FRONTIERS IN COMMUNICATIONS

Advances in Local Area Networks

Edited by
Karl Kümmerle
John O. Limb
Fouad A. Tobagi

IEEE PRESS

Published for the IEEE Communications Society by the IEEE PRESS

FRONTIERS IN COMMUNICATIONS

Advances in Local Area Networks

Edited by

Karl Kümmerle
IBM Research Laboratory, Zurich, Switzerland

Fouad A. Tobagi
Stanford University

John O. Limb
Hewlett Packard Laboratories, Great Britain

IEEE PRESS

COMSOC IEEE

Published for the IEEE Communications Society by the IEEE PRESS

The Institute of Electrical and Electronics Engineers, Inc., New York

IEEE PRESS

1987 Editorial Board

R. F. Cotellessa, *Editor in Chief*

J. K. Aggarwal, *Editor, Selected Reprint Series*

Glen Wade, *Editor, Special Issue Series*

James Aylor
F. S. Barnes
J. E. Brittain
B. D. Carrol
Aileen Cavanagh
D. G. Childers
H. W. Colborn
J. F. Hayes
W. K. Jenkins
A. E. Joel, Jr.
Shlomo Karni
R. W. Lucky
R. G. Meyer
Seinosuke Narita
J. D. Ryder
A. C. Schell
L. G. Shaw
M. I. Skolnik
P. W. Smith
M. A. Soderstrand
M. E. Van Valkenburg
John Zaborsky

W. R. Crone, *Managing Editor*
H. P. Leander, *Technical Editor*
Laura Kelly, *Administrative Assistant*

Carolyne Tamney, *Publications Manager*
Valerie Cammarata, *Associate Editor*

Copyright © 1987 by

THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.
345 East 47th Street, New York, NY 10017-2394

All rights reserved

PRINTED IN THE UNITED STATES OF AMERICA

IEEE Order Number: PC02105

Library of Congress Cataloging-in-Publication Data

Advances in local area networks.

(Frontiers in communications)
Includes index.
1. Local area networks (Computer networks)
I. Kümmerle, Karl. II. Limb, J. O. (John O.)
III. Tobagi, Fouad A. IV. IEEE Communications
Society. V. Series.
TK5105.7.A38 1987 004.6′8 87-4220
ISBN 0-87942-217-3

The Contributors

J. L. Adams
British Telecom Research Laboratories

A. Albanese
Bell Communications Research, Inc.

Flaminio Borgonovo
Politecnico di Milano

Peter I. P. Boulton
University of Toronto

Thomas H. Brunner
Swiss Federal Institute of Technology

Werner Bux
IBM Corporation

H. Che
Dana Computer, Inc.

Felix H. Closs
IBM Corporation

Ronald C. Crane
3Com Corporation

Yogen K. Dalal
Metaphor Computer Systems

John D. DeTreville
Digital Equipment Corporation

Jeremy Dion
Cambridge University

Robert A. Donnan
IBM Corporation—Consultant

Ken-Ichi Donuma
Toshiba Corporation

R. M. Falconer
British Telecom Research Laboratories

Marion R. Finley, Jr.
University of Montreal

Lois E. Flamm
Bell Communications Research, Inc.

C. Flores
AT&T Bell Laboratories

Alexander G. Fraser
AT&T Bell Laboratories

Luigi Fratta
Politecnico di Milano

Yohji Fujii
Nippon Telegraph and Telephone Corporation

Timothy A. Gonsalves
Worcester Polytechnic Institute

Larry Green
Communication Machinery Corporation

Davide Grillo
Fondazione Ugo Bordoni

J. P. Haggerty
AT&T Bell Laboratories

Bill Hawe
Digital Equipment Corporation

Jeremiah F. Hayes
Concordia University

Peter L. Heinzmann
Swiss Federal Institute of Technology

Gregory T. Hopkins
Ungermann-Bass, Inc.

