as it transits the node; transponders are needed only for the add/drop traffic. The through traffic is said to *optically bypass* the node. Studies have shown that in typical carrier networks, on average, over 50% of the traffic entering a node is through traffic; thus, the amount of transponders that can be eliminated with optical bypass is significant. While the OADM itself costs more than two optical terminals, the reduction in transponders results in an overall lower nodal cost, assuming the level of traffic is high enough. The economics of optical bypass are explored further in Chapter 8.

OADMs have been commercially available since the mid 1990s, although significant deployment did not start until after 2000. The name of the element derives from a SONET/SDH Add/Drop Multiplexer (ADM), which is capable of adding/extracting lower-rate SONET/SDH signals to/from a higher-rate signal without terminating the entire higher-rate signal. Similarly, the OADM adds/extracts wavelengths to/from a fiber without having to electronically terminate all of the wavelengths comprising the WDM signal.




Fig. 2.5 OADM at a degree-two node. Transponders are required only for the add/drop traffic. The through traffic remains in the optical domain as it transits the node. (Adapted from [Simm05]. © 2005 IEEE)

Networks equipped with elements that support optical bypass are referred to here as *optical-bypass-enabled* networks. Other terms used in the literature to describe this type of network are ‘all-optical’ or ‘transparent’. However, given that O-E-O-based regenerators are typically not entirely eliminated from all end-to-end paths (see Section 2.9) and that networks supporting protocol-and-format transparency have not materialized in a major way (see Section 1.3), these terms are not completely accurate.

The advantages and disadvantages of the optical-bypass-enabled architecture are diametrically opposite to those that were discussed for the O-E-O architecture. First, as the network traffic level increases, optical-bypass technology is potentially more scalable in cost, space, power, and heat dissipation because much of the electronics is eliminated. Second, optics is more agnostic to the system bitrate as compared to electronics. For example, an OADM should typically function whether the wavelengths are carrying 10-Gb/s signals or 40-Gb/s signals (as long as the wavelength spacing and the signal spectrum are compatible with the OADM). Third, provisioning a connection is operationally simpler; in many scenarios, a new connection requires that a technician visit just the source and destination nodes to install the add/drop transponders. Fourth, the elimination of much of the electronics also improves the overall reliability; though the optical-bypass equipment itself may have higher failure rates than optical terminals, the removal of much of the equipment in the signal path typically leads to an overall lower failure rate for the connection [MaLe03].

Conversely, removing the transponders from the through path eliminates the functions that they provided. First, the optical signal of the through traffic is not regenerated, thereby requiring extended optical reach, as described in Section 2.9. Second, removing the transponders from some or all of the intermediate nodes of a path also eliminates the node-by-node error-checking functionality they provided. In the absence of electronic performance monitoring, optical monitoring techniques are needed, as discussed in Chapter 7. Third, it may be challenging to support a multi-vendor environment in an optical-bypass-enabled network because not all intra-nodal traffic is converted to a standard optical signal as it is in an O-E-O network. Standards for extended-reach WDM transmission have not been defined yet, such that interoperability between multiple vendors is not guaranteed. For example, the transmission system of one vendor may not be compatible with the OADM of another vendor. Furthermore, without standard performance guidelines, it may be difficult to isolate which vendor’s equipment is malfunctioning under a failure condition. Thus, in optical-bypass-enabled networks, it is common for a single vendor to provide both the transmission system and the optical networking elements. (The notion of ‘islands of transparency’, where designated vendors operate within non-overlapping subsets of the network, is discussed in Chapter 4.)

Finally, and most important, a transiting connection enters and exits the node on the same wavelength; there is not the same opportunity for wavelength conversion as with the O-E-O architecture. Given that two signals on the same fiber cannot be assigned the same wavelength, this implies that the assignment of wave-

lengths on one link potentially affects the assignment on other links in the network. This *wavelength continuity constraint* is the major reason why advanced algorithms are required to efficiently operate a network based on optical-bypass technology. Such algorithms are covered in Chapters 3, 4 and 5. (Note that in the future, technology may be commercially available that all-optically performs wavelength conversion. With this technology, a path that remains in the optical domain may not necessarily be carried on the same wavelength. However, all-optical wavelength converters represent an additional cost and are not likely to be deployed at all nodes on all wavelengths. Thus, the wavelength continuity constraint is likely to remain a relevant factor in optical-bypass-enabled networks.)

There are numerous important properties that characterize the various types of OADMs. These properties, which affect how the element is used in a network, are discussed in the following sections.

2.4.1 Configurability

One of the most important properties of an OADM is its degree of reconfigurability. The earliest commercial OADMs were not configurable. Carriers needed to specify up front which particular wavelengths would be added/dropped at a particular node, with all remaining wavelengths transiting the node. Once installed, the OADM was fixed in that configuration. Clearly, this rigidity limits the ability of the network to adapt to changing traffic patterns.

Today, however, most OADMs are configurable. This implies that any wavelength can be added/dropped at any node, and that the choice of add/drop wavelengths can be readily changed without impacting any of the other connections terminating at or transiting the node. Furthermore, it is highly desirable that the OADM be remotely configurable through software as opposed to requiring manual intervention. Such fully configurable OADMs are often called Reconfigurable OADMs, or ROADMs; however, there were fully configurable OADMs deployed in carrier networks prior to this term being coined.

One limitation of some OADMs is that they are fully configurable provided that the amount of add/drop does not exceed a given threshold. A typical threshold in such OADMs is a maximum of 50% of the wavelengths supportable on a fiber can be added/dropped (e.g., a maximum of 40 add/drops from a fiber that can support 80 wavelengths). For many nodes in a network, this is sufficient flexibility. However, there is typically a small subset of nodes that would ideally add/drop a higher percentage than this (this is examined further in Chapter 8). These nodes would be forced to employ the traditional O-E-O architecture. While the cost savings of using an OADM at a node that drops over 50% of the wavelengths may not be that significant, it is still desirable to have the option to use an OADM at such a node. An OADM would provide more agility than an O-E-O architecture with two optical terminals and a patch-panel. Moreover, OADMs that support up to 100% add/drop allow the flexibility of using an OADM at any node

without having to estimate the maximum percentage drop that will ever occur at the node.

2.4.2 Wavelength vs. Waveband Granularity

The previous section described configurability on a per-wavelength basis, where the choice of add/drop versus through can be made independently for each wavelength. Alternatively, OADMs can be fully configurable on the basis of a *waveband*. A waveband is a set of wavelengths that are treated as a single unit; either the whole waveband is added/dropped or the whole waveband transits the node. Wavebands are usually composed of wavelengths that are contiguous in the spectrum. In most implementations, the wavebands are of equal size; however, non-uniform waveband sizes may be more efficient depending on the traffic [IGKV03].

Clearly, waveband granularity is not as flexible as wavelength granularity. The chief motivation for using a waveband-based OADM is the potential for reduced cost and complexity. Waveband technology is more common in metro-core networks, where sensitivity to cost is greater.

Wavebands are most effective when many connections are being routed over the same paths in the network. Otherwise, inefficiencies can arise due to some of the bandwidth being ‘stranded’ in partially filled bands. Studies have shown that under reasonable traffic conditions, and through the use of intelligent algorithms, the inefficiencies resulting from wavebands are small [BuWW03]. Nevertheless, some carriers are averse to using waveband technology because of the somewhat diminished flexibility.

2.4.3 Wavelength Reuse

Another key OADM property that potentially affects network efficiency is whether or not the OADM supports *wavelength reuse*. With wavelength reuse, if a particular wavelength (i.e., optical frequency) is dropped at a node from one of the network fibers, then that same wavelength can be added at the node on the other network fiber. This is illustrated in Fig. 2.6(a). In the figure, a particular wavelength (λ_2) enters the node on the East fiber and is dropped. After being dropped, the wavelength does not continue to be routed through the OADM to the West fiber. This allows the node to add the same wavelength to the West fiber, as is shown in the figure (i.e., the add traffic is ‘reusing’ the same wavelength that was dropped). If the dropped wavelength had continued through the OADM, then traffic could not have been added to the West fiber on this wavelength because the two signals would interfere.




Fig. 2.6 (a) OADM with wavelength reuse. A wavelength (λ_2) is dropped from the East link. The OADM is configured to not pass this wavelength through the node, such that the same wavelength can be added on the West link. (b) OADM without wavelength reuse. The wavelength that is dropped from the East link continues to be routed through the node so that the same wavelength cannot be added on the West link.

Figure 2.6(b) illustrates an OADM that does not have wavelength reuse. The wavelength entering from the East network port is dropped but also continues through the OADM to the West fiber. This clearly wastes bandwidth, as this particular wavelength cannot be used to carry useful traffic on the West fiber (i.e., it is wasted bandwidth because it is carrying traffic that has already reached its destination). If there are multiple consecutive OADMs without reuse, the signal will continue to be routed through each one of them, preventing the wavelength from being reused to carry useful traffic on all of the intermediary links. (Note that the nodes on a ring cannot all be populated by no-reuse OADMs, as the optical signals will continue to wrap around.)

Wavelength reuse is a desirable trait in OADMs to maximize the useful capacity of the network. Nevertheless, OADMs without reuse can be useful network elements and are part of the commercial offerings of several system vendors. First, no-reuse OADMs are significantly lower cost than OADMs with reuse, making them a cost-effective option for nodes that drop a small amount of traffic.

Most vendors allow just a small percentage of the wavelengths on a fiber to be added/dropped with a no-reuse OADM; e.g., a maximum of 8 add/drops from a fiber with 80 wavelengths. This is in consonance with how the no-reuse OADM should be used. Additionally, the element is ideally deployed at nodes located on lightly loaded links, so that the wasted bandwidth is inconsequential.

In many implementations, a no-reuse OADM is little more than an optical amplifier equipped with a coupler and splitter to add and drop traffic. Because of this design, it is often possible to upgrade from an optical amplifier to a no-reuse OADM without affecting the traffic already passing through the amplifier. When a system is first installed, a network site that is not generating any traffic may be equipped with just an optical amplifier. As the network grows, the site may need to source/terminate traffic; when this occurs, the optical amplifier can be upgraded to the no-reuse OADM (upgrading from an optical amplifier to an OADM *with* reuse is generally not possible without a major overhaul of the equipment).

2.4.4 Automatic Power Equalization

In addition to network efficiency, there is often another important difference between OADMs with and without reuse. In many implementations, the technology that is used to achieve wavelength reuse can also be used to automatically equalize the power levels of the individual wavelengths that are transiting the node. The wavelengths comprising a WDM signal may have originated at different nodes, such that the power levels entering a node are unequal. Unbalanced power levels also result from uneven amplifier gain across the wavelength spectrum. The need for periodic power equalization is often another factor that comes into play when determining whether a no-reuse OADM is suitable for a node.

2.4.5 Edge Configurability

OADM configurability was discussed in Section 2.4.1 with respect to flexibly allowing any wavelength to bypass or drop at the node. Here, a different type of configurability is covered.

As discussed in Section 2.3.2, one of the advantages of deploying a switch with the O-E-O architecture is that the traffic from the higher layers (e.g., IP) can be automatically directed to any of the network links at the node. This flexibility, known as *edge configurability*, is especially useful in a dynamic environment, where connections are continually set up and torn down. Edge configurability allows any connection associated with a particular transponder to be automatically routed via any of the links at the node. It is also useful for protection in the optical layer, where at the time of failure, a connection can be sent out on an alternative path using the same transponder.

Edge configurability is possible with some OADM implementations, as is covered later in the chapter. This configurability is illustrated in Fig. 2.7, where the

transponder can access either the East link or the West link, depending on how the OADM is configured. Typically, fewer transponders need to be predeployed on an OADM that supports edge configurability because the transponders are not tied to a particular network link.

While having OADMs with edge flexibility is desirable, there are alternative means of achieving this configurability, as is discussed in Section 2.8.




Fig. 2.7 Edge configurability in an OADM, where one transponder can access either network link. In some implementations, one transponder can access both links simultaneously to provide optical multicast.

2.4.6 Multicast

Some OADMs can be configured to support optical multicast, where a given signal is sent to multiple destinations rather than just one, and the signal replication occurs in the optical domain as opposed to in the electronic domain. This can be accomplished by implementing *drop-and-continue*, where the intermediate destination OADMs are configured to both drop a wavelength and allow that same wavelength to pass through the node. The drop-and-continue concept was illustrated in Fig. 2.6(b), where λ_2 both dropped from and passed through the OADM. (Note that any OADM not capable of wavelength reuse is performing drop-and-continue for any dropped wavelength, whether it is desired or not.)

Additionally, some OADMs support a second means of optical multicast, where a single transponder can be configured to transmit simultaneously on the East port and the West port. This launches the same signal on two different network links. (This feature is also beneficial for optical-layer protection with an active backup path, as is discussed further in Chapter 7.)

2.4.7 Slot Flexibility

The discussion regarding the benefits of slot flexibility for optical terminals (Section 2.2.1) holds for the OADM add/drop slots as well. Because of the wave-

length continuity constraint in optical-bypass-enabled systems, it is often important to use a specific wavelength to carry a given connection. It is desirable to be able to insert the corresponding transponder in any add/drop slot of the OADM; i.e., colorless slots.

2.4.8 East/West Separability

East/West separability relates to the failure and repair modes of the OADM. If one side of the OADM fails, it is desirable that any add/drop traffic on that side be able to be directed to the other side for protection. It would be undesirable to have an architecture where, for example, a failure of the East link causes traffic on the West link to be shut down as well. Furthermore, if the East add/drop port fails, it would be undesirable if the process of repairing the port also requires that the West add/drop port be taken down.

2.4.9 Broadcast-and-Select and Wavelength-Selective Architectures

In this section, two common architectures for building an OADM are presented at a high level. *Broadcast-and-select* is a prevalent OADM architecture as it is suitable for optically bypassing several consecutive OADMs [BSAL02]. One common broadcast-and-select implementation is shown in Fig. 2.8. In this implementation, as the WDM signal enters from the East network fiber, some of the signal power is tapped off and directed to the East add/drop port. The dropped WDM signal is demultiplexed, and transponders are deployed only for those wavelengths that need to be dropped at the node (i.e., the entire WDM signal is ‘broadcast’ to the port, and transponders are deployed to each ‘select’ one particular wavelength). The add/drop port is a ‘*multi-wavelength port*’, as the WDM signal is dropped there. As mentioned above, the add/drop slots are ideally colorless.

Only a relatively small portion of the signal power is tapped off on the add/drop port, say 10%, which is a major reason the architecture is suitable for optical bypass. The remainder of the optical signal power continues on through the OADM to the West fiber. In an OADM with wavelength re-use, the technology of the OADM is capable of blocking any wavelengths that have been dropped so that they do not continue to the West fiber. (Examples of such technology are covered in Section 2.5.2 and 2.5.3.) Wavelengths can be added from the West add/drop port to be multiplexed with the signals passing through the OADM (assuming there are no lambda conflicts). The operation in the West to East direction is similar.

Depending on the technology that is used, the broadcast-and-select OADM is capable of any of the properties discussed in the previous sections. However, the specific implementation shown in Fig. 2.8, with the add/drop ports being tapped off from the network fibers, does not provide edge configurability. In this figure,

note that the transponders on the East add/drop port can only enter/exit from the East link; they cannot access the West link. (Other broadcast-and-select OADM implementations do provide edge configurability, however, as is discussed in Section 2.6.3.)




Fig. 2.8 One example of a broadcast-and-select OADM architecture. This particular implementation does not support edge configurability. The trapezoids represent general colorless or fixed add/drop ports.

An alternative to broadcast-and-select is the *wavelength-selective* architecture, an implementation of which is shown in Fig. 2.9. The WDM signal from the East network fiber is demultiplexed and the constituent wavelengths are each fed into a port on an optical switch. The switch is configured to direct some of the wavelengths to the add/drop ports on the switch; these add/drop ports are *single-wavelength* ports. The remaining wavelengths are directed to the West side, where they are multiplexed together, along with any add traffic, and sent out on the West network fiber. The operation in the West to East direction is similar.




Fig. 2.9 Wavelength-selective OADM architecture.

The switch size in this wavelength-selective architecture is $M \times M$, where M must be large enough to accommodate all of the wavelengths on the nodal fibers as well as all of the add/drop wavelengths. For example, if there are 40 wavelengths per fiber, and the OADM allows 100% drop from each fiber, then the switch must be of size 160×160 . The fabric of the switch must be optical so that there is no need for O-E-O conversion in the node. MEMS technology is often used in such OADM. (As noted in Section 2.1, the MEMS elements themselves are not wavelength-selective; however when combined with multiplexers and demultiplexers as shown, the combination is a wavelength-selective architecture. To be more specific, the device shown in Fig. 2.9 is a WSXC.)

This type of OADM provides wavelength reuse, as dropped signals cannot also pass through the node. In fact, this architecture does not readily support drop-and-continue for optical multicast. In contrast, the broadcast-and-select architecture can support drop-and-continue simply by not blocking a wavelength that has been dropped.

The wavelength-selective OADM architecture is also compatible with any of the other properties described above (including waveband granularity [KZJP04]). As shown in the figure, edge configurability is possible because any transponder can access either network fiber by reconfiguring the switch. One potential issue with the architecture is the cost of the switch. Furthermore, for reliability reasons, the OADM is often deployed with a spare switching fabric, which adds to the cost. Another possible downside to this architecture is that it is not as suitable as the broadcast-and-select architecture for optically bypassing multiple consecutive OADMs because of higher loss on the through path [TzZT03].

2.5 Multi-Degree OADMs

The OADM network element was designed to provide optical bypass at nodes of degree two. While all nodes in a ring architecture have degree two, a significant number of nodes in interconnected-ring and mesh topologies have a degree greater than two. Table 2.1 shows the percentage of nodes of a given degree averaged over several typical United States (U.S.) mesh backbone networks. Table 2.2 and Table 2.3 show the nodal-degree percentage averaged over several U.S. metro-core networks for interconnected-ring and mesh topologies, respectively.

Table 2.1 Nodal Degree Percentage Averaged over Several Typical U.S. Backbone Networks

Nodal Degree	Percentage
2	55%
3	35%
4	7%
5	2%
6	1%

Table 2.2 Nodal Degree Percentage Averaged over Several Typical U.S. Metro-Core Networks with Interconnected-Ring Topologies

Nodal Degree	Percentage
2	85%
4	9%
6	4%
8	2%

Table 2.3 Nodal Degree Percentage Averaged over Several U.S. Metro-Core Networks with Mesh Topologies

Nodal Degree	Percentage
2	30%
3	25%
4	15%
5	15%
6	10%
7	3%
8	2%

The question is what type of equipment to deploy at nodes of degree three or higher. One thought is to continue using the optical-terminal-based O-E-O architecture at these nodes, while using OADMs at degree-two nodes. This is sometimes referred to as the ‘O-E-O-at-the-hubs’ architecture, where nodes of degree three or more are considered hubs. As all traffic must be regenerated at the hubs, it implies that electronic performance monitoring can be performed at the junction sites of the network, which may be advantageous for localizing faults. However, it also implies that the scalability issues imposed by O-E-O technology will still exist at a large percentage of the nodes.

Another option is to deploy OADMs, possibly in conjunction with an optical terminal, at the hubs. Figure 2.10(a) depicts a degree-three node equipped with one OADM and one optical terminal. Optical bypass is possible only for traffic transiting between the East and West links. Traffic transiting between the South and East links or the South and West links must undergo O-E-O conversion via transponders, as shown in the figure. Figure 2.10(b) depicts a degree-four node equipped with two OADMs. Optical bypass is supported between the East and South links, and between the North and West links, but not between any other link pairs. The design strategy with these quasi-optical-bypass architectures is to deploy the OADM(s) in the direction where the most transiting traffic is expected. However, if the actual traffic turns out to be very different from the forecast, then there may be an unexpectedly large amount of regeneration; i.e., the architectures of Fig. 2.10 are not ‘forecast-tolerant’.




Fig. 2.10 (a) Degree-three node with one OADM and one optical terminal. (b) Degree-four node with two OADM. In these quasi-optical-bypass architectures, some of the transponders are used for transiting traffic that crosses two different network elements at a node.

With either of the above strategies, there is potentially still a large amount of electronics needed for transiting traffic, as is explored quantitatively in Chapter 8. Another alternative is to deploy the network element known as the multi-degree OADM (OADM-MD), which extends the functionality of an OADM to higher-degree nodes. A degree-three OADM-MD, with a broadcast-and-select architecture, is shown in Fig. 2.11. With an OADM-MD, optical bypass is supported in *all* directions through the node to maximize the amount of transponders that can be eliminated. As with OADM, transponders are needed only for the add/drop traffic. All of the properties that were discussed for the OADM are relevant for this element as well. The economic benefits of an OADM-MD relative to the ‘O-E-O-at-the-hubs’ and the ‘OADM-only’ architectures are studied in Chapter 8.

With the combination of the OADM and the OADM-MD, optical bypass can be provided in a network of arbitrary topology (subject to the maximum degree of the OADM-MD, which is discussed later in this section). For example, in Fig. 2.12(a), a degree-six OADM-MD deployed in the node at the junction of the three rings allows traffic to pass all-optically from one ring to another; the remainder of

the nodes have an OADM. In the arbitrary mesh of Fig. 2.12(b), a combination of OADMs, degree-three OADM-MDs, and degree-four OADM-MDs is deployed to provide optical bypass in any direction through any node.




Fig. 2.11 Degree-three OADM-MD. Optical bypass is possible in all three directions through the node. Transponders are needed only for the add/drop traffic. The trapezoids represent general colorless or fixed add/drop ports. (Adapted from [Simm05]. © 2005 IEEE)

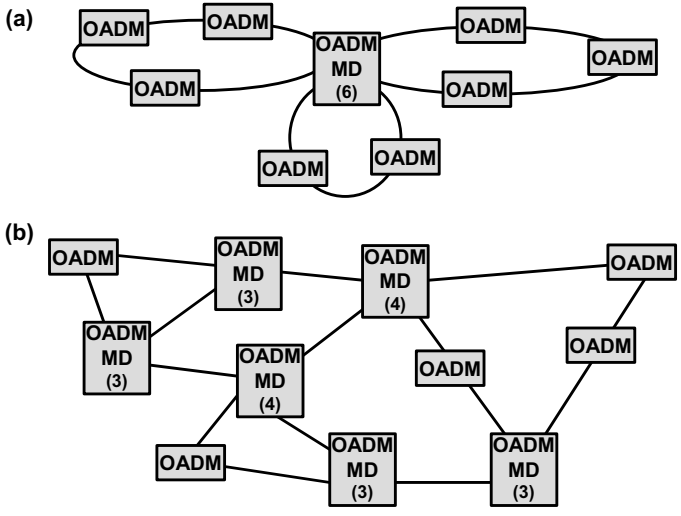


Fig. 2.12 (a) A degree-6 OADM-MD is deployed at the junction site of three rings, allowing traffic to transit all-optically between rings. The remaining nodes have OADMs. (b) In this arbitrary mesh topology, a combination of OADMs, degree-3 OADM-MDs, and degree-4 OADM-MDs are deployed according to the nodal degree. Optical bypass is supported in all directions through any node.

2.5.1 Optical Terminal to OADM to OADM-MD Upgrade Path

OADM-MDs can be part of a graceful network growth scenario. Carriers may choose to roll-out optical-bypass technology in stages in the network, where initially they deploy optical-bypass elements on just a few links to gradually grow their network. Consider deploying optical-bypass technology in just the small portion of the network shown in Fig. 2.13(a). An OADM would be deployed at Node B to allow bypass, with optical terminals deployed at Nodes A and C. As a next step, assume that the carrier wishes to extend optical bypass to the ring shown in Fig. 2.13(b). Ideally, the optical terminals at Nodes A and C are in-service upgradeable (i.e., existing traffic is not affected) to OADMs, so that all five nodes in the ring have an OADM. In the next phase, shown in Fig. 2.13(c), a link is added between Nodes A and E to enhance network connectivity. It is desirable that the OADMs at Nodes A and E be in-service upgradeable to a degree-3 OADM-MD. This element upgrade path, from optical terminal to OADM to OADM-MD, is desirable for network growth, and is supported by several commercial offerings. Furthermore, upgrading to higher-degree OADM-MDs is typically possible, up to the limit of the technology, as is discussed next.

