MTP
MACMILLAN TECHNICAL PUBLISHING U·S·A

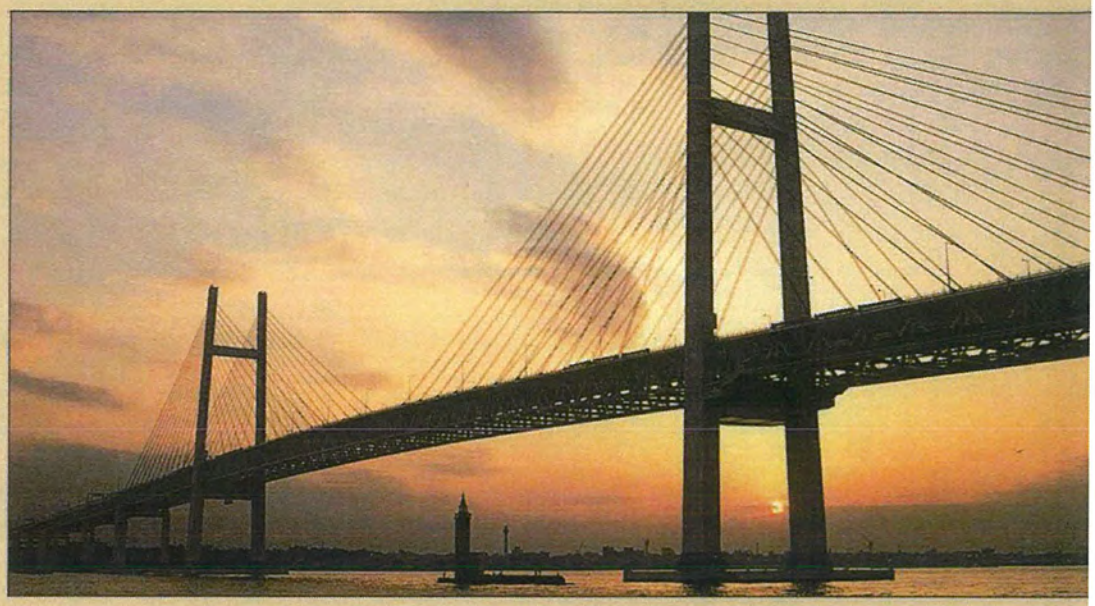

MACMILLAN TECHNOLOGY SERIES

ADSL/VDSL PRINCIPLES

A Practical and Precise Study of

Asymmetric Digital Subscriber Lines and

Very High Speed Digital Subscriber Lines

Dr. Dennis J. Rauschmayer

Comcast, Ex. 1031
Comcast v. Entropic
IPR2024-00431

PENGAD 800-631-6989
PLAINTIFF'S EXHIBIT
1031

ADSL/VDSL PRINCIPLES:

A Practical and Precise Study of Asymmetric Digital Subscriber Lines and Very High Speed Digital Subscriber Lines

Dr. Dennis J. Rauschmayer

MTP
MACMILLAN TECHNICAL PUBLISHING U·S·A

ADSL/VDSL Principles

By Dr. Dennis J. Rauschmayer

Published by:
Macmillan Technical Publishing
201 West 103rd Street
Indianapolis, IN 46290 USA

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

Copyright © 1999 by Macmillan Technical Publishing

Printed in the United States of America 1 2 3 4 5 6 7 8 9 0

02 01 00 99 4 3 2 1

Library of Congress Cataloging-in-Publication Data

Library of Congress Cataloging Number: 96-80467

ISBN: 1-57870-015-9

Universal Product Code: 619472001592

Warning and Disclaimer

This book is designed to provide information about Digital Subscriber Line technology. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an "as is" basis. The authors and Macmillan Technical Publishing shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the disks or programs that may accompany it.

Trademark Acknowledgments

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Macmillan Technical Publishing cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Kurt K. Wendt Library
University of Wisconsin-Madison
215 N. Randall Avenue
Madison, WI 53706-1688

Development Editor
Brad Miser
Christopher Cleveland

Project Editor
Dayna Isley

Copy Editors
Geneil Breeze
June Waldman

Technical Reviewer
Jacky Chow

Team Coordinator
Amy Lewis

Manufacturing Coordinator
Brook Farling

Book Designer
Anne Jones

Cover Designer
Sandra Schroeder

Cover Production
Aren Howell

Director of Production
Larry Klein

Production Manager
Laurie Casey

Graphics Image Specialist
Oliver Jackson

Production Team
Louis Porter, Jr.

Indexer
Bront Davis

Proofreader
Mary Ellen Stephenson

OVERVIEW

1 Introduction and Motivation, 1

2 Twisted Pair Environment, 23

3 Loop Analysis, 63

4 Power Spectral Densities and Crosstalk Models, 89

5 DSL Theoretical Capacity in Crosstalk Environments, 131

6 DSL Modulation Basics, 157

7 ADSL Modulation Specifics, 207

8 ADSL in WAN Networks, 253

9 VDSL Overview, 283

Index, 307

TABLE OF CONTENTS

1 *Introduction and Motivation* 1

Telephone and Data Systems 2
- Coexistence 2
- Subscriber Lines 4

Motivation for Higher Bit Rates 5
- Growth of the PC 5
- Bulletin Board Services 6
- Internet 6
- World Wide Web 6
- Internet Service Providers 7
- Video on Demand 8
- Small Office/Home Office 8

Modem Evolution 9

Why DSL Technology? 10

The Market Demand 11

The DSLs 12
- ADSL 13
- HDSL 13
- HDSL2 14
- IDSL 14
- MDSL 14
- RDSL and RADSL 14
- SDSL 15
- VDSL 15

Other High-Speed Data Options 15
- Cable Modems 15
- Wireless Technologies 16

End-to-End System Diagram 16
- Customer Interface 16
- The POTS Splitter 17
- The Twisted Pair 17
- xTU-R 17
- xTU-C 17
- DSLAM 18
- NSP Network 18
- Line Codes 18
- Data Link Framing and Encapsulation 19
- Backbone/ATM Cloud 19

ADSL and VDSL Standardization 19
ADSL Standardization 19
VDSL Standardization 22
Summary 22

2 Twisted Pair Environment 23

A Brief History of Twisted Pairs 23
Introduction to Twisted Pairs 25
Differential Mode and Common Mode Signals 25
Two-Wire Full-Duplex Systems 27
Physical Characteristics 28
Cable Bundles 28
Gauges and Line Characteristics 29
Electrical Characteristics of Twisted Pairs 30
Analytic RLCG Models 30
Numeric RLCG Models 31
Transmission Line Parameters 35
Propagation Constant 36
Characteristic Impedance 41
Crosstalk 42
Crosstalk Types 43
Unbalance Models 44
General Unbalance Expression 53
Near-end Crosstalk (NEXT) 53
Far-end Crosstalk (FEXT) 57
Summary 59
Exercises 60
Endnotes 60

3 Loop Analysis 63

ABCD Parameters 63
ABCD Parameters of a Lumped Impedance 65
ABCD Parameters of a Uniform Twisted Pair 66
Bridged Taps 70
Loop Analysis 71
Input Impedance, Transfer Functions, and Insertion Loss 72
Basic Loop Configurations 76
ADSL Loops 76
VDSL Loops 82

Summary 86
Exercises 86
Endnotes 87

4 *Power Spectral Densities and Crosstalk Models* 89

Introduction 90
Crosstalk Equations 91
Number of Crosstalkers 92
Types of Disturbers 92
ISDN and HDSL 94
T1 Lines 99
DMT ADSL (Downstream) 104
DMT ADSL (Upstream) 109
CAP/QAM ADSL 112
VDSL PSDs 121
Summary of PSDs 123
NEXT and FEXT Expressions 123
NEXT Disturbers 124
FEXT Disturbers 127
Summary 128
Exercises 129

5 *DSL Theoretical Capacity in Crosstalk Environments* 131

A Review of Capacity 132
Entropy and Information 132
Gaussian Channels 135
Channels with Memory and Interference 136
Framework for Capacity Calculations for ADSL and VDSL 141
ADSL Capacities 142
Loop Capacities with Identical Crosstalk 142
Comparison of Different Types of Crosstalk 143
Industry Standard Loop and Crosstalk Configurations 145
VDSL Capacities 145
Usable Frequencies on Crosstalk Channels 147
Summary 150
Exercises 150
Endnotes 151
Appendix 5A 152

6 DSL Modulation Basics 157

Modulation/Demodulation Blocks 158
- QAM Modulation 158
- CAP Modulation 164
- Multitone Modulation 166

Forward Error Correction 172
- Reed-Solomon Coding 173
- Convolution Codes 174

Interleaving 182
- Block Interleaving 182
- Convolutional Interleaving 184

Scrambling 185

Equalization 187
- Folded Spectrums 189
- Linear Equalization 191
- Decision Feedback Equalization 194
- Tomlinson-Harashima Precoding 199
- Frequency-Domain Equalization 201

Shaping 203

Summary 203

Exercises 203

Endnotes 204

7 ADSL Modulation Specifics 207

DMT ADSL Systems 207
- Transmitter and Receiver Blocks 209
- DMT ADSL Overhead 228
- ADSL Reduced Overhead Modes 233
- Physical Layer Training 235
- ATM over ADSL 240

CAP/QAM ADSL Systems 240

Summary 249

Exercises 249

Endnotes 250

8 ADSL in WAN Networks 253

Protocol Basics 254
- Ethernet 255

IP 256
TCP 263
ATM 263
PPP 265
Tunneling 266

End-to-End Players *268*
Users 269
Network Access Providers 269
Network Transport Providers 269
Network Service Providers 270
ILECs and CLECs 270

Architectures for DSL Networks *270*
Bridging Architectures 270
ATM End-to-End Architectures 273
Tunneling Architectures 276
Premise Network Considerations 278

Summary *280*

Exercises *280*

Endnotes *280*

9 VDSL Overview ***283***

VDSL Deployment Locations *283*
Crosstalk Issues 286
RFI Issues 288

VDSL Requirements *291*
Rate and Reach Goals 291
Frequency Allocation 292

Modulation Methods for VDSL *295*
TDD Approaches 295
FDD Approaches 302

VDSL in WAN Networks *304*
Video Services 304
Data Services 305
The Full-Service Network 305

Summary *305*

Endnotes *305*

Index ***307***

This material may be protected by Copyright law (Title 17 U.S. Code)

CHAPTER 2

Twisted Pair Environment

In this chapter:

- History of twisted pairs
- Physical and electrical characteristics of twisted pairs
- Transmission line parameters
- Derivation of basic crosstalk equations

This chapter is devoted to the physical and electrical characteristics of twisted pairs. A twisted pair channel is noisy, lossy, and prone to crosstalk. ADSL and VDSL modulation schemes are designed to overcome and mitigate these impairments. Thorough understanding of ADSL and VDSL begins with thorough understanding of the environment in which they must exist.

This chapter discusses the physical and electrical properties of common twisted pairs. The transmission line characteristics and basic crosstalk equations in twisted pair bundles will be derived. This information becomes useful in subsequent chapters dealing with transmission performance. But first, an introduction to twisted pairs is in order.

A Brief History of Twisted Pairs

The history of telephone wires is more than a century old and has roots in the wires originally used to carry information from the telephone's predecessor, the telegraph.

Samuel F. B. Morse constructed the first commercial telegraph line between Baltimore, Maryland, and Washington, D.C., in 1844. After failing to successfully install underground lines due to high leakage from the environment, Morse decided that an aerial approach

would be better. Morse proceeded to install 700 wooden poles spaced about 300 feet apart for the 40-mile span. A 2-foot cross arm was nailed to the top of each pole, and a wire was run through a slot at the end of each cross arm. Thus, the groundwork was laid for the deployment of twisted pairs.

After the telephone began to gain popularity, the need for many more wires was necessary, and the telephone cable—consisting of many wires bound in a sheath—was developed. One of the biggest problems at first was crosstalk between wires, as the two-wire metallic circuit had not yet been employed. In New York City, in 1879, cable was laid across the Brooklyn Bridge, years before it was even open. In 1880, an electrician speaking for the Metropolitan Telephone and Telegraph Company reported:

> We have over the East River Bridge at the present time, four cables, 3800 feet long, each cable with seven conductors.... In using the cables and talking on one wire, you could hear whatever was said on another wire.[1]

In the following years, many schemes were employed to reduce crosstalk, most dealing with shielding the individual wires or providing grounding at intervals. Finally in 1881, a patent for the metallic twisted pair circuit, which worked well in suppressing crosstalk, was given to Alexander Graham Bell, who testified to the Honorable Commissioner of Patents about his experiments as follows:

> The methods by which I sought to accomplish this result, on the 26th of November, 1876 was to employ a metallic circuit, the direct and return wires which should be parallel to the disturbing wires and to one another.... The improvements in the arrangement that I devised in England were:
>
> To place the direct and return lines close together...and
>
> To twist the direct and return lines around one another so that they should be absolutely equidistant from the disturbing wires.[1]

Thus, the twisted pair was born. Since that time, the deployment of twisted pair lines has been extensive and includes many types and sizes of cable.

Several surveys have been conducted on deployed twisted pair subscriber loop statistics with some results including average lengths, impedances, and losses.[2,3]

Note

For an interesting history of all methods of mass communications in the United States, including the rise of the telegraph and telephone and the development of the twisted pair, I highly recommend visiting the communications exhibits at the Smithsonian American History Museum in Washington, D.C.

Introduction to Twisted Pairs

Before getting into the technical and mathematical details regarding twisted pairs, let's review exactly what the term twisted pair means and why, in nontechnical terms, twisted pairs are used for communications.

As described in the previous section, a twisted pair is made up of two wires twisted around one another. The two wires are electrically driven and sensed differentially, and neither is connected to a common ground point. Because the wires are physically close and geometrically symmetric to one another, any outside signal influencing one of the wires influences the other wire in the same way. If a receiver senses only the difference between the two wires, the outside influence subtracts out. Historically, the two wires making up a twisted pair are called tip and ring.

Note

The names *tip* and *ring* come from connector cables used by operators in which one wire terminated at the tip of the connector, and one wire terminated in a ring around the connector. The names are still widely used today.

Differential Mode and Common Mode Signals

An important concept to understand about twisted pairs is the difference between differential mode and common mode signals. A *differential mode signal* is applied between the two wires making up the twisted pair. A *common mode signal* is applied between the two individual wires of the twisted pair and ground. Differential mode signals are sometimes called *metallic signals*, and common mode signals are sometimes called *longitudinal signals*. Figure 2.1 illustrates differential mode signals and common mode signals.

Figure 2.1 Differential mode and common mode signals.

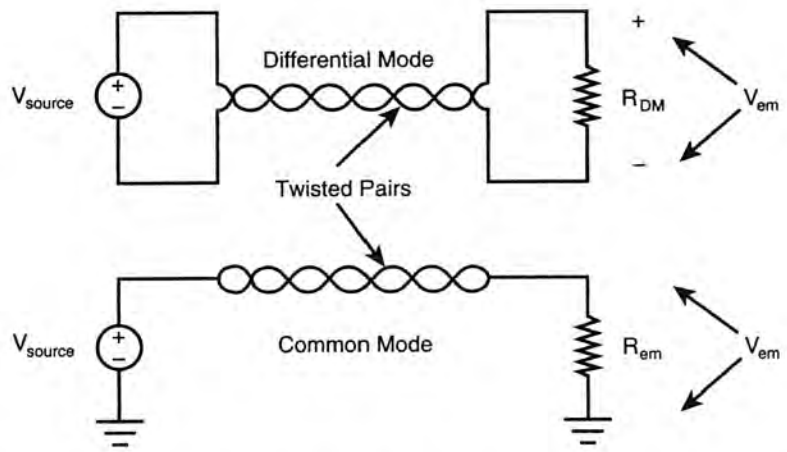

Data and voice signals on a twisted pair are always transmitted as differential signals. Common mode signals are typically due to outside influences such as 60 Hz induction from nearby power lines, RF interference from radio stations and other active operating bands, and disturbances from signals on other twisted pairs.

Note

In some systems, such as HDSL powering units, the common mode is used to transmit D.C. power for remote unit powering.

At the receiver of a twisted pair system, only the differential voltage of the twisted pair is processed for information. The common mode signal is rejected by transformers and well-balanced differential amplifiers, which effectively subtract the voltage on the individual wires of the twisted pair.

For example, consider the twisted pair shown in Figure 2.2.

Voltages of each wire are shown along the length of the pair with respect to ground. Note that the voltage on the twisted pair appears as a 1000 Hz, 1 V peak sine wave riding on a 60 Hz wave. The 60 Hz is in phase on each wire, and the 1000 Hz wave is 180 degrees out of phase on each wire. After the differential amplifier, which essentially subtracts the voltage on one wire from the voltage on the other, the 60 Hz voltage subtracts away and is no longer present, and only the 1000 Hz waveform remains (in this simple example, loss on the twisted pair is ignored). This is sometimes referred to as common mode rejection.

Figure 2.2 An example of common mode rejection on a twisted pair.

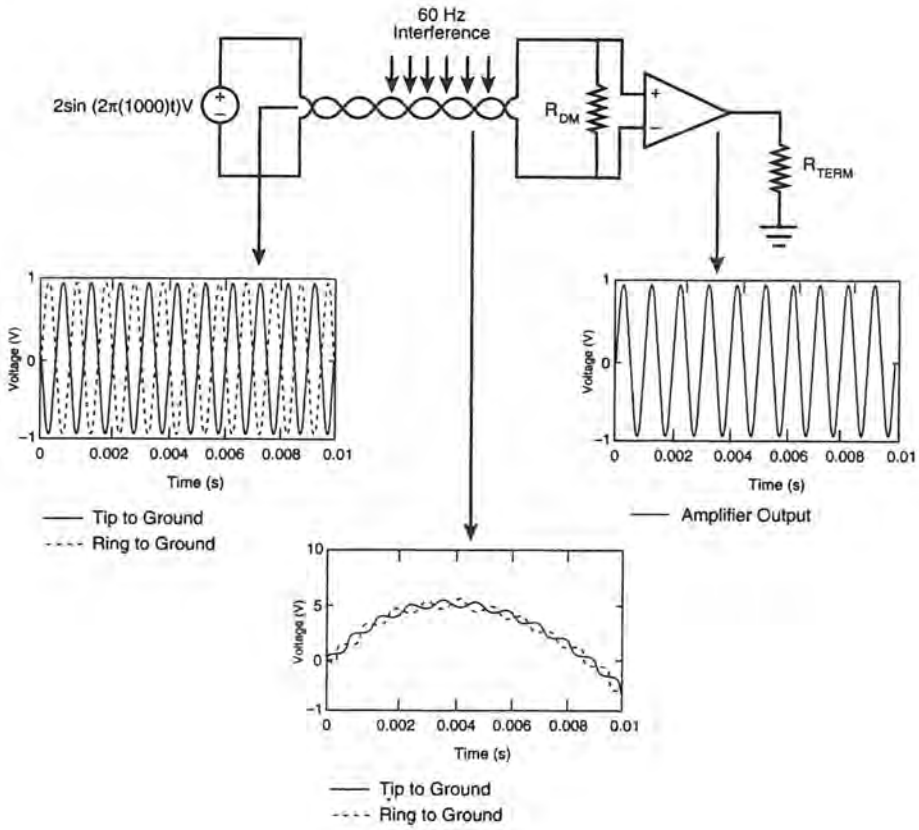

Two-Wire Full-Duplex Systems

If you don't have experience with twisted pair applications, the use of two wires to send signals in both directions simultaneously, or full-duplex operation, might seem extraordinary. If fact, full duplex use of a twisted pair is common and simple. Figure 2.3 shows a basic scheme.

Here, the twisted pair part of the circuit is called the *two-wire path*, and the transmit and receive portion on each end is called the *four-wire path*. The conversion at each end of the twisted pair between two-wire and four-wire paths is called *the two-wire to four-wire hybrid*, or simply the *hybrid*. Signals arriving at point A pass through TX1 and on through the transformer to the twisted pair. Some portion of the signal at point B will enter the receive path of side 1. This signal is called the *echo signal*. The echo signal is canceled by a balance filter B1 and thus will not pass to point C. The original signal will arrive and pass through side 2's transformer and into the receive path of side 2. If side 2 is simultaneously transmitting a signal, any echo from that signal will be canceled by B2 and will not affect the signal proceeding to point D.

Figure 2.3 A full-duplex twisted pair circuit with two-wire to four-wire hybrids on each end.

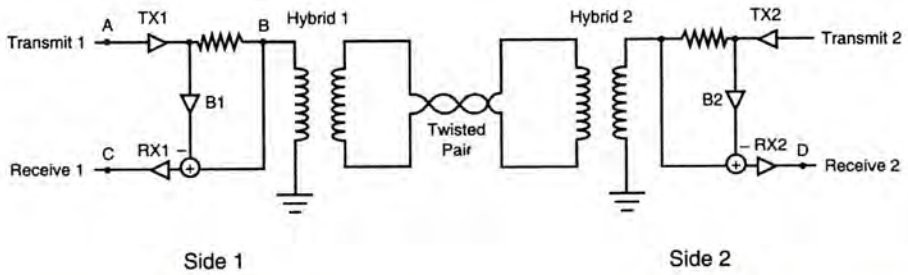

Note

Take note that different hybrid designs exist, including designs using transformers to cancel the anticipated echo, using analog or digital amplifiers to cancel the "average expected" echo, and using adaptive filters to measure the echo and adjust to do cancellation. Most high-speed data systems use the latter method.

Using a hybrid to achieve full-duplex communications on a twisted pair is typical of voice systems and some *echo canceled systems* including ISDN, HDSL, and some ADSL systems. Other systems, especially high bit rate systems, do not use echo cancellation because they often require complex echo cancellation filters and front-end analog-to-digital converters. Instead, these systems may use *frequency division multiplexing* (FDM) or *time division multiplexing* (TDM). Such systems are discussed in detail in later chapters.

Physical Characteristics

With the massive number of twisted pairs deployed over many years, care must be taken not to mix generalizations of the characteristics of common twisted pairs with the characteristics of all deployed twisted pairs. In other words, some twisted pairs have characteristics different from those described in this chapter. This fact is especially important when troubleshooting "real world" issues and problems.

Cable Bundles

Twisted pairs are usually deployed in *cable bundles* containing multiple twisted pair lines in a cable sheath. The number of pairs contained in the cable varies from as little as one pair to as many as 4,200 pairs. Common cables used in the subscriber loops contain 25 to 100 twisted pairs. A color coding scheme is used to identify specific twisted pairs within a bundle.[4,5]

Each wire in a twisted pair has an insulator coating. Older wires used a paper-based insulation—these are called *pulp-insulated twisted pairs.* Newer twisted pairs use various types of plastics as insulators. The plastics used include polyethylene, polypropylene, and PVC. These types of twisted pairs are generically referred to as *plastic-insulated cables* or PIC. Some cables might also be packed with gel to increase resistance to moisture.

Cable bundles normally have an outer metallic shield. This shield is grounded at the cable ends to help reduce interference from outside sources. The type of shield varies from cable to cable. Cables may also contain support structures such as coated steel armor and galvanized steel supporting strands.

Individual twisted pairs encountered in the outside plant normally do not have dedicated shields. The twisting along the length of a specific pair is constant, but the amount of twist varies between pairs from 2 to 6 inches in length. The varying twist length of neighboring twisted pairs helps further reduce crosstalk between the pairs.

Gauges and Line Characteristics

Twisted pairs are most often characterized by their gauge using the American Wire Gauge (AWG) designations; the gauge is indicative of the diameter of the copper wires making up the twisted pair. Typical twisted pair gauges include #19, #22, #24, and #26 with diameters ranging from 0.03589 inches (0.912 mm) to 0.01594 inches (0.404 mm) with the smaller gauge number corresponding to the larger diameter. Gauges #24 and #26 are most common in DSL applications. Table 2.1 shows the preceding gauges along with the diameters and DC resistance per 1,000 feet of each.

Table 2.1 Common Twisted Pair Wire Gauges and Characteristics

AWG	Diameter (in)	DC Resistance 20° (ohms/kft)
19	0.03589	16.9
22	0.02535	33.8
24	0.02010	53.4
26	0.01594	85.8

Twisted pairs are also defined by categories specifying electrical tolerances of the pair. Installations for high-speed data are typically required to be category 3 or higher cables. Higher category numbers correspond to higher performance twisted pairs in terms of the loss and crosstalk properties of each. Table 2.2 lists the different cable categories, the approximate bit rates supported, and the target applications of each category.

Table 2.2 Twisted Pair Categories and Characteristics

Category	Bit Rates	Applications
1	Unspecified	
2	1 Mbps	Some low-speed data circuits (DDS)
3	16 Mbps	16 Mbps Used for 10BaseT and 4 Mbps Token Ring
4	20 Mbps	10BaseT and 16 Mbps Token Ring
5	100 Mbps	10/100BaseT, other high-speed copper technologies

For DSL technologies, category 3 or higher twisted pair is preferred. As discussed in Chapter 1, reuse of the existing copper plant is a major driver for ADSL and VDSL. Existing cable is normally as good as or better than category 3 specifications.[6] Detailed specifications for typical wiring categories can be found in *Telecommunications Wiring for Commercial Buildings, A Practical Guide.*[7]

Electrical Characteristics of Twisted Pairs

For any transmission line, the basic parameters of interest are resistance, inductance, capacitance, and admittance. These parameters are also called the *RLCG parameters.* In the voice band, the RLCG parameters are close to constant for #24 and #26 gauge twisted pairs. At frequencies above the voice band, resistance, inductance, and admittance vary with frequency. These three parameters also vary with gauge. Capacitance tends to be constant over both frequency and gauge.

Depending on the situation, one of two assumptions can be made when discussing line parameters. The first assumption ignores the effects of admittance and assumes that inductance and capacitance are constant with frequency. Also, resistance is assumed to be proportional to $f^{1/2}$ where f is frequency in Hz. This approximation is used when deriving near-end and far-end crosstalk models. Later in this chapter, and in Chapter 3, you will see that this approximation makes for a simple and accurate model for crosstalk on DSLs.

The second approximation is used for actual numeric calculations of twisted pair characteristics. This approximation will be used when finding the transfer function and insertion losses of a twisted pair as well as the characteristic impedance. This approximation was numerically determined for #24 and #26 twisted pairs.[8]

Analytic RLCG Models

Analytic RLCG models have been accepted and used in literature for many years. The parameters are given in Eqtn. 2.1–2.4 and the constants in Table 2.3. The resulting RLCG

parameters are with respect to 1,000 feet, or 1 kft, a common measure of distance for twisted pairs.

Eqtn. 2.1

$$R = R_0 f^{\frac{1}{2}}$$

Eqtn. 2.2

$$L = L_0$$

Eqtn. 2.3

$$C = C_0$$

Eqtn. 2.4

$$G = 0$$

Table 2.3 Parameters for the Analytic Twisted Pair Model

Parameter	#24 Gauge	#26 Gauge
R^0 (ohms/kft)	0.15	0.195
L^0 (mH/kft)	0.188	0.205
C^0 (nF/kft)	15.7	15.7

Numeric RLCG Models

The numeric RLCG models for #24 and #26 gauge twisted pairs are valid over the frequency range from D.C. to 10 MHz. The model parameters were found by curve fitting to measured cable.[8] Eqtns. 2.5–2.8 are the general equations for all four parameters; the constant values applicable to #24 and #26 gauge pairs are given in Table 2.4. The resulting RLCG parameters in Table 2.4 are with respect to 1 km.

