

SECOND EDITION

Rajiv Ramaswami
Kumar N. Sivarajan

Optical Networks

A PRACTICAL PERSPECTIVE

Nokia of America Corporation

IPR2023-00960, Ex. 1013

Ex. 1013.001

AT&T Ex. 1013

The readable style, balanced coverage of transmission and networking issues, and inclusion of real-world issues has been carried forward from the first edition into this second edition. This new version provides timely updates on emerging technologies that will shape the deployment of optical networks for years to come.

—Thomas S. Afferton, District Manager, Advanced Transport Technology and Architecture Planning, AT&T Network Services

The second edition of Optical Networks impressively covers the new technologies that have taken flight since the release of the first edition. The result is a fine text that provides comprehensive and detailed information on optical technologies and networks for students, professionals, and researchers.

—Joseph Bannister, Director, Computer Networks Division, Information Sciences Institute, USC

. . . In the second edition, Ramaswami and Sivarajan have continued to provide a concise approach to this broad and multidisciplinary field and to strike an important balance between device and transport physics and transmission engineering, network protocols, and architectures. Readers will value the important updates on technology advances and their impact on future network architectures.

—Daniel J. Blumenthal, Professor of ECE, University of California, Santa Barbara and Cofounder, Calient Networks

The horizons of optical networks are much more than high speed physical layer transport. An intelligent optical network design must include higher network layer considerations. This is the only book currently on the market that addresses optical networks from the physical layer to the network layer and should be valuable for those who try to understand the intricacies of what optical networks can be.

—Vincent Chan, Professor, MIT Department of Electrical Engineering and Computer Science

This book is not only essential reading for anyone in the optical networks industry, it is important. It provides the necessary foundation of learning for anyone hoping to contribute to this technology's rapid evolution.

—Scott Clavenna, President, PointEast Research

The authors' grasp of what is truly workable and worthwhile in optical networks is fundamental, and they have effectively packaged this knowledge in an easy-to-comprehend text that will be valued to both veterans and those new to optical networking.

—Scott Grout, President and CEO, Chorum Technologies

This is a comprehensive and authoritative work on optical networks, ranging in scope from components and systems to overall design principles. I find the book well organized and easy to use, and I particularly like the treatment of network design and operation. An essential book for anyone seriously interested in optical networks.

—Goff Hill, Chief Network Architect, Altamar Networks, UK

The second edition of this widely read text will be a valuable resource for students and engineers alike. It provides extensive coverage not only of topics related to current practice, but also techniques being researched, which may be deployed in the future.

—Dr. David Hunter, Senior Researcher, Marconi Labs, Cambridge, UK

Nokia of America Corporation

IPR2023-00960, Ex. 1013

AT&T Ex. 1013

Ex. 1013.002

The authors have done an excellent job of updating the content with all the new developments in optical networking in this second edition. Optical networking is a rapidly changing technology that can support the enormous Internet traffic growth demands required by bandwidth intensive applications of the 21st century.

—Rao Lingampalli, Senior Technical Marketing Manager, Calient Networks

I really enjoy the bottoms-up approach taken by the authors to address fundamentals of optical components as the enablers, optical transmission system design and engineering as the building blocks, and network architecture and its management features that deliver applications to the network operators and services providers at the top of the food chain.

—Shoa-Kai Liu, Director of Advanced Technology, Worldcom

This book not only provides the fundamentals and details of photonics, but the pragmatic perspective presented enables the service provider, the equipment manufacturer, and the academician to view light from a real-life standpoint.

—Mathew Oommen, Vice President, Network Architecture, Williams Communications Group

This book functions as both an introduction to optical networking and as a text to reference again and again. Great for system designers as well as those marketing and selling those systems. Optical Networks provides theory and applications. While no text can be truly state-of-the-art in the fast moving area of optical networking, this one comes as close as possible.

—Alan Repech, System Architect, Cisco Systems Optical Transport

This book provides the most comprehensive coverage of both the theory and practice of optical networking. Its up-to-date coverage makes it an invaluable reference for both practitioners and researchers.

—Suresh Subramaniam, Assistant Professor, Department of Electrical and Computer Engineering, George Washington University

This book definitely provides the most comprehensive coverage of the technologies, architectures, and applications essential for the study of the optical networking field. This second edition is very well written and organized, and the new topics covered appropriately reflect many of the latest technology and standards developments within our industry.

—Mark R. Wilson, Ph.D., Adjunct Professor of Electrical Engineering, University of Pennsylvania; Senior Manager, Optical Network Architectures, Agere Systems

This book provides an excellent overview of the complex field of optical networking. I especially like how it ties the optical hardware functionality into the overall networking picture. Everybody who wants to be a player in the optical networking space should have this book within easy reach.

—Martin Zirngibl, Director, Photonics Network Research, Lucent Technologies, Bell Laboratories

Optical Networks

A Practical Perspective

Second Edition

The Morgan Kaufmann Series in Networking
Series Editor, David Clark, MIT

Optical Networks: A Practical Perspective, 2e
Rajiv Ramaswami and Kumar N. Sivarajan

Telecommunications Law in the Internet Age
Sharon K. Black

Internet QoS: Architectures and Mechanisms
Zheng Wang

TCP/IP Sockets in Java: Practical Guide for Programmers
Michael J. Donahoo and Kenneth L. Calvert

TCP/IP Sockets in C: Practical Guide for Programmers
Kenneth L. Calvert and Michael J. Donahoo

Multicast Communication: Protocols, Programming, and Applications
Ralph Wittmann and Martina Zitterbart

MPLS: Technology and Applications
Bruce Davie and Yakov Rekhter

High-Performance Communication Networks, 2e
Jean Walrand and Pravin Varaiya

Computer Networks: A Systems Approach, 2e
Larry L. Peterson and Bruce S. Davie

Internetworking Multimedia
Jon Crowcroft, Mark Handley, and Ian Wakeman

Understanding Networked Applications: A First Course
David G. Messerschmitt

Integrated Management of Networked Systems: Concepts, Architectures, and their Operational Application
Heinz-Gerd Hegering, Sebastian Abeck, and Bernhard Neumair

Virtual Private Networks: Making the Right Connection
Dennis Fowler

Networked Applications: A Guide to the New Computing Infrastructure
David G. Messerschmitt

Modern Cable Television Technology: Video, Voice, and Data Communications
Walter Ciciora, James Farmer, and David Large

Switching in IP Networks: IP Switching, Tag Switching, and Related Technologies
Bruce S. Davie, Paul Doolan, and Yakov Rekhter

Wide Area Network Design: Concepts and Tools for Optimization
Robert S. Cahn

Practical Computer Network Analysis and Design
James D. McCabe

Frame Relay Applications: Business and Technology Case Studies
James P. Cavanagh

For further information on these books and for a list of forthcoming titles,
please visit our Web site at www.mkp.com.

Optical Networks

A Practical Perspective

Second Edition

Rajiv Ramaswami
Kumar N. Sivarajan



MORGAN KAUFMANN PUBLISHERS

An Imprint of Elsevier

SAN FRANCISCO SAN DIEGO NEW YORK BOSTON
LONDON SYDNEY TOKYO

Nokia of America Corporation
IPR2023-00960, Ex. 1013
AT&T Ex. 1013
Ex. 1013.006

AT&T Ex. 1013

To our parents.

This Page Intentionally Left Blank



Foreword

by Paul E. Green, Jr.

Director, Optical Network Technology (retired)

Tellabs, Inc.

If anyone needed evidence of the rapid pace of innovation in WDM optical telecommunications, a year-by-year perusal of the innovations reported in the journals and in the activities of start-up companies in the optical space would provide all needed persuasion. While the field is old—some fifteen years to be exact—it is also young in the sense that almost nothing in the technology has reached steady-state maturity. New optical bands are being opened up, unprecedented bitrates are being carried per wavelength over all-optical distances that were not thought possible not long ago, new forms of fiber are being introduced, and point-point links are now being organized by their owners into rings and other real network structures, using new optical switching and routing technologies.

Thus, it is a most welcome development to see that Ramaswami and Sivarajan have reorganized and updated their outstanding 1998 traversal of this fascinating and important field and have done so in a way that accurately portrays what is going on. The first edition reflected the authors' backgrounds, straddling the academic and industrial research community on one hand and the school of commercial hard knocks on the other, and it is still true that this new book provides a view of optical communication that is lacking in other frequently updated volumes in this field, whose authors lack the authority possessed by Ramaswami and Sivarajan.

At this juncture, there is a pause in the headlong rush to provide ever more bandwidth to everybody, and the reasons are not just economic. The bottleneck between the gigabit environment on everyone's desktop and the terabit environment of the telco and cable fiber infrastructure is a consequence of the failure of the available "last mile" communication technologies to seriously meet the challenge of providing

enough bandwidth. Between now and the next edition of this book, one may hope that fiber last mile solutions will finally begin to proliferate, building on the technical ideas captured here.

There can be no doubt that fiber optic communication can only play an ever widening, if invisible, role in our lives, and once again, I cannot recommend the book highly enough to anyone interested in the technical underpinnings.

Foreword to the First Edition

by Paul E. Green, Jr.

Director, Optical Network Technology (retired)

Tellabs, Inc.

Not too many years ago, whenever one wanted to send messages effectively, there were really only two choices—send them by wire or send them by radio. This situation lasted for decades until the mid-1960s, when the fiber optics revolution began, quietly at first, and then with increasing force as people began to appreciate that sending pulses of light through tiny strands of glass wasn't so crazy after all. This revolution is now in full cry, with 4000 strand miles of fiber being installed per day, just in the United States alone. Fiber has been displacing wire in many applications, and gradually it is emerging as one of the two dominant Cinderella transmission technologies of today, wireless being the other. One of these (wireless) goes anywhere but doesn't do much when it gets there, whereas the other (fiber) will never go everywhere but does a great deal indeed wherever it reaches. From the earliest days of fiber communication, people realized that this simple glass medium has incredible amounts of untapped bandwidth capacity waiting to be mined, should the day come when we would actually need it, and should we be able to figure out how to tap it. That day has now come. The demand is here and so are the solutions.

This book describes a revolution within a revolution, the opening up of the capacity of the now-familiar optical fiber to carry more messages, handle a wider variety of transmission types, and provide improved reliabilities and ease of use. In many places where fiber has been installed simply as a better form of copper, even the gigabit capacities that result have not proved adequate to keep up with the demand. The inborn human voracity for more and more bandwidth, plus the growing realization that there are other flexibilities to be had by imaginative use of the fiber, have led people to explore all-optical networks, the subject of this book.

xi

Such networks are those in which either wavelength division or time division is used in new ways to form entire network structures where the messages travel in purely optical form all the way from one user location to another.