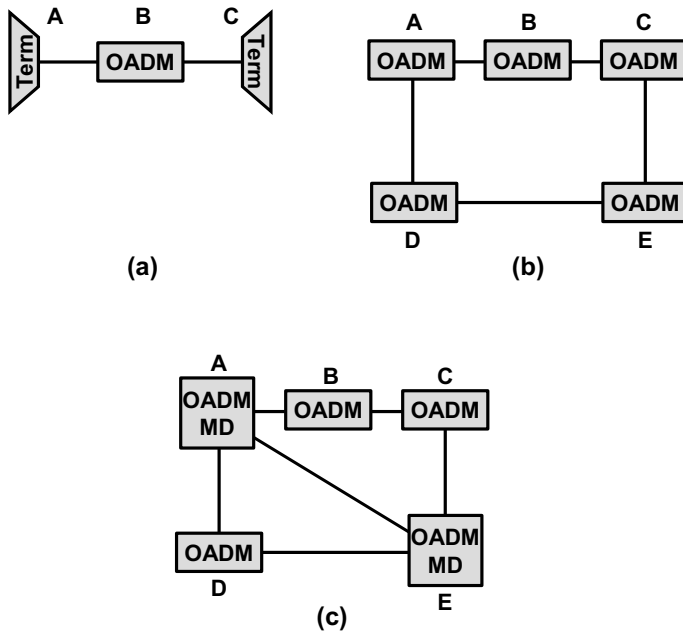


Fig. 2.13 In this network evolution, Node A is equipped with an optical terminal in (a), an OADM in (b), and a degree-three OADM-MD in (c). Ideally, these upgrades are performed in-service, without affecting any existing traffic at the node.

2.5.2 First-Generation OADM-MD Technology

The maximum degree that can be supported with an OADM-MD is very dependent on the underlying technology. The first generation of OADM-MD was built using liquid crystal wavelength blockers with a broadcast-and-select architecture, as shown in Fig. 2.14 [Tomk02, PCHN03]. The liquid crystal blockers are also called ‘dynamic spectral equalizers’ (DSEs) or ‘dynamic channel equalizers’ (DCEs); the term DSE is used here. This figure is one implementation of the OADM-MD represented in Fig. 2.11. In order to illustrate more clearly the technology in Fig. 2.14, the input network links are shown on the left of the figure and the output network links are shown on the right. The transponder receivers are shown on the left of the figure and the transponder transmitters are shown on the right.

Consider a degree- D OADM-MD, with D being three in Fig. 2.14 (for ease of discussion, D is also assumed to be the degree of the node). The signal from each incoming network fiber is initially split into a drop path and a through path. The through path is further split into $D-1$ paths, corresponding to the $D-1$ other network fibers, with a DSE along each of these through paths. (If port loopback is desired, then the through path is split into D paths.) There is typically amplification in the nodes, which can make up for the splitting loss. The DSEs can be dynamically configured to selectively pass or block any wavelength, depending on whether the wavelength is optically bypassing or being dropped at the node, respectively.

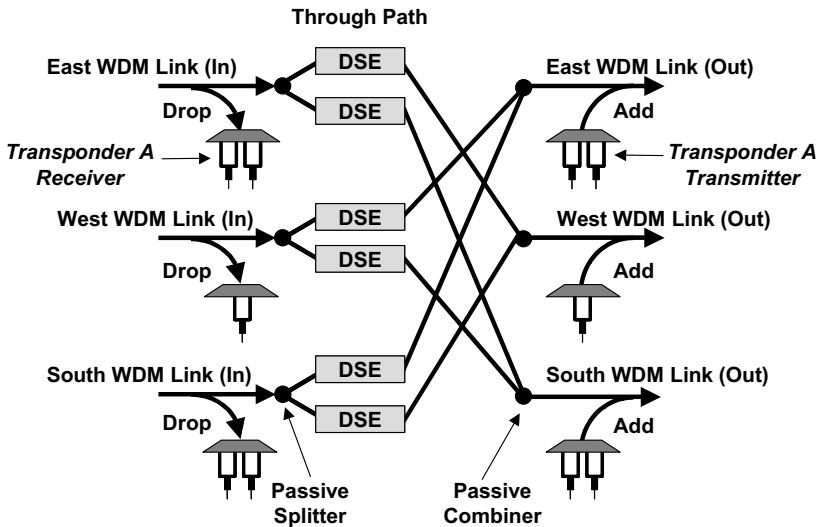


Fig. 2.14 First-generation degree-three OADM-MD built with Dynamic Spectral Equalizers (DSEs). The number of required DSEs grows quadratically with the degree of the OADM-MD. The trapezoids represent general colorless or fixed add/drop ports.

Consider a wavelength that enters from the East link that is being dropped at the node. The two DSEs associated with the incoming East link (i.e., the top two DSEs in the figure) would be configured to block this wavelength so that it is not routed to either the West or the South links. The DSE technology thus provides wavelength reuse. Next, consider a wavelength that transits the node from the East link to the South link. The upper DSE blocks the wavelength from being routed to the West link, while the second DSE allows the wavelength to pass through to the South link.

Optical multicasting on the network fibers is supported by this architecture. For example, if both DSEs associated with the incoming East link allow a wavelength to pass, then the signal is sent to both the West and South links. Furthermore, drop-and-continue is supported. For example, it is possible for a signal to drop from the East link and be transmitted to either, or both of, the West and South links.

This architecture does not support edge configurability. For example, the transponder labeled ‘A’ in the figure can add/drop only to/from the East link. Moreover, this architecture does not support the form of optical multiplexing where a sourced optical signal is sent out on multiple links.

A total of $D*(D-1)$ DSEs are needed in this architecture; thus, the size and cost of the switch fabric scale quadratically with the degree D . (D^2 DSEs are needed if port loopback is desired.) In practice, this architecture generally is limited to a maximum degree of four.

2.5.3 Second-Generation OADM-MD Technology

Second-generation OADM-MDs, also based on the broadcast-and-select architecture, make use of $N \times 1$ wavelength selective switches (WSS) as shown in Fig. 2.15 [MMMT03]. For a degree- D OADM-MD, N equals $D-1$, unless port loopback is desired, in which case N equals D . This OADM-MD is functionally equivalent to the first-generation device; e.g., it supports wavelength reuse and optical multicast. As with the DSE, the WSS can be configured to selectively pass or block a particular wavelength. The difference is that one WSS operates on N incoming paths. Thus, a total of D WSSs are needed for a degree- D OADM-MD; i.e., the amount of equipment scales linearly with D , not quadratically. (The required size of the WSS increases with the degree; however, the cost of the WSS increases sublinearly with N .) The improvement in scalability is especially useful in metro-core networks where multiple rings may be interconnected at a single node. For example, a degree-eight OADM-MD deployed at the intersection of four rings allows all intra-ring and inter-ring transiting traffic to remain in the optical domain.

In a variation of this OADM-MD architecture, each of the N -way passive splitters on the input side in Fig. 2.15 is replaced by a $1 \times N$ WSS. The chief motivation for this design is to decrease the overall loss through the device, especially at high-degree nodes where the passive splitting loss is large. The disadvantages are the need for twice as many WSSs and the lack of support for multicast.

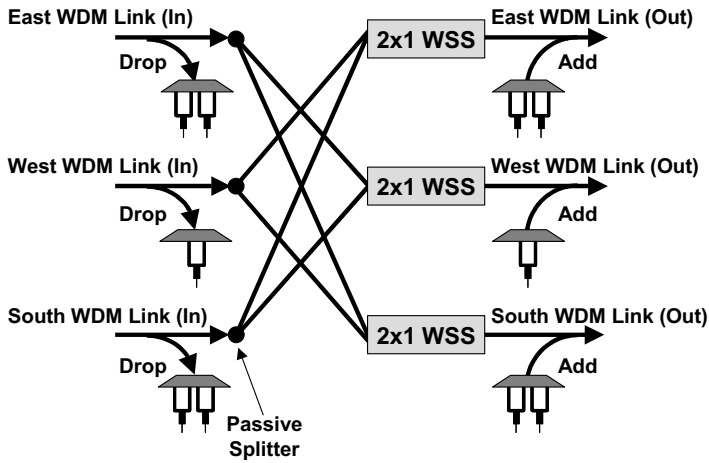


Fig. 2.15 Second-generation degree-three OADM-MD, based on $N \times 1$ Wavelength Selective Switches (WSSs). The number of required WSSs grows linearly with the degree of the OADM-MD. The trapezoids represent general colorless or fixed add/drop ports.

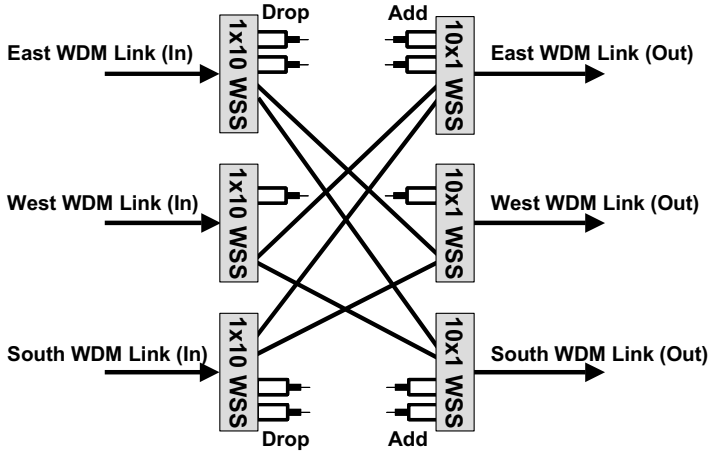


Fig. 2.16 An OADM-MD architecture based on $N \times 1$ WSSs, where some of the WSS ports are used for single-wavelength colorless add/drop. With the configuration shown, the maximum number of add/drops per fiber is eight (but only one or two add/drops per fiber are shown).

A third OADM-MD configuration based on $N \times 1$ WSSs is shown in Fig. 2.16. Again, the passive splitters of Fig. 2.15 are replaced by WSSs; in addition, some of the ports on the WSSs are used as single-wavelength add/drop ports. This provides colorless add/drop; i.e., any λ can be added/dropped to/from any port. The amount of add/drop is limited by the port size of the WSS; thus, this configuration may be more appropriate for the metro-core area where the amount of

add/drop at a node is more likely to be small. With the 1×10 WSSs shown in Fig. 2.16 for a degree-three node, the maximum number of add/drops is eight. This architecture does not support multicast.

(In contrast, the architecture of Fig. 2.15 has multi-wavelength add/drop terminals. As discussed in Section 2.2.1, multi-wavelength add/drop with colorless slots is also possible, depending on the technology used for the ‘trapezoids’. For example, one method is to use a passive coupler/splitter for the add/drop. However, the configuration of Fig. 2.16 likely yields a lower-loss add/drop path. A second method of providing colorless terminals in Fig. 2.15 is to deploy a $1 \times S$ WSS at the drop and a $S \times 1$ WSS at the add, where S is the maximum number of add/drop wavelengths at the node. This is closer to the configuration of Fig. 2.16, except a $1 \times S$ WSS and a $1 \times (D-1)$ splitter are used in place of one $1 \times (S+D-1)$ WSS on the input fiber, and a $S \times 1$ and a $(D-1) \times 1$ WSS are used instead of one $(S+D-1) \times 1$ WSS on the output fiber, where D is the OADM-MD degree.)

As with the first-generation architecture, the OADM-MDs of Fig. 2.15 and Fig. 2.16 do not provide edge configurability. In fact, when the term OADM-MD was first coined, it was specifically intended to denote a device that did *not* have edge configurability. The term ‘all-optical switch’ is the proper classification of a higher-degree optical-bypass element that can switch all nodal traffic, including the add/drop traffic, thereby providing edge configurability. These types of switches are covered in the next section. As a historical note, all-optical switches actually existed prior to the development of the OADM-MD; however, they tended to be high-cost. Optical bypass was not introduced on a large scale at the hub nodes of carrier networks until the lower-cost OADM-MD was developed. (The industry sometimes blurs the difference between OADM-MDs and all-optical switches and uses the terms interchangeably; however, the intent was that they be distinguished based on whether or not they provide edge configurability.)

2.6 Optical Switches

The all-optical switch is one member of a broader class of switches called *optical switches*. Optical switches are also referred to as optical cross-connects. As discussed in Section 2.1, the term optical switch indicates that the switch ports operate on the granularity of a wavelength or a waveband. It does not imply that the switch supports optical bypass, nor does it imply that the switch fabric is optical. There are several flavors of optical switches as described below. For more details of the underlying optical switch technologies, see [Sala02, PaPP03, ElBa06].

2.6.1 O-E-O Optical Switch

An optical switch based on O-E-O technology is shown in Fig. 2.17(a) (it also was shown as part of Fig. 2.4). The switch fabric is electronic, and each of the switch ports is equipped with a short-reach interface to convert the incoming 1310-nm

optical signal to an electronic signal. These types of switches present scalability challenges in cost, power, and heat dissipation due to the amount of electronics. Consider using such a switch to provide configurability at a degree-four O-E-O node. Assume that each fiber carries 160 wavelengths and assume that the node needs to support 50% add/drop. There needs to be a port for each wavelength on the nodal fibers as well as each add/drop wavelength. This requires a 960x960 switch, with each switch port having a short-reach interface.

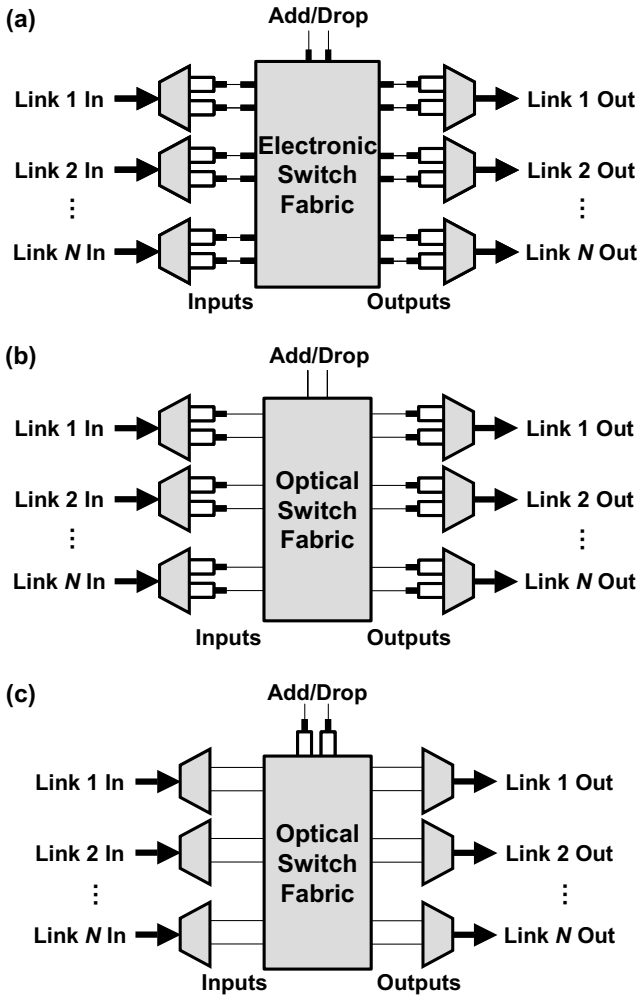


Fig. 2.17 Examples of optical-switch architectures. (a) O-E-O architecture with electronic switch fabric and electronic interfaces on all ports (Section 2.6.1). (b) Photonic switch that switches the 1310-nm optical signal (Section 2.6.2). (c) A wavelength-selective all-optical switch (Section 2.6.3).

One special type of O-E-O switch is a grooming switch, which processes the substrate signals carried on a wavelength in order to better pack the wavelengths. Grooming switches are discussed further in Chapter 6.

2.6.2 Photonic Switch

The term *photonic switch* refers to an optical switch where the switch fabric is optical so that the incoming optical signal does not have to be converted to the electrical domain. MEMS technology is often used to build photonic switches. However, a photonic switch does not necessarily imply optical bypass through a node. For example, Fig. 2.17(b) illustrates a photonic switch that is being used to switch 1310-nm optical signals. There are WDM transponders on both the input and output fibers, and thus, optical bypass is not supported.

A carrier may choose to implement the configuration of Fig. 2.17(b) in order to isolate various technologies in the network. Because each wavelength is terminated on a WDM transponder in the node, the links are isolated from each other. This would allow, for example, the transmission technology used on each link to be supplied by different vendors. Furthermore, the switch vendor may be independent from the transmission vendor because the switch is operating on the standard 1310-nm signal as opposed to a WDM-compatible signal.

This same vendor independence is provided by the O-E-O switch of Fig. 2.17(a); however, the photonic switch configuration of Fig. 2.17(b) requires significantly less electronics. In terms of port count, however, the photonic switch is no smaller than the O-E-O switch. Using the same example of a degree-four node with 160 wavelengths per fiber and 50% add/drop, the required switch size, assuming wavelength granularity, is again 960x960. The advantage is that the switch fabric is not electronic and electronic interfaces are not required on the ports.

2.6.3 All-Optical Switch

All-optical switch is the term for a switch with an optical switch fabric that is being used to support optical bypass. There have been many technologies proposed over the years for building all-optical switches, although some tended to be costly. One cost-effective all-optical switch implementation, developed in the 2003 timeframe, is based on the same Nx1 WSS technology as the second-generation OADM-MD. Because of the scalability of the WSS-based architecture, it is possible to pass the add/drop traffic through the switch fabric, as shown in Fig. 2.18, rather than having the add/drop ports be simple taps off of the network fibers. In contrast to the OADM-MD of Fig. 2.15 (and Fig. 2.16), the all-optical switch of Fig. 2.18 allows any transponder to access any network link. The tradeoff is that the all-optical switch has greater complexity and cost. The degree of the element shown is six, with three of the six ports designated for local access (when the

add/drop traffic passes through the switch fabric in this type of architecture, the ports are referred to as local access ports). Thus, six 5x1 WSSs are needed.

The edge flexibility of an all-optical switch provides similar benefits as described in Section 2.4.5 for the OADM. It can flexibly route dynamic traffic, as any transponder can access any of the network links at the node. It supports optical-layer protection, where a single transponder is used for either the working or protect paths, with the switch toggling between paths at the time of a failure. Because the transponders are not tied to a particular network link, typically fewer transponders are predeployed at a node with an all-optical switch as compared to an OADM-MD.

The architecture of Fig. 2.18 is considered broadcast-and-select (even though it contains wavelength-selective switches) because the passive splitters on the input link broadcast the signal to multiple lines and the WSSs select which wavelengths should pass through to a network link or to a local access port. As with the broadcast-and-select OADM-MD, this all-optical switch architecture supports network multicast where a signal from an incoming fiber is sent to multiple outgoing fibers. In addition, it allows an added signal to be multicast to multiple outgoing fibers; this is not supported with an OADM-MD.




Fig. 2.18 All-optical switch with three network ports and three local access ports. This architecture provides edge configurability. The trapezoids represent general colorless or fixed local access ports.

This type of all-optical switch can be part of an element migration path, similar to the OADM-MD. One can potentially grow in-service from an optical terminal to an OADM to an all-optical switch. Note that the OADM based on this architecture would pass the add/drop traffic through the switch fabric, such that edge configurability is provided; i.e., the OADM in this migration path is composed of two network link ports, two local access ports, and four 3x1 WSSs.

With this type of all-optical switch, demultiplexing of the WDM signal typically occurs internal to the switch, so that the ports are multi-wavelength. There are D network ports, where D is the nodal degree, and L local access ports. An interesting question arises as to how large L should be relative to D . In Fig. 2.18, D and L are both three, thereby potentially supporting up to 100% drop at the node. However, because transponders can access any network port, and because the percentage of drop at a node is typically less than 100%, it is often possible to reduce the number of local access ports. However, while this strategy saves cost, it can lead to wavelength conflicts. In general, if there are L multi-wavelength local access ports, then a particular wavelength can only be dropped from L of the network ports; otherwise there would be wavelength clashes on the local access ports. This limitation becomes more of a problem when the network is heavily loaded and there are few free wavelengths on a link. While advanced routing and wavelength assignment algorithms can be used to minimize wavelength conflicts, contention may still arise if there are not enough local access ports.

Note that the DSE technology used in the first-generation OADM-MD, where the total number of ports is limited to four, could conceivably be used for an all-optical switch at degree-three nodes. Three ports would be used for network links and one port would be used for local access. This would allow transponders on the one local access port to access any of the three network links. However, this design limits the maximum add/drop at the node to 33%, which is not sufficient for many nodes. Furthermore, any given wavelength can drop from at most one network link, which would likely lead to wavelength conflicts as the network fills. Thus, the DSE technology is not favorable for the all-optical switch configuration.

Even if the number of local access ports equals the nodal degree (i.e., L equals D), wavelength conflicts may arise due to the multi-wavelength local access ports. Assume that a transponder tuned to λ_1 is in use on a local access port, where the corresponding signal is added/dropped to/from the East link. Further assume that an IP router feeds a second transponder that is plugged into the same local access port. Assume that the IP router wishes to initiate a connection via the West link and the only available wavelength (due to congestion on other links of the path) is λ_1 . This connection would be blocked because there is already a transponder on this local access port using λ_1 . While such a scenario is possible, its occurrence is rare and would likely only occur if the network is very heavily loaded.

Another type of all-optical switch is based on the wavelength-selective architecture, as illustrated in Fig. 2.17(c) (this is similar to the wavelength-selective OADM architecture of Fig. 2.9). A MEMS-based switch, in combination with multiplexers/demultiplexers, is often used for this application. The WDM signal entering from a fiber is demultiplexed into its constituent wavelengths, where each

wavelength feeds an input port on the switch fabric. Some of the wavelengths are directed to add/drop ports while others are directed to output ports corresponding to other network fibers. In contrast to the broadcast-and-select architecture of Fig. 2.18, the add/drop ports are single-wavelength so that wavelength conflicts do not arise. However, this type of architecture does not readily support multicast, either from one input fiber to multiple output fibers or from one add port to multiple output fibers. It also tends to be more expensive than the architecture of Fig. 2.18.

Note that while the switch of Fig. 2.17(c) supports optical bypass, such that transponders are not required for the through traffic, the number of ports on the switch fabric is no smaller than that required for Fig. 2.17(a) or Fig. 2.17(b). Again, for a degree-four node supporting 50% add/drop, where each fiber carries 160 wavelengths, the switch fabric needs to be 960x960, assuming wavelength granularity. Scaling technologies such as MEMS to this size can be challenging.

2.7 Hierarchical Switches

One means of fabricating scalable optical switches is through the use of a hierarchical architecture where multiple switch granularities are supported [HSKO99, SaSi99]. Coarse switching is provided to alleviate requirements for switches with very large port counts; a limited amount of finer granularity switching is also provided. A functional illustration of a three-level hierarchical switch is shown in Fig. 2.19, where switching can be performed on a per-fiber basis, a per-waveband basis, or a per-wavelength basis. While the three levels are shown as distinct switches (or cross-connects), it is possible to build such a multi-granular switch with a single switching fabric, e.g., with MEMS technology [LiVe02].

If the traffic at a node is such that all the traffic entering from one fiber is directed to another fiber, then that traffic is passed through the fiber-level switch only. The switch provides configurability if the traffic pattern changes, while providing optical bypass for all traffic on the fiber. If fiber-level bypass is not appropriate for a network, then that level of the switching hierarchy can be removed.

For traffic that needs to be processed on a finer granularity than a fiber, the band-level switch demultiplexes the WDM signal into its constituent wavebands. Some of the wavebands are switched without any further demultiplexing, providing band-level bypass, whereas some of the wavebands have to be further demultiplexed into their constituent wavelengths so that individual wavelengths can be dropped or switched. When equipped with wavelength converters, whether electronic or optical, the wavelength-level switch can also be used to better pack the wavebands (i.e., waveband grooming). Changing the lambda of a wavelength allows the wavelength to be shifted from one waveband to another.

The hierarchical approach addresses the port count issue while providing more flexibility than a single-layer switch of a coarse granularity. Studies have shown a significant number of ports can be saved through the use of a multi-granularity switch, with the percentage of ports saved increasing with the level of traffic carried in the network [NoVD01, CaAQ04].