Eqtn. 2.5

$$R(f) = \sqrt[4]{r_{oc}^4 + a_c f^2}$$

Eqtn. 2.6

$$L(f) = \frac{l_o + l_\infty \left(\frac{f}{f_m}\right)^b}{1 + l_\infty \left(\frac{f}{f_m}\right)^b}$$

Eqtn. 2.7

$$G(f) = g_o f^{g_e}$$

Eqtn. 2.8

$$C(f) = c_\infty$$

Table 2.4 Parameters for the Numeric Twisted Pair Model

Parameter	#24 Gauge	#26 Gauge
r_{oc} (ohms/km)	174.55888	286.17578
a_c (ohms^4/km^{4}Hz^2)	0.053073481	0.14769620
l_o (H/km)	$6.1729593*10_{-6}$	$675.36888*10_{-6}$
l_∞ (H/km)	$478.97099*10_{-6}$	$488.95186*10_{-6}$
f_m (Hz)	553760.63	806338.63
b	1.1529766	0.92930728
g_o (Siemen/Hz*km)	$0.23487476*10_{-12}$	$4.3*10_{-8}$
g_e	1.38	0.70
c_∞ (nF/km)	$50*10_{-9}$	$49*10_{-9}$

The RLCG parameters for both the analytic models and the numeric models are graphed in Figures 2.4–2.7.

In Figure 2.7, only the numeric model is shown because the analytic admittance is assumed to be zero. As frequency increases past 100 kHz, note the closeness of the numeric resistance model to the analytic model dependent on $f^{1/2}$. This is largely due to skin effect of the copper wires at high frequency, which for a cylindrical conductor, is known to exhibit such a dependence on frequency.[9]

Figure 2.4 Resistance models for #24 and #26 gauge twisted pairs.

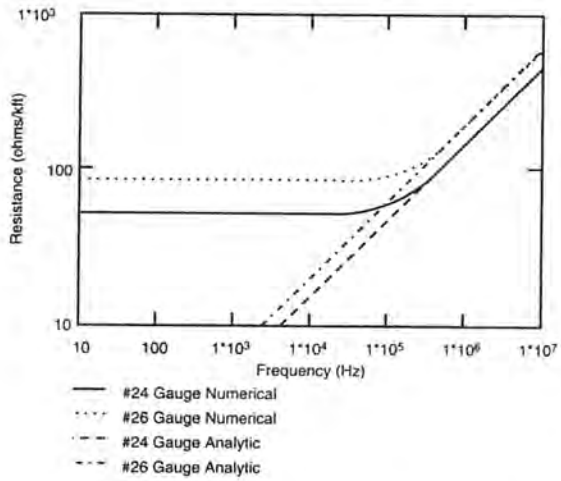

Figure 2.5 Inductance models for #24 and #26 gauge twisted pairs.

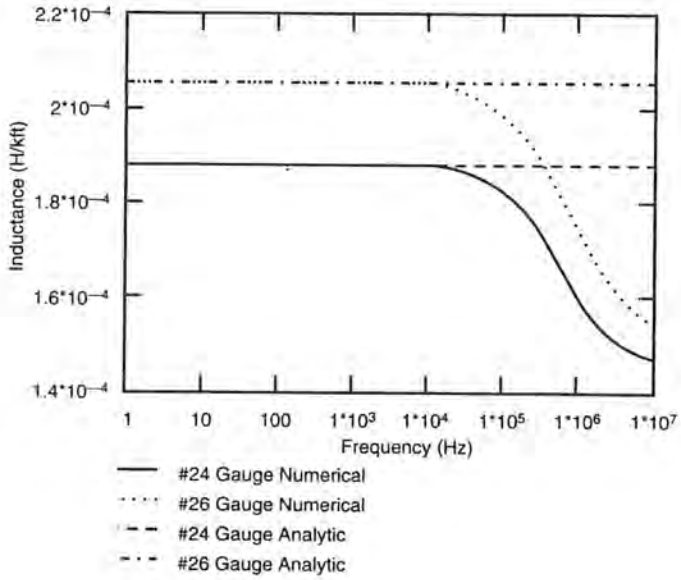

Figure 2.6 Capacitance models for #24 and #26 gauge twisted pairs.

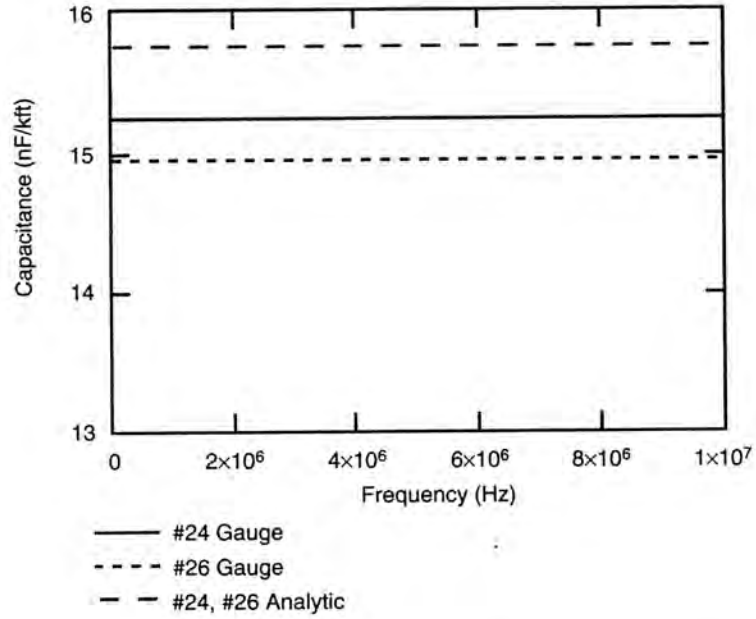

Figure 2.7 Admittance models for #24 and #26 gauge twisted pairs.

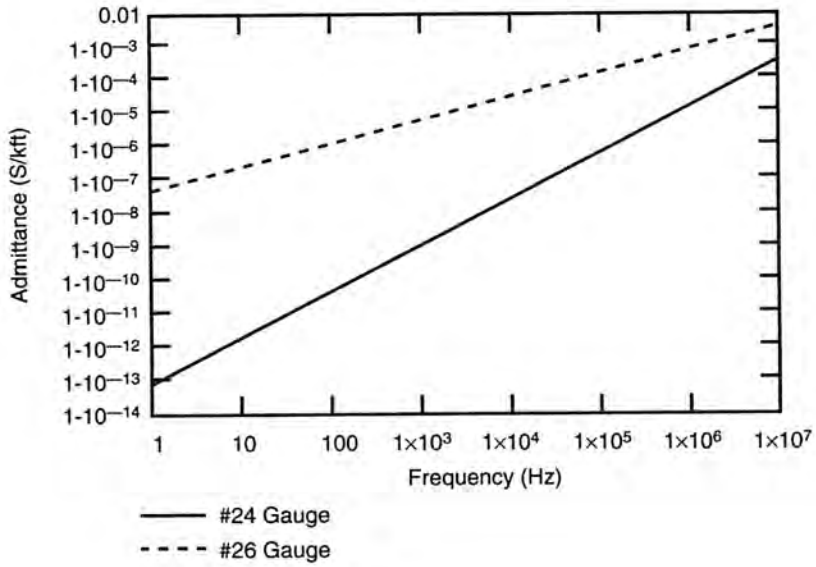

Note that in Figure 2.5, the inductance for both #24 and #26 wires varies by about 20 percent in the region above 100 kHz. Assuming that inductance is constant in the analytic model is certainly incorrect, but the choice to trade off error for simplicity is reasonable. The following sections show that the assumption of zero admittance in the analytic model,

although not accurate according to the numeric model, is reasonable at frequencies above 100 kHz. At these frequencies, the admittance needs to compare to $j\omega C$ where j is the imaginary unit vector and ω is the radian frequency variable ($2\pi f$).

Transmission Line Parameters

The RLCG parameters of a transmission line define two common parameters useful in characterizing voltage and current along the line. These are the *propagation constant* (γ) and *characteristic impedance* ($Z_o(f)$). With these parameters, the voltage at some point z on a transmission line of length l with a load $Z_L(f)$ is given by Eqtn. 2.9. (The assumption is made that z=0 at the source voltage and increases positively until z=l at the termination load.)

Eqtn. 2.9

$$V(f,z) = V_0(f)\left(e^{-\gamma(f)z} + \Gamma_L(f)e^{\gamma(f)z}\right)$$

In Eqtn. 2.9, $V_0(f)$ is the source voltage, and $\Gamma_L(f)$ is the reflection coefficient of the line given by Eqtn. 2.10.

Eqtn. 2.10

$$\Gamma_L(f) = \frac{Z_L(f) - Z_o(f)}{Z_L(f) + Z_o(f)}$$

Note that for a known source voltage, termination impedance, and position, the voltage on the line is a function of only propagation constant and characteristic impedance.

Maximum power is delivered to a termination load on a transmission line when the characteristic impedance of the line and the load impedance are equal. This is desired in the case of a DSL because the termination load is the front of the channel receiver; delivering the maximum possible power to the receiver is necessary to optimize performance.

Under this condition, the expression for voltage on the line is simplified as the reflection coefficient in Eqtn. 2.10 becomes zero and the second exponential term in Eqtn. 2.9 drops out of the equation. When the load impedance is the same as the transmission line characteristic impedance, we say that the transmission line is matched. Hereafter, unless otherwise stated, it will be assumed that the twisted pair transmission lines under investigation are always matched, and the expression for voltage at some point z along the line is given by Eqtn. 2.11.

Eqtn. 2.11

$$V(f,z) = V_0(f)e^{-\gamma(f)z}$$

Note

Later in the chapter, you will see that this is a reasonable assumption over the frequencies of interest. Even when the characteristic impedance and load impedance are slightly different, the second term in Eqtn. 2.9 is negligible compared to the first term.

Knowledge of the propagation constant of a twisted pair is key to understanding the effect that a twisted pair has on signals. The propagation constant can further be broken down resulting in Eqtn. 2.12.

Eqtn. 2.12

$$\gamma = \alpha + j\beta = \sqrt{(R + j\omega L)(G + j\omega C)}$$

In Eqtn. 2.12, dependence of all quantities on frequency has been dropped only for notational simplicity. The terms α and β are the real and imaginary parts of the propagation constant, respectively. Substituting these into Eqtn. 2.11 yields the following in Eqtn. 2.13.

Eqtn. 2.13

$$V(f, z) = V_0(f)e^{-(\alpha + j\beta)z} = V_0(f)e^{-\alpha z}e^{-j\beta z}$$

The transfer function between the source voltage and some point z on the twisted pair is given by H*(f)* as defined in Eqtn. 2.14.

Eqtn. 2.14

$$H(f) = \frac{V(f, z)}{V_0(f)} = e^{-\alpha z}e^{-j\beta z} = \left|e^{-\alpha z}\right| \angle \beta z$$

The last part of Eqtn. 2.14 is valid because the exponential involving β has unity magnitude and only contributes to phase shifting the voltage. On the other hand, the term involving α affects only the magnitude of the voltage. Thus, from Eqtn. 2.14, you can see that the real part of the propagation constant determines by how much the voltage will change over a length of twisted pair and the imaginary part how much the phase will change.

Propagation Constant

Later in this chapter, you will see that crosstalk between two twisted pairs becomes worse as frequency increases. Also, propagation loss of a twisted pair increases as frequency increases. In this section, an expression will be derived for the propagation constant of a twisted pair at these high frequencies. Two basic assumptions will aid in characterizing the propagation constant: first, $j\omega C >> G$, or equivalently, $G=0$; and second, $j\omega L > R$. Note that the second condition is less strict than the first. These approximations are valid for #24 and #26 twisted pairs at frequencies above 100 kHz.

Using the approximations set forth, the propagation constant of a twisted pair can be written as shown in Eqtn. 2.15.

Eqtn. 2.15

$$\gamma = \sqrt{(R + j\omega L)(G + j\omega C)} \approx \sqrt{j\omega C}\sqrt{(R + j\omega L)}$$

In Eqtn. 2.15, the first assumption was used to eliminate the admittance term. Factoring the equation yields the result in Eqtn. 2.16.

Eqtn. 2.16

$$\gamma = \sqrt{j\omega C}\sqrt{(R + j\omega L)} = j\omega\sqrt{LC}\sqrt{\left(\frac{R}{j\omega L} + 1\right)}$$

$$\approx j\omega\sqrt{LC}\left(\frac{R}{j\omega L} + 1\right)$$

$$= R\sqrt{\frac{C}{L}} + j\omega\sqrt{LC}$$

In this equation, the final approximation is based on the second assumption $(x+1)^{1/2}=(x+1)$ for small x.

In terms of α and β, the real and imaginary parts of the propagation constant, the result from Eqtn. 2.16 yields these terms as shown in Eqtn. 2.17 and Eqtn. 2.18.

Eqtn. 2.17

$$\alpha = R\sqrt{\frac{C}{L}} = \alpha_o f^{1/2}$$

Eqtn. 2.18

$$\beta = \omega\sqrt{LC} = \beta_o f$$

In Eqtns. 2.17 and 2.18, α_o and β_o are introduced as constants. The final result for α comes from using the analytic approximation for the RLCG parameters. α is proportional to the square root of frequency because capacitance and inductance are both assumed constant, and R is assumed proportional to $f^{1/2}$ due to skin effect. The final result for β shows that β is proportional to frequency due to the $\omega(2\pi f)$ term.

Figures 2.8 and 2.9 show the real and imaginary part of propagation constant of a #24 and #26 gauge twisted pair. At frequencies above 100 kHz, both plots are consistent with Eqtn. 2.17 and Eqtn. 2.18. Note that the real part of the propagation constant, which was already established to have a direct effect on loss, increases with frequency and is higher for #26 wire than for #24.

Figure 2.8 Real part of the propagation constant for #24 and #26 twisted pairs.

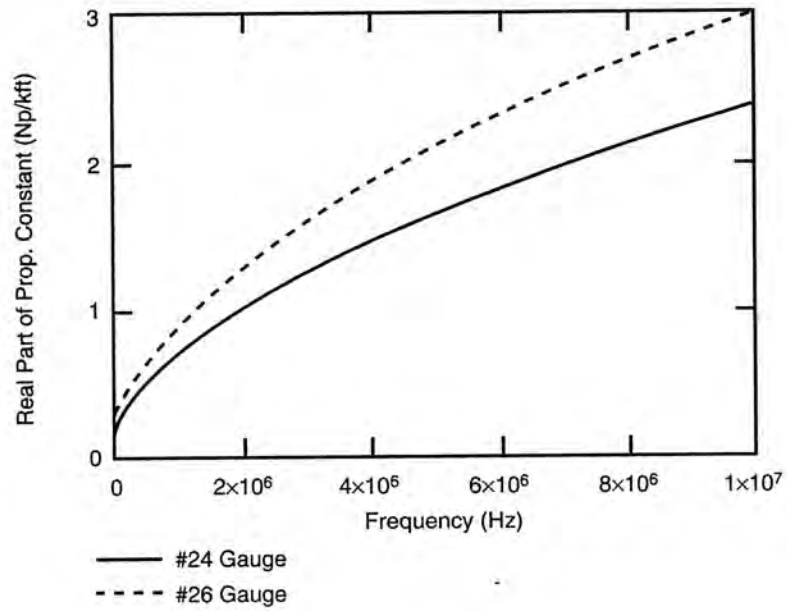

Figure 2.9 Imaginary part of the propagation constant for #24 and #26 twisted pairs.

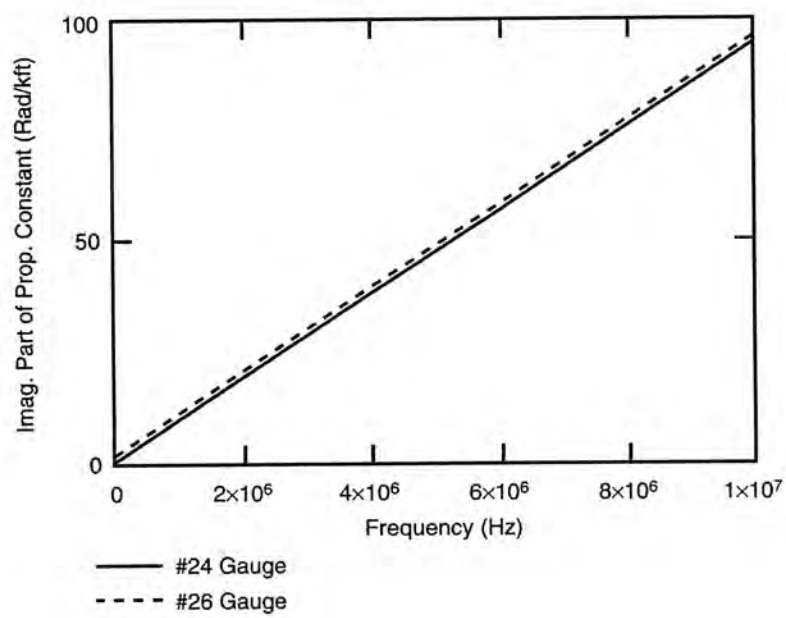

In Eqtn. 2.11, an expression for the transfer function of a twisted pair was given in terms of the propagation constant. Figures 2.10–2.13 show the transfer function magnitude and phase for both #24 and #26 gauge cables.

Figure 2.10 Transfer function magnitude for various #24 gauge twisted pairs.

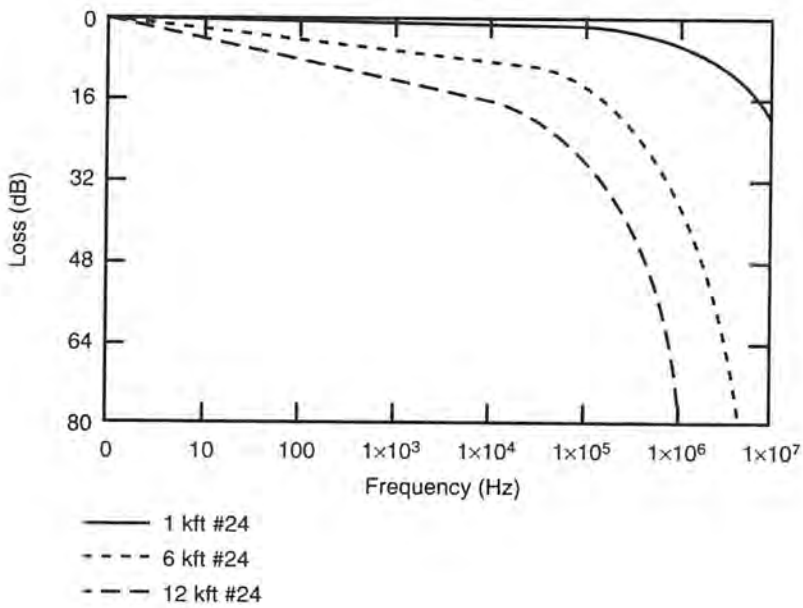

Figure 2.11 Transfer function phase for various #24 gauge twisted pairs.

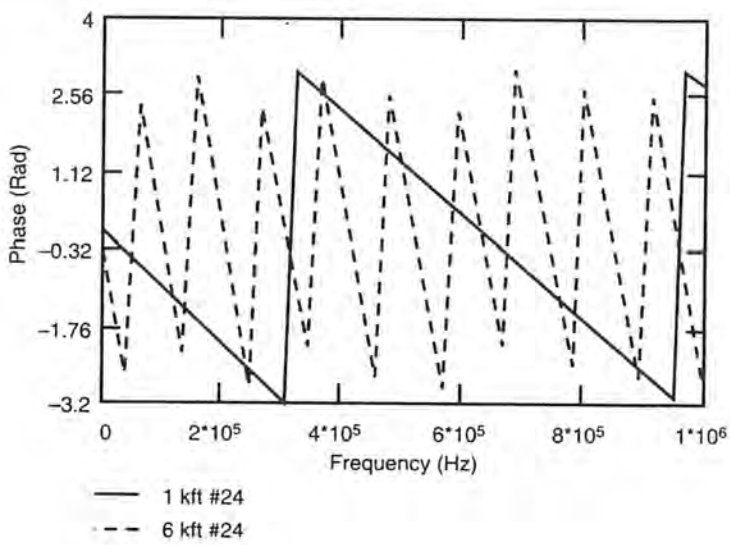

Figure 2.12 Transfer function magnitude for various #26 gauge twisted pairs.

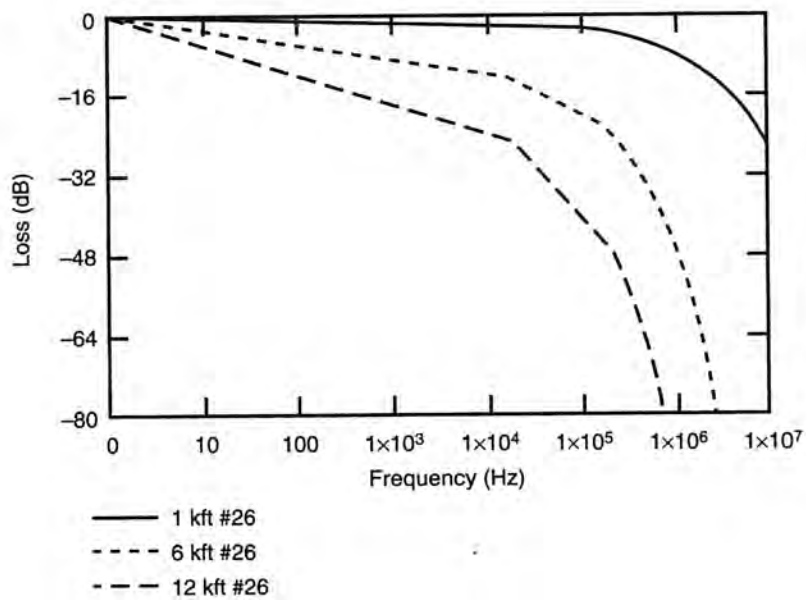

Figure 2.13 Transfer function phase for various #26 gauge twisted pairs.

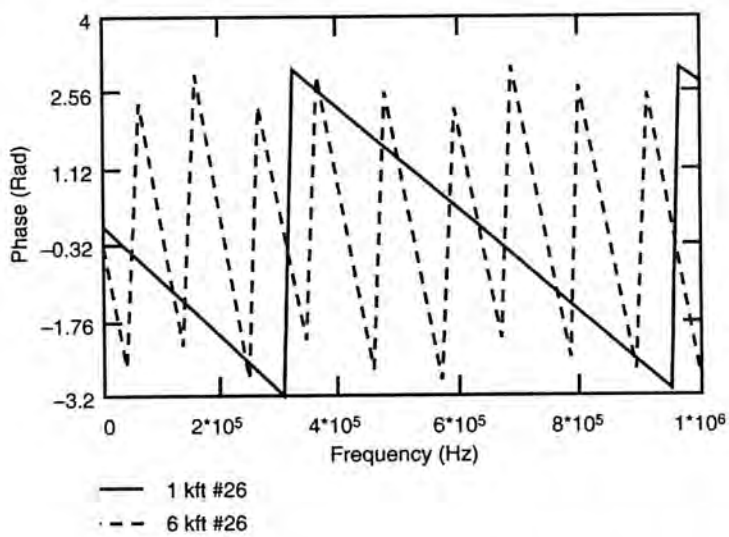

Figures 2.10–2.13 show various cable line lengths. Note that the loss of the cable increases as line length increases and that it also increases with higher frequencies. Intuitively, these behaviors make sense. Longer lines have more loss than shorter lines, and higher frequencies are attenuated more by a pair of wires. Note also in Figures 2.11 and 2.13, where

phase is shown for only a 1,000 ft. and 6,000 ft. loop, the greater phase slope of the longer line. This trend continues as the loop length is further increased.

Characteristic Impedance

The preceding analysis assumes that the twisted pair was terminated in its characteristic impedance, Z_o. The characteristic impedance of any transmission line is defined in terms of the transmission line RLCG parameters and is given by Eqtn. 2.19.

Eqtn. 2.19

$$Z_o = \sqrt{\frac{R + j\omega L}{G + j\omega C}}$$

Note that the characteristic impedance does not depend on the length of the twisted pair. Figure 2.14 shows the characteristic impedance magnitude for a #24 and #26 gauge twisted pair. The respective phases are shown in Figure 2.15.

Figure 2.14 The characteristic impedance magnitude for #24 and #26 gauge twisted pairs.

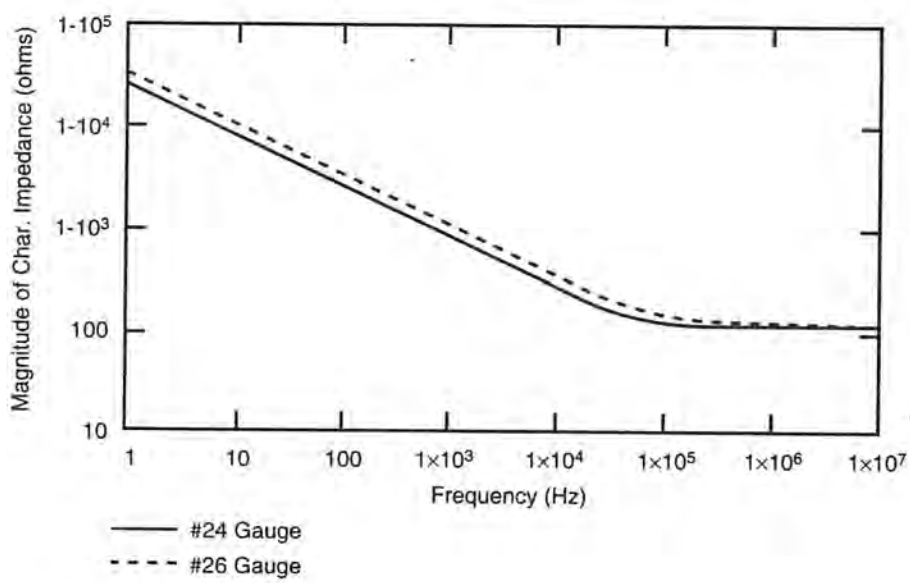

Figure 2.15 The characteristic impedance phase for #24 and #26 gauge twisted pairs.