When I attempted the same kind of book in 1993, nobody was quite sure whether optical networking would be a roaring success or disappear into the annals of “whatever happened to . . .” stories of technology that had once sounded great on paper, but that had somehow never panned out in the real world. My book (*Fiber Optic Networks*, Prentice Hall) spent most of its pages talking about technology building blocks and lamenting their limitations since there was little to say about real networks, the architectural considerations underlying them, and what good they had ever done anybody.

In the last four years, optical networking has indeed really happened, essentially all of it based on wavelength division multiplexing, and with this book Ramaswami and Sivarajan, two of the principal architects of this success, have redressed the insufficiencies of earlier books such as mine. Today, hundreds of millions of dollars of wavelength division networking systems are being sold annually, major new businesses have been created that produce nothing but optical networks, and bandwidth bottlenecks are being relieved and proliferating protocol zoos tamed by this remarkably transparent new way of doing networking; what’s more, there is a rich architectural understanding of where to go next. Network experts, fresh from the novelties of such excitements as the Web, now have still another wonderful toy shop to play in. The whole optical networking idea is endlessly fascinating in itself—based on a medium with thousands of gigabits of capacity yet so small as to be almost invisible, transmitters no larger than a grain of salt, amplifiers that amplify vast chunks of bandwidth purely as light, transmission designs that bypass 50 years of hard-won but complex coding, modulation and equalization insights, network architectures that subsume many functions usually done more clumsily in the lower layers of classical layered architectures—these are all fresh and interesting topics that await the reader of this book.

To understand this new networking revolution within a revolution, it is necessary to be led with a sure hand through territory that to many will be unfamiliar. The present authors, with their rare mixture of physics and network architecture expertise, are eminently qualified to serve as guides. After spending some time with this book, you will be more thoroughly conversant with all the important issues that today affect how optical networks are made, what their limitations and potentialities are, and how they fit in with more classical forms of communication networks based on electronic time division. Whether you are a computer network expert wondering how to use fiber to break the bandwidth bottlenecks that are limiting your system capabilities, a planner or implementer trying to future-proof your telephone network, a teacher planning a truly up-to-date communication engineering curriculum, a student

looking for a fun lucrative career, or a midcareer person in need of a retread, this volume will provide the help you need.

The authors have captured what is going on and what is going to be going on in this field in a completely up-to-date treatment unavailable elsewhere. I learned a lot from reading it and expect that you will too.

This Page Intentionally Left Blank



Contents

Foreword ix

Foreword to the First Edition xi

Preface xxvii

1 Introduction to Optical Networks 1

1.1 Telecommunications Network Architecture 3

1.2 Services, Circuit Switching, and Packet Switching 6

1.2.1 The Changing Services Landscape 9

1.3 Optical Networks 10

1.3.1 Multiplexing Techniques 12

1.3.2 Second-Generation Optical Networks 14

1.4 The Optical Layer 16

1.5 Transparency and All-Optical Networks 24

1.6 Optical Packet Switching 26

1.7 Transmission Basics 28

1.7.1 Wavelengths, Frequencies, and Channel Spacing 28

1.7.2 Wavelength Standards 30

1.7.3 Optical Power and Loss 31

1.8 Network Evolution 32

1.8.1 Early Days—Multimode Fiber 33

1.8.2 Single-Mode Fiber 35

1.8.3	Optical Amplifiers and WDM	37
1.8.4	Beyond Transmission Links to Networks	39
	Summary	40
	Further Reading	41
	References	42
I Technology		47
2	Propagation of Signals in Optical Fiber	49
2.1	Light Propagation in Optical Fiber	50
2.1.1	Geometrical Optics Approach	50
2.1.2	Wave Theory Approach	55
2.2	Loss and Bandwidth	65
2.2.1	Bending Loss	67
2.3	Chromatic Dispersion	68
2.3.1	Chirped Gaussian Pulses	69
2.3.2	Controlling the Dispersion Profile	74
2.4	Nonlinear Effects	76
2.4.1	Effective Length and Area	77
2.4.2	Stimulated Brillouin Scattering	79
2.4.3	Stimulated Raman Scattering	80
2.4.4	Propagation in a Nonlinear Medium	81
2.4.5	Self-Phase Modulation	83
2.4.6	SPM-Induced Chirp for Gaussian Pulses	87
2.4.7	Cross-Phase Modulation	89
2.4.8	Four-Wave Mixing	90
2.4.9	New Optical Fiber Types	93
2.5	Solitons	98
2.5.1	Dispersion-Managed Solitons	100
	Summary	101
	Further Reading	102
	Problems	103
	References	104
3	Components	107
3.1	Couplers	108
3.1.1	Principle of Operation	110
3.1.2	Conservation of Energy	111
3.2	Isolators and Circulators	112
3.2.1	Principle of Operation	113

3.3	Multiplexers and Filters	115
3.3.1	Gratings	118
3.3.2	Diffraction Pattern	122
3.3.3	Bragg Gratings	123
3.3.4	Fiber Gratings	126
3.3.5	Fabry-Perot Filters	130
3.3.6	Multilayer Dielectric Thin-Film Filters	133
3.3.7	Mach-Zehnder Interferometers	135
3.3.8	Arrayed Waveguide Grating	139
3.3.9	Acousto-Optic Tunable Filter	143
3.3.10	High Channel Count Multiplexer Architectures	148
3.4	Optical Amplifiers	151
3.4.1	Stimulated Emission	152
3.4.2	Spontaneous Emission	153
3.4.3	Erbium-Doped Fiber Amplifiers	154
3.4.4	Raman Amplifiers	159
3.4.5	Semiconductor Optical Amplifiers	161
3.4.6	Crosstalk in SOAs	165
3.5	Transmitters	165
3.5.1	Lasers	166
3.5.2	Light-Emitting Diodes	176
3.5.3	Tunable Lasers	178
3.5.4	Direct and External Modulation	186
3.5.5	Pump Sources for Raman Amplifiers	190
3.6	Detectors	192
3.6.1	Photodetectors	192
3.6.2	Front-End Amplifiers	197
3.7	Switches	199
3.7.1	Large Optical Switches	201
3.7.2	Optical Switch Technologies	207
3.7.3	Large Electronic Switches	215
3.8	Wavelength Converters	216
3.8.1	Optoelectronic Approach	217
3.8.2	Optical Gating	218
3.8.3	Interferometric Techniques	220
3.8.4	Wave Mixing	223
	Summary	224
	Further Reading	225
	Problems	226
	References	232

4	Modulation and Demodulation	239
4.1	Modulation	239
4.1.1	Signal Formats	240
4.2	Subcarrier Modulation and Multiplexing	242
4.2.1	Clipping and Intermodulation Products	243
4.2.2	Applications of SCM	245
4.3	Spectral Efficiency	245
4.3.1	Optical Duobinary Modulation	246
4.3.2	Optical Single Sideband Modulation	248
4.3.3	Multilevel Modulation	249
4.3.4	Capacity Limits of Optical Fiber	249
4.4	Demodulation	250
4.4.1	An Ideal Receiver	252
4.4.2	A Practical Direct Detection Receiver	253
4.4.3	Front-End Amplifier Noise	254
4.4.4	APD Noise	255
4.4.5	Optical Preamplifiers	255
4.4.6	Bit Error Rates	258
4.4.7	Coherent Detection	263
4.4.8	Timing Recovery	265
4.4.9	Equalization	266
4.5	Error Detection and Correction	267
4.5.1	Reed-Solomon Codes	270
4.5.2	Interleaving	271
	Summary	272
	Further Reading	273
	Problems	274
	References	279
5	Transmission System Engineering	283
5.1	System Model	283
5.2	Power Penalty	284
5.3	Transmitter	287
5.4	Receiver	288
5.5	Optical Amplifiers	289
5.5.1	Gain Saturation in EDFAs	290
5.5.2	Gain Equalization in EDFAs	292
5.5.3	Amplifier Cascades	293
5.5.4	Amplifier Spacing Penalty	294

5.5.5	Power Transients and Automatic Gain Control	296
5.5.6	Lasing Loops	298
5.6	Crosstalk	299
5.6.1	Intrachannel Crosstalk	299
5.6.2	Interchannel Crosstalk	301
5.6.3	Crosstalk in Networks	303
5.6.4	Bidirectional Systems	303
5.6.5	Crosstalk Reduction	305
5.6.6	Cascaded Filters	307
5.7	Dispersion	308
5.7.1	Chromatic Dispersion Limits: NRZ Modulation	309
5.7.2	Chromatic Dispersion Limits: RZ Modulation	311
5.7.3	Dispersion Compensation	314
5.7.4	Polarization-Mode Dispersion (PMD)	320
5.8	Fiber Nonlinearities	323
5.8.1	Effective Length in Amplified Systems	323
5.8.2	Stimulated Brillouin Scattering	325
5.8.3	Stimulated Raman Scattering	326
5.8.4	Four-Wave Mixing	329
5.8.5	Self-/Cross-Phase Modulation	333
5.8.6	Role of Chromatic Dispersion Management	335
5.9	Wavelength Stabilization	335
5.10	Design of Soliton Systems	336
5.11	Design of Dispersion-Managed Soliton Systems	338
5.12	Overall Design Considerations	341
5.12.1	Fiber Type	341
5.12.2	Transmit Power and Amplifier Spacing	342
5.12.3	Chromatic Dispersion Compensation	343
5.12.4	Modulation	343
5.12.5	Nonlinearities	344
5.12.6	Interchannel Spacing and Number of Wavelengths	344
5.12.7	All-Optical Networks	345
5.12.8	Wavelength Planning	346
5.12.9	Transparency	348
	Summary	348
	Further Reading	348
	Problems	349
	References	356