Fig. 2.19 A three-level hierarchical switch, allowing fiber-bypass, band-bypass, and wavelength-bypass. In the figure, each fiber contains two wavebands, and each waveband contains four wavelengths. Wavelength conversion is used to groom the wavebands. (Adapted from [SaSi99]. © 1999 IEEE)

2.8 Adding Edge Configurability to a Node

As discussed previously, OADM-MDs and some OADM implementations do not have built-in edge configurability. Without manual intervention, any transponder is restricted to accessing just a single network link at the node. A greater amount of automatic configurability is desirable to meet the requirements of dynamic services. Two methods for achieving the desired edge configurability with these network elements are discussed here. Whether to deploy one of these architectures, or alternatively deploy an all-optical switch that has built-in edge configurability, depends on overall nodal cost, complexity, reliability, and multicast requirements. Note that some telecommunications carriers have installed OADM-MDs in their networks due to low-cost all-optical switches not being available when the networks were deployed.

2.8.1 Adjunct Edge Switch

One option for increasing the configurability of the node is to deploy an optical switch in conjunction with the OADM-MD or OADM, as shown with a degree-three OADM-MD in Fig. 2.20. The added switch is often referred to as an ‘*adjunct edge switch*’ or simply an ‘*edge switch*’. The traffic from higher layers (e.g., IP) is fed through the edge switch so that it can be automatically directed to the desired network port. Because the edge switch is only for the add/drop traffic, and not for the through traffic, the size can be smaller as compared to a switch that operates on all of the nodal wavelengths. Returning to the example of a degree-four node with 160 wavelengths per fiber and 50% add/drop at the node, the edge switch needs to be of size 640x640, as opposed to 960x960 for a ‘core’ switch. (Note that the IP router itself could be used to provide edge configurability by directing traffic to a different router port that is connected to the desired transponder; however, this requires more router ports, which tend to be costly.)

With the configuration of Fig. 2.20(a), the client 1310-nm optical signal is passed through the edge switch. The edge switch can be O-E-O-based or photonic; the latter is shown in the figure. The transponders are tied to a particular port, and the edge switch directs the client signal to the desired transponder. In Fig. 2.20(b), the edge switch operates on the WDM-compatible signal, and thus, must be photonic. This configuration provides more agility than Fig. 2.20(a) because the transponders can access any port. Thus, fewer transponders need to be predeployed with this configuration. There are some applications where the configuration of Fig. 2.20(b) must be used; e.g., see Section 2.10 and Section 4.5.2.

Note that deploying an edge switch at a node can be useful for other reasons, e.g., protection, as is discussed in later chapters.

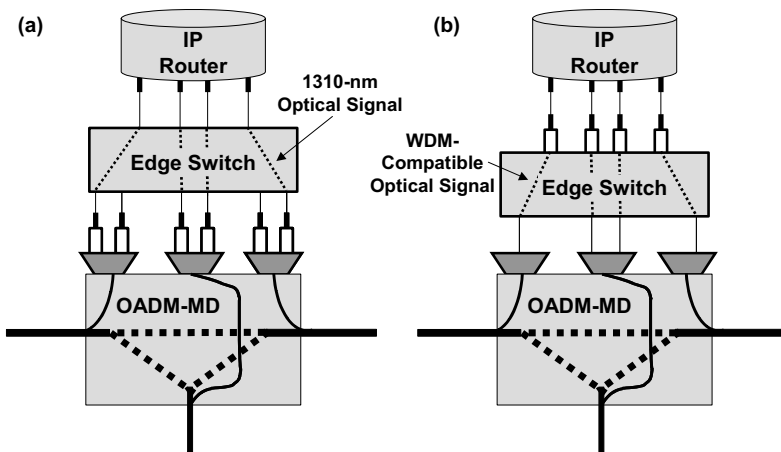


Fig. 2.20 An edge switch used in conjunction with an OADM-MD in order to provide edge configurability. In (a), the 1310-nm signal is switched; in (b), the WDM-compatible signal is switched. There is more agility with (b), so that fewer transponders need to be predeployed.

2.8.2 Flexible Transponders

Rather than incurring the cost of an edge switch, another option is to use *flexible WDM transponders*, where the output of the transponder is split into multiple paths and each path feeds into a different add port on the OADM or OADM-MD. In the reverse direction, the transponder is equipped with a switch to select a signal from one of the drop ports. This is illustrated in Fig. 2.21, where a transponder with three-way flexibility is able to access any of the network links at the degree-three node. Consider the general scenario of a degree- N node. Passively splitting the transponder output over N paths to access any of the links results in $10 \cdot \log_{10} N$ dB of optical loss, which is undesirable as N increases. One option is to have the transponder output split over M paths, where $M < N$; such a transponder is limited to accessing only M of the N nodal links. However, as proposed in [SiSa07], there is an alternative transponder architecture that uses small switches on the transmit side, so that the transponder can access any of the N nodal links while suffering no more than a nominal 3-dB optical loss. Furthermore, this transponder design allows it to simultaneously access any combination of two of the N nodal links, to support two-way multicast of an added signal, or to support optical-layer protection with an active backup path as is discussed further in Chapter 7.

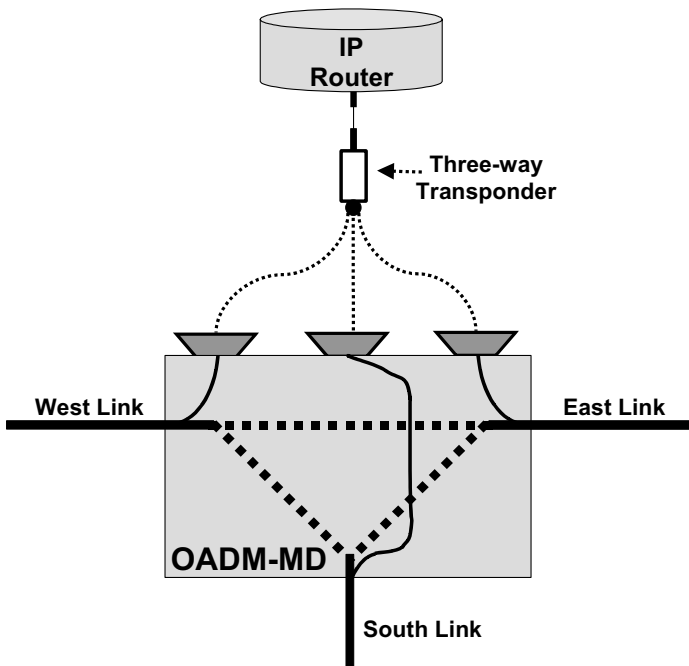


Fig. 2.21 A three-way transponder is able to access any of the network links at the degree-three node. It is desirable to use an optical backplane to simplify the cabling.

One advantage of the flexible transponder solution is that fewer transponders need to be deployed at a node as compared to an OADM-MD with regular transponders. There is one pool of transponders that can access any of the N nodal links, rather than N pools of transponders, each associated with a particular link. One possible drawback of the N -way flexible transponder is that one transponder occupies N add/drop slots. However, not all transponders at the node have to be equipped with the flexible architecture; regular transponders can still be used for the non-dynamic traffic.

2.9 Optical Reach

Network elements capable of allowing transiting traffic to remain in the optical domain, such as the OADM, OADM-MD, and the all-optical switch, is one requirement for optical bypass. In addition, the underlying transmission system must be compatible with a signal remaining in the optical domain as it traverses one or more nodes. An important property of a transmission system is the optical reach, which is the maximum distance an optical signal can be transmitted before it degrades to a level that requires the signal be regenerated, where regeneration occurs in the electronic domain. (All-optical regeneration, while feasible, has not been widely implemented. This is discussed further in Chapter 4.)

Consider the four-node linear network in Fig. 2.22. Nodes A and D are equipped with optical terminals, whereas Nodes B and C are equipped with OADMs. The distance of each of the three links is 1,000 km. Assume that the connection of interest is between Node A and Node D.

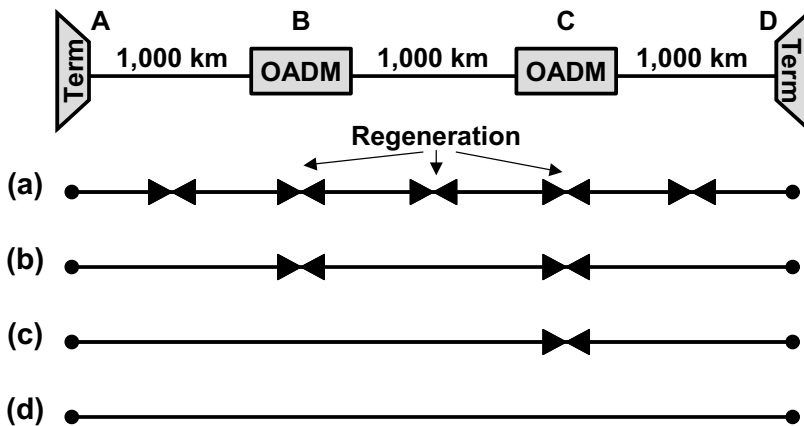


Fig. 2.22 Connection between Nodes A and D with: (a) 500-km optical reach; (b) 1,000-km optical reach; (c) 2,000-km optical reach; and (d) 3,000-km optical reach.

In Fig. 2.22(a), the optical reach is assumed to be 500 km. With this reach, not only must the connection be regenerated at Nodes B and C, it must also be regenerated at intermediate dedicated regeneration sites along the links. These sites would otherwise be equipped with just an optical amplifier, but due to the limited reach, need to regenerate all traffic that passes through them. The OADMs at Nodes B and C, while capable of optical bypass, cannot be used for that purpose in this scenario because of the limited reach. In Fig. 2.22(b), the optical reach is 1,000 km; this removes the regeneration at intermediate sites along the links, but still is not sufficient to allow the signal to optically bypass either Node B or C. With an optical reach of 2,000 km, as shown in Fig. 2.22(c), the connection can make use of the OADM at Node B to optically bypass that node, but it still must be regenerated at Node C (equivalently, it could be regenerated at Node B and optically bypass Node C). With an optical reach of 3,000 km, the connection is able to remain in the optical domain over the whole path, as shown in Fig. 2.22(d).

As this example shows, the optical reach is critical in determining how much optical bypass is achieved in a network. In legacy transmission systems based on EDFA technology, the optical reach is on the order of 500 to 600 km; with newer EDFA systems, the maximum reach is on the order of 1,000 to 1,500 km. To obtain significantly longer reach, Raman amplification is used (see Chapter 4). Raman systems generally have an optical reach in the range of 1,500 to 4,000 km, depending on the equipment vendor. Such technology is sometimes referred to as ‘ultra-long-haul technology’ to emphasize the extended optical reach. In typical networks, the combination of optical-bypass elements and extended optical reach may eliminate on the order of 90% of the regenerations as compared to a system with no optical bypass and 500-km reach.

Many factors go into developing a transmission system that supports extended reach, some of which are touched on here; the subject is revisited in Chapter 4. As noted above, Raman amplification, sometimes in conjunction with EDFA amplification, is generally used in such systems. It is also necessary to deal with a host of optical impairments, such as chromatic dispersion, polarization-mode dispersion, crosstalk, and four-wave mixing, all of which can degrade an optical signal (these are discussed in Chapter 4). Some of these impairments can be mitigated with special compensating equipment, while some of the problems from these impairments can be avoided by transmitting signals at a low enough power level.

Advanced transmitters and receivers, as well as robust transmission schemes, are also important for attaining extended optical reach. Furthermore, the optical network elements themselves must be compatible with extended reach. For example, they must be low loss, and not cause excessive distortion of the signal. Another key system component is advanced forward error correction (FEC). FEC allows errors picked up in transmission to be corrected at the destination. The stronger the FEC coding, the more errors can be corrected. This allows the signal to degrade further before it needs to be regenerated.

This list is not meant to be a comprehensive discussion of what goes into achieving extended optical reach. Suffice it to say it is quite an engineering accomplishment.

The desirable optical reach for a network depends on the geographic tier of where the system is deployed. In the metro-core, the longest connection paths are on the order of a few hundred kilometers. Thus, an optical reach of 500 km is sufficient to remove most, if not all, of the regeneration in the network. In regional networks, the longest connection paths are typically in the range of 1,000 to 1,500 km, requiring an optical reach of that order to eliminate regeneration. In backbone networks, the longest connection paths can be several thousand kilometers. For example, in the United States, the longest backbone connections, when including protection paths, are on the order of 8,000 km. A system with 4,000-km optical reach will eliminate much of the regeneration, but not all of it.

Note that the key factor in determining whether extended optical reach is useful in a particular network is the distribution of connection path distances. A common mistake is to focus on the link distances instead, where some may assume that extended optical reach provides no benefit unless the link distances are very long. In fact, optical bypass can be more effective in networks with high nodal density and relatively short links because a particular connection is likely to transit several intermediate nodes. Chapter 8 investigates the optimal optical reach of a network from a cost perspective.

Regarding terminology, it is perhaps not clear whether traffic that is regenerated at a node should be considered through traffic or add/drop traffic. As regeneration is typically accomplished via O-E-O means, it is usually considered add/drop traffic because it is dropping from the optical domain. This is the convention that is adopted here.

Chapter 4 discusses various regeneration architectures and strategies.

2.10 Integrating WDM Transceivers in the Client Layer

Much of the discussion so far has focused on removing the transponders for the nodal through traffic by implementing optical bypass. However, the add/drop traffic also provides an opportunity to remove some of the electronics from the node. The electronic higher layers and the optical layer typically communicate via a standard 1310-nm optical signal. In Fig. 2.23(a), the IP router is equipped with an interface to generate the 1310-nm signal, and the WDM transponder plugged into the OADM converts the 1310-nm signal to a WDM-compatible signal.

The equipment can be simplified by having a WDM-compatible transceiver deployed directly on the IP router, as shown in Fig. 2.23(b). This is referred to as *integrating* the transceivers with the IP router. (The term transceiver is used rather than transponder because the input is an electrical signal from the IP router rather than a 1310-nm optical signal; i.e., there is no transponding of an optical signal.) The IP router and OADM would then communicate via a WDM-compatible signal, which would be added to the WDM signal exiting the OADM. This eliminates the electronic interfaces between the router and the OADM. Clearly, the transceiver output must meet the specifications of the WDM transmission system for this configuration to work.




Fig. 2.23 In (a), the IP router communicates with the OADM via a standard 1310-nm optical signal. In (b), the electrical interfaces can be removed as WDM transceivers are deployed directly on the IP router.

This same configuration is possible with an IP router in combination with an OADM-MD or an all-optical switch as well. Furthermore, transceivers can be integrated with other electronic switching elements. For example, referring back to Fig. 2.17(a), the combination of a WDM transponder and an electronic interface on each switch port of the O-E-O switch can be replaced by a transceiver that is integrated with the O-E-O switch, thereby eliminating a lot of electronics. Note, however, this integrated architecture negates one of the advantages of O-E-O switches, namely the independence of the switch and transmission-system vendors. With integrated transceivers, either one vendor supplies both the switch and the transmission system, or separate vendors collaborate to ensure compatibility.

Consider the integrated-transceiver scenario where the optical network element is not capable of edge configurability. Assume that an adjunct edge switch is deployed at the node as discussed in Section 2.8.1 in order to achieve edge flexibility. The edge switch would be deployed between the transceivers and the optical network element, and would need to be capable of switching WDM-compatible signals.

2.11 Photonic Integrated Circuits

One of the motivations for removing electronics from a node is to reduce cost and physical space requirements. This led to the development of optical-bypass technology, which is now deployed in several major telecommunications networks. A more recent development is *photonic integrated circuit* (PIC) technology, used in combination with the more traditional O-E-O architecture [MDLP05, Welc06]. The idea is not to eliminate transponders but to make them lower cost and smaller through integration. For example, ten transponders may be integrated onto a sin-

gle small chip. One advantage is that because the architecture is based on O-E-O at every node, there are no wavelength continuity constraints, thereby simplifying the algorithms needed to run the network. It also allows for SONET/SDH performance monitoring at every node. If there is widespread adoption of the PIC technology, this would change the direction of optical network evolution. However, it is not clear that this technology addresses operational issues such as required power, heat dissipation, reliability, and full nodal configurability.

2.12 Multi-Fiber-Pair Systems

Thus far, the discussion has implicitly assumed that each network link is populated by one fiber-pair. In this section, the multi-fiber-pair scenario is considered.

The O-E-O architecture does not change with multiple fiber-pairs per link. There is an optical terminal for every fiber-pair entering a node, and the signal from every incoming fiber is demultiplexed. In order to reduce the electronics, it may be possible to have an incoming fiber from one link be directly connected to an outgoing fiber on another link, so that fiber-bypass is achieved, though this arrangement is not readily reconfigurable.

For a network with optical bypass, there are a few options for architecting the node. First, assume that every link has N fiber-pairs. One option is to deploy N copies of a network element at a node. For example, Fig. 2.24 shows a degree-two node with two fiber-pairs per link. The node is equipped with two OADMs. Optical bypass is supported for traffic that bypasses the node using fiber-pair 1 or fiber-pair 2. However, traffic routed on fiber-pair 1 of one link and fiber-pair 2 of the other link requires O-E-O conversion.




Fig. 2.24 Two OADMs deployed in parallel at a degree-two node with two fiber-pairs per link.

Alternatively, if optical bypass is desired regardless of how the traffic is routed, a degree-four OADM-MD (or all-optical switch) could be used. This may provide more flexibility than is required, however, as it allows a signal to be all-optically routed between fiber-pair 1 of a link and fiber-pair 2 of that same link. In a broadcast-and-select architecture, the OADM-MD design could be simplified so that paths from an incoming fiber are split into two paths rather than three, to allow optical bypass in just the two desired directions. For example, the OADM-MD architecture of Fig. 2.15 can be modified for a ‘quasi-degree-four’ node as shown in Fig. 2.25. (A similar type of simplification is possible with the all-optical switch architecture of Fig. 2.18.)




Fig. 2.25 A simplified OADM-MD at a ‘quasi-degree-four’ node where routing between fiber-pairs on the same link is not required. As compared to a full degree-four OADM-MD, each through path is split two ways rather than three, and there are 2x1 WSSs rather than 3x1 WSSs.

If the number of fiber-pairs on the links is unequal, due to some links being more heavily utilized, then either one large OADM-MD or all-optical switch could be deployed at a node to provide full optical bypass, or some combination of optical bypass and O-E-O could be used. For example, if a degree-three node has one link with two fiber-pairs and one link with one fiber-pair, a degree-three OADM-MD would provide optical bypass in all directions. Again, this is actually more flexibility than is typically needed as it provides an all-optical path between the

two fiber-pairs that are on the same link; the OADM-MD can be simplified, analogous to what is shown in Fig. 2.25. Another option is to deploy an OADM in combination with an optical terminal, in which case the network planning process should favor routing the through traffic on the fiber-pairs interconnected by the OADM.

When considering the maximum desired size of an OADM-MD or all-optical switch it is important to take into account the possibility of multiple fiber-pairs. Thus, the maximum desired sizes of the various network elements may be larger than what is indicated by simply looking at nodal-degree distribution.

Chapter 3

Routing Algorithms

Telecommunications networks are generally so large and complex that manually designing a network in a reasonable amount of time is prohibitively difficult. Network designers primarily rely on automated algorithms to determine, for example, how to route traffic through the network, how to protect the traffic, and how to bundle the traffic into wavelengths. In systems with optical bypass, additional algorithms are needed to handle regeneration and to ensure that wavelength contention issues are minimal. The fact that networks have reached a stage where algorithms are essential in producing cost-effective and efficient network designs can be daunting. The good news is that extensive research has been done in this area and much expertise has been gained from live network deployments, resulting in the development of relatively straightforward algorithms that produce very effective network designs.

When designing network algorithms, it is important to consider the size of the problem in terms of the number of network nodes, the amount of traffic carried in the network, and the system specifications. Metro-core networks have tens of nodes whereas backbone networks may have as many as 100 nodes or more. The size of the demand set depends on whether the traffic requires grooming or not. First, consider a backbone network. If all of the traffic is at the wavelength line-rate (no grooming needed), there are typically a few hundred to a couple of thousand demands in the network. If all of the traffic is subrate, such that grooming is needed, there could be tens of thousands of demands. The number of demands in a metro-core network is significantly lower than this. A key system parameter is the number of wavelengths per fiber. Metro-core WDM networks generally have no more than 40 wavelengths per fiber, whereas backbone networks typically have 80 to 160 wavelengths per fiber. Any algorithms used in the network planning process must be scalable under these conditions.

The run-time of the network planning algorithms is very important. In a highly dynamic real-time environment, a new connection may need to be established in less than one second. The process of planning the route of the connection and de-

termining which network resources should be used to carry it ideally should be completed in less than ~250 milliseconds to allow time for the network to be configured appropriately to carry the connection. Furthermore, real-time design may be implemented in a distributed manner at the network nodes, where processing and memory capabilities may be limited.

In long-term network planning, design time is not as critical, although it is still important. In some long-term planning exercises, a large number of demands (e.g., thousands of substrate demands) may need to be processed at once. Furthermore, if working with a ‘greenfield’ network (i.e., a completely new network), numerous design scenarios typically are considered. For example, the process may include comparing different network topologies, different line-rates, or different protection strategies. Due to the number of designs that need to be run, and the often short time allocated to the network design process (especially when the design exercise is performed by a system vendor in response to a carrier request), it is desirable that the planning process for each scenario take no more than a couple of minutes.

To produce a network design that is optimal relative to a set of metrics is typically computationally complex and would require an inordinate amount of time for realistic networks. Thus, many steps in the planning process rely on heuristics, which experience has shown produce very good, though not always optimal, results, and which run in a reasonable amount of time.

A brief overview of the network planning process can be found in [Simm06]. This chapter focuses on the routing component of the algorithms, whereas Chapters 4 and 5 discuss regeneration and wavelength assignment, respectively. The initial emphasis is on treating these three components as sequential steps in the planning process; however, performing routing, regeneration, and wavelength assignment in a single step is covered in Chapter 5.

Routing is the process of selecting a path through the network for a traffic demand, where there are typically many possible paths to get from the demand source to the demand destination. It is important to take into account several factors when selecting a route. First, cost is a key consideration. The selected route should require adding minimal cost to the network when possible. The path distance and number of links in the path may also be relevant, as these are indicators of the bandwidth occupied by the path; these factors also may affect the reliability of the connection. However, this does not imply that demands should always be routed over the shortest possible path or the path with the fewest links. In fact, such a strategy may lead to inefficient designs. It is also necessary to consider the total potential capacity of the network, where a particular path may be chosen because it leaves the network in a better state to accommodate future traffic.

Several routing strategies are discussed in this chapter, with a focus on relatively straightforward methodologies that are effective in network planning for practical optical networks. The chapter is not meant to be a review of all known routing algorithms and strategies. Much of the chapter is equally applicable to a network with optical bypass as to a network based on O-E-O technology.

The first few sections of the chapter specifically examine routing a demand with one source and one destination over a single path. Section 3.1 introduces shortest path routing algorithms and Section 3.2 covers how such algorithms are used depending on the underlying technology of the network. Sections 3.3 and 3.4 describe some effective routing strategies that take into account network cost and network utilization. Section 3.5 considers the more detailed network modeling that may be needed when routing demands in real-time, where only the equipment that is already deployed in the network can be used. Routing with protection is covered in Section 3.6, where two or more diverse paths are required for a demand to enable recovery from a failure. Section 3.7 is relevant to planning scenarios where multiple demand requests are processed at one time. Assuming that the routing policy is adaptive, where the selected path is dependent on the state of the network, then the order in which the demands are considered is important. Various effective ordering strategies are presented in this section. Section 3.8 discusses multicast traffic, where there is one source and multiple destinations. The final section in the chapter discusses some of the challenges of real-time network planning when the current state of the network is not fully known.

3.1 Shortest Path Algorithms

Most routing strategies incorporate some type of shortest path algorithm to determine which path minimizes a particular metric. A general discussion of shortest path algorithms can be found in [CoLR90]. In the shortest path algorithms discussed here, it is assumed that the metric for an end-to-end path is the sum of the metrics of the links comprising the path. Any such additive metric can be used, depending on the goal of the routing process. For example, to find the path with the shortest geographic distance, each link is assigned a metric equal to its own distance. As another example, assume that each link is assigned a metric of unity. The shortest path algorithm then finds the path that traverses the fewest hops. (The term *hop* is often used to refer to each link in a path.) As a third example, assume that each network link has a certain probability of being available (i.e., not failed), and assume that each link in the network fails independently such that the availability of a path is the product of the availabilities of each link in the path (ignoring node failures). The link metric can be chosen to be the negative of the logarithm of the link availability, where the logarithm function is used in general to convert a multiplicative metric to an additive metric. Higher link availability corresponds to a smaller link metric; thus, running the shortest path algorithm with this metric produces the path with the highest availability. As this last example demonstrates, the metric may be unrelated to distance; thus, the term ‘shortest path’ is in general a misnomer. Nevertheless, this term will be used here to represent the path that minimizes the desired metric.

