


网络新技术系列丛书 影印版

GIGABIT ETHERNET
NETWORKING

千兆以太网

David G. Cunningham, Ph.D.
William G. Lane, Ph.D.



清华大学出版社
<http://www.tup.tsinghua.edu.cn>



**M
TIP**
MACMILLAN
TECHNICAL
PUBLISHING
U.S.A.

网络新技术系列丛书

影印版

GIGABIT ETHERNET
NETWORKING

千兆以太网

David G. Cunningham, Ph.D.

William G. Lane, Ph.D.

江苏工业学院图书馆
藏书章

清华大学出版社

<http://www.tup.tsinghua.edu.cn>

**M
TIP**
MACMILLAN
TECHNICAL
PUBLISHING
U.S.A.

(京)新登字 158 号

Gigabit Ethernet Networking

By David G. Cunningham, Ph.D. & William G. Lane, Ph.D.

“Authorized reprint from the English language edition published by Macmillan Technical Publishing
Copyright © 1999”

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 retrieval system, without permission from the Publisher.

English language reprint edition published by Tsinghua University Press
Copyright © 1999

This edition is authorized for sale only in the following Territory: The People's Republic of China, with the exception of the Hong Kong and Macau territories.

本书封面贴有清华大学出版社激光防伪标签,无标签者不得销售。

北京市版权局著作权合同登记号: 01-2000-0013

书 名: 千兆以太网

作 者: David G. Cunningham, Ph.D. & William G. Lane, Ph.D.

出版者: 清华大学出版社(北京清华大学学研楼,邮编 100084)

[http:// www.tup.tsinghua.edu.cn](http://www.tup.tsinghua.edu.cn)

印刷者: 清华大学印刷厂

发行者: 新华书店总店北京发行所

开 本: 787 × 960 1/16 印张: 37

版 次: 2000 年 1 月第 1 版 2000 年 1 月第 1 次印刷

书 号: ISBN 7-302-01092-7 / TP·2209

印 数: 0001 ~ 5000

定 价: 49.00 元

出版前言

21 世纪人类面对的将是一个网络化的新时代,网络化程度的高低将是衡量一个国家现代化水平和综合国力的重要标志。考虑到我国广大科技工作者面临着网络技术飞速发展的挑战,我们精选了一些反映网络技术最新发展的、且具有权威性的图书,组成“网络新技术系列丛书(影印版)”,奉献给广大读者。既表达对我国广大科技工作者的一种支持,也是我社为我国实施“科教兴国”的战略应尽的义务。

这套丛书包括:千兆以太网、移动 IP、虚拟局域网、交换式局域网、IP 组播技术、虚拟专用网、网络安全技术以及目录使能的网络等一系列先进技术。由于我们水平有限,希望各界专家和广大读者提出建议和要求,促使这套丛书出得更好。

清华大学出版社

1999.12

About the Authors

David G. Cunningham holds a B.S. in Physics and a Ph.D. in Laser Physics from The Queen's University in Belfast, Northern Ireland. After graduation, he spent two years at British Telecom Research Laboratories before joining Hewlett-Packard Laboratories in Bristol, England, where he is a principal scientist and project manager responsible for Link Technology. David has more than 14 years experience with communication systems design and has personally conducted much of the basic research on the use of optical fibers for longer distance Gigabit transmission. He is cited as a key inventor on many patents and has authored numerous journal papers. He has been involved with the development of LAN/MAN standards for more than 10 years, including serving as a sub-task force editor for the IEEE 802.12 Demand Priority standard. David was given an award by the Gigabit Ethernet task force (IEEE 802.3z) for his critical contributions to the Gigabit Ethernet standard.