II Networks 361

6 Client Layers of the Optical Layer	363
6.1 SONET/SDH	364
6.1.1 Multiplexing	367
6.1.2 SONET/SDH Layers	370
6.1.3 SONET Frame Structure	371
6.1.4 SONET/SDH Physical Layer	375
6.1.5 Elements of a SONET/SDH Infrastructure	378
6.2 ATM	381
6.2.1 Functions of ATM	382
6.2.2 Adaptation Layers	385
6.2.3 Quality of Service	386
6.2.4 Flow Control	387
6.2.5 Signaling and Routing	387
6.3 IP	388
6.3.1 Routing and Forwarding	390
6.3.2 Quality of Service	392
6.3.3 Multiprotocol Label Switching (MPLS)	392
6.3.4 Whither ATM?	394
6.4 Storage-Area Networks	395
6.4.1 ESCON	396
6.4.2 Fibre Channel	397
6.4.3 HIPPI	397
6.5 Gigabit and 10-Gigabit Ethernet	398
Summary	398
Further Reading	399
Problems	399
References	401
7 WDM Network Elements	403
7.1 Optical Line Terminals	406
7.2 Optical Line Amplifiers	408
7.3 Optical Add/Drop Multiplexers	408
7.3.1 OADM Architectures	411
7.3.2 Reconfigurable OADMs	417
7.4 Optical Crossconnects	419
7.4.1 All-Optical OXC Configurations	425
Summary	428
Further Reading	430

Problems	431
References	433
8 WDM Network Design	437
8.1 Cost Trade-Offs: A Detailed Ring Network Example	441
8.2 LTD and RWA Problems	448
8.2.1 Lightpath Topology Design	449
8.2.2 Routing and Wavelength Assignment	454
8.2.3 Wavelength Conversion	457
8.2.4 Relationship to Graph Coloring	461
8.3 Dimensioning Wavelength-Routing Networks	462
8.4 Statistical Dimensioning Models	464
8.4.1 First-Passage Model	466
8.4.2 Blocking Model	467
8.5 Maximum Load Dimensioning Models	475
8.5.1 Offline Lightpath Requests	475
8.5.2 Online RWA in Rings	481
Summary	482
Further Reading	484
Problems	484
References	488
9 Control and Management	495
9.1 Network Management Functions	495
9.1.1 Management Framework	497
9.1.2 Information Model	499
9.1.3 Management Protocols	500
9.2 Optical Layer Services and Interfacing	502
9.3 Layers within the Optical Layer	504
9.4 Multivendor Interoperability	505
9.5 Performance and Fault Management	507
9.5.1 The Impact of Transparency	508
9.5.2 BER Measurement	509
9.5.3 Optical Trace	509
9.5.4 Alarm Management	510
9.5.5 Data Communication Network (DCN) and Signaling	512
9.5.6 Policing	513
9.5.7 Optical Layer Overhead	514
9.6 Configuration Management	519
9.6.1 Equipment Management	519

9.6.2	Connection Management	520
9.6.3	Adaptation Management	524
9.7	Optical Safety	526
9.7.1	Open Fiber Control Protocol	528
	Summary	530
	Further Reading	531
	Problems	532
	References	534
10	Network Survivability	537
10.1	Basic Concepts	539
10.2	Protection in SONET/SDH	542
10.2.1	Point-to-Point Links	544
10.2.2	Self-Healing Rings	546
10.2.3	Unidirectional Path-Switched Rings	549
10.2.4	Bidirectional Line-Switched Rings	550
10.2.5	Ring Interconnection and Dual Homing	555
10.3	Protection in IP Networks	558
10.4	Why Optical Layer Protection	560
10.4.1	Service Classes Based on Protection	566
10.5	Optical Layer Protection Schemes	567
10.5.1	1 + 1 OMS Protection	570
10.5.2	1:1 OMS Protection	571
10.5.3	OMS-DPRing	571
10.5.4	OMS-SPRing	571
10.5.5	1:N Transponder Protection	573
10.5.6	1 + 1 OCh Dedicated Protection	574
10.5.7	OCh-SPRing	575
10.5.8	OCh-Mesh Protection	576
10.5.9	Choice of Protection Technique	582
10.6	Interworking between Layers	582
	Summary	583
	Further Reading	584
	Problems	584
	References	587
11	Access Networks	591
11.1	Network Architecture Overview	593
11.2	Enhanced HFC	598
11.3	Fiber to the Curb (FTTC)	599
11.3.1	PON Evolution	609

Summary	610
Further Reading	611
Problems	611
References	612
12 Photonic Packet Switching 615	
12.1 Optical Time Division Multiplexing	619
12.1.1 Bit Interleaving	620
12.1.2 Packet Interleaving	623
12.1.3 Optical AND Gates	625
12.2 Synchronization	631
12.2.1 Tunable Delays	632
12.2.2 Optical Phase Lock Loop	633
12.3 Header Processing	634
12.4 Buffering	635
12.4.1 Output Buffering	637
12.4.2 Input Buffering	639
12.4.3 Recirculation Buffering	639
12.4.4 Using Wavelengths for Contention Resolution	641
12.4.5 Deflection Routing	645
12.5 Burst Switching	649
12.6 Testbeds	650
12.6.1 KEOPS	650
12.6.2 NTT's Optical ATM Switches	652
12.6.3 BT Labs Testbeds	654
12.6.4 Princeton University Testbed	654
12.6.5 AON	655
12.6.6 CORD	656
Summary	657
Further Reading	658
Problems	659
References	660
13 Deployment Considerations 667	
13.1 The Evolving Telecommunications Network	667
13.1.1 The SONET/SDH Core Network	669
13.1.2 Architectural Choices for Next-Generation Transport Networks	673
13.2 Designing the Transmission Layer	681
13.2.1 Using SDM	682
13.2.2 Using TDM	683
13.2.3 Using WDM	684

13.2.4	Unidirectional versus Bidirectional WDM Systems	686
13.2.5	Long-Haul Networks	688
13.2.6	Long-Haul Network Case Study	689
13.2.7	Long-Haul Undersea Networks	697
13.2.8	Metro Networks	698
13.2.9	Metro Ring Case Study	700
13.2.10	From Opaque Links to Agile All-Optical Networks	704
	Summary	705
	Further Reading	706
	Problems	706
	References	710
A	Acronyms	711
B	Symbols and Parameters	717
C	Standards	721
C.1	International Telecommunications Union (ITU-T)	721
C.1.1	Fiber	721
C.1.2	SDH (Synchronous Digital Hierarchy)	721
C.1.3	Optical Networking	722
C.1.4	Management	722
C.2	Telcordia	723
C.2.1	Physical and Environmental	723
C.2.2	SONET	723
C.2.3	Optical Networking	724
C.3	American National Standards Institute (ANSI)	724
C.3.1	SONET	724
C.3.2	ESCON and Fibre Channel	724
D	Wave Equations	727
E	Pulse Propagation in Optical Fiber	731
E.1	Propagation of Chirped Gaussian Pulses	734
E.2	Nonlinear Effects on Pulse Propagation	735
E.3	Soliton Pulse Propagation	738
	Further Reading	739
	References	739
F	Nonlinear Polarization	741

G Multilayer Thin-Film Filters	743
G.1 Wave Propagation at Dielectric Interfaces	743
G.2 Filter Design	747
References	750
H Random Variables and Processes	751
H.1 Random Variables	751
H.1.1 Gaussian Distribution	752
H.1.2 Maxwell Distribution	753
H.1.3 Poisson Distribution	753
H.2 Random Processes	754
H.2.1 Poisson Random Process	755
H.2.2 Gaussian Random Process	756
Further Reading	756
References	756
I Receiver Noise Statistics	757
I.1 Shot Noise	759
I.2 Amplifier Noise	760
References	762
Bibliography	763
Index	797




Figure 1.15 Evolution of optical fiber transmission systems. (a) An early system using LEDs over multimode fiber. (b) A system using MLM lasers over single-mode fiber in the $1.3 \mu\text{m}$ band to overcome intermodal dispersion in multimode fiber. (c) A later system using the $1.55 \mu\text{m}$ band for lower loss, and using SLM lasers to overcome chromatic dispersion limits. (d) A current-generation WDM system using multiple wavelengths at $1.55 \mu\text{m}$ and optical amplifiers instead of regenerators. The P - λ curves to the left of the transmitters indicate the power spectrum of the signal transmitted.

1.8.1 Early Days—Multimode Fiber

Early experiments in the mid-1960s demonstrated that information encoded in light signals could be transmitted over a glass fiber *waveguide*. A waveguide provides a medium that can *guide* the light signal, enabling it to stay focused for a reasonable distance without being scattered. This allows the signal to be received at the other

end with sufficient strength so that the information can be decoded. These early experiments proved that optical transmission over fiber was feasible.

An optical fiber is a very thin cylindrical glass waveguide consisting of two parts: an inner *core* material and an outer *cladding* material. The core and cladding are designed so as to keep the light signals *guided* inside the fiber, allowing the light signal to be transmitted for reasonably long distances before the signal degrades in quality.

It was not until the invention of low-loss optical fiber in the early 1970s that optical fiber transmission systems really took off. This silica-based optical fiber has three low-loss windows in the 0.8, 1.3, and 1.55 μm infrared wavelength bands. The lowest loss is around 0.25 dB/km in the 1.55 μm band, and about 0.5 dB/km in the 1.3 μm band. These fibers enabled transmission of light signals over distances of several tens of kilometers before they needed to be *regenerated*. A regenerator converts the light signal into an electrical signal and retransmits a fresh copy of the data as a new light signal.

The early fibers were the so-called multimode fibers. Multimode fibers have core diameters of about 50 to 85 μm . This diameter is large compared to the operating wavelength of the light signal. A basic understanding of light propagation in these fibers can be obtained using the so-called geometrical optics model, illustrated in Figure 1.16. In this model, a light ray bounces back and forth in the core, being reflected at the core-cladding interface. The signal consists of multiple light rays, each of which potentially takes a different path through the fiber. Each of these different paths corresponds to a *propagation mode*. The length of the different paths is different, as seen in the figure. Each mode therefore travels with a slightly different speed compared to the other modes.

The other key devices needed for optical fiber transmission are light sources and receivers. Compact semiconductor lasers and light-emitting diodes (LEDs) provided practical light sources. These lasers and LEDs were simply turned on and off rapidly to transmit digital (binary) data. Semiconductor photodetectors enabled the conversion of the light signal back into the electrical domain.