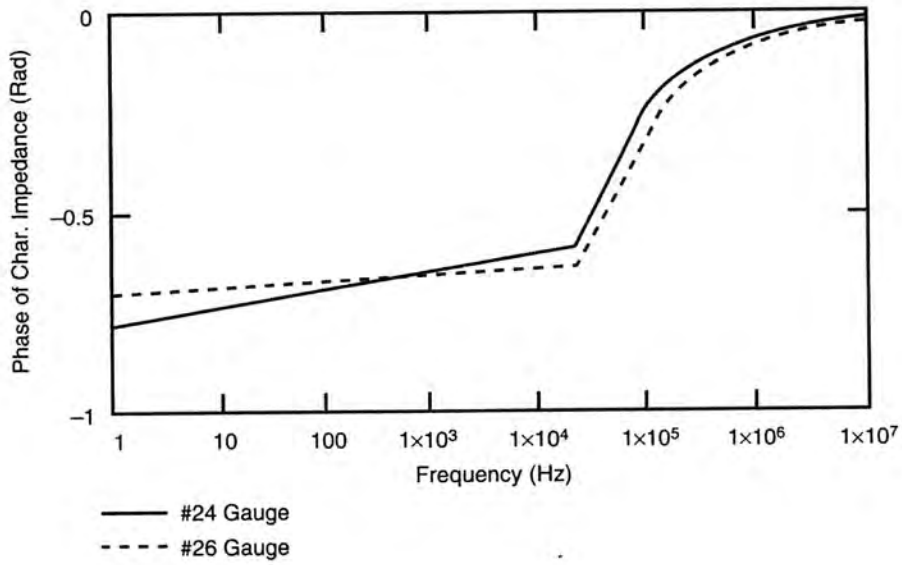

Although the characteristic impedances vary with frequency, a 100 ohm purely resistive impedance is often assumed. (In some cases, especially in Europe, 135 ohms is assumed.) The error of this assumption is negligible in many cases because the reflection coefficient further reduces the effect of small mismatches as indicated in Eqtn. 2.10.

Note
In an actual implementation, it might be reasonable to design a terminating impedance to better favor either a #24 gauge or a #26 gauge twisted pair.

Crosstalk

This section is dedicated to defining crosstalk on twisted pairs and creating a mathematical model for it.

Crosstalk refers simply to disturbance on one twisted pair due to signals on another twisted pair. Usually, the twisted pairs are in the same cable or *binder group*. This text refers to the twisted pair being disturbed as the *disturbed pair* and to the twisted pair causing the disturbance as the *disturber* or *disturbing pair*. In cables with many pairs, it is common to have multiple disturbers. This also is addressed.

Note
If you have ever been on the telephone and faintly heard someone else's conversation on your line, you have experienced crosstalk firsthand.

Crosstalk Types

Crosstalk is generally characterized as either *near-end crosstalk* (NEXT) or *far-end crosstalk* (FEXT). NEXT is characterized by the disturbing pair's source being local to the disturbed pair's receiver. In this case, the disturbing signal starts down the disturbing pair, couples into the disturbed pair, and then propagates back to the disturbed pair's receiver. FEXT is characterized by the disturbing pair's source being distant from disturbed pair's receiver. In this case, the disturbing signal propagates down the disturbing pair, crosstalks into the disturbed pair, and propagates the rest of the distance along the disturbed pair into the disturbed pair's receiver. Figure 2.16 illustrates NEXT and FEXT.

Figure 2.16 Near-end crosstalk (NEXT) and far-end crosstalk (FEXT).

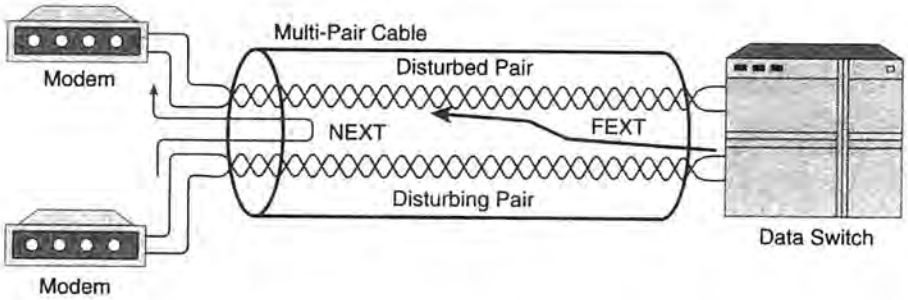

As shown in Figure 2.16, disturbing signals due to NEXT usually have to travel over less length of the twisted pair before entering a disturbed pair's receiver as compared to FEXT. For this reason, NEXT from a particular disturber is usually more damaging than FEXT. Note that either NEXT or FEXT or both might exist for a disturbed pair, and furthermore, more than one NEXT or FEXT disturber might exist.

NEXT or FEXT from a disturber whose line code is the same as that of the disturbed pair creates self-NEXT or self-FEXT, respectively. This type of disturbance is often present as similar services are often deployed from a central location. Later, you will see that ADSL and VDSL systems avoid self-NEXT, although self-FEXT is inevitable. Other common disturbers (both NEXT and FEXT) for ADSL systems include ISDN, HDSL, and T1 systems.

Unbalance Models

The following sections break the derivation of a mathematical model for NEXT and FEXT into two parts. The first part derives a mathematical model for crosstalk between two very short twisted pair sections. The second part takes that generic result and applies it to integrals specific to either NEXT or FEXT. It is interesting to note that the models stem from work done just after the turn of the 20th century by George Campbell[10,11] and also from work done in the 1960s by Cravis and Carter.[12] Both works are considered classics in crosstalk modeling and referenced in many journal articles on crosstalk.

For two twisted pairs located near one another, capacitive and inductive components exist between each of the four wires and also between each wire and ground. Analysis shows that crosstalk occurs when an unbalance exists among the capacitive and inductive components. The following assumptions are made while deriving the crosstalk models:

- The total crosstalk is a sum of the crosstalk due to capacitive and inductive crosstalk.
- Secondary effects (or crosstalk due to crosstalk) can be ignored.

Capacitive Unbalance Model

The capacitive unbalance analysis follows the derivation done by Miller.[13] Figure 2.17 shows a model of a portion of length of two twisted pair circuits in reasonably close proximity (for example, in a cable or binder group).

The four wires that compose the model each have some self-impedance, an admittance to one another, and a capacitance to ground. Furthermore, it is assumed that resistance between wires is negligible due to the plastic insulation. Thus, all admittances are purely capacitive (consistent with the analytic RLCG model).

Figure 2.17 Capacitive model of a short section of two twisted pairs.

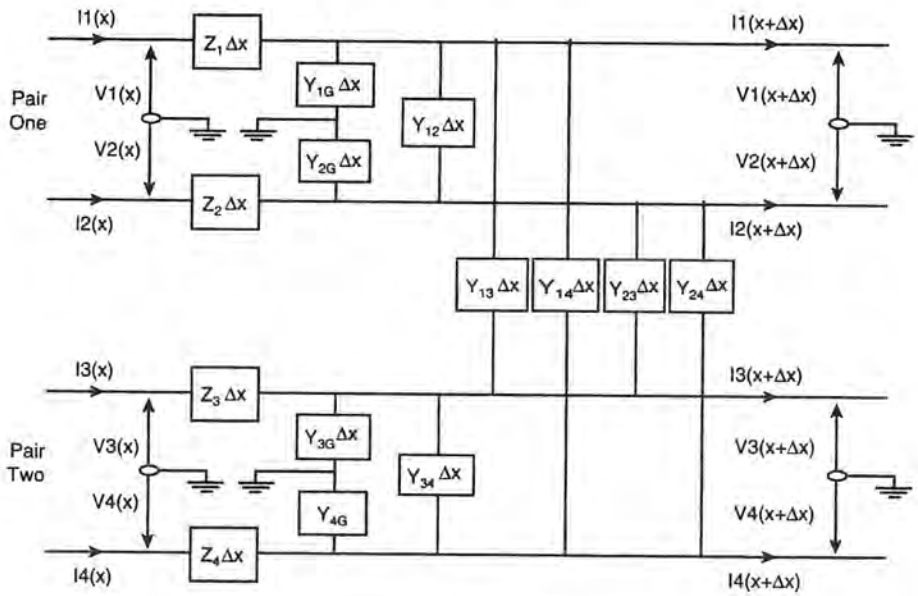

For convenience, you can write all capacitances as admittances ($Y=j\omega C$). You can then write equations for the voltage and current on each of the four wires at the location $x+\Delta x$ using basic circuit theory as in Eqtns. 2.20–2.27.

Eqtn. 2.20

$$V_1(x+\Delta x)=V_1(x)-I_1(x)Z_1\Delta x$$

Eqtn. 2.21

$$V_2(x+\Delta x)=V_2(x)-I_2(x)Z_2\Delta x$$

Eqtn. 2.22

$$V_3(x+\Delta x)=V_3(x)-I_3(x)Z_3\Delta x$$

Eqtn. 2.23

$$V_4(x+\Delta x)=V_4(x)-I_4(x)Z_4\Delta x$$

Eqtn. 2.24

$$I_1(x+\Delta x) = I_1(x) - V_1(x)Y_{1G}\Delta x - [V_1(x) - V_2(x)]Y_{12}\Delta x$$
$$-[V_1(x) - V_3(x)]Y_{13}\Delta x - [V_1(x) - V_4(x)]Y_{14}\Delta x$$

Eqtn. 2.25

$$I_2(x+\Delta x) = I_2(x) - V_2(x)Y_{2G}\Delta x - [V_2(x) - V_1(x)]Y_{12}\Delta x$$
$$-[V_2(x) - V_3(x)]Y_{23}\Delta x - [V_2(x) - V_4(x)]Y_{24}\Delta x$$

Eqtn. 2.26

$$I_3(x+\Delta x) = I_3(x) - V_3(x)Y_{3G}\Delta x - [V_3(x) - V_1(x)]Y_{13}\Delta x$$
$$-[V_3(x) - V_2(x)]Y_{23}\Delta x - [V_3(x) - V_4(x)]Y_{34}\Delta x$$

Eqtn. 2.27

$$I_4(x+\Delta x) = I_4(x) - V_4(x)Y_{4G}\Delta x - [V_4(x) - V_1(x)]Y_{14}\Delta x$$
$$-[V_4(x) - V_2(x)]Y_{24}\Delta x - [V_4(x) - V_3(x)]Y_{34}\Delta x$$

In Eqtns. 2.20–2.27, you can divide through by Δx, and as Δx approaches zero, the definition of the derivative can be invoked. The eight simultaneous equations can then be rearranged and placed in matrix form as shown in Eqtn. 2.28.

Eqtn. 2.28

$$\frac{d}{dx}\begin{bmatrix} V_1 \\ V_2 \\ V_3 \\ V_4 \\ I_1 \\ I_2 \\ I_3 \\ I_4 \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 0 & -Z_1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -Z_2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -Z_3 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -Z_4 \\ -\begin{pmatrix} Y_{1G}+Y_{12} \\ +Y_{13}+Y_{14} \end{pmatrix} & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ Y_{12} & -\begin{pmatrix} Y_{2G}+Y_{12} \\ +Y_{23}+Y_{24} \end{pmatrix} & 0 & 0 & 0 & 0 & 0 & 0 \\ Y_{13} & 0 & -\begin{pmatrix} Y_{3G}+Y_{13} \\ +Y_{23}+Y_{34} \end{pmatrix} & 0 & 0 & 0 & 0 & 0 \\ Y_{14} & 0 & 0 & -\begin{pmatrix} Y_{4G}+Y_{14} \\ +Y_{24}+Y_{34} \end{pmatrix} & 0 & 0 & 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} V_1 \\ V_2 \\ V_3 \\ V_4 \\ I_1 \\ I_2 \\ I_3 \\ I_4 \end{bmatrix} = A \cdot S$$

In Eqtn. 2.28, for convenience of notation only, the dependence on x from all voltages and currents has been dropped, and S and A have been introduced to denote the 8×8 impedance matrix and the 8×1 voltage and current vector, respectively.

At this point, it is helpful to transform the voltages and currents on the four lines to metallic and longitudinal voltages and currents. A simple set of linear transformations shown in Eqtns. 2.29–2.36 accomplishes this.

Eqtn. 2.29

$$V_{1M} = V_1 - V_2$$

Eqtn. 2.30

$$V_{2M} = V_3 - V_4$$

Eqtn. 2.31

$$V_{1L} = \frac{V_1 + V_2}{2}$$

Eqtn. 2.32

$$V_{2L} = \frac{V_3 + V_4}{2}$$

Eqtn. 2.33

$$I_{1M} = \frac{I_1 - I_2}{2}$$

Eqtn. 2.34

$$I_{2M} = \frac{I_3 - I_4}{2}$$

Eqtn. 2.35

$$I_{1L} = I_1 + I_2$$

Eqtn. 2.36

$$I_{2L} = I_3 + I_4$$

In these equations, the first subscript $_{(1 \text{ or } 2)}$ identifies the twisted pair, and the second subscript $_{(M \text{ or } L)}$ indicates metallic or longitudinal signal type, respectively. Thus for example, V_{1M} would represent the metallic voltage on twisted pair 1.

To transform Eqtn. 2.28 to be in terms of metallic and longitudinal signals, Eqtns. 2.30–2.36 must be solved for the original voltages and currents. By letting S be a column vector of the metallic and longitudinal voltages and currents, it can easily be shown that Eqtn. 2.37 is true.

Eqtn. 2.37

$$S = T \cdot \begin{bmatrix} V_{1M} \\ V_{1L} \\ V_{2M} \\ V_{2L} \\ I_{1M} \\ I_{1L} \\ I_{2M} \\ I_{2L} \end{bmatrix} = T \cdot S'$$

In Eqtn. 2.37, S is the column vector of voltage and currents, and T is a transformation matrix given by Eqtn. 2.38.

Eqtn. 2.38

$$T = \begin{bmatrix} \frac{1}{2} & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ \frac{-1}{2} & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{2} & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & \frac{-1}{2} & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & \frac{1}{2} & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & \frac{1}{2} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & \frac{1}{2} \\ 0 & 0 & 0 & 0 & 0 & 0 & -1 & \frac{1}{2} \end{bmatrix}$$

You can then transform the system of equations given in Eqtn. 2.28 as in Eqtn. 2.39.

Eqtn. 2.39

$$\frac{d}{dx} S' = T^{-1} ATS'$$

The result of this transformation shows that the derivatives with respect to x of the metallic and longitudinal voltages depend only on the metallic and longitudinal currents, and the derivative with respect to x of the metallic and longitudinal currents depends only on the metallic and longitudinal voltages. Concentration will be placed only on the latter, simplifying the system of equations to four equations and four unknowns. Thus, the simplified set of equations in terms of metallic and longitudinal signals is shown in Eqtn. 2.40.

Eqtn. 2.40

$$\frac{d}{dx}\begin{bmatrix} I_{1M} \\ I_{1L} \\ I_{2M} \\ I_{2L} \end{bmatrix} = \frac{-j \cdot \omega}{4} \begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{14} \\ a_{21} & a_{22} & a_{23} & a_{24} \\ a_{31} & a_{32} & a_{33} & a_{34} \\ a_{41} & a_{42} & a_{43} & a_{4} \end{bmatrix} \cdot \begin{bmatrix} V_{1M} \\ V_{1L} \\ V_{2M} \\ V_{2L} \end{bmatrix}$$

The elements of the four-by-four matrix are given by Eqtns. 2.41–2.50.

Eqtn. 2.41

$$a_{11} = C_{1G} + C_{2G} + 4C_{12} + C_{13} + C_{14} + C_{23} + C_{24}$$

Eqtn. 2.42

$$a_{21} = a_{12} = 2C_{1G} - 2C_{2G} + 2C_{13} + 2C_{14} - 2C_{23} - 2C_{24}$$

Eqtn. 2.43

$$a_{31} = a_{31} = -C_{13} + C_{14} + C_{23} - C_{24}$$

Eqtn. 2.44

$$a_{41} = a_{14} = -2C_{13} - 2C_{14} + 2C_{23} + 2C_{24}$$

Eqtn. 2.45

$$a_{22} = 4C_{1G} + 4C_{2G} + 4C_{13} + 4C_{14} + 4C_{23} + 4C_{24}$$

Eqtn. 2.46

$$a_{23} = a_{32} = -2C_{13} + 2C_{14} - 2C_{23} + 2C_{24}$$

Eqtn. 2.47

$$a_{24} = a_{42} = -4C_{13} - 4C_{14} - 4C_{23} - 4C_{24}$$

Eqtn. 2.48

$$a_{33} = C_{3G} + C_{4G} + C_{13} + C_{14} + C_{23} + C_{24} + 4 \cdot C_{34}$$

Eqtn. 2.49

$$a_{34} = a_{43} = 2C_{3G} - 2C_{4G} + 2C_{13} - 2C_{14} + 2C_{23} - 2C_{24}$$

Eqtn. 2.50

$$a_{44} = 4C_{3G} + 4C_{4G} + 4C_{13} + 4C_{14} + 4C_{23} + 4C_{24}$$

In these equations, capacitance has been used instead of admittance to further stress that the impedances between all wires is capacitive. For crosstalk between twisted pairs, you should be most concerned with the coupling from the metallic signal of the disturbing pair into the metallic signal of the disturbed pair. The parameter a_{31} from Eqtn. 2.43 defines this coupling, metallic voltage in the disturbing pair (V_{1M}) to metallic current in the disturbed pair (I_{2M}). Parameter a_{31} is referred to as the *capacitance unbalance* of a twisted pair. For notational purposes, a_{31} will be renamed as C_{M1M2}. Thus, crosstalk resulting from capacitance unbalance is given by Eqtn. 2.51.

Eqtn. 2.51

$$I_{2M} = \frac{-d}{dx}\frac{j\omega}{4}C_{M1M2}V_{1M}$$

Inductive Unbalance Model

Crosstalk due to mutual inductance unbalance is a bit easier to derive than capacitive unbalance. Consider a small section of two twisted pairs as shown in Figure 2.18.

Figure 2.18 Mutual inductance model of a short section of two twisted pairs.

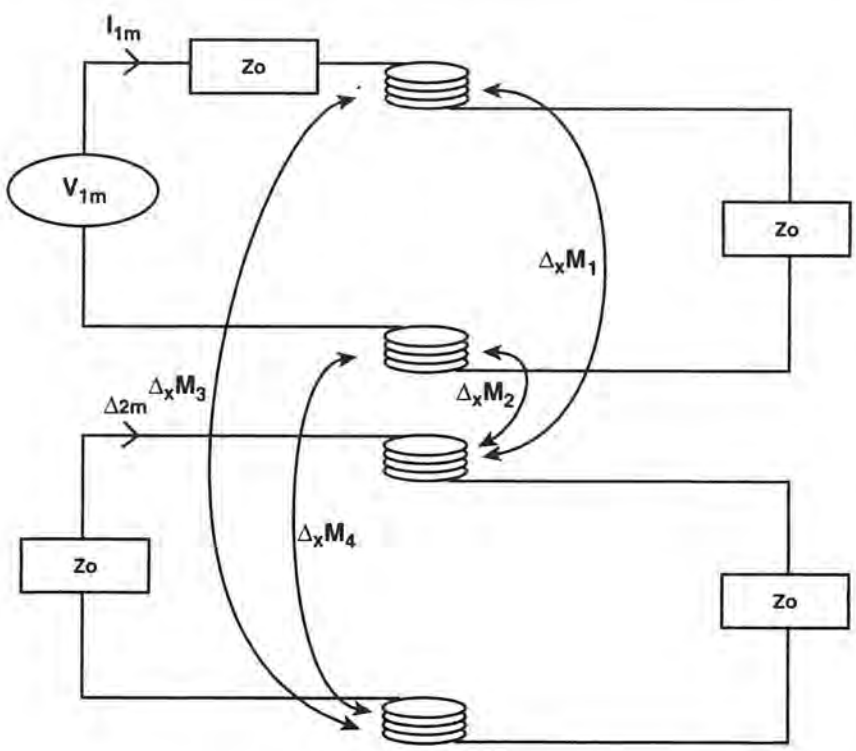

Mutual inductance exists pair-wise between each of the wires in the two twisted pairs. The effective termination at each end of the two pairs is the characteristic impedance. No

crosstalk occurs if the couplings are balanced with respect to one another. For example, if M_1 and M_3 produce exactly opposite currents in the disturbed pair and M_2 and M_4 do the same, then all the coupled currents will cancel one another, and no net crosstalk will be present on the disturbed pair.

To mathematically analyze Figure 2.18, it is helpful to draw the equivalent circuit shown in Figure 2.19.

Figure 2.19 Equivalent circuit for the mutual inductance unbalance of a short section of two twisted pairs.

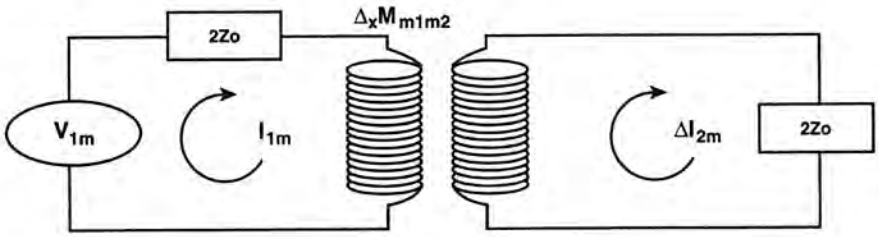

The mutual inductance denoted by M_{M1M2} is given by Eqtn. 2.52.

Eqtn. 2.52

$$M_{M1M2} = M_1 + M_2 + M_3 + M_4$$

Note that M_{M1M2} will be zero when the mutual inductances are balanced. In Figure 2.19, the current ΔI_{2M} is the current generated on the disturbed twisted pair on the small section shown.

The circuit in Figure 2.19 can be analyzed by writing the following loop equations (Eqtn. 2.53 and Eqtn. 2.54) for I_{1M} and ΔI_{2M}:

Eqtn. 2.53

$$2ZI_{1M} - M\Delta x I_{2M} = V_{1M}$$

Eqtn. 2.54

$$-M\Delta x I_{1M} + 2ZI_{2M} = 0$$

These equations can then be solved simultaneously, and the resulting ΔI_{2M} is in Eqtn. 2.55.

Eqtn. 2.55

$$I_{2M} = \frac{j\omega M \Delta x V_{1M}}{4Z^2 - M^2\omega^2\Delta x^2}$$

Because M is usually very small and Δx is a very small portion of line, the second term in the denominator is small with respect to the first term and can be dropped. Also, the definition of derivative can be invoked as Δx approaches zero, so Eqtn. 2.55 can be written as Eqtn. 2.56.

Eqtn. 2.56

$$\frac{dI_{2M}}{dx} \approx \frac{j\omega M V_{1M}}{4Z^2}$$

Thus, the result of inductive coupling here gives a similar result to the earlier result for capacitive coupling in that the current induced on a disturbed line due to a voltage on a disturbing line is proportional to the frequency. (This assumes that the characteristic impedance is constant with frequency as was shown earlier in this chapter.)

General Unbalance Expression

From the capacitive and inductive unbalance expressions, the total crosstalk on a small portion of twisted pair can be written as in Eqtn. 2.57.

Eqtn. 2.57

$$\frac{dI_{2M}}{dx} = \left(\frac{j\omega M}{4Z^2} + \frac{j\omega C_{M1M2}}{4}\right)V_{1M}$$

If the characteristic impedance of the line is a constant over frequency, then the current in Eqtn. 2.57 is proportional to frequency for a given voltage. With this in mind, all the constants of Eqtn. 2.57 can be grouped together to form a new unbalance constant, Q_{M1M2} such that Eqtn. 2.58 holds.

Eqtn. 2.58

$$\frac{dI_{2M}}{dx} = \left(\frac{j\omega M}{4Z^2} + \frac{j\omega C_{M1M2}}{4}\right)V_{1M} = j\omega Q_{M1M2}V_{1M}$$

Near-end Crosstalk (NEXT)

Now that a generic expression has been obtained for the crosstalk from a very short segment of a disturbing pair to a very short segment of a disturbed pair, attention can be

turned to the specific cases of NEXT and FEXT. Consider a signal V traveling along a disturbing pair. At some distance x1 along a disturbing pair, it couples into a disturbed pair due to unbalance and then travels back to the receiver of the disturbed pair as shown in Figure 2.20.

Figure 2.20 An example of near-end crosstalk due to only one unbalance point (x1).

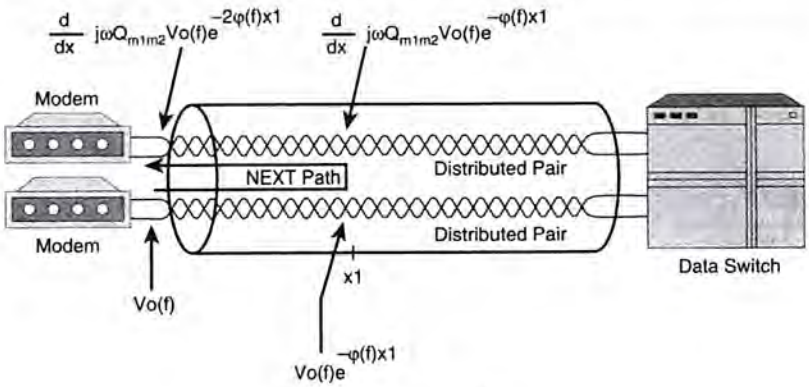

If the signal exciting the disturbing pair is V_o, then from Eqtn. 2.11 the signal on the disturbing pair at x1 is as shown in Eqtn. 2.59.

Eqtn. 2.59

$$V_1(f, x1) = V_o(f) e^{-\gamma(f)x1}$$

This signal couples due to unbalance onto the disturbed pair at x1. From the unbalance term of Eqtn. 2.58, the resulting current on the disturbed pair is shown by Eqtn. 2.60.

Eqtn. 2.60

$$I_2(f, x1) = \frac{d}{dx} j\omega Q_{M1M2} V_1(f, x1) = \frac{d}{dx} j\omega Q_{M1M2} V_o(f) e^{-\gamma(f)x1}$$

To reach the disturbed pair's receiver, the signal must propagate the distance x1 along the disturbed pair. This introduces further attenuation of $e_{ms\gamma x1}$. The current arriving at the disturbed pair's receiver can thus be written as in Eqtn. 2.61.

Eqtn. 2.61

$$I_2(f, 0) = I_2(f, x1) e^{-\gamma(f)x1} = \frac{d}{dx} j\omega Q_{M1M2} V_o(f) e^{-2\gamma(f)x1}$$

The expression in Eqtn. 2.61 gives the current that NEXT will produce if only one unbalance point existed along the twisted pairs. Usually, unbalance will exist at many points along the pairs, and furthermore, the unbalance coefficient, Q_{M1M2} might be different at

each point. To find the total current at the disturbed pair's receiver due to all the unbalance points along the pairs, you can let the unbalance coefficient be a function of position on the line, Q_{M1M2}, and integrate along the length of the line. The total current is then given by Eqtn. 2.62.

Eqtn. 2.62

$$I_2(f) = \int_0^L j\omega Q_{M1M2}(x) V_o(f) e^{-2\gamma(f)x]} dx$$

For performance analysis, the concern is more with the power due to NEXT than with the current. Given a current and a termination impedance, Z_{term}, power can be found by Eqtn. 2.63.

Eqtn. 2.63

$$P(f) = I(f) I^*(f) Z_{term} \; P(f) = I(f) I^*(f) Z_{term}$$

In Eqtn. 2.63, * is the complex conjugate operator. For the NEXT expression, still in integral form, explicit calculation of the current, and therefore power, would be impossible because knowledge of Q_{M1M2} at every point x on the pairs is not known. Assuming that the termination impedance is equal to the line's characteristic impedance, from Eqtn. 2.62, the NEXT power at the disturbed receiver can be written as in Eqtn. 2.64.