One very well known shortest path algorithm is the Dijkstra algorithm, where the inputs to the algorithm are the network topology and the source and the desti-

nation. This is a ‘greedy’ algorithm that is guaranteed to find the shortest path from source to destination, assuming a path exists. Greedy algorithms proceed by choosing the optimal option at each step without considering future steps. In the case of the Dijkstra algorithm, this strategy produces the optimal overall result (in general, however, greedy algorithms do not always yield the optimal solution).

Another shortest path algorithm is the Breadth-First-Search (BFS) Shortest Path algorithm, which proceeds by considering all one-hop paths from the source, then all two-hop paths from the source, etc., until the shortest path is found from source to destination [Bhan99]. If there is a unique shortest path from source to destination, the BFS and Dijkstra algorithms produce the same result. However, if there are multiple paths that are tied for the shortest, the BFS algorithm finds the shortest path with the fewest number of hops. This can be helpful in network design because fewer hops can potentially translate into lower cost or less wavelength contention, as is discussed in the next section. The Dijkstra algorithm does not in general have this same tie-breaking property. Furthermore, the BFS algorithm works with negative link metrics, as long as there are no cycles in the network where the sum of the link metrics is negative. This is relevant for one of the graph transformations that is commonly used as part of an algorithm to find two or more diverse routes, as discussed in Section 3.6. The Dijkstra algorithm needs a small modification to be used with negative link metrics. Overall, then, the BFS shortest path algorithm is somewhat better suited for network design; code for this algorithm is provided in Appendix B.

The Dijkstra and BFS algorithms can be applied whether or not the links in the network are bi-directional (a link is bi-directional if traffic can be routed in either direction over the link). Furthermore, the algorithms work if different metrics are assigned to the two directions of a bi-directional link. However, if the network is bi-directionally symmetric, such that the traffic flow is always two-way and such that the metrics are the same for the two directions, then a shortest path from source to destination also represents, in reverse, a shortest path from destination to source. (This is also called an *undirected* network.) In this scenario, which is typical of telecommunications networks, it does not matter which node is designated as the source and which is designated as the destination.

The shortest path algorithm can be incorporated as part of a larger procedure to find the K-shortest paths. K-shortest path routines find the shortest path between the source and destination, the second shortest path, etc., until the K^{th} shortest path is found or until no more paths exist. Note that the paths that are found are not necessarily completely disjoint from each other; i.e., the paths may have links and/or nodes in common. Many K-shortest paths algorithms exist, e.g., [Yen71, Epps94], where the ones that find only simple paths (i.e., paths without loops) are the most relevant for network design. The code for one such K-shortest paths algorithm is provided in Appendix B (the code follows the procedure described in [HeMS03]).

A variation of the shortest-path problem arises when one or more constraints are placed on the desired path; this is known as the *constrained shortest path* (CSP) problem. Some constraints are straightforward to handle. For example, if

one is searching for the shortest path subject to all links of the path having at least N wavelengths free, then prior to running a shortest path algorithm, all links with less than N free wavelengths are removed from the topology. As another example, the intermediate steps of the BFS shortest path algorithm can be readily used to determine the shortest path subject to the number of path hops being less than H , for any $H > 0$ (similar to [GuOr02]). However, more generally, the CSP problem can be difficult to solve; for example, determining the shortest path subject to the availability of the path being greater than some threshold (where it is assumed the availability is based on factors other than distance). Various heuristics have been proposed to address the CSP problem; e.g., [KoKr01, LiRa01] (the latter reference addresses the constrained K-shortest paths problem). Some heuristics have been proposed to specifically address the scenario where there is just a single constraint; this is known as the *restricted shortest path* (RSP) problem. Additionally, a simpler version of the multi-constraint problem arises when *any* path satisfying all of the constraints is desired, not necessarily the shortest path; this is known as the *multi-constrained path* (MCP) problem. An overview, including a performance comparison, of various heuristics that address the RSP and MCP problems can be found in [KKKV04].

3.2 Routing Metrics

As discussed in the previous section, a variety of metrics can be used with the shortest path algorithm. Two common search strategies are: find the path with the fewest hops and find the path with the shortest distance. With respect to minimizing network cost, the optimal routing strategy to use is dependent on the underlying system technology, as discussed next. (In this section, issues such as grooming and shared protection that also may have an impact on cost are not considered; these issues are discussed in Chapters 6 and 7, respectively.)

3.2.1 Fewest-Hops Path vs. Shortest-Distance Path

In a pure O-E-O network, a connection is electronically terminated (i.e., regenerated) at every intermediate node along the path, where the electronic terminating equipment is a major component of the path cost. Thus, searching for the path from source to destination with the fewest hops is generally favored as it minimizes the amount of required regeneration. This is illustrated in Fig. 3.1 for a connection between Nodes A and Z. Path 1 is the shortest-distance path at 900 km, but includes four hops. Path 2, though it has a distance of 1,200 km, is typically lower cost because it has only two hops and thus requires less intermediate electronic termination.




Fig. 3.1 Path 1, A-B-C-D-Z, is the shortest-distance path between Nodes A and Z, but Path 2, A-E-Z, is the fewest-hops path. In an O-E-O network, where the signal is regenerated at every intermediate node, Path 2 is typically the lower-cost path.

In networks with optical bypass, regeneration is determined by the system optical reach, which is typically based on distance. For example, an optical reach of 2,000 km indicates the connection can travel no further than 2,000 km before it needs to be regenerated. (In reality, the optical reach is determined by many factors as is discussed in Chapter 4, but for simplicity, it is usually specified in terms of a distance.) This favors searching for the shortest-distance path between the source and destination. However, with optical-bypass systems, there is a wavelength continuity constraint, such that the connection must remain on the same wavelength (i.e., λ) as it optically bypasses nodes. Finding a wavelength that is free to carry the connection is potentially more difficult as the number of links in the path increases. This implies that the number of path hops should be considered as well.




Fig. 3.2 Assume that this is an optical-bypass-enabled network with an optical reach of 2,000 km. Path 1, A-B-C-Z, has the fewest hops but requires one regeneration. Path 2, A-D-E-F-G-Z, and Path 3, A-H-I-J-Z, require no regeneration. Of these two lowest-cost paths, Path 3 is preferred because it has fewer hops.

Overall, a good strategy for optical-bypass systems is to search for a route based on distance, but of the paths that meet the minimal regeneration, favor the one with the fewest hops. This is illustrated by the three paths between Nodes A and Z shown in Fig. 3.2, where the optical reach is assumed to be 2,000 km. Path 1, with a distance of 2,500 km, has the fewest hops but requires one regeneration. Path 2, with a distance of 1,500 km, and Path 3, with a distance of 1,600 km, do not require any regeneration and are thus lower-cost paths. Of these two paths, Path 3 is generally more favorable (all other factors, e.g., link load, being equal), even though it is somewhat longer than Path 2, because it has only four hops compared to the five hops of Path 2.

3.2.2 Shortest-Distance Path vs. Minimum-Regeneration Path

While path distance is clearly related to the amount of regeneration in an optical-bypass-enabled network, it is important to note that the path with shortest physical distance is not necessarily the path with minimum regeneration. In addition to considering the distance over which a signal has traveled in determining where to regenerate, carriers generally require that any regeneration occur in a network node (i.e., the add/drop and switching locations of a network), as opposed to at an arbitrary site along a link. First, sites along a link, e.g., amplifier huts⁷, may not be large enough to house regeneration equipment. Second, from a maintenance perspective, it is beneficial to limit the number of locations where regeneration equipment is deployed. (If the optical reach of the system happens to be shorter than the distance between two nodes, then there is no choice but to deploy a dedicated regeneration site along the link where all transiting traffic is regenerated. In this scenario, the dedicated regeneration site can be considered a network node.)

This additional design constraint can lead to more regeneration than would be predicted by the path length. Consider the example shown in Fig. 3.3, with a connection between Node A and Node Z, and assume that the optical reach is 2,000 km. Path 1 is the shortest path for this connection, with a length of 3,500 km. Based on the path length and the optical reach, it is expected that one regeneration would be required. However, because regeneration occurs only at nodes, two regenerations are actually required (e.g., the regenerators could be at Nodes C and E). Path 2, while longer, with a length of 3,800 km, requires just one regeneration (at Node G).

Another factor that affects the amount of regeneration is whether the network fully supports optical bypass in all directions at all nodes. There may be some nodes not fully equipped with optical-bypass equipment. For example, a degree-three node may be equipped with an OADM and an optical terminal (see Section 2.5); any transiting traffic entering via the optical terminal needs to be regenerated regardless of the distance over which it has been transmitted.

⁷ Typically small buildings that house the optical amplifiers.




Fig. 3.3 Assume that this is an optical-bypass-enabled network with an optical reach of 2,000 km. Path 1 is 3,500 km, but requires two regenerations. Path 2 is longer at 3,800 km but requires only one regeneration. (Adapted from [Simm06]. © 2006 IEEE)

As the examples of this section illustrate, simply running a shortest path algorithm may not produce the most desirable route. Rather, it is often advantageous to generate a *set* of candidate paths, and then select a particular path to use based on other factors. For example, a particular route may be selected because of its lower cost or because it avoids a ‘bottleneck’ link that is already heavily loaded.

The next section discusses how to generate a good set of candidate paths, and Section 3.4 discusses how the path set is used in the routing process.

3.3 Generating a Set of Candidate Paths

Two different strategies are presented for generating a set of candidate paths. The first is a more formal strategy that uses the K-shortest paths algorithm to find minimum-cost paths. The second is a somewhat ad-hoc methodology to find a set of paths with good link diversity.

3.3.1 K-Shortest Paths Strategy

The marginal cost of adding a new demand to the network is largely a function of the number of electronic terminations (i.e., regenerations) needed to support the demand. Much of the other network costs, e.g., amplification, are incurred when the network is first installed and are amortized over the whole demand set.

Thus, roughly speaking, in an O-E-O network, paths that have the same number of hops have an equivalent cost because the number of intermediate electronic terminations is the same. To find a set of lowest-cost paths, one can run a K-shortest paths algorithm where the metric is set to unity for all links. The first N of the paths returned by the algorithm will satisfy the fewest-hops criterion, for some N where $1 \leq N \leq K$. Setting K to a number around 10 will typically ensure that all lowest-cost paths are found; i.e., N is usually less than or equal to 10 in practical optical networks.

Similarly, in an optical-bypass-enabled network, paths with the same amount of regeneration can be considered equivalent-cost paths. First, consider networks where any loopless path is shorter than the optical reach, as may be the case in a metro-core network. All paths can be considered lowest-cost because there is no need for regeneration in the network. In this scenario, one can run a K-shortest paths algorithm with a link metric of unity to generate a set of lowest-cost paths that have the fewest, or close to the fewest, number of hops.

In more general optical-bypass-enabled networks where there is regeneration, it is not quite as straightforward to find a set of paths that meet the minimum regeneration. A K-shortest paths algorithm can be run with distance as the link metric. However, as illustrated by the example of Fig. 3.3, each of the returned paths must be examined to determine the actual number of required regenerations. The ones with minimum regeneration make up a set of least-cost paths. There is no guarantee that this procedure produces a path with the lowest possible amount of regeneration. (In Chapter 4, an alternative metric that is more tied to the underlying optical system is presented that may be a better predictor of regeneration.) However, if the minimum number of regenerations found is R , and at least one of the paths found has a distance greater than $(R+1)*[\text{Optical Reach}]$, then the set of minimum-regeneration paths must have been found. Practically speaking, setting K to about 10 in the K-shortest paths algorithm is sufficient to generate a good set of least-cost paths.

Alternatively, one can use a more complex topology transformation, as described in Section 3.5.1, to ensure a minimum-regeneration path is found. However, the simpler method described above is generally sufficient.

(Note: One could use the intermediate steps of the BFS algorithm to find the shortest-distance path subject to a maximum number of hops to assist in finding the minimum-regeneration path with the minimum number of hops. Again, because distance does not directly translate to regeneration, this strategy does not necessarily always succeed. Furthermore, with K large enough, such paths are generally found by the above process anyway.)

3.3.2 Bottleneck-Avoidance Strategy

While the K-shortest paths technique can generate a set of lowest-cost paths, the paths that are found may not exhibit good link diversity. For example, a link that is expected to be heavily loaded may appear in every lowest-cost path found between a particular source and destination. If the link diversity is not sufficient for a particular source/destination pair, then an alternative strategy can be used to generate a candidate path set, as described here.

The first step is to determine the links in the network that are likely to be highly loaded (i.e., the ‘hot spots’). One methodology for estimating load is to perform a preliminary routing where each demand in the forecasted traffic set is routed over its fewest-hops path (O-E-O networks) or its shortest-distance path (optical-bypass-enabled networks). While a traffic forecast may not accurately predict the

traffic that will actually be supported in the network, it can be used as a reasonably good estimator of which links are likely to be heavily loaded. (Alternatively, one can combine the traffic forecast with the maximum-flow method of [KaKL00] to determine the critical links.)

It is important to consider not just single links that are likely to be bottlenecks, but also sequences of consecutive links that may be heavily loaded, and find routes that avoid the whole sequence of bad links. This is illustrated in Fig. 3.4, where Links BC, CD, and DZ are assumed to be likely bottlenecks. The shortest-distance path between Nodes A and Z is Path 1, which is routed over all three of the problem links. If one were to look for a path between these nodes that avoids just Link BC, then Path 2 is the remaining shortest-distance path. This is not satisfactory as the path still traverses Links CD and DZ. It would be better to simultaneously avoid all three bottleneck links and find Path 3.




Fig. 3.4 Links BC, CD, and DZ are assumed to be bottleneck links. Path 1, A-B-C-D-Z, crosses all three of these links. If Link BC is eliminated from the topology, the resulting shortest path, Path 2, A-E-F-C-D-Z, still crosses two of the bottleneck links. All three bottleneck links must be simultaneously eliminated to yield Path 3, A-E-F-G-H-Z.

After identifying the top 10 to 20 ‘hot spots’ in the network, the next step is to run the shortest path algorithm multiple times, where in each run, one bad link or one bad sequence of links is removed from the topology. (If a particular ‘hot spot’ does not appear in the lowest-cost paths for a source/destination pair, then that hot spot can be skipped. Furthermore, it is not desirable to remove *all* potentially bad links at once when running the shortest path algorithm, as the resulting set of paths may be circuitous and may shift the hot spots to other locations in the network.) This process finds paths that avoid particular bottleneck areas, if possible.

The bottleneck-avoidance strategy method can be combined with the K-shortest paths strategy such that, overall, the candidate path set includes lowest-cost paths and paths that are diverse with respect to the expected hot spots. There are clearly other methods one can devise to generate the candidate paths; however, the strategy described above is simple and is effective in producing designs for practical optical networks with relatively low cost and good load balancing.

3.4 Routing Strategies

The previous section discussed methods of producing a good set of candidate paths for all relevant source/destination pairs. This section considers strategies of selecting one of the paths to use for a given demand. These strategies are general such that they hold for scenarios where demand requests enter the network one at a time, as well as scenarios where there is a whole set of demands that needs to be routed but it is assumed that the demands have been ordered so that they are considered one at a time. (Ordering the demand set is covered in Section 3.7.) In either real-time or long-term planning, a particular candidate path may not be feasible as the network evolves because there is no free bandwidth on one or more of the path links. Furthermore, with real-time operation, there is the additional constraint that all necessary equipment to support the connection must already be deployed. Thus, if a particular path requires regeneration at a node, and the node does not have the requisite available equipment, the path is considered infeasible.

In optical-bypass-enabled networks, selecting a wavelength for the route is an important step of the planning process. This section focuses on selecting a route independent of wavelength assignment, where wavelength assignment is performed as a separate step later in the process. Chapter 5 considers treating both of these aspects of the planning process in a single step.

3.4.1 *Fixed-Path Routing*

In the strategy known as fixed-path routing, the candidate set of paths is generated prior to any demands being added to the network. For each source/destination pair, one path is chosen from the associated candidate path set. Ideally, the path is a lowest-cost path, although load balancing can also be a consideration for selecting a particular path. Whenever a demand request is received for a given source/destination pair, the selected path for that pair is used to route the demand; no other candidate paths for that source/destination pair are considered.

This is clearly a very simple strategy, with any calculations performed up-front, prior to any traffic being added. However, the performance of this strategy can be very poor, as it usually results in certain areas of the network becoming unnecessarily congested. The same path is always used for a given source/destination pair, providing no opportunity to adapt to the current network state. This often results in premature blocking of a demand even though feasible paths do exist for it. Overall, this strategy is not recommended.

3.4.2 *Alternative-Path Routing*

Alternative-path routing also relies on generating a candidate set of paths prior to any demands being added to the network. However, in this strategy, the candidate

set is narrowed down to M paths for each source/destination pair, for some small number M . When a demand request is received for a given source/destination, one of the M paths is selected to be used for that particular demand. In practice, selecting about three paths per each source/destination pair is a good strategy, although in real-time planning, where equipment availability is an issue, it may be desirable to select somewhat more paths.

In some research regarding alternative-path routing, it is assumed that the M paths must have no links in common; however, this is unnecessarily restrictive. (Selecting paths with no links in common, however, is important for protection, as is covered in Section 3.6.) The goal is to select the M paths such that the same expected ‘hot spots’ do not appear in all of the paths. Furthermore, it is not necessary to pick the M paths such that no ‘hot spot’ appears in any of the paths; this would have the effect of simply shifting the heavy load to other links as opposed to balancing the load across the network. Ideally, the M paths are lowest-cost paths; however, in order to get enough ‘hot-spot’ diversity among the paths, it may be necessary to include a path that does not meet the lowest cost; e.g., one of the M paths may have an additional regeneration. (Selecting a path with a small amount of extra cost is not ideal; however, it is typically preferable to blocking a demand request due to poor load balancing.) In addition, all other factors being equal (e.g., cost, expected load), paths with shorter distance and fewer hops should be favored for inclusion in the set of M paths.

When a demand request arrives for a particular source/destination, any of the M paths can be potentially used to carry the connection, assuming the path is feasible (i.e., it has the necessary available bandwidth and equipment). Typically, the current state of the network is used in determining which of the M paths to use. A common strategy is to select the feasible path that will leave the network in the ‘least-loaded’ state. Assume that the most heavily loaded link in the i^{th} path has W_i wavelengths already routed on it; the selected path is the one with the minimum W_i . If multiple paths are tied for the lowest W_i , then the load on the second most heavily loaded link in these paths is compared, and so on. (This is also known as *Least Congested Path* routing [ChYu94].) If multiple paths continue to be tied with respect to load, or if the tie is broken only when comparing links with load much less than the maximum, then one can consider other factors; e.g., in an optical-bypass-enabled network, the path with fewest hops can be used to break the tie. Furthermore, if one of the M paths has a higher cost than the other paths, then this path should not be selected unless its W_i is significantly lower than that of the other paths, or unless it is the only feasible path.

In another strategy, congestion and hops are considered together. For example, when choosing between two candidate paths, a path with H more hops is selected only if its most heavily loaded link has L fewer wavelengths routed on it, where the parameters H and L can be tuned as desired.

In real-time routing, it may also be beneficial to consider the available equipment at the nodes when selecting one of the M paths. If a particular path requires a regeneration at a node and there is very little free regeneration equipment at the

node, then that path may not be favored, especially if there are other paths that have similar link loading and greater equipment availability.

The alternative-path routing strategy works very well in practice. It uses the traffic forecast to assist in generating the initial candidate path set, whereas it uses current network conditions to select one of the M paths for a particular demand. One can add a larger dynamic component to the algorithm by allowing the set of M paths to be updated periodically as the network evolves. With a good choice of paths, the network is generally fairly well loaded before all M paths for a particular source/destination pair are infeasible. When this occurs, one can revert to dynamically searching for a path, as is covered next.

(A more restrictive form of alternative-path routing is known as *fixed-alternate routing*, where the M candidate paths are considered in a fixed order, and the first such path that has available capacity is selected. This is simpler than more general alternative-path routing because it only needs to track whether a path is available or not; it does not need to track the load on every link. However, this method is not as effective at load balancing.)

3.4.3 Dynamic-Path Routing

In dynamic-path routing (also called *adaptive unconstrained routing* [MoAz98]), there is no predetermination of which paths to use for a particular source/destination combination. The path calculation is performed at the time of each demand request, based on the current state of the network. The first step is to determine if there are any links in the network with insufficient available bandwidth to carry the new demand. Any such links should be temporarily eliminated from the network topology. In addition, in real-time design with an O-E-O network, any link for which either of its endpoints does not have an available transponder should be temporarily removed from the topology, because any link in the path would require a transponder at either end for O-E-O conversion. (Recall that in O-E-O networks, transponders are plugged into optical terminals, where each optical terminal is associated with a link.)

After the topology has been pruned based on the current network state (more advanced topology transformations are discussed in Section 3.5.1), the procedure for generating a candidate set of paths can be followed; i.e., the K-shortest paths algorithm can be run and/or the bottleneck-avoidance strategy can be used where the current hot-spots in the network are systematically eliminated from the already-pruned topology. (Note that the links with no available capacity have already been eliminated, so these hot spots are the relatively heavily loaded links that still have available capacity.) This process takes tens of milliseconds to complete, so it is possible to generate a candidate set of paths every time a new demand request is received (assuming the total provisioning time is on the order of one second or more). One can use a smaller 'K' in the K-shortest paths algorithm or consider fewer hot-spots in the bottleneck-avoidance methodology in order to

reduce the processing time further. Many dynamic implementations simply look for a single shortest path.

A variety of metrics can be used for dynamic routing. Typically, the metric reflects number of hops, distance, or current congestion; e.g., [BhSF01]. One method suggested in [ZTTD02] for optical-bypass-enabled networks uses a metric based on the number of wavelengths that are free on consecutive links, as this is an indicator of the likelihood of being able to assign wavelengths to the links. This was shown in [ZTTD02] to provide better performance than simply considering link congestion; however, it does involve a graph transformation in order to capture the relationship between adjacent links in the shortest paths algorithm.

After the candidate paths are generated, one path is selected for the new demand based on the current network state. (In implementations where only a single candidate path is generated, then clearly this step of choosing one of the candidate paths is not needed.) Note that in real-time planning, some of the candidate paths may be infeasible due to a lack of resources, even though some amount of topology pruning occurred up front. For example, in an optical-bypass-enabled network, a candidate path may require regeneration to occur at a node that does not have available regeneration equipment. (With optical bypass it is typically not known ahead of time whether an intermediate node in a path will require regeneration, thus the node is not pruned from the topology during the pre-processing step.) After eliminating any of the candidate paths that are infeasible, a technique such as selecting the least-loaded path, as described in Section 3.4.2, can be used to pick among the remaining paths.

The dynamic path selection methodology provides the greatest adaptability to network conditions. While this may sound good in theory, studies have shown that routing strategies that consider a more global design based on forecasted traffic can be advantageous [EMSW03]. Thus, a strictly adaptive algorithm is not necessarily ideal. Furthermore, the dynamic methodology potentially results in many different paths followed by the demands. This has the effect of decreasing the network ‘interference length’, which can potentially lead to more contention in the wavelength assignment process for optical-bypass-enabled networks. (The interference length is the average number of hops shared by two paths that have at least one hop in common, as defined in [BaHu96].) In addition, if the network makes use of wavebands, where groups of wavelengths are treated as a single unit, then the diversity of paths produced by a purely dynamic strategy can be detrimental from the viewpoint of efficiently packing the wavebands. Finally, the dynamic methodology is the slowest of the three routing strategies discussed and involves the most computation. It would likely not be suitable, for example, if the demand request must be provisioned in a sub-second time-frame.

Given the good results that are produced by the simpler alternative-path routing strategy of Section 3.4.2, it is often favored over a purely dynamic routing strategy. When the network is so full that none of the alternative paths are feasible, the dynamic strategy can be used instead.

3.5 Avoiding Infeasible Paths

3.5.1 Capturing the Available Equipment in the Network Model

In real-time planning, some, or even all, of the candidate paths may be infeasible due to a lack of available equipment in the appropriate nodes. As described above, the dynamic routing process first prunes out the links and nodes that would clearly be infeasible for a new demand. In a scenario where there is little available equipment, it may be necessary to perform more involved topology transformations to model the available equipment in more detail. In the examples below, it is assumed that there is available equipment at the source and destination nodes; otherwise, the demand is rejected without further analysis.