The early telecommunication systems (late 1970s through the early 1980s) used multimode fibers along with LEDs or laser transmitters in the 0.8 and 1.3 μm wavelength bands. LEDs were relatively low-power devices that emitted light over a fairly wide spectrum of several nanometers to tens of nanometers. A laser provided higher output power than an LED and therefore allowed transmission over greater distances before regeneration. The early lasers were *multilongitudinal mode* (MLM) Fabry-Perot lasers. These MLM lasers emit light over a fairly wide spectrum of several nanometers to tens of nanometers. The actual spectrum consists of multiple spectral lines, which can be thought of as different longitudinal modes, hence the




Figure 1.16 Geometrical optics model to illustrate the propagation of light in an optical fiber. (a) Cross section of an optical fiber. The fiber has an inner core and an outer cladding, with the core having a slightly higher refractive index than the cladding. (b) Longitudinal view. Light rays within the core hitting the core-cladding boundary are reflected back into the core by total internal reflection.

term MLM. Note that these longitudinal laser modes are different from the propagation modes inside the optical fiber! While both LEDs and MLM lasers emit light over a broad spectrum, the spectrum of an LED is continuous, whereas the spectrum of an MLM laser consists of many periodic lines.

These early systems had to have regenerators every few kilometers to regenerate the signal. Regenerators were expensive devices and continue to be expensive today, so it is highly desirable to maximize the distance between regenerators. In this case, the distance limitation was primarily due to a phenomenon known as *intermodal dispersion*. As we saw earlier, in a multimode fiber, the energy in a pulse travels in different modes, each with a different speed. At the end of the fiber, the different modes arrive at slightly different times, resulting in a smearing of the pulse. This smearing in general is called *dispersion*, and this specific form is called intermodal dispersion. Typically, these early systems operated at bit rates ranging from 32 to 140 Mb/s with regenerators every 10 km. Such systems are still used for low-cost computer interconnection at a few hundred megabits per second over a few kilometers.

1.8.2 Single-Mode Fiber

The next generation of systems deployed starting around 1984 used *single-mode* fiber as a means of eliminating intermodal dispersion, along with MLM Fabry-Perot lasers in the $1.3\ \mu\text{m}$ wavelength band. Single-mode fiber has a relatively small core diameter of about 8 to $10\ \mu\text{m}$, which is a small multiple of the operating wavelength

range of the light signal. This forces all the energy in a light signal to travel in the form of a single mode. Using single-mode fiber effectively eliminated intermodal dispersion and enabled a dramatic increase in the bit rates and distances possible between regenerators. These systems typically had regenerator spacings of about 40 km and operated at bit rates of a few hundred megabits per second. At this point, the distance between regenerators was limited primarily by the fiber loss.

The next step in this evolution in the late 1980s was to deploy systems in the 1.55 μm wavelength window to take advantage of the lower loss in this window, relative to the 1.3 μm window. This enabled longer spans between regenerators. At this point, another impairment, namely, *chromatic dispersion*, started becoming a limiting factor as far as increasing the bit rates was concerned. Chromatic dispersion is another form of dispersion in optical fiber (we looked at intermodal dispersion earlier). As we saw in Section 1.7, the energy in a light signal or pulse has a finite bandwidth. Even in a single-mode fiber, the different frequency components of a pulse propagate with different speeds. This is due to the fundamental physical properties of the glass. This effect again causes a smearing of the pulse at the output, just as with intermodal dispersion. The wider the spectrum of the pulse, the more the smearing due to chromatic dispersion. The chromatic dispersion in an optical fiber depends on the wavelength of the signal. It turns out that without any special effort, the standard silica-based optical fiber has essentially no chromatic dispersion in the 1.3 μm band, but has significant dispersion in the 1.55 μm band. Thus chromatic dispersion was not an issue in the earlier systems at 1.3 μm .

The high chromatic dispersion at 1.55 μm motivated the development of *dispersion-shifted fiber*. Dispersion-shifted fiber is carefully designed to have zero dispersion in the 1.55 μm wavelength window so that we need not worry about chromatic dispersion in this window. However, by this time there was already a large installed base of standard single-mode fiber deployed for which this solution could not be applied. Some carriers, particularly NTT in Japan and MCI (now part of Worldcom) in the United States, did deploy dispersion-shifted fiber.

At this time, researchers started looking for ways to overcome chromatic dispersion while still continuing to make use of standard fiber. The main technique that came into play was to reduce the width of the spectrum of the transmitted pulse. As we saw earlier, the wider the spectrum of the transmitted pulse, the greater the smearing due to chromatic dispersion. The bandwidth of the transmitted pulse is at least equal to its modulation bandwidth. On top of this, however, the bandwidth may be determined entirely by the width of the spectrum of the transmitter used. The MLM Fabry-Perot lasers, as we said earlier, emitted over a fairly wide spectrum of several nanometers (or, equivalently, hundreds of gigahertz), which is much larger than the modulation bandwidth of the signal itself. If we reduce the spectrum of the transmitted pulse to something close to its modulation bandwidth, the penalty due

to chromatic dispersion is significantly reduced. This motivated the development of a laser source with a narrow spectral width—the *distributed-feedback* (DFB) laser. A DFB laser is an example of a *single-longitudinal mode* (SLM) laser. An SLM laser emits a narrow single-wavelength signal in a single spectral line, in contrast to MLM lasers whose spectrum consists of many spectral lines. This technological breakthrough spurred further increases in the bit rate to more than 1 Gb/s.

1.8.3 Optical Amplifiers and WDM

The next major milestone in the evolution of optical fiber transmission systems was the development of *erbium-doped fiber amplifiers* (EDFAs) in the late 1980s and early 1990s. The EDFA basically consists of a length of optical fiber, typically a few meters to tens of meters, doped with the rare earth element erbium. The erbium atoms in the fiber are *pumped* from their ground state to an excited state at a higher energy level using a pump source. An incoming signal photon triggers these atoms to come down to their ground state. In the process, each atom emits a photon. Thus incoming signal photons trigger the emission of additional photons, resulting in optical amplification. Due to a unique coincidence of nature, the difference in energy levels of the atomic states of erbium line up with the 1.5 μm low-loss window in the optical fiber. The pumping itself is done using a pump laser at a lower wavelength than the signal because photons with a lower wavelength have higher energies and energy can be transferred only from a photon of higher energy to that with a lower energy. The EDFA concept was invented in the 1960s but had to wait for the availability of reliable high-power semiconductor pump lasers in the late 1980s and early 1990s before becoming commercially viable.

EDFAs spurred the deployment of a completely new generation of systems. A major advantage of EDFAs is that they are capable of amplifying signals at many wavelengths simultaneously. This provided another way of increasing the system capacity: rather than increasing the bit rate, keep the bit rate the same and use more than one wavelength; that is, use wavelength division multiplexing. EDFAs were perhaps the single biggest catalyst aiding the deployment of WDM systems. The use of WDM and EDFAs dramatically brought down the cost of long-haul transmission systems and increased their capacity. At each regenerator location, a single optical amplifier could replace an entire array of expensive regenerators, one per fiber. This proved to be so compelling that almost every long-haul carrier has widely deployed amplified WDM systems today. Moreover WDM provided the ability to turn on capacity quickly, as opposed to the months to years it could take to deploy new fiber. WDM systems with EDFAs were deployed starting in the mid-1990s and are today achieving capacities over 1 Tb/s over a single fiber. At the same time, transmission bit rates on a single channel have risen to 10 Gb/s. Among the earliest

5 chapter

Transmission System Engineering

OUR GOAL IN THIS CHAPTER is to understand how to design the physical layer of an optical network. To this end, we will understand the various impairments that we must deal with, how to allocate margins for each of these impairments, how to reduce the effect of these impairments, and finally all the trade-offs that are involved between the different design parameters.

5.1 System Model

Figure 5.1 shows a block diagram of the various components of a unidirectional WDM link. The transmitter consists of a set of DFB lasers, with or without external modulators, one for each wavelength. The signals at the different wavelengths are combined into a single fiber by means of an optical multiplexer. An optical power amplifier may be used to increase the transmission power. After some distance along the fiber, the signal is amplified by an optical in-line amplifier. Depending on the distance, bit rate, and type of fiber used, the signal may also be passed through a dispersion-compensating module, usually at each amplifier stage. At the receiving end, the signal may be amplified by an optical preamplifier before it is passed through a demultiplexer. Each wavelength is then received by a separate photodetector.

Throughout this chapter, we will be focusing on digital systems, although it is possible to transmit analog signals over fiber as well. The physical layer of the system must ensure that bits are transmitted from the source to their destination reliably. The measures of quality are the bit error rate (BER) and the additional power budget




Figure 5.1 Components of a WDM link.

margin provided in the system. Usually the required bit error rates are of the order of 10^{-9} to 10^{-15} , typically 10^{-12} . The BER depends on the amount of noise as well as other impairments that are present in the system. Unless otherwise stated, we will assume that non-return-to-zero (NRZ) modulation is used. In some specific cases, such as chromatic dispersion, we consider both NRZ and return-to-zero (RZ) modulation.

The physical layer is also responsible for the link initialization and link take-down procedures, which are necessary to prevent exposure to potentially harmful laser radiation. This aspect is dealt with in Chapter 9.

We will look at the different components that are part of a system, including the transmitters, receivers, optical amplifiers, wavelength multiplexers, demultiplexers and switches, and the fiber itself, and we will discuss various forms of system impairments that arise from each of these components. Table B.1 in Appendix B summarizes the large number of parameters that are used in this chapter.

5.2 Power Penalty

The physical layer design must take into account the effect of a number of system impairments as previously discussed. Usually each impairment results in a *power penalty* to the system. In the presence of an impairment, a higher signal power will be required at the receiver in order to maintain a desired bit error rate. One way to define the power penalty is as the increase in signal power required (in dB) to maintain the same bit error rate in the presence of impairments. Another way to define the power penalty is as the reduction in signal-to-noise ratio as quantified by the value of γ (the argument to the $Q(\cdot)$ function as defined in Section 4.4.6) due

to a specific impairment. We will be using the latter definition since it is easier to calculate and consistent with popular usage.

Let P_1 denote the optical power received during a 1 bit, and P_0 the power received during a 0 bit without any system impairments. The corresponding electrical currents are given by $\mathcal{R}P_1$ and $\mathcal{R}P_0$, respectively, where $\mathcal{R}$ is the responsivity of the photodetector.

Let σ_1 and σ_0 denote the noise standard deviations during a 1 bit and a 0 bit, respectively. Assume that the noise is Gaussian. The bit error rate, assuming equally likely 1s and 0s, is obtained from (4.14) as

$$\text{BER} = Q\left(\frac{\mathcal{R}(P_1 - P_0)}{\sigma_1 + \sigma_0}\right). \quad (5.1)$$

This expression assumes that the receiver's decision threshold is set to the optimal value indicated by (4.12).

