

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 Subrate Demands	224
7.10.1 Wavelength-Level Protection	224
7.10.2 Subrate-Level Protection	226
7.10.3 Wavelength-Level vs. Subrate-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.

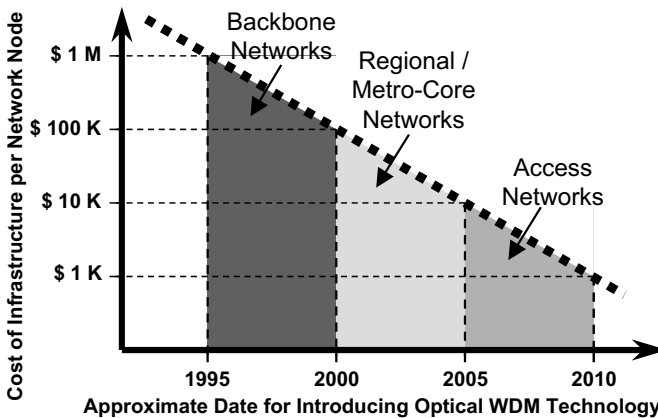


Fig. 1.2 As the cost of WDM infrastructure decreases over time, it is introduced closer to the network edge. (Adapted from [Sale98b].)

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 *wavelength* is frequently used to refer to the particular wavelength on which an optical channel is carried; *wavelength 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].