Computer Network Architectures and Protocols

SECOND EDITION

Computer Network Architectures and Protocols

SECOND EDITION

江苏工业学院图书馆
藏书章

Edited by

Carl A. Sunshine

Unisys West Coast Research Center
Santa Monica, California

PLENUM PRESS • NEW YORK AND LONDON

Library of Congress Cataloging in Publication Data

Computer network architectures and protocols / edited by Carl A. Sunshine. — 2nd ed.
p. cm. — (Applications of communications theory)
Includes bibliographical references.
ISBN 0-306-43189-0
1. Computer network protocols. 2. Computer network architectures. I. Sunshine, Carl A. II. Series.
TK5105.5.C638 1989 89-35230
004.6 — dc20 CIP

© 1989, 1982 Plenum Press, New York
A Division of Plenum Publishing Corporation
233 Spring Street, New York, N.Y. 10013

All rights reserved

No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, microfilming, recording, or otherwise, without written permission from the Publisher

Printed in the United States of America

Contributors

Paul Bartoli
AT&T Bell Laboratories, Holmdel, New Jersey 07733

H. V. Bertine
AT&T Bell Laboratories, Holmdel, New Jersey 07733

Abhay K. Bhushan
Xerox Corporation, El Segundo, California 90245

Gregor V. Bochmann
Département d'IRO, University of Montreal, Montreal H3C 3J7, Canada

David E. Carlson
AT&T Bell Laboratories, Holmdel, New Jersey 07733

James W. Conard
Conard Associates, Costa Mesa, California 92626

John Day
Motorola, Inc., Schaumburg, Illinois 60196

Harold C. Folts
Omnicom, Inc., Vienna, Virginia 22180

Dennis G. Frahmann
Xerox Corporation, El Segundo, California 90245

Mario Gerla
Computer Science Department, UCLA, Los Angeles, California 90024

James P. Gray
IBM Corporation, Research Triangle Park, North Carolina 27709

Leonard Kleinrock
Computer Science Department, UCLA, Los Angeles, California 90024

Daniel A. Pitt
IBM Corporation, Research Triangle Park, North Carolina 27709

Diane P. Pozefsky
IBM Corporation, Research Triangle Park, North Carolina 27709

Antony Rybczynski
Data Networks Division, Northern Telecom, Ottawa, Ontario K2C 3T1, Canada

Mischa Schwartz
Center for Telecommunications Research, Columbia University, New York, New York 10027

Thomas E. Stern
Center for Telecommunications Research, Columbia University, New York, New York 10027

Carl A. Sunshine
Unisys, West Coast Research Center, Santa Monica, California 90406

Fouad A. Tobagi
Computer Systems Laboratory, Stanford University, Stanford, California 94305

Ronald P. Uhlig
Northern Telecom Inc., Richardson, Texas 75081

Charles E. Young
AT&T Bell Laboratories, Holmdel, New Jersey 07733

此为试读，需要完整PDF请访问：www.ertongbook.com

Preface

This is a book about the bricks and mortar from which are built those edifices that will permeate the emerging information society of the future—computer networks. For many years such computer networks have played an indirect role in our daily lives as the hidden servants of banks, airlines, and stores. Now they are becoming more visible as they enter our offices and homes and directly become part of our work, entertainment, and daily living.

The study of how computer networks function is a combined study of communication theory and computer science, two disciplines appearing to have very little in common. The modern communication scientist wishing to work in this area soon finds that solving the traditional problems of transmission, modulation, noise immunity, and error bounds in getting the signal from one point to another is just the beginning of the challenge. The communication must be in the right form to be routed properly, to be handled without congestion, and to be understood at various points in the network.

As for the computer scientist, he finds that his discipline has also changed. The fraction of computers that belong to networks is increasing all the time. And for a typical single computer, the fraction of its execution load, storage occupancy, and system management problems that are involved with being part of a network is also growing.

It is the objective of this book to provide a comprehensive text and reference volume that can be used in education, research, and development in this combined field of computer networks. The aim is to cover both theory and practice in a style that is instructive but highly readable. To this end, the majority of the volume is devoted to a presentation of structural principles and architectural concepts, with emphasis on the OSI international standards, and illustrated with brief examples from currently operating systems. This is followed by a detailed description of network systems developed by two leading vendors: Xerox and IBM.