In the presence of impairments, let P'_1 , P'_0 , σ'_1 , σ'_0 denote the received powers and noise standard deviations, respectively. Assuming an optimized threshold setting, the power penalty is given by

$$\text{PP} = -10 \log \left(\frac{\frac{\mathcal{R}(P'_1 - P'_0)}{\sigma'_1 + \sigma'_0}}{\frac{\mathcal{R}(P_1 - P_0)}{\sigma_1 + \sigma_0}} \right). \quad (5.2)$$

Calculating the power penalty in general for the simple AC-coupled receiver discussed in Section 4.4.6 is somewhat more complicated, but we will see that it is the same as the penalty for the optimized receiver for two important cases of interest.

The first case of interest is when the dominant noise component is receiver thermal noise, for which $\sigma_0 = \sigma_1 = \sigma_{\text{th}}$. This is usually the case in unamplified direct detection *pin* receivers. In this case, or in any situation where the noise is independent of the signal power, the power penalty is given by

$$\text{PP}_{\text{sig-indep}} = -10 \log \left(\frac{P'_1 - P'_0}{P_1 - P_0} \right) \quad (5.3)$$

and the best threshold setting corresponds to the setting of a simple AC-coupled receiver.

The other case of interest is amplified systems, or systems with APD receivers. In amplified systems, the dominant noise component is usually the amplifier signal-spontaneous beat noise (see Section 4.4.5). In APD receivers, the dominant noise component is the shot noise, which is enhanced because of the APD gain (see

Section 3.6.1). In amplified systems, and in systems with APD receivers, we can assume that $\sigma_1 \propto \sqrt{P_1}$; that is, the noise variance depends on the signal power. Assume also that $P_0 \ll P_1$. In this case, we can assume that $\sigma_1 \gg \sigma_0$. Here an optimized receiver would set its threshold close to the 0 level, whereas the simple receiver would still set its threshold at the average received power and would have a somewhat higher bit error rate. However, the power penalties turn out to be the same in both cases. This penalty is given by

$$PP_{\text{sig-dep}} = -5 \log \left(\frac{P'_1}{P_1} \right). \quad (5.4)$$

Finally, it must be kept in mind that polarization plays an important role in many system impairments where signals interfere with each other. The worst case is usually when the interfering signals have the same state of polarization. However, the state of polarization of each signal varies slowly with time in a random manner, and thus we can expect the power penalties to vary with time as well. The system must be designed, however, to accommodate the worst case, usually identical polarizations.

System design requires careful budgeting of the power penalties for the different impairments. Here we sketch out one way of doing such a design for a transmission system with optical amplifiers. First we determine the ideal value of the parameter γ (see Section 4.4.6) that is needed. For a bit error rate of 10^{-12} typically assumed in high-speed transmission systems, we need $\gamma = 7$, or $20 \log \gamma = 17$ dB. This would be the case if there were no transmission impairments leading to power penalties. In practice, the various impairments result in power penalties that must be added onto this ideal value of γ , as shown in Table 5.1, to obtain the required value of γ that the system must be designed to yield. For instance, in the table, we allocate a 1 dB power penalty for an imperfect transmitter and a 2 dB power penalty for chromatic dispersion. (We will study these and several other impairments in the rest of this chapter.) The required value of γ after adding all these allocations is 31 dB. This is the value that we must obtain if we assume an ideal system to start with and compute γ based on only optical amplifier noise accumulation. The power penalty due to each impairment is then calculated one at a time assuming that the rest of the system is ideal. In practice, this is an approximate method because the different impairments may be related to each other, and we may not be able to isolate each one by itself. For example, the power penalties due to a nonideal transmitter and crosstalk may be related to each other, whereas chromatic dispersion may be treated as an independent penalty.

11

chapter

Access Networks

IN PREVIOUS CHAPTERS, we have explored the use of optical networks for metro and long-haul network applications. The *access network* is the “last leg” of the telecommunications network that runs from the service provider’s facility to the home or business. With fiber now directly available to many office buildings in metropolitan areas, networks based on SONET/SDH or Ethernet-based technologies are being used to provide high-speed access to large business users. Business users are big consumers of data services, many of which are delivered in the form of leased lines at various speeds ranging from 1.5 Mb/s to several gigabits per second. While this is happening, the telephone and cable companies are also placing a significant emphasis on the development of networks that will allow them to provide a variety of services to individual homes and small to medium businesses. This is the focus of this chapter.

Today, homes get essentially two types of services: plain old telephone service (POTS) over the telephone network and broadcast analog video over the cable network. Recently added to this mix are data services for Internet access using either digital subscriber line (DSL) technology over the telephone network or cable modem service over the cable network.

Early efforts on developing high-capacity access networks were devoted to developing networks that would accommodate various forms of video, such as video-on-demand and high-definition television. However, the range of services that users are expected to demand in the future is vast and unpredictable. Today, end users

Table 11.1 Different types of services that must be supported by an access network. The bandwidth requirements are given for each individual stream.

Service	Type	Downstream Bandwidth	Upstream Bandwidth
Telephony	Switched	4 kHz	4 kHz
ISDN	Switched	144 kb/s	144 kb/s
Broadcast video	Broadcast	6 MHz or 2–6 Mb/s	0
Interactive video	Switched	6 Mb/s	Small
Internet access	Switched	A few Mb/s	Small initially
Videoconferencing	Switched	6 Mb/s	6 Mb/s
Business services	Switched	1.5 Mb/s–10 Gb/s	1.5 Mb/s–10 Gb/s

are interested in both Internet access and other high-speed data access services, for such applications as telecommuting, distance learning, and eventually entertainment video and videoconferencing. Future, unforeseen applications are also sure to arise and make ever-increasing demands on the bandwidth available in the last mile. The term *full service* encompasses the variety of services that are expected to be delivered via access networks. A sampling of the different services and their characteristics is given in Table 11.1. Both telephone and cable companies are striving to become full-service providers.

At a broad level, these services can be classified based on three major criteria. The first is the bandwidth requirement, which can vary from a few kilohertz for telephony to several megahertz per video stream and hundreds of megabits per second for high-speed leased lines. The second is whether this requirement is *symmetric* (two way), for example, videoconferencing, or *asymmetric* (one way), for example, broadcast video. Today, while most business services are symmetric, other services tend to be asymmetric, with more bandwidth needed from the service provider to the user (the downstream direction) than from the user to the service provider (the upstream direction). The last criterion is whether the service is inherently broadcast, where every user gets the same information, for example, broadcast video, or whether the service is switched, where different users get different information, as is the case with Internet access.

In the next section, we provide an overview of the different types of existing and emerging access network architectures. We then provide a more detailed description of the two most promising access architectures—the hybrid fiber coax (HFC) network and the fiber to the curb (FTTC) approach and its variants.

11.1 Network Architecture Overview

In broad terms, an access network consists of a hub, remote nodes (RNs), and network interface units (NIUs), as shown in Figure 11.1. In the case of a telephone company, the hub is a *central office* (also called a *local exchange* in many parts of the world), and in the case of a cable company, it is called a *head end*. Each hub serves several homes or businesses via the NIUs. An NIU either may be located in a subscriber location or may itself serve several subscribers. The hub itself may be part of a larger network, but for our purposes, we can think of the hub as being the source of data to the NIUs and the sink of data from the NIUs. In many cases, rather than running cables from the hub to each individual NIU, another hierarchical level is introduced between the hub and the NIUs. Each hub may be connected to several RNs deployed in the field, with each RN in turn serving a separate set of NIUs. The network between the hub and the RN is called the *feeder network*, and the network between the RN and the NIUs is called the *distribution network*.

We saw that services could be either broadcast or switched. In the same way, the distribution network could also be either broadcast or switched. Note that in the context of services, we are using the terms *broadcast* and *switched* to denote whether all users get the same information or not. In the context of the network, we are referring to the network topology. Different combinations of services and network topologies are possible—a broadcast service may be supported by a broadcast or a switched network, and a switched service may be supported by a broadcast or a switched network. In a broadcast network, an RN broadcasts the data it receives from the feeder network to all its NIUs. In a switched network, the RN processes the data coming in and sends possibly separate data streams to different NIUs. The telephone network that we will study later is a switched network, whereas the cable television network is a broadcast network. Broadcast networks may be cheaper than switched networks, are well suited for delivering broadcast services, and have the advantage that all the NIUs are identical, making them easier to deploy. (In some switched networks that we will study, different NIUs use different wavelengths, which makes it more complicated to manage and track the inventory of NIUs in the network.) Switched networks, as their name suggests, are well suited for delivering switched services and provide more security. For example, it is not possible for one subscriber to tap into another subscriber's data, and it is more difficult for one subscriber to corrupt the entire network. Fault location is generally easier in a switched network than in a broadcast network. In broadcast networks, the “intelligence” is all at the NIUs, whereas in switched networks, it is in the network. Thus NIUs in switched networks may be simpler than in broadcast networks.




Figure 11.1 Architecture of an access network. It consists of a hub, which is a telephone company central office or cable company head end, remote nodes deployed in the field, and network interface units that serve one or more individual subscribers.

Another way of classifying access networks is based on the type of feeder network, which is the network between the hub and the RN. In one scenario, the feeder network could assign each NIU its own *dedicated* bandwidth. By dedicated bandwidth, we mean that different NIUs are assigned different frequency (or wavelength) bands in the frequency (or wavelength) domain. In another scenario, the feeder network could have a total bandwidth that is *shared* by all the NIUs. By shared bandwidth, we mean that multiple NIUs share a given bandwidth in the time domain. In this case, each NIU could potentially access the entire bandwidth for short periods. For upstream transmission from the NIUs back to the hub, we will need some form of media access control to coordinate access to the shared bandwidth by the NIUs. If the traffic from/to the NIUs is bursty, it is more efficient to share a large total amount of bandwidth among many NIUs rather than assign each NIU its own dedicated bandwidth. On the other hand, with dedicated bandwidth, each NIU can be guaranteed a certain quality of service, which is more difficult to do with shared bandwidth. A disadvantage of the shared bandwidth approach is that each NIU must have optics/electronics that operate at the total bandwidth of the network as opposed to the bandwidth needed by the NIU.

Table 11.2 classifies the different types of access networks that we will be studying in this chapter according to whether their distribution network is broadcast or switched, and whether they use dedicated or shared bandwidth in the feeder network. For example, the telephone network is a switched network with each NIU getting its own dedicated bandwidth of 4 kHz. The cable network is a broadcast network with

Table 11.2 Classification of different types of access networks, from [FRI96]. The acronyms refer to the following: HFC—hybrid fiber coax network; DSL—digital subscriber loop; and PON—passive optical network, with the T standing for telephony, W for wavelength, and WR for wavelength routed.