William G. Lane holds a B.S. in Electrical Engineering from Stanford University, an M.S. in Electrical Engineering from the University of Southern California, and a Ph.D. in Electrical Engineering from the University of California at Davis. He was a design engineer in industry for eight years before joining the California State University, Chico, Electrical Engineering faculty in 1960. Bill was the founder of their computer science program and later was a co-founder of their computer engineering program. Bill served nine years as department chair, six years as dean, and is currently Professor Emeritus of computer science and computer engineering. He has 32 years of teaching experience and has been a computer and communications systems consultant for both government and industry, including approximately 10 years with one of the larger U.S. Navy laboratories. Most recently, Bill has been a participant in the IEEE 802 standards development process for more than five years, during which he co-authored and served as technical editor of the Ethernet-compatible IEEE 802.12 Demand Priority standard. He has also closely followed development of the Gigabit Ethernet standards.

Dedication

To Shirley, whose belief, support, encouragement, and tolerance enabled me to complete this book. Also, to my son Adam, for the many times I didn't play with him.

David

To Jeanne, who after 52 years of marriage, still has the patience, love, and sense of humor to enable me to take on yet another project before our house remodeling is finished and the landscaping is completed.

Bill

Acknowledgments

Writing this book has been an interesting experience in long-distance communications. David Cunningham lives near Bristol, in the west of England, and Bill Lane lives near Sacramento, in northern California. We couldn't have done it without the help of a number of people on both sides of the Atlantic.

We would like to extend our sincere thanks to Hewlett-Packard's Laboratories in Bristol (HPLB) for their encouragement, time, equipment, and support during the preparation of this book. Special thanks go to Steve Wright, David's manager at HPLB, for making it possible for us to have regular telephone conversations and several face-to-face visits on weekends following David's company meetings in California.

Special thanks also to Del Hanson at HP's Fiber Components Division for believing we could write this book, for putting us in contact with Macmillan, for encouraging us throughout the writing, and for being a brilliant reviewer.

To Jim LeValley of Macmillan Technical Publishing (MTP) for accepting Del's seemingly mad suggestion that we could write the book and then convincing us that we should. Chris Cleveland, development editor at MTP, deserves a great deal of credit for the quality of this book. Amy Lewis and Karen Wachs, also of MTP, helped greatly by taking care of many details. Thanks also for providing part of the funding for Bill to attend the Spring 1998 IEEE 802 plenary meeting.

Ken Deroucher, Bill's dean at CSU, Chico, provided the remainder of the funding for Bill to attend the Spring 1998 plenary meeting and made a processor available for his use while this book was being written.

Chris DiMinico and Charlie Catlett provided excellent technical reviews for the book. Their insight and suggestions really made a fantastic difference to the quality.

Mark Nowell, formerly of HPLB, provided most of the experimental data for the graphs and figures, and Alistair Coles of HPLB supplied electronic versions of the 3B4B, 5B6B code tables in Chapter 8 and the 8B10B code tables in Appendix B. Thanks to both of them for numerous helpful suggestions and discussions.

And finally, to the members of the IEEE 802.3z, the EMBI, the MBI, and the TIA FO 2.2 committees: it was a great privilege working with you all, and we hope this book is a true reflection of your great contribution to Gigabit Ethernet and LAN/MAN technology.

About the Technical Editors

Del Hanson is Principal Engineer in the Fiber Optic Communications Division of Hewlett-Packard Company, responsible for network standards development and related strategic component product planning. He worked on microwave subsystem development at Bell Laboratories and HP prior to starting fiber optic LAN technology development at HP Laboratories in 1974. He has participated in the development of numerous fiber optic LAN link standards and has presented many papers and seminars on fiber optic networks and link performance. He received a B.S. in electrical engineering from the University of Wisconsin and a Ph.D. in electrical engineering from the University of Michigan. Dr. Hanson is a Senior Member of IEEE and was co-chairman of the 1987 Optical Fiber Communication (OFC) Conference. He was chairman of the OFC Networks and Switching Subcommittee and the IEEE 802.3z Gigabit Ethernet Optical PMD Link Subcommittee.