First, consider an O-E-O network with the topology shown in Fig. 3.5(a) and assume that a new demand request arrives where the source is Node A and the destination is Node Z. Node B and Node D of the network are illustrated in more detail in Fig. 3.5(b) and Fig. 3.5(c), respectively. In this example, it is assumed that pairs of transponders are interconnected via patch-cables rather than through a flexible switch (i.e., the nodal architecture is that of Fig. 2.3, not Fig. 2.4). Thus, in order for a new demand to transit Node B from Link j to Link k , ($1 \leq j, k \leq 4$), there must be an available transponder on the optical terminal for Link j , an available transponder on the optical terminal for Link k , and the two transponders must be interconnected. Assuming that Fig. 3.5(b) depicts all of the available equipment at Node B, then the only possible paths through the node for a new demand are between Links 1 and 2, Links 1 and 3, and Links 2 and 4. Similarly, the only possible paths through Node D are between Links 3 and 5 and between Links 5 and 6. It is assumed that the remaining nodes have sufficient available equipment to support any path; i.e., Node A has an available transponder on Link 1, Node Z has available transponders on both Links 4 and 6, and there are available transponders at Node C to support a path between Link 2 and Link 5.

To capture the path restrictions imposed by the limited amount of available equipment at Nodes B and D, one can perform a graph transformation where each link in the original topology becomes a node in the new topology. To be more precise, because each link shown in Fig. 3.5(a) actually represents bi-directional communication, each direction of a link becomes a node. These nodes are interconnected in the new topology only if there is equipment available in the real network to allow a new path to be routed between them. Nodes also have to be added to represent the source and destination of the new demand, i.e., Node A and Node Z, respectively.

The resulting transformed graph is illustrated in Fig. 3.6, where the node numbers in this graph correspond to the link numbers of Fig. 3.5. The single-prime nodes represent the links in the direction from the (alphabetically) lower letter to the higher letter, and the double-prime nodes represent the reverse link direction. Thus, Node 2' represents Link 2 in the original graph in the direction from Node B to Node C; Node 2'' represents Link 2 in the direction from Node C to Node B.

There is no need to add a node representing Link 1" because this link enters the demand source; similarly, there is no need to add a node representing Link 4" or Link 6" because these links exit the demand destination. Note that the node representing Link 1' is connected to the nodes representing Link 2' and Link 3', but not Link 4' due to the lack of a transponder pair connecting these links (in Node B). Similarly, there is no link connecting the node representing Link 3' and the node representing Link 6'.

A shortest path algorithm is run on the transformed topology to find a feasible path, using unity as the link metric to minimize the number of O-E-O conversions. The desired path from Node A to Node Z in the transformed graph is A-1'-2'-5'-6'-Z, which corresponds to path A-B-C-D-Z in the original graph.




Fig. 3.5 (a) Network topology, where a new demand is requested from Node A to Node Z. (b) The available equipment at Node B. (c) The available equipment at Node D.




Fig. 3.6 Graph transformation to represent the available equipment in the network of Fig. 3.5. The numbered nodes correspond to the links in Fig. 3.5 with the same number, with the prime and double-prime representing the two directions of the link. Nodes A and Z are the demand endpoints.

Routing constraints such as those imposed by the available transponder pairs, where only certain directions through a node are possible, are known as *turn constraints* (this term arises from vehicular routing, where only certain turns are permissible). The graph transformation described above is one means of solving such problems, which allows a standard shortest path algorithm to be run. An alternative is to use the original graph but modify the Dijkstra algorithm or the BFS algorithm to take into account the turn constraints when building out the path from source to destination [BoUh98, SoPe02]. With this methodology, an explicit graph transformation is not required.

Note that if the nodes are equipped with switches then turn constraints do not arise, because the switch can interconnect any two transponders in the node. Additionally, if there are many available transponders at each node, then this level of modeling is not needed as typically any path through the node can be supported.

In an optical-bypass-enabled network, with edge configurability, a different graph transformation can be used in real-time planning, similar to [GeRa04]. (The methodology is described first, followed by an example.) A new topology is created, comprised of only those nodes that have available regeneration equipment, along with the source and destination of the new demand. Any two of these nodes are interconnected by a link in the new topology if a regeneration-free path with available bandwidth exists between the two nodes. Even if there are multiple regeneration-free paths between a pair of nodes, only one link is added in the new topology. A shortest path algorithm is run on this transformed topology to find a feasible path. A link metric such as $[LARGE + NumHops]$ can be used, where *LARGE* is some number greater than any possible number of end-to-end hops, and *NumHops* is the number of hops in the minimum-hop regeneration-free path be-

tween the two nodes interconnected by the link. This finds a minimum-regeneration feasible path of minimum number of hops, assuming one exists.

An example of such a graph transformation is shown in Fig. 3.7. The full network is shown in Fig. 3.7(a). Assume that this is an optical-bypass-enabled network with edge configurability and an optical reach of 2,000 km. Furthermore, assume that Nodes A and Z are the demand endpoints, and that only Nodes B and D are equipped with available regeneration equipment. Figure 3.7(b) illustrates the associated transformed network; Nodes C and E do not appear in the transformed graph because they do not have available regeneration equipment. Each of the links in the transformed network represents a regeneration-free path in the real network. For example, Link AD in the transformed network represents path A-E-D in the real network. Running a shortest path algorithm on the transformed graph yields the path A-D-Z, corresponding to A-E-D-Z in the true network.




Fig. 3.7 In the network shown in (a), Nodes A and Z are assumed to be the endpoints of a new demand, and Nodes B and D are assumed to be the only nodes with available regeneration equipment. The optical reach is 2,000 km. This results in the transformed graph shown in (b). LARGE represents a number larger than any possible path hop count.

Such graph transformations as described above for O-E-O and optical-bypass-enabled networks would need to be performed every time there is a new demand request to ensure that the current state of the network is taken into account, which adds to the complexity of the routing process. Ideally, sufficient equipment is predeployed in the network and the candidate paths are sufficiently diverse that one can avoid these transformations, at least until the network is heavily loaded.

3.5.2 Predeployment of Equipment

Infeasible paths due to lack of equipment can be minimized by having equipment judiciously predeployed in the network in anticipation of future demands. Predeployed equipment in large-part refers to transponders (or other regeneration equipment) and the racks for housing them. Selecting how much equipment to

predeploy, as well as where to place the equipment, is strategically very important. Predeploying too little equipment leads to sub-optimal routing or excessive blocking of demand requests; predeploying too much equipment is an unnecessary expense. Clearly, the amount of equipment to predeploy depends on the underlying system characteristics; e.g., one would expect to predeploy more transponders in an O-E-O network than in an optical-bypass-enabled network. Furthermore, the different levels of configurability provided by the various system architectures mandate different levels of accuracy in the estimation process. For example, in an O-E-O architecture, one must calculate the number of transponders to predeploy on each optical terminal (i.e., a per-link estimate); in an optical-bypass architecture with edge configurability, it is necessary to estimate only the total number of transponders to predeploy at each node (i.e., a per-node estimate) [GeRa04].

Several strategies exist for estimating the amount of equipment to predeploy. A common strategy is to run simulations based on forecasted traffic, where the simulation results can be used to estimate the amount of equipment that should be predeployed at each node in order to reduce the blocking probability below a given threshold. Alternatively, one can use standard queuing theory, where each source/destination demand pair is associated with requiring particular equipment in the network. The arrival and departure processes of the demands can be modeled to estimate the required equipment to reduce the blocking probability below the desired threshold (e.g., [MoBB04]).

Another interesting strategy for optical-bypass-enabled networks, proposed in [BaLe02], involves estimating for each node the probability that any new demand will require regeneration at that node. The probabilities are determined based on the nodal position within the network and the lengths of the links feeding into the node, where being closer to the center of the network and being an endpoint of a long link increases the likelihood of regeneration at a node. This analysis can be used to assist in determining the amount of regeneration equipment to predeploy at each node.

The predeployment strategy to use may depend on the level of detail in the traffic forecast. If the forecast is very specific, then running a simulation is probably the simplest strategy to determine where to predeploy equipment. If there are only approximate models of the demand arrivals and departures, then queuing analysis can be used. If only a forecast of the total number of demands in the network is available, then the method based on nodal regeneration probabilities can be used.

In addition to intelligent predeployment, having network elements with a high degree of flexibility can be very helpful in maximizing the utility of the equipment that is deployed, as was described in Chapter 2.

3.6 Diverse Routing for Protection

The previous sections were focused on finding a single path between a source and a destination. If any of the equipment supporting a connection fails, or if the fiber over which the connection is routed is cut, the demand is brought down. Thus, it

is often desirable to provide protection for a demand to improve its availability. Numerous possible protection schemes are covered in Chapter 7. Here, it is simply assumed that two paths are required from source to destination, where the two paths should be disjoint to ensure that a single failure does not bring down the demand. If one is concerned only with link failures, then the two paths can be simply link-disjoint, where nodes can be common to both paths. If one is concerned with both link and node failures, then the two paths should be both link-and-node disjoint (except for the source and destination nodes). The question is how to find the desired disjoint paths in a network.

Network designers sometimes resort to the simple strategy of first searching for a single path using a shortest path algorithm. The links in the returned path are then pruned from the topology and the shortest path algorithm is invoked a second time. (If node-disjointness is required, then the intermediate nodes from the first found path are also pruned from the topology.) If a path is found with this second invocation, then it is guaranteed to be disjoint from the first path.

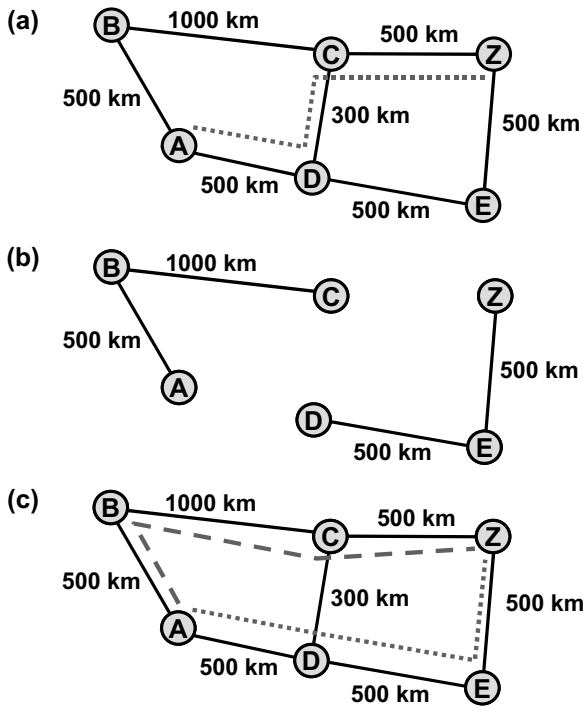


Fig. 3.8 The shortest pair of disjoint paths, with distance as the metric, is desired between Nodes A and Z. (a) The first call to the shortest path algorithm returns the path shown by the dotted line. (b) The network topology after pruning the links comprising the shortest path. The second call to the shortest path algorithm fails as no path exists between Nodes A and Z in this pruned topology. (c) The shortest pair of disjoint paths between Nodes A and Z, shown by the dotted and dashed lines.

While a simple strategy, it unfortunately fails in some circumstances. Consider the network topology shown in Fig. 3.8(a), and assume that two link-diverse paths are required from Node A to Node Z. The first invocation of the shortest path algorithm returns the path shown by the dotted line. Removing the links of this path from the topology yields the topology of Fig. 3.8(b). It is not possible to find a path between Nodes A and Z on this pruned topology, causing the strategy to fail. In fact, two diverse paths can be found in the original topology as shown in Fig. 3.8(c). This type of scenario is called a ‘trap topology’, where two sequential calls to the shortest path algorithm fail to find disjoint paths even though they do exist.

Even if the simple two-call strategy succeeds in finding two disjoint paths, the paths may not be optimal. In the network of Fig. 3.9(a), it is assumed O-E-O technology is used such that minimizing the number of hops is desirable. The fewest-hops path from Node A to Node Z is shown by the dotted line. The links of this path are pruned from the topology resulting in the topology shown in Fig. 3.9(b). The second call to the shortest path algorithm returns the path shown by the dashed line in this figure. While indeed disjoint, the two paths of Figs. 3.9 (a) and (b) cover a total of ten hops. However, the lowest-cost pair of disjoint paths has a total of only eight hops, as shown in Fig. 3.9(c).




Fig. 3.9 The shortest pair of disjoint paths, with hops as the metric, is desired between Nodes A and Z. (a) The first call to the shortest path algorithm returns the path shown by the dotted line. (b) The network topology after pruning the links comprising the shortest path. The second call to the shortest path algorithm finds the path indicated by the dashed line. The total number of hops in the two paths is ten. (c) The shortest pair of disjoint paths between Nodes A and Z, shown by the dotted and dashed lines; the total number of hops in the two paths is only eight.

As these examples show, the two-call strategy may not be desirable. It is preferable to use an algorithm specifically designed for finding the optimal pair of disjoint paths, as described next.

3.6.1 Shortest Pair of Disjoint Paths

The two best known Shortest Pair of Disjoint Paths (SPDP) algorithms are the Suurballe algorithm [Suur74, SuTa84] and the Bhandari algorithm [Bhan99]. Both algorithms involve calls to a regular shortest path algorithm; however, they require different graph transformations (e.g., removing links, changing the link metrics) to ensure that the shortest pair of disjoint paths is found. The graph transformations of the Bhandari algorithm may generate links with negative metrics, which is why it requires an associated shortest path algorithm such as BFS, which can handle graphs with negative link metrics.

Both the Suurballe and the Bhandari algorithms are guaranteed to find the pair of disjoint paths between a source and destination where the sum of the metrics on the two paths is minimized (assuming that at least one pair of disjoint paths exists). As illustrated by the examples of Fig. 3.8 and Fig. 3.9, the shortest single path may not be a part of the shortest-disjoint-paths solution. The run-times of the Suurballe and Bhandari algorithms are about the same; however, the latter may be more readily extensible to other applications [Bhan99]. Appendix B provides code for the Bhandari algorithm.

The SPDP algorithms can be used to find either the shortest pair of link-disjoint paths or the shortest pair of link-and-node-disjoint paths. To illustrate the difference, Fig. 3.10(a) shows the shortest pair of link-disjoint paths between Nodes A and Z, where the paths have Node B in common; together, the paths cover seven links and 700 km. Figure 3.10(b) shows the shortest link-and-node-disjoint paths between A and Z for the same topology; these paths cover eight links and 800 km.

Furthermore, the SPDP algorithms can be modified to find the shortest *maximally* link-disjoint (and optionally node-disjoint) paths when totally disjoint paths do not exist. Consider the topology shown in Fig. 3.11, where a protected connection is required between Nodes A and Z. A pair of completely disjoint paths does not exist between these two nodes. However, the maximally disjoint pair of paths, with one common link (Link DG) and two common nodes (Nodes D and G), is shown by the dotted and dashed lines in the figure. This pair of paths minimizes the number of single points of failures for the connection.

If a demand is very susceptible to failure, or the availability requirements (i.e., the fraction of time the demand must be in-service) are very stringent, then it may be desirable to establish more than two disjoint paths for the demand. The SPDP algorithms can be extended to search for the N shortest disjoint paths between two nodes, for any N , where the N paths are mutually disjoint. In most optical networks, there are rarely more than just a small number of disjoint paths between a given source and destination (say two to four); if N is larger than this, the algorithms can be used to return the N shortest maximally disjoint paths.




Fig. 3.10 (a) The shortest link-disjoint pair of paths is shown by the dotted and dashed lines. (b) The shortest link-and-node-disjoint pair of paths is shown by the dotted and dashed lines.




Fig. 3.11 There is no completely disjoint pair of paths between Nodes A and Z. The set of paths shown by the dotted and dashed lines represent the shortest maximally disjoint pair of paths. The paths have Nodes D and G, and the link between them, in common.

3.6.2 Shortest Pair of Disjoint Paths: Dual-Sources/Dual-Destinations

Another interesting twist to the problem of finding the shortest pair of disjoint paths arises when there are two sources and/or two destinations [Bhan99]. Consider the scenario shown in Fig. 3.12, where the source is in one regional network, the destination is in another regional network, and the two regional networks are interconnected by a backbone network. As shown in the figure, there are two nodes that serve as the gateways between each regional network and the backbone network. Each network is a separate domain, where routing occurs separately within each domain. It is desired that a protected connection be established between the source and destination. Thus, the paths in Regional Network 1 between the source and Gateways 1 and 2 should be diverse. Similarly, the paths in Regional Network 2 between the destination and Gateways 3 and 4 should be diverse. Thus, the paths in the backbone network between Gateways 1 and 2 and between Gateways 3 and 4 should be diverse.




Fig. 3.12 A protected connection between the source and destination is routed from Regional Network 1, through the backbone network, to Regional Network 2. The gateways are the nodes at the boundaries of the regional networks and the backbone network. It is desirable to have diverse paths from the source to Gateways 1 and 2, diverse paths between Gateways 1 and 2 and Gateways 3 and 4, and diverse paths from Gateways 3 and 4 to the destination.

This is one application for an algorithm that finds the shortest pair of disjoint paths between one source and two destinations, or, equivalently, two sources and one destination. (A similar problem arises in backhauling substrate traffic from a node that is not equipped with a grooming switch to two nodes that are equipped with grooming switches, where the two backhaul paths should be diverse, as discussed further in Chapter 6 in Section 6.5. Backhauling refers to the general process of transporting traffic from a minor site to a major site for further distribution.)

The problem setup is illustrated in Fig. 3.13(a), where Node A is the source, and Nodes Y and Z are the two destinations. (This figure is a general illustration;

it is not related to Fig. 3.12.) Finding the shortest pair of disjoint paths is quite simple. A ‘dummy’ destination node is added to the topology as shown in Fig. 3.13(b); Nodes Y and Nodes Z are connected to the dummy node via links that are assigned a metric of zero. An SPDP algorithm is then run using Node A as the source and the dummy node as the destination. This implicitly finds the shortest pair of disjoint paths from Node A to Nodes Y and Z.

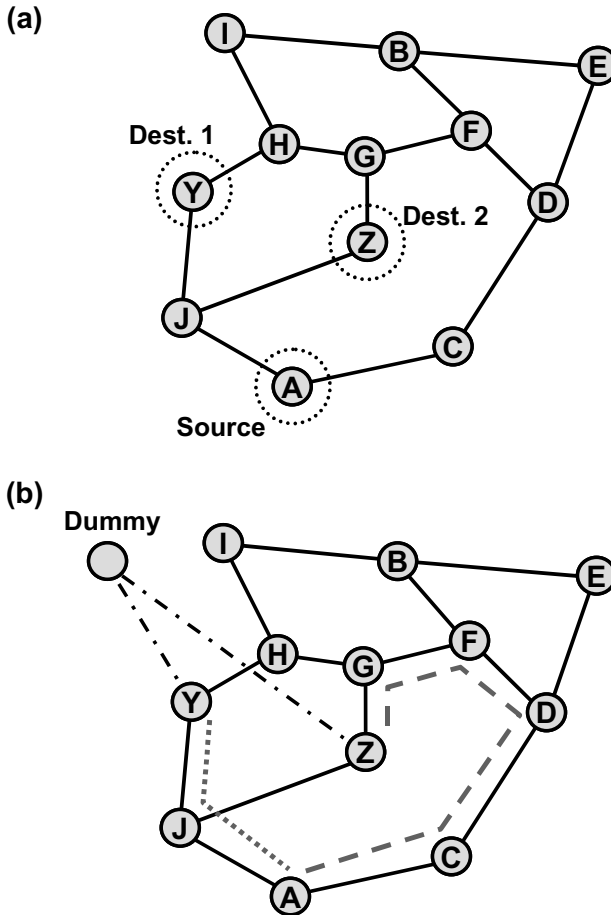




Fig. 3.13 (a) A disjoint path is desired between one source (Node A) and two destinations (Nodes Y and Z). (b) A dummy node is added to the topology and connected to the two destinations via links that are assigned a metric of zero. An SPDP algorithm is run between Node A and the dummy node to implicitly generate the desired disjoint paths, as shown by the dotted line and the dashed line.

(a) Dummy Dest.



(b)



(c)




Fig. 3.14 Diverse paths are required from two sources (Nodes A and B) to two destinations (Nodes Y and Z). (a) One dummy node is connected to the two sources and one dummy node is connected to the two destinations via links that are assigned a metric of zero. (b) In this solution, one path is between Nodes A and Y and the other path is between Nodes B and Z. (c) In this solution, one path is between Nodes A and Z and the other path is between Nodes B and Y.

If disjoint paths are desired between two sources and two destinations, then both a dummy source and a dummy destination are added, and the SPDP algorithm is run between the two dummy nodes. (This application is relevant to Fig. 3.12, where it is desired to find disjoint paths in the backbone network between the two sets of gateway nodes.) This procedure is illustrated in the example network of Fig. 3.14(a) where it is assumed a shortest pair of disjoint paths from Nodes A and B to Nodes Y and Z is required. One dummy node is connected to Nodes A and B, and the other to Nodes Y and Z, via links with a metric of zero. After running the SPDP algorithm between the two dummy nodes, it is interesting that the two paths that are found may be between Nodes A and Y and between Nodes B and Z, as shown in Fig. 3.14(b); or, the two paths may be between Nodes A and Z and between Nodes B and Y, as shown in Fig. 3.14(c).

3.6.3 Shared Risk Link Groups (SRLGs)

When searching for disjoint paths for protection, it may be necessary to consider the underlying physical topology of the network in more detail. Two links that appear to be disjoint when looking at the network from the link level may actually overlap at the physical fiber level. Portions of the two links may lie in the same fiber conduit such that a failure along that conduit would simultaneously disrupt both links. Links that are part of the same failure group comprise what is known as a Shared Risk Link Group (SRLG). (There may be network resources other than links that fail as a group. Thus, the more general term is Shared Risk Group (SRG).)

A common SRLG configuration, known as the fork configuration, occurs when multiple links lie in the same conduit as they exit/enter a node. This is illustrated in Fig. 3.15. The link-level view of the network is shown in Fig. 3.15(a), where it appears Links AB, AD, and AG are mutually disjoint. The fiber-level view is shown in Fig. 3.15(b), where it is clear that these three links lie in the same conduit at Node A. Thus, it would not be desirable to have a protected demand where, for example, one path includes Link AB and the other path includes Link AG, because the common conduit is a single point of failure. To find paths that are truly diverse requires that the SPDP algorithm be modified to account for the SRLGs, as described next.

In SPDP algorithms such as the Bhandari algorithm, the first step is to find the single shortest path from source to destination, which is then used as a basis for a set of graph transformations. If the source or destination is part of an SRLG fork configuration, and one of the links included in the SRLG lies along the shortest path that is found in the first step of the SPDP, then an additional graph transformation such as the one shown in Fig. 3.16 for Node A is required [Bhan99]. A dummy node is temporarily added to the topology and the SRLG links with an endpoint of Node A are modified to have the dummy node as an endpoint. The link metrics are kept the same. Another link, with a metric of zero, is added be-

tween Node A and the dummy node. The SPDP algorithm can then proceed. Because of the presence of the added node and link and the fact that disjoint paths are required, the shortest pair of disjoint paths will not include two links from the same SRLG fork configuration.




Fig. 3.15 (a) In the link-level view of the topology, Links AB, AD, and AG appear to be diverse. (b) In the fiber-level view, these three links lie in the same conduit exiting Node A, and thus are not diverse. A single cut to this section of conduit can cause all three of these links to fail.




Fig. 3.16 Graph transformation performed on Fig. 3.15 to account for the SRLG extending from Node A. A dummy node is added, and each link belonging to the SRLG is modified to have this dummy node as its endpoint instead of Node A. A link is added between Node A and the dummy node, where this link is assigned a metric of zero.

In addition to links with common conduit resulting in SRLGs, certain graph transformations may produce the same effect. Consider the graph transformation for capturing available regeneration equipment in an optical-bypass-enabled network that was presented in Section 3.5.1. In this transformation, links are added to a transformed graph where the links represent regeneration-free paths in the original graph. Links in the new graph may appear to be diverse that actually are not. This effect is illustrated in Fig. 3.17. The true network is shown in Fig. 3.17(a). Assume that the optical reach is 2,000 km and assume that a protected path is desired between Nodes A and Z. Assume that the nodes with available regeneration equipment are Nodes B, D, H, J, K, L, and M. The transformed graph is shown in Fig. 3.17(b). A link is added to this transformed graph if a regeneration-free path exists in the true network between the link endpoints, and the two endpoints have available regeneration equipment. The link metrics in the transformed graph are based on the number of hops in the regeneration-free path. The next step is to find the shortest pair of diverse paths in the transformed graph, which corresponds to the minimum-regeneration diverse pair of feasible paths in the true network (i.e., feasible with respect to the available equipment).