Distribution Network	Feeder Network	
	Shared	Dedicated
Broadcast	Cable TV (HFC), TPON	WPON
Switched		Telephony, DSL, WRPON

all NIUs sharing the total cable bandwidth. A broadcast star WDM passive optical network (WPON), with each NIU assigned a separate wavelength, is an example of a broadcast network but with dedicated bandwidth to each NIU. We will study this architecture in Section 11.3.

Today, two kinds of access networks reach our homes: the telephone network and the cable network. The telephone network runs over twisted-pair copper cable. It consists of point-to-point copper pairs between the telco central office and the individual home. The two wires in a pair are twisted together to reduce the crosstalk between them, hence the name *twisted pair*. This plant was designed to provide 4 kHz bandwidth to each home, although we will see that much higher bandwidths can be extracted out of it using contemporary signal-processing techniques. Wires from individual homes are aggregated as shown in Figure 11.2. The telephone network is a switched network that provides dedicated bandwidth to each user.

A typical cable network is shown in Figure 11.3. It consists of fibers between the cable company head end (analogous to a telco central office) and remote (fiber) nodes. Usually, the channels from the head end are broadcast to the remote nodes by using subcarrier multiplexing (SCM) on a laser (see Section 4.2 to understand how SCM works). From the remote node, coaxial cables go to each home. A remote node serves between 500 and 2000 homes. Such a network is called a hybrid fiber coax (HFC) network. The cable bandwidth used is between 50 and 550 MHz, and the cable carries up to 78 AM-VSB (amplitude-modulated vestigial sideband) television signals in channels placed 6 MHz apart in the American NTSC (National Television System Committee) standard. A return path in the 5 to 40 MHz window is available as well. Many cable companies have now upgraded their networks to carry the video channels in digital format. The cable network is a broadcast network where all users share a common total bandwidth. The same set of signals from the head end is delivered to all the homes.




Figure 11.2 The twisted-pair telephone access network, which consists of individual twisted pairs routed from the central office (CO) to the individual subscribers.




Figure 11.3 The hybrid fiber coax cable television network. The head end broadcasts signals over fiber to the remote node, which then distributes it to individual subscribers via coaxial cable drops.

The telephone and cable networks are vastly different. The telephone network provides very little bandwidth per home but incorporates sophisticated switching equipment and operations and management systems. The cable network provides a lot of bandwidth to each home, but it is all unidirectional and broadcast, with no switching and very simple management.

Several approaches have been used to upgrade the access network infrastructure to support the emerging set of new services. The *integrated services digital network* (ISDN) today provides 144 kb/s of bandwidth over the existing twisted-pair infrastructure and is available in many metropolitan areas. *Digital subscriber loop* (DSL) is another technique that works over the existing twisted-pair infrastructure but provides significantly more bandwidth than ISDN. DSL uses sophisticated modulation and coding techniques to realize a capacity of a few megabits per second over twisted pair, which is sufficient to transmit compressed video. This requires the central office (CO) and the home to each have a DSL modem. However, DSL has some limitations. The realizable bandwidth is inversely proportional to the distance between the CO and the home, and with today's technology, we can achieve several hundred kilobits per second to a few megabits per second over this infrastructure. The existing twisted-pair infrastructure incorporates several 4 kHz filters that must be removed. The bandwidth on the upstream (return) path is severely limited to a few hundred kilobits per second. Many variations and enhancements of DSL have been proposed. As in the conventional telephone network, ISDN and DSL can be classified as switched networks with dedicated bandwidth per NIU.

Satellites provide another way of delivering access services. The direct broadcast satellite system uses a geosynchronous satellite to broadcast a few hundred channels to individual homes. A satellite may provide more bandwidth than a terrestrial coaxial cable system. However, the main problem is that, unlike terrestrial systems, the amount of spatial reuse of bandwidth possible is quite limited, since a single satellite has a wide coverage area within which it broadcasts the signals. Also there is no easy way to handle the upstream traffic. Today, it is possible to have high-speed Internet access delivered via satellite, with the upstream direction carried over a regular telephone line.

Wireless access is yet another viable option. Although it suffers from limited bandwidth and range, it can be deployed rapidly and allows providers without an existing infrastructure to enter the market. Among the variants are the *multichannel multipoint distribution service* (MMDS) and the *local multipoint distribution service* (LMDS), both of which are terrestrial line-of-sight systems. MMDS provides thirty-three 6 MHz channels in the 2–3 GHz band with a range of 15 to 55 km, depending on the transmit power. LMDS operates in the 28 GHz band with 1.3 GHz of bandwidth and is suitable for short-range (3–5 km) deployment in dense metropolitan areas (the distance is also dependent on the amount of rainfall, as rain attenuates signals in this band). Optical fiberless systems using lasers transmitting over free space into the home are also being developed as an alternative approach. These systems can provide about 622 Mb/s of capacity over a line-of-sight range of 200 to 500 m.

In the context of the next-generation access network, the two main architectures being considered today are the so-called hybrid fiber coax (HFC) approach and the fiber to the curb (FTTC) approach. The HFC approach is still a broadcast architecture, whereas the FTTC approach incorporates switching.

11.2 Enhanced HFC

Although we have used the term HFC to describe the existing cable infrastructure, HFC is also the term used to describe an upgraded version of this architecture, which we will refer to as an *enhanced* HFC architecture. Since both the fiber and the coax cable carry multiple subcarrier modulated streams, and it is a broadcast network, a better term to describe the HFC architecture is *subcarrier modulated fiber coax bus* (SMFCB). The network architecture is essentially the same as that shown in Figure 11.3. In order to provide increased bandwidth per user, the network is being enhanced using a combination of several techniques. First, the transmitted frequency range can be increased, for example, up to 1 GHz from the 500 MHz in conventional HFC systems. Enhanced HFC systems being deployed today in larger metropolitan areas are already delivering up to 862 MHz of bandwidth. Within each subcarrier channel, we can use spectrally efficient digital modulation techniques, such as 256 QAM (quadrature amplitude modulation), which provides a spectral efficiency of 8 bps/Hz. In addition, we can drive fiber deeper into the network and reduce the number of homes served by a remote node down to about 50 homes, from the 500 homes typically served by an HFC network. This is being done today as well. We can also use multiple fibers and multiple wavelengths to increase the overall capacity.

In a typical enhanced HFC architecture, like the existing cable network, downstream data is broadcast from the head end to remote (fiber) nodes by using a passive optical star coupler. In recent deployments, it is common to use high-power 1.55 μm transmitters in conjunction with booster amplifiers to achieve a high split ratio. In addition, signals at 1.3 μm can be multiplexed on the same set of fibers. These 1.3 μm signals can be used in a *narrowcasting* mode. That is, these signals can be transmitted only to a selected set of users, rather than to all users. This feature can be used to provide additional bandwidth for selected groups of users.

From a remote node, several coax trees branch out to the network interface units. An NIU may serve one or more homes. Its function is to separate the signals into telephone signals and broadcast video signals, and to send the telephone signal on twisted pair and the video signal on coax to each home that it serves. Each coax leg serves about 50–500 homes. Logically, the architecture is a broadcast bus, although it is implemented as a combination of optical stars and coax trees/buses. Downstream




Figure 11.4 Bandwidth allocation in an enhanced HFC network.

broadcast video to the home would be sent on analog subcarrier channels. Video signals could be sent as analog AM-VSB streams, compatible with existing equipment inside homes. Digital video, as well as telephony and data services, can be carried over the same infrastructure. In addition, upstream channels can be provided in the 5–40 MHz band, which is not used for downstream traffic. Figure 11.4 shows the bandwidth usage in an enhanced HFC network.

The cable infrastructure has already been upgraded in many cities to provide Internet access services through the use of a specific modem developed for this application, called a cable modem, at the head end and at the home. The modems use a shared media Ethernet-type media access control protocol to provide this service. The peak rate of this service is on the order of a few megabits per second, but is shared among all the users in a neighborhood as the HFC network is fundamentally a broadcast network. The amount of bandwidth available per user depends on how many other users are accessing the network and the traffic generated by the other users.

Clearly, enhanced HFC is the natural evolution path for the cable service providers. It maintains compatibility with existing analog equipment and is an efficient approach to deliver broadcast services. On the other hand, it has the disadvantages of a coax-based solution, such as limited upstream bandwidth, limited reliability, and powering needed for the many amplifiers in the path.

11.3 Fiber to the Curb (FTTC)

In contrast to HFC, in FTTC, data is transmitted digitally over optical fiber from the hub, or central office, to fiber-terminating nodes called *optical network units* (ONUs). The expectation is that the fiber would get much closer to the subscriber with this architecture. Depending on how close the fiber gets to an individual subscriber, different terms are employed to describe this architecture (see Figure 11.5). In the most optimistic scenario, fiber would go to each home, in which case this architecture is called *fiber to the home* (FTTH), and the ONUs would perform the function of the NIUs. For the case where ONUs serve a few homes or buildings, say, 8–64, this can




Figure 11.5 Different types of fiber access networks, based on how close the fiber gets to the end user. In many cases, the remote node may be located at the central office itself. The ONUs terminate the fiber signal, and the links between the ONUs and the NIUs are copper based.

be thought of as *fiber to the curb* (FTTC) or *fiber to the building* (FTTB). Typically, in FTTC, the fiber is within about 100 m of the end user. In this case, there is an additional distribution network from the ONUs to the NIUs. With the fiber to the cabinet (FTTCab) approach, the fiber is terminated in a cabinet in the neighborhood and is within about 1 km of the end user.

To make the FTTC architecture viable, the network from the CO to the ONU is typically a *passive optical network* (PON). The remote node is a simple passive device such as an optical star coupler, and it may sometimes be colocated in the central office itself rather than in the field. Although many different architectural alternatives can be used for FTTC, the term FTTC is usually used to describe a version where the signals are broadcast from the central office to the ONUs, and the ONUs share a common total bandwidth in time division multiplexed fashion.

In the context of FTTC, the feeder network is the portion of the network between the central office and the remote node, and the distribution network is between the remote node and the ONUs. We will see that a variety of different types of architectures can be realized by using different types of sources at the central office combined with different types of remote nodes.