Eqtn. 2.64

$$\begin{aligned} E[P(f)] &= E\left[|Z_o| I_2(f) I_2(f)^*\right] \\ &= E\left[|Z_o| \int_0^L j\omega Q_{M1M2}(x) V_o(f) e^{-2\gamma(f)x} dx \int_0^L - j\omega Q_{M1M2}(y) V_o(f) e^{-2\gamma^*(f)y} dy\right] \end{aligned}$$

The operator E* is used here to denote the expected value. All terms that do not involve x or y can be pulled outside the integral so that Eqtn. 2.64 simplifies to Eqtn. 2.65.

Eqtn. 2.65

$$\begin{aligned} E[P(f)] &= E\left[|Z_o| I_2(f) I_2(f)^*\right] \\ &= E\left[|Z_o| \omega^2 V_o^2(f) \int_0^L \int_0^L Q_{M1M2}(x) Q_{M1M2}(y) e^{-2\gamma(f)x} e^{-2\gamma^*(f)y} dx dy\right] \end{aligned}$$

The double integration in Eqtn. 2.65 can be simplified by assuming that the unbalance at some point x1 on the twisted pair is independent of the unbalance at another point x2. Using this assumption leads to the expression in Eqtn. 2.66.

Eqtn. 2.66

$$E\left[Q_{M1M2}(x)Q_{M1M2}(y)\right] = k \cdot \delta(x-y)$$

Here, k is the second moment of the random variable Q_{M1M2}, and δ is the khroniker delta function such that Eqtn. 2.67 is true.

Eqtn. 2.67

$$\delta(x) = \begin{cases} 1 & \text{if } x = 0 \\ 0 & \text{otherwise} \end{cases}$$

Because Q_{M1M2} is the only random variable in Eqtn. 2.65, the expectation can be taken inside, and the first integration performed such that the result in Eqtn. 2.68 is obtained.

Eqtn. 2.68

$$\begin{aligned} E[P(f)] &= |Z_o|\omega^2 V_o^2(f)\int_0^L\int_0^L E\left[Q_{M1M2}(x)Q_{M1M2}(y)\right]e^{-2\gamma(f)x}e^{-2\gamma^*(f)y}dxdy \\ &= |Z_o|\omega^2 V_o^2(f)\int_0^L\int_0^L k\delta(x-y)e^{-2\gamma(f)x}e^{-2\gamma^*(f)y}dxdy \\ &= |Z_o|\omega^2 V_o^2(f)\int_0^L ke^{-2\gamma(f)x}e^{-2\gamma^*(f)x}dx \end{aligned}$$

The sifting property was used to obtain the last line of Eqtn. 2.68.

If the exponential terms in Eqtn. 2.68 are combined and the propagation constant expressed in terms of α and β, Eqtn. 2.69 results.

Eqtn. 2.69

$$\begin{aligned} E[P(f)] &= |Z_o|\omega^2 V_o^2(f)\int_0^L ke^{-2\gamma(f)x}e^{-2\gamma^*(f)x}dx \\ &= |Z_o|\omega^2 V_o^2(f)\int_0^L ke^{-2x[(\alpha(f)+j\beta(f))+(\alpha(f)-j\beta(f))]}dx \\ &= |Z_o|\omega^2 V_o^2(f)\int_0^L ke^{-4x\alpha(f)}dx \end{aligned}$$

The final integration can now be carried out. The result is in Eqtn. 2.70.

Eqtn. 2.70

$$E[P(f)] = |Z_o|\omega^2 V_o^2(f)\int_0^L ke^{-4x\alpha(f)}dx$$
$$= |Z_o|\omega^2 V_o^2(f)k\frac{e^{-4x\alpha(f)}}{-4\alpha(f)}\Big|_0^L$$
$$= \frac{|Z_o|\omega^2 V_o^2(f)k}{-4\alpha(f)}\left(e^{-4L\alpha(f)} - 1\right)$$

For long lines, the exponential in the result will be small compared to 1, and Eqtn. 2.70 can be simplified. In addition, Eqtn. 2.70 can be divided by the power of the disturbing transmitter ($|V_o|^2/Z_o$) to obtain the power transfer function between the disturbing transmitter and the disturbed receiver, denoted as H_{NEXT} as in Eqtn. 2.71.

Eqtn. 2.71

$$H_{NEXT}(f) = \frac{|Z_o|^2\omega^2 k}{4\alpha(f)} = \frac{k' f^2}{\alpha_o f^{1/2}} = K_{NEXT} f^{3/2}$$

In Eqtn. 2.71, K_{NEXT} is a constant of proportionality encompassing all the constants in the first two terms of Eqtn 2.71. Thus, for a disturber with a *flat power spectral density* (PSD), you have an expression showing that NEXT is proportional to $f^{3/2}$. Eqtn. 2.71 is sometimes called a *crosstalk loss function*. It is more generic than a general NEXT power equation because it does not depend on the PSD of the disturber. Instead, it can be multiplied by any PSD to give the expected NEXT power that the disturber would cause.

Note that NEXT does not depend on the length of the twisted pair. This is because the interference due to NEXT, although a sum of crosstalk at every point along the pair, will be dominated by the unbalance points nearest to the receiver. The crosstalk from the nearer unbalance points will dominate the crosstalk from distant unbalance points due simply to the length of twisted pair that each must propagate through.

Note also that K_{NEXT} is a random variable. For many different disturbers into a disturbed pair, K_{NEXT} will tend toward a Gaussian distribution due to the central limit theorem.[14,15,16]

Far-end Crosstalk (FEXT)

FEXT will be derived using a similar method as used in the NEXT derivation. Figure 2.21 shows an example of FEXT due to an unbalance point at x1 only.

Figure 2.21 An example of far-end crosstalk due to only one unbalance point (x1).

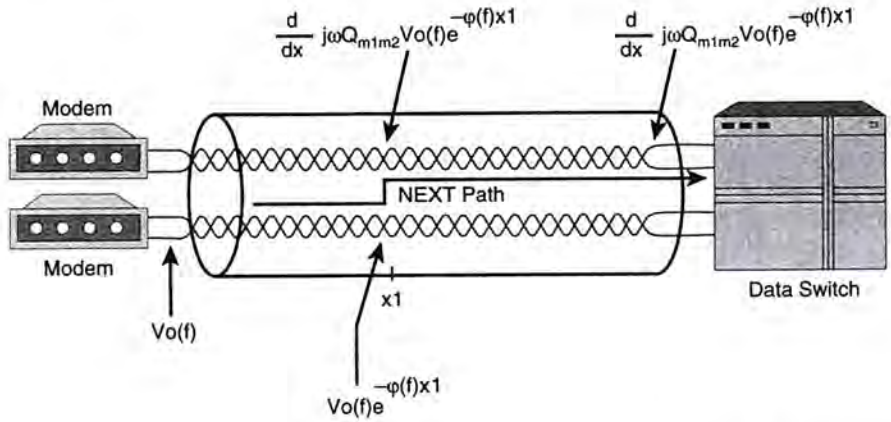

The signal expressions on the disturbing pair are the same as for the NEXT case, and the current at x1 on the disturbed pair is also the same as in the NEXT derivation. Repeating the expression for this current from Eqtn. 2.60 yields Eqtn. 2.72.

Eqtn. 2.72

$$I_2(f, x1) = \frac{d}{dx} j\omega Q_{M1M2} V_1(f, x1) = \frac{d}{dx} j\omega Q_{M1M2} V_o(f) e^{-\gamma(f)x1}$$

To reach the disturbed pair's receiver, the signal must propagate over the remainder of the disturbed pair (a distance of L-x1). Upon reaching the disturbed pair's receiver, the resulting expression for current on the disturbed line is shown in Eqtn. 2.73.

Eqtn. 2.73

$$I_2(f, L) = \frac{d}{dx} j\omega Q_{M1M2} V_o(f) e^{-\gamma(f)x1} e^{-\gamma(f)(L-x1)}$$
$$= \frac{d}{dx} j\omega Q_{M1M2} V_o(f) e^{-\gamma(f)L}$$

Just as in the NEXT case, the total crosstalk signal on the disturbed pair will be a summation of crosstalk due to unbalance points all along the twisted pairs and not just the one at x1. Eqtn. 2.73 can be integrated to find this expression. Moving one step farther, one can write an expression for the total FEXT power similar to Eqtn. 2.64 as in the following Eqtn. 2.74.

Eqtn. 2.74

$$E[P(f)] = E\left[|Z_o| I_2(f) I_2(f)^*\right]$$
$$= E\left[|Z_o| \int_0^L j\omega Q_{M1M2}(x) V_o(f) e^{-\gamma(f)L} dx \int_0^L - j\omega Q_{M1M2}(y) V_o(f) e^{-\gamma^*(f)L} dy\right]$$
$$= |Z_o| \omega^2 V_o^2(f) e^{-\gamma(f)L} e^{-\gamma^*(f)L} \int_0^L \int_0^L E[Q_{M1M2}(x) Q_{M1M2}(y)] dx dy$$

The relationship given in Eqtn. 2.66 can again be used to help carry out the integral in Eqtn. 2.74 so that the expression for FEXT simplifies to Eqtn. 2.75.

Eqtn. 2.75

$$E[P(f)] = |Z_o| \omega^2 V_o^2(f) e^{-2\alpha(f)L} \int_0^L k dx = |Z_o| \omega^2 V_o^2(f) e^{-2\alpha(f)L} L$$

The constants in Eqtn. 2.75 can be grouped together in a single constant, and the average FEXT transfer function can be written as Eqtn. 2.76.

Eqtn. 2.76

$$H_{FEXT}(f) = K_{FEXT} f^2 e^{-2\alpha(f)L} L$$

Note that K_{FEXT}, like K_{NEXT} is also a random variable. When multiple disturbers exist, K_{FEXT} can be expected to tend toward a Gaussian random variable due to the central limit theorem. Note also that FEXT, unlike NEXT, is a function of the length of the twisted pair. This fact is intuitively satisfying because FEXT disturbances due to any unbalance point along the twisted pairs must propagate the entire length of the pair to disturb the receiver. (At this point, the reader should realize that both NEXT and FEXT increase with frequency, and thus a twisted pair designed for low frequency signals presents challenges when high frequency signals are involved.)

Summary

This chapter was dedicated to better understanding twisted pair physical and electrical properties. The primary line constants, RLCG, of a twisted pair were discussed. Numeric and analytic models of a twisted pair's propagation constant and characteristic impedance were derived. The attenuation of various twisted pair cables was also illustrated.

Crosstalk between twisted pairs was discussed in some detail. The basic cause of crosstalk, unbalance between neighboring pairs, was derived from basic circuit theory. From the unbalance expressions, expressions for near-end crosstalk (NEXT) power and far-end crosstalk (FEXT) power were derived.

Understanding the challenges that a twisted pair environment poses to an ADSL or a VDSL environment will help the reader understand the motivation for some of the modulation and coding schemes presented in later chapters. Also, the models developed here will be used directly in later chapters to predict and bound the performance of certain ADSL and VDSL examples.

Exercises

1. For #24 and #26 gauge twisted pairs, plot over frequency the characteristic impedance and propagation constant using the analytic and numeric models.
2. Find the loss at 500 kHz of a 10 kft, #26 gauge twisted pair using both the analytic and numeric models.
3. For a #26 loop, plot Eqtn. 2.70 for various lengths. At what length does the approximation of Eqtn. 2.71 become valid?
4. Graph the NEXT and FEXT loss functions on the same set of axes for different lengths of #24 and #26 twisted pair. At what length does NEXT become dominant over FEXT?
5. Plot the loss of a #26 twisted pair at 500 kHz against the length of the pair. At what length is the loss equal to the NEXT loss at 500 kHz?

Endnotes

1. Rhodes, Frederick Leland, *Beginnings of Telephony.* New York: Harper and Brothers, 1929.
2. Manhire, L. M., "Physical and Transmission Characteristics of Customer Loop Plant." *Bell System Technical Journal* 57 (January 1978): 35–39.
3. Chin, Wu-Jhy, et al, "Loop Survey in the Taiwan Area and Feasibility Study for HDSL." IEEE Selected Areas in Communications 9, no. 6 (August 1991): 801–809.
4. Reeve, Whitham D., *Subscriber Loop Signalling and Transmission Handbook.* New York: IEEE Press, 1992.
5. Bellamy, John, *Digital Telephony.* New York: John Wiley and Sons, 1991.
6. Gang, Huang, and J.J. Werner, "Cable Characteristics." *ANSI Contribution T1E1.4/97-169*, Clearwater, FL, May 12, 1997.
7. Gilbert, Elliot and John Nellist, *Telecommunications Wiring for Commercial Buildings, A Practical Guide.* New York: IEEE Press, 1996.

8. Pickering, Ashley, "Primary Line Constants for US 24 and 26 AG Twisted Pair Cables." *ANSI Contribution T1E1.4/96-074.* Colorado Springs, CO, April 22, 1996.

9. Sadiku, M., *Elements of Electromagnetics.* Philadelphia: Saunders College Publishing, 1989.

10. Campbell, G., "Dr. Campbell's Memoranda of 1907 and 1912." *Bell System Technical Journal* 14, no. 4 (October 1935): 558–472.

11. AT&T Company, *Collected Papers of George A. Campbell.* New York: AT&T, 1937.

12. Cravis, H. and T.V. Carter, "Engineering of T1 Carrier System Repeated Lines." *Bell System Technical Journal* 42, no. 1 (March 1963): 431–486.

13. Miller, G., "The Effect of Longitudinal Imbalance on Crosstalk." *Bell System Technical Journal* 54, no. 7 (September 1975): 1227–1251.

14. Kerpez, K., "Near End Crosstalk Is Almost Gaussian." *IEEE Transactions on Communications* 41, no. 1 (January 1993): 670–672.

15. Werner, J.J., "The HDSL Environment." *IEEE Journal on Selected Areas in Communications* 9, no. 6 (August 1991): 785–800.

16. Gallager, Robert, *Information Theory and Reliable Communications.* New York: John Wiley and Sons, 1968.

This material may be protected by Copyright law (Title 17 U.S. Code)

CHAPTER 8

ADSL in WAN Networks

In this chapter:

- An overview of common protocols found in LANs and WANs relevant to networks utilizing ADSL access
- Several candidate architectures for WANs that use ADSL as an access technology
- A discussion of options and issues for ATU-R connectivity to the network

Until now, ADSL and other DSL technologies have been covered at the *physical layer*. Sometimes the physical layer is also called *layer 1*. Both refer to the first layer in the seven-layer OSI protocol stack shown in Figure 8.1.

Figure 8.1 The seven-layer OSI protocol model; each layer is independent and passes data to and from neighboring layers.

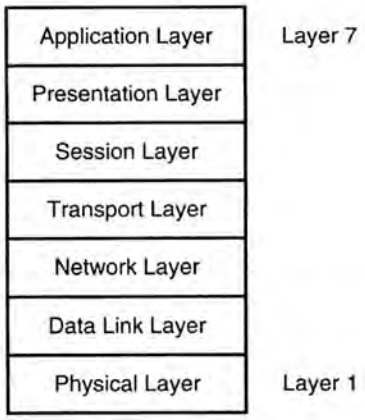

The physical layer is characterized by the type of transmission used, including the parameters of the transmitter and receiver blocks discussed in Chapter 7, "ADSL Modulation

Specifics." These parameters include framing structure, scrambling, FEC method, and modulation type.

When looking at a network as a whole, the layers above the physical layer play a large role in how the network is set up, provisioned, and managed as well as how data is sent over the network. ADSL, when used in the access link of a network, should not affect the network's operation and in most cases the underlying line code should be irrelevant.

Note

In some cases, the actual access speed and asymmetry, but not the line code, of the DSL might affect the design of the higher layer protocols.

This chapter discusses protocols used in networks that ADSL is used as an access medium. Several of the more common types of architectures being designed and deployed are covered along with the protocols. Many of these architectures assume that the Internet standard TCP/IP protocols are used somewhere in the protocol stack. Before jumping directly into the architectures, several sections discuss the basic features of the various protocols. (This presentation is not a complete and thorough treatment of each protocol—in many cases a book on each subject would be necessary for complete treatment.)

Protocol Basics

The first two protocols discussed are *Ethernet* and *Internet Protocol* (IP). These are the most common protocols in LANs. IP is also the protocol used by the public Internet. Figure 8.2 shows how IP and Ethernet fit into the seven-layer OSI model. Additionally, TCP is shown in this figure as an example wherever a Web browser is being used. The common physical layer of 10BaseT is used in the example. After discussing TCP/IP and Ethernet, asynchronous transfer mode (ATM), Point-to-Point Protocol (PPP), tunneling, and other protocols are discussed.

Figure 8.2 Mapping a typical TCP/IP implementation into the seven-layer OSI protocol model.

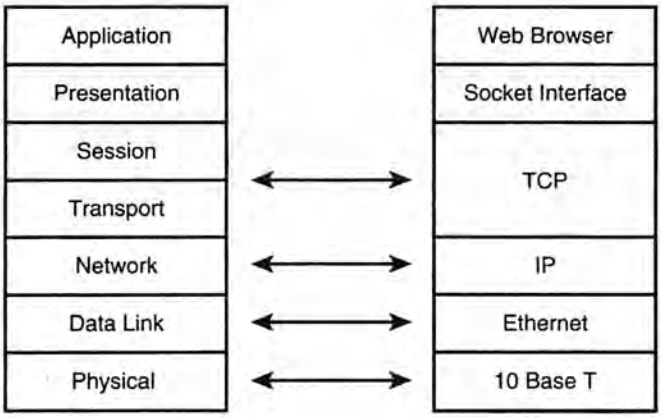

Ethernet

Ethernet is a layer 2 protocol. It is commonly subdivided into a *multiple access control* (MAC) sublayer and a *logical link control* (LLC) sublayer. Figure 8.3 shows two Ethernet frame formats carrying IP traffic. The first is consistent with IEEE 802.2/802.3, and the second with RFC 894.[1,2] RFC 894 encapsulation is more commonly used.

Figure 8.3 Ethernet frame formats according to (a) IEEE 802.2/802.3 and (b) RFC 894.

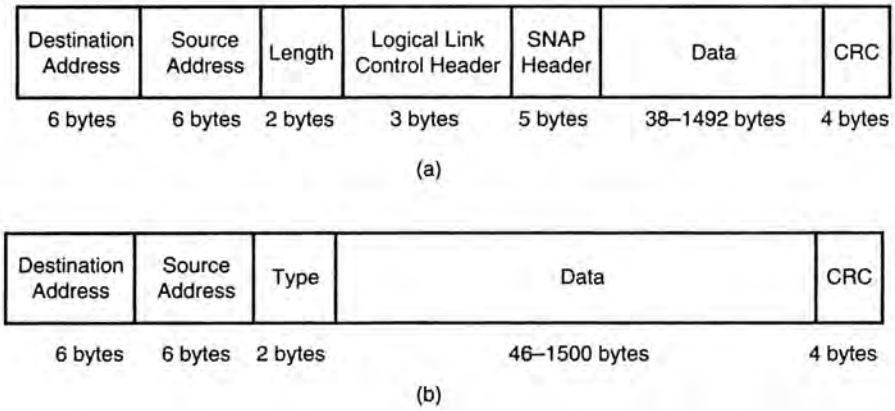

The Ethernet frame using RFC 894 encapsulation contains four fields in addition to the IP data field. A *Source Address field* and a *Destination Address field* identify the Ethernet port sending (source) and receiving (destination) the Ethernet packet. The six-byte addresses used for the source and destination address fields are also called *Ethernet addresses* and sometimes *hardware addresses.* Typically a manufacturer of Ethernet cards preconfigures the Ethernet address for each card (and each card or port has a different Ethernet

address). Typically, Ethernet addresses do not have local significance. That is, ports on the same network have Ethernet addresses with no common portion among the six address bytes. This scheme is in contrast to IP addresses on the same network that have a common portion.

Note

Some Ethernet cards allow a user to change some or all of the six bytes.

The two-byte Type field in part b of Figure 8.3 identifies the type of layer 3 protocol being encapsulated in the frame. If the public Internet is being accessed, the protocol is IP and the type value assigned for IP encapsulation is 0x0800. It is possible to encapsulate other types of protocols in an Ethernet frame as well, resulting in a different type value. The Type field is sometimes called the *ethertype*.

The final field in the Ethernet frame in part b of Figure 8.3 is the CRC field. This four-byte field covers the entire frame (except for the CRC itself). Note that this CRC has no relationship to the CRC discussed for ADSL (which is purely a physical layer entity). The CRC field determines whether the layer 2 packet was completely and successfully received from the physical layer.

IP

The IP is a layer 3, or network layer, protocol. Figure 8.4 shows the IP packet format.

Figure 8.4 A breakdown of the fields in an IP header.

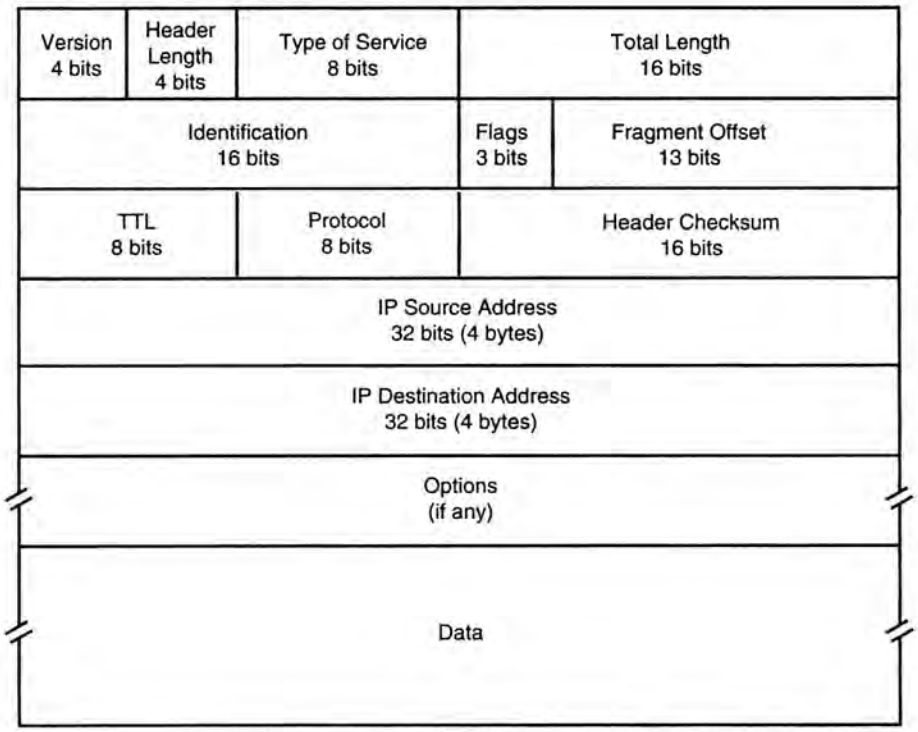

Table 8.1 describes the many fields in an IP packet. Of these fields, the two most important to our discussion are the *Source IP Address* field and the *Destination IP Address* field.

Table 8.1 IP Packet Fields

Field Name	Length	Description
Version	4 bits	The IP protocol version number (currently 4).
Header Length	4 bits	The number of 32-bit words in the header.
Type of Service (TOS)	8 bits	Contains four fields specifying the precedence (usually ignored), delay requirement (high or low), throughput maximization (high or low), and a reliability parameter.
Total Length	16 bits	Specifies the number of bytes in the IP packet, including the header.
Identification	16 bits	Modulo 16 counter that identifies an IP packet.

continues

Table 8.1 Continued

Field Name	Length	Description
Flags	3 bits	The DF bit signifies whether the IP datagram may be fragmented, and the MF bit indicates whether the packet is the final fragment. One of these three bits is not used and always set to zero.
Fragment Offset	13 bits	If the data is a fragment, this bit indicates where it belongs in the entire datagram. The unit of measurement is 64-bit units.
Time to Live (TTL)	8 bits	Indicates how many router hops this packet may pass through before being discarded. Normally, routers decrement this value. The packet is discarded when this value hits zero, keeping a packet from eternally existing in a network.
Protocol	8 bits	Identifies the protocol that follows the IP header (TCP, UDP, IGMP, and so on).
Header Checksum	16 bits	A checksum to ensure the reliability of the header.
Source IP address	32 bits	The IP address of the host that sent the IP packet.
Destination IP Address	32 bits	The address of the host for which this packet is intended.
Options	Variable	May set routing or security restrictions on the packet. Options are usually not used.
Data	Variable	The IP data payload.

Each of the fields documented in the preceding table is four bytes (32 bits) long. A network administrator normally assigns four-byte IP addresses to each host (computer) on a network. All computers on a LAN typically have a common portion of their IP addresses that identifies the network followed by a unique portion representing each individual host.

Consider the example in Figure 8.5.

Figure 8.5 An example of two subnets with multiple hosts connected by a router.

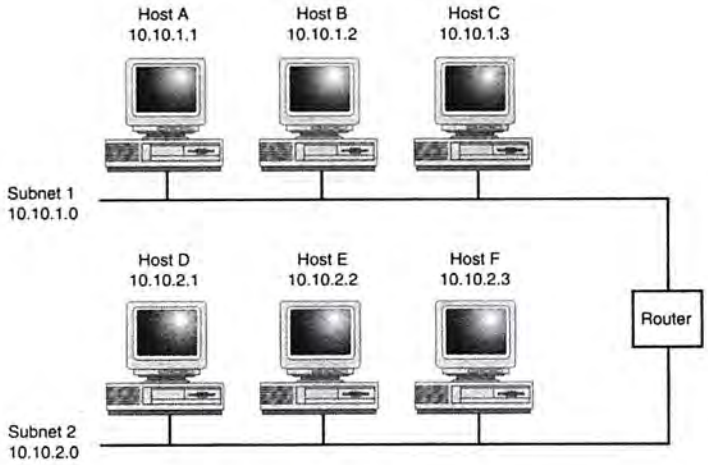

This figure shows two *subnets*, or logical networks: 10.10.1.0 and 10.10.2.0. The first item to note is the method used to write addresses. This method is called *dotted-decimal notation*. Each decimal number represents one of the four bytes making up the IP address. Because each number represents eight bits, the number may be any positive integer between 0 and 255. In this example, each subnet has three hosts. Hosts on the same subnet have the same first three bytes. However, the common part of a network address (representing the subnet) may differ by three bytes depending on the network.

Note

Associated with a network is a 32-bit subnet mask containing 1s for the N MSBs and 0s for the 32-N LSBs (N may be any number between 1 and 31). A bitwise AND function performed on the mask and IP address from a network yields the network part of the IP address. A bitwise NAND function yields the host portion of the IP address.

You can find very detailed and informative discussions of IP in *TCP/IP Illustrated, Volume 1: The Protocols*; *TCP/IP Illustrated, Volume 2: The Implementation*; *and TCP/IP: Architectures, Protocols and Implementation.*[3,4,5]

IP networks can be viewed as a number of connected LANs. Switches, bridges, and routers are devices that connect hosts together on Ethernet and IP networks. Switching and bridging are layer 2 functions, and routing is a layer 3 function. Figure 8.6 shows a typical bridging application.

Figure 8.6 A bridging application using two half-bridges and modems to connect hosts on the same subnet.