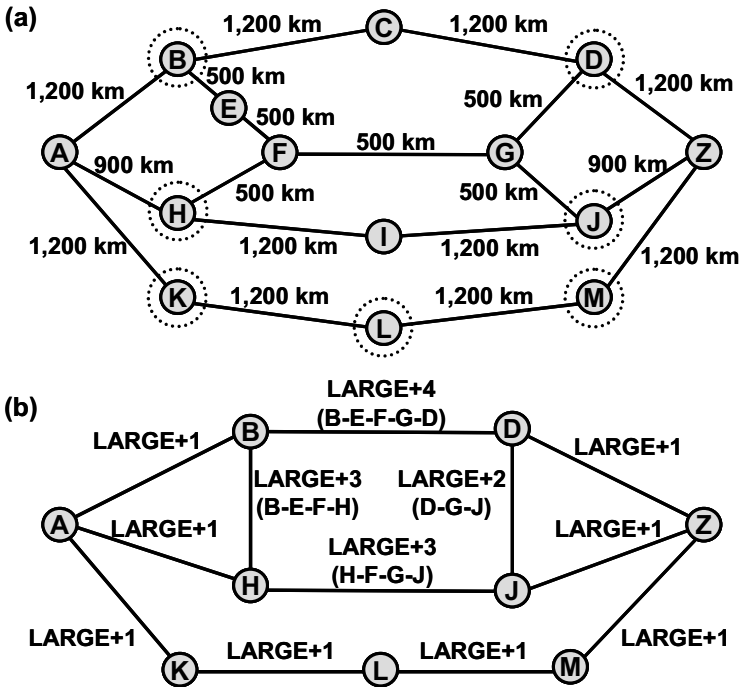


Fig. 3.17 (a) Example network, where the source and destination are A and Z, and the optical reach is 2,000 km. Only Nodes B, D, H, J, K, L, and M have available regeneration equipment. (b) Resulting transformed graph, where links represent regeneration-free paths. Links BD and HJ appear to be diverse but actually represent paths with a common link.

Links BD and HJ appear to be diverse in the transformed graph of Fig. 3.17(b). However, Link BD corresponds to path B-E-F-G-D in the real network and Link HJ corresponds to path H-F-G-J in the real network. Thus, both links correspond to paths that contain the link FG in the real network, implying that Links BD and HJ in the transformed graph comprise an SRLG. This type of SRLG formation, where only the middle portions of the links overlap, is referred to as the bridge configuration. (It is so-named because the configuration occasionally arises in actual network topologies due to a portion of the conduits of multiple links being deployed along the same bridge in order to cross a body of water.) To handle this SRLG scenario, the SPDP algorithm is run twice, each time eliminating one of the links comprising the bridge SRLG. The better result of the two runs is taken as the solution. In this example, this yields A-H-J-Z and A-K-L-M-Z in the transformed graph, corresponding to A-H-F-G-J-Z and A-K-L-M-Z in the real network.

To find the optimal pair of diverse paths in a network with multiple bridge configurations, all possible combinations have to be considered where only one link from each bridge is present. As the number of possibilities grows exponentially, this becomes intractable if the number of bridges is large. Furthermore, there are other classes of SRLGs, though not very common, for which finding an optimal routing solution is difficult, as described in [Bhan99]. In general, there are no computationally efficient algorithms that are guaranteed to find the optimal pair of disjoint paths in the presence of any type of SRLG; thus, heuristics are generally used when 'difficult' SRLGs are present in a network.

One such heuristic to handle arbitrary SRLGs is proposed in [XXQL03]. The first step in this heuristic is to find the single shortest path from source to destination; call this Path 1. The links in Path 1 are temporarily pruned from the topology. Furthermore, any link that belongs to an SRLG to which at least one of the links in Path 1 also belongs is assigned a very large metric to discourage its use. The shortest path algorithm is called again on the modified graph.

This second call may fail to find a diverse path. First, it may fail to find a path because of Path 1's links having been removed from the topology. In this case, the process restarts with a different Path 1 (one could use a K-shortest paths algorithm to generate the next path to use as Path 1). Second, it may fail because the path that is found, call it Path 2, includes a link that is a member of an SRLG that also includes a link in Path 1, in which case Path 1 and Path 2 are not diverse. (While the large metric discourages the use of such links in Path 2, it does not prevent it.) In this scenario, the most 'risky' link in Path 1 is determined; this is the link that shares the most SRLGs with the links in Path 2. The links in Path 1 are restored to the topology, with the exception of the risky link, and the process restarts; i.e., a new Path 1 is found on this reduced topology. The process continues as above, until a diverse pair of paths can be found, or until no more paths between the endpoints remain in the topology (due to risky links being sequentially removed), in which case it fails. If the procedure does find a diverse pair of paths, there is no guarantee that they are the shortest such paths; however, [XXQL03] reports achieving good results using this strategy, with reasonable run-time.

3.6.4 Routing Strategies With Protected Demands

Sections 3.3 and 3.4 covered generating candidate paths and using various routing strategies for unprotected demands (i.e., demands with a single path between the source and destination). In this section, these same topics are revisited, except for protected demands. Above, various methodologies were described to find a pair of disjoint paths between two nodes. These methodologies are used in this section to generate candidate paths. The link metric is again assumed to be unity for O-E-O networks and distance for optical-bypass-enabled networks (as noted previously, an alternative metric that is more closely tied to regeneration is discussed in Chapter 4).

As in Section 3.3, the initial step is to determine the links that are expected to be heavily loaded. This can be done by performing a preliminary design with all demands in the forecasted traffic set routed over their shortest paths, where in general the forecasted traffic will include both unprotected and protected demands. (An alternative means of estimating the critical links for protected traffic, based on maximum-flow theory, is detailed in [KaKL03]; this method can be combined with the traffic forecast to determine the links expected to be the most heavily loaded.) One can then generate a set of candidate paths for the protected demands by using the bottleneck-avoidance strategy, where the most heavily loaded links or sequence of links are systematically removed from the topology, with the SPDP algorithm run on the pruned topology. The goal is to generate a set of lowest-cost (or close to lowest-cost) disjoint path pairs that do not all contain the same expected ‘bad’ links. Note that a given source/destination pair may support both unprotected and protected demands. A candidate path set is independently generated for each source/destination/protection combination.

As noted in Section 3.2.2, the paths of minimum distance in an optical-bypass-enabled network do not necessarily translate to the paths of minimum regeneration. With single paths, this phenomenon is because regeneration typically must occur in network nodes as opposed to at arbitrary sites along the links (and because some nodes may not be equipped with elements that support optical bypass in all directions). This also applies to determining the regeneration locations for a pair of disjoint paths. Furthermore, the fact that regeneration is determined independently on the disjoint paths may also cause the minimum-distance pair of disjoint paths to have an extra regeneration. In Fig. 3.18, assume that a pair of link-and-node-disjoint paths is required between Nodes A and Z, and assume that the network supports optical bypass with an optical reach of 2,000 km. The shortest-distance pair of disjoint paths, shown in Fig. 3.18(a), is A-B-C-Z and A-D-E-Z. These two paths have a combined distance of 6,700 km, and require a total of three regenerations (at Nodes C, D and E). However, the minimum-regeneration pair of disjoint paths, shown in Fig. 3.18(b), is A-B-Z and A-D-C-Z, which covers a total of 7,000 km, but requires a total of only two regenerations (at Nodes B and D). As this example illustrates, each of the candidate pair of disjoint paths must be explicitly examined to determine the number of required regenerations.




Fig. 3.18 Assume that the optical reach is 2,000 km. (a) For a protected connection from Node A to Node Z, the combination of diverse paths A-B-C-Z and A-D-E-Z is the shortest (6,700 km), but requires three regenerations. (b) The combination of diverse paths A-B-Z and A-D-C-Z is longer (7,000 km), but requires two regenerations. (Adapted from [Simm06]. © 2006 IEEE)

Any of the three routing strategies discussed in Section 3.4 – fixed-path routing, alternative-path routing, and dynamic-path routing – are applicable to protected demands. (Variations of these strategies may be used for demands that share protection bandwidth, as covered in Chapter 7.) As with unprotected demands, alternative-path routing is an effective strategy for routing protected demands in practical optical networks.

3.7 Routing Order

In real-time network planning, demand requests generally are received and processed one at a time. In long-term network planning, however, there may be a set of demands to be processed at once. If the routing strategy used is adaptive, such that the network state affects the choice of route, then the order in which the demands in the set are processed will affect how they are routed. This in turn can affect the cost of the network design, the loading in the network, and the blocking probability. Thus, some attention should be paid to the order in which the de-

mands are processed. (Note that with global optimization techniques such as *integer linear programming*, the order of the demands is not relevant; however, such methodologies often have a very long run-time.)

One common strategy is to order the demands based on the lengths of the shortest paths for the demands, where the demands with longer paths are processed first. The idea is that demands with longer paths are harder to accommodate and thus should be handled earlier to ensure that they are assigned to optimal paths. This criterion can be combined with whether or not the demand requires protection, where protected demands are routed earlier as they require more bandwidth and there is generally less flexibility in how they can be routed.

This scheme often yields better results when combined with a round-robin strategy. Within a given ‘round’, the ordering is based on the required protection and the path length. However, at most one instance of a particular source/destination/protection combination is routed in each round. For example, if there are two protected demands between Node A and Node B, and three protected demands between Node A and Node C, where the A-B path is longer than the A-C path, plus one unprotected demand between Nodes A and D, then the routing order is: A-B, A-C, A-D, A-B, A-C, A-C.

Another strategy that is compatible with alternative-path routing is to order the demands based on the quality of the associated path set. If a particular source/destination/protection combination has few desirable path options, then the demands between this source and destination with this level of protection are routed earlier. For example, a particular path set may have only one path that meets the minimum cost. It is advantageous to route the associated demands earlier to better ensure that the minimum cost path can be utilized. In addition, the expected load on the links comprising the path set should be considered; the heavier the projected congestion for a particular source/destination/protection combination, the earlier it is routed.

Other ordering strategies are clearly possible. In general, no one strategy yields the best results in all network planning exercises. However, when using alternative-path routing, the routing process is so fast that multiple routing orders can be tested to determine which yields the best results. For example, routing thousands of demands with three different ordering strategies takes on the order of a couple of seconds. This is acceptable for long-term network planning.

3.8 Multicast Routing

Multicast traffic involves one source communicating with multiple destinations, where the communications is one-way. Multicast is also referred to as point-to-multipoint communications, in contrast to a point-to-point connection between a single source and single destination. (A tutorial on multicast routing can be found in [SaMu00].) The need for multicast could arise, for example, if the optical network is being used to distribute video simultaneously to multiple cities. Rather than setting up a separate unicast connection between the source and each of the

destinations, where multiple copies of the signal may be transmitted on a link, a multicast tree is constructed to reduce the amount of required bandwidth. The multicast tree connects the source to each of the destinations (without any loops), such that just one copy of the signal is sent on any link. This is illustrated in Fig. 3.19, where Node Q is the source and Nodes W, X, Y, and Z are the destinations. In Fig. 3.19(a), four separate unicast connections are established between Node Q and each of the destinations. Note that four connections traverse the link between Nodes Q and R. In Fig. 3.19(b), a single multicast tree is established, as shown by the dotted line, which requires significantly less bandwidth. Nodes R and T are branching nodes of this multicast tree.




Fig. 3.19 (a) Four unicast connections are established between the source, Q, and the destinations, W, X, Y and Z. (b) One multicast connection is established between Node Q and the four destinations.

A tree that interconnects the source and all of the destinations is known as a Steiner tree (where it is assumed all links are bi-directionally symmetric; i.e., two-way links with the same metric in both directions). The weight of the tree can be considered the sum of the metrics of all links that comprise the tree. Finding the Steiner tree of minimum weight is in general a difficult problem to solve (unless the source is broadcasting to every node in the network); however, several heuristics exist to find good approximate solutions to the problem [Voss92].

The heuristic of [KoMB81] combined with the enhancement of [Waxm88] is presented here, followed by a small example to illustrate the various steps. The original network topology is referred to here as A . The first step is to create a new topology, B , composed of just the source and destination nodes. All of the nodes are fully interconnected in B , where the metric of the link connecting a pair of nodes equals the metric of the shortest path between the two nodes in A . A minimum spanning tree is found on B using an algorithm such as Prim's [CoLR90]. (A minimum spanning tree is a tree that touches all nodes in the topology where the sum of the metrics of the links comprising the tree is minimized. This is different from a minimum Steiner tree in that *all* nodes are included, and is easier to solve.) Each link in the resulting minimum spanning tree, where a link connects

two nodes i and j , is expanded into the shortest path in the true topology between nodes i and j (e.g., if the shortest path in the true topology between nodes i and j has three links, then the link between nodes i and j in the new topology is replaced by the three links). Call the resulting topology B' . Another new topology is then formed, C , composed of all of the nodes in B' , with all of the nodes fully interconnected. The operations performed on topology B are repeated for topology C , resulting in the topology C' . A minimum spanning tree is then found on C' . Any links in the resulting tree that are not needed to get from the source to the destinations, if any, are removed, leaving the approximation to the minimum Steiner tree.




Fig. 3.20 (a) Original topology A , where the source and destination nodes are circled. (b) Topology B formed by interconnecting all source and destination nodes. (c) The minimum spanning tree on B expanded into paths, forming topology B' . (d) Topology C formed by interconnecting all nodes of B' . (e) The minimum spanning tree on C expanded into paths, forming topology C' .

This heuristic is illustrated with the small example of Fig. 3.20. Using the notation from above, the topology shown in Fig. 3.20(a) is the A topology. The source is Node W and the multicast destinations are Nodes X, Y , and Z . (The procedure is the same regardless of which of the four nodes is considered the source.) A fully interconnected topology composed of these four nodes is shown in Fig. 3.20(b). This is the B topology. The link metrics are based on distance here; e.g., the metric of Link WY , 150 km, is the shortest distance between Nodes

W and Y in the A topology. The minimum spanning tree on B is shown by the thick lines. Each link in this tree is expanded into its associated shortest path in the original topology. This produces the B' topology shown in Fig. 3.20(c). For example, Link XY in B is expanded into X-T-Y in B' . A new fully interconnected topology is formed with the five nodes of B' , as shown in Fig. 3.20(d). This is the C topology, where the minimum spanning tree is shown by the thick lines. The links of the tree are expanded into their associated shortest paths to form the C' topology, shown in Fig. 3.20(e). This topology is already a minimum spanning tree, thus, no further operations are needed, leaving C' as the final solution. In this example, the solution is optimal but that is not always the case.

A heuristic such as the one described can be used to generate an approximate minimum-cost multicast tree, where the link metric is set to unity for O-E-O networks and is set to link distance for optical-bypass-enabled networks. If it is known in advance the possible sets of nodes that will be involved in multicast demand requests, then one can pre-calculate the associated multicast trees. Furthermore, alternative multicast trees can be generated for each set of nodes where certain bottleneck links are avoided in each tree. The best tree to use would be selected at the time the multicast demand request arrives. If the multicast groups are unknown in advance, then the tree can be generated dynamically as the requests arrive. If it is necessary to add nodes to an existing multicast group, then greedy-like heuristics can be used to determine how to grow the tree [Waxm88].

As with unicast connections, the actual amount of regeneration in an optical-bypass multicast tree will likely depend on factors other than distance. For example, the branching points in the tree may be favored for regeneration. Refer to Fig. 3.21, which shows just the links included in the multicast tree of example Fig. 3.19, where the source is Node Q and the multicast destinations are Nodes W, X, Y, and Z. Assume that the optical reach is 2,000 km. In Fig. 3.21(a), the signal is regenerated at the furthest possible node from the source without violating the optical reach. This results in regenerations at Nodes S and T. If, however, the regeneration occurs at the branching point Node R as in Fig. 3.21(b), then no other regeneration is needed.




Fig. 3.21 Assume that the optical reach is 2,000 km. (a) Regeneration at both Nodes S and T. (b) Regeneration only at Node R.

An all-optical switch with multicast capability (Section 2.6.3) is very useful in a multicast tree. At a branching node where regeneration does not occur, e.g., Node T in Fig. 3.21(b), the switch can be used to all-optically multicast the signal onto the outgoing branches (the all-optical switch includes amplification such that the outgoing power level is not cut in half). Second, at a node with regeneration, e.g., Node R in the same figure, the all-optical switch can be used to multicast the regenerated signal onto both Links RS and RT. Finally, the drop-and-continue feature of the all-optical switch can be used at Node Y, where the signal drops and continues on to Node Z. (An OAMD-MD with multicast capability can perform the first and third of these functions. To allow multicast after a regeneration, a flexible regenerator card can be used, as is described in Section 4.5.2.)

3.8.1 Multicast Protection

Consider providing protection for a multicast connection such that the source remains connected to each of the destinations under a failure condition. One simplistic approach is to try to provision two disjoint multicast trees between the source and the destinations. However, because of the number of links involved in a multicast tree, it is often difficult, if not impossible, to find two completely disjoint trees.

Another approach is to find the minimum-cost ring path that traverses the source and each of the destinations [RaEl05]. The source can communicate to any destination using either the clockwise or counter-clockwise direction of the ring; thus, the source and each destination remain connected under any single failure. This idea is generalized to more complex protection topologies in [MFBG99].

A third approach is to arbitrarily order the destinations, and then sequentially find a disjoint pair of paths between the source and each destination, using an SPDP algorithm [SiSM03]. When looking for a disjoint pair of paths between the source and the $(i+1)^{\text{st}}$ destination, the links that comprise the path pairs between the source and the first i destinations are favored. The resulting sub-graph, which may not have a tree structure, provides protection against any single failure.

3.9 Routing with Inaccurate Information

The above sections on routing have presented numerous schemes for selecting a good path for each incoming demand. In long-term network planning, these strategies can be encompassed in a design tool that is run in a centralized location such as an operations center. In an on-line system, where the connection set-up time may play an important role in meeting customer requirements, an important consideration is where to locate the routing decision machinery. One strategy is to direct all incoming demand requests to a centralized controller that tracks the utili-

zation of network resources. For example, the IETF has defined a Path Computation Element (PCE)-based architecture, where the PCE is a powerful computing platform that is capable of performing constraint-based routing [FaVA06]. In one model, the PCE is located at a centralized node or server, and all routing requests are directed to it.

While centralized routing is advantageous from the standpoint of operating from a consistent picture of the network, this approach may incur excessive signaling delays in transmitting the request to the centralized location and disseminating the set-up information to the appropriate nodes. Furthermore, if the rate of demand requests is too high, the centralized controller could become a computational bottleneck.

Decentralized route computation is also possible. For example, in the IETF model, there could be multiple active PCEs in a network, where each PCE handles the route requests from a particular network region or from a particular domain. Alternatively, multiple PCEs, one per domain, could work together to select a multi-domain end-to-end route. In an extreme decentralized approach, each node is capable of calculating a route; e.g., the source node of a new demand selects the path to use. While such a distributed route selection process may be more responsive, this approach requires state information to be maintained by each of the network nodes.

The distributed routing process may be complicated by the inaccuracy of the state information. For example, consider the scenario where all nodes are capable of selecting a route. There are a variety of factors that can result in stale state information existing at one or more nodes of a dynamic network. First, updates on the current state may be distributed throughout the network only at certain time intervals as opposed to whenever any network change occurs, in order to reduce the amount of data that is flooded in the network. There is also a delay from the time the information is initially flooded until the time it is received by all of the relevant network nodes. Furthermore, in a large network, it may be desirable to distribute only summarized information about different portions of the network as opposed to the full details regarding resource utilization. All of these factors contribute to the nodes having to select a route with less than full knowledge of the current network state.

Routing with inaccurate information may require adding a probabilistic component to the routing strategy. For example, consider a relatively heavily loaded network where a subset of the links and nodes have little or no available resources. It is assumed that the state information used in the routing process may be somewhat stale such that the non-availability of a link or node is not known with certainty. In this scenario, it may be beneficial to search for a path based on the probability of it being feasible [GuOr99]. After performing a graph transformation as described in Section 3.5.1 to better model the network resources, each link in the transformed topology can be assigned a probability that it is available to carry the new demand. This probability can be estimated using, for example, the information regarding network resources (which may be out of date) and knowledge of the rate of resources being requested in certain areas of the network based

on prior history. The probability of a path being feasible can be estimated as the product of the probabilities assigned to each of the links comprising the path (resource usage on the links is unlikely to be completely independent such that the product of the probabilities is only an estimate of the probability of feasibility for the whole path). By taking the link metric to be the negative of the logarithm of its associated probability, and running a shortest path algorithm, the path with the highest estimated probability of being feasible is found.

Other methodologies for optical-bypass-enabled networks not only consider the number of wavelengths free on the links, but also consider which particular wavelengths are free to account for the wavelength continuity constraint (e.g., [MSSD03]).

Distributed routing schemes must also deal with resource contention due to concurrent demand requests being handled by multiple nodes. A path could be selected by a node that requires a particular wavelength on a given link only to find that by the time the connection establishment is initiated, a demand sourced by another node has also reserved that wavelength. One method of handling this complexity is through *contention resolution* schemes [GoLY02]; i.e., after the contention has occurred, an agreed upon rule is invoked to resolve the deadlock. For example, the GMPLS signaling specification stipulates that such contention be resolved based on the IDs of the nodes that are involved [Berg03]. Alternatively, *contention avoidance* methodologies can be implemented such that contention does not occur, or at least is minimized; this is generally preferable as it can result in smaller delays for time-critical functions such as dynamic failure restoration [GoLY02]. For example, various contention-avoidance schemes have been proposed to reduce the probability of reserving the same wavelength on the same link for two concurrent demand requests [ADHN01, OzPJ03, LiWM07]. Note that the development of intelligent reservation schemes that can be used to minimize contention without unduly blocking other demands is still an area of current research.

Chapter 4

Regeneration

Using the routing techniques of Chapter 3, a path is selected for each demand request entering the network. The next step in the planning process is selecting the regeneration sites for the demand, if any. (Chapter 5 considers techniques for accomplishing routing, regeneration, and wavelength assignment in a single step.) Regeneration ‘cleans up’ the optical signal, typically by reamplifying, reshaping, and retiming it; this is referred to as ‘3R’ regeneration. As discussed in Chapter 3, paths are usually selected to minimize the amount of required regeneration, as regeneration adds to the network cost.

If the network is based on O-E-O technology, then determining the regeneration locations for a selected path is straightforward because the demand is regenerated at every intermediate node along the path. For an optical-bypass-enabled network, where it is possible to transit an intermediate node in the optical domain depending on the quality of the optical signal, determining the regeneration locations for a path may be more challenging. Numerous factors affect when an optical signal must be regenerated, including the underlying transmission technology of the system, the properties of the network elements, and the characteristics of the fiber plant on which the system is deployed. When these factors are considered together, one can estimate the nominal distance over which an optical signal can travel before requiring regeneration (i.e., the optical reach). However, while quoting the optical reach in terms of physical distance is expedient for summarizing the system performance, it is not sufficient for determining the necessary regeneration locations in an actual network. Many factors other than distance need to be considered.

Section 4.1 presents a high-level discussion of some of the optical impairments and system properties that have an impact on when a signal must be regenerated. Ultra-long-reach technology has succeeded because it is possible to mitigate many of these impairments or design a system such that their effects are negligible. In Section 4.2, the discussion focuses on one of the more important impairments, noise, which leads to a metric that can be used in the routing process instead of distance to improve the likelihood of finding minimum-regeneration paths.

Given that regeneration, and thus implicitly network cost, are dependent on the physical layer design, it may be beneficial to integrate algorithms for physical network design (e.g., algorithms to optimize the amplifier configuration) with the architectural design and planning tool. Some of the benefits of this approach are discussed in Section 4.3.

In Section 4.4, the discussion moves from the physical layer aspects of regeneration to the architectural facets. Several regeneration strategies are presented, where the tradeoff is between operational simplicity and cost. The strategy employed will likely affect how much regeneration is required and where a connection must be regenerated; thus, it is an important aspect of the network design.