Charlie Catlett is Chief Technology Officer for the National Center for Supercomputing Applications (NCSA), located at the University of Illinois at Urbana-Champaign. He was a member of the original NSFNET backbone project in 1987 and has designed, built, and supported large wide area, campus, and local area networks since the early 1980s. During the

past decade, he has worked on the DARPA/NSF Gigabit Testbeds initiative and the NSF vBNS project, and is a founding principal investigator for the National Laboratory for Applied Network Research. During the last several years, he has focused on long-range strategic planning for NCSA in supercomputing, storage, component software technologies, information security, and networking. Charlie holds a B.S. in computer engineering from the University of Illinois at Urbana-Champaign.

Christopher T. DiMinico is the Director of Network Systems Technology with Cable Design Technology (CDT) Corporation. Christopher is a member of the IEEE, TIA TR41.8.1, and the U.S. advisory group for international cabling standards development.

INTRODUCTION

The primary focus of this book is Gigabit Ethernet and its use in multispeed LANs. The Ethernet protocol quickly rose to the forefront when it was first proposed and has since become the preferred LAN protocol for computer users. This does not mean that there has not been competition. A number of other LAN technologies, including Token Ring, FDDI, 100VG-AnyLAN, and most notably ATM, have all attempted to eclipse Ethernet within the LAN marketplace. However, they have all been spectacular in their failure to achieve that goal. Ethernet is now used on 80 to 85 percent of the world's LAN-connected PCs and workstations because of its demonstrated capability to respond to the ever-increasing need for more and more bandwidth.

The Ethernet protocol has been able to successfully incorporate a number of new capabilities because its developers, manufacturers, and standards committees have remained true to the initial Ethernet goals and design principles of simplicity, interoperability, robustness, and low cost. In short, Ethernet has remained viable precisely because it has been adapted to meet the changing needs of its customers while remaining easy to use, easy to install, easy to maintain, and economical to own.

Gigabit Ethernet is the most important recent addition to the Ethernet protocol. The capability for priority-based transmission at 1,000 Mbps in both directions (full-duplex operation) over a single link plus associated improvements in network addressing methodology (VLAN tagging) and in switching/routing hardware will ensure that Ethernet remains well ahead of other LAN technologies.

Gigabit Ethernet has a very important role to play in existing LANs that are implemented with a combination of twisted pair copper and multimode optical fiber cabling. However, transmission through these cables at Gigabit rates presented some very challenging problems and at one point even threatened the Gigabit Ethernet goal of being able to operate over long-established standard link lengths. In contrast with other books on Gigabit Ethernet, we have not shied away from these issues, but have included in-depth discussions of both the problems encountered and their solution development.

We took the time to write this book because we passionately believe that you should understand the recent capability additions to Ethernet: why they were developed, how they work, what the cabling limitations are, how Gigabit Ethernet fits into existing Ethernet (and other) LANs, how it should be installed, and how all of this will affect future LAN/MAN development.

Like other books on Gigabit Ethernet, we have included a discussion of the protocol, but at greater depth than most, and we have also included discussions on the following:

- Layered architectures (a concept that we have found is not well understood by many of the people in the field)
- The concepts and nuances of signal transmission through optical fiber (explained in a diagrammatic way that does not require higher level mathematics to be understood)
- Systems considerations and practical planning methods for new networks and network upgrades (how you would actually develop a network plan)
- Network topologies that provide improved chances for rapid recovery in case of equipment or link failure (where and why you should install alternate paths and redundant equipment or cables)
- Cable plant installation and management (including a few anecdotal examples of what not to do)
- Integrating Gigabit Ethernet with other LAN technologies (such as ATM)

These are all areas where experience has shown that there is a definite need, but that are often not well-covered in technical books on the Ethernet protocol.

Our Approach

Several people we talked to when we first considered writing this book told us that they didn't need "just another Gigabit Ethernet book." What they did need was a book that contained technical information in sufficient detail to provide them with the background knowledge that would allow them not only to understand the current state of the protocol, but to be better prepared for future changes, as well.

Our approach throughout the book has been to provide an essentially non-mathematical treatment to each subject and to use conceptual diagrams and graphs to explain what is going on in each part of the system. For example, the coverage of baseband transmission through optical fiber links is very diagrammatic and, based on reactions from people who have seen it, should be readily understandable, even to readers not previously familiar with the physics of light transmission through optical fiber.