Andrew Hopper
University of Cambridge

Daniel E. Huber
Swiss PTT

Willi Huber
Swiss Federal Institute of Technology

V. Bruce Hunt
Forthright Systems, Inc.

Takao Ito
Toshiba Corporation

Dittmar Janetzky
Siemens AG

J. Richard Jones
FiberLAN, Inc.

Heinz J. Keller
IBM Corporation

Yukio Kimura
Nippon Telegraph and Telephone Corporation

Alan Kirby
Digital Equipment Corporation

P. A. Kirslis
AT&T Bell Laboratories

Felix Kugler
Swiss Federal Institute of Technology

Karl Kümmerle
IBM Research Laboratory

Chin-Tau A. Lea
Georgia Institute of Technology

E. Stewart Lee
University of Toronto

John O. Limb
Hewlett-Packard Laboratories

G. W. R. Luderer
AT&T Bell Laboratories

W. T. Marshall
AT&T Bell Laboratories

James S. Meditch
University of Washington

Norman B. Meisner
Linkware Corporation

Heinrich Meyr
Technical University of Aachen

James G. Mitchell
Xerox Corporation

Hans R. Mueller
IBM Corporation

Shigeru Ohshima
Toshiba Corporation

Yoshinori Oikawa
Nippon Telegraph and Telephone Corporation

David L. Oster
Communication Machinery Corporation

Daniel Avery Pitt
IBM Corporation

David D. Redell
Digital Equipment Corporation

Jeffrey W. Reedy
FiberLAN, Inc.

Martin Reiser
IBM Corporation

John F. Shoch
Asset Management Company

此为试读，需要完整PDF请访问：www.ertongbook.com

W. David Sincoskie
Bell Communications Research, Inc.

Walter Steinlin
Swiss PTT

Bob Stewart
Xyplex Corporation

Hisayoshi Sugiyama
Nippon Telegraph and Telephone Corporation

Kian-Bon K. Sy
IBM Corporation

Dale Taylor
Communication Machinery Corporation

Fouad A. Tobagi
Stanford University

Nobuyuki Tokura
Nippon Telegraph and Telephone Corporation

Kym S. Watson
Fraunhofer-Institute of Information and Data Processing

Peter J. Wild
Hasler Ltd.

Robin C. Williamson
IBM Corporation

Willy Zwaenepoel
Rice University

Preface
Advances in Local Area Networks

LOCAL AREA NETWORKS (LAN's) are still evolving rapidly. While rings were first invented in the late 1960's it was not until the mid to late 1970's that local area networks became anything more than a laboratory curiosity. First-generation LAN's operating at speeds up to 10 Mbits/s are widely deployed in office environments today. Second-generation LAN's and metropolitan area networks (MAN's) operating at about 100 Mbits/s are the subject of current standardization efforts, and third-generation LAN's operating in the Gbit/s region are the object of current research.

LAN's can take many forms, so that at times it becomes difficult to define just what constitutes a LAN. Indeed, it is not infrequently suggested that a digital PABX handling voice and data is a form of LAN. We have taken a somewhat more restricted definition. We believe local area networks differ from classical communications channels in that firstly, the transmission medium is shared by a number of users and, secondly, that at least some of the information is sent in packets containing a destination address. These two properties are also shared by many MAN's. LAN's and MAN's exploit similar concepts, the primary difference being in the geographical coverage provided.

The aim of this volume was to collect in one place a description of the most important principles of LAN's, and to emphasize some of the new directions that LAN technology is taking. In selecting contributions for the volume, we placed emphasis on concepts rather than accurately reflecting the current state of the technology; there are a number of available books that guide the data communications manager in his/her search for an appropriate network.

As part of COMSOC's Frontiers in Communications Series, the source of our material was primarily from the JOURNAL ON SELECTED AREAS IN COMMUNICATIONS. Where appropriate, other sources have been tapped. The chapters presented herein are not just reprints of the original papers, corrections have been made, repetitive sections have been deleted, and, where appropriate, new material has been added.

Part I provides an introduction to the subject for readers who have no previous knowledge of the area. This chapter provides a good foundation for understanding any of the chapters in Part II. In Part II an attempt is made to describe the most common types of LAN's. Bus, star, and ring topologies are presented; token, slotted, and buffer insertion rings and CSMA/CD are covered; coaxial cable, twisted pair, and optical fiber are referenced.