An effort of this scope is beyond the capabilities of any single author. Building on the framework created in the First Edition by Paul Green, the "theory" chapters have been revised to reflect current applications. Chapters dealing with the more stable lower levels of the protocol architecture have been updated. Entirely new chapters were solicited covering the emerging higher-layer OSI standards and the IBM and Xerox network systems. The result represents a collection of tutorials by outstanding experts in each area, providing a more extensive coverage of the subject than any single author could.

Any summary of the present status of the computer network art must draw upon three main sources: research networks built by universities, often operating under government support; private networks provided by the computer manufacturers; and public network offerings provided by common carriers. In this volume all three sources of expertise have been tapped. Cooperation among these communities has led to the establishment of a high level of worldwide standardization on many aspects of network architectures and protocols, and these form the core of the new edition.

The book is organized along the lines of the now familiar "layered" view of networks, which is introduced in Part I. According to this scheme, the structure within any network "node" or "machine" may be broken into layers, with the raw transmission facilities of classical communications (for example, wires or satellite links) at the bottom, and the "users" (human or program) at the top. Part II discusses the lowest layer, the physical layer, by which data are transmitted between physically connected nodes. Part III presents the data link layer, which operates to deliver entire messages or "packets" from a particular node to one of its neighboring nodes. This includes the operation of "multiaccess links" and high-speed local area networks such as Ethernet, which are in such widespread use today.

In Part IV we see how packets make their way from the originating node to the destination node within the network layer, a process that can be a complex one when there are intermediate nodes and when there are many simultaneous users of network resources. Problems of routing, congestion control, and interconnecting multiple networks into "internets" are addressed. When we get to Part V, the fact that the path of the messages has been a sequence of nodes and links is no longer apparent. The higher layers deal with functions handled between end user nodes, such as error recovery, checkpointing, agreeing on a common format for exchanging data, and separating streams of data from different "conversations."

Part VI presents in detail the network systems developed by Xerox and IBM. These examples illustrate how two major vendors have chosen to apply the principles described earlier in the book. Finally, Part VII summarizes another increasingly important area: how to ensure the correctness of computer network protocols using formal specification and validation methods. This can help guarantee that the protocol designs are free of bugs and that specific products properly implement the protocols so that they will interoperate with other equipment.

Assembling this collection has taken long hours by many busy individuals. Some of the authors have borne the burden of revising their chapters several times to reflect new developments over the two-year period that the Second Edition has been underway. I wish to specially thank the Unisys Corporation and my family for supporting my own efforts on this edition. I am also grateful for the energy and patience of my colleagues, and of Sy Marchand at Plenum Press, in bearing with me through the lengthy gestation period of this volume, whose birth I hope will reflect the growing maturity and importance of computer network architectures and protocols.

Carl A. Sunshine

Santa Monica, California

Contents

I

Introduction

此为试读,需要完整PDF请访问: www.ertongbook.com

1

A Brief History of Computer Networking

Carl A. Sunshine

Computer networking as we know it today may be said to have gotten its start with the ARPANET development in the late 1960s and early 1970s. Prior to that time there were computer vendor "networks" designed primarily to connect terminals and remote job entry stations to a mainframe. But the notion of networking between computers viewing each other as equal peers to achieve "resource sharing" was fundamental to the ARPANET design [1]. The other strong emphasis of the ARPANET work was its reliance on the then novel technique of packet switching to efficiently share communication resources among "bursty" users, instead of the more traditional message or circuit switching.

Although the term "network architecture" was not yet widely used, the initial ARPANET design did have a definite structure and introduced another key concept: protocol layering, or the idea that the total communications functions could be divided into several layers, each building upon the services of the one below. The original design had three major layers, a network layer, which included the network access and switch-to-switch (IMP-to-IMP) protocols, a host-to-host layer (the Network Control Protocol or NCP), and a "function-oriented protocol" layer, where specific applications such as file transfer, mail, speech, and remote terminal support were provided [2].