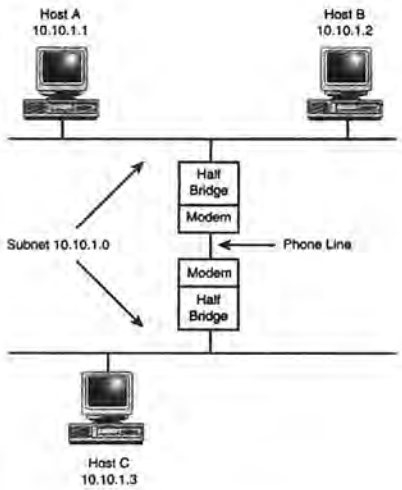

The two half-bridges in Figure 8.6 make up a bridge connecting two Ethernet segments. Ethernet packets destined for ports on the other Ethernet segment are transparently sent over the bridge. To function properly, the bridge must know which Ethernet addresses are on each segment. This information can be manually configured in the bridge but is more often learned by the half-bridges. To learn the location of different Ethernet ports, each half-bridge monitors the Source Ethernet Addresses of packets on its Ethernet segment. All source addresses are entered into a table. If a half bridge sees a packet with a Destination Ethernet Address not in its table, that host must be on the other side of the bridge and the Ethernet frame will be forwarded. A bridge does not change the contents of a packet, including the Ethernet source and destination addresses. Bridges allow two Ethernet segments to appear to be connected. Segments separated by a bridge are typically part of the same IP subnet.

An Ethernet switch connects to multiple Ethernet segments and simply switches packets between segments depending on the destination Ethernet address. Switches are convenient for separating segments to reduce traffic on each segment. Segments connected by an Ethernet switch normally are part of the same IP subnet. A switching table kept by the Ethernet switch keeps a mapping between the Ethernet Destination address and the port to which the frame should be switched. Figure 8.7 shows multiple Ethernet segments connected by a switch and the corresponding switching table.

Figure 8.7 Multiple Ethernet segments connected by a switch.

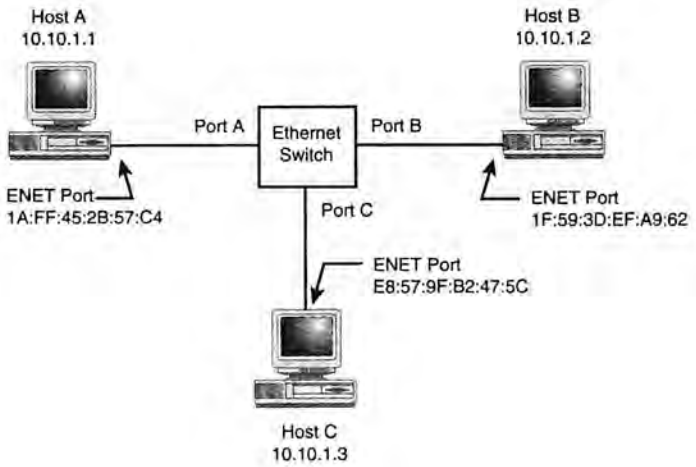

Input Port	Ethernet Destination Address	Output Port
A	1F:59:3D:EF:A9:62	B
A	E8:57:9F:B2:47:5C	C
B	1A:FF:45:2B:57:C4	A
B	EF:57:9F:B2;47:5C	C
C	1A:FF:45:2B:57:C4	A
C	1F:59:3D:EF:A9:62	B

A router is a layer 3, or network layer, device that transfers IP packets, based on the destination IP address. Routers connect different local IP networks. Figure 8.8 shows a router connecting two networks. Overlaid on Figure 8.8 is an example of a packet being sent from host A to host B.

The following events transpire to make the transfer in Figure 8.8 happen:

- Host A has an IP packet to send to host B. The destination IP address of the packet identifies host B and the source IP address of the packet identifies host A.
- Host A knows that host B is not on the same subnet. Host A thus sends the packet to host A's default router. The source and destination IP addresses of the packet still identify host A and host B respectively. The source Ethernet address identifies host A's Ethernet port, and the destination Ethernet address identifies the router's Ethernet port.
- The router receives the packet, recognizing the destination Ethernet address as its own. The router's IP layer recognizes that the destination IP address does not belong to the router. The router then looks into a routing table and identifies an output port where the IP packet (or IP datagram) should be sent to reach host B.

Figure 8.8 An example of routing between several subnets; the addresses show the IP and Ethernet addresses present in packets on each side of the router when routed from host A to host B.

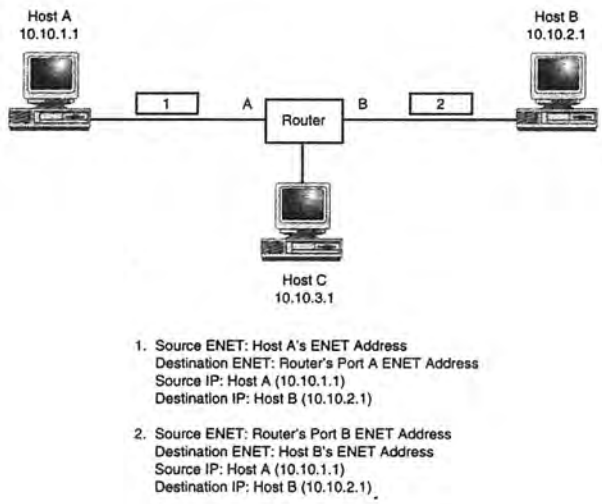

Note

In many cases, the destination host may not be directly connected to the first router. In such a case, the router forwards the packet to another router. Eventually, barring any errors, the host will be reached.

- The router does not modify the IP part of the packet. It reformats the Ethernet part of the message to include the router's Ethernet port on subnet B as the source Ethernet address and the Ethernet address of host B as the destination Ethernet address. The newly formatted Ethernet packet is then sent on to port B's Ethernet segment.
- Host B's Ethernet port recognizes an incoming packet with its Ethernet address in the Destination Ethernet Address field. Recognizing that this packet is for itself, host B's Ethernet port forwards the IP portion of the packet to its IP layer. The IP layer recognizes its own IP address in the Destination IP Address field and knows that the IP packet is for host B. The packet will then be sent further up host B's protocol stack for handling.

Note from this example the difference in usage of IP and Ethernet addresses. If a host in California needs to send an IP packet to a host in New York via the Internet, the sending host must know the IP address of the host in New York, but does not need to know the Ethernet address. (Most probably, the IP address of the destination host would be found using a *Domain Name Server* [DNS].) Public IP addresses on the Internet are administered and maintained by several organizations. Understanding the difference between IP

addresses and Ethernet addresses, as well as the difference between layer 2 and layer 3 devices, is necessary to really understand the networking issues discussed later in this chapter.

TCP

The IP and Ethernet protocols do not ensure that a packet sent to a destination actually arrives at the destination. A channel is not set up, and the receiving host does not send an acknowledgment. For this reason, IP is called a *connectionless* protocol. TCP, on the other hand, is a *connection-oriented* protocol. When data is to be transferred to another host, TCP uses IP to contact the destination and tell the host at the destination that it wants to set up a *session*; a session is basically a connection made between two hosts and acknowledged by both endpoints. When a session is set up, TCP sees to it that data is transferred between the endpoints and that the receiving host sends acknowledgment packets. If an acknowledgment is not received for a certain packet, the TCP protocol resends the packet.

ATM

Asynchronous transfer mode (ATM) is newer than the IP and Ethernet protocols. It is sometimes referred to as *cell transfer*, whereas IP and Ethernet are thought of as packet transfer technologies. ATM is a layer 2 protocol. Figure 8.9 shows a protocol stack of ATM carrying IP.

Figure 8.9 A protocol stack with IP running over ATM.

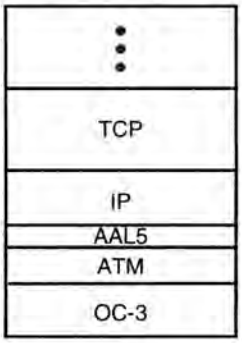

Here, an OC-3 (155 Mbit/s optical fiber) is assumed as the physical layer. The *ATM adaptation layer 5* (AAL5) layer shown adapts the higher-layer protocol (IP) so that it will fit nicely into ATM cells.

In ATM, instead of variable-length packets, fixed-length cells are created and sent. The cell length is 53 bytes, 48 of which are payload bytes and 5 of which make up a header. Figure 8.10 shows the format of an ATM cell. Of particular interest in the ATM cell header are the *virtual path identifier* (VPI) and *virtual channel identifier* (VCI) values. ATM switches use the VPI and VCI values to distinguish local connections and direct cells from an input to an output. The fixed, short size of ATM cells allows switches to run faster than most IP routers.

Figure 8.10 Format of an ATM cell at a user-to-network interface.

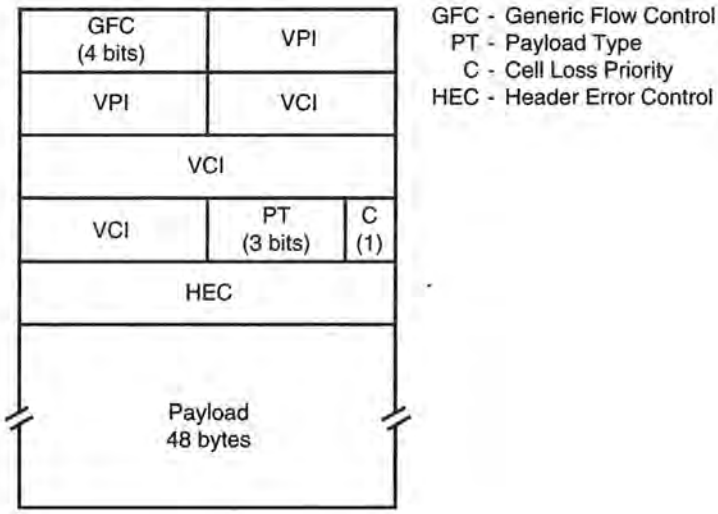

ATM is a connection-oriented protocol. A connection is set up between the ATM endpoints before payload data is transmitted. A connection either can be provisioned into the equipment by an operator or can be dynamically set up using signaling. The former type of connection is called a *permanent virtual circuit* (PVC), and the latter is a *switched virtual circuit* (SVC). PVCs exist until an administrator unprovisions the connection. SVCs exist for the duration of the ATM session and are set up and torn down by the signaling protocol. ATM SVC setup is very analogous to dialing a telephone in that an end station's ATM address is specified and sent to the network. ATM adds the capability of specifying how much bandwidth and what type of traffic will flow through the configured path. It is said that a *Quality of Service* (QoS) is requested for the call. All switches through which the virtual circuit will pass have the ability to deny the connection if they cannot support the requested QoS. Supporting QoS is a major advantage of ATM over a protocol such as Ethernet. Some disadvantages of ATM are the high cost compared to Ethernet and the relative amount of deployed equipment with respect to Ethernet (especially equipment supporting SVCs). This chapter discusses cell architectures, packet architectures, and mixed architectures for networks including ADSL as an access mechanism.

For detailed discussions of ATM, refer to *ATM: Foundations for Broadband Networks* and *Asynchronous Transfer Mode: ATM Architecture and Implementation.*[6,7]

PPP

PPP is used on many dial-up modem connections. It is specified in RFC 1661.[8] PPP is a layer 2 protocol and provides a simple method to encapsulate any protocol inside. PPP consists of a link-negotiation phase called *Link Control Protocol* (LCP) and a network negotiation phase called *Network Control Protocol* (NCP).

PPP adds several bytes of overhead to a packet. The format of a PPP packet is shown in Figure 8.11.

Figure 8.11 The format of a PPP packet.

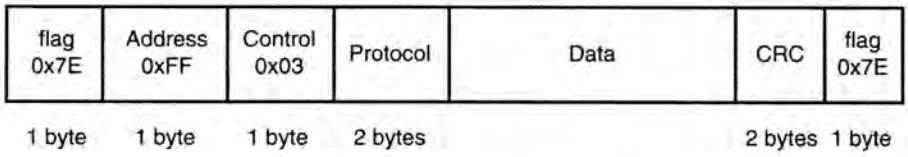

The type of information in the Data field is specified by the value in the Protocol field. Thus two bytes are used to indicate what type of protocol is being encapsulated by PPP. The value most relevant for ADSL access is the protocol value for the IP protocol, 0x0021.

When a link is first brought up, LCP negotiation is performed. During LCP, many different data-link parameters are negotiated. Some of these include the maximum PPP packet size, the authentication protocol to be used, and protocol and header compression options. LCP packets are identified by the protocol field value of 0xC021. These packets are sent when the link is first being made.

After LCP negotiation, NCP negotiation occurs. For IP, the NCP negotiation used is called *Internet Protocol Control Protocol* (IPCP), which is defined in RFC 1332.[9] IPCP allows an IP address to be assigned to an end station. NCP packets are identified by the protocol field value of 0x8021.

Two important authentication protocols can be used with PPP. They are Point-to-Point Authentication Protocol (PAP) and Challenge Authentication Protocol (CHAP). PAP consists simply of one of the endpoints (normally the dial-up user) sending a username and a password to the other end (normally the network being accessed). CHAP consists of the one end of the connection (normally the network end) sending a challenge to the other end and that end replying to the challenge.

PAP or CHAP is normally run between the LCP and NCP negotiations. When used with IP, PPP provides a powerful method to accomplish the following:

- Configure a link so that packets can be reliably sent between the endpoints (LCP)
- Make sure, with reasonable certainty, that the user accessing the network is an authorized user (PAP or CHAP)
- After having confidence in the user's identity, assign that user an address on the accessed network (NCP)

When all of the negotiation and authentication activities are complete, the connection is ready to send user data between the two endpoints.

Tunneling

PPP provides an effective way to manage a session between two endpoints. However, one drawback of PPP is that it has no addressing mechanism, and so a single logical connection existing between the two endpoints participating in the session can carry only one PPP session. Because of the lack of addressing mechanism, PPP does not allow a group of PPP sessions to be aggregated at an access node and sent over the same channel to a common destination. Tunneling methods do allow this however.

A *tunnel* simply refers to a single logical channel over which sessions that normally do not share a logical channel are sent. Three of the more common tunneling methods are *Point-to-Point Tunneling Protocol* (PPTP), *Layer Two Forwarding* (L2F), and *Layer Two Tunneling Protocol* (L2TP). L2TP is a superset of PPTP and L2F developed by proponents of both of these protocols and is intended to replace them.

L2TP is designed to multiplex multiple PPP sessions over a common channel. Essentially, this is done by assigning a call number to each PPP session at one end of the tunnel before multiplexing so that the sessions can be differentiated at the other end of the tunnel. Figure 8.12 shows an example of a number of PPP sessions between subscribers and a service provider.

In part a of Figure 8.12, PPP aggregation does not occur. In this situation, each session requires a separate physical channel between each subscriber and a service provider. Part b of Figure 8.12 shows an example in which PPP sessions are aggregated, but tunneling is not used. In this case, separate logical channels between the aggregation point and the service provider are necessary to keep the PPP sessions identifiable. Part c of Figure 8.12 shows aggregation with tunneling. In this case, only one logical connection is needed between the aggregation point (commonly called the *L2TP access concentrator* or LAC) and the other tunnel endpoint (commonly called the *L2TP network server* or LNS) at the service provider.

Figure 8.12 Comparison of access systems (a) without aggregation, (b) with aggregation but without tunneling requiring, and (c) with aggregation and tunneling.

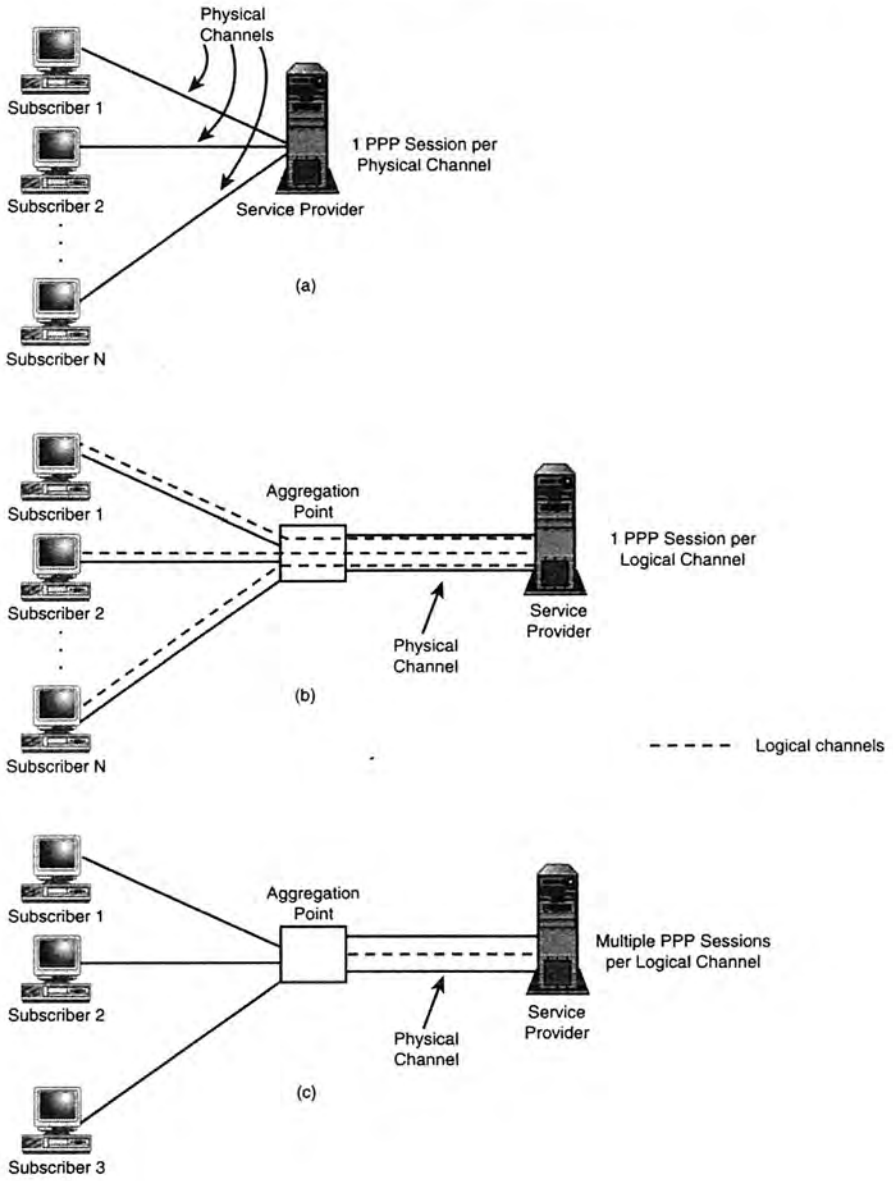

Note

For example, in an ATM network, each PPP session would have a unique VPI/VCI assigned to it. In a frame relay network, each PPP session would have a unique data link connection identifier (DLCI) assigned to it.

The value of L2TP is illustrated in Figure 8.13.

Figure 8.13 An access network using L2TP tunneling between access points and a service provider.

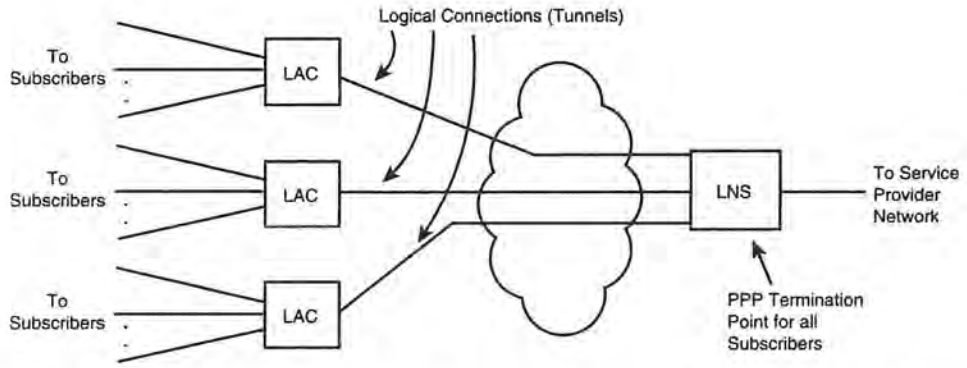

In this figure, access to some network (the Internet, for example) is located at a central point (sometimes called a *point of presence*, or POP). Access nodes (LACs) are located at various points away from the LNS. Because PPP provides methods to authenticate users (from which billing can be derived) and assign addresses to users, terminating the PPP sessions at a common point (the LNS), instead of at all the access points, is convenient. On the other hand, provisioning a logical connection for every subscriber between all the LACs and the LNS is usually not convenient. L2TP solves this issue by allowing a single logical connection between each LAC and the LNS to transport the PPP sessions. You can find more detailed information on L2TP in the Internet Draft, "Layer Two Tunneling Protocol (L2TP)."[10]

End-to-End Players

A firm understanding of networks with ADSL access involves understanding the different players that take part in making the access possible and providing services. These include end users or subscribers, network access providers (NAPs), network transport providers (NTPs), and network service providers (NSPs). The next few sections discuss these players and where they fit into the end-to-end picture. Figure 8.14 illustrates the different players. In many cases, the NAP, NTP, and NSP may not be separate companies; a single company may do any or all of these functions.

Figure 8.14 Illustration of the end-to-end players relevant to a DSL access network.

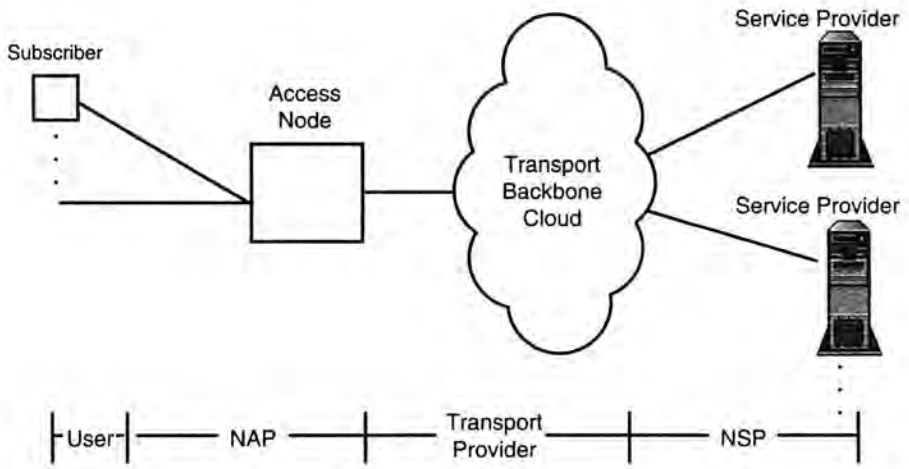

Users

Users or subscribers are one endpoint of an end-to-end connection. Users include residential subscribers wanting Internet access, small business offices wanting connectivity to other offices, and telecommuters. A user might have a full-time dedicated connection to a single location or may want to connect to various places at various times (connect to the office during the day and to the Internet in the evening).

Network Access Providers

NAPs normally own and operate the equipment that terminates the physical ADSL link and aggregates traffic. Traffic from a NAP may go to one or more NSPs. NAPs may terminate more than simply ADSL, including other DSL flavors or even voice-band modem access.

Network Transport Providers

A *network transport provider* moves traffic between NAPs and NSPs. The transport provider typically owns the transport cloud shown in Figure 8.14. (Other arrangements are possible as well. For example, the transport provider may lease some of the cloud.) Note that the cloud is typically made up of various types of interconnected switches and routers. A transport provider might carry traffic from various NAPs or NSPs. In addition, the transport provider might carry traffic totally unrelated to ADSL networks, including long-distance voice service or private, dedicated connections between corporate locations.

Network Service Providers

NSPs are typically an endpoint of a PPP session originating from a user. Examples of NSPs are Internet service providers, corporate networks providing access to corporate telecommuters, or banking institutions providing online banking functions. In many cases, NSPs provide users access to some type of network (the Internet or a corporate LAN). The NSP frequently assigns IP addresses to users via IPCP, authenticates users to make sure that they are allowed to access the network, and keeps track of billing and usage information.

ILECs and CLECs

Incumbent local exchange carriers (ILECs) and competitive local exchange carriers (CLECs) are companies that can be NAPs, transport providers, and NSPs. Traditionally, ILECs are Regional Bell Operation Companies (RBOCs). Competitive carriers have emerged in the 1990s since the Telecom Act ruled that the ILECs must fairly lease their copper facilities to competitors to open up the market of providing basic services. ILECs generally own central offices, the bulk of the copper in their area's network, and the bulk of the equipment in the central office. CLECs typically lease space from ILECs in a central office, lease copper on which to provide service, and in some cases, lease equipment and service from the ILEC.

Note

Often, central office space is literally a 10-foot by 10-foot cage to which the CLEC receives a key. Copper pairs are run to the cage for the CLEC by the ILEC, and all of the CLEC's equipment must be stationed inside the cage.

Architectures for DSL Networks

Many different architectural options exist for networks involving DSL access. Depending on the application of the network, some architectures may work better than others. For example, a virtual LAN network in which some traffic might be meant for all users would have different requirements than a dial-up network meant to privately connect subscribers to the corporate networks or the Internet. The following sections discuss some of the more common architecture approaches. In many cases, these architectures might coexist (at least in part) and hybrid approaches may be implemented.

Bridging Architectures

Earlier types of ADSL networks were based on bridging across the loop with routing in the network backbone. Such a network is shown in Figure 8.15.

Figure 8.15 A bridging architecture for ADSL access.

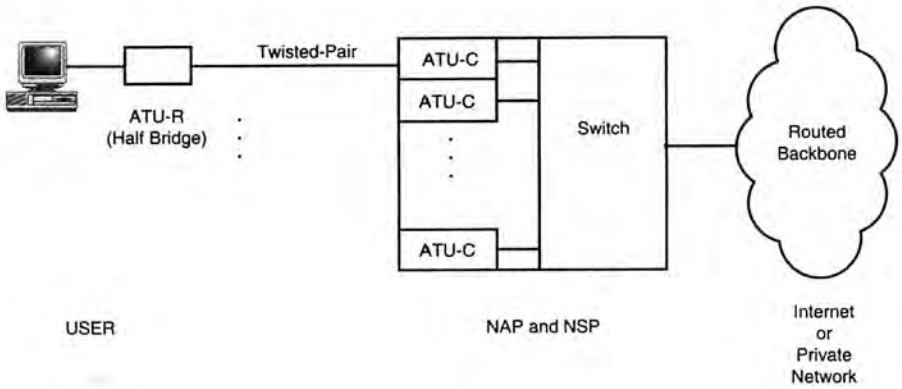

In this figure, the end users are connected to the network through external modems with Ethernet interfaces. (The ATU-R and ATU-C are considered half-bridges.) It is also possible to form a switching backbone; however, the provisioning needs and scalability of the network might be affected because the switches would require very large switch tables. In addition, MAC addresses can normally not be assigned to end stations because they are commonly permanently configured during the ports' manufacturing.

Note

Even if the MAC address were not factory configured, no common protocols exist to assign addresses.

In this type of network, the NAP typically is the NSP allowing a user to access a network. That network might be the Internet, a private network with gateways to the Internet, or other locations such as corporate networks.