Our Target Audience

This book was written for several broad audiences, including but not limited to:

- Networking systems product developers: engineers, computer scientists, and marketing personnel working for companies that design and manufacture network equipment (such as switches and routers) and components (such as laser transceivers, LAN interfaces, networking software, and optical fiber cables and connectors)
- Network implementers: network planners, managers, buyers, installers, and maintenance personnel who work for organizations that currently need, or soon will need, to upgrade or better understand their particular networks
- Students and faculty in electrical and computer engineering and computer science for coursework at the senior or first-year graduate level

While we assume that you may be familiar with LANs and Ethernet, we have provided enough information in the book to enable you to understand all aspects of Gigabit Ethernet networking without prior experience with high speed digital communications and network bridges, routers, and switches.

In addition, this book can be used as a companion to the Gigabit Ethernet standards documentation. In this role, it contains a lot of previously unpublished information that provides context and will help you understand why the specification is how it is.

Organization

The book is organized in six parts, in a fashion that allows you to skim or even skip individual sections or chapters, depending on your interest or need at the moment. We recommend, however, that you review the table of contents to determine the topics included in each chapter before you decide to go on. Backward and forward references are included throughout the text to help you find related information on specific topics.

Part I: Introduction

Chapter 1, “Ethernet Development: The Need for Speed,” documents the creation, development, and evolution of Gigabit Ethernet networking. This chapter also introduces the CSMA/CD transmission and reception procedures and provides an overview of the components and topology of a Gigabit Ethernet network.

Chapter 2, “The OSI and Gigabit Ethernet Standard Reference Models,” introduces both the ISO and IEEE 802.3 Ethernet reference models as a precursor to detailed discussion

of layered architectures and the logical requirements of the interfaces between the different protocol layers of the model.

Part II: Network Access and Control

Chapter 3, “Media Access Control,” details the Ethernet media access control (MAC) layer functions and responsibilities, the MAC frame format and format options, frame generation and transmission in both full- and half-duplex mode (including collision filtering, carrier extension, and frame bursting in CSMA/CD mode), and frame reception and error control.

Chapter 4, “Gigabit Repeaters, Bridges, Routers, and Switches,” introduces reference models for Gigabit repeaters and bridges and describes functions/responsibilities of repeaters, bridges, routers, and switches, including port buffering, speed handling, and flow control.

Part III: Transmission Fundamentals

Chapter 5, “Fundamentals of Baseband Transmission,” is essentially a non-mathematical introduction to baseband digital transmission. All key concepts required to understand the underlying thinking, design procedures, and trade-offs that led to the final specifications for the 1000BASE-X Gigabit Ethernet physical layers are covered.

Chapter 6, “Fundamentals of Fiber Optic Communication,” is also a self-contained, essentially non-mathematical tutorial that introduces the basic optical fiber definitions, concepts, components, and technology relevant to Gigabit Ethernet. This chapter provides important background information for anyone implementing optical fiber networks.

Part IV: The Common 1000BASE-X Physical Layers

Chapter 7, “The Common Physical Sublayers: Reconciliation and the GMII,” introduces the Gigabit Media Independent Interface that provides a simple, inexpensive, and easy-to-implement interconnection between the Gigabit Ethernet media-independent layer and the media-dependent physical layers, and between the physical layers and the network management processor.

Chapter 8, “Physical Coding, Physical Medium Attachment, and Auto-Negotiation for 1000BASE-X,” introduces the concepts and operation of the Physical Coding Sublayer (PCS), Physical Medium Attachment Sublayer (PMA), and Auto-Negotiation unit that are common to all 1000BASE-X systems.

Part V: The 1000BASE-X Media Dependent Layers

Chapter 9, “The Gigabit Ethernet Optical Link Model,” documents the Gigabit Ethernet optical link model, a tool developed by the Gigabit Ethernet task force to aid development of the 1000BASE-X physical sublayer specifications. The link model is currently (1999) the state of the art as far as laser-based data communication standards are concerned.