Optical fiber is a medium with many desirable characteristics on which to build a LAN, particularly where very high speeds are required. However, obtaining optical access to a fiber presents some problems and the chapters in Part III explore this. While the logical operation of a LAN may frequently be described in a handful of well-crafted sentences, they are very difficult to analyze in detail. In Part IV we present some landmark results describing the performance of some types of LAN's. It is still necessary to make some approximations in order to get solutions, but the models being analyzed are coming close to real systems.

One of the exciting aspects of LAN's is that they can be used for transmitting not just computer data but voice, facsimile, graphics, video, and visual information. Part V describes three studies in combining multiple types of traffic; however, there is still much work to be done in this area. The only commercial systems that we are aware of that combine voice, video, and data traffic do so using separate frequency bands—scarcely integration.

In order to expand the size of a LAN, to accommodate more users, or to get from one type of network to another requires a bridge or a gateway; a whole new set of problems is encountered with gateways and bridges and some discussion of these takes place in Part VI.

The purpose of LAN's, of course, is to provide the interconnection necessary to support distributed processing and resource sharing. In many cases, the multiple access property of shared-medium LAN's and the inherent broadcast capability permit approaches to supporting these applications that are very different from those permitted by more classical communications systems. Examples presented in Part VII describe distributed computing applications, servers providing simultaneous access from many clients, and distributed call processing for an integrated data and voice system.

We would particularly like to acknowledge the contributions of the authors. To an individual they worked conscientiously to maintain a high quality. Finally, we thank Reed Crone, Carolyne Tamney, and Valerie Cammarata of the IEEE Press for their committed support.

Contents

Part IV: Performance

Part V: Integrated Traffic

Part VI: Bridges and Gateways

Part VII: Software and Applications

Part I
Introduction

此为试读,需要完整PDF请访问: www.ertongbook.com

1
Local Area Networks—Major Technologies and Trends

KARL KÜMMERLE AND MARTIN REISER

In this chapter, we first review the origins of local area networks, discuss the major factors characterizing them, and outline the driving application scenarios. Particular emphasis is put onto bus and ring systems with both standardized and other medium-access control mechanisms. Specifically, we discuss the delay-throughput characteristic of carrier-sense multiple access with collision detection and token passing, and review the relevant engineering considerations for CSMA/CD baseband and broadband systems and for token rings. Finally, we summarize several new concepts and standardization efforts which aim at higher data rates than the IEEE 802 LAN's, and use optical-fiber media.

I. Introduction

During the early 1960's two developments took place which proved to be important for the later emergence of distributed local area networks (LAN's): the introduction of packet switching [1] and the realization of the ALOHA packet-radio network [2].

For a better understanding of the unique characteristics of the local area network, let us take a look at a typical cluster of terminals attached to a data-processing machine, also called a host. Figure 1(a) shows the structure of a local cluster. A set of terminals, typically up to 32, is attached to a control unit which at its backend connects to the I/O channel of the host. The necessary control functions are performed in the control unit and shared among the terminals. At the time when the cost of logic was a substantial factor of the overall system cost, this shared logic design obviously improved the economy.

Figure 1(b) depicts the case of remote terminals attached via modems and leased TP lines. Often, several terminals (or groups of terminals) share the line forming a so-called multidrop configuration. To avoid interference through multiple simultaneous transmissions, access to the multidrop line must be controlled. This is typically achieved through polling. The host, called the primary station, polls the terminals, the secondary stations. A terminal ready to send has to wait for its poll signal until allowed to initiate the data transmission. Note that control of the system is centralized and executed in the processor of the primary station.