Similar ideas were being pursued in several other research projects around the world, including the Cyclades network in France [3], the

Carl A. Sunshine • Unisys, West Coast Research Center, Santa Monica, California 90406.

National Physical Laboratory Network in England [4], and the Ethernet system [5] at Xerox PARC in the USA. Some of these projects focused more heavily on the potential for high-speed local networks such as the early 3-Mbps Ethernet. Satellite and radio channels for mobile users were also a topic of growing interest.

By 1973 it was clear to the networking vanguard that another protocol layer needed to be inserted into the protocol hierarchy to accommodate the interconnection of diverse types of individual networks. Cerf and Kahn published their seminal paper describing such a scheme [6], and development of the new Internet Protocol (IP) and Transmission Control Protocol (TCP) to jointly replace the NCP began. Similar work was being pursued by other groups meeting in the newly formed IFIP WG 6.1, called the Internetwork Working Group [7].

The basis for the network interconnection approach developing in this community was to make use of a variety of individual networks each providing only a simple "best effort" or "datagram" transmission service. Reliable virtual circuit services would then be provided on an end-to-end basis with the TCP (or similar protocol) in the hosts. During the same time period, public data networks (PDNs) were emerging under the auspices of CCITT, aimed at providing more traditional virtual circuit types of network service via the newly defined X.25 protocol. The middle and late 1970s saw networking conferences dominated by heated debates over the relative merits of circuit versus packet switching and datagrams versus X.25 virtual circuits [8]. The computer vendors continued to offer their proprietary networks, gradually supporting the new X.25 service as links under their own protocols. Digital Equipment Corporation (DEC) was the notable exception, adopting the research community approach of peer-to-peer networking at an early date, and coming out with its own new suite of protocols (DECNET) [9].

By the late 1970s, a new major influence was emerging in the computer network community. The computer manufacturers realized that multivendor systems could no longer be avoided, and began to take action to satisfy the growing user demand for interoperability. Working through their traditional international body, the ISO, a new group (SC16) was created to develop standards in the networking area. Their initial charter was to define an explicit "architecture" for "Open Systems Interconnection" (OSI).

By the early 1980s there were three major players in the networking game: the ARPANET-style research community, the carriers with their PDNs in CCITT, and the manufacturers in ISO. The conference circuit became more acrimonious, with the research community lambasting the slow progress, ponderousness (7 layers!), lack of experimental support, and all-inclusiveness (five classes of transport protocol) of the ISO workers, while still taking occasional shots at the PDNs and X.25. The CCITT and

ISO had the big players and the dollars on their side, but TCP/IP protocols were included in the increasingly popular UNIX* operating system.

The outgrowth of the ARPANET, the DoD Internet, was beginning to face its own problems of success by the mid-1980s. With the hundreds of LANs and tens of thousands of workstation hosts, serious performance problems were emerging, and it was beginning to look like the critics of "stateless" datagram networking might have been right on some points. After the first ten years, the DoD had greatly reduced research funding for networks, and found that it had to hurriedly perform some engineering studies to maintain operations.

Now in the late 1980s, much of the battling seems over. CCITT and ISO have aligned their efforts, and the research community seems largely to have resigned itself to OSI. Bob Metcalfe, the creator of Ethernet and founder of 3Com, summed up the sentiment recently by stating that "it has not been worth the ten years wait to get from TCP to TP4, but OSI is now inevitable." OSI has created an internet sublayer within the network layer to accommodate the datagram internetting approach beside the CCITT X.25 approach. However, with the inevitable tendency of committee work toward all-inclusiveness, there remain some serious potential pitfalls in interoperability among those following different "profiles" of OSI protocols (e.g., how will TP2 users on X.25 networks talk to TP4 users on datagram networks?).

In hindsight, much of the networking debate has resulted from differences in how to prioritize the basic network design goals such as accountability, reliability, robustness, autonomy, efficiency, and cost effectiveness [10]. Higher priority on robustness and autonomy led to the DoD Internet design, while the PDNs have emphasized accountability and controllability.