Upstream Traffic

A packet generated by an end user's computer will be seen by the ATU-R. The ATU-R half-bridge can learn what Ethernet addresses exist on its Ethernet segment by monitoring and tabulating the source MAC address of all packets that it does not generate. It will bridge to the ATU-C any packet received on its local Ethernet having an Ethernet Destination Address field that does not match an entry in its address table. At the ATU-C, the Ethernet packets bridged from the ATU-R would be aggregated by either a hub or an Ethernet switch and sent to a router. The router is at the edge of an IP routed backbone connecting the access node to various gateways including, for example, the Internet and corporate networks.

Downstream Traffic

Traffic destined for a subscriber is routed through the backbone to the proper access node. The access node router must route the packet to the proper port interfacing the switch. Note that the access router must insert the end user's Ethernet address into the destination Ethernet field of the packet before sending the packet to the switch. The switch can then switch the packet to the appropriate ATU-C to be bridged over to the user's ATU-R and Ethernet segment.

Architecture Evaluation

Ethernet bridging over the loop combined with routing in the backbone offers a very simple architectural option for applications where the NAP and NSP are the same entity. In this case, a true transport provider does not necessarily exist. This type of network may have security and scaling limitations, depending on the implementation. Handling of broadcast messages as well as address assignment also raises security issues discussed in the second bullet below. Although some of the problems can be overcome with special implementations, the solutions engender other problems. Among the issues that arise with this type of network are the following:

- Secure address management—If PPP is not used (and thus PPP authentication methods are not used), dynamic address assignment can be troublesome, as most dynamic address assignment protocols do not include authentication. Addresses can be permanently assigned to users; however, this task would be difficult and might consume too many IP addresses. (Normally, an NSP has fewer IP addresses to assign than it has users, as all users are usually not active simultaneously.)
- Broadcast packet handling—Ethernet bridges and switches can sometimes be configured to block broadcast packets. However, some broadcasts must be passed between the NAP router and the end PC's (address resolution protocol (ARP) packets, for example). This process opens the door for packets from one private LAN to be present on another private LAN.
- Session establishment—In this type of network, it can be difficult to monitor the actual network usage by a user. Such information is often useful for billing purposes.
- Flexibility—This type of network is meant for the joint NAP/NSP arrangement. If the NSP is located at the other side of the backbone, more problems of security and scalability arise. Certain relationships between the NSP and NAP can alleviate some of these problems; however, these are not always practical in real life.

ATM End-to-End Architectures

A more popular architecture for ADSL access systems is the ATM end-to-end architecture. Several variations of this architecture have appeared, including those based on true PVCs, those based on SVCs, and some hybrid approaches. PVC- and SVC-based networks are discussed in the next few sections. Although different types of protocol models have been implemented, the most popular involve PPP. For this reason, models involving PPP are the primary focus. ATM architectures are discussed in the ADSL Forum contributions, "TR-002: ATM over ADSL Recommendations" and "Enhancements to Core Networks Architectures for ADSL Access."[11,12]

PVC-Based Models

The block diagram in Figure 8.16 shows a network with PVCs end-to-end. The figure shows several subscribers along with several NSPs. The NAP, transport provider, and all NSPs are treated as independent entities.

Figure 8.16 An end-to-end PVC access architecture with PPP over ATM.

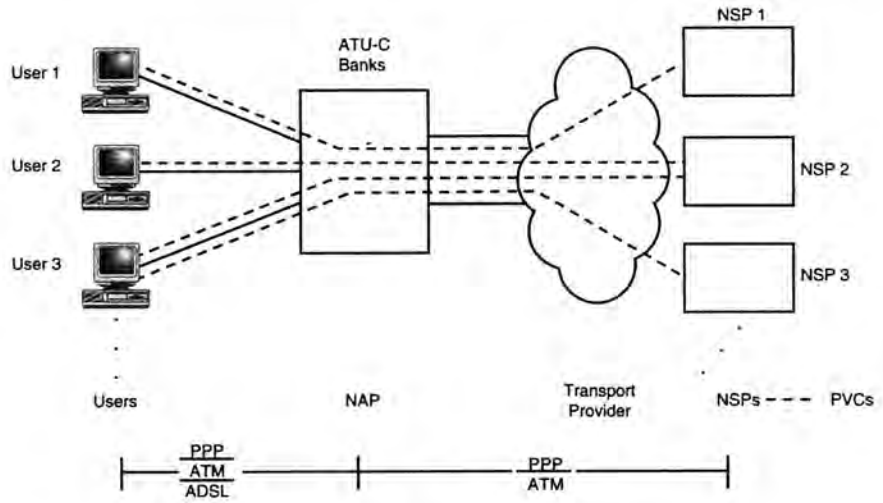

Figure 8.16 also shows the protocol model for the connection. It is important to note that both ATM and PPP are present end-to-end. Over the loop, these protocols are carried over ADSL and in the transport backbone over whatever physical layer is being used. (Often this is DS3 or OC-3.) Typically, the NSP to which traffic is forwarded is a gateway to another network. This network could be a private corporate network, the Internet, a banking institution's network, or perhaps some type of entertainment provider.

The model assumes neither a specific type of network at the user location (the premise network) nor a specific type of NSP network. It should be noted that to the access and

transport domains, the protocols in the premise network and NSP do not matter, as they operate at the ATM layer only and do not interact with the higher-layer protocols. A good portion of the time, the higher-layer protocol above ATM will be IP. Much of this discussion is based on that assumption.

Note

The protocols in the NSP network and the premise network must be compatible with one another, or one of the two entities must have the proper protocol conversion equipment to ensure proper operation.

Shown in Figure 8.16 are the PVCs provisioned for this sample network. A PVC is provisioned between user 1 and NSP 1, between user 2 and NSP 2, and one each between user 3 and NSPs 2 and 3. These PVCs would normally be provisioned when a user signs up for service with a specific NSP. If user 1 decides to terminate service with NSP 1 in favor of NSP 2, the existing PVC between user 1 and NSP 1 must be removed and a new PVC provisioned between user 1 and NSP 2. Note that user 3 has a PVC provisioned to both NSP 2 and NSP 3 and can access service from either entity. The key idea being illustrated here is that to access more than one NSP, multiple PVCs must be provisioned end-to-end. In Figure 8.16, user 1 can access only NSP 1 and cannot access any of the others. Likewise, user 2 can access only NSP 2. In an end-to-end PVC environment, dynamically selecting an NSP, also known as *service selection*, is not possible.

Note in Figure 8.16 that PPP is terminated at the user and at the NSP gateways. Further, the NAP and transport providers play no part in the PPP session. Thus the NSP can control all PPP negotiation parameters, can implement the authentication method and verification, and can issue IP addresses without intervention or knowledge on the part of the NAP or transport providers. This feature is important to the practical business practices of the NSP (an NSP normally would not want to share its private password databases with another party).

Because the end-to-end PVC for each user is provisioned at subscription time and remains available until the user wants to change service providers (which in many cases could be years), no information short of counting actual user cells can be used to determine the actual user's usage of the service. Thus ATM alone provides limited options for usage tracking. An advantage of using PPP over ATM is that PPP can help determine the amount of time a user is actually accessing the NSP. Entries can be made in a database detailing when a user logs on to an NSP's network as well as when a user logs off. This information opens up more options on how billing policies can be implemented (flat monthly fees versus usage-time fees, for example).

In all, an ATM PVC implementation has many benefits including the capability to provide a guaranteed bit rate through the access and transport domains for every user. A drawback of this type of access network is the coordination between NSPs, transport providers, and NAPs and provisioning needs necessary to connect all users to the desired NSPs (one PVC per user per destination). This problem is amplified as the number of users grows and as users change from one NSP to another (often called user churn). Another drawback of an end-to-end PVC architecture is the lack of service selection. If a user wants to access more than one NSP, multiple PVCs must be provisioned for the user, further complicating the provisioning problem.

SVC-Based Model

An SVC model overcomes many of the issues associated with a PVC model. Figure 8.17 shows an end-to-end ATM SVC architecture.

Figure 8.17 An end-to-end SVC access system with PPP over ATM.

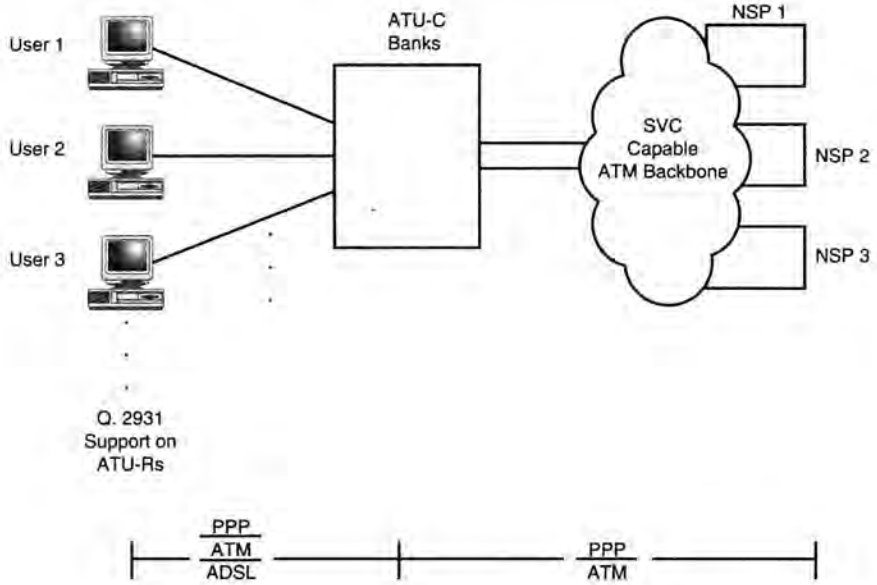

The basic diagram is the same as that shown in Figure 8.16 except that the transport backbone is now SVC capable. Therefore, PVCs do not have to be provisioned at subscription time. Instead, virtual circuits are dynamically created when a subscriber wishes to access an NSP (begin a session) and are torn down when a subscriber ends the session. Service selection is done by a signaling specification known as Q.2931.[13] This signaling method transmits all the information necessary to indicate a desired destination as well as the bandwidth characteristics of the session. Assuming that an adequate virtual circuit can be set up

between the user and the desired NSP, the connection will be made. Service selection in an SVC network is sometimes simply referred to as dialing, as it mimics the functionality of dialing in the telephone system.

Note that the data protocol stacks in Figure 8.17 did not change from those in Figure 8.16. Although PPP is no longer necessary to delineate sessions (because call starts and stops can be recorded), it is still convenient for authentication and address assignment. Thus PPP end-to-end would remain in SVC access networks.

The major drawback of an SVC ADSL access network is that the access node, backbone, and all endpoints must be ATM based and support SVCs. For many different reasons, this requirement is not widely met. For example, some transport networks are based on technologies other than ATM. Also, many existing NSPs do not have ATM-based gateways and for business reasons are not motivated to upgrade their gateways to ATM. Although an end-to-end SVC network is very enticing, true deployment based on it will have to wait until basic upgrades are made to the network.

Tunneling Architectures

In an attempt to bridge the gap between PVC networks and SVC networks, tunneling architectures have been proposed. Figure 8.18 shows a network employing tunneling.

Figure 8.18 A tunneling access system with PPP end-to-end.

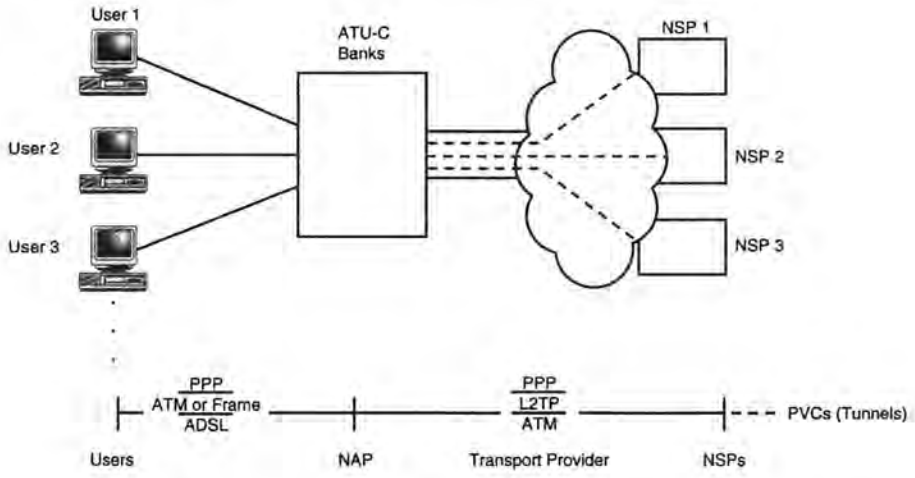

The notable characteristic of this network is that L2TP is present in the transport backbone between the NAP and the NSP. In addition, a single PVC is provisioned between the NAP and each NSP. Each PVC represents a single tunnel. All users' traffic passing through the NAP in Figure 8.18 destined for NSP 1 would be sent on the tunnel between the

NAP and NSP 1. Thus, as compared to a PVC end-to-end network, a single PVC can carry traffic from multiple users. Likewise, all traffic between NSP 2 and the NAP is sent on the tunnel between the two. PPP is again present as an end-to-end protocol and provides the methods for authentication, address assignment, and session delineation. Over the actual access loop, both cell (ATM) and frame (FUNI) implementations are possible.

Users in a tunneling network must use a dialing method to access an NSP. Two dialing methods are possible. The first uses Q.2931 signaling. In this case, the signaling is terminated at the NAP, and the NAP forwards the user's traffic onto the appropriate NSP. The desired destination will be known after decoding the Q.2931 messages. This NAP model differs from the use of Q.2931 in the SVC end-to-end network in which the switches as well as the NSP endpoint participate in the session set up. Upon receiving a Q.2931 request for a connection to an NSP, the NAP (LAC in L2TP terms) initiates a call setup with the NSP (LNS in L2TP terms). Essentially, this call setup assigns a call ID to the user's PPP session so that the session can be uniquely identified at each end of the tunnel.

A second possible method of service selection in a tunneling environment is to encode the desired destination name into the authentication name sent by the user. When a user initiates a PPP session, the LCP negotiation is done by the LAC, not the LNS. A call ID is requested and received for the user after the LCP phase. PPP is "temporarily terminated" at the LAC, and then the PPP endpoint is handed off to the LNS. In addition to LCP negotiation, authentication information may be received by the LAC. Authentication information often contains a user name. If, as a rule, the username is formatted as *user_name@nsp.com*, where *user_name* identifies the user and *nsp.com* identifies the desired NSP, this field can have the dual purpose of identifying the user as well as serving as a dialing string. For example, if the author wanted to access NSP 1, the username field might be `Dennis.Rauschmayer@NSP1.com`. The value in using this type of dialing method is its simplicity compared to supporting Q.2931.

Tunneling networks with some type of service selection method mimic an SVC network without needing ubiquitous support of SVC signaling. Though provisioning in the backbone is still necessary, the number of PVCs needing provisioning scales with the number of NSPs, instead of with the number of subscribers as in the PVC end-to-end case. In addition, churn of PVCs is less likely. Because L2TP is independent of the protocol over which it runs, deployment in backbone networks consisting of technologies other than ATM is possible as well as in hybrid networks. These factors have made tunneling architectures a practical compromise for xDSL access networks.

Tunneling networks have some less attractive features as well. Compared to ATM end-to-end systems, a guaranteed bandwidth can not be maintained for individual subscribers (even if the backbone is all ATM) because multiple subscribers share the same PVC. This

situation can be somewhat overcome by offering differently loaded tunnels between NAPs and NSPs. For example, a regular and a "gold" tunnel where a premium rate is charged for the "gold" tunnel, but the ratio of bandwidth to subscribers is higher than for the regular tunnel. However, the true QoS offered by ATM cannot be matched. Also, L2TP requires that PPP packets be reassembled at the NAP before being placed into the tunnel. This step is inconvenient if cells are used over the loop.

Premise Network Considerations

In the architectures discussed in the previous sections, no mention was made of considerations in the premise network. Several popular options are common in the premise network, including an external ATU-R connected by Ethernet to a computer (or computer network), an external ATU-R connected though an ATM-25 port to a computer, and an internal ATU-R, sometimes called a *network interface card* (NIC).

When a NIC or an ATM-25 connection is used, much of the protocol processing including PPP can be implemented on the host PC. This approach greatly simplifies the ATU-R design as it essentially has only physical layer functionality. Address assignment to the user is straightforward as well because an IP address assigned though PPP is expected. Figure 8.19 shows premise network configurations.

Figure 8.19 Premise network configurations where (a) the ATU-R is connected via an ATM 25 interface and (b) the ATU-R is inside the PC; in both cases, ATM and PPP are terminated inside the PC.

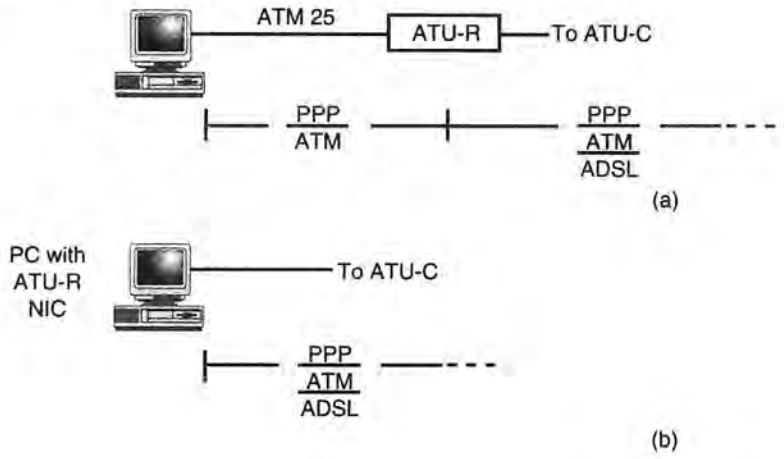

When an ATU-R is connected to a user's computer or computer network via Ethernet, address handling and PPP termination are a bit more involve. This case is illustrated in Figure 8.20.

Figure 8.20 An example of a premise network in which the ATU-R is connected to the PC via Ethernet; in this case, the PPP termination point is in the ATU-R and not the PC.

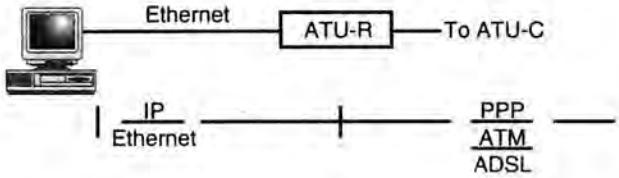

Because Ethernet ports and drivers are usually not equipped to handle PPP, the PPP endpoint is typically found in the ATU-R instead of in the PC. Also, the PC will typically already have an IP address assigned to its Ethernet port before being connected to the NSP. This IP address is sometimes called the *private* IP address. This situation raises a problem because the NSP normally requires the user to take on an IP address consistent with the NSP's address plan. One method to solve this problem is *network address translation* (NAT). In this case, the ATU-R obtains an IP address on behalf of the PC though IPCP. When traffic from the PC destined for the NSP passes through the ATU-R, it substitutes the obtained IP address for the private IP address in the source IP address field. For traffic passing in the other direction, the ATU-R substitutes the private IP address for the obtained IP address in the destination IP address field. When this step is properly done (which includes recalculating IP and TCP checksum values), the operation is transparent to both the PC and the NSP. NAT can run into problems with certain protocols such as file transfer protocol (FTP). Some datagrams sent using this protocol embed the IP address of the user into the packet in ASCII form. For FTP to work through an NAT device, FTP packets must be recognized, decoded, and adjusted to reflect the translated IP address. In addition, some IP security methods do not work with NAT. NAT is discussed in further detail in RFC 1631, "The IP Network Address Translator (NAT)."[14]

Several other methods exist to combat the issues of an Ethernet attached ATU-R. One method involves creating a tunnel through the premise network and extending the PPP session all the way to the PC so that protocol processing can be done locally. Other protocols are also being proposed to handle this situation. No one method of dealing with the premise network dominates, and various methods will continue to exist with the important quality that they all look the same to the NAP.

Summary

The first part of the chapter reviewed the basics of networks and protocols. Then the discussion turned to the way in which these protocols fit together in WANs incorporating ADSL. Several types of architectures, including bridging architectures, ATM end-to-end architectures, and tunneling architectures, were considered, and their advantages and disadvantages were examined. The role of PPP in access networks was discussed along with the many benefits it provides. The chapter concluded with a discussion of issues and options in the premise network. Different common implementations of an ATU-R were described, along with the address-assignment and protocol-processing requirements of each.

Exercises

1. Find RFC1660 on the Internet. Explain the LCP parameters available for negotiation.
2. ARP and Ping commands are commonly available on workstations. Use the Ping command to verify connectivity to various hosts. View the ARP table after each Ping. What type of information is given in the ARP table?
3. Find the IETF specification for L2TP on the Internet. Draw the frame structure for an L2TP packet.
4. Draw the L2TP frame structure of an L2TP packet carrying an IP frame. Do the same for an Ethernet frame.
5. Assume that 10,000 ADSL users are accessing 20 NSPs through a transport provider. If 50 access nodes exist and each user accesses only one NSP, how many PVCs are necessary in an end-to-end PVC model? How many are necessary if each user accesses two NSPs? Four NSPs?
6. Repeat exercise 5 assuming a tunneling model is employed.
7. What is the maximum number of sessions that can be multiplexed over an L2TP tunnel?

Endnotes

1. W. Stallings, *Handbook of Computer Communications Standards, Volume 2: Local Network Standards.* New York, NY: Macmillan, 1990.

2. C. Hornig, "Standard for Transmission of IP Datagrams over Ethernet Networks," RFC 894, 1984.

3. W. Richard Stevens, *TCP/IP Illustrated Volume 1: The Protocols.* Reading, MA: Addison-Wesley, 1994.

4. Gary Wright and W. Richard Stevens, *TCP/IP Illustrated Volume 2: The Implementation.* Reading, MA: Addison-Wesley, 1995.

5. S. Feit, *TCP/IP: Architectures, Protocols and Implementation.* Washington, DC: McGraw Hill, 1993.

6. U. Black, *ATM: Foundations for Broadband Networks*. New York: Prentice Hall, 1995.

7. J. Martin, et al., *Asynchronous Transfer Mode: ATM Architecture and Implementation.* New York: Prentice Hall, 1996.

8. W. Simpson, "The Point-to-Point Protocol (PPP)," RFC 1661. July 1994.

9. G. Mcgregor, "The PPP Internet Protocol Control Protocol (IPCP)," RFC 1332. May 1992.

10. Kory Hamzeh, et al., "Layer Two Tunneling Protocol (L2TP)," Internet Draft draft-ietf-pppext-l2tp-01.txt. October 1996.

11. ADSL Forum, "TR-002: ATM over ADSL Recommendations." March 1997.

12. J. Wojewoda, et al., "Enhancements to Core Networks Architectures for ADSL Access," ADSL Forum Contribution 97-072. June 1998.

13. ITU-T, "Recommendation Q.2931: B-ISDN Application Protocols for Access Signaling." February 1995.

14. K. Egevang and P. Francis, "The IP Network Address Translator (NAT)," RFC 1631. May 1994.

This material may be protected by Copyright law (Title 17 U.S. Code)

CHAPTER 9

VDSL Overview

In this chapter:

- Basic deployment locations of VDSL
- Requirements for VDSL systems
- Proposed modulation methods for VDSL
- VDSL in WAN networks

ADSL satisfies the need to send up to 8 Mbps or so over twisted pairs. However, there is still a need to send more data at even higher speeds. VDSL fills this void by delivering up to 50 Mbps. Unlike ADSL, VDSL offers both symmetric and asymmetric configurations. Typically, VDSL uses far more bandwidth than ADSL uses and has a much shorter reach due to the higher bit rates supported. This chapter covers the basic goals of this relatively new technology and presents some candidate implementations. As with ADSL, some implementations promote DMT, and others use CAP/QAM techniques. The chapter also includes a brief discussion of how WAN networks might implement VDSL.

VDSL Deployment Locations

Like ADSL, VDSL can provide frequency separation from baseband, leaving room for Plain Old Telephone Service (POTS). (As you will see later, in addition to POTS, other services such as ISDN can also be run on the same twisted pair as VDSL.) Thus each end of a twisted pair carrying VDSL and POTS uses a splitter to separate the two signals. At the subscriber's location, the modem terminating VDSL is called the *VTU-R*. The very high speeds of VDSL suggest that some sort of high-speed optical network will feed a bank of centralized VDSL modems. This central point might be a CO of some remote optical

node. For this reason, the peer modem to the VTU-R modem is called the VTU-O (as opposed to VTU-C, which seems to follow from the ADSL notation), denoting an optical feeder.

Note

However, other methods of feeding VDSL banks are certainly possible. For example, VDSL can be used as a means to multiplex and carry multiple T1 streams over a short distance.

Figure 9.1 shows an example of VDSL range and reach from a central deployment location. If the location is a CO, the actual area to which VDSL could be deployed, or the coverage area, would be much smaller than the area served by the CO (typically extending well beyond 12 kft). Limiting deployment to this percentage of the overall serving area severely limits the business potential of VDSL.

Figure 9.1 Range and reach of VDSL from a central location.

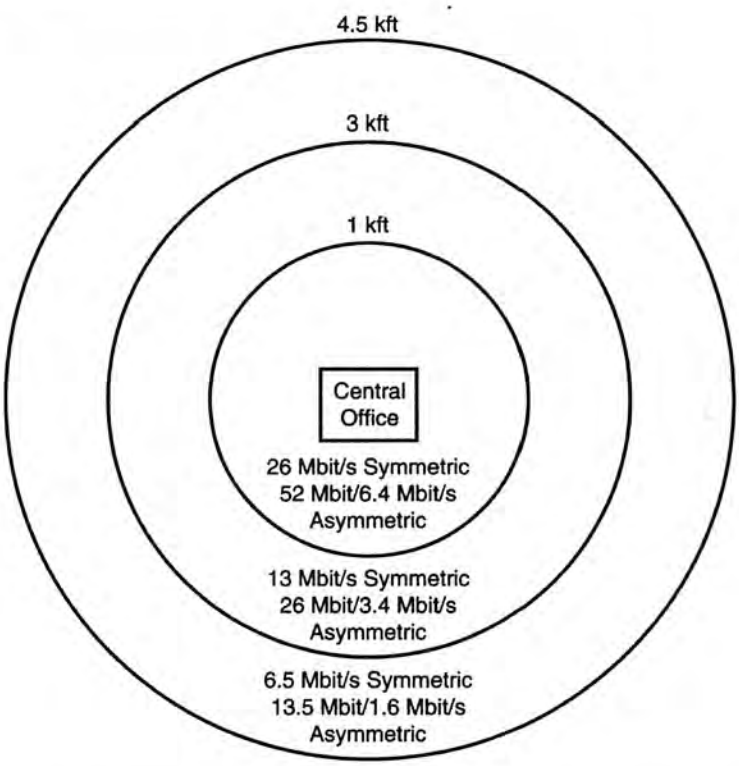

Instead, VDSL deployment will likely take place from the CO as well as from *optical network units* (ONUs). These nodes typically are deployed in, and serve, neighborhoods, industrial parks, and other areas with high telecommunications traffic patterns. ONUs are sometimes used in context when discussing technologies such as fiber to the curb (FTTC),

fiber in the loop (FITL), and fiber to the neighborhood (FTTN). The common feature of these technologies is that high-capacity fiber is deployed to an ONU that serves a small area. The media connecting subscribers to the ONU can be coaxial cable, some sort of wireless link, or more probably, a twisted pair. The combination of a high-capacity link to a serving node, the small area served by each node, and the popularity of twisted pairs connecting subscribers to the serving node make a network utilizing ONUs a perfect fit for a technology like VDSL.