Practically speaking, it is quite expensive today to transmit analog video signals over an all-fiber infrastructure; this may necessitate an analog hybrid fiber coax overlay that carries the analog video signals. The FTTC architecture is sometimes also called *baseband modulated fiber coax bus* (BMFCB) or *switched digital video* (SDV).

Table 11.3 Comparison of different PON architectures. N denotes the number of ONUs in the network. An ONU bit rate of 1 indicates that the ONU operates at the bit rate corresponding to the traffic it terminates rather than the aggregate traffic of N . Node sync refers to whether the nodes in the network must be synchronized to a common clock or not. CO sharing relates to whether the equipment is shared among multiple users or whether separate equipment is required to service each user.

Architecture	Fiber Sharing	Power Splitting	ONU Bit Rate	Node Sync	CO Sharing
All fiber	No	None	1	No	No
TPON	Yes	$1/N$	N	Yes	Yes
WPON	Yes	$1/N$	1	Yes	No
WRPON	Yes	None	1	Yes	Yes

In what follows, we shall concentrate on different alternatives for realizing the portion of the access network that is optical. Optical access network architectures must be simple, and the network must be easy to operate and service. This means that passive architectures, where the network itself does not have any switching in it and does not need to be controlled, are preferable to active ones. Passive networks also do not need to be powered, except at the end points, which provide significant cost savings to operators. Moreover, the ONU itself must be kept very simple in order to reduce cost and improve reliability. This rules out using sophisticated lasers and other optical components within the ONU. Preferably, the components used in the ONU must be capable of operating without any temperature control. The CO equipment can be somewhat more sophisticated, since it resides in a controlled environment, and its cost can be amortized over the many subscribers served out of a single CO.

The optical networks proposed for this application are commonly called PONs (passive optical networks)—all of them use passive architectures. They use some form of passive component, such as an optical star coupler or static wavelength router, as the remote node. The main advantages of using passive architectures in this case come from their reliability, ease of maintenance, and the fact that the field-deployed network does not need to be powered. Moreover, the fiber infrastructure itself is transparent to bit rates and modulation formats, and the overall network can be upgraded in the future without changing the infrastructure itself. Table 11.3 compares the different architectures.

The simplest PON architecture, shown in Figure 11.6(a), uses a separate fiber pair from the CO to each ONU. The main problem with this approach is that the cost of CO equipment scales with the number of ONUs. Moreover, the operator needs




Figure 11.6 (a) The point-to-point fiber approach. (b) In practice the fibers could be laid in the form of a ring.

to install and maintain all these fiber pairs. This approach is being implemented on a limited scale today, primarily to provide high-speed services to businesses. In Japan, NTT is operating such a system at bit rates from 8 to 32 Mb/s over each fiber. Although logically there is a separate fiber pair to each ONU, physically the fibers could be laid in a ring configuration, as shown in Figure 11.6.

Instead of providing a fiber pair to each ONU, a single fiber can be used with bidirectional transmission. However, the same wavelength cannot be used to transmit data simultaneously in both directions because of uncontrolled reflections in the fiber. One way is to use time division multiplexing so that both ends don't transmit simultaneously. Another is to use different wavelengths (1.3 and 1.55 μm , for example) for the different directions.

More commonly, rather than dedicating a fiber pair per user, the fiber pair is shared by many users. The most common example of such networks are the SONET/SDH rings, which are now widely deployed to provide high-speed services to large business customers. These rings operate at speeds ranging from 155 Mb/s to 10 Gb/s. In this case, an ONU is a SONET add/drop multiplexer (ADM), and multiple ONUs can be present on the same ring. However, these rings are not considered part of the PON family. Rather they can be viewed as an alternative fiber access solution.




Figure 11.7 A broadcast and select TPON. The CO broadcasts its signal downstream to all the ONUs using a passive star coupler. The ONUs share an upstream channel in a time-multiplexed fashion. In this case, upstream and downstream signals are carried using different wavelengths over a single fiber.

While SONET/SDH rings are suitable for delivering the higher-speed services and addressing the needs of large business customers, the PON architectures that we will study here can provide a more cost-effective solution for addressing the needs of small- and medium-sized businesses and homes, which require a few DS1 (1.5 Mb/s) lines, DSL lines, or 10 Mb/s Ethernet connections.

The most common PON architecture is the TPON (originally called PON for telephony) architecture [Ste87], shown in Figure 11.7. The downstream traffic is broadcast by a transmitter at the CO to all the ONUs by a passive star coupler. Though the architecture is a broadcast architecture, switched services can be supported by assigning specific time slots to individual ONUs based on their bandwidth demands. For the upstream channel, the ONUs share a channel that is combined using a coupler, again via fixed time division multiplexing (TDM) or some other multiaccess protocol. In the TDM approach, the ONUs need to be synchronized to a

common clock. This is done by a process called *ranging*, where each ONU measures its delay from the CO and adjusts its clock such that all the ONUs are synchronized relative to the CO. The CO then assigns time slots to each ONU as needed.

This architecture allows the relatively expensive CO equipment to be shared among all the ONUs and makes use of fairly mature low-cost optical components. The CO transmitter can be an LED or a Fabry-Perot laser, and cheap, uncooled pinFET receivers and LEDs/Fabry-Perot lasers can be used within the ONUs. The number of ONUs that can be supported is limited by the splitting loss in the star coupler. Each ONU must have electronics that run at the *aggregate* bit rate of all the ONUs. There is a trade-off between the transmit power, receiver sensitivity, bit rate, and number of ONUs (which determines the splitting loss) and the total distance covered.

As we mentioned earlier, TPONs may be more cost-effective at offering lower-speed services compared to SONET/SDH rings or Ethernet-based offerings. TPON vendors claim that it is easier to provision bandwidth in a flexible manner remotely by changing the number of time slots an individual subscriber is assigned. However, the current generation of SONET/SDH products is being enhanced to provide dynamic bandwidth provisioning as well. In a TPON, a failure of one subscriber's equipment does not affect other subscribers, whereas a SONET/SDH ring node failing affects all the nodes on the ring. However, SONET/SDH has built-in protection mechanisms to reroute traffic in the event of both equipment failures and fiber cuts and restore services rapidly. In contrast, dealing with fiber cuts is not easy in the TPON architecture, without doubling up on the fiber plant. By the same token, with the TPON architecture, additional subscribers can be added without affecting any of the other subscribers. In SONET/SDH rings, this is a more complex process.

TPON development has accelerated with the establishment of the *full service access network* standard by a large group of service providers and equipment companies. The standard specifies an ATM-based TPON architecture with a downstream bandwidth of up to 622 Mb/s and an upstream bandwidth of up to 155 Mb/s. The targeted distance is 20 km with a total fiber attenuation in the 10–30 dB range. Practical link budgets using lasers at the CO and ONUs allow a 16- to 32-way split with this approach. For example, a TPON operating at 622 Mb/s using a 32-way splitter can provide each subscriber with about 20 Mb/s of bandwidth. The TPON can operate over a single fiber pair by using different wavelengths in the upstream (1.3 μm) and downstream (1.55 μm) directions. Alternatively, it can also operate over a fiber pair using 1.3 μm transmitters. As of this writing, FSAN-based TPONs are beginning to be deployed in the field. Some carriers, notably BT, have deployed early versions of TPONs for several years, primarily for telephony. Their deployment may become more widespread as optical component costs come down.




Figure 11.8 A broadcast-and-select WDM PON (WPON), which is an upgraded version of the basic PON architecture. In this case, the CO broadcasts multiple wavelengths to all the ONUs, and each ONU selects a particular wavelength. As in a conventional TPON, the ONUs time-share an upstream channel at a wavelength different from the downstream wavelengths.

Many enhancements have been proposed to the basic TPON architecture to increase its capacity and flexibility. The next step, shown in Figure 11.8, is to replace the single transceiver at the CO with a WDM array of transmitters or a single tunable transmitter to yield a WDM PON (WPON). This approach allows each ONU to have electronics running only at the rate it receives data, and not at the aggregate bit rate. However, it is still limited by the power splitting at the star coupler.

Introducing wavelength routing solves the splitting loss problem while retaining all the other advantages of the WDM PON. In addition, it allows point-to-point dedicated services to be provided to ONUs. This leads to the WRPON architecture shown in Figure 11.9.

Several types of WRPONs have been proposed and demonstrated. They all use a wavelength router, typically an arrayed waveguide grating (AWG) for the




Figure 11.9 A wavelength-routing PON (WRPON). In this case, a passive arrayed waveguide grating (AWG) is used to route different wavelengths to different ONUs in the downstream direction, without incurring a splitting loss. As in the TPON and WPON architectures, the ONUs time-share a wavelength for upstream transmission.

downstream traffic, but vary in the type of equipment located at the CO and ONUs, and in how the upstream traffic is supported. The router directs different wavelengths to different ONUs. The earliest demonstration was the so-called passive photonics loop (PPL) [WKR⁺88, WL88]. It used 16 channels in the 1.3 μm band for downstream transmission and 16 additional channels in the 1.55 μm band for upstream transmission. However, this approach is not economical because we need two expensive lasers for each ONU—one inside the ONU and one at the central office. We describe several variants of this architecture that provide more economical sharing of resources at the CO and ONUs.

The RITENET architecture [Fri94] (see Figure 11.10) uses a tunable laser at the CO. A frame sent to each ONU from the CO consists of two parts: a *data* part, wherein data is transmitted by the CO, and a *return traffic* part, wherein no data




Figure 11.10 The RITENET WRPON architecture. The ONUs use an external modulator to modulate an unmodulated signal transmitted from the CO.

is transmitted but the CO laser is left turned on. Each ONU is provided with an external modulator. During the return traffic part of the frame, the ONU uses the modulator to modulate the light signal from the CO. This avoids the need for having a laser at the ONU. The upstream traffic from the ONUs is also sent to the router. The router combines all the different wavelengths and sends them out on a common port to a receiver in the CO. If a single receiver is used in the CO, then the ONUs must use time division multiplexing to get access to that receiver. Alternatively, if a separate receiver is used for each wavelength at the CO, each ONU gets a dedicated wavelength to transmit upstream back to the CO. This architecture avoids the need for having a laser at each ONU. Instead, each ONU has an external modulator.

A lower-cost alternative to RITENET is the LARNET architecture [ZJS⁺95] (see Figure 11.11), which uses an LED at the ONU instead of an external modulator for transmission in the upstream direction. The LED emits a broadband signal that gets “sliced” upon going through the wavelength router, as shown in Figure 11.12. Only




Figure 11.11 The LARNET WRPON architecture. A broadband signal from the LED at the CO is split into individual wavelength components by the AWG and broadcast to all the ONUs.

the power in the part of the LED spectrum corresponding to the passband of the wavelength router is transmitted through to the receiver at the CO. Note, however, that with N ONUs, this imposes a splitting loss of at least $1/N$ —only a small fraction of the total power falls within the passband of the router.