The centralized polling mechanism used in the TP world was questioned by Abramson, who was faced with the task of interconnecting various terminals dispersed in a wide geographical area to a central computer via a radio network [2]. This network, known as

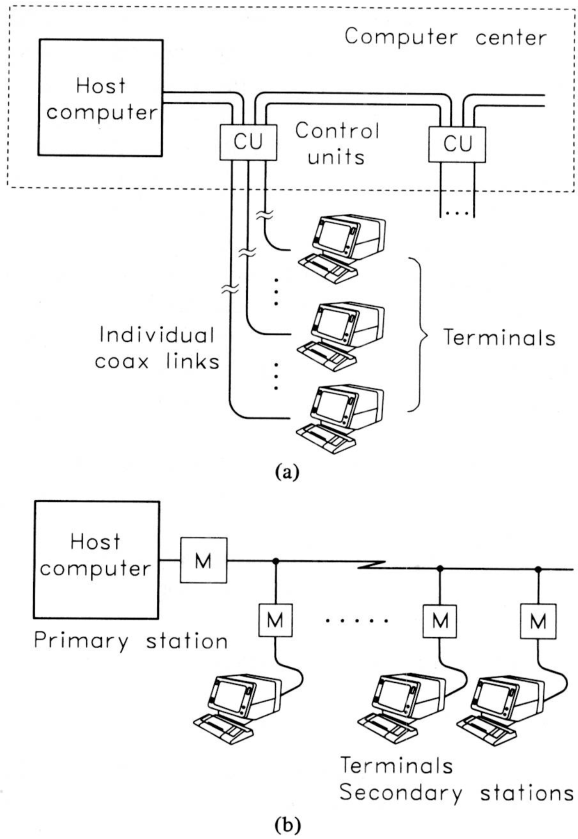

Fig. 1. Terminal-to-host communication. (a) Star system for local data communication. (b) TP cluster on multidrop link.

the ALOHA system, uses the simplest form of random access. A terminal ready to send simply sends its message, disregarding the state of the channel. Clearly, collisions may take place, in which case the terminals affected retransmit after a randomly chosen time-out.

The ALOHA system had a significant impact, especially on the academic community. The major benefits envisioned by the inventors were elimination of the access delay owing to waiting for the poll signal and simplicity of the system. There was no central control needed, the stations operated in a purely distributed fashion. However, simulation studies and analytical analysis showed that the elegance of random access was tarnished by a reduction of the usable bandwidth (18 percent in the case of ALOHA). A great deal of effort to improve the efficiency of random access commenced [2], [3].

During the late 1960's, R. Metcalfe was faced with the task of constructing a fast network linking small computers. He drew on the ideas of random access thus far discussed in the domain of packet radio and satellite communication. Metcalfe envisaged a "packet-radio network" on a coaxial cable, which he termed suggestively "Ethernet." One of Metcalfe's major contributions was carrier sensing, a technique which significantly improved the achievable performance of random-access systems. Running at 3 Mbits/s, a prototype "Ethernet" was implemented at the Xerox Palo Alto Research Center and used in the distributed office system prototype [4]. From this highly influential research into systems composed of intelligent workstations and specialized servers [5], the notion was lodged that the LAN was a key component of future office systems.

A substantial amount of research and development work aiming at the optimal LAN was launched and a rich body of literature developed, both advancing theory of access protocols as well as describing research prototypes [6], [7]. Alternatives to the original "Ethernet" design emerged. Re-examination of the theory of more traditional telecommunication systems revealed that there were designs which could be adopted to fit the new requirements of the distributed LAN. For example, a group of researchers at Cambridge University generalized the slotted-system approach, well known in PCM systems, and arrived at the slotted ring, also known as the "Cambridge Ring" [8].

Within the theory of polling systems, there is a symmetric version called hub-polling [9]. The token-access method was derived from the hub-polling scheme. Token rings were investigated independently and implemented at M.I.T. [10] and at IBM [11].

The mid-1970's and the early 1980's were characterized by vigorous research activities at universities and in industry, leading to a great number of competing system designs. Now, small companies, semiconductor vendors, and major providers of data-processing equipment offer local area network products and components. However, it is fair to say that the market explosion predicted by some analysts has not yet occurred.

II. Taxonomy and Application Scenarios

A. Definition

The object of this survey is the local area network, more precisely the distributed version of such networks which is schematically depicted in Fig. 2, and consists of the following elements:

1) a transmission medium shared among the attaching stations, and providing a broadcast capability,

2) a distributed protocol, the medium-access protocol, which controls access to the transmission system and prevents simultaneous transmissions, or recovers from collisions,

3) a set of cooperating adaptors through which stations attach to the network and which execute the access protocol, provide necessary data buffers, and interface with the processors of the station.