Chapter 10, “The Gigabit Ethernet Modal Bandwidth Investigation,” discusses the resolution of some long-standing fundamental issues regarding laser-based multimode optical fiber transmission by Gigabit Ethernet: modal noise and unpredictable bandwidth performance.

Chapter 11, “1000BASE-X: Optical Fiber and Copper PMDs,” discusses the functions, components, and specifications for baseband Gigabit Ethernet optical and short-haul copper PMDs. The Gigabit Ethernet optical link model is used with techniques discussed in Chapter 5 to illustrate and derive the PMD specifications.

Part VI: Network Installation and System Considerations

Chapter 12, “The Cable Plant: Installation and Management,” introduces basic concepts of structured cabling relevant to Gigabit Ethernet and discusses installing conduit and cable. As-built drawing and database content requirements for effective cable plant management are also defined.

Chapter 13, “Upgrading Ethernet LANs: System and Topology Considerations,” reviews the Ethernet system and compatibility considerations and defines planning procedures for implementing Gigabit Ethernet in network upgrades and/or new networks.

Chapter 14, “Gigabit Ethernet in Context with Other LAN Technologies,” reviews some non-Ethernet technologies currently deployed in LANs (Token Ring, FDDI, 100VG-AnyLAN, and ATM) and then considers the use of Gigabit Ethernet as a network upgrade.

Chapter 15, “The Future: Gigabit Ethernet and Beyond,” begins with a review of 1000BASE-T Gigabit Ethernet and link aggregation. It then moves on to discuss the use of Gigabit Ethernet in MAN/WAN networks and as part of WDM-based virtual MANs/WANs. Ten Gigabit Ethernet is also discussed.

Reader Guidelines

We are very aware that not everyone needs, nor should they be expected, to know and understand all the information that is contained in each chapter of this book. We also realize that while the table of contents provides some insight into each chapter's contents, a few suggestions might also be in order.

First, because development of 1000BASE-T (Gigabit Ethernet over four pair Category-5 or better, unshielded twisted pair copper cable) was not complete until most of the chapters had been written, we included it in Chapter 15. We recommend that you go to section 15.1 after you have completed Chapter 11, "1000BASE-X: Optical Fiber and Copper PMDs." The 1000BASE-T and 1000BASE-X physical sublayers are very different from each other, and should be considered in sequence.

Second, the following table suggests possible approaches for several different categories of readers.

Reader Work Assignment or Interest	Chapters Containing Background Material	Chapters Containing Essential Material
Network consultants, planners, and implementers	1, 2	3-15
Network managers	1, 2, 7-11	3-6, 12-15
Network installers and maintenance personnel	1-4, 7-11, 14-15	5-6, 12-13
Designers of network equipment and components	1, 12	2-11, 13-15
Components designers (for example, network interface cards, transceivers)	1, 12, 14	2-11, 13, 15
Cable and cable components designers, (for example, cable, connectors, cross-connect panels)	1-4, 7-8	5-6, 9-15
Electrical/Computer Engineering students	1	2-15
Computer Science students	1, 8-11	2-7, 12-15

Course Materials

Bill Lane is planning to use this book as the text for a High Speed Data Communications Networks course for Computer Science and Electrical/Computer Engineering majors (senior/graduate level) during the fall of 1999. Course materials will be available after the end of the course and will include a syllabus, suggested assignments, and additional discussion topics. You can receive this material by sending an email to *bill_lane@ieee.org*.

Contacting the Authors

If you have comments or suggestions about the book, you may contact either of us by email or at the address shown below. While we cannot guarantee an immediate response to questions and comments, we will certainly try to get back to you in a timely fashion.

William G. Lane, Prof. Emeritus
Dept. of Electrical and Computer Engineering
California State University, Chico
Chico, CA 95929
bill_lane@ieee.org

FOREWORD

Over the past five years, we have witnessed two dramatic improvements in LAN technology. Ethernet, already well-established as the world's most popular network, was reborn as Fast Ethernet. The idea of providing a tenfold increase in bandwidth while retaining all of the important characteristics of a mature and well-understood technology proved to be an irresistible combination for a market that perpetually demands ever faster and ever larger capacity for information exchange. Thus, when Fast Ethernet appeared on the scene in 1995, it rapidly became the network of choice.