More important, an LED can be used at the CO as well [IFD95] for downstream transmission. In this case, the signal sent by the CO LED effectively gets broadcast to all the ONUs. It is in fact possible to have two transmitters within the CO: an LED, say, at $1.3 \mu\text{m}$, broadcasting to all the ONUs, and a tunable laser at $1.55 \mu\text{m}$ selectively transmitting to the ONUs. This is an important way to carry broadcast analog video signals over the digital switched fiber infrastructure at low cost without having to use a separate overlay network for this purpose.

WDM components for PONs are not yet mature and are more expensive than the components required for simple broadcast PONs. However, WRPONs offer much




Figure 11.12 Spectral slicing: if a broadband LED signal is sent through a filter, only the portion of the LED spectrum that is passed by the filter comes out.

higher capacities than the simple broadcast PONs, and simple PONs can be upgraded to WRPONs as the need arises.

11.3.1 PON Evolution

We have studied a number of PON variants in this section. It is important to realize that there is a nice evolution path from a very simple TPON architecture to some of the more complex WRPON architectures. The evolution can be performed with minimal disruption of existing services and without wasting already-deployed equipment. In general, the terminal equipment can be upgraded as additional capacity and services are needed, without having to upgrade the outside fiber plant, which is a true long-term investment. The upgrade scenario for PONs could go as follows. The operator can start by deploying a simple broadcast TPON, which is a broadcast star network with shared bandwidth, according to the classification of Table 11.2. If more ONUs need to be supported, the operator can upgrade the network to a WDM broadcast PON, which is a broadcast network with dedicated bandwidth provided to each ONU. This can be done by upgrading the transmitters at the CO to WDM transmitters, and the operator may be able to reuse the existing ONUs. If higher capacities per ONU are needed, the operator can further upgrade the network to a wavelength-routed PON, which is a switched network with dedicated

bandwidth. Moreover, this wavelength-routed PON can also support broadcast services efficiently using the spectral slicing technique described earlier. Thus there is an upgrade path starting from a broadcast network with shared bandwidth to a broadcast network with dedicated bandwidth and eventually to a switched network with dedicated bandwidth.

Summary

Service providers, both telephone operators and cable companies, are actively looking to deploy broadband access networks to provide a variety of new services. Fiber-based services are now available for many businesses in metropolitan areas. When it comes to residential access, however, fiber is yet to reach the home. SONET/SDH ring-based architectures have been deployed to support the needs of large business customers, but they are not as suited for supporting the needs of residential users and small business customers. The two main architectures for broadband access networks are the hybrid fiber coax (HFC) architecture, which is based on evolving the current plant deployed by cable television operators, and the fiber to the curb (FTTC) architecture, or equivalently a passive optical network (PON) architecture. Compared to the HFC approach, FTTC has a higher initial cost, but provides bandwidth deeper in the network and may prove to be a better longer-term solution. Although FTTC refers to a simple broadcast TDM star PON architecture, we also explored several upgrade options of the PON approach that provide higher capacities by making clever use of wavelength division multiplexing techniques. A number of major telephone carriers and manufacturers in the world have gotten together and defined the requirements for an FTTC-based architecture to enable them to deploy the full service access network [FSA98]. As of this writing, PONs based on this architecture are just beginning to be deployed.

FTTC is attractive in places where coaxial cable is not already deployed, which is the case in many countries other than the United States. FTTC also makes sense for telephone companies who lack a cable infrastructure.

Variants of FTTC have been around a long time, but deployment has been slow for several reasons. First, there is significant cost associated with building and deploying a new access network, which can take several years to pay back. Therefore there is a big barrier toward making the investment in the first place. Second, this is coupled with the uncertain outlook in terms of the revenue that can be generated from the investment. Third, optical component costs are only now starting to decline, with the development of components especially optimized for PON applications, such as low-cost, uncooled semiconductor lasers and transceivers.

The HFC approach, on the other hand, is attractive in places where coaxial cable is already deployed to the home, such as the United States. It is the logical evolution

choice for cable companies who have already deployed a simpler version of the HFC architecture to provide basic cable television service.

As optical component costs come down and bandwidth needs increase, it is clear that optical fiber will play a major role in access networks—the question is how close will it get to our homes?

Further Reading

There is a vast body of literature on access networks, and several conferences have sessions devoted to it. There is an informative Web page maintained by the DSL forum (<http://www.adsl.com>). See also [Bha99] for a nice overview of the different types of DSL.

The papers in [Fra98, Aar95, Kob94, KKS00, SKY89] describe plans for deploying fiber in the access network and compare different architectural approaches. TDM PONs were first proposed in [Ste87]. At this time, a group of telephone carriers and manufacturers have gotten together and agreed upon a standard for a fiber-based access architecture, called the full service access network (FSAN) initiative. Details may be found on the Web at www.fsanet.net and in [FSA98, Qua98]. The International Telecommunications Union (ITU) has published a requirements standard based on FSAN [ITU98].

[FRI96, FHJ⁺98, VMVQ00] describe some possible evolutions of the basic TPN architecture by making clever use of WDM and optical amplifiers. A variety of WDM PONs are described in [WKR⁺88, WL88, Fri94, ZJS⁺95, IFD95, IRF96].

Problems

- 11.1 Do a power budget calculation for the different types of PON architectures considered in this chapter and determine the number of ONUs that can be supported in each case, assuming the following parameters:

Laser output power	−3 dBm
LED output power	−20 dBm
Transmit bit rate	155 Mb/s
Receiver sensitivity	−40 dBm
Fiber loss, including connectors	10 dB
1 × 8 wavelength router loss	5 dB
1 × 32 wavelength router loss	9 dB
1 × 64 wavelength router loss	12 dB
Excess splitter loss	1 dB

The normal wavelength router losses are indicated above. However, with spectral slicing, an additional loss is also incurred as only a small fraction of the spectrum is transmitted out of each port on the wavelength router. Assume that in addition to the standard loss, we get only $1/2N$ of the transmitted power in each channel, where N is the number of ONUs.

- 11.2 Consider the RITENET architecture shown in Figure 11.10. Suppose the laser speed at the CO is limited to 155 Mb/s. The network needs to support 20 ONUs and provide each ONU with 10 Mb/s bandwidth from the CO to the ONU and 2 Mb/s from the ONU to the CO. How could you modify the architecture to support this requirement?

References

- [Aar95] R. Aaron, editor. *IEEE Communications Magazine: Special Issue on Access to Broadband Services*, volume 33, Aug. 1995.
- [Bha99] V. K. Bhagavath. Emerging high-speed xDSL access services: Architectures, issues, insights, and implications. *IEEE Communications Magazine*, 37(11):106–114, Nov. 1999.
- [FHJ⁺98] R. D. Feldman, E. E. Harstead, S. Jiang, T. H. Wood, and M. Zirngibl. An evaluation of architectures incorporating wavelength division multiplexing for broad-band fiber access. *IEEE/OSA Journal on Lightwave Technology*, 16(9):1546–1559, Sept. 1998.
- [Fra98] P. W. France, editor. *BT Technology Journal—Special Issue on Local Access Technologies*, volume 16, Oct. 1998.
- [Fri94] N. J. Frigo et al. A wavelength-division-multiplexed passive optical network with cost-shared components. *IEEE Photonics Technology Letters*, 6(11):1365–1367, 1994.
- [FRI96] N. J. Frigo, K. C. Reichmann, and P. P. Iannone. WDM passive optical networks: A robust and flexible infrastructure for local access. In *Proceedings of International Workshop on Photonic Networks and Technologies*, pages 201–212, 1996.
- [FSA98] *Full Services Access Network Requirements Specification*, 1998. Available on the Web at www.fsanet.net.
- [IFD95] P. P. Iannone, N. J. Frigo, and T. E. Darcie. WDM passive optical network architecture with bidirectional optical spectral slicing. In *OFC'95 Technical Digest*, pages 51–53, 1995. Paper TuK2.

- [IRF96] P. P. Iannone, K. C. Reichmann, and N. J. Frigo. Broadcast digital video delivered over WDM passive optical networks. *IEEE Photonics Technology Letters*, 8(7):930–932, 1996.
- [ITU98] ITU-T. *Recommendation G.983: Broadband Optical Access Systems Based on Passive Optical Networks*, 1998.
- [KKS00] D. Kettler, H. Kafka, and D. Spears. Driving fiber to the home. *IEEE Communications Magazine*, 38(11):106–110, Nov. 2000.
- [Kob94] I. Kobayashi, editor. *IEEE Communications Magazine: Special Issue on Fiber-Optic Subscriber Loops*, volume 32, Feb. 1994.
- [Qua98] J. A. Quayle et al. Achieving global consensus on the strategic broadband access network—the full service access initiative. *BT Technology Journal*, 16(4):58–70, Oct. 1998.
- [SKY89] P. W. Shumate, O. Krumpholz, and K. Yamaguchi, editors. *IEEE/OSA JLT/JSAC Special Issue on Subscriber Loop Technology*, volume 7, Nov. 1989.
- [Ste87] J. Stern et al. Passive optical local networks for telephony applications. *Electronics Letters*, 23:1255–1257, 1987.
- [VMVQ00] I. Van de Voorde, C. M. Martin, J. Vandewege, and X. Z. Qiu. The superPON demonstrator: An exploration of possible evolution paths for optical access networks. *IEEE Communications Magazine*, 38(2):74–82, Feb. 2000.
- [WKR⁺88] S. S. Wagner, H. Kobrinski, T. J. Robe, H. L. Lemberg, and L. S. Smoot. Experimental demonstration of a passive optical subscriber loop architecture. *Electronics Letters*, 24:344–346, 1988.
- [WL88] S. S. Wagner and H. L. Lemberg. Technology and system issues for the WDM-based fiber loop architecture. *IEEE/OSA Journal on Lightwave Technology*, 7(11):1759–1768, 1988.
- [ZJS⁺95] M. Zirngibl, C. H. Joyner, L. W. Stulz, C. Dragone, H. M. Presby, and I. P. Kaminow. LARnet, a local access router network. *IEEE Photonics Technology Letters*, 7(2):1041–1135, Feb. 1995.