The service offered by such a LAN is packet switching, i.e., data blocks contain an explicit address recognized by the system, and used to deliver the data block to its destination.

B. Taxonomy

Local area networks may be characterized by the following factors:

1) Transmission medium:
- coaxial cable
- twisted-pair wire (unshielded and shielded)
- optical fiber

2) Modulation scheme:
- baseband transmission
- carrier-modulated transmission termed "broadband" in the LAN literature

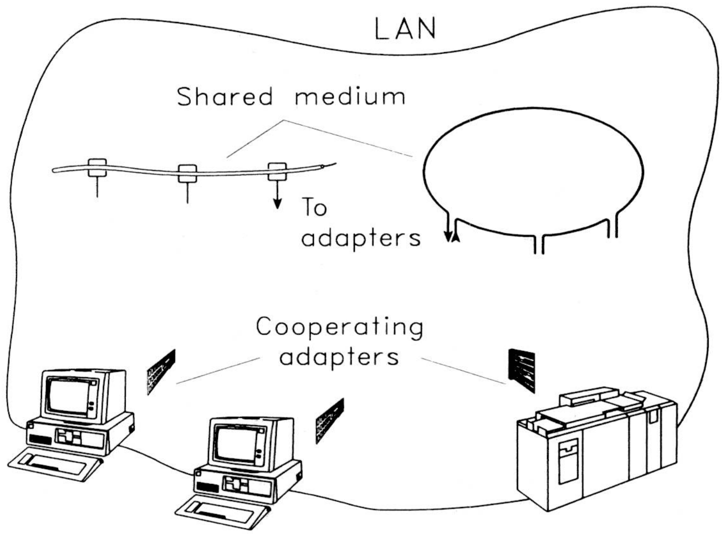

Fig. 2. LAN definition.

3) Wiring scheme or wiring topology:
   - bus or tree of buses
   - ring
   - star or set of connected stars
4) Logical topology:
   - bus
   - ring
   - star
5) Medium-access control:
   - random-access schemes, most importantly carrier-sense multiple access with collision detection (CSMA/CD) on bus topology
   - controlled access schemes, most importantly token passing on both bus and ring topologies; in addition, slotted ring and buffer/register insertion ring.

Some major system design points are the following:

- baseband bus with CSMA/CD, wiring: tree of bus segments, coaxial cable. This system is widely known as ''Ethernet'' [12], [13] [see Fig. 3(a)],
- broadband bus with token passing, wiring: tree, coaxial cable [14] [see Fig. 3(b)],
- baseband ring with token passing, star wiring from distribution frames to wall outlets, twisted-pair cable [15], [16] [see Fig. 3(c)],
- slotted ring, baseband transmission on twisted pairs [8].

These systems will be discussed in more technical detail in subsequent sections.

C. Driving Applications

1) Distributed Office System: Historically, the distributed office system was the first major application foreseen for the LAN [6]. Each user is provided with a workstation with a powerful microprocessor. The high function version of such a workstation is equipped with