It is ironic that while a consensus has developed that OSI is indeed inevitable, the TCP/IP protocol suite has achieved widespread deployment, and now serves as a de facto interoperability standard. Research is underway again, with new results on how to optimize TCP/IP performance over variable delay and/or very-high-speed networks. The TCP/IP Interoperability Conference (itself a new event started only in 1986) now features slick booths by the major computer and network vendors rather than the small niche vendors that once dominated the DoD networking marketplace. It appears that the vendors were unable to bring OSI products to market quickly enough to satisfy the demand for interoperable systems, and TCP/IP were there to fill the need.

The majority of this book deals with the OSI architecture and protocols because these will be dominant in the future. But efforts have been made wherever possible to show their roots in previous work, and the

*UNIX is a trademark of AT & T Bell Laboratories.

relation of OSI protocols to research and vendor efforts. The past decade has seen a tremendous rate of change in computing technology and in the networks for interconnecting computers. The architecture for the lower layers appears to be stabilizing, and the material in this book should carry us well into the next decade. The application layer is the new frontier, and will require another book to deal with in the future.

References

[1] L. Roberts and B. Wessler, "Computer network development to achieve resource sharing," *AFIPS Conf. Proc. (SJCC)*, vol. 36, 1970, pp. 543–549.

[2] S. Crocker *et al.*, "Function-oriented protocols for the ARPA computer network," *AFIPS Conf. Proc.*, vol. 40, 1972, pp. 271–279.

[3] L. Pouzin, "Presentation and major design aspects of the Cyclades computer network," *Proc. 3rd Data Commun. Symp.*, St. Petersburg, Nov. 1973, pp. 80–85.

[4] R. Scantlebury and P. Wilkinson, "The National Physical Laboratory data communication network," *Proc. ICCC*, Stockholm, Aug. 1974.

[5] R. Metcalfe and D. Boggs, "ETHERNET: Distributed packet switching for local computer networks," *Commun. ACM*, vol. 19, July 1976, pp. 395–404.

[6] V. Cerf and R. Kahn, "A protocol for packet network intercommunication," *IEEE Trans. Commun.*, vol. COM-22, 1974, pp. 637–648.

[7] V. Cerf, A. McKenzie, R. Scantlebury, and H. Zimmermann, "Proposal for an international end-to-end protocol," *Computer Commun. Rev.*, vol. 6, ACM SIGCOMM, 1974, pp. 68–89.

[8] V. DiCiccio, C. Sunshine, J. Field, and E. Manning, "Alternatives for interconnection of public packet switching data networks," *Proc. 6th Data Commun. Symp.*, Nov. 1979, ACM/IEEE, pp. 120–125.

[9] S. Wecker, "DNA—The Digital Network Architecture," *Computer Network Architectures and Protocols*, P. Green, Jr., ed., New York: Plenum Press, 1982.

[10] D. Clark, "The design philosophy of the DARPA internet protocols," *Proc. SIGCOMM '88 Symp.*, Stanford, Aug. 1988, pp. 106–114.

2

The Reference Model for Open Systems Interconnection

John Day

I. Introduction

In the mid-1970s, the development and use of resource-sharing computer networks began to achieve considerable attention. The early successes of the ARPANET [1] and CYCLADES [2], the immediate commercial potential of packet switching, satellite, and local network technology, and the declining cost of hardware made it apparent that computer networking was quickly becoming an important area of innovation and commerce. It was also apparent that to utilize the full potential of such computer networks, international standards would be required to ensure that any system could communicate with any other system anywhere in the world.

In 1978, the International Organization for Standardization (ISO) Technical Committee 97 on Information Processing, recognizing the standards for networks of heterogeneous systems were urgently required, created a new subcommittee (SC16) for "open systems interconnection" (OSI). The term "open" was chosen to emphasize that by conforming to OSI standards, a system would be open to communication with any other system anywhere in the world obeying the same standards.

It was clear that the commercial endeavors to exploit the emerging communication technology would wait neither for SC16 to leisurely develop communication standards nor for the research community to answer many of the outstanding questions. If there was to be a consistent set of

JOHN DAY • Motorola, Inc., Schaumburg, Illinois 60196.

此为试读,需要完整PDF请访问: www.ertongbook.com