Finally, in Section 4.5, different options for actually implementing regeneration in a node are presented. Again, there is a tradeoff of operational flexibility versus cost. While most of this chapter is relevant only to optical-bypass-enabled networks, much of this final section applies to O-E-O networks as well.

There are a few key points to be emphasized in this chapter. First, while having to encompass physical layer phenomena in a network planning tool may seem imposing, there are known methodologies for tackling this problem that have been successfully implemented in live networks. Furthermore, when optical-bypass technology is developed by system vendors, there needs to be close collaboration between the systems engineers and the network architects. If the requirements of the physical layer are so exacting that the ensuing complexities in the network planning tool are unmanageable (e.g., if the wavelengths have to be assigned in a precise order on all links), then it may be necessary to modify the technological approach. Finally, while occasional regeneration may appear to be undesirable because of its cost, it does provide an opportunity to change the wavelength on which an optical signal is carried. This can be advantageous in the wavelength assignment step of the planning process, which is covered in Chapter 5.

4.1 Factors That Affect Regeneration

4.1.1 Optical Impairments

One of the major impairments that an optical signal encounters is accumulated noise. The principal source of noise is the spontaneous emissions of optical amplifiers, which are amplified along with the signal as they propagate together through the network. (This is suitably referred to as *amplified spontaneous emissions (ASE) noise*.) The strength of the signal compared to the level of the noise is captured by the signal's *optical-signal-to-noise-ratio (OSNR)*, where signals with lower OSNR are more difficult to receive without errors.

Many other optical impairments arise from the physical properties of light propagating in a fiber. For example, the propagation speed of light within a fiber depends on the optical frequency. This causes the optical signal pulses, which have a finite spectral width, to be distorted as they propagate along a fiber (gener-

ally, the effect is that the pulses spread out in time). This phenomenon is known as *chromatic dispersion*, or simply *dispersion*. Dispersion accumulates as a linear function of the propagation distance. Higher-bit-rate signals, where pulses are closer together, are more susceptible to errors due to this effect; the amount of tolerable dispersion decreases with the square of the bit-rate.

A different type of dispersion, known as *polarization-mode dispersion* (PMD), stems from different light polarizations propagating in the fiber at different speeds. In simplistic terms, an optical signal can be locally decomposed into two orthogonal principal states of polarization, each of which propagates along the fiber at a different speed, resulting in distortion. PMD accumulates as a function of the square root of the propagation distance. As with chromatic dispersion, PMD is a larger problem as the bit-rate of the signal increases; the PMD-limited reach decreases with the square of the bit-rate.

Several nonlinear optical effects arise as a result of the fiber refractive index being dependent on the optical intensity. (The refractive index governs the speed of light propagation in a fiber.) As the optical signal power is increased, these nonlinearities become more prominent. One such nonlinearity is *self-phase modulation* (SPM), where the intensity of the light causes the phase of the optical signal to vary with time. This potentially interacts with the system dispersion to cause significant pulse distortion. *Cross-phase modulation* (XPM) is a similar effect except that it arises from the interaction of two signals, which is more likely to occur when signals are closely packed together in the spectrum. Another nonlinear effect is *four-wave mixing* (FWM). This arises when signals carried on three particularly spaced optical frequencies interact to yield a stray signal at a fourth frequency, or two frequencies interact to generate two stray signals. These stray signals can potentially interfere with desired signals near or at these frequencies.

There are various other mechanisms of fiber nonlinearities that can distort the signal. For more detailed treatment of optical impairments, see [FoTC97, BaKi02].

4.1.2 Mitigation of Optical Impairments

There are well known techniques for mitigating some of these optical impairments. For example, dispersion compensation is typically used to combat chromatic dispersion, where the level of compensation needed depends on the amount of dispersion in the transmission fiber and the dispersion tolerance of the system. Note that it is not desirable to reduce the dispersion to zero, as its presence helps reduce some of the harmful nonlinear effects [Kurt93, TkCh94, TCFG95]. In one commonly used technique, dispersion compensating fiber (DCF) having inverse dispersion relative to the transmission fiber is installed at various sites along each link. DCF is generally expensive, high loss, and provides only a static means of compensation. Further problems with DCF may result from the dispersion level of the transmission fiber not being constant across the transmission band; typi-

cally, the dispersion level of the fiber has a particular slope across the band. The DCF may not have precisely the same inverse dispersion slope, leading to different levels of residual dispersion depending on the transmission wavelength.

As networks evolve, electronic dispersion compensation (EDC) is likely to be used as an enhancement to, or a replacement of, the DCF strategy [KaSG04]. EDC can be deployed on a per-wavelength basis (e.g., as part of the WDM transponder), and can be dynamically tuned over a range of dispersion levels to better match the compensation requirements of a given connection. For example, receivers based on maximum likelihood sequence estimation (MLSE) are an active area of research as a means of combating chromatic dispersion, as well as possibly other impairments [CaCH04, ChGn06]. (MLSE operates on a sequence of bits rather than a single bit at a time, and selects the data sequence that is statistically most likely to have generated the detected signal.) In another EDC strategy, pre-compensation is used at the transmitter, based on feedback from the receiver [MORC05]. Alternatively, post-compensation can be implemented at the receiver, which may be more suitable for dynamic networking.

PMD compensation is more challenging because the level of PMD may vary as a function of time. Cost-effective adaptive PMD compensators are an active area of research. There has typically been less of a need for PMD compensation compared to dispersion compensation. However, as wavelength rates continue to increase, it may require more attention, especially on older fibers. New fiber types being developed tend to have very low PMD.

Many of the problems from nonlinear effects can be avoided by maintaining the signal power at a low enough level (but still sufficiently higher than the noise level). In addition, as mentioned above, maintaining a small amount of residual system dispersion can be effective in reducing some of the nonlinear effects.

4.1.3 Network Element Effects

In addition to impairments that accumulate due to propagation in a fiber, there are a number of potential deleterious effects that a signal may suffer when transiting an optical-bypass-enabled network element. For example, a network element may utilize optical filters to internally separate the wavelengths entering from a WDM network port. Each time a signal passes through such a filter, the bandwidth of the channel through which the signal propagates ‘narrows’ to some degree, distorting the signal. Another source of signal degradation is the crosstalk caused by ‘leakage’ within a switching element. This occurs when a small portion of the input signal power appears at outputs other than the desired output. Additionally, the optical loss of a network element may depend on the state of polarization of the signal; this is known as *polarization dependent loss (PDL)*. Since the signal polarization may vary with time, the loss may also vary over time, which is undesirable. Furthermore, the network element may contribute to dispersion, and the dispersion level may not be flat across the transmission band, making compensation more difficult.

Factors such as filter narrowing, crosstalk, PDL, and dispersion contribute to a limit on the number of network elements that can be optically bypassed before needing to regenerate the signal. However, as optical-bypass technology has matured, the performance of the network elements has significantly improved. Many commercial systems support optical bypass of up to 10 (backbone) or 16 (metro) network elements prior to requiring regeneration. With this capability, the number of network elements bypassed is not usually the limiting factor in determining where regeneration must occur, especially in backbone networks where the distance between nodes may be very long; i.e., other limiting effects ‘kick-in’ prior to ten nodes being traversed. (However, the network elements do have an impact on the OSNR, as is discussed in Section 4.2.1.)

4.1.4 Transmission System Design

The characteristics of the transmission system clearly influence the optical reach of the system. One of the most important system design choices is the type of amplification, where Raman technology is generally used to attain an extended optical reach. Distributed Raman amplification uses the fiber itself to amplify the optical signal, so that the rate of OSNR degradation is less steep as compared to EDFA amplification. This trend is illustrated in Fig. 4.1, which depicts the OSNR level as a function of transmission distance, in a hypothetical system, for both distributed Raman and lumped EDFA amplification (lumped indicates the amplification occurs only at the amplifier sites).




Fig. 4.1 OSNR as a function of transmission distance for distributed Raman amplification and for lumped EDFA amplification, for both 80-km and 100-km amplifier spacings. The OSNR degrades more slowly with Raman amplification. For a given amplification type, the OSNR degrades more slowly with amplifiers spaced closer together. Many carrier networks have an amplifier spacing on the order of 80 km; however, there are some carrier networks with an average amplifier spacing closer to 100 km.

The acceptable OSNR level depends on the receiver sensitivity (i.e., the minimum average optical power necessary to achieve a specified bit error rate [RaSi01]) and the desired system margin (a network is typically designed to initially perform better than the minimum acceptable level to account for degradation as the system ages). As shown in Fig. 4.1, for a desired level of OSNR at the receiver and for a given amplifier spacing, Raman amplification supports a significantly longer transmission distance. More details on Raman technology can be found in [RoSt02].

Another important design choice is the signal modulation format, which is the format used for coding the data on a lightstream. A common modulation scheme is simple on-off keying (OOK), where the presence of light indicates a ‘1’ and the relative absence of light indicates a ‘0’. Two common signal formats used in conjunction with OOK are non-return-to-zero (NRZ) and return-to-zero (RZ), as shown in Fig. 4.2. The RZ format typically allows a longer optical reach. There are other formats of greater complexity (e.g., multi-level amplitude and/or phase modulation) that offer different tradeoffs in capacity, reach, complexity, and cost (e.g., see [Conr02]).

In addition to amplification and modulation, there are many other important system properties that affect optical reach, including: the spacing between channels (e.g., closer spacing reduces the reach); the initial launched powers of the optical signals (increasing the launched power increases the optical reach, up to a point; however, if the signal power is too high, the nonlinear optical impairments will have a negative impact on the reach, implying that there is an optimum value, which is system dependent); and the FEC coding strength (the stronger the FEC code, the greater its ability to detect and correct errors, which allows a longer reach; FEC was discussed in Section 2.9).




Fig. 4.2 (a) Non-Return-to-Zero (NRZ) signal format. (b) Return-to-Zero (RZ) signal format.

4.1.5 Fiber Plant Specifications

The characteristics of the physical fiber plant also have a large impact on system regeneration. Often, the optical reach of a given system depends on the fiber plant on which it is installed. For example, if the amplifier huts are spaced further apart, then the OSNR degrades more quickly leading to more frequent regeneration. This is illustrated in Fig. 4.1, where the OSNR decreases more sharply for 100-km amplifier spacing as compared to 80-km spacing, for a given amplification type.

The type of fiber in the network may have an impact as well, where the most commonly used fiber types are classified as either *non dispersion-shifted fiber* (NDSF) or *non-zero dispersion-shifted fiber* (NZ-DSF). (SMF-28[®] is an example of NDSF fiber; LEAF[®] and TrueWave[®] REACH LWP are examples of NZ-DSF fiber⁸.) As the names imply, these classes of fiber differ in the dispersion level in the portion of the spectrum occupied by WDM systems, thereby requiring different levels of dispersion compensation and having different mitigating effects on the nonlinear optical impairments. (For example, SMF-28 fiber has roughly four times the level of dispersion as LEAF fiber in the WDM region of interest.) In addition to dispersion differences, fiber types may differ in their Raman gain efficiency, where higher efficiency may lead to longer reach.

4.1.6 System Regeneration Rules

As the above discussion indicates, there are numerous factors to take into account when determining where an optical signal needs to be regenerated. Every vendor has a different approach to the problem, such that there is no uniform set of rules for regeneration across systems. It is generally up to the individual vendor to analyze their own particular implementation and develop a set of rules that govern routing and regeneration in a network. While there are many potential aspects to consider, in some systems it is possible to come up with a minimal set of rules that are sufficient for operating a network in real-time. For example, the system rules may be along the lines of: regenerate a connection if the OSNR is below N , or if the accumulated dispersion is above D , or if the accumulated PMD is above P , or if the number of network elements optically bypassed is greater than E (where N , D , P , and E depend on the system). When determining these rules, vendors usually factor in a system margin to account for aging of the components, splicing losses (i.e., optical losses that arise when fiber cuts are repaired), and other effects. Furthermore, it is important that the rules be immune to the dynamics of the traffic in the network. Connections are constantly brought up and down in a network, either due to changing traffic patterns or due to the occurrence of, or recovery from,

⁸ SMF-28 and LEAF are registered trademarks of Corning Incorporated; TrueWave is a registered trademark of Furukawa Electric North America, Inc.

failures; thus, it is important that the regeneration rules be independent of the number of active channels on a fiber.

While this may appear to be quite challenging, it is important to point out that there are optical-bypass-enabled networks with optical reaches over 3,000 km that have operated over a number of years using a set of relatively simple rules. It may be necessary to sacrifice a small amount of optical reach in order to come up with straightforward rules; however, as is examined more fully in Chapter 8, the marginal benefits of increasing the optical reach beyond a certain point are small anyway.

As pointed out in Section 4.1.5, the performance of a system depends on the characteristics of the fiber plant in the carrier's network. In the early stages of evaluating a system for a given carrier, prior to any equipment deployment, the exact specifications of the fiber plant may not be known. Thus, in the beginning stages of network design, one may have to rely on selecting regeneration locations based on path distance, where the optical reach in terms of a nominal distance is used. For purposes of system evaluation and cost estimation, this is generally acceptable. After the fiber spans have been fully characterized, the network planning tools can implement the more precise system regeneration rules.

4.2 Routing with Noise Figure as the Link Metric

Consider a system where the signal power levels are low enough that nonlinearities can be neglected (this is often the case in Raman-based systems). Assume that the fiber PMD is very low, such that the most important factors to consider with respect to regeneration are the OSNR, chromatic dispersion, and the number of network elements optically bypassed (due to the factors discussed in Section 4.1.3). The focus in this section is on OSNR; the other two factors are considered in Section 4.3.1 to ensure a cohesive system design.

As an optical signal propagates down a fiber link, its OSNR degrades. The *noise figure* (NF) of a link is defined as the ratio of the OSNR at the start of the link to the OSNR at the end of a link; i.e.:

$$NF_{Link} = \frac{OSNR_{Link\ Start}}{OSNR_{Link\ End}} \quad (4.1)$$

The noise figure is always greater than or equal to unity. Low noise figure is desirable as it indicates less signal degradation. Noise figure is a quantity that can be measured in the field for each link once the amplifiers have been deployed.

Numerous factors affect the noise figure of a link. For example, the type of amplification is very important, where Raman amplification generally produces a lower noise figure than EDFA amplification. Large fiber attenuation or large splicing losses can contribute to a higher noise figure. Longer span distances (i.e., the distances between amplifier huts) also generally increase the noise figure.

OSNR is important in determining when an optical signal needs to be regenerated. Transponder receivers generally have a minimum acceptable OSNR threshold, below which the signal cannot be detected properly. Thus, the evolution of OSNR as the signal traverses an optical path is critical. Consider the two consecutive links shown in Fig. 4.3, where each link i has an associated noise figure, NF_i , and a net gain, G_i . Net gain refers to the total amplification on the link compared to the total losses. The cumulative noise figure for an optical signal traversing Link 1 followed by Link 2 is given by:

$$NF_{Total} = NF_1 + \frac{(NF_2 - 1)}{G_1} \quad (4.2)$$

In most systems, the net gain on a link is unity, because the total amplification is designed to exactly cancel the total loss. In addition, typical values for the link noise figure are in the hundreds (using linear units), such that the '1' term is negligible. The formula then simplifies to:

$$NF_{Total} \approx NF_1 + NF_2 \quad (4.3)$$

Extending this formula to multiple links, the noise figure of an end-to-end path is the sum of the noise figures on each link of the path, where it is desirable to minimize this sum. Link noise figure is thus a suitable additive link metric that can be used in the shortest path algorithms of Chapter 3. Using this as a metric yields the path with least noise figure, or equivalently, the path with the highest overall OSNR (assuming other impairments are properly managed). Noise figure may be a better metric than distance in finding paths that minimize the number of required regenerations. However, it is still true that the minimum-noise-figure path may not be the minimum-regeneration path, e.g., due to regenerations occurring only at network nodes. Thus, in the process of generating candidate paths, as described in Chapter 3, each path must be evaluated to determine the actual regeneration locations, to ensure minimum-regeneration paths are chosen.

When working with the formulas for noise figure, it is important to use the correct units. The noise figure of a link is typically quoted in decibels (dB). However, the additive formula above requires that the noise figure be in linear units. The following formula is used to convert from dBs to linear units:

$$Linear Units = 10^{Decibel Units / 10} \quad (4.4)$$




Fig. 4.3 Two consecutive links, each with their respective noise figures and net gains.

(Note: If the system is more complex such that OSNR is not the primary determinant of where to regenerate, then one can use a performance factor, which is referred to as the Q-factor, as the link metric [KTMT05].)

4.2.1 Network Element Noise Figure

In addition to each link having a measurable noise figure, the nodal network elements contribute to OSNR degradation as well. Thus, each network element, such as an OADM or an all-optical switch, has an associated noise figure. To account for this effect in the routing process, the link metric should be adjusted based on the type of equipment deployed at either end of the link. (This is simpler than modeling the network elements as additional ‘links’ in the topology.) One strategy is to add half of the noise figure of the elements at the endpoints to the link noise figure. (This adjustment is used only for routing purposes; it does not imply that the noise figure of the element add/drop path is half that of the through path.)

Consider Link 2 shown in Fig. 4.4, which is equipped with an OADM at one end and an OADM-MD at the other end. Assume that the noise figure of the OADM is about 16 dB and the noise figure of the OADM-MD is about 17 dB. Halving these values yields roughly 13 dB and 14 dB, respectively (subtracting 3 dB is roughly equivalent to dividing by two). These amounts should be added to the noise figure of Link 2, where the additions must be done using linear units. For example, if Link 2 has a noise figure of 25 dB, then it should be assigned a link metric of: $10^{25/10} + 10^{13/10} + 10^{14/10} = 361.3$. By adding half of the element noise figure to each link connected to the element, the full noise figure of the element is accounted for regardless of the direction in which the element is traversed. (If, instead, the full amount of the element noise figure were added to each link entering the element, then the element noise figure would be double-counted along a path. If the full amount of the element noise figure were added to just one link entering the element, then some paths may not count the element penalty at all; e.g., if the penalty of the OADM-MD were added to Link 2 only, then a path from Link 3 to Link 4 would not include any OADM-MD penalty.)




Fig. 4.4 The noise figure of each link needs to be adjusted to account for the noise figure of the network elements at either end of the link. For example, the noise figure of Link 2 is incremented by half of the noise figure of an OADM and half of the noise figure of an OADM-MD.

4.2.2 Impact of the OADM without Wavelength Reuse

There is one network element that warrants special consideration with respect to noise and regeneration: the OADM that does not have wavelength reuse (see Section 2.4.3). Such elements are typically little more than an optical amplifier with a coupler/splitter for adding/dropping traffic at the node. A simplified illustration of such an OADM is shown in Fig. 4.5. An optical signal added at this node is coupled to the light passing through the node from the network links. Assume that the added signal in the figure is carried on λ_1 and assume that it is sent out on the East link. While there is no signal at λ_1 entering the OADM from the West link, the noise in the λ_1 region of the spectrum has propagated down the fiber and undergone amplification along with the rest of the spectrum. This noise will be coupled with the signal being added on λ_1 at the node. From an OSNR perspective, it appears as if the added signal has effectively been transmitted over some distance, thus affecting where it needs to be regenerated.




Fig. 4.5 A connection carried on λ_1 is added at a node equipped with an OADM without wavelength reuse. The added signal is combined with the WDM signal entering the node from the West link. The noise in the λ_1 region of the spectrum from the West link is added to the new connection's signal.

The effect on regeneration is illustrated more clearly in Fig. 4.6. Node B in the figure is equipped with an OADM without wavelength reuse, whereas the remaining nodes have an OADM with reuse. Assume that the nominal optical reach of the system is 2,000 km (for simplicity, we will still discuss reach in terms of a distance), and assume that a connection is established from Node B to Node Z. As the path distance from B to Z is 1,500 km, it would appear that no regeneration is required. However, the connection added at Node B has accumulated noise as if it originated at Node A. Thus, from an OSNR perspective, it is equivalent to a connection from Node A to Node Z, which covers 2,500 km. Therefore, a regeneration is required at Node C in order to clean up the signal.

This same effect does not occur with elements that have wavelength reuse because such elements are equipped with a means of blocking any wavelength along with the noise in the region of the spectrum around it. Thus, the noise is not combined with a wavelength being added at the node.




Fig. 4.6 Assume that the optical reach is 2,000 km. The OADM at Node B does not have wavelength reuse. A signal from Node B to Node Z has a level of noise as if it originated at Node A, and thus, needs to be regenerated at Node C.

Note that regenerating a connection at a node with a no-reuse OADM is not desirable. Consider the setup shown in Fig. 4.7, where Node C has a no-reuse OADM and the remaining nodes have OADMs with reuse. Assume that the optical reach is 1,000 km, and assume that a connection between Nodes A and Z is launched from Node A using wavelength λ_1 . If this connection is dropped at Node C for regeneration, λ_1 will continue on through the node because of the inability of the OADM at Node C to block the wavelength. Thus, after the signal is regenerated at Node C, it must be relaunched on a different wavelength, say λ_5 . This one connection will then ‘burn’ two wavelengths on the link between Nodes C and D (λ_1 and λ_5). Furthermore, because of noise being combined with the added signal, as described above, the signal on λ_5 will have a noise level as if it had been added at Node B. The added noise will cause the signal to require regeneration at Node D. Thus, the regeneration at Node C is not effective because one regeneration at Node D would have been sufficient for the whole end-to-end path. Additionally, the amount of add/drop supported at a no-reuse OADM is usually very limited. Using it for regeneration may prevent the node from sourcing/terminating additional traffic in the future. For these reasons, regeneration is generally not recommended at a node with a no-reuse OADM.

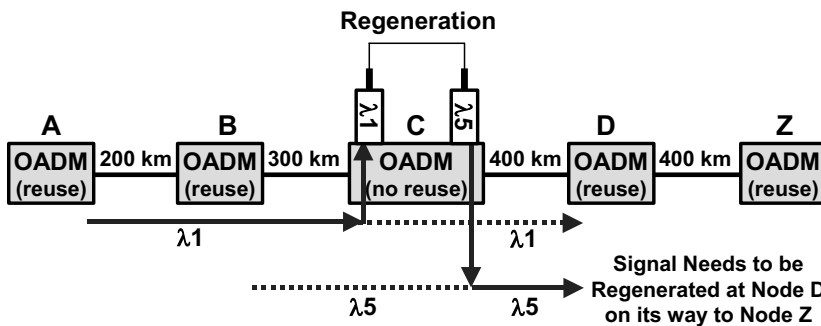


Fig. 4.7 The desired connection is between Nodes A and Z. Assume that the optical reach is 1,000 km. If the connection is regenerated at Node C, which has a no-reuse OADM, it will need to be regenerated again at Node D.

4.3 Link Engineering

Link engineering (or span engineering) is the process of designing the physical infrastructure for a particular network. For example, this may include determining the amplifier type for each site (e.g., pure Raman or hybrid Raman/EDFA), the gain setting of each amplifier, the amount of dispersion compensation, and the locations of the dispersion compensation. Algorithms specifically designed to optimize the system performance can be very helpful in this process. For example, *simulated annealing* [LaAa87] has proved to be one good technique for selecting the locations of the dispersion compensation modules. For a given set of amplifiers, proper setting of the gain and proper distribution of dispersion compensation can provide extra system margin (say on the order of 0.25 to 0.5 dB).

Clearly, the design choices in the physical layer have an impact on the network design at higher layers. It can be very advantageous to have a unified network design tool that incorporates the physical-layer design. For example, consider a system that has two types of amplifiers, where the Type A amplifier provides greater gain than the Type B amplifier, but it costs more. Installing the Type A amplifier as opposed to the Type B amplifier in certain sites would reduce the noise figure of the corresponding links. However, the effect on the overall amount of regeneration in the network may be small. By integrating physical-layer design with network planning, one can evaluate whether the extra cost of the Type A amplifier is justified by the expected reduction in regeneration. This allows better performance and cost optimization of the network as a whole.

4.3.1 Cohesive System Design

In Section 4.2, a network was considered where the major factors with respect to regeneration are assumed to be the OSNR, dispersion, and the number of network elements optically bypassed. (It is assumed that the OSNR of the network elements is taken into account in the system OSNR allowance.) In that section, the focus was on the OSNR, where routing is performed using noise figure as the link metric. Here, the dispersion and the network elements are considered to ensure there is a cohesive system design.