The example in Figure 9.2 shows an area with VDSL deployment using ONUs to serve more remote areas and VDSL deployment from the CO (CO) to serve less remote areas. In the figure, simple optical links serve each ONU. In reality, optical rings or some other type of optical distribution might be more practical.

Figure 9.2 Coverage area of VDSL when deployed from remote ONUs.

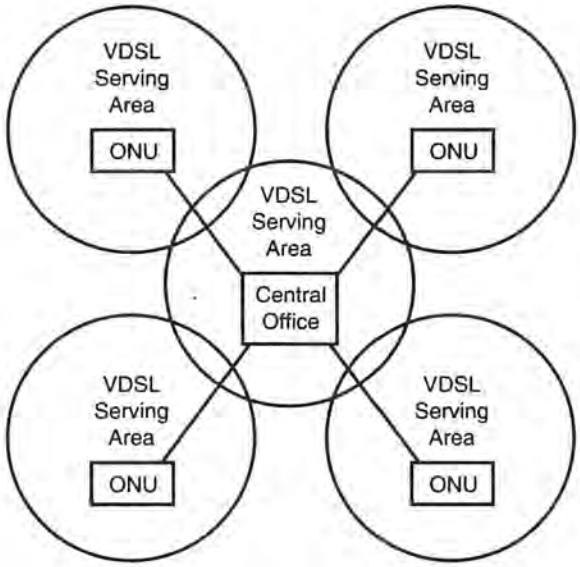

The total optical bandwidth available to an ONU is usually not greater than the sum of potential bandwidths of all ONU users. For example, if an ONU serves 20 subscribers and each has a 50 Mbps VDSL link, the total potential bandwidth for that ONU is 1 Gbps, much more than the ONU would normally be fed. The ratio of the optical bandwidth available to the ONU to the potential aggregate bandwidth of all users is called the *oversubscription ratio*. Oversubscription ratios are engineered so that performance remains reasonable for all users.

Crosstalk Issues

As with FDM ADSL, VDSL does not have self-NEXT. The trend toward technologies with no self-NEXT, or at least with limited self-NEXT, has evolved because of the increased performance achievable and, in some cases, the design simplifications possible with such implementations. Modulation schemes for VDSL and how they avoid self-NEXT is discussed later in this chapter.

Many of the crosstalk analyses in previous chapters are applicable to VDSL. In addition, the short deployment range of VDSL poses several new crosstalk possibilities. Consider the two deployment configurations in Figure 9.3 and Figure 9.4.

Figure 9.3 Crosstalk scenario with VDSL and another DSL technology in the CO mix.

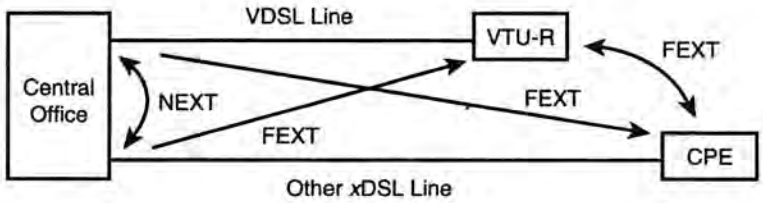

Figure 9.4 Crosstalk scenario with VDSL and another DSL technology in the customer premise (CP) mix.

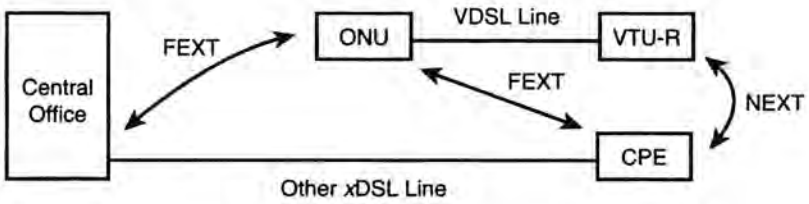

Figure 9.3 is sometimes called CO mix.[1] In Figure 9.3, VDSL and another xDSL technology are both provided from the CO and share a binder group between the CO and some point at which the VDSL is removed. (The VDSL might be terminated at this point, or it might be routed into a different binder group.) The various types of NEXT and FEXT that can exist are also shown in Figure 9.3. The FEXT between the two endpoints is normally not a factor because the distance is usually large enough to attenuate the disturbing signal. However, other FEXT paths may exist as well. Whether these FEXT signals will be significant or not depends on the lengths of the two channels as well as the frequencies used by the upstream and downstream channels.

In the configuration shown in Figure 9.4, the VDSL signal is terminated at a remote ONU and the other xDSL signal at the CO. At one point, the twisted pair segments carrying the signals share a common binder. This configuration is sometimes called the CP mix.[2] Note that it is not necessary for the VTU-R and the other xDSL CPE modem to be located in the same building, but the twisted pairs carrying them must share a binder group for some length near the respective endpoints. Figure 9.4 also shows the possible significant NEXT and FEXT paths for the CP mix configuration. The significance of the FEXT paths depends on the location of the ONU as well as the frequencies used by the channels. Note that the endpoints of the VDSL circuit can contribute either NEXT or FEXT into the endpoints of the other xDSL technology, depending on the deployment configuration. This flexibility adds a new degree of freedom in analyzing crosstalk.

Of particular interest is the coexistence of ADSL and VDSL. In the CO mix configuration, ADSL and VDSL are both provided from the CO and share a binder group between the CO and the point at which the VDSL is removed. (The VDSL might be terminated at this point, or it might be routed into a different binder group.) For no NEXT to exist between these services at the CO, the following two conditions must exist:

- The downstream VDSL signal must not overlap the upstream ADSL signal (the VDSL downstream signal must start higher than 138 kHz).
- The downstream ADSL signal must not overlap the upstream VDSL signal (the VDSL upstream signal must not overlap 138 kHz to 1.1 MHz).

In addition, the downstream ADSL signal may cause a reasonable level of FEXT into the downstream VDSL signal since the distance between the ADSL transmitter at the CO and the VTU-R receiver is small. Likewise, FEXT from the VDSL signal may pose a problem to the upstream ADSL signal.

For no NEXT to exist between the ADSL and VDSL in the CP mix configuration, the upstream and downstream frequency spectrums of the respective technologies must not overlap. Additionally, VDSL FEXT into the downstream ADSL signal might occur, and ADSL FEXT can occur into VDSL depending on the ONU location. Owing to the short loop lengths supported by VDSL, the FEXT, though less severe than NEXT, might have a detrimental impact on the ADSL downstream.

Because of the need to eliminate NEXT between ADSL and VDSL systems and the possibility that FEXT may indeed cause performance degradation, the spectrum allotted to VDSL falls above that used by either the upstream or downstream ADSL channels when ADSL and VDSL are both present. In addition, when dealing with other technologies, the deployment configuration of VDSL must be taken into account to decide which points

contribute NEXT or FEXT to other points. The specific frequency allocations for VDSL are discussed later in this chapter (after other influencing factors are discussed).

RFI Issues

VDSL receivers must deal with the issue of radio frequency interference (RFI). Included in RFI issues are *ingress* and *egress*. The cause of RFI ingress is in-band radio waves from nearby antennae incident upon a twisted pair carrying the VDSL. An amateur radio antenna is a good example of an RFI ingress disturber. Figure 9.5 illustrates this basic situation.

Figure 9.5 RFI ingress into VDSL due to a local transmitter.

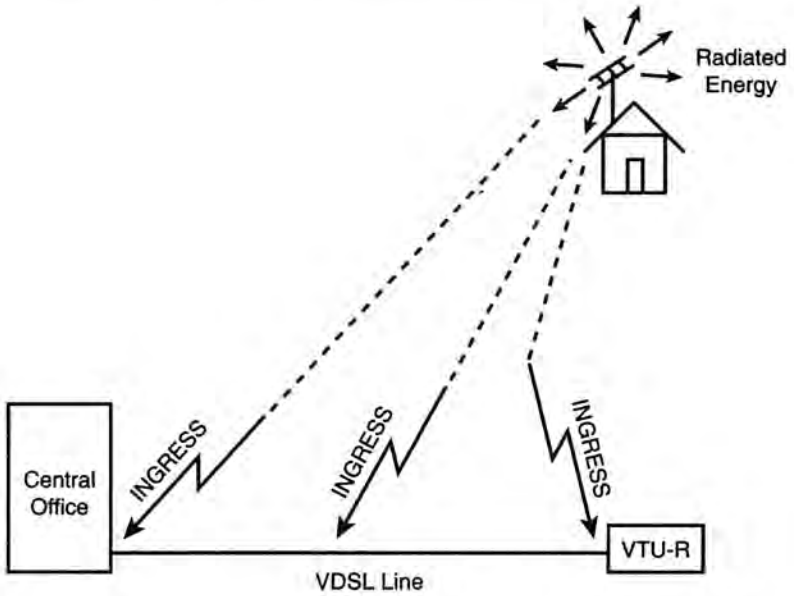

The factors influencing the amount of ingress include the power output of the antenna, the distance between the antenna and the twisted pair, the relative orientation and shielding properties of the binder group, and the balance of the twisted pair itself. Normally, the RFI ingress excites each wire in the twisted pair, thus creating a longitudinal ingress signal on the pair. Because the balance of the twisted pair is not ideal (normally it is 30 dB to 35 dB for high frequencies), some of the ingress leaks into the differential signal.

RFI ingress signals are often very narrow in bandwidth when compared to the VDSL signal. This characteristic is helpful, as the ingress will affect only a small portion of the usable bandwidth.

Another issue is that the ingress signal may be very large compared to the received VDSL signal. In this case, the analog front end of the receiver must be carefully designed so as not to saturate. In addition, some action must be taken so that the analog-to-digital converter (ADC) performs with the proper precision. The issue here is that a large ingress signal will swamp the converter, causing it to waste bits of precision on this unusable signal instead of on the received VDSL signal. Because the VDSL signal is quantized with less precision, the quantization noise, or imprecision due to the conversion process, increases, effectively decreasing the achievable bit rate on the channel. The ADC must have enough range to handle the ingress signal as well as enough precision to adequately quantize the VDSL received signal. Even after quantization, further RFI cancellation may be necessary in the digital domain as documented in Jeong and Yoo's article, "An Efficient Digital RFI Canceler for DMT-Based VDSL," and Wiese and Bingham's article, "Digital Radio Frequency Cancellation for DMT VDSL."[3,4]

To facilitate the design of the ADC, the analog front end of a VDSL receiver might employ circuitry to reduce RFI ingress. Recall that ingress is normally present both longitudinally and differentially on a twisted pair. Figure 9.6 shows a method to use the longitudinal ingress signal to reduce the metallic ingress signal.

Figure 9.6 A method to reduce ingress that uses a longitudinal signal.

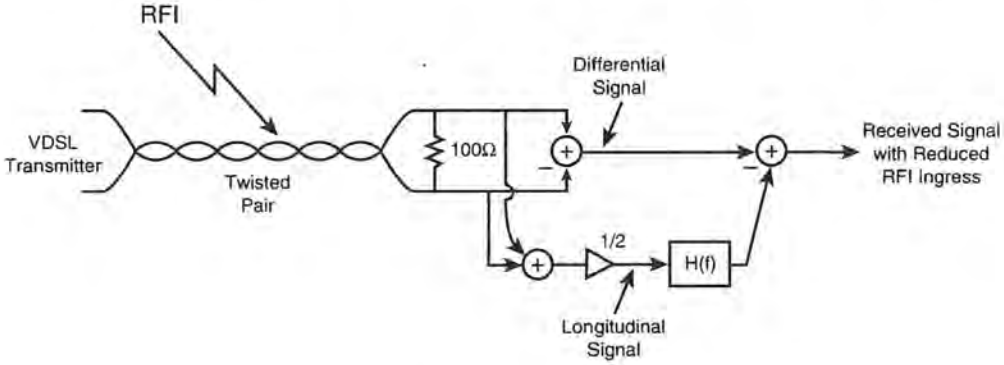

The filter tries to match the unbalance of the twisted pair that is essentially the mechanism by which the ingress reaches the differential signal. If the filter matched the unbalance perfectly, the signal at the output of the summation would consist only of the received VDSL signal (as well as crosstalk and background noise). Figure 9.7 illustrates this concept via a block diagram.

Figure 9.7 A specific implementation of an RFI ingress cancellation circuit.

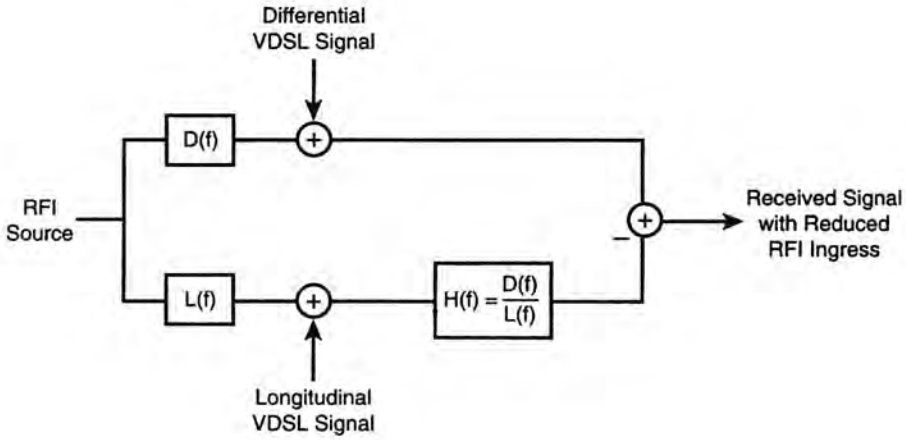

A specific implementation is described in Cioffi, et al.'s article, "Analog RF Cancellation with SDMT."[5] Ingress reduction on the order of 35 dB can be achieved by this method. Note that this would be equivalent to gaining more than five bits of precision in the ADC.

RFI egress is also a concern with VDSL. Egress is illustrated in Figure 9.8.

Figure 9.8 An example of RFI egress.

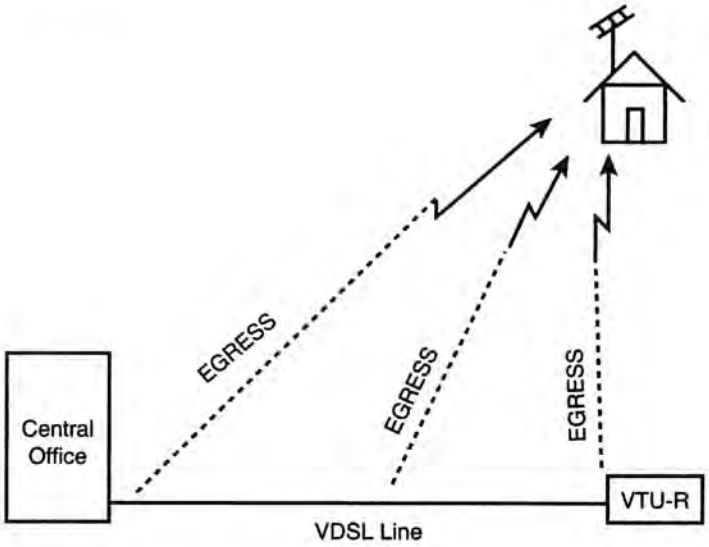

Here the VDSL signal is radiated from the twisted pair and can disturb the signal received by local antennae if these received signals overlap the VDSL spectrum. To combat this problem, the VDSL transmit power in frequency regions reserved for wireless or radio

services must be lowered (these regions were discussed in Chapter 4, "Power Spectral Densities and Crosstalk Models"). Normally, 20 dB reductions in these areas will be adequate to mitigate problems from VDSL RFI egress.

VDSL Requirements

This section discusses the more fundamental requirements put forth for VDSL systems. These requirements are captured in the technical reports "Very High Speed Digital Subscriber Line System Requirements"[6] and "Transmission and Multiplexing : Access Transmission Systems on Metallic Access Cables; Very High Speed Digital Subscriber Line (VDSL); Part I: Functional Requirements."[7] It should be noted that at the time of publication, the requirements were not completely static. Also, some aspects of the candidate implementations in this chapter may not completely satisfy all requirements in the aforementioned reports.

Rate and Reach Goals

In contrast to ADSL, VDSL technology provides either symmetric or asymmetric data rates to a subscriber. The reason that both need to be supported is that some VDSL applications are best served by an asymmetric connection, whereas others are best served by a symmetric connection. Video on demand (VoD), for example, is an application that benefits from an asymmetric connection. Another application, ATM LAN extension, might better be served by a symmetric connection. In addition, applications that do not even exist might turn out to be application drivers for VDSL. Defining symmetric and asymmetric modes of operation up front does not limit the types of application that might one day use VDSL.

Note

This situation developed with ADSL. The initial application driver was VoD; however, Internet/LAN access later became an even stronger application driver.

Table 9.1 shows asymmetric data rate goals for VDSL along with the approximate reach goals for each configuration. How the rates are provisioned (hard configuration as compared to automatically adapting to the loop) is yet to be determined. At the very low end of the asymmetric rates (6.5 Mbps/ 0.8 Mbps), VDSL approaches the performance limits of ADSL. Though it might seem logical to define even lower rates for VDSL and not bother with ADSL, this approach would most probably cause modems to be more complex than necessary and would impact cost effectiveness.

Table 9.1 Asymmetric Data Rate Goals for VDSL

Length Classification	Downstream Rate (Mbps)	Upstream Rate (Mbps)	Reach Goal (kft)
Short	52	6.4	1
Medium	26	3.2	3
Long	13	1.6	4.5
Long	6.5	0.8	6

Table 9.2 shows rate and reach goals for symmetric VDSL operation. Many of these rates are primarily intended to support ATM access systems (for example, 26 Mbps would support 25.6 Mbps ATM adapters along with some overhead). It is expected that some of the symmetric as well as asymmetric rate and reach goals will churn and evolve as VDSL goes through the standardization process. However, Tables 9.1 and 9.2 provide good insight into the probable operating ranges of VDSL.

Table 9.2 Rate and Reach Goals for Symmetric VDSL Operation

Length Classification	Downstream Rate (Mbps)	Upstream Rate (Mbps)	Reach Goal (kft)
Short	26	26	1
Medium	13	13	3
Long	6.5	6.5	4.5

Frequency Allocation

The frequency ranges available for VDSL to use are limited by the need for compatibility with other DSL technologies as well as RFI concerns discussed previously. The base power spectral density (PSD) mask for VDSL appears in Figure 9.9.

Figure 9.9 The basic PSD requirements of a VDSL transmitter.

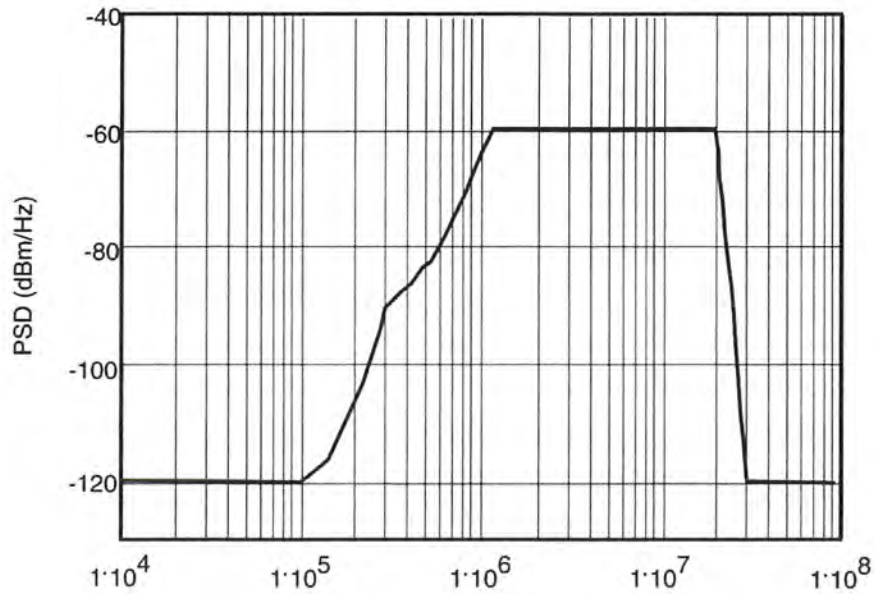

This basic template has a PSD of –60 dBm/Hz from 1.1 MHz to 20 MHz. This range would normally be considered the usable part of the spectrum, and the rest would be considered sidebands. In the ADSL band, the PSD rolls off so as to minimize the crosstalk degradation into adjacent ADSL lines. A POTS or ISDN link could be run on the same line as a VDSL system using this mask, given that a splitter was used to separate the signals.

The mask shown in Figure 9.9 may cause excess egress and interfere with local antennas. In this case, a notching option can be applied resulting in the mask shown in Figure 9.10. Here bands known to be prone to causing egress into existing radio/wireless services are notched down to a level of –80 dBm/Hz (essentially reducing the egress power in the notched areas by a factor of 10). The notched frequency regions can still transmit information if the received SNRs across the regions are favorable.

Note

Of course, if egress is a problem that causes a band to be notched, ingress may likely exist in the same band as well (the local antennas may transmit as well as receive), further degrading the SNR and causing the notched bands to be useless for transmitting data.

Figure 9.10 A VDSL PSD with notches to reduce RFI egress.

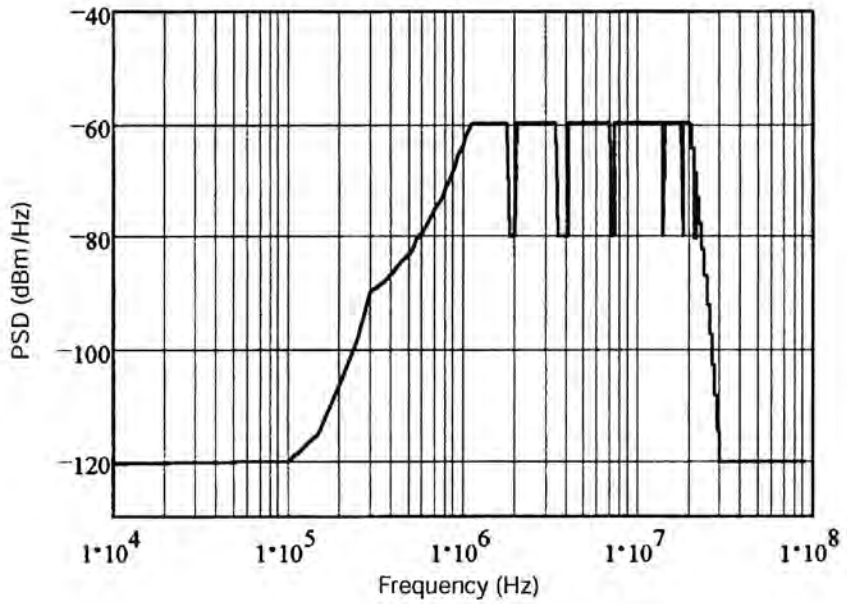

A power boost option has also been proposed for VDSL. When ADSL compatibility is not necessary (perhaps because a service provider did not deploy ADSL), this option would not only boost the VDSL level in the 1.1 MHz to 20 MHz region but also allow the usable end to stretch down to about 300 kHz. Below 300 kHz, the PSD would roll off and eventually intersect with that shown in Figure 9.9 so that ISDN and POTS could be supported.

It may take trial deployments of VDSL to determine which spectral mask properties are more desirable. Real-life deployment may also flag new problems needing to be overcome. In addition, VDSL spectrum requirements may depend on territory. In any case, Figure 9.9 provides at least a starting point to understanding VDSL PSD development.

One final note on the VDSL PSDs discussed is that the masks do not assume anything about symmetric versus asymmetric operation nor upstream and downstream channels. As I discussed previously, implementations are free to use the frequencies within the defined mask in any way they want, although self-NEXT avoidance will likely be necessary for performance reasons. The following sections discuss different proposed implementations of VDSL. You will see that several methods are employed to abide by the PSD requirements and, at the same time, avoid self-NEXT.

Modulation Methods for VDSL

Several different modulation methods have been developed to satisfy the requirements of VDSL. At the time of this writing, VDSL was still very early in the standardization process, and many details of the proposals are yet to be specified (for example, framing content is not yet known). Though many proposals currently exist for VDSL, most fall into two categories: a time-domain duplexing (TDD) category and a frequency-domain duplexing (FDD) category. Different modulation methods within each category exist as well. The ideal method by which standardization can be achieved is to choose a modulation method and fill in all of the implementation details necessary for a standards document. In reality, one or more of the methods are likely to become standardized (or at least implemented by service providers), mainly because of the lack of consensus or compromise on behalf of the standardization body members. The following sections describe some of the candidate implementations.

TDD Approaches

TDD is a half-duplex approach for sending signals between two endpoints. Basically, one modem transmits while its peer modem only receives. After a set amount of time, the process reverses, and the first modem only receives while its peer transmits. This basic concept is shown in Figure 9.11.

Figure 9.11 An example of TDD.

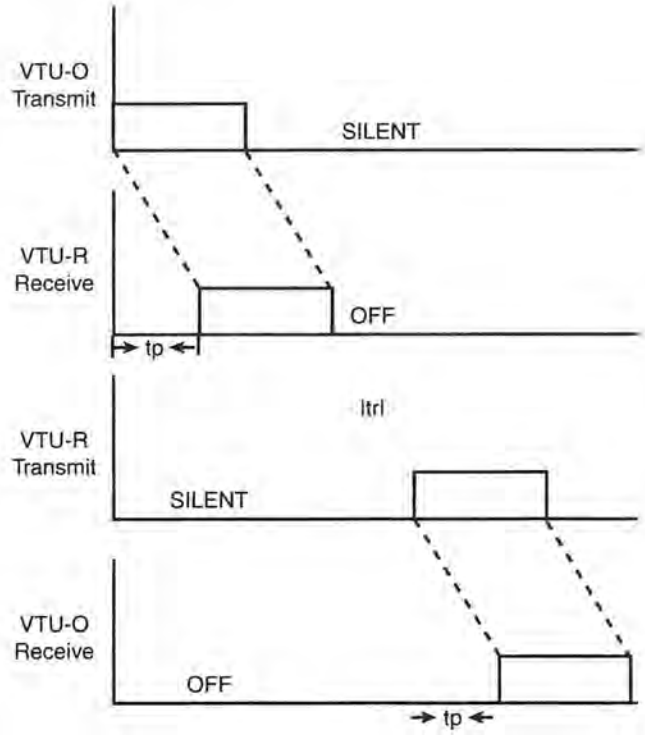

Sometimes this method is known as "ping-pong" for obvious reasons. In each direction, the entire usable frequency band can be used to send information, and thus no separate upstream and downstream PSD need be defined.

The most popular TDD approach proposed for VDSL is the synchronous DMT approach (SDMT). Essentially, this technology employs many of the same fundamental blocks used in ADSL along with the capability to gate the output of the transmitter. (Though many of the same blocks are used, the parameters, such as FFT/IFFT size, will be different.) Figure 9.12 shows a block diagram of an SDMT transmitter and receiver.

Figure 9.12 Block diagram of an SDMT transmitter and receiver.

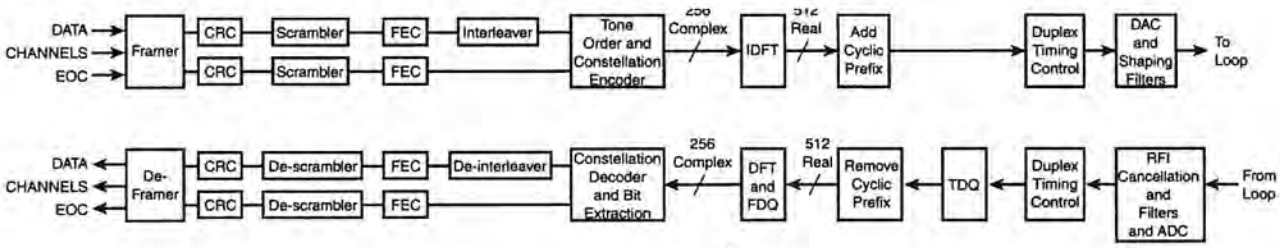

Figure 9.12 is very similar to the DMT-based ADSL diagram discussed in Chapter 7, "ADSL Modulation Specifics." The actual implementation of these blocks will no doubt undergo changes in the early stages of VDSL. A discussion of the blocks in order to explain the main concepts of SDMT follows.

Non-Modulation-Specific Blocks

Some of the blocks in Figure 9.12 should be well understood by now, and implementation details are not necessary for an understanding of SDMT. These blocks include the scrambler and descrambler, the FEC blocks, and the interleaver and de-interleaver blocks. In addition, block implementation is not yet static. Some of the common implementations seen for other DSL systems, including self-synchronizing scramblers and descramblers, Reed-Solomon FEC blocks in GF(256), and convolutional interleaving, are likely implementation candidates.

To a lesser extent, framing and tone ordering/constellation encoding are not modulation-specific blocks. Of course, the tone-ordering and constellation-encoding algorithms must know the number of subchannnels in the modulation path. For SDMT, the number of subchannels supported is 256. As with DMT-based ADSL, some of these subchannels must not be used so that the SDMT output spectrum conforms to the VDSL PSD mask. Because SDMT is a half-duplex modulation scheme, each transceiver will use 256 tones, as opposed to DMT ADSL, where upstream and downstream transmitters used a different number of tones.

SDMT Modulator

The SDMT modulator accepts a vector of 256 complex values from the constellation encoder. Each of these values represents a constellation point for a subchannel or equivalently, the cosine wave magnitude (real part) and sine wave magnitude (imaginary part) of the time-domain waveform produced for a subchannel. The complex conjugates of the 256 complex values are then appended to the vector so that the new vector has complex conjugate symmetry. An IFFT is performed, as given by Eqtn. 9.1.

Eqtn. 9.1

$$x_k = \sum_{i=0}^{511} Z_i e^{\frac{j\pi ki}{256}} \quad \text{for k} = 0..511$$

Here the vector Z represents the vector with complex conjugate symmetry. The output vector of the IFFT contains 512 time-domain values (all real). This IFFT transform is identical to that done in a downstream DMT ADSL transmitter.

Cyclic Prefix Block

The cyclic prefix block adds the last L sample points of the 512-point time-domain vector to the beginning of the vector. L can be any integer less than or equal to 64. The new time-domain vector will be given by (^{x}511–L,^{x}511–L+1,^{x}511–L+2,Ö,^{x}511,^{x}0,^{x}1,Ö,^{x}511). Adding a cyclic prefix makes a channel appear to be circular so that equalization can be done in the frequency domain. The 512 time-domain samples along with the L sample cyclic prefix represent one SDMT symbol. The choice of L is further discussed in the next section.

Note

If the channel's impulse response is longer than the cyclic prefix, the channel will not look circular. A time-domain filter in the receiver must be used in this case.

Superframe Structure

SDMT uses a superframe structure with duration of 500 µs. Thus the superframe frequency is 2 kHz. A superframe consists of a number of downstream symbols followed by a quiet period and then a number of upstream symbols followed by a quiet period. Figure 9.13 illustrates an SDMT frame structure.

Figure 9.13 Sample SDMT frame structure.

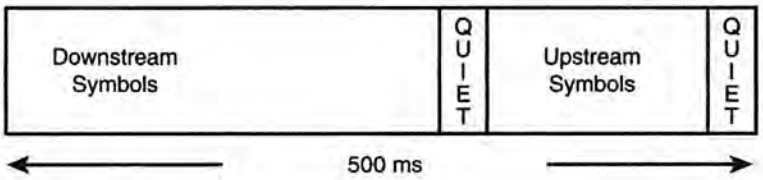

By varying the number of downstream symbols sent and the number of upstream symbols sent, the average upstream and downstream bit rates as well as the bit rate ratio can be controlled.

Consider the examples shown in Figure 9.14. The first shows an 8:1 downstream to upstream ratio, the second a 2:1 ratio, and the third a 1:1 ratio (symmetric example).

Figure 9.14 Possible implementations of an SDMT frame for different upstream and downstream bit-rate ratios.

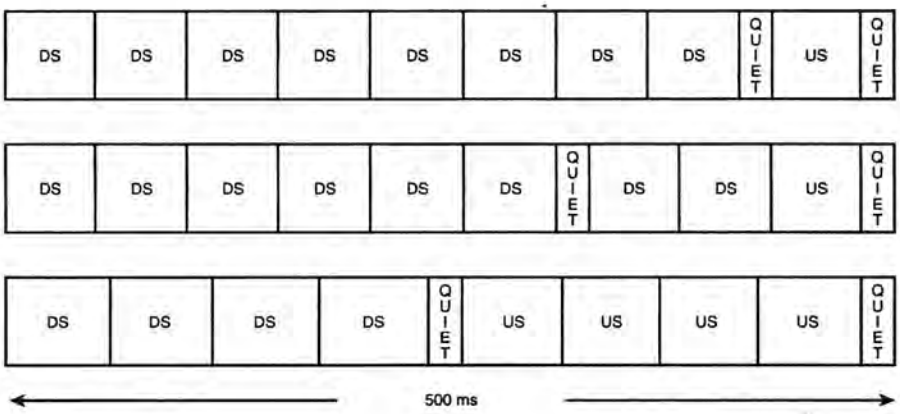

Self-NEXT avoidance is relatively straightforward when lines sharing a binder group use the same symmetry ratio. In such a case, the frame structure of all lines must be aligned so that all VTU-Os transmit simultaneously and likewise receive simultaneously. If no VTU-O transmits while another VTU-O receives, then self-NEXT does not exist. Aligning the superframes of all VDSL lines in a binder group is not trivial, but is certainly possible. For the 2:1 ratio case, self-NEXT elimination is illustrated in Figure 9.15.

Figure 9.15 Self-Next avoidance in an SDMT system.

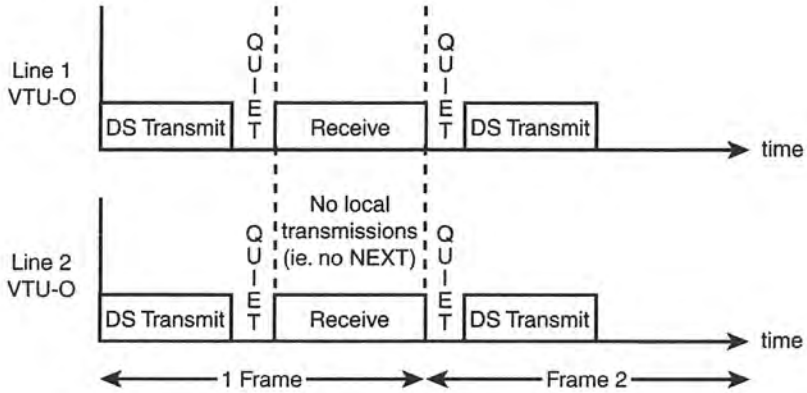

Note

A proposal on how to align the superframes of all VDSL lines in a binder group includes using a common network clock for synchronization, using stuff and delete methods,[8] and even using GPS signals.[9]

When mixing different data rate ratio configurations within a binder group, complete self-NEXT elimination is not possible. However, frames can be arranged to reduce the occurrence of self-NEXT. See Figure 9.16 for an example of limited self-NEXT due to frame alignment of a symmetric and an asymmetric VDSL mix.

Figure 9.16 Two SDMT frames with different upstream-to-downstream ratios aligned to limit the amount of mutual NEXT.

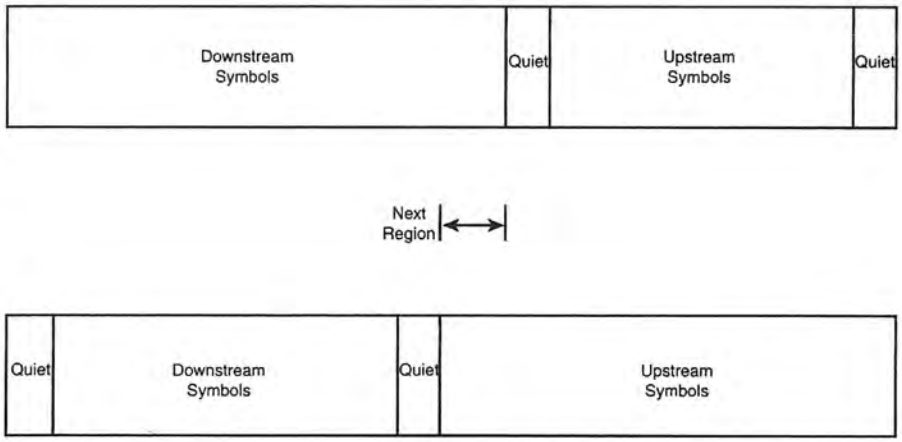

Though it is not possible with this scheme to perfectly eliminate self-NEXT, the self-NEXT levels are predictable and line rates can be adjusted accordingly when self-NEXT is present.

Sampling Considerations

Up to now, the number of samples in a frame has been discussed, as well as the length of a superframe. Not discussed yet is the duration of a sample (or the inverse of that being the sample rate). Following from the sample rate is the number of frames in a superframe, the highest frequency used in the system, and the width of each frequency bin. Remember that all blocks up to the output of the modulator have functionality that does not depend on the output sample rate. Of course, the reason is that the length of a superframe is fixed (500 μs); the higher number of frames per superframe implies that the blocks leading up to the modulator must perform their functions more quickly.

Three sample rates are proposed for SDMT: 11.04 MHz, 22.08 MHz, and 44.16 MHz. For these rates, Table 9.3 shows the maximum operating frequency, bin width, number of samples per superframe, and approximate number of frames per superframe. When more frames exist per superframe, the net data rate in each direction can be higher. Thus short lines supporting the extremes of the required VDSL rates use a sample rate of 44.16 MHz.

Table 9.3 Maximum Operating Frequency, Bin Width, Number of Samples per Superframe, and Approximate Number of Frames per Superframe

Sampling Frequency (MHz)	Maximum Frequency (MHz) (approximate)	Bin Width (kHz)	Samples per Superframe	Frames per Superframe
11.04	5.52	21.5625	5520	10
22.08	11.04	43.125	11040	21
44.16	22.08	86.25	22080	43

For the lowest sample rate, the total number of samples per superframe is 5520. Given that each frame contains 512 base samples along with a cyclic prefix of between 0 and 64 samples, the maximum possible number of frames per superframe (assuming a 40 sample cyclic prefix) is 10, leaving no samples to be used for the quiet period (30.4 us). This scenario is not very practical, as some line turnaround time will be necessary for proper half-duplex operation.

The length of the quiet time can be increased by reducing the number of samples used in the cyclic prefix (which may increase the front-end filtering needs of the receiver) or by

decreasing the number of frames per superframe to nine. The latter method would leave, using the maximum length cyclic prefix, a minimum of 336 samples for the quiet times, or approximately 30.4 us. This time would be divided between the two required quiet periods (though it might not be evenly divided). For the higher sample rates, similar calculations can be done to determine the lengths of the cyclic prefixes and quiet times. It is intended that many of the parameters including the sample rate, the number of frames per superframe, and length of the cyclic prefix and quiet times be negotiated during startup.

RFI Cancellation Circuitry

Generally, the front end of an SDMT receiver contains an RFI cancellation block. As discussed earlier, this block is necessary to combat ingress from local RFI sources. Normally, some type of cancellation is necessary before the analog-to-digital conversion process to adequately digitize the received signal with minimal quantization noise.

Receive Blocks

Much of the functionality in the receive path of an SDMT modem follows from the discussion of the receive path of a DMT modem. Time-domain equalization is done to shorten the channel impulse response to less than the length of the cyclic prefix. An FFT is done on a block of 512 received samples, yielding a complex vector with complex conjugate symmetry. Frequency-domain equalization is performed on the lower 256 values of the FFT output, each complex value representing the received component from a single bin. Constellation decoding and error correction is done before the data is demultiplexed by a framer.

Final Words on SDMT

As a candidate for VDSL, SDMT has several advantages. First, the use of TDD allows for easy control over the upstream and downstream data rate ratio. Second, because SDMT is half-duplex, a hybrid is not needed on VDSL modems, and the sharing of digital circuitry between the transmitter and receiver (the FFT/IFFT circuitry for example) is possible, allowing for complexity reduction. Third, the nature of multicarrier modulation allows for flexibility in dealing with narrowband interference from RFI and can easily reduce egress in a band by simply reducing the power on a tone.

One issue related to SDMT is the requirement to properly synchronize frames for VDSL lines in a common binder group. If all lines emanate from a single ONU, this synchronization might be feasible by using a common clock. However, when independent ONUs feed VDSL into the same binder group, achieving frame synchronization is much more difficult.

FDD Approaches

FDD approaches to transmission are similar to those discussed for ADSL. This approach uses full-duplex transmission and separates upstream and downstream transmissions by using different frequency bands for each. CAP/QAM-based techniques lend themselves nicely to an FDD approach, as the basic spectrum can be positioned anywhere, depending on the hilbert pair filter responses in the case of CAP and center frequency in the case of QAM. A second method discussed here is an FDD method utilizing DMT. This method is sometimes known as Zipper.

Note

FDD approaches to transmission include CAP/QAM-based ADSL and FDM ADSL. Echo-canceled ADSL is partially an FDD approach, but not completely because some of the upstream and downstream bands overlap.

CAP/QAM-Based FDD Systems

CAP and QAM technologies offer another method of implementing VDSL. The basic blocks in the modulation system are similar to those used in CAP/QAM ADSL implementations. For VDSL, the proposed CAP/QAM implementation can run in symmetric mode as well as asymmetric mode. In asymmetric mode, the upstream channel's frequency range starts at about 1.12 MHz and can go up as high as about 2.12 MHz. Different symbol rates are defined as well as different constellation sizes to control the upstream data rate. The downstream channel begins at around 2 MHz and can range up to about 18 MHz. Again, different baud rates and constellation sizes can be used to provide different downstream data rates. Self-NEXT is avoided simply by using different frequency bands for transmission in the two directions.

Symmetrical service is handled in a very interesting way by CAP/QAM-based VDSL. For symmetric operation, two upstream channels exist: one from 1.12 MHz to 1.6 MHz, and one from 11.30 MHz extending up to 19.8 MHz. The downstream channel sits between the two upstream channels and ranges between 2.4 MHz and up to 10.7 MHz. Again no self-NEXT exists because the upstream and downstream channels do not overlap in frequency.

Figure 9.17 shows the various spectrums utilized by CAP/QAM VDSL for symmetric and asymmetric service.

Figure 9.17 Upstream and downstream spectra for CAP/QAM-based VDSL for (a) asymmetric rates and (b) symmetric rates.

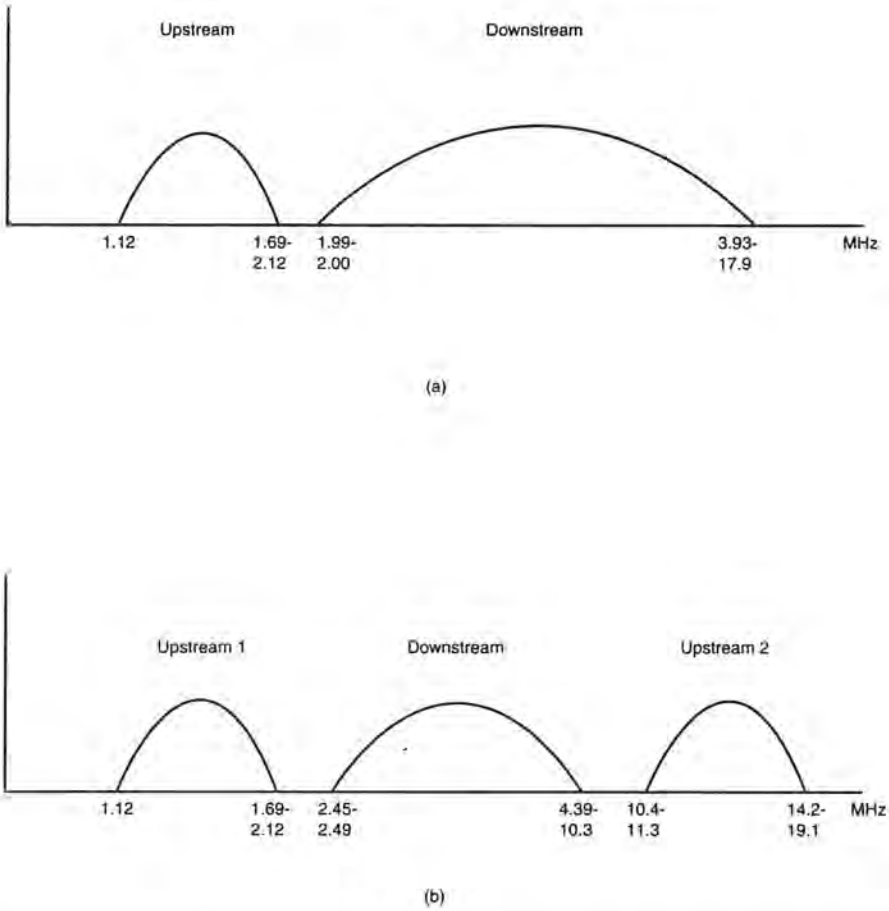

The figure shows different downstream baud rate options (and thus frequency ranges). When mixing symmetric and asymmetric services, in some cases self-NEXT will occur. Note that the highest baud rate downstream carrier overlaps the higher upstream regions in the symmetric frequency plot, which would cause a problem. However, some of the other downstream symmetric options would not, as they end below the higher upstream channel. As CAP/QAM VDSL modems mature, some of the frequencies may change; however, the basic implementation strategy is likely to remain the same.

RF egress control is a bit more difficult for CAP/QAM-based modems. Recall that SDMT-based modems could simply lower the power in a subcarrier at frequencies where power reduction was necessary. CAP/QAM modems must use time-domain filters to do this. Digital filters can normally be implemented for this task, but their flexibility with respect to changes in the power reduced frequencies and the amount of power reduction is somewhat reduced.

DMT-Based FDD Systems

A second FDD method for VDSL uses DMT and is a relatively new modulation scheme named Zipper. Zipper is a time-synchronized version of DMT using 2,048 subchannels going up to 11.04 MHz. Each subchannel has a width of about 5.309 kHz. A subchannel in Zipper can be used for upstream transmission or for downstream transmission, but not for both. If all lines in a binder use a similar bin assignment, then no self-NEXT will be present. (The frames must be time aligned as well so that crosstalk is orthogonal to the received signal.) In addition, subchannel usage can be chosen in a way that minimizes NEXT interaction with other services. Zipper has many of the same benefits as SDMT with regard to ingress and egress handling flexibility. Though a relatively new addition to VDSL implementation, this modulation method may develop into a powerful tool for high-speed access.

VDSL in WAN Networks

The speeds that VDSL supports make it suitable for many types of applications. Many applications that exist today may use VDSL as a transport mechanism; some applications yet to be developed may also use VDSL. A good example of VDSL usage is the full-service network described later in the chapter. The following sections describe video services and data services that utilize VDSL.

Video Services

The high-rate options of VDSL make it a very good access technology for VoD. Figure 9.18 shows a system using VDSL in this application.

Figure 9.18 VDSL used in a network providing video services.

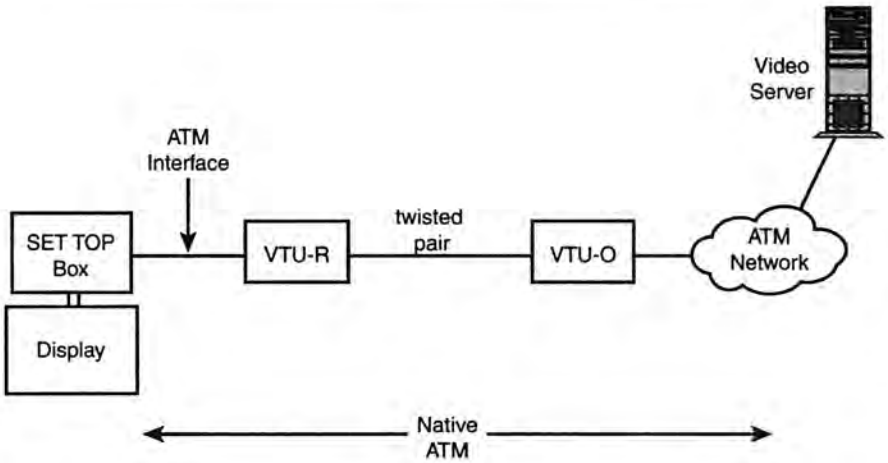

The video server in this example sends digital video over ATM through access equipment to the VDSL loop. At the customer end, some type of video codec converts the digital video to the proper signal for the display screen. The digital video would likely be compressed, and the codec would perform decompression. It is also likely that PPP would not be used on this connection, as it might not be the most efficient protocol for carrying high-speed video. Another control channel might use PPP for authentication and billing.

Data Services

Many data service possibilities exist for VDSL. In the near future, VDSL might be overkill for residential Internet access and Web surfing. However, back-hauling data traffic between locations using twisted pairs might be more reasonable. VDSL could also be used to connect larger offices and corporations to data networks in lieu of fiber connections. As time goes on, more and more opportunities are likely to exist for VDSL to carry data.

The Full-Service Network

Because of the high bit rates supported by VDSL, it is often thought of as an access mechanism to a full-service network (FSN). This type of network is intended to serve all the communication needs of a subscriber, including voice, video, and data applications. Such an all encompassing network could replace today's telephone system and cable system, and it could add extra features such as video telephony. A full-service access network (FSAN), consisting primarily of telephone and service providers, is working to piece together existing standards where available and stimulate new standards to define a FSN.

Summary

This chapter focused on VDSL. The basic needs that VDSL is intended to fulfill were discussed along with several candidate VDSL implementations. The approach differences of these implementations were highlighted, along with some pros and cons of each. Finally, VDSL network applications were discussed. The true direction and application drivers for VDSL are yet to be seen, but any technology promising the delivery speeds of VDSL has great potential.

Endnotes

1. K. S. Jacobsen, "Spectral Compatibility of ADSL and VDSL, Part I: The Impact of VDSL on ADSL Performance," ANSI Contribution T1E1.4/97-404, December 1997.

2. K. S. Jacobsen, "Spectral Compatibility of ADSL and VDSL, Part I: The Impact of ADSL on VDSL Performance," ANSI Contribution T1E1.4/98-035, March 1998.

3. Byung-Jang Jeong and Kyung-Hyun Yoo, "An Efficient Digital RFI Canceler for DMT-Based VDSL," ANSI Contribution T1E1.4/97-422, December 1997.

4. Brian Wiese and John Bingham, "Digital Radio Frequency Cancellation for DMT VDSL," ANSI Contribution T1E1.4/97-460, December 1997.

5. J. Cioffi, et al., "Analog RF Cancellation with SDMT," ANSI Contribution T1E1.4/96-084, April 1996.

6. "Very High Speed Digital Subscriber Line System Requirements," Revision 13, ANSI Contribution T1E1.4/98-043, March 1998.

7. K. Foster, "Transmission and Multiplexing : Access Transmission Systems on Metallic Access Cables; Very High Speed Digital Subscriber Line (VDSL); Part I: Functional Requirements," ETSI TM6 Draft Technical Report, DTS/TM-06003-1, V0.0.7 February 1998.

8. J. Cioffi, "Add/Delete SDMT Solution for Unsynchronized VDSL Lines," ANSI Contribution T1E1.4/96-247, September 1996.

9. Sven-Rune Olofsson, et al., "Synchronizing Optical Networks Using GPS," ANSI Contribution T1E1.4/97-283, September 1997.

MACMILLAN TECHNOLOGY SERIES

ADSL/VDSL PRINCIPLES

A Practical and Precise Study of Asymmetric Digital Subscriber Lines and Very High Speed Digital Subscriber Lines

89065112484
B89065112484A

Asymmetric Digital Subscriber Lines (ADSL) and Very High Speed Digital Subscriber Lines (VDSL) are the future of access technologies in high-speed communication networks. These technologies offer the promise of significantly faster data transfer over existing, copper-based telephone lines. Detailed knowledge of the ADSL/VDSL technologies enables access providers and service providers to plan and evaluate networks involving ADSL and VDSL access. This book provides the communications and networking engineer with practical explanations, technical detail, and in-depth insight needed to fully implement ADSL and VDSL. *ADSL/VDSL Principles* covers the topics that are essential to the successful implementation of these technologies, including the following:

- The fundamentals of the transmission theory and crosstalk in the outside plant, including the details of modeling and simulating the expected performance of ADSL and VDSL under different operation conditions.
- Comprehensive analysis techniques for any type of loop.
- Modulation methods for different types of ADSL and VDSL implementations, including DMT, CAP, TDD, and FDD systems.
- Explanations of network aspects involved in implementing networks with subscribers using DSL access techniques, including protocol issues for ATM networks, Ethernet and LAN networks, and PPP and tunneling techniques.

With the detailed and comprehensive information provided by a noted expert in ADSL/VDSL technologies, the communication and networking engineer will be able to understand and successfully implement ADSL/VDSL technologies. *ADSL/VDSL Principles* provides practical explanations and discussions for the casual reader as well as technical details for the more advanced reader in the following areas:

- The outside plant
- Crosstalk analysis
- ADSL and VDSL modulation
- Protocol schemes for access networks

Dr. Dennis J. Rauschmayer received a B.S. degree (summa cum laude) in electrical engineering from Wilkes University, an M.S.E.E. degree from George Washington University, and a Ph.D. in electrical engineering from Temple University.

Dr. Rauschmayer was a design engineer with Pulsecom, where he designed voice band line cards with automatic gain control and adaptive hybrid equalization to optimize performance against voice band data modems and on hook data transmission devices. At present, he is employed as Principle Scientist, Broadband Copper Access at Pulsecom. His work includes research, design, and development of high-speed copper access devices as well as network design of systems utilizing xDSL access devices.

He is an active member of ANSI Working Group T1E1.4, ITU-T, and the ADSL Forum, making numerous contributions to assist in standardization of ADSL and VDSL technologies.

$44.99 USA/ $64.95 CAN

ISBN 1-57870-015-9

6 19472 00159 2
9 781578 700158 94499

Category: ADSL/VDSL Technologies

www.macmillantech.com