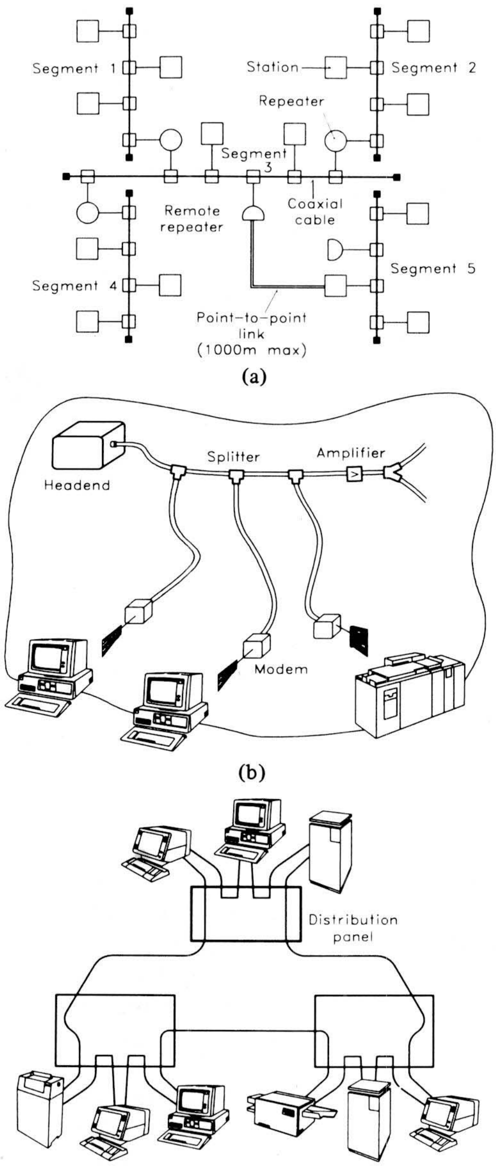

Fig. 3. Typical LAN's. (a) "Ethernet." (b) Broadband bus. (c) Token ring.

an all-points addressable full screen display, a pointing device such as a mouse, and a local hard disk. The user interface (e.g., editors, graphics spreadsheets, etc.) is executed on the local processor. However, there are expensive resources such as laser printers which must be shared for cost reasons. Also, users may prefer a reliable central storage facility which will alleviate the need for data security and management.

In [5], the term clients was introduced for workstations and servers for shared resources, e.g., file server, printer server, communication server, mail server. The local area network is the glue which holds the distributed system together (see Fig. 4). From the need for favorable human factors, the data rate of such a local area network needs to be high, i.e., 1 to 10 Mbits/s.

2) Terminal to Host Communications, Front-End Network: The local cluster design of Fig. 1(a) served the industry well and is widely used by mainframe manufacturers. The cost advantages which accrued from the shared logic design were vital in the 1960's and early 1970's, and only now lose importance owing to the low cost of VLSI microprocessors of suitable power.

However, leading establishments witness growing problems in maintaining their terminal or front-end networks (see Fig. 5).

1) The number of terminals approaches the number of employees. Terminals are installed and frequently moved.

2) The fact that controllers need to be placed in close proximity to the data hosts since I/O channels are of limited length (typically 200 ft.). Space is at a premium in the machine room, and placement of a few tens of controllers may become difficult.

3) Cables converge in the machine room, and may crowd cable ducts and cable termination racks.

4) There are different hosts and a variety of smaller machines; a variety of types and grades of cables are in use.

5) Terminals are physically bound to one processor. If the user needs services from another host, he has to go to a different terminal or he needs more than one cable strung to his office, thus aggravating the duct crowding mentioned above.

Various tactical fixes are in use, such as link multiplexors, channel extenders, and wide-band circuit switches. It is the local area network, however, which may lead to a truly

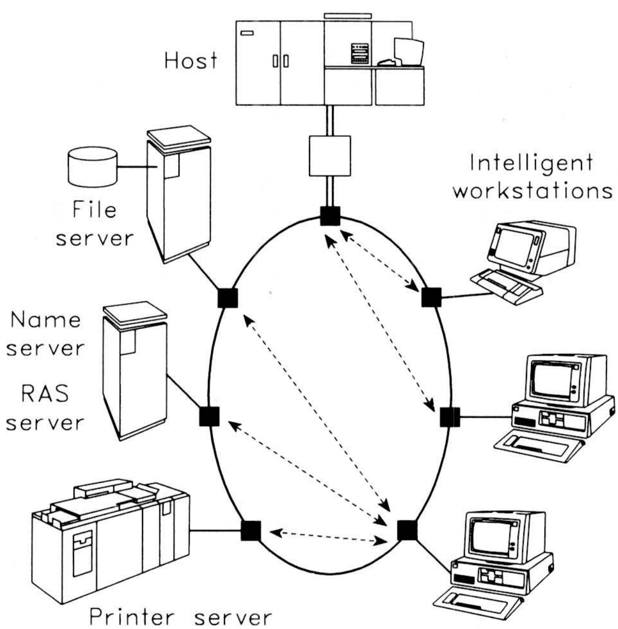

Fig. 4. Client-server network.

此为试读,需要完整PDF请访问: www.ertongbook.com