This phenomenon hardly went unnoticed. Indeed, with many other technologies contending for the honor of succeeding Ethernet, every step of the evolution of Fast Ethernet was monitored and reported with interest and no small amount of sensationalism. The spectacular success enjoyed by those who invested in Fast Ethernet attracted the attention of the venture capital community, which was smarting from the poor returns it had received on investments in other, more complicated, networking technologies. Fast Ethernet quickly acquired a golden reputation as a simple and reliable technology, with broad appeal and tangible benefits. It required a fairly small investment, both to develop and to deploy, yet it yielded dramatic returns. Whether you were a venture capitalist putting money into a startup company, or a Management Information Systems manager purchasing the network infrastructure for your enterprise, Fast Ethernet was the smart place to put your money.

And then, just five months after the Institute of Electrical and Electronic Engineers completed its work on the formal standard for Fast Ethernet, they began work on a standard for Gigabit Ethernet. Against the backdrop of Fast Ethernet's success, the ordinarily dull and dry exercise of writing an IEEE standard became a closely watched event. Once again, the ink and the money started flowing. In the space of just one year, twenty startup companies were funded to develop Gigabit Ethernet products. The founders of these companies became celebrities, and the companies achieved lofty valuations on the basis of little more than the market potential of Gigabit Ethernet.

While all this attention raged about them, a relatively small group of engineers applied themselves to the task of defining just what Gigabit Ethernet was, and how it would work. This

group banded together under the auspices of the IEEE 802.3 Working Group, the committee responsible for the development of Ethernet standards. For the most part, they were volunteers, taking time away from their normal jobs to work side-by-side with engineers from rival companies. The standards writing process took two and a half years from start to finish, and along the way there were numerous technical challenges to address and resolve. This process took place under the harsh glare of attention from the industry press, and the constant pressure to deliver the standard on schedule. This last point is significant in that few, if any, standards projects commit to and publish a target schedule for completion.

No one took on a larger responsibility or produced a larger contribution to the IEEE standard for Gigabit Ethernet than Dr. David Cunningham. At a critical point in the development of the standard, Dr. Cunningham assumed the leadership role in addressing a critical technical issue which threatened to delay the progress of the standards project. This issue was referred to as Differential Mode Delay, and it was resolved by the Modal Bandwidth Investigation, which was chaired by Dr. Cunningham. By applying his extensive knowledge of laser physics and fiber optics, and his exceptional leadership and communication skills, Dr. Cunningham became one of the heroes of Gigabit Ethernet, for which the industry owes him profound thanks.

Together, Dr. Cunningham and Dr. William Lane have authored a definitive work on the subject of Gigabit Ethernet. Their long experience with the development of standards in the IEEE 802 LAN/MAN Standards Committee gives them unique insight into the inner workings of this exciting technology. Both of these gentlemen are acknowledged experts in the development of networking standards, and both share an insider's view of the process by which an industry standard turns a concept into a formal specification. *Gigabit Ethernet Networking* provides more than just a comprehensive explanation of the contents of the IEEE 802.3z standard for Gigabit Ethernet. It provides the background knowledge necessary to understand how Gigabit Ethernet works, why it works, and where to use it to the best advantage.

Technical professionals at all levels will benefit from reading this work. Whether one seeks to learn the physical principles that underlie the technology of high-speed fiber optic communications, or wishes to acquire a detailed knowledge of the protocols and algorithms specified in the IEEE 802.3z standard for Gigabit Ethernet, or desires to understand how Gigabit Ethernet can and should be deployed in a modern data communications network infrastructure, *Gigabit Ethernet Networking* by Dr. Cunningham and Dr. Lane is the authoritative source of information.

Howard M. Frazier, Jr.

Former Chairman, IEEE P802.3z Gigabit Task Force