Assume that based solely on the OSNR analysis, it is anticipated that the optical reach of a system is on the order of D km. It is beneficial to co-ordinate the dispersion management and the network element performance to align with this number. For example, if the OSNR forces a path to be regenerated after roughly D km, then there is no need to add dispersion compensation suitable for transmission well beyond D km.

To be more concrete, assume that, based on the OSNR analysis, the optical reach is nominally on the order of 2,000 km. In addition, assume that the dispersion tolerance of the system is 10,000 ps/nm; i.e., after accumulating this much dispersion, a signal needs to be regenerated. Assume that the fiber has a disper-

sion level of 15 ps/nm·km, and assume that the network elements have negligible dispersion. Assuming DCF is used for dispersion compensation, then enough DCF should be added such that after 2,000 km, the 10,000 ps/nm limit is not exceeded. The required level of DCF can be determined by solving for C in the following equation:

$$2,000 \text{ km} \cdot (15 \text{ ps / nm} \cdot \text{km} - C) = 10,000 \text{ ps / nm} \quad (4.5)$$

This yields a C of 10 ps/nm·km, which is a guideline as to how much average dispersion compensation is needed per km of fiber.

In addition, assume that the nodes in the network are spaced such that the typical distance between nodes is 250 km. This implies that in traversing 2,000 km, it is not likely that more than seven or so network elements will be optically bypassed. This can serve as a guideline in designing the network elements, in terms of how many consecutive elements will need to be traversed optically.

As this example illustrates, by coordinating the different aspects of system design, the overall system can be made more cost effective.

4.4 Regeneration Strategies

The previous sections addressed some of the physical-layer factors that affect where regeneration is required. In this section, some of the architectural issues related to regeneration are considered.

There are several approaches to managing regeneration in a network. Three strategies are presented here, where the three differ with respect to flexibility, operational complexity, and cost. Although the above discussion has elucidated several factors other than distance that affect optical reach, the illustrative examples here will continue to refer to optical reach in terms of distance, for simplicity.

4.4.1 Islands of Transparency

In the architecture known as ‘islands of transparency’ [Sale98a, Sale00], a network is partitioned into multiple ‘islands’. The geographic extents of the islands are such that any intra-island (loopless) path can be established without requiring regeneration. However, a regeneration is required whenever a path crosses an island boundary, regardless of the path distance.

An example of a network partitioned into three islands is shown in Fig. 4.8. Island 1 is composed of Nodes A through H; regeneration is not required for a connection between any two of these nodes. Node A is also a member of Island 2, and serves as the regeneration point for any connection routed between Island 1 and Island 2. It is assumed Node A is equipped with two OADMs: one OADM is

Note that the isolation provided by the island paradigm typically exists with respect to the different geographic tiers of a network. For example, it is not common for traffic to be routed all-optically from a metro network into a backbone (or regional) network. The metro WDM system usually has coarser wavelength spacing and lower-cost components. The tolerances of the metro optics may not be stringent enough to be compatible with the system of the backbone network. Thus, the metro traffic typically undergoes O-E-O conversion prior to being carried on a backbone network regardless of the connection distance. (However, the metro and backbone tiers do not represent islands according to the definition above because regeneration may be required within either tier.)

4.4.2 Designated Regeneration Sites

A second architecture designates a subset of the nodes as regeneration sites, and allows regeneration to occur only at those sites. If an end-to-end connection is too long to be carried solely in the optical domain, then it must be routed through one or more of the regeneration sites. Either the designated regeneration sites are equipped with O-E-O equipment such that all traffic that transits them needs to be regenerated, or they have optical-bypass equipment so that regeneration occurs only when needed.

To ensure that a path is found that transits the necessary regeneration sites, one can use a graph transformation similar to the one discussed for optical-bypass-enabled networks in Section 3.5.1. (Other strategies are presented in [CMSG04, YaRa05a].) In that section, the transformation was discussed in the context of real-time planning, where only certain nodes may have available regeneration equipment. The transformed graph is comprised of only those nodes with available equipment; a link between two of these nodes is added only if a regeneration-free path exists between them in the true topology. Finding a path in this transformed graph guarantees regeneration feasibility. Restricting regeneration to certain nodes in a network is an equivalent problem, where only those nodes that are designated as regeneration sites appear in the transformed graph.

One strategy for judiciously selecting the regeneration sites is to first route all of the forecasted traffic over its shortest path. A greedy type of strategy can then be employed where nodes are sequentially picked to be regeneration sites based on the number of paths that become feasible with regeneration allowed at that site [CSGJ03]. Enough nodes are picked until all paths are feasible.

Another strategy for selecting the regeneration sites, proposed in [CSGJ03], is based on *connected dominating sets (CDSs)*. A dominating set of a graph is a subset of the nodes, S , such that all nodes not in S are directly connected to at least one of the nodes in S . The dominating set is connected if there is a path between any two nodes in S that does not pass through a node not in S . To use this methodology, the first step is to create a new topology that contains all network nodes, where two nodes are connected by a link only if there is a regeneration-free path between them in the true topology. On this new topology, a minimal CDS is

found, using heuristics such as those described in [GuKh98]. (The minimal CDS is the CDS with the fewest number of nodes.) The nodes in the minimal CDS are designated as the regeneration sites. By the definition of a CDS and by virtue of how links are added to this topology, any node not in the CDS is able to reach a node in the CDS without requiring intermediate regeneration. This guarantees that for any source/destination combination, a path exists that is feasible from a regeneration standpoint.

Because this architecture limits regeneration to a subset of the nodes, extra regeneration may occur, as illustrated in Fig. 4.9. Assume that the optical reach is 1,000 km, and assume that regeneration is permitted only at Nodes B, D, and E. Assume that the nodes are equipped with optical-bypass elements, including the regeneration-capable sites. A connection between Nodes A and Z is ideally regenerated at just Node C. However, because regeneration is not permitted at Node C, the connection is regenerated at both Nodes B and E (or at Nodes B and D), resulting in an extra regeneration.




Fig. 4.9 Assume that the optical reach is 1,000 km. A connection between Nodes A and Z is ideally regenerated in just Node C. However, because it is assumed that regeneration is not permitted at this site, the connection is regenerated at both Nodes B and E instead.

Moreover, if the designated regeneration sites are equipped with O-E-O network elements (e.g., back-to-back optical terminals), then the amount of excess regeneration is likely to be significantly higher, as any connection crossing such a node needs to be regenerated. The extra regeneration cost may be partially offset by somewhat lower network element costs. Consider two adjacent degree-two network nodes where the two nodes are close enough together such that any required regeneration could equivalently occur in either node. Assume that all of the transiting traffic needs to be regenerated in one of the two nodes. In Scenario 1, assume that both nodes are equipped with OADMs and assume that they both regenerate 50% of the traffic that traverses them. In Scenario 2, one of the nodes has an OADM and regenerates none of the traffic, and the other node has two optical terminals and regenerates 100% of the transiting traffic. The total amount of regeneration is the same in either scenario. However, because two optical terminals cost less than an OADM, Scenario 2 is overall less costly. (This is an extreme example because it assumed 100% of the transiting traffic needed to be regenerated.) Setting up dedicated regeneration sites (i.e., regeneration sites with O-E-O elements) in a network may have this same effect to a degree, but the overall cost is still likely to be higher because of the extra regeneration.

A possible benefit of designating only certain nodes as regeneration sites is more streamlined equipment predeployment. With regeneration occurring at a limited number of sites, the process of predeploying equipment is more economical; i.e., with fewer pools of regeneration equipment, it is likely that less regeneration equipment needs to be predeployed across the network to reduce the blocking probability below a given threshold.

Overall, given that regeneration can be accomplished without any special equipment (i.e., the same transponders used for traffic add/drop can be used for regeneration), it is not clear if it is necessary to severely limit the number of nodes at which regeneration is supported. The biggest disadvantage is it makes the routing process more challenging and less flexible, and it likely costs more depending on the amount of extra regenerations it produces. While it may make sense to eliminate some nodes as possible regeneration sites, because the nodal offices are not large enough to house a lot of terminating equipment or because the nodes are equipped with OADMs without reuse, in general, designating only a subset of the nodes for regeneration may not be optimal.

4.4.3 Selective Regeneration

The third architecture is *selective regeneration*, which allows any (or almost any) node to perform regeneration, and regenerates a demand only when needed. The decision as to whether regeneration is needed, and if so, where to implement it, is made on a per-demand basis. This strategy is the one most commonly used in actual network deployments. Given the freedom in selecting regeneration locations, this approach yields the fewest regenerations (assuming enough available regeneration equipment is deployed at the nodes) and also allows the most flexibility when routing.

Often, there will be several options as to where the regeneration can occur for a given connection. Consider a connection between Nodes A and Z in Fig. 4.10. Assume that the optical reach is 1,000 km, and assume that regeneration is possible in any of the nodes. The minimum number of regenerations for the connection is two. As shown in the figure, there are three possible scenarios that yield two regenerations: regenerate at Nodes B and D, regenerate at Nodes C and D, or regenerate at Nodes C and E.




Fig. 4.10 Assume that the optical reach is 1,000 km, and assume that regeneration is permitted at any node. A connection between Nodes A and Z can be regenerated at Nodes B and D, at Nodes C and D, or at Nodes C and E.

There are several factors that should be considered when selecting one of these regeneration scenarios for the A-Z connection. In real-time planning, the amount of available equipment at each node should be considered, where regeneration is favored in the nodes that have more free equipment. Additionally, if the nodes are equipped with network elements that have a limit on the total add/drop (see Section 2.4.1), then it is important to favor regeneration at the nodes that are not close to reaching this limit. (Reaching the maximum amount of add/drop at a node could severely impact future growth, as that node will not be able to source/terminate more traffic.)

One needs to also consider the subconnections that will result from a particular regeneration scenario. The term *subconnection* is used here to refer to the portions of the connection that fall between two regeneration points or between an endpoint and a regeneration point. For example, if the connection in the figure is regenerated at Nodes C and D, the resulting subconnections are A-C, C-D, and D-Z. By aligning the newly formed subconnections with those that already exist in the network (i.e., producing subconnections with the same endpoints, on the same links), the wavelength assignment process may encounter less contention. Furthermore, in a waveband system where bands of wavelengths are treated as a single unit, creating subconnections with similar endpoints yields better packing of the wavebands.

One could also consider the system margin of the resulting subconnections when selecting where to regenerate. Regenerating the connection at Nodes C and E produces a subconnection with a length of 900 km. With the other two regeneration options, no subconnection is longer than 800 km, such that there is somewhat greater system margin. Of course, if the optical reach is specified as 1,000 km, then any of these options should work. Thus, while taking the resulting distances of each subconnection into account may produce extra margin, adding this consideration to the algorithms should not be necessary.

Note that some of the aforementioned factors may be at odds with one another. The desire to align the subconnections favors continuing to regenerate at the same nodes, but the add/drop limits of the equipment may require regeneration to be more dispersed. Given that reaching the add/drop limit at a node could restrict network growth, the limits of the network elements, if any, should dominate as the network becomes more full.

4.5 Regeneration Architectures

This chapter has thus far covered many of the physical effects, as well as the architectural strategies, that affect where a given connection must be regenerated. All of the relevant factors must be incorporated in the network planning tool to ensure that regeneration sites are selected as needed for each demand. This section examines how regeneration is actually implemented within a node.

4.5.1 Back-to-Back WDM Transponders

As has been pointed out several times already, one means of regenerating a signal is to have it exit the optical domain on one WDM transponder and re-enter the optical domain on a second WDM transponder. Figure 4.11 illustrates this architecture where the pairs of transponders used for regeneration are connected via a patch cable. The process of O-E-O conversion typically achieves full 3R regeneration, where the signal is reamplified, reshaped, and retimed. It usually provides an opportunity to change the wavelength of the optical signal as well.




Fig. 4.11 Regeneration via back-to-back transponders (TxRx) that are interconnected by a patch cable. (a) The edge configurability provided by the all-optical switch allows any regeneration pair to access any two network links. (b) The OADM-MD does not have edge configurability; thus, in this example, regeneration is possible only in the South/West and East/South directions. (Adapted from [Simm05]. © 2005 IEEE)

The flexibility of the back-to-back transponder architecture largely depends on the capabilities of the equipment. Consider regeneration of a connection entering from the East link and exiting on the West link. If the corresponding transponders connected to the East and West links are fully tunable, then each transponder can be independently tuned such that any (East/West) input/output wavelength combination is supported. If the transponders are not tunable, then the possible input/output wavelength combinations depend on the wavelengths of the transponders that are cabled together.

If the network element has edge configurability, as illustrated by the all-optical switch of Fig. 4.11(a), then any pair of back-to-back transponders can access any two of the network links. For example, in the figure, any pair of transponders can be used to regenerate between the East/West links, the East/South links, or the South/West links. Contrast this with Fig. 4.11(b), where the OADM-MD does not have edge configurability. With this network element, the possible regeneration directions are determined by which pairs of transponders are interconnected. Thus, in the figure, regeneration is possible between the South/West links and between the East/South links, but not between the East/West links. Manual intervention is required to rearrange the patch cables to allow for different configurations.

In Fig. 4.11, one limitation is that the transponders must be partitioned between the ‘true’ add/drop function and the regeneration function. (As mentioned previously, a regeneration can be considered add/drop traffic because the signal drops from the optical layer; the term ‘true add/drop’ is used here to distinguish those signals that actually originate from, or terminate at, the node.) Only those transponders that are cabled together can be used for regeneration, where manual intervention is required to adjust the apportionments. To address this, one can use an adjunct edge switch, as shown in combination with an OADM-MD in Fig. 4.12.

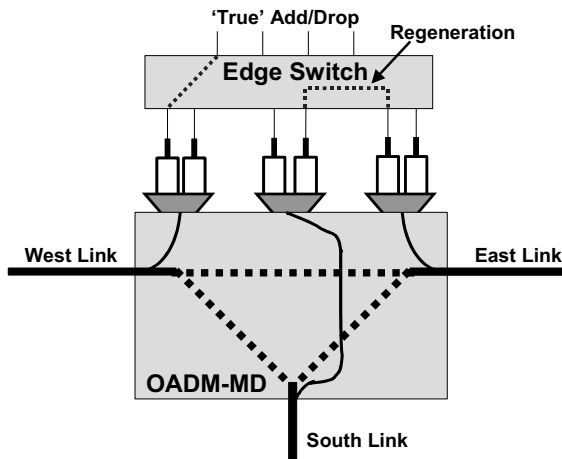


Fig. 4.12 An adjunct edge switch provides additional flexibility at the node. Any transponder can be used for either ‘true’ add/drop or regeneration; any regeneration direction through the node is supported. (Adapted from [Simm05]. © 2005 IEEE)

Adding the edge switch allows any transponder to be used for either ‘true’ add/drop or regeneration, depending on the switch configuration. While this architecture incurs the cost of the edge switch, it reduces the number of transponders that have to be predeployed at a node and/or it reduces the amount of manual intervention required to configure the node. Note that the presence of the adjunct switch also provides edge configurability for the OADM-MD. (Using an adjunct switch with an OADM-MD to provide edge configurability was discussed in Section 2.8.1.) Thus, it allows the transponders to be used to regenerate in any direction through the node, as could already be achieved with the all-optical switch.

4.5.2 Regenerator Cards

A WDM transponder converts an incoming WDM-compatible optical signal to a 1310-nm optical signal. When two transponders are cabled together for regeneration, they communicate via the 1310-nm signal. However, for regeneration purposes, this conversion is unnecessary. A simpler device capable of 3R regeneration, referred to as a *regenerator*, is shown in the architectures of Fig. 4.13. The received optical signal on one side of the regenerator is converted to an electronic signal that directly modulates the optical transmitter on the other side, eliminating the need for short-reach interfaces (i.e., the 1310-nm interfaces). The motivation for this device is cost; the cost of one regenerator card is roughly 70% of the cost of two transponder cards.

Regenerator cards cannot be used for ‘true’ add/drop, as they lack the short-reach interface; thus, there is a clear division between the add/drop equipment and the regeneration equipment. The attributes of the regenerator card can have a profound impact on network operations. Consider using a regenerator card for a connection entering from the East link and exiting on the West link. It is desirable for the regenerator to allow the incoming wavelength from the East link to be different from the outgoing wavelength on the West link (the same also applies for traffic going in the reverse direction). Otherwise, wavelength conversion would be prohibited from occurring in concert with a regeneration, which is a significant restriction. Furthermore, the regenerator card is ideally fully tunable, such that any combination of incoming and outgoing wavelengths can be accommodated with a single card. If the regenerator cards are not tunable, then inventory issues become problematic if every combination of input and output wavelengths is potentially desired. (Typically, system vendors and/or carriers maintain an inventory of equipment to support new traffic or replace failed equipment. Storing thousands of different regenerator combinations would be impractical.)

Regenerator cards can be used with an all-optical switch, as shown in Fig. 4.13(a), or with an OADM-MD, as shown in Fig. 4.13(b). As with the back-to-back transponder architecture, the all-optical switch allows a regenerator to be used for regeneration in any direction through the node. In the OADM-MD, the regenerator is tied to a particular regeneration direction (e.g., East/West in the figure).

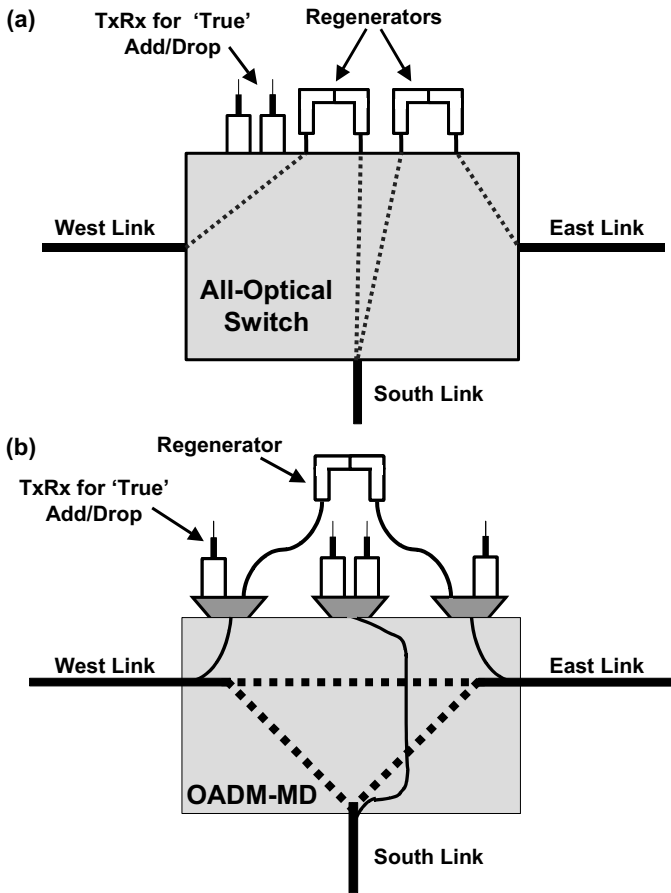


Fig. 4.13 (a) Regenerators used in conjunction with an all-optical switch. (b) Regenerator used in conjunction with an OADM-MD. In the latter, regeneration is supported only in the East/West direction. (Adapted from [Simm05]. © 2005 IEEE)

Deploying an adjunct edge switch with an OADM-MD to gain flexibility with the regenerator card is inefficient with respect to switch port utilization. The configuration is shown in Fig. 4.14. The incoming WDM signal is demultiplexed and the constituent wavelengths are fed into the edge switch. The signals that need to be regenerated are directed by the edge switch to regenerator cards, which ideally are tunable. (The signals that are truly dropping at the node are fed into transponders.) The edge switch allows wavelengths from any two network ports to be fed into a particular regenerator, thereby providing flexibility in the regeneration direction. For example, in the figure, the switch is configured to enable a regeneration in the East/South direction. However, to accomplish this, note that four

ports are utilized on the edge switch for a regeneration. With the configuration shown in Fig. 4.12, only two ports are utilized on the edge switch per regeneration. Additionally, in Fig. 4.14, the edge switch must be capable of switching a WDM-compatible signal; e.g., it could be a MEMS-based switch. An electronic-based switch is not suitable for this application.

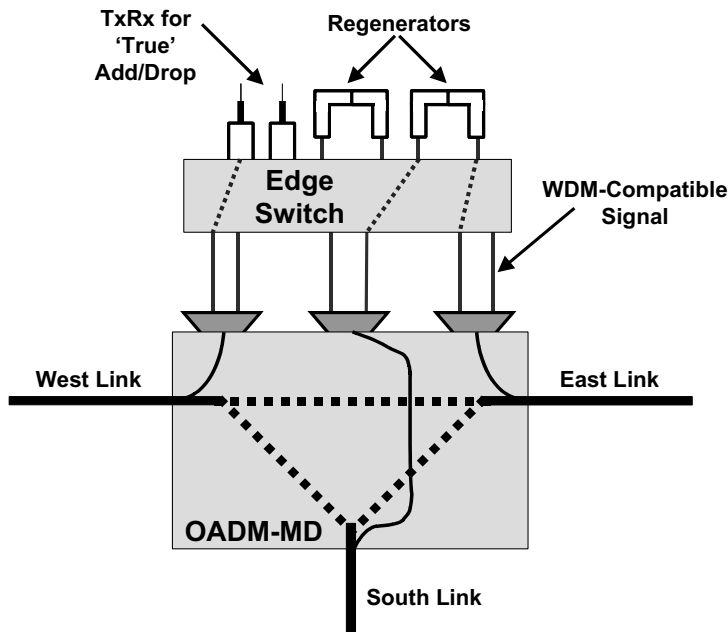


Fig. 4.14 With an edge switch used in combination with regenerator cards, four ports on the switch are utilized for each regeneration. Additionally, the edge switch must be capable of switching a WDM-compatible signal.

Rather than using an edge switch, a degree of configurability can be attained with the OADM-MD architecture through the use of flexible regenerators [SiSa07], similar to the flexible transponders that were discussed in Section 2.8.2. Refer to the flexible regenerator shown in Fig. 4.15 with a degree-four OADM-MD. One side of the regenerator is connected to the East and North links, and the other side is connected to the West and South links. Thus, this one regenerator allows regeneration to occur in either the East/West, East/South, North/West, or North/South directions; i.e., four of the six possible directions through the node are covered. Furthermore, it allows a regenerated signal to be sent out on two simultaneous links; e.g., a signal entering from the East link can be regenerated and sent out on both the West and South links. This is useful for multicast connections, as was discussed in Section 3.8.




Fig. 4.15 A degree-four OADM-MD combined with a flexible regenerator that allows regeneration in either the East/West, East/South, North/West, or North/South directions. An optical backplane can be used to eliminate complex cabling. (Adapted from [SaSi06]. © 2006 IEEE)

4.5.3 All-Optical Regenerators

An all-optical regenerator card is a relatively new regeneration alternative [LLBB03, LeJC04]. Due to the scalability of optics, it is expected that the cost of all-optical regenerators will scale well with increasing line-rate. For example, there may be only a small price premium for a 40-Gb/s regenerator as compared to a 10-Gb/s regenerator. Furthermore, all-optical regenerators are expected to consume less power as compared to their electronic counterparts, which should improve nodal scalability.

All-optical regenerators may not fully replace electronic regeneration. Some all-optical regenerators provide only 2R regeneration, i.e., reamplification and reshaping, as opposed to 3R, which includes retiming as well. Thus, a combination of all-optical and electronic regeneration may be needed; the bulk of the regeneration can be performed all-optically, with electronic regenerators used intermittently to clean up the timing jitter.

The first commercial all-optical regenerators are likely to be compatible only with relatively simple modulation formats such as OOK. However, in order to

meet the capacity requirements of future networks, more advanced modulation formats are needed. All-optical regeneration that is compatible with advanced modulation formats is an area of current research; e.g., [SMCS05, Mats05]. Some of the solutions that have been proposed are fiber-based, which will result in a bulky design. It is desirable that solutions that are more integratable be developed.

All-optical regeneration, at least initially, is likely to operate on a per-wavelength basis, similar to electronic regenerators; i.e., Fig. 4.13 holds for all-optical regenerators as well. It is expected that these all-optical regenerators will provide complete flexibility with respect to wavelength conversion, allowing any input/output wavelength combination. Another potential technology being researched is multi-wavelength regenerators, where a whole band of wavelengths is regenerated at once [CXBT02, HXGB05]. This potentially improves the economics of regeneration even further, although it is not clear whether such a technology can support wavelength conversion.

Overall, all-optical regeneration is a relatively new technology that is still an area of active research on